VMware vSphere:
Operate, Scale and Secure [V8]

Lecture Manual

vmware®

VMware® Education Services
VMware, Inc.
www.vmware.com/education

VMware vSphere: Operate, Scale and Secure [V8]
Lecture Manual
Part Number EDU-EN-VSOSS8-LEC (11-JUL-2023)

Copyright © 2023 VMware, Inc. All rights reserved. This manual and its accompanying materials are protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. VMware vSphere® vMotion®, VMware vSphere® Web Client, VMware vSphere® Virtual Volumes™, VMware vSphere® Update Manager™, VMware vSphere® Trust Authority™, VMware vSphere® Storage vMotion®, VMware vSphere® Storage DRS™, VMware vSphere® Storage APIs - Data Protection, VMware vSphere® Storage APIs - Array Integration, VMware vSphere® Replication™, VMware vSphere® PowerCLI™, VMware vSphere® Native Key Provider™, VMware vSphere® Lifecycle Manager™, VMware vSphere® High Availability, VMware vSphere® Fault Tolerance, VMware vSphere® ESXi™ Shell, VMware vSphere® ESXi™ Image Builder CLI, VMware vSphere® ESXi™ Dump Collector, VMware vSphere® Distributed Switch™, VMware vSphere® Distributed Services Engine™, VMware vSphere® Data Protection™, VMware vSphere® Distributed Resource Scheduler™, VMware vSphere® Distributed Power Management™, VMware vSphere® Configuration Profiles™, VMware vSphere® Client™, VMware vSphere® Command-Line Interface, VMware vSphere® Auto Deploy™, VMware vSphere® Authentication Proxy, VMware vSphere® API for Storage Awareness™, VMware vSphere® API, VMware vSphere® 2015, VMware vSphere®, VMware vShield Endpoint™, VMware vSAN™, VMware vRealize® Operations Manager™, VMware vRealize® Operations Manager™ for Horizon®, VMware vRealize® Operations™ Enterprise, VMware vRealize® Operations™, VMware vRealize® Operations™ Standard, VMware vRealize® Operations™ Advanced, VMware vRealize®, VMware vCloud®, VMware vCenter® Single Sign-On, VMware vCenter® Server Appliance™, VMware vCenter Server®, VMware vCenter®, VMware Workstation™, VMware Horizon® 7, VMware Verify™, VMware Validated Design™ for Software-Defined Data Center, VMware Validated Design™ for IT Automation Cloud, VMware Skyline™ Health, VMware Skyline™, VMware Site Recovery™ for VMware Cloud™ on AWS, VMware® Postgres, VMware Host Client™, VMware Customer Connect™, VMware Cloud Foundation™, VMware Cloud Foundation™ for Amazon EC2, VMware Cloud Foundation™ for Remote Office Branch Office, VMware Cloud™ on AWS GovCloud (US), VMware Cloud™ on AWS Outposts, VMware Certificate Authority. No trademark., VMware vSphere® VMFS, VMware vSphere® Distributed Switch™, vSphere Storage vMotion, VMware vSphere® Storage I/O Control, VMware Skyline™ Health Diagnostics™, VMware Skyline Collector™, VMware Skyline Advisor™ Pro, VMware Site Recovery™, VMware PowerCLI™, Project Photon OS™, VMware vSphere® Network I/O Control, VMware NSX® Manager™, VMware NSX®, VMware NSX® Professional, VMware NSX® for Remote Office Branch Office, VMware NSX® for Desktop, VMware NSX® Enterprise Plus, My VMware®, vCenter Linked Mode, VMware Lab Connect™, VMware Horizon® Standard Edition, VMware Pivotal Labs® Health Check™, VMware Go™, Express Storage Architecture™, VMware ESXi™, VMware ESX®, VMware Cloud Foundation™ Distributed Services Engine™, VMware vSphere® Distributed Services Engine™, and VMware vSphere® Distributed Resource Scheduler™ are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

The training material is provided “as is,” and all express or implied conditions, representations, and warranties, including any implied warranty of merchantability, fitness for a particular purpose or noninfringement, are disclaimed, even if VMware, Inc., has been advised of the possibility of such claims. This material is designed to be used for reference purposes in conjunction with a training course.

The training material is not a standalone training tool. Use of the training material for self-study without class attendance is not recommended. These materials and the computer programs to which it relates are the property of, and embody trade secrets and confidential information proprietary to, VMware, Inc., and may not be reproduced, copied, disclosed, transferred, adapted or modified without the express written approval of VMware, Inc.

www.vmware.com/education

VMware vSphere: Operate, Scale and Secure [V8]
Lecture Manual
Part Number EDU-EN-VSOSS8-LEC (11-JUL-2023)

Copyright © 2023 VMware, Inc. All rights reserved. This manual and its accompanying materials are protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. VMware vSphere® vMotion®, VMware vSphere® Web Client, VMware vSphere® Virtual Volumes™, VMware vSphere® Update Manager™, VMware vSphere® Trust Authority™, VMware vSphere® Storage vMotion®, VMware vSphere® Storage DRS™, VMware vSphere® Storage APIs - Data Protection, VMware vSphere® Storage APIs - Array Integration, VMware vSphere® Replication™, VMware vSphere® PowerCLI™, VMware vSphere® Native Key Provider™, VMware vSphere® Lifecycle Manager™, VMware vSphere® High Availability, VMware vSphere® Fault Tolerance, VMware vSphere® ESXi™ Shell, VMware vSphere® ESXi™ Image Builder CLI, VMware vSphere® ESXi™ Dump Collector, VMware vSphere® Distributed Switch™, VMware vSphere® Distributed Services Engine™, VMware vSphere® Data Protection™, VMware vSphere® Distributed Resource Scheduler™, VMware vSphere® Distributed Power Management™, VMware vSphere® Configuration Profiles™, VMware vSphere® Client™, VMware vSphere® Command-Line Interface, VMware vSphere® Auto Deploy™, VMware vSphere® Authentication Proxy, VMware vSphere® API for Storage Awareness™, VMware vSphere® API, VMware vSphere® 2015, VMware vSphere®, VMware vShield Endpoint™, VMware vSAN™, VMware vRealize® Operations Manager™, VMware vRealize® Operations Manager™ for Horizon®, VMware vRealize® Operations™ Enterprise, VMware vRealize® Operations™, VMware vRealize® Operations™ Standard, VMware vRealize® Operations™ Advanced, VMware vRealize®, VMware vCloud®, VMware vCenter® Single Sign-On, VMware vCenter® Server Appliance™, VMware vCenter Server®, VMware vCenter®, VMware Workstation™, VMware Horizon® 7, VMware Verify™, VMware Validated Design™ for Software-Defined Data Center, VMware Validated Design™ for IT Automation Cloud, VMware Skyline™ Health, VMware Skyline™, VMware Site Recovery™ for VMware Cloud™ on AWS, VMware® Postgres, VMware Host Client™, VMware Customer Connect™, VMware Cloud Foundation™, VMware Cloud Foundation™ for Amazon EC2, VMware Cloud Foundation™ for Remote Office Branch Office, VMware Cloud™ on AWS GovCloud (US), VMware Cloud™ on AWS Outposts, VMware Certificate Authority. No trademark., VMware vSphere® VMFS, VMware vSphere® Distributed Switch™, vSphere Storage vMotion, VMware vSphere® Storage I/O Control, VMware Skyline™ Health Diagnostics™, VMware Skyline Collector™, VMware Skyline Advisor™ Pro, VMware Site Recovery™, VMware PowerCLI™, Project Photon OS™, VMware vSphere® Network I/O Control, VMware NSX® Manager™, VMware NSX®, VMware NSX® Professional, VMware NSX® for Remote Office Branch Office, VMware NSX® for Desktop, VMware NSX® Enterprise Plus, My VMware®, vCenter Linked Mode, VMware Lab Connect™, VMware Horizon® Standard Edition, VMware Pivotal Labs® Health Check™, VMware Go™, Express Storage Architecture™, VMware ESXi™, VMware ESX®, VMware Cloud Foundation™ Distributed Services Engine™, VMware vSphere® Distributed Services Engine™, and VMware vSphere® Distributed Resource Scheduler™ are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

The training material is provided “as is,” and all express or implied conditions, representations, and warranties, including any implied warranty of merchantability, fitness for a particular purpose or noninfringement, are disclaimed, even if VMware, Inc., has been advised of the possibility of such claims. This material is designed to be used for reference purposes in conjunction with a training course.

The training material is not a standalone training tool. Use of the training material for self-study without class attendance is not recommended. These materials and the computer programs to which it relates are the property of, and embody trade secrets and confidential information proprietary to, VMware, Inc., and may not be reproduced, copied, disclosed, transferred, adapted or modified without the express written approval of VMware, Inc.

www.vmware.com/education

VitalSour...

Contents

Module 1 Course Introduction 1

1-2 Course Introduction 1
1-3 Importance 1
1-4 Learner Objectives (1) 1
1-5 Learner Objectives (2) 2
1-6 Course Outline 2
1-7 Typographical Conventions 3
1-8 References (1) 4
1-9 References (2) 5
1-10 VMware Online Resources 6
1-11 VMware Learning Overview 7
1-12 VMware Certification Overview 8
1-13 VMware Credentials Overview 9
1-14 Lab 1: Accessing the Lab Environment 10

Module 2 Virtual Machine Operations 11

2-2 Importance 11
2-3 Module Lessons 11
2-4 Lesson 1: VMware Tools Repository 12
2-5 Learner Objectives 12
2-6 Centralized Repository for VMware Tools 12
2-7 Configuring a VMware Tools Repository 12
2-8 Configuring a VMware Tools Repository using the vSphere Client 13
2-9 Configuring a VMware Tools Repository using the ESXi Shell 13
2-10 Updating /productLocker Symlink (1) 14
2-11 Updating /productLocker Symlink (2) 15
2-12 Configuring a VMware Tools Repository using PowerCLI 16

2-13 Lab 2: Configuring a Centralized VMware Tools Installation Repository 17
2-14 Review of Learner Objectives 17
2-15 Lesson 2: VM Backups 18
2-16 Learner Objectives 18
2-17 Backup and Restore Solution for VMs 18
2-18 vSphere Storage APIs - Data Protection: Offloaded Backup Processing 19
2-19 vSphere Storage APIs - Data Protection: Changed-Block Tracking 20
2-20 Virtual Disk Transport Methods 21
2-21 SAN Transport Mode 21
2-22 HotAdd Transport Mode (1) 22
2-23 HotAdd Transport Mode (2) 23
2-24 NBDSSL Transport Mode (1) 23
2-25 NBDSSL Transport Mode (2) 23
2-26 Local File Access 24
2-27 Review of Learner Objectives 24
2-28 Lesson 3: vSphere Replication 25
2-29 Learner Objectives 25
2-30 About vSphere Replication 25
2-31 vSphere Replication Benefits 26
2-32 About the vSphere Replication Appliance 26
2-33 vSphere Replication Appliance Components 27
2-34 vSphere Replication Management Service (VRMS) Appliance Management Interface (VAMI) 27
2-35 Replication Functions 28
2-36 Replication Between Two Sites 29
2-37 Replication in a Single vCenter 30
2-38 Replication to a Shared Target Site 31
2-39 Deploying the vSphere Replication Appliance 32
2-40 Cloud Provider 33
2-41 Lab 3: (Simulation) Deploying vSphere Replication and Site Recovery Instances 33
2-42 Review of Learner Objectives 33
2-43 Lesson 4: VM Replication 34
2-44 Learner Objectives 34
2-45 Configuring vSphere Replication for a Single VM 34
2-46 Configuring Recovery Point Objective and Point in Time Instances 35
2-47 Recovering Replicated VMs (1) 36

2-48 Recovering Replicated VMs (2) 36
2-49 Lab 4: (Simulation) Configuring Replication for a Single VM 37
2-50 Review of Learner Objectives 37
2-51 Key Points 37

Module 3 vSphere Cluster Operations 39

3-2 Importance 39
3-3 Module Lessons 39
3-4 Lesson 1: Resource Pools 40
3-5 Learner Objectives 40
3-6 Reservations, Limits, and Shares (1) 40
3-7 Reservations, Limits, and Shares (2) 41
3-8 About Resource Pools 42
3-9 Reasons to Use Resource Pools 43
3-10 Resource Pool Attributes 44
3-11 Resource Pool Example 45
3-12 Resource Pool Example: CPU Shares 46
3-13 Resource Pool Example: CPU Contention 47
3-14 About Expandable Reservations 48
3-15 Example of Expandable Reservation (1) 49
3-16 Example of Expandable Reservation (2) 50
3-17 Resource Pool Summary Tab 51
3-18 Resource Reservation Details 52
3-19 Resource Pool Shares 53
3-20 Diluted Shares 54
3-21 About Scalable Shares (1) 55
3-22 About Scalable Shares (2) 56
3-23 Enabling Scalable Shares 56
3-24 Lab 5: Managing Resource Pools 57
3-25 Review of Learner Objectives 57
3-26 Lesson 2: vSphere Cluster Service (vCLS) 58
3-27 Learner Objectives 58
3-28 About vCLS (1) 58
3-29 About vCLS (2) 59
3-30 About vCLS (3) 59
3-31 vCLS Components 60
3-32 About vCLS VMs (1) 62

3-33 About vCLS VMs (2) 63
3-34 About EAM Agency 63
3-35 vCLS Cluster Creation Workflow 64
3-36 vCLS Cluster Upgrade Workflow 65
3-37 Moving ESXi Hosts Between Clusters 66
3-38 vCLS Functions 67
3-39 Maintaining Health of vCLS (1) 68
3-40 Maintaining Health of vCLS (2) 69
3-41 Retreat Mode 69
3-42 Activating Retreat Mode 70
3-43 Troubleshooting Log Files 71
3-44 Lab 6: Enabling vCLS Retreat Mode 71
3-45 Review of Learner Objectives 72
3-46 Key Points 72

Module 4 Network Operations 73

4-2 Importance 73
4-3 Module Lessons 73
4-4 Lesson 1: Managing vSphere Distributed Switches 74
4-5 Learner Objectives 74
4-6 About Distributed Switches 75
4-7 About the Traffic Filtering and Marking Policy 76
4-8 Creating Network Traffic Rules 77
4-9 Network Traffic Rule: Example 1 78
4-10 Network Traffic Rule: Example 2 79
4-11 Marking Network Traffic (1) 80
4-12 Marking Network Traffic (2) 81
4-13 Access to NSX Virtual Distributed Switch (1) 82
4-14 Access to NSX Virtual Distributed Switch (2) 83
4-15 About Health Check for Distributed Switches 84
4-16 Health Check Example 85
4-17 Enabling Health Check 86
4-18 Monitoring Health Check Results 87
4-19 Backing Up and Restoring Configurations 88
4-20 Exporting Distributed Switch Configurations 89
4-21 Restoring and Importing Distributed Switch Configurations 90
4-22 Rollback and Recovery of the Host Management Network 91

4-23 Rollback Details 92
4-24 Recovery Through the DCUI 93
4-25 Lab 7: Configuring vSphere Distributed Switches 94
4-26 Lab 8: Managing vSphere Distributed Switches 94
4-27 Review of Learner Objectives 94
4-28 Lesson 2: Using Network I/O Control in Distributed Switches 95
4-29 Learner Objectives 95
4-30 About Network I/O Control (1) 96
4-31 About Network I/O Control (2) 97
4-32 About the Bandwidth Allocation Model for System Traffic 98
4-33 Configuring Bandwidth Allocations for System Traffic 99
4-34 Reserving Bandwidth for System Traffic (1) 100
4-35 Reserving Bandwidth for System Traffic (2) 101
4-36 Bandwidth Aggregation for Network Resource Pools 102
4-37 Creating Network Resource Pools 103
4-38 Bandwidth Allocation for VM Traffic 104
4-39 Bandwidth Allocation for Individual VMs 105
4-40 Bandwidth Admission Control in vSphere DRS (1) 106
4-41 Bandwidth Admission Control in vSphere DRS (2) 106
4-42 Bandwidth Admission Control in vSphere HA 107
4-43 Activity: Network I/O Control (1) 107
4-44 Activity: Network I/O Control (2) 108
4-45 Review of Learner Objectives 108
4-46 Lesson 3: NetFlow and Port Mirroring in Distributed Switches 109
4-47 Learner Objectives 109
4-48 About NetFlow 109
4-49 About Network Flows 110
4-50 Network Flow Analysis (1) 111
4-51 Network Flow Analysis (2) 112
4-52 Configuring NetFlow on Distributed Switches 113
4-53 Enabling NetFlow on a Distributed Port Group 114
4-54 About Port Mirroring 115
4-55 Port Mirroring Session Types (1) 116
4-56 Port Mirroring Session Types (2) 117
4-57 Port Mirroring Session Properties 118
4-58 Source and Destination in a Port Mirroring Session 119

4-59 Lab 9: Using Port Mirroring 120
4-60 Review of Learner Objectives 120
4-61 Lesson 4: vSphere Distributed Services Engine 121
4-62 Learner Objectives 121
4-63 SmartNICs 121
4-64 Traditional Infrastructure Challenges 122
4-65 About vSphere Distributed Services Engine 123
4-66 Next Generation Infrastructure with vSphere Distributed Services Engine 124
4-67 vSphere Distributed Services Engine Benefits 125
4-68 vSphere Distributed Services Engine Ecosystem 125
4-69 Architecture Changes with vSphere Distributed Services Engine 126
4-70 vSphere Distributed Services Engine Hardware Components 127
4-71 Review of Learner Objectives 128
4-72 Key Points 128

Module 5 Storage Operations 129

5-2 Importance 129
5-3 Module Lessons 129
5-4 Lesson 1: NVMe Storage Systems 130
5-5 Learner Objectives 130
5-6 vSphere Support for NVMe 130
5-7 Supported NVMe Topologies 131
5-8 FC-NVMe and NVMe over RDMA (RoCE v2) 131
5-9 Support for FC-NVMe 132
5-10 Requirements for FC-NVMe 133
5-11 Support for NVMe over RDMA (RoCE v2) 134
5-12 Requirements for NVMe over RDMA (RoCE v2) 135
5-13 Configuring ESXi for NVMe over RDMA (RoCE v2) (1) 135
5-14 Configuring ESXi for NVMe over RDMA (RoCE v2) (2) 136
5-15 Configuring ESXi for NVMe over RDMA (RoCE v2) (3) 137
5-16 Configuring ESXi for NVMe over RDMA (RoCE v2) (4) 138
5-17 Pluggable Storage Architecture 139
5-18 High-Performance Plug-In for NVMe 140
5-19 Support for High-Performance Plug-In Multipathing 141
5-20 High-Performance Plug-In Best Practices 142
5-21 Review of Learner Objectives 142
5-22 Lesson 2: iSCSI Extension for RDMA (iSER) 143

5-23 Learner Objectives 143
5-24 About iSER 143
5-25 Configuring iSER 144
5-26 Requirements for iSER 145
5-27 Running iSER 146
5-28 Review of Learner Objectives 146
5-29 Lesson 3: VMFS Datastores 147
5-30 Learner Objectives 147
5-31 About VMFS Datastores 147
5-32 VMFS5 and VMFS6 Datastores 148
5-33 VMFS6 Datastores 148
5-34 Storage Device Formats 149
5-35 Storage Devices Supported by VMFS Datastores 150
5-36 Limitations of Using 4Kn Drives 150
5-37 Snapshot Formats for VMFS 151
5-38 Migrating Your Data from VMFS5 to VMFS6 152
5-39 Activity: VMFS6 Features (1) 152
5-40 Activity: VMFS6 Features (2) 153
5-41 Review of Learner Objectives 153
5-42 Lesson 4: vSAN Datastores 154
5-43 Learner Objectives 154
5-44 vSAN Storage Architecture 154
5-45 About vSAN Datastores 155
5-46 Disk Groups 156
5-47 vSAN Hardware Requirements 157
5-48 Viewing the vSAN Datastore Summary 158
5-49 Objects in vSAN Datastores 159
5-50 VM Storage Policies 160
5-51 Viewing VM Settings for vSAN Information 161
5-52 About vSAN Fault Domains 162
5-53 Fault Domain Configurations 163
5-54 Lab 10: Viewing a vSAN Datastore Configuration 163
5-55 Review of Learner Objectives 163
5-56 Lesson 5: vSphere Storage APIs - Array Integration 164
5-57 Learner Objectives 164
5-58 About vSphere Storage APIs - Array Integration 164

5-59 Hardware Acceleration APIs for Block Storage 165
5-60 Hardware Acceleration for NAS 166
5-61 Array Thin Provisioning APIs 167
5-62 About Space Reclamation (1) 168
5-63 About Space Reclamation (2) 169
5-64 Space Reclamation with VMFS 170
5-65 Performing Manual Space Reclamation on a VMFS5 Datastore 171
5-66 Enabling Automatic Space Reclamation on a VMFS6 Datastore 172
5-67 Configuring the Automatic Unmap Rate 173
5-68 Space Reclamation Support for Guest Operating Systems (1) 173
5-69 Space Reclamation Support for Guest Operating Systems (2) 174
5-70 Space Reclamation Support for Guest Operating Systems (3) 174
5-71 Extended Copy (XCOPY) 175
5-72 Activity: vSphere Storage APIs – Array Integration (1) 175
5-73 Activity: vSphere Storage APIs – Array Integration (2) 176
5-74 Review of Learner Objectives 176
5-75 Lesson 6: vSphere Storage APIs for Storage Awareness and vSphere APIs for I/O Filtering (VASA and VAIO) 177
5-76 Learner Objectives 177
5-77 About vSphere API for Storage Awareness 178
5-78 Benefits of Storage Providers 179
5-79 Storage Provider Categories (1) 179
5-80 Storage Provider Categories (2) 180
5-81 Registering a Storage Provider 181
5-82 About vSphere APIs for I/O Filtering 182
5-83 Types of I/O Filters 183
5-84 How I/O Filters Work (1) 184
5-85 How I/O Filters Work (2) 185
5-86 I/O Filter Providers 186
5-87 Displaying I/O Filter Provider Information 187
5-88 Review of Learner Objectives 187
5-89 Lesson 7: Storage Policy Based Management Overview 188
5-90 Learner Objectives 188
5-91 About Storage Policy Based Management (1) 188
5-92 About Storage Policy Based Management (2) 189
5-93 About VM Storage Policies 190

5-94 About Storage Policy Rules 191
5-95 Storage Policy Example (1) 192
5-96 Storage Policy Example (2) 193
5-97 About Storage Policy Components 194
5-98 Storage Policy Component Example 195
5-99 Activity: Storage Policy Rules (1) 195
5-100 Activity: Storage Policy Rules (2) 196
5-101 Review of Learner Objectives 196
5-102 Lesson 8: Using VM Storage Policies 197
5-103 Learner Objectives 197
5-104 Creating and Managing Storage Policies 197
5-105 Step 1: Configuring Storage to Be Used for Storage Policies 197
5-106 Step 2: Creating Storage Policy Components 198
5-107 Step 3: Creating VM Storage Policies 198
5-108 Example: Creating a Silver Tier Storage Policy (1) 199
5-109 Example: Creating a Silver Tier Storage Policy (2) 199
5-110 Example: Creating a Silver Tier Storage Policy (3) 200
5-111 Example: Creating a Silver Tier Storage Policy (4) 200
5-112 Step 4: Applying the Storage Policy to the VM 201
5-113 Step 5: Checking Compliance for the VM Storage Policy 202
5-114 Lab 11: Using Policy-Based Storage 203
5-115 Review of Learner Objectives 203
5-116 Lesson 9: Storage Policies for vSAN 204
5-117 Learner Objectives 204
5-118 vSAN Storage Providers 204
5-119 vSAN Storage Policies 205
5-120 Creating vSAN Storage Policies 206
5-121 vSAN Capabilities (1) 206
5-122 vSAN Capabilities (2) 207
5-123 Lab 12: Creating vSAN Storage Policies 207
5-124 Review of Learner Objectives 207
5-125 Lesson 10: Storage Policies for vSphere Virtual Volumes 208
5-126 Learner Objectives 208
5-127 About vSphere Virtual Volumes 208
5-128 vSphere Virtual Volumes Architecture 209
5-129 vSphere Virtual Volumes Storage Providers 210

5-130 Protocol Endpoints 211
5-131 Storage Containers and vSphere Virtual Volumes Datastores 212
5-132 Virtual Volumes 213
5-133 Mapping VM Files to Virtual Volumes 214
5-134 vSphere Virtual Volumes Replication 215
5-135 Virtual Volume Storage Policies 216
5-136 Virtual Volume Storage Policy: Examples 216
5-137 Creating Virtual Volumes Storage Policies (1) 217
5-138 Creating Virtual Volumes Storage Policies (2) 218
5-139 Activity: Displaying Storage Capabilities (1) 218
5-140 Activity: Displaying Storage Capabilities (2) 219
5-141 Review of Learner Objectives 219
5-142 Lesson 11: Storage I/O Control 220
5-143 Learner Objectives 220
5-144 About Storage I/O Control 220
5-145 Storage I/O Control Resource Shares and Limits 221
5-146 Enabling Storage I/O Control (1) 221
5-147 Enabling Storage I/O Control (2) 222
5-148 Monitoring Latency 223
5-149 Automatic Threshold Detection 224
5-150 Storage Policy Components for Storage I/O Control 225
5-151 Configuring Storage I/O Control 226
5-152 Storage I/O Control Requirements 227
5-153 Activity: Storage I/O Control (1) 228
5-154 Activity: Storage I/O Control (2) 229
5-155 Review of Learner Objectives 229
5-156 Lesson 12: Datastore Clusters and vSphere Storage DRS 230
5-157 Learner Objectives 230
5-158 About Datastore Clusters 231
5-159 Datastore Cluster Requirements 232
5-160 Host Cluster to Datastore Cluster Relationships 233
5-161 vSphere Storage DRS Overview 234
5-162 Initial Disk Placement 235
5-163 vSphere Storage DRS Affinity Rules 236
5-164 Storage Migration Recommendations 237
5-165 Datastore Correlation Detector 238

5-166 Configuring vSphere Storage DRS Runtime Settings 239
5-167 vSphere Storage DRS Maintenance Mode 240
5-168 Backups and vSphere Storage DRS 241
5-169 vSphere Storage DRS and Storage I/O Control 242
5-170 Activity: vSphere Storage DRS (1) 243
5-171 Activity: vSphere Storage DRS (2) 244
5-172 Review of Learner Objectives 245
5-173 Key Points 245

Module 6 vCenter and ESXi Operations 247

6-2 Importance 247
6-3 Module Lessons 247
6-4 Lesson 1: Backing Up and Restoring vCenter 248
6-5 Learner Objectives 248
6-6 About vCenter Backup and Restore 249
6-7 Method for vCenter Backup and Restore 250
6-8 Scheduling Backups 250
6-9 Viewing the Backup Schedule 251
6-10 File-Based Backup of vCenter 251
6-11 File-Based Restore of vCenter 252
6-12 Lab 13: Backing Up vCenter Appliance 253
6-13 Review of Learner Objectives 253
6-14 Lesson 2: vCenter High Availability 254
6-15 Learner Objectives 254
6-16 Importance of Keeping vCenter Highly Available 255
6-17 About vCenter High Availability 256
6-18 Scenario: Active Node Failure 257
6-19 Scenario: Passive Node Failure 258
6-20 Scenario: Witness Node Failure 259
6-21 Benefits of vCenter High Availability 260
6-22 vCenter High Availability Requirements 260
6-23 Demonstration: Configuring vCenter High Availability 261
6-24 Review of Learner Objectives 261
6-25 Lesson 3: vSphere Configuration Profiles 262
6-26 Learner Objectives 262
6-27 About vSphere Configuration Profiles (1) 262
6-28 About vSphere Configuration Profiles (2) 263

6-29 Requirements for vSphere Configuration Profiles 264
6-30 Limitations to vSphere Configuration Profiles 264
6-31 Cluster-Level Configuration 265
6-32 Using vSphere Configuration Profiles 266
6-33 View Host Settings in the Desired Configuration 267
6-34 Configuration Schema and the Configuration Document 268
6-35 Check Host Compliance Against the Desired Configuration for The Cluster 268
6-36 Run a Remediation Pre-Check 269
6-37 Remediate a Cluster Against the Desired Configuration 270
6-38 Consistent Configuration Across Multiple Sites 271
6-39 vSphere Distributed Switch Support 272
6-40 Transitioning a Cluster with a vSphere Distributed Switch 272
6-41 vSphere Configuration Profiles Service Health 273
6-42 About Host Profiles 274
6-43 Implementing Host Profiles 276
6-44 Monitoring for Compliance 277
6-45 Standardization Across Multiple vCenter Instances 278
6-46 Lab 14: Using vSphere Configuration Profiles 279
6-47 Review of Learner Objectives 279
6-48 Lesson 4: VMware Certificate Management 280
6-49 Learner Objectives 280
6-50 Certificate Use Case 280
6-51 About Certificate Authorities 281
6-52 About VMware CA 282
6-53 vSphere Certificate Types 283
6-54 Chain of Trust (1) 284
6-55 Chain of Trust (2) 285
6-56 Chain of Trust (3) 286
6-57 VMware CA Modes 287
6-58 VMware CA: Root CA Mode 288
6-59 Replacing the VMware CA Certificate in Root CA Mode (1) 288
6-60 Replacing the VMware CA Certificate in Root CA Mode (2) 289
6-61 VMware CA: Subordinate CA Mode 290
6-62 Replacing the VMware CA Certificate with a Subordinate CA Certificate 291
6-63 VMware CA: Hybrid Mode 291
6-64 Hybrid Mode: Replacing a vCenter Machine Server Certificate with a Custom Certificate 292

6-65 Certificate Replacement Options 292
6-66 Certificate Manager (1) 293
6-67 Certificate Manager (2) 294
6-68 Replacing All VMware CA Certificates with Custom Certificates 295
6-69 ESXi Certificate Expiration 296
6-70 ESXi Certificate Replacement Options 296
6-71 Renewing ESXi Certificates 297
6-72 Refreshing ESXi Certificates 297
6-73 VMware CA Availability 298
6-74 Lab 15: Working with Certificates 298
6-75 Review of Learner Objectives 298
6-76 Key Points 299

Module 7 vSphere Monitoring 301
7-2 Importance 301
7-3 Module Lessons 301
7-4 Lesson 1: Resource Monitoring Tools 302
7-5 Learner Objectives 302
7-6 Performance-Tuning Methodology 302
7-7 Resource-Monitoring Tools 303
7-8 Guest Operating System Monitoring Tools 304
7-9 Using Perfmon to Monitor VM Resources 305
7-10 Using esxtop to Monitor VM Resources 306
7-11 Monitoring Inventory Objects with Performance Charts 306
7-12 Working with Overview Performance Charts 307
7-13 Working with Advanced Performance Charts 308
7-14 Chart Options: Real-Time and Historical 309
7-15 Chart Types: Bar and Pie 310
7-16 Chart Types: Line 311
7-17 Chart Types: Stacked 312
7-18 Chart Types: Stacked Per VM 313
7-19 Saving Charts 314
7-20 About Objects and Counters 315
7-21 About Statistics Types 316
7-22 About Rollup 316
7-23 Rollup Types 317
7-24 Review of Learner Objectives 318

7-25 Lesson 2: Monitoring Resource Use 319
7-26 Learner Objectives 319
7-27 Interpreting Data from Tools 319
7-28 CPU-Constrained VMs (1) 320
7-29 CPU-Constrained VMs (2) 321
7-30 Memory-Constrained VMs (1) 322
7-31 Memory-Constrained VMs (2) 323
7-32 Memory-Constrained Hosts 324
7-33 Disk-Constrained VMs 325
7-34 Monitoring Disk Latency 326
7-35 Network-Constrained VMs 327
7-36 Lab 16: Monitoring Virtual Machine Performance 328
7-37 Review of Learner Objectives 328
7-38 Lesson 3: Monitoring vCenter and Performance 329
7-39 Learner Objectives 329
7-40 Monitoring vCenter Services 329
7-41 Monthly Patch Updates for vCenter 330
7-42 vCenter Performance Considerations 330
7-43 vCenter CPU Performance 331
7-44 vCenter Memory Performance 332
7-45 Monitoring CPU and Memory Use with the VAMI 333
7-46 Using vimtop to Monitor vCenter Appliance Resources 334
7-47 vimtop Commands and Columns 335
7-48 Monitoring CPU Use with vimtop 336
7-49 Activity: CPU and Memory Use (1) 336
7-50 Activity: CPU and Memory Use (2) 337
7-51 Monitoring Memory Use with vimtop 337
7-52 Viewing Service Memory Allocation with cloudvm-ram-size 338
7-53 Viewing Service Heap Size 339
7-54 Changing Service Heap Size 340
7-55 vCenter Network Performance 340
7-56 Monitoring Network Activity with the VAMI 341
7-57 Monitoring Network Activity with vimtop 342
7-58 vCenter Database Performance (1) 342
7-59 vCenter Database Performance (2) 343
7-60 Effects of Changing Statistics Level on Database Traffic (1) 344

7-61 Effects of Changing Statistics Level on Database Traffic (2) 345
7-62 vCenter Disk Use (1) 346
7-63 vCenter Disk Use (2) 347
7-64 Monitoring Disk Usage with the VAMI 348
7-65 Monitoring Disk Usage with the df Command 349
7-66 Monitoring Disk Activity with vimtop 350
7-67 Monitoring Disk usage with the Appliance Shell 351
7-68 Increase Disk Space for vCenter 352
7-69 Monitoring Database Activity with the VAMI 353
7-70 Track Your Progress Towards Sustainability Goals 353
7-71 NEW: Enhanced vSphere Green Metrics 354
7-72 Activity: Network Performance (1) 354
7-73 Activity: Network Performance (2) 355
7-74 Activity: vCenter Tools (1) 355
7-75 Activity: vCenter Tools (2) 356
7-76 Review of Learner Objectives 356
7-77 Lesson 4: Using Alarms 357
7-78 Learner Objectives 357
7-79 About Alarms 357
7-80 Predefined Alarms 358
7-81 Creating a Custom Alarm 359
7-82 Defining the Alarm Target Type 360
7-83 Defining the Alarm Rule: Condition-based Trigger 361
7-84 Defining the Alarm Rule: Event-based Trigger 362
7-85 Defining the Alarm Rule: Setting the Notification 363
7-86 Defining the Alarm Rule: Add Advanced Options 364
7-87 Defining the Alarm Reset Rules 364
7-88 Activating the Alarm 365
7-89 Triggered Alarms 365
7-90 Configuring vCenter Notifications 366
7-91 Lab 17: Using Alarms 367
7-92 Review of Learner Objectives 367
7-93 Lesson 5: Proactive Management with VMware Skyline 368
7-94 Learner Objectives 368
7-95 About VMware Skyline 368
7-96 VMware Skyline Family 369

7-97 Skyline Privacy and Security Policy 371
7-98 Skyline Health for vSphere (1) 371
7-99 Skyline Health for vSphere (2) 372
7-100 Skyline Health for vSAN 373
7-101 Skyline Components: Collector and Advisor 374
7-102 VMware Skyline Networking Design 375
7-103 How VMware Skyline Works 376
7-104 Skyline Advisor Pro 377
7-105 Skyline Advisor Pro: Proactive Findings 378
7-106 Log Assist 379
7-107 Skyline Features Comparison by Support Level 380
7-108 Integrated Skyline Health Diagnostics Plug-in 381
7-109 Benefits of Using VMware Skyline Health Diagnostics 381
7-110 VMware Skyline Health Diagnostics Architecture 382
7-111 Review of Learner Objectives 383
7-112 Key Points 383

Module 8 vSphere Security and Access Control 385

8-2 Importance 385
8-3 Module Lessons 385
8-4 Lesson 1: Securing vSphere 386
8-5 Learner Objectives 386
8-6 About Securing vSphere 386
8-7 About the vSphere Security Configuration Guide 387
8-8 Controlling Access to vCenter Systems 388
8-9 vCenter Access Control 389
8-10 Securing vCenter 390
8-11 Securing vSphere Management Networks 391
8-12 Securing ESXi Hosts 392
8-13 VM Protection Overview 393
8-14 Review of Learner Objectives 393
8-15 Lesson 2: Configuring ESXi Host Access and Authentication 394
8-16 Learner Objectives 394
8-17 About ESXi Host Security Properties 394
8-18 Configuring the ESXi Firewall 395
8-19 Configuring ESXi Services 396
8-20 About Lockdown Mode (1) 396

8-21 About Lockdown Mode (2) 397
8-22 About Lockdown Mode (3) 398
8-23 Normal Lockdown Mode Behavior 399
8-24 Strict Lockdown Mode Behavior 399
8-25 Integrating ESXi with Active Directory 400
8-26 VMware vSphere Authentication Proxy Service 401
8-27 Lab 18: Configuring Lockdown Mode 402
8-28 Review of Learner Objectives 402
8-29 Lesson 3: vSphere with Single Sign-On and LDAP Authentication 403
8-30 Learner Objectives 403
8-31 vCenter Single Sign-On 403
8-32 vCenter Single Sign-On Services 403
8-33 VMware Directory Service (VMDir) 404
8-34 VMware Certificate Authority (VMCA) 404
8-35 Lookup Service (lookupsvc) 405
8-36 Identity Management Service (vIDM) 405
8-37 VMware Authentication Framework Daemon (VMAFD) 406
8-38 About Identity Sources 407
8-39 LDAP Server Identity Source Settings 407
8-40 Identity Sources for vCenter with vCenter Single Sign-On 408
8-41 vCenter Two-Factor Authentication 409
8-42 Smart Card Authentication 409
8-43 Logging in with Smart Card Authentication 410
8-44 RSA SecurID Authentication 410
8-45 Review of Learner Objectives 410
8-46 Lesson 4: vCenter Identity Federation 411
8-47 Learner Objectives 411
8-48 About vCenter Identity Provider Federation for ADFS 411
8-49 vCenter Identity Provider Federation Use Cases 411
8-50 Prerequisites for vCenter Identity Provider Federation 412
8-51 vCenter Identity Provider Federation Components 412
8-52 User Login with the vCenter Built-In Identity Provider 413
8-53 vSphere Identity Sources with Multifactor Authentication (MFA) 414
8-54 Okta Federated Identity Management for vCenter 415
8-55 Okta Prerequisites 416
8-56 vCenter Prerequisites 416

8-57 Configure vCenter Identity Provider Federation for Okta (1) 417
8-58 Configure vCenter Identity Provider Federation for Okta (2) 417
8-59 Configure vCenter Identity Provider Federation for Okta (3) 418
8-60 Configure vCenter Identity Provider Federation for Okta (4) 418
8-61 Configure vCenter Identity Provider Federation for Okta (5) 419
8-62 Login Workflow of the vCenter Identity Provider Federation for ADFS 420
8-63 Configuring vCenter Identity Provider Federation in vSphere (1) 421
8-64 Configuring vCenter Identity Provider Federation in vSphere (2) 422
8-65 Configuring vCenter Identity Provider Federation in vSphere (3) 423
8-66 Configuring vCenter Identity Provider Federation in vSphere (4) 423
8-67 Enabling ADFS on an Existing Enhanced Linked Mode Configuration 424
8-68 Lab 19: (Simulation) Configuring Identity Federation to Use Microsoft ADFS 424
8-69 Review of Learner Objectives 425
8-70 Key Points 425

Module 9 vSphere Trusted Environments and VM Encryption 427

9-2 Importance 427
9-3 Module Lessons 427
9-4 Lesson 1: Booting ESXi Securely 428
9-5 Learner Objectives 428
9-6 UEFI Secure Boot for ESXi Hosts 428
9-7 UEFI Secure Boot Sequence for ESXi Hosts 429
9-8 Troubleshooting UEFI Secure Boot 430
9-9 UEFI Secure Boot Validation Script 430
9-10 TPM 2.0 on ESXi Hosts 431
9-11 ESXi Quick Boot Support on Servers with TPM 2.0 Processor 432
9-12 About Remote Attestation 433
9-13 Activity: ESXi Secure Boot (1) 434
9-14 Activity: ESXi Secure Boot (2) 434
9-15 Review of Learner Objectives 434
9-16 Lesson 2: Securing VMs 435
9-17 Learner Objectives 435
9-18 VM Protection with UEFI Secure Boot 435
9-19 Virtual TPM on ESXi Hosts 436
9-20 Using vTPM in a VM (1) 437
9-21 Using vTPM in a VM (2) 438
9-22 Prerequisites for Adding vTPM 439

9-23 vSphere Fault Tolerance with vTPM 440
9-24 Virtualization-Based Security for Windows Guest Operating Systems 441
9-25 VBS Requirements and Limitations 442
9-26 Enabling VBS 443
9-27 About SGX 444
9-28 About AMD SEV-ES 444
9-29 About Encrypted vSphere vMotion 445
9-30 Configuring Encrypted vSphere vMotion 446
9-31 Review of Learner Objectives 446
9-32 Lesson 3: VM Encryption 447
9-33 Learner Objectives 447
9-34 About VM Encryption 448
9-35 Overview of VM Encryption 449
9-36 VM Encryption Interoperability (1) 449
9-37 VM Encryption Interoperability (2) 451
9-38 VM Encryption Enhancements 452
9-39 Advantages of VM Encryption 452
9-40 Business Use Case: Securing VMs 453
9-41 VM Encryption Architecture 454
9-42 Supported vSphere KMS Servers 454
9-43 VM Encryption Prerequisites 455
9-44 Enabling Encryption During VM Creation 456
9-45 Encrypting Existing VMs 457
9-46 Unlocking Encrypted VMs (1) 457
9-47 Unlocking Encrypted VMs (2) 458
9-48 About the KMS 458
9-49 KMIP Client Certificate 459
9-50 Key Management Clusters 460
9-51 Summary of Keys 461
9-52 Role of vCenter in VM Encryption 462
9-53 KMS and vCenter Communication 463
9-54 Ensuring the ESXi Host is Cryptographically Safe 464
9-55 VM Encryption Process Flow (1) 465
9-56 VM Encryption Process Flow (2) 466
9-57 Activity: VM Keys (1) 466
9-58 Activity: VM Keys (2) 467

9-59 Managing VM Encryption 467
9-60 vCenter Role: No Cryptography Administrator 468
9-61 Activity: Cryptographic Privileges (1) 468
9-62 Activity: Cryptographic Privileges (2) 469
9-63 Cross vCenter Migration and Cloning of an Encrypted VM 469
9-64 Performing a Shallow Recrypt of VMs with Snapshots 470
9-65 Cloning an Encrypted VM with Instant Clone Technology 470
9-66 Cloning a VM: Encrypting or Decrypting the Destination VM 471
9-67 System Requirements When Cloning a VM 471
9-68 VM Encryption Operations When Cloning a VM 472
9-69 Backing Up Encrypted VMs (1) 473
9-70 Backing Up Encrypted VMs (2) 474
9-71 VM Encryption Events and Alarms 475
9-72 Review of Learner Objectives 475
9-73 Lesson 4: Key Providers 476
9-74 Learner Objectives 476
9-75 Overview of Standard Key Provider 476
9-76 Available Options 476
9-77 Implementation Considerations 477
9-78 Overview of vSphere Native Key Provider 477
9-79 Features of vSphere Native Key Provider 478
9-80 vSphere Native Key Provider Requirements 478
9-81 vSphere Native Key Provider vMotion and Cross-vCenter vMotion 478
9-82 vSphere Native Key Provider and Enhanced Linked Mode 479
9-83 Review of Learner Objectives 479
9-84 Lesson 5: vSphere Trust Authority 480
9-85 Learner Objectives 480
9-86 About vSphere Trust Authority 480
9-87 About vSphere Trust Authority Administrator 481
9-88 Use Case for vSphere Trust Authority 482
9-89 vSphere Trust Authority Components 483
9-90 ESXi Components 484
9-91 Attestation Service 485
9-92 Attestation Service Trust 486
9-93 Key Provider Service 486
9-94 Trusted Infrastructure Agent 487

9-95 vSphere Trust Authority Workflow 488
9-96 vCenter Topologies 489
9-97 About Enabling vSphere Trust Authority 489
9-98 About VM Encryption with vSphere Trust Authority 490
9-99 VM Encryption with a Trusted Key Provider 491
9-100 Identifying VM Encryption Key Providers 492
9-101 Migrating a VM Encrypted with a Trusted Key Provider (1) 493
9-102 Migrating a VM Encrypted with a Trusted Key Provider (2) 494
9-103 Lab 20: Configuring vCenter to work with an external KMS 494
9-104 Lab 21: Creating an Encrypted Virtual Machine 495
9-105 Review of Learner Objectives 495
9-106 Key Points 495

Module 1

Course Introduction

1-2 Course Introduction

1-3 Importance

As a vSphere administrator, you require skills to configure and maintain a highly available and scalable vSphere environment. By developing knowledge and skills in vSphere, you can operate, scale and secure ESXi hosts and vCenter instances in your environment.

1-4 Learner Objectives (1)

After completing this course, you should be able to meet the following objectives:

- Configure and manage a VMware Tools Repository
- Configure vSphere Replication and recover replicated VMs
- Manage VM resource usage with resource pools
- Configure and manage vSphere networking and storage for a large and sophisticated enterprise
- Configure vCenter High Availability
- Use vSphere Configuration Profiles to manage VMware ESXi host compliance
- Use the vSphere Client and the command line to manage certificates
- Monitor the performance of vCenter, ESXi, and VMs in the vSphere Client
- Secure vCenter, ESXi, and VMs in your vSphere environment

1-5 Learner Objectives (2)

After completing this course, you should be able to meet the following objectives:

- Use Identity Federation to configure the vCenter to use external identity sources
- Use VMware vSphere Trust Authority to secure the infrastructure for encrypted VMs

1-6 Course Outline

1. Course Introduction
2. Virtual Machines Operations
3. vSphere Cluster Operations
4. Network Operations
5. Storage Operations
6. vCenter and ESXi Operations
7. vSphere Monitoring
8. vSphere Security and Access Control
9. vSphere Trusted Environments and VM Encryption

1-7 Typographical Conventions

The following typographical conventions are used in this course.

Conventions	Usage and Examples
`Monospace`	Identifies command names, command options, parameters, code fragments, error messages, filenames, folder names, directory names, and path names: • Run the `esxtop` command. • ... found in the `/var/log/messages` file.
`Monospace Bold`	Identifies user inputs: • Enter **`ip a`**.
Boldface	Identifies user interface controls: • Click the **Configuration** tab.
Italic	Identifies book titles: • *vSphere Virtual Machine Administration*
< >	Indicates placeholder variables: • <ESXi_host_name> • ... the `Settings/<Your_Name>.txt` file

1-8 References (1)

Title	Location
vSphere ESXi Installation and Setup	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-801-installation-setup-guide.pdf
vCenter Installation and Setup	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-vcenter-801-installation-guide.pdf
vSphere Virtual Machine Administration	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-vm-administration/GUID-55238059-912E-411F-A0E9-A7A536972A91.html
VMware Configuration Maximums	https://configmax.vmware.com/home
VMware Compatibility Guide	http://www.vmware.com/resources/compatibility
vSphere Central	https://vspherecentral.vmware.com
Storage and Availability Technical Documents	https://storagehub.vmware.com

1-9 References (2)

Title	Location
vSphere Networking	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-801-networking-guide.pdf
vSphere Storage	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-801-storage-guide.pdf
vSphere Resource Management	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-801-resource-management-guide.pdf
vSphere Host Profiles	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-801-host-profiles-guide.pdf
vSphere Security	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-801-security-guide.pdf
vSphere Monitoring and Performance	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-801-monitoring-performance-guide.pdf

Technet24

1-10 VMware Online Resources

Documentation for vSphere: https://docs.vmware.com/

VMware Communities: http://communities.vmware.com

- Start a discussion.
- Access the knowledge base.
- Access documentation, technical papers, and compatibility guides.
- Access communities.
- Access user groups.

VMware Support: http://www.vmware.com/support

VMware Hands-on Labs: http://hol.vmware.com

VMware Blogs: http://blogs.vmware.com/

VMware Learning: http://www.vmware.com/learning

- Access course catalog and worldwide course schedule.
- Access worldwide course schedule.

VMware Education and Certification Blog: http://blogs.vmware.com/education

1-11 VMware Learning Overview

You can access the following Education Services:

- VMware Learning Paths:
  - Help you find the course that you need based on the product, your role, and your level of experience
  - Can be accessed at https://vmware.com/learning
- VMware Customer Connect Learning, which is the official source of digital training, includes the following options:
  - On-Demand Courses: Self-paced learning that combines lecture modules with hands-on practice labs
  - VMware Lab Connect: Self-paced, technical lab environment where you can practice skills learned during instructor-led training
  - Certification Exam Prep: Comprehensive video-based reviews of exam topics and objectives to help you take your certification exam Prep
- For more information, see https://vmware.com/learning/connect-learning.

VMware Customer Connect Learning offers different subscriptions. For a comparison of the subscriptions, go to https://vmware.com/learning/connect-learning and click **WHICH SUBSCRIPTION IS RIGHT FOR YOU?.**

Technet24

1-12 VMware Certification Overview

VMware certifications validate your expertise and recognize your technical knowledge and skills with VMware technology.

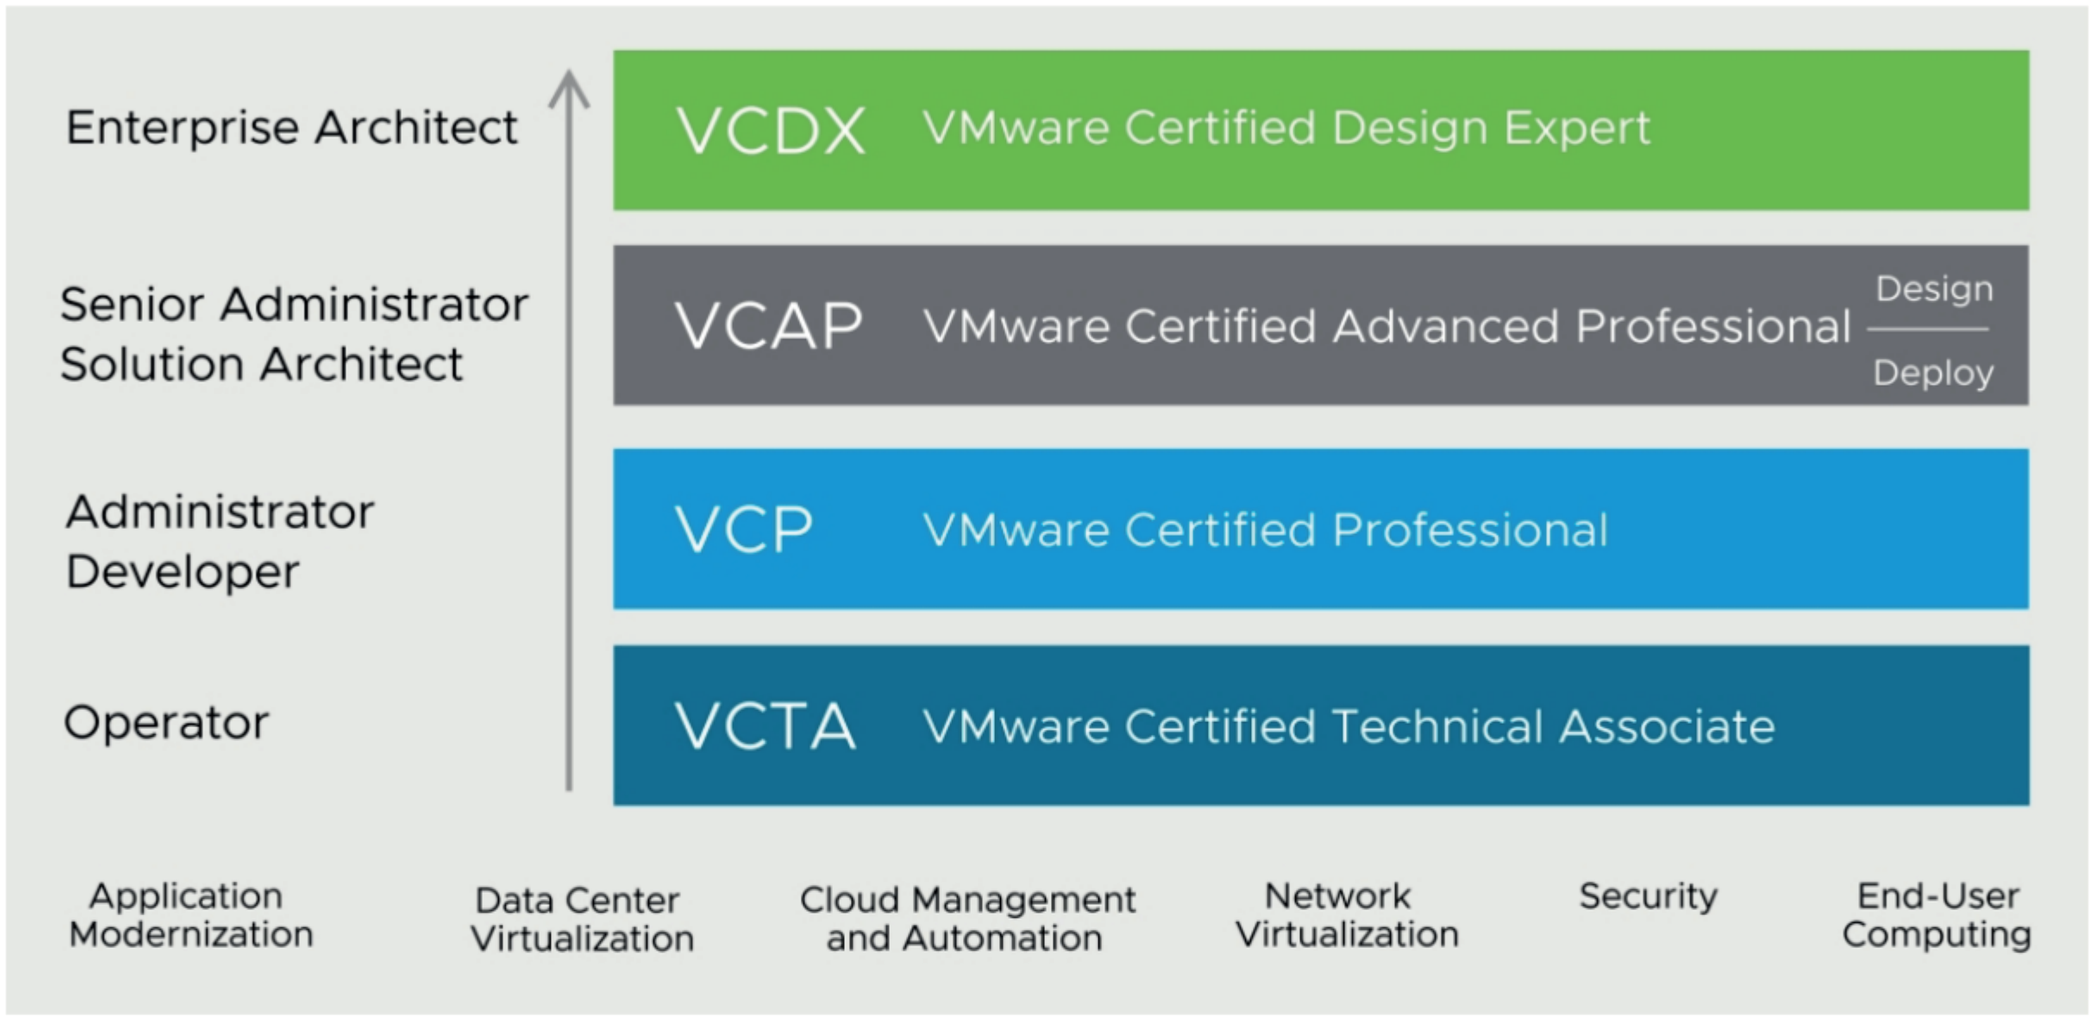

VMware certification sets the standards for IT professionals who work with VMware technology. Certifications are grouped into technology tracks. Each track offers one or more levels of certification (up to four levels).

For the complete list of certifications and details about how to attain these certifications, see https://vmware.com/certification.

1-13 VMware Credentials Overview

VMware badges are digital emblems of skills and achievements. Career certifications align to job roles and validate expertise across a solution domain. Certifications can cover multiple products in the same certification.

Specialist certifications and skills badges align to products and verticals and show expanded expertise.

Digital badges have the following features:

- Easy to share in social media (LinkedIn, Twitter, Facebook, blogs, and so on)
- Validate and verify achievement
- Contain metadata with skill tags and accomplishments
- Based on Mozilla's Open Badges standard

For the complete list of digital badges, see http://www.pearsonvue.com/vmware/badging.

Technet24

1-14 Lab 1: Accessing the Lab Environment

Access the lab environment and verify that vSphere licenses are valid:

1. Access Your Lab Environment
2. Verify That the vSphere Licenses Are Valid
3. (Optional) Assign Valid vSphere Licenses

Module 2

Virtual Machine Operations

2-2 Importance

Virtual machines are a fundamental part of your virtual infrastructure. Managing VMs effectively requires skills in protecting the VMs through replication and backups and resource management.

2-3 Module Lessons

1. VMware Tools Repository
2. VM Backups
3. vSphere Replication
4. VM Replication

Technet24

2-4 Lesson 1: VMware Tools Repository

2-5 Learner Objectives

- Recognize the role of a VMware Tools Repository
- Configure a VMware Tools Repository

2-6 Centralized Repository for VMware Tools

Using a centralized repository, provides the following benefits:

- Streamline the updating and installation of VMware Tools in your environment.
- Single location where VMware Tools is located, meaning you only have one location to update when you want to upgrade to the latest version of tools.
- Any ESXi host that has access to shared storage can be easily pointed to the central repository for a VMware Tools Update.
- Standardization of VMware Tools ISO version across multiple hosts.
- Faster consumption of the latest VMware Tools releases.
- Broader coverage of diverse guest operating systems.

Introduced in vSphere 6.7 Update 1, the `UpdateProductLockerLocation` is a vSphere API which allows repointing to a centralized repository without the need for entering maintenance mode and rebooting your ESXi host

2-7 Configuring a VMware Tools Repository

There are three steps required to implement a shared repository, also known as the "product locker":

1. Copy desired VMware Tools software bundles to a folder on an existing or a new shared datastore.
2. Update the `UserVars.ProductLockerLocation` on each VMware ESXi host.
3. Reboot or manually make the changed configuration effective by removing the `/productLocker` symbolic link (symlink) and recreating it in the vCenter MOB to point to the extracted directory under the datastore.

2-8 Configuring a VMware Tools Repository using the vSphere Client

In the vSphere Client you will need to update the value of `UserVars.ProductLockerLocation` from the ESXi Advanced System Settings with the desired datastore.

To update the new value of `UserVars.ProductLockerLocation` you will need to reboot the ESXi or manually set up `/productLocker` symlink.

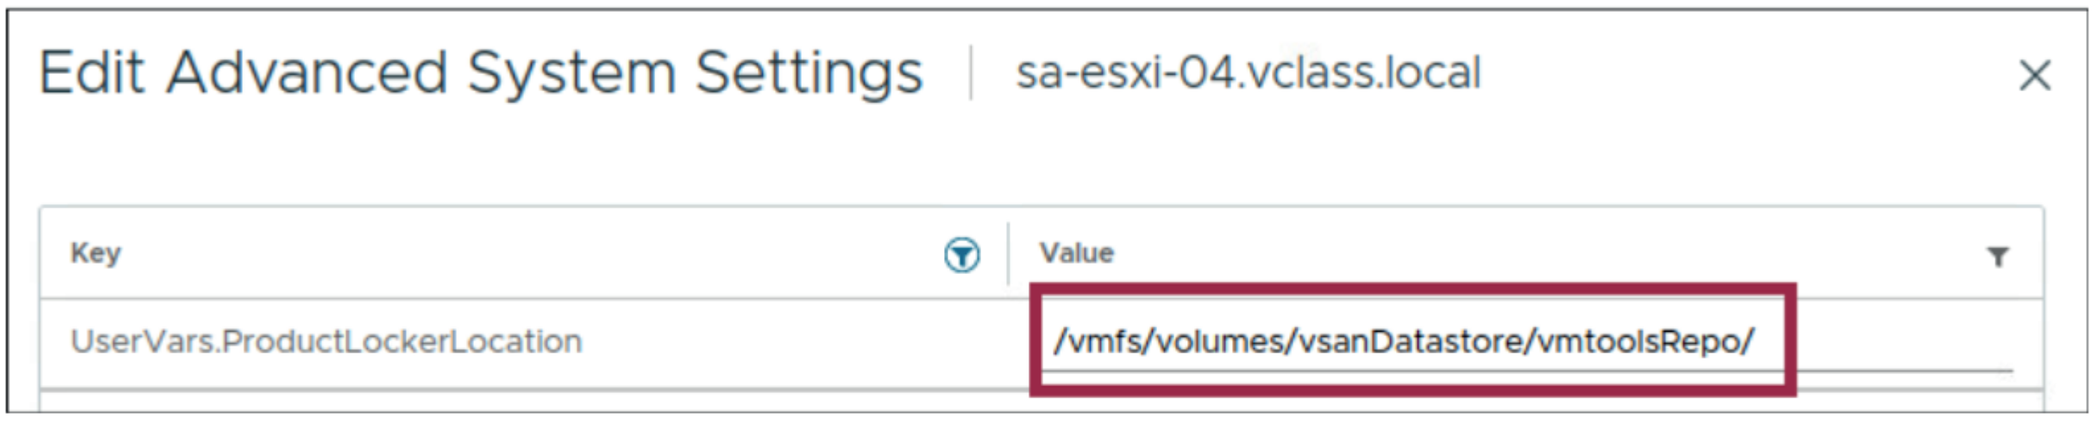

2-9 Configuring a VMware Tools Repository using the ESXi Shell

Using ESXi Shell, you can also change the `UserVars.ProductLockerLocation` variable to point it to the directory containing the VM Tools.

You will need to Reboot the ESXi host to update the new value of `UserVars.ProductLockerLocation` or manually set up `/productLocker` symlink.

```
[root@sa-esxi-04:~] esxcli system settings advanced set -o /UserVars/ProductLockerLocation -s /vmfs/volumes/vsanDatastore/vmtoolsRepo
[root@sa-esxi-04:~] esxcli system settings advanced list -o /UserVars/ProductLockerLocation
   Path: /UserVars/ProductLockerLocation
   Type: string
   Int Value: 0
   Default Int Value: 0
   Min Value: 0
   Max Value: 0
   String Value: /vmfs/volumes/vsanDatastore/vmtoolsRepo
   Default String Value: /locker/packages/vmtoolsRepo/
   Valid Characters: *
   Description: Path to VMware Tools repository
   Host Specific: true
   Impact:
[root@sa-esxi-04:~]
```


Technet24

2-10 Updating /productLocker Symlink (1)

To invoke the API, you can use the vCenter Managed Object Browser (MOB).

You can log in with the administrator@vsphere.local account, then browse to **content > rootFolder > childEntity > hostFolder**.

Using the MOB:

- Allows you to update the /productLocker on each VMware ESXi host by invoking `updateProductLockerLocation` vSphere API with the desired datastore as the argument
- This method does not require a reboot of the ESXi host.

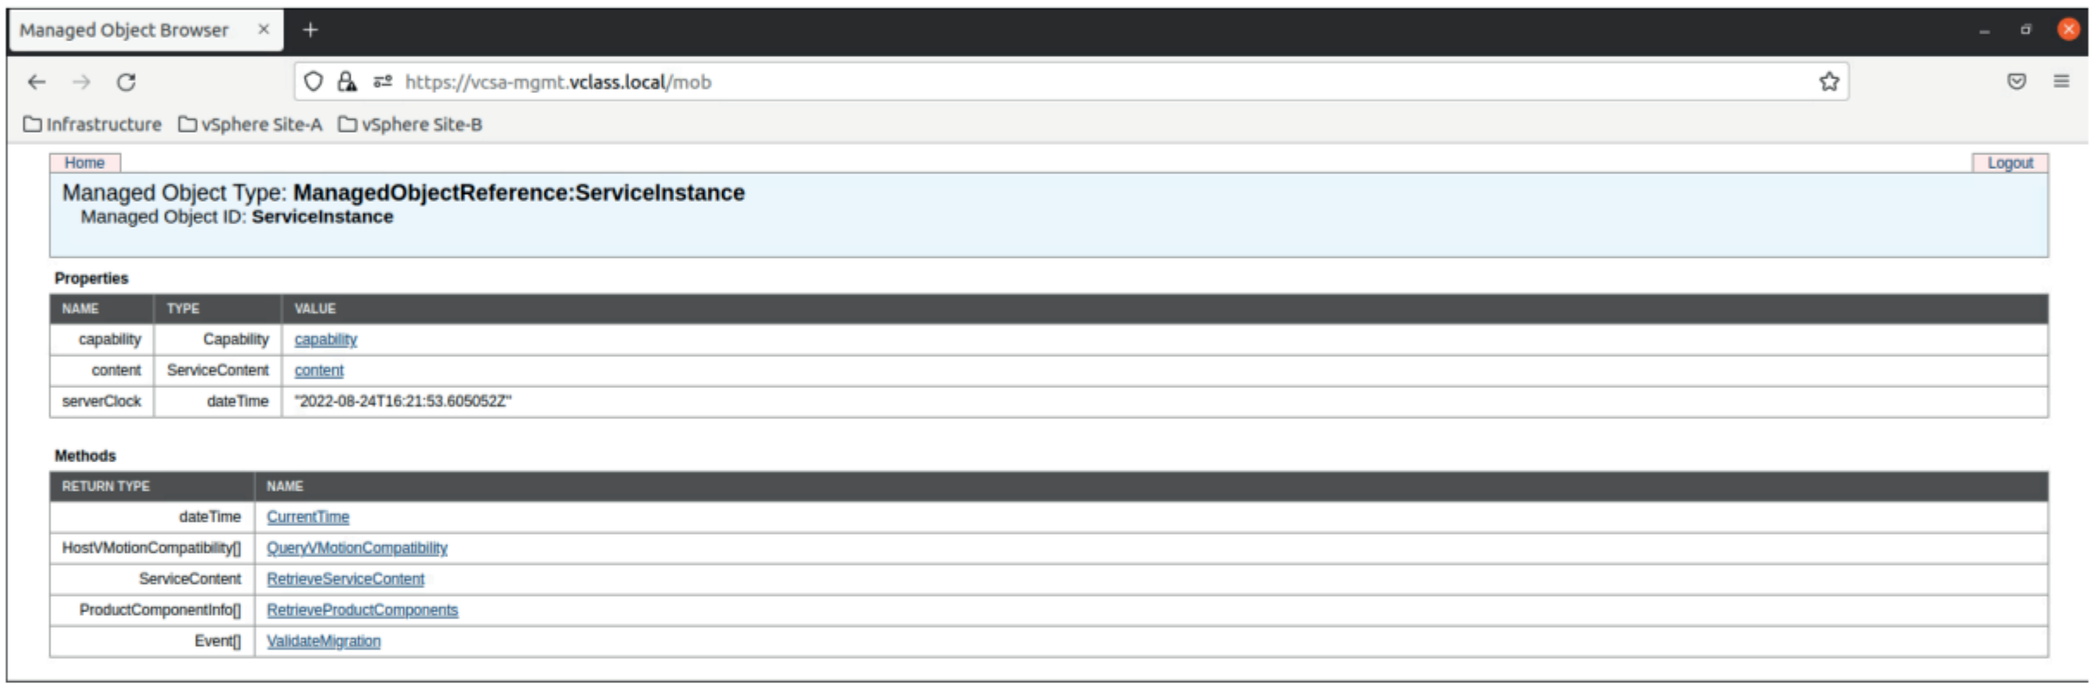

In vSphere 6.7 Update 1, updateProductLockerPolicy tool has been deprecated by the updateProductLockerLocation vSphere API.

2-11 Updating /productLocker Symlink (2)

You can open the API directly using the Host ID. For this, select the host in vSphere Client and the URL will include the host ID.

For example:

```
https://<vcenter_fqdn>/mob/?moid=host-
4007&method=updateProductLockerLocation
```

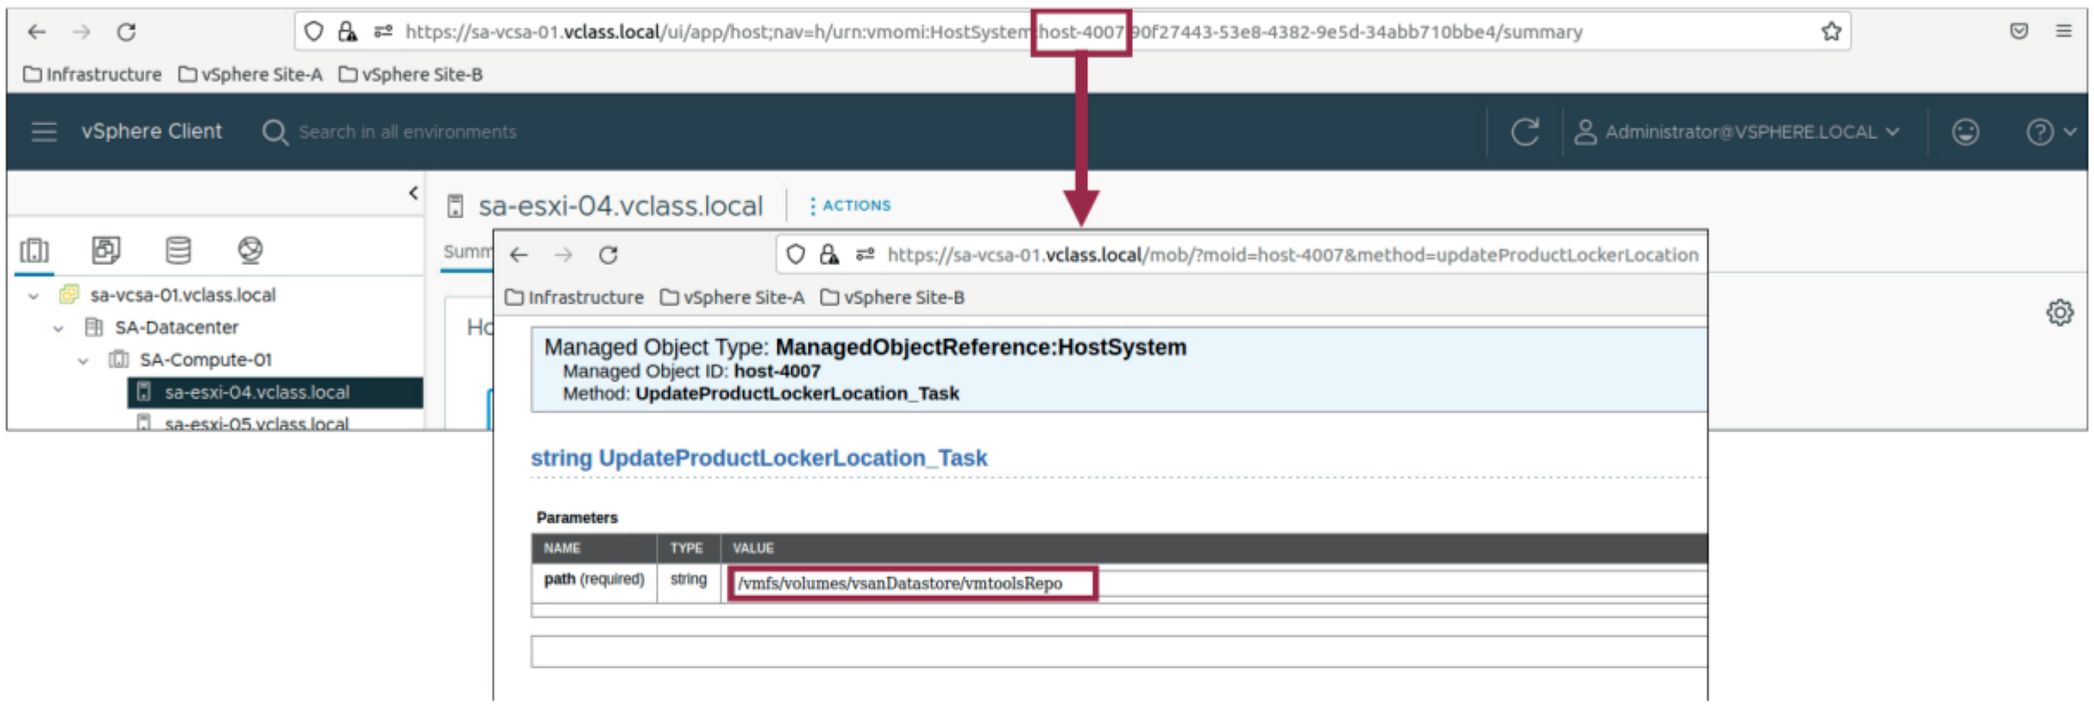

Technet24

2-12 Configuring a VMware Tools Repository using PowerCLI

A PowerCLI script can be used to update the productLocker symlink on all the hosts in a cluster.

```powershell
# Import PowerCLI module
Import-Module VMware.PowerCLI

# VMware vCenter name
$VCserver = read-host "Enter the vCenter name"

# Connect to the vCenter
Connect-VIServer -server $VCserver

$cluster = 'SA-Compute-01'
$hosts = Get-Cluster -Name $cluster | Get-VMHost
$datastore = '/vmfs/volumes/vsanDatastore/vmwtoolsRepo'

# Display current VMware Location
$hosts | Get-AdvancedSetting -Name "UserVars.ProductLockerLocation" | Select-Object Entity,Value

# Change VMware Tools location
Get-cluster -name $cluster | Get-VMhost | %{$_.ExtensionData.UpdateProductLockerLocation($datastore)}

# Display current VMware Location
$hosts | Get-AdvancedSetting -Name "UserVars.ProductLockerLocation" | Select-Object Entity,Value

# Disconnect vCenter
Disconnect-VIserver -server * -Confirm:$false
```

The following variables are unique to your vSphere environment:

- $cluster
- $datastore

Similar to manually updating the /productLocker symlink, this does not require a reboot of the ESXi host.

2-13 Lab 2: Configuring a Centralized VMware Tools Installation Repository

Create a shared VMware Tools repository:

1. Preconfigure the Environment
2. Copy the VMware Tools Packages to a Datastore
3. View the ESXi Advanced Settings
4. Re-create the `/productLocker` Symlink

2-14 Review of Learner Objectives

- Recognize the role of a VMware Tools Repository
- Configure a VMware Tools Repository

Technet24

2-15 Lesson 2: VM Backups

2-16 Learner Objectives

- Explain the backup and restore solution for VMs
- Describe the benefits of vSphere Storage APIs - Data Protection
- Recognize the supported Virtual Disk Transport Methods

2-17 Backup and Restore Solution for VMs

To protect your VM's data, you can use a backup solution based on vSphere Storage APIs - Data Protection.

With vSphere Storage APIs - Data Protection, backup products can perform centralized, efficient, off-host, LAN-free backups of vSphere VMs.

vSphere Storage APIs – Data Protection is VMware's data protection framework, which was introduced in vSphere 4.0. A backup product that uses this API can back up VMs from a central backup system (physical or virtual system). The backup does not require backup agents or any backup processing to be done inside the guest operating system.

Backup processing is offloaded from the ESXi host. In addition, vSphere snapshot capabilities are used to support backups across the SAN without requiring downtime for VMs. As a result, backups can be performed non-disruptively at any time of the day without requiring extended backup windows.

For frequently asked questions about vSphere Storage APIs - Data Protection, see VMware knowledge base article 1021175 at https://kb.vmware.com/s/article/1021175.

2-18 vSphere Storage APIs - Data Protection: Offloaded Backup Processing

Configure the storage environment so that the backup server can access the storage volumes that are managed by the ESXi hosts.

Backup processing is offloaded from the ESXi host to the backup server, which prevents local ESXi resources from becoming overloaded.

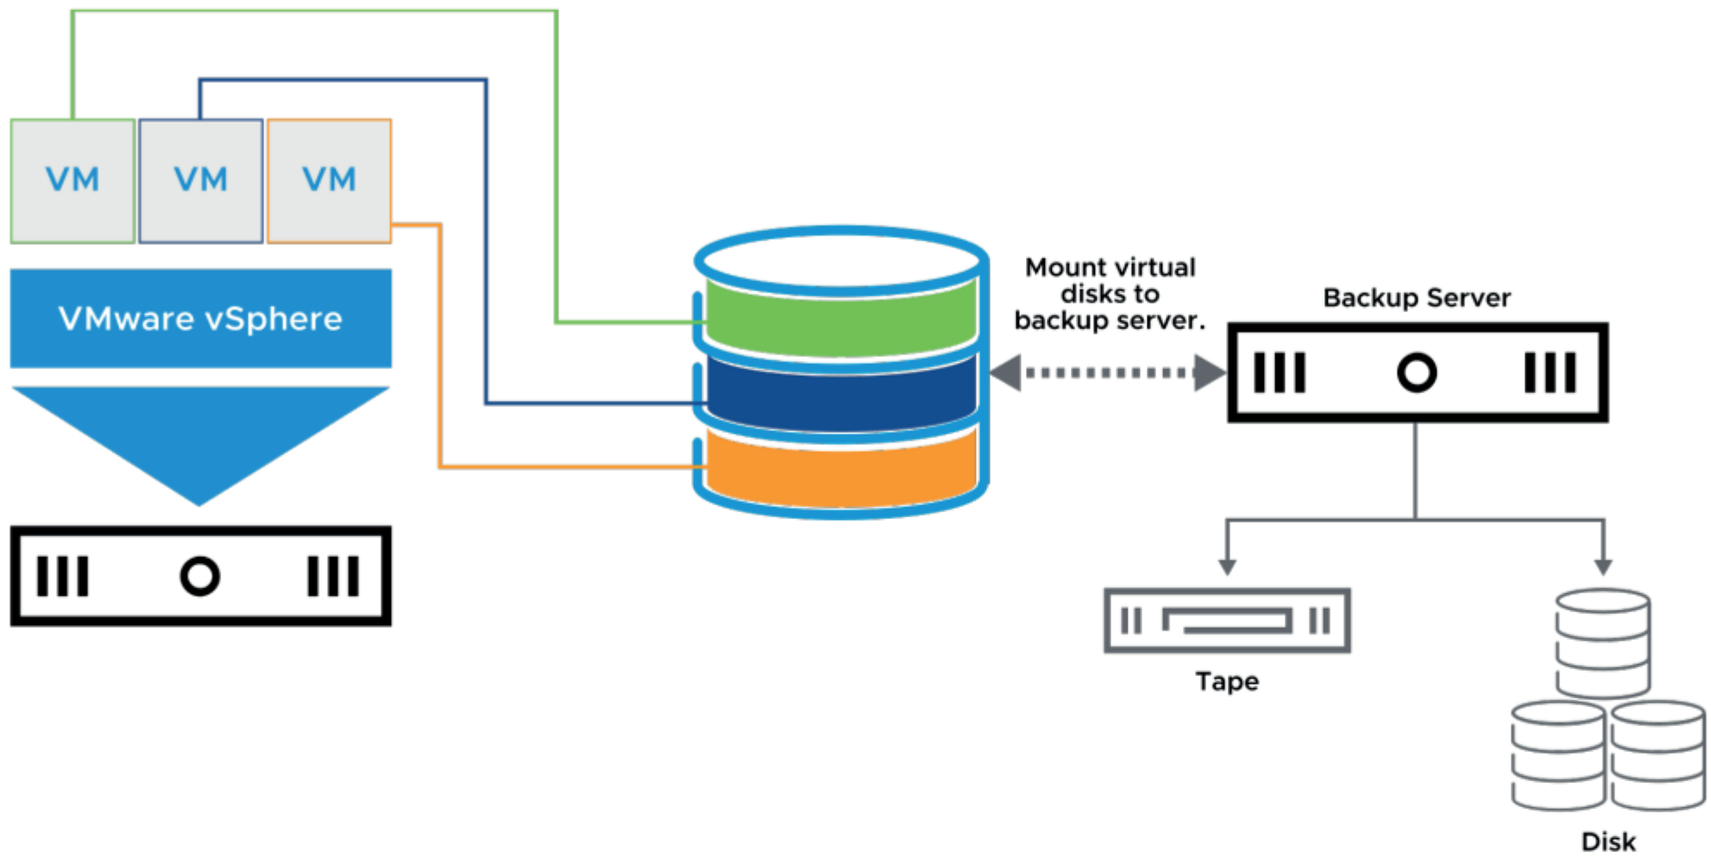

One of the biggest bottlenecks that limits backup performance is the backup server that is handling all the backup coordination tasks. One of these backup tasks is copying data from point A to point B. Other backup tasks do much CPU processing. For example, tasks are performed to determine what data to back up and what not to back up. Other tasks are performed to deduplicate data and compress data that is written to the target.

A server with insufficient CPU resources can greatly reduce backup performance. Provide enough resources for your backup server. A physical server or VM with an ample amount of memory and CPU capacity is necessary for the best backup performance possible.

The motivation to use LAN-free backups is to reduce the stress on the physical resources of the ESXi host when VMs are backed up. LAN-free backups reduce the stress by offloading backup processing from the ESXi host to a backup proxy server.

You can configure your environment for LAN-free backups to the backup server, also called the backup proxy server. For LAN-free backups, the backup server must be able to access the storage managed by the ESXi hosts on which the VMs for backup are running.

If you use NAS or direct-attached storage, ensure that the backup proxy server accesses the volumes with a network-based transport. If you run a direct SAN backup, zone the SAN and configure the disk subsystem host mappings. The host mappings must be configured so that all ESXi hosts and the backup proxy server access the same disk volumes.

Technet24

2-19 vSphere Storage APIs - Data Protection: Changed-Block Tracking

With changed-block tracking, the backup solution copies only file blocks that changed since the last backup.

Changed-block tracking supports faster incremental backups.

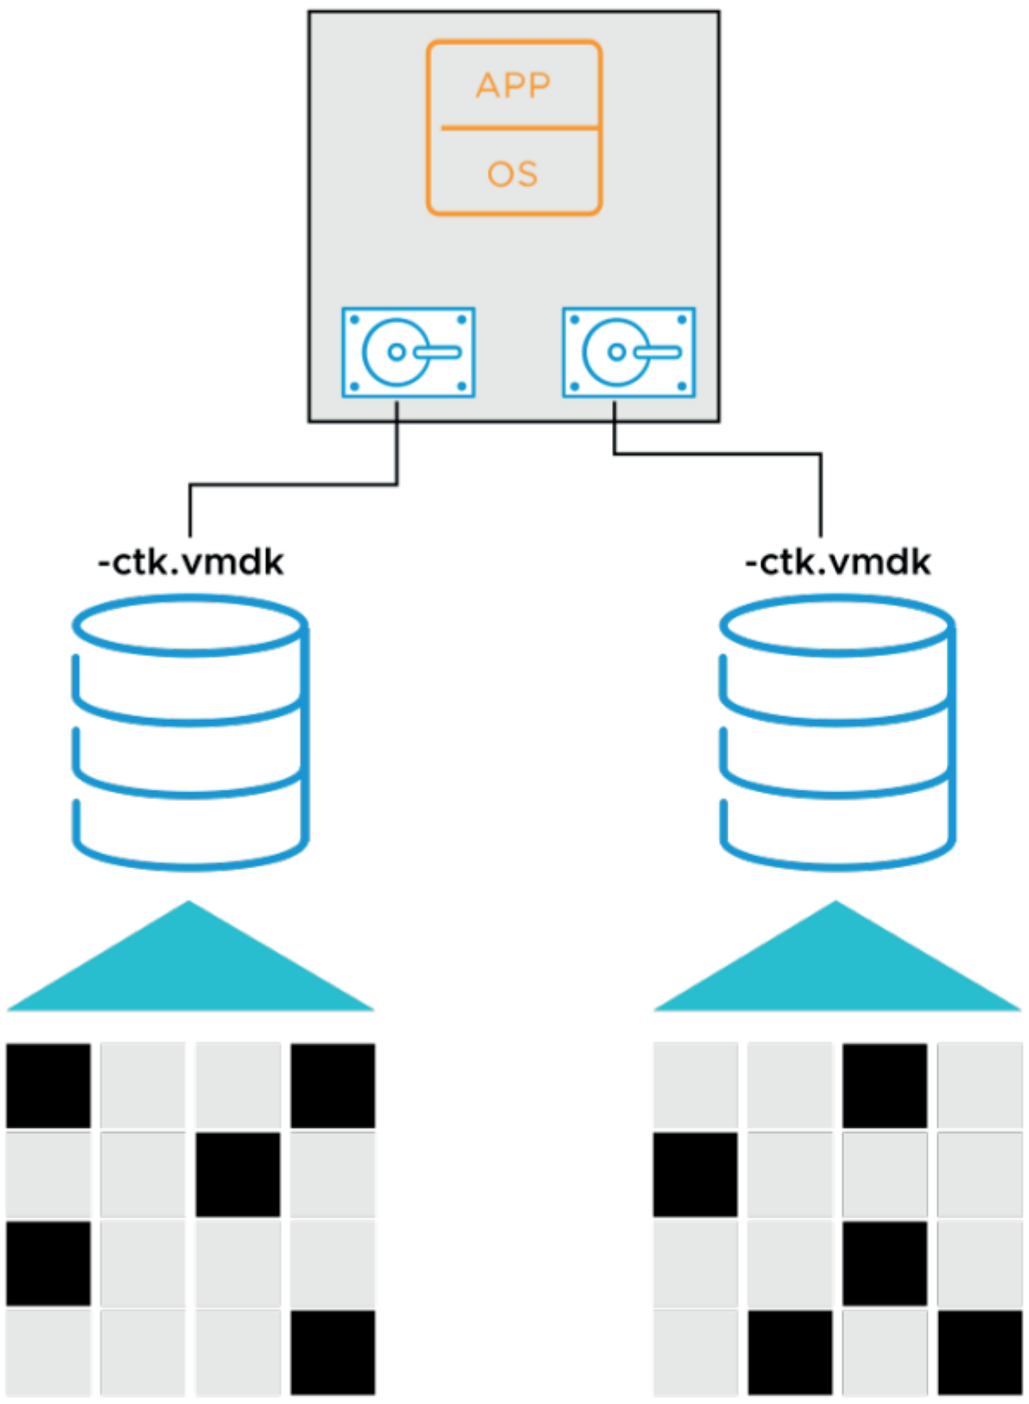

Changed-block tracking (CBT) is a VMkernel feature that tracks the storage blocks of VMs as they change over time. The VMkernel tracks block changes on VMs, enhancing the backup process for applications that are developed to exploit vSphere Storage APIs - Data Protection.

By using CBT during restores, vSphere Data Protection offers fast and efficient recoveries of VMs to their original location. During a restore process, the backup solution uses CBT to determine which blocks changed since the last backup. The use of CBT reduces data transfer within the vSphere environment during a recovery operation and, more important, reduces the recovery time.

2-20 Virtual Disk Transport Methods

Virtual Disk Transport Methods (VADP) offers interfaces for integration of storage-aware applications, including backup, with efficient access to storage clusters.

Backup vendors may use different transport methods, which provide efficient I/O methods to maximize backup performance and flexible deployment options.

VADP supports four access methods:

- SAN
- SCSI HotAdd
- NBD (network block device) over LAN with SSL encryption (NBDSSL)
- Local File Access

2-21 SAN Transport Mode

SAN mode requires applications to run on a backup server with access to SAN storage (Fibre Channel, iSCSI, or SAS connected) containing the virtual disks to be accessed.

The SAN backup server must be a physical machine.

If it has optical media or tape drive connected, backups can be made entirely LAN-free.

SAN mode is not supported for backup or restore of virtual machines:

- Residing on vSAN datastores
- Residing on Virtual Volume datastores
- Configured with VM encryption

Technet24

2-22 HotAdd Transport Mode (1)

Backup software running in a virtual appliance can utilize the HotAdd transport mode in vSphere.

It can take a snapshot and create a linked clone of the target virtual machine, then attach and read the linked clone's virtual disks for backup.

Virtual disks of the linked clone are HotAdded to the backup proxy while both VMs are on the same ESXi host.

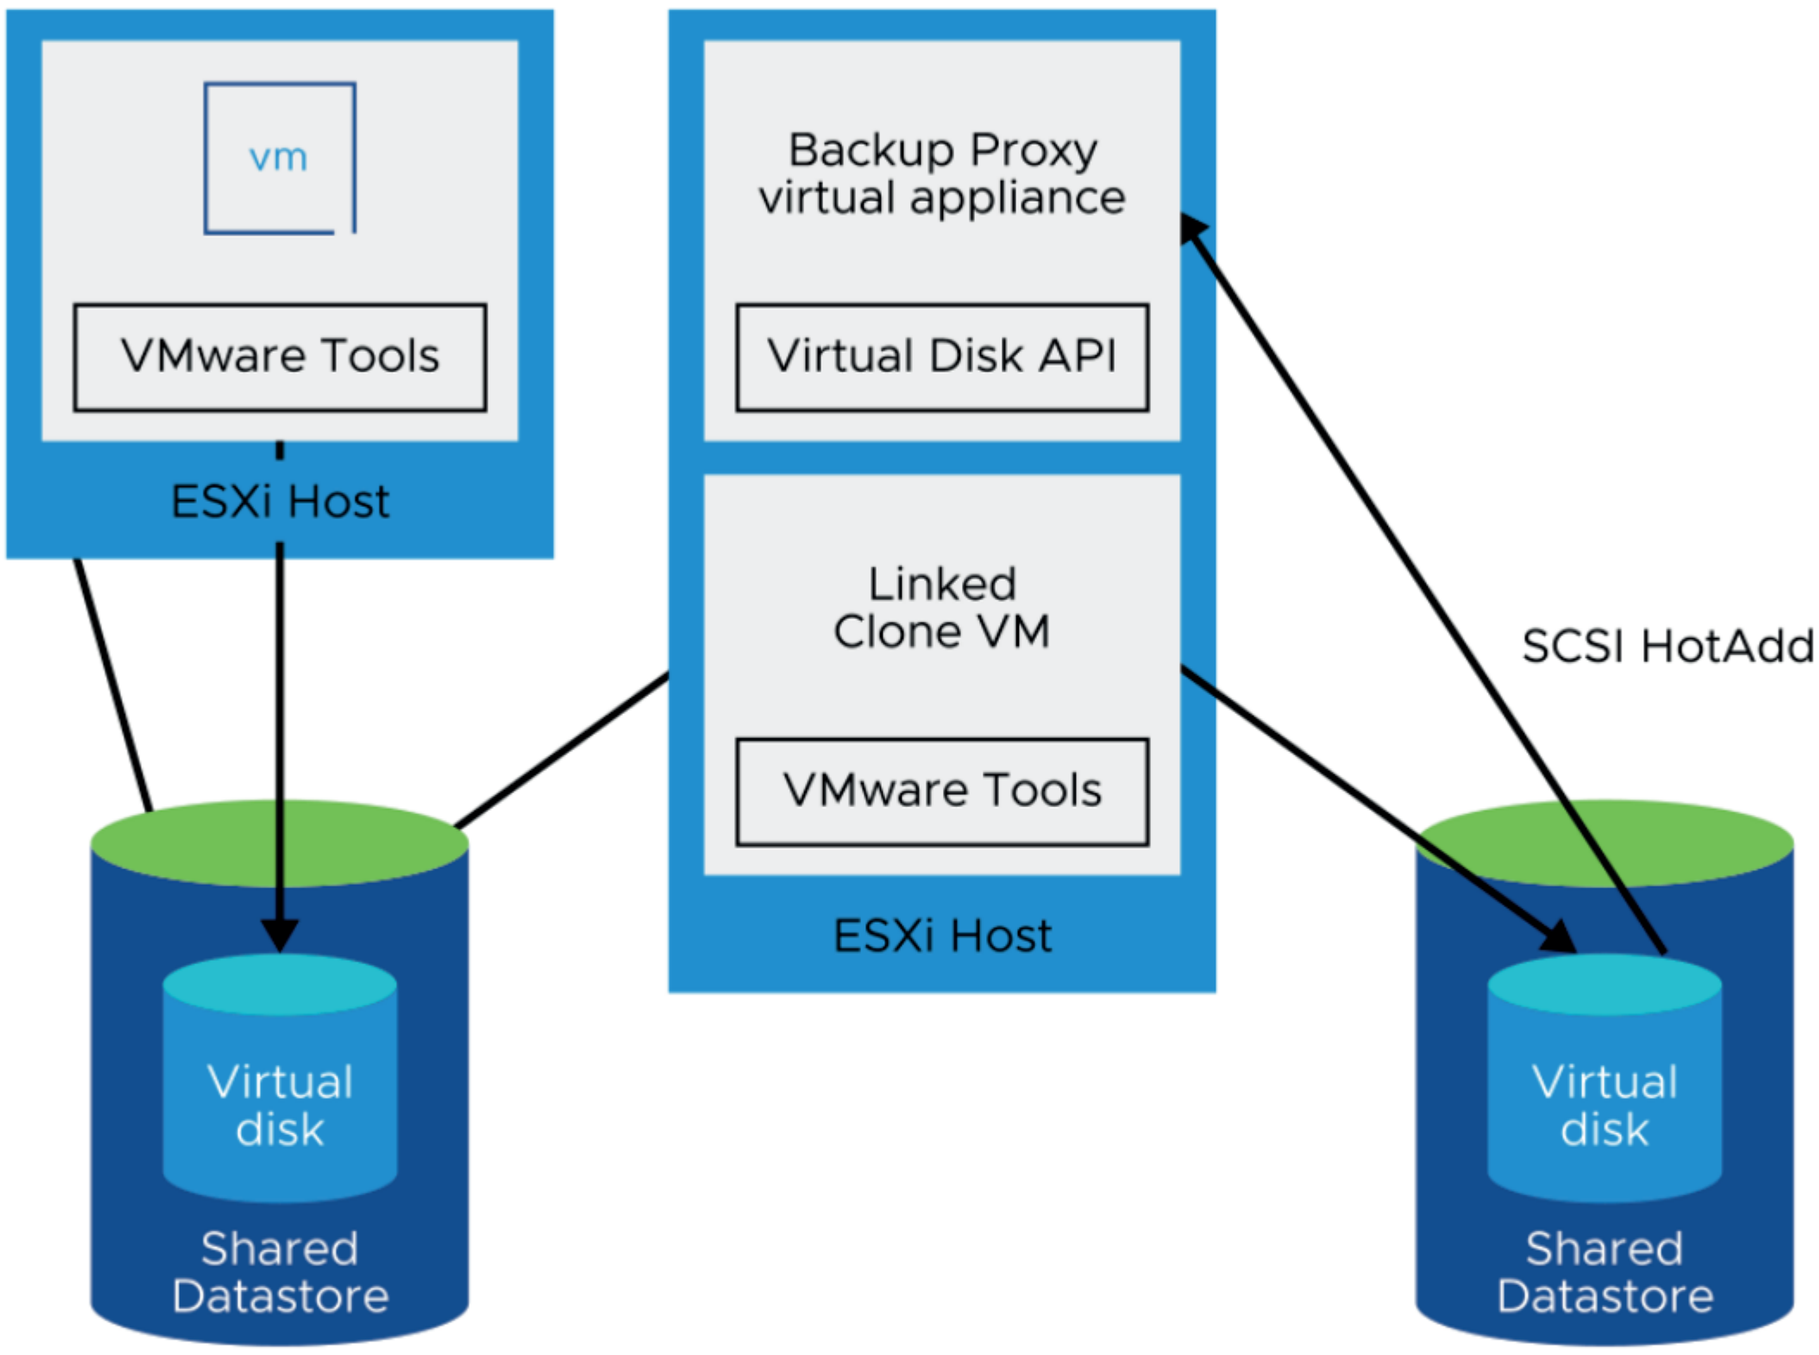

2-23 HotAdd Transport Mode (2)

Benefits of HotAdd Transport Mode include:

- It is easy to move a virtual machine to a new media server
- It can back up local storage without using the LAN

A HotAdd proxy backs up either Windows or Linux virtual machines, but not both. For parallel backup, sites can deploy multiple proxies.

2-24 NBDSSL Transport Mode (1)

When no other transport is available, networked storage applications can use LAN transport for data access, with NBD (network block device) protocol and SSL encryption, called NBDSSL.

NBD is a Linux-style kernel module that treats storage on a remote host as a block device.

NBDSSL is the VMware variant that uses SSL to encrypt all data passed over the TCP connection.

The NBDSSL transport method is built into the virtual disk library, so it is always available, and is the fall-back when no other transport method is available.

In this mode, the ESXi host reads data from storage and sends it across a network to the backup server.

With LAN transport, large virtual disks can take a long time to transmit.

2-25 NBDSSL Transport Mode (2)

NBDSSL transport mode adds traffic to the LAN, unlike SAN and HotAdd transport, but NBDSSL transport offers the following advantages:

- The ESXi host can use any storage device, including local storage or remote-mounted NAS.
- The backup proxy can be a virtual machine, so customers can use vSphere resource pools to minimize the performance impact of backup.

  For example, the backup proxy can be in a lower-priority resource pool than the production VMs.

Technet24

2-26 Local File Access

The virtual disk library can read virtual disk data from `/vmfs/volumes` on ESXi hosts, or from the local file system on hosted products.

The local file access method is built into the virtual disk library, so it is always available on local storage.

This is not a network transport method, and is only used for vSphere backup in few instances.

2-27 Review of Learner Objectives

- Explain the backup and restore solution for VMs
- Describe the benefits of vSphere Storage APIs - Data Protection
- Recognize the supported Virtual Disk Transport Methods

2-28 Lesson 3: vSphere Replication

2-29 Learner Objectives

- Identify the components in the vSphere Replication architecture
- Deploy and configure vSphere Replication and Site Recover instances

2-30 About vSphere Replication

vSphere Replication is an extension for vCenter.

It provides hypervisor-based VM replication and recovery.

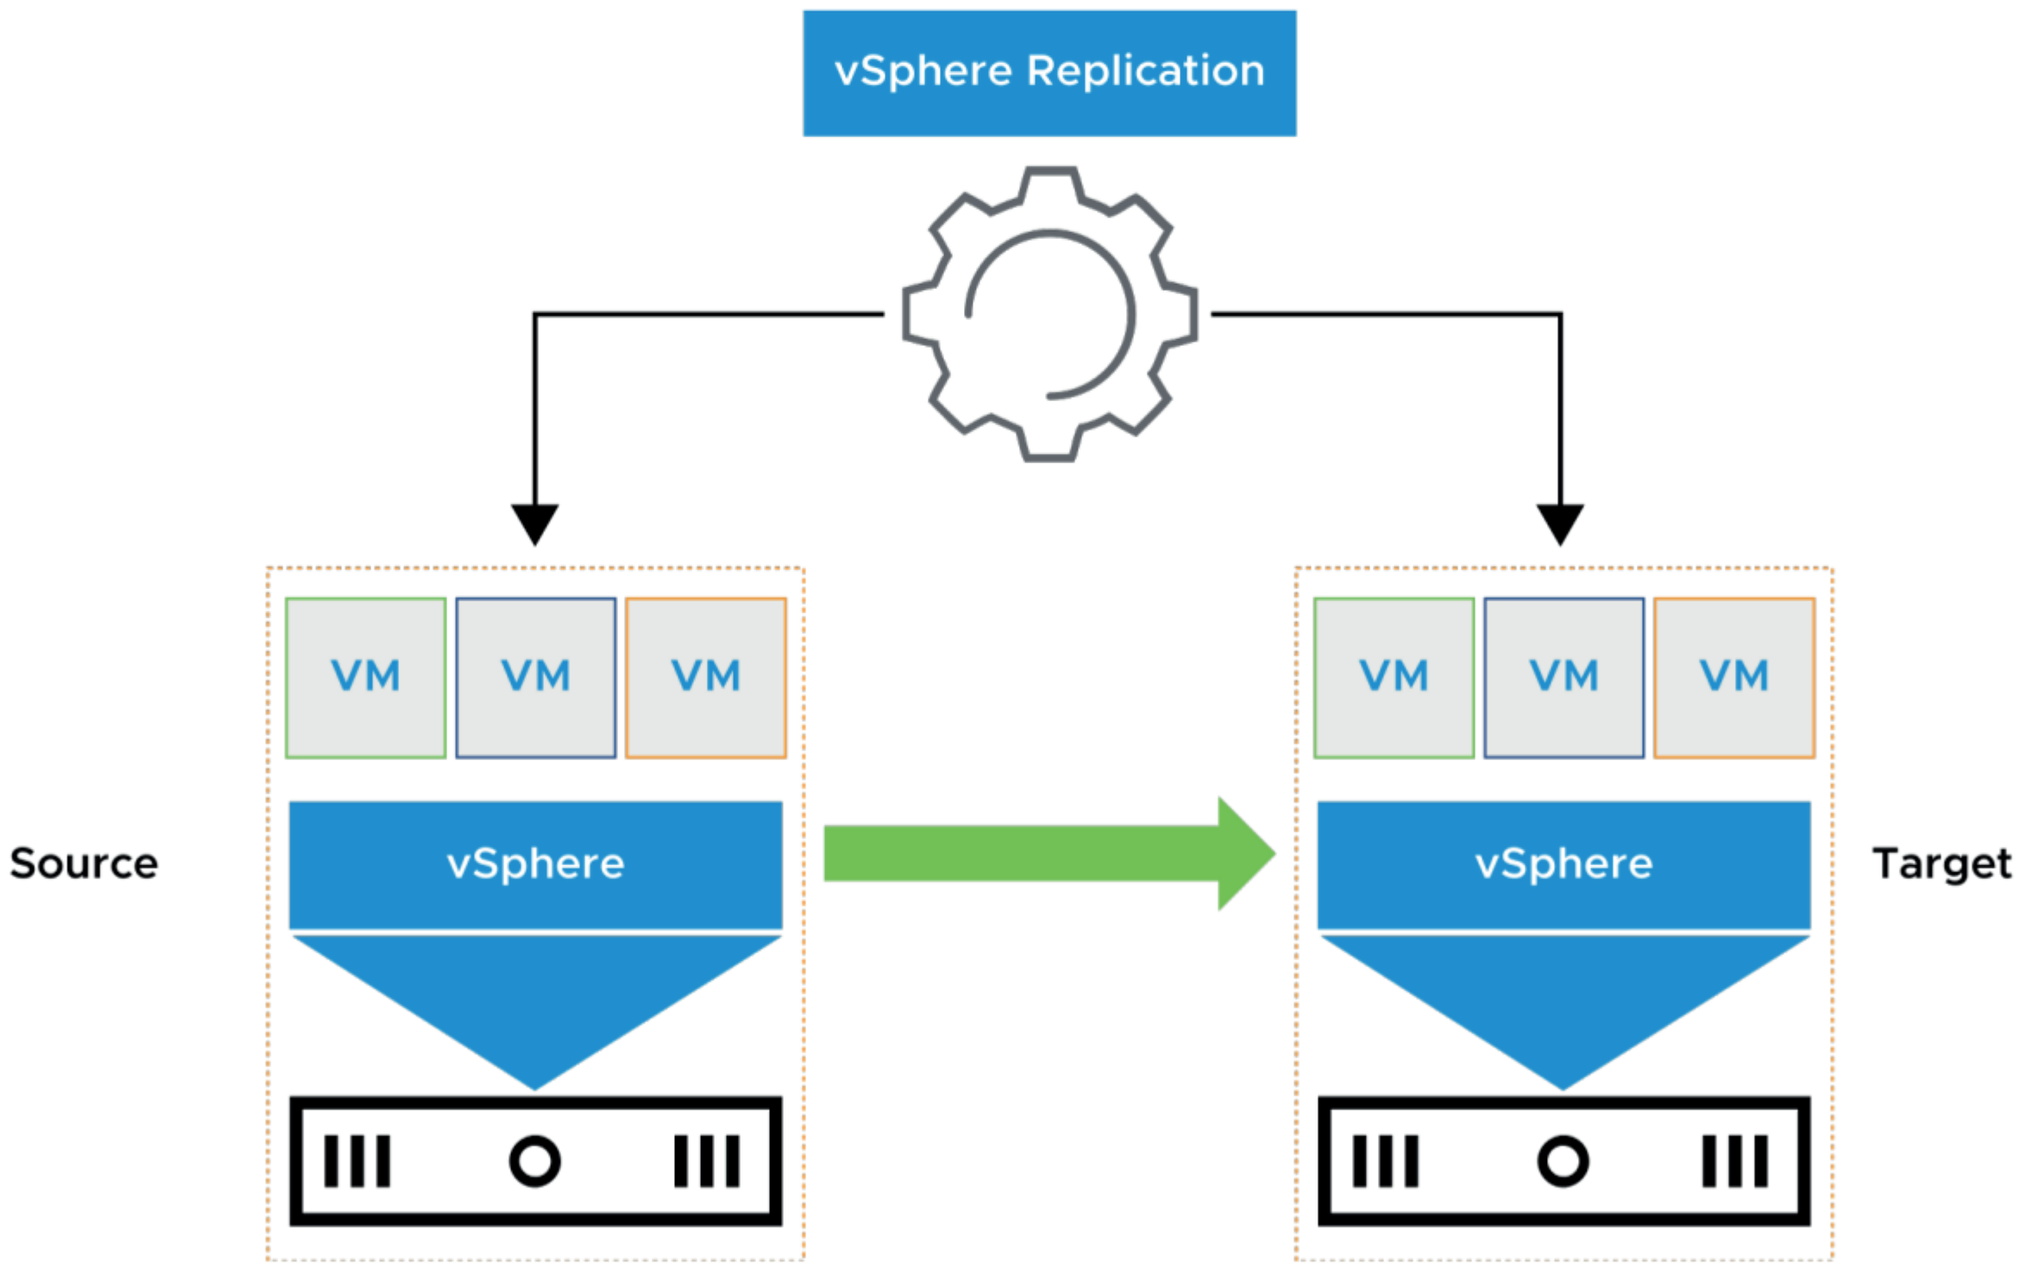

Technet24

2-31 vSphere Replication Benefits

vSphere Replication provides several benefits as compared to array-based replication:

- Data protection at lower cost per VM
- A replication solution that supports flexibility in storage vendor selection at the source and target sites
- Overall lower cost per replication

2-32 About the vSphere Replication Appliance

The vSphere Replication appliance provides all the components that are required to perform VM replication.

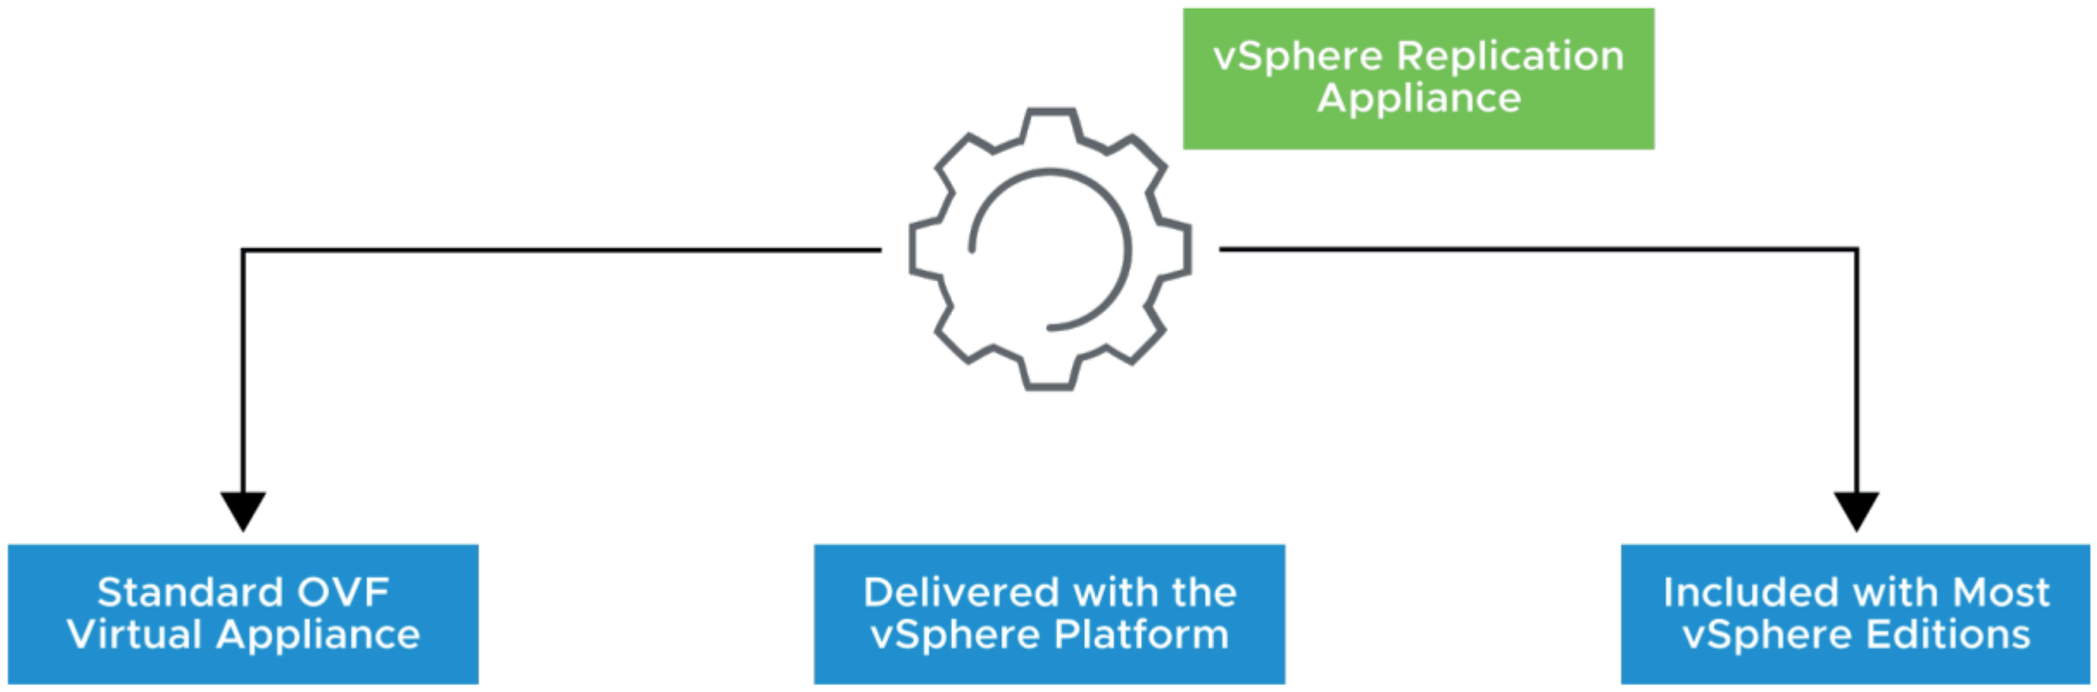

2-33 vSphere Replication Appliance Components

The vSphere Replication appliance includes the following components.

vSphere Replication server

- Provides the core of the vSphere Replication infrastructure

Embedded Database

- Stores replication configuration and management information

vSphere Replication Management Server

- Configures the vSphere Replication
- Activates, manages and monitors replications
- Authenticates users and checks their permissions to perform vSphere Replication operations

vSphere Client Plug-in

- Provides a user interface for vSphere Replication

2-34 vSphere Replication Management Service (VRMS) Appliance Management Interface (VAMI)

You can use vSphere Replication immediately after you deploy the appliance.

The vSphere Replication appliance provides the vSphere Replication Management Service (VRMS) appliance management interface that is used to reconfigure the appliance after deployment.

You can use the vSphere Replication Management Service (VRMS) Appliance Management Interface (VAMI) to:

- Make changes to the appliance security settings
- Make changes to the network settings
- Deploy additional vSphere Replication servers by using a separate Open Virtualization Format (OVF) package

OVF is an open source standard for packaging and distributing software applications and services for virtual machines (VMs).

Technet24

2-35 Replication Functions

With vSphere Replication, you can replicate a VM from a source site to a target site, monitor and manage the replication status, and recover the VM at the target site.

As administrator, you can manage the type of replication:

- Replication Between Two Sites
- Replication in a Single vCenter
- Replication to a Shared Target Site

You can replicate a VM between two sites. vSphere Replication is installed on both source and target sites. Only one vSphere Replication appliance is deployed on each vCenter. The vSphere Replication (VR) appliance contains an embedded vSphere Replication server that manages the replication process. To meet the load-balancing needs of your environment, you might need to deploy additional vSphere Replication servers at each site.

When you configure a VM for replication, the vSphere Replication agent sends changed blocks in the VM disks from the source site to the target site according to the RPO. The changed blocks are applied to the copy of the VM. This process occurs independently of the storage layer. vSphere Replication performs an initial full synchronization of the source VM and its replica copy. You can use replication seeds to reduce the network traffic that is generated by data transfer during the initial full synchronization.

2-36 Replication Between Two Sites

You can replicate a virtual machine between two sites.

vSphere Replication is installed on both source and target sites.

Only one vSphere Replication appliance is deployed on each vCenter.

You can deploy additional vSphere Replication Servers

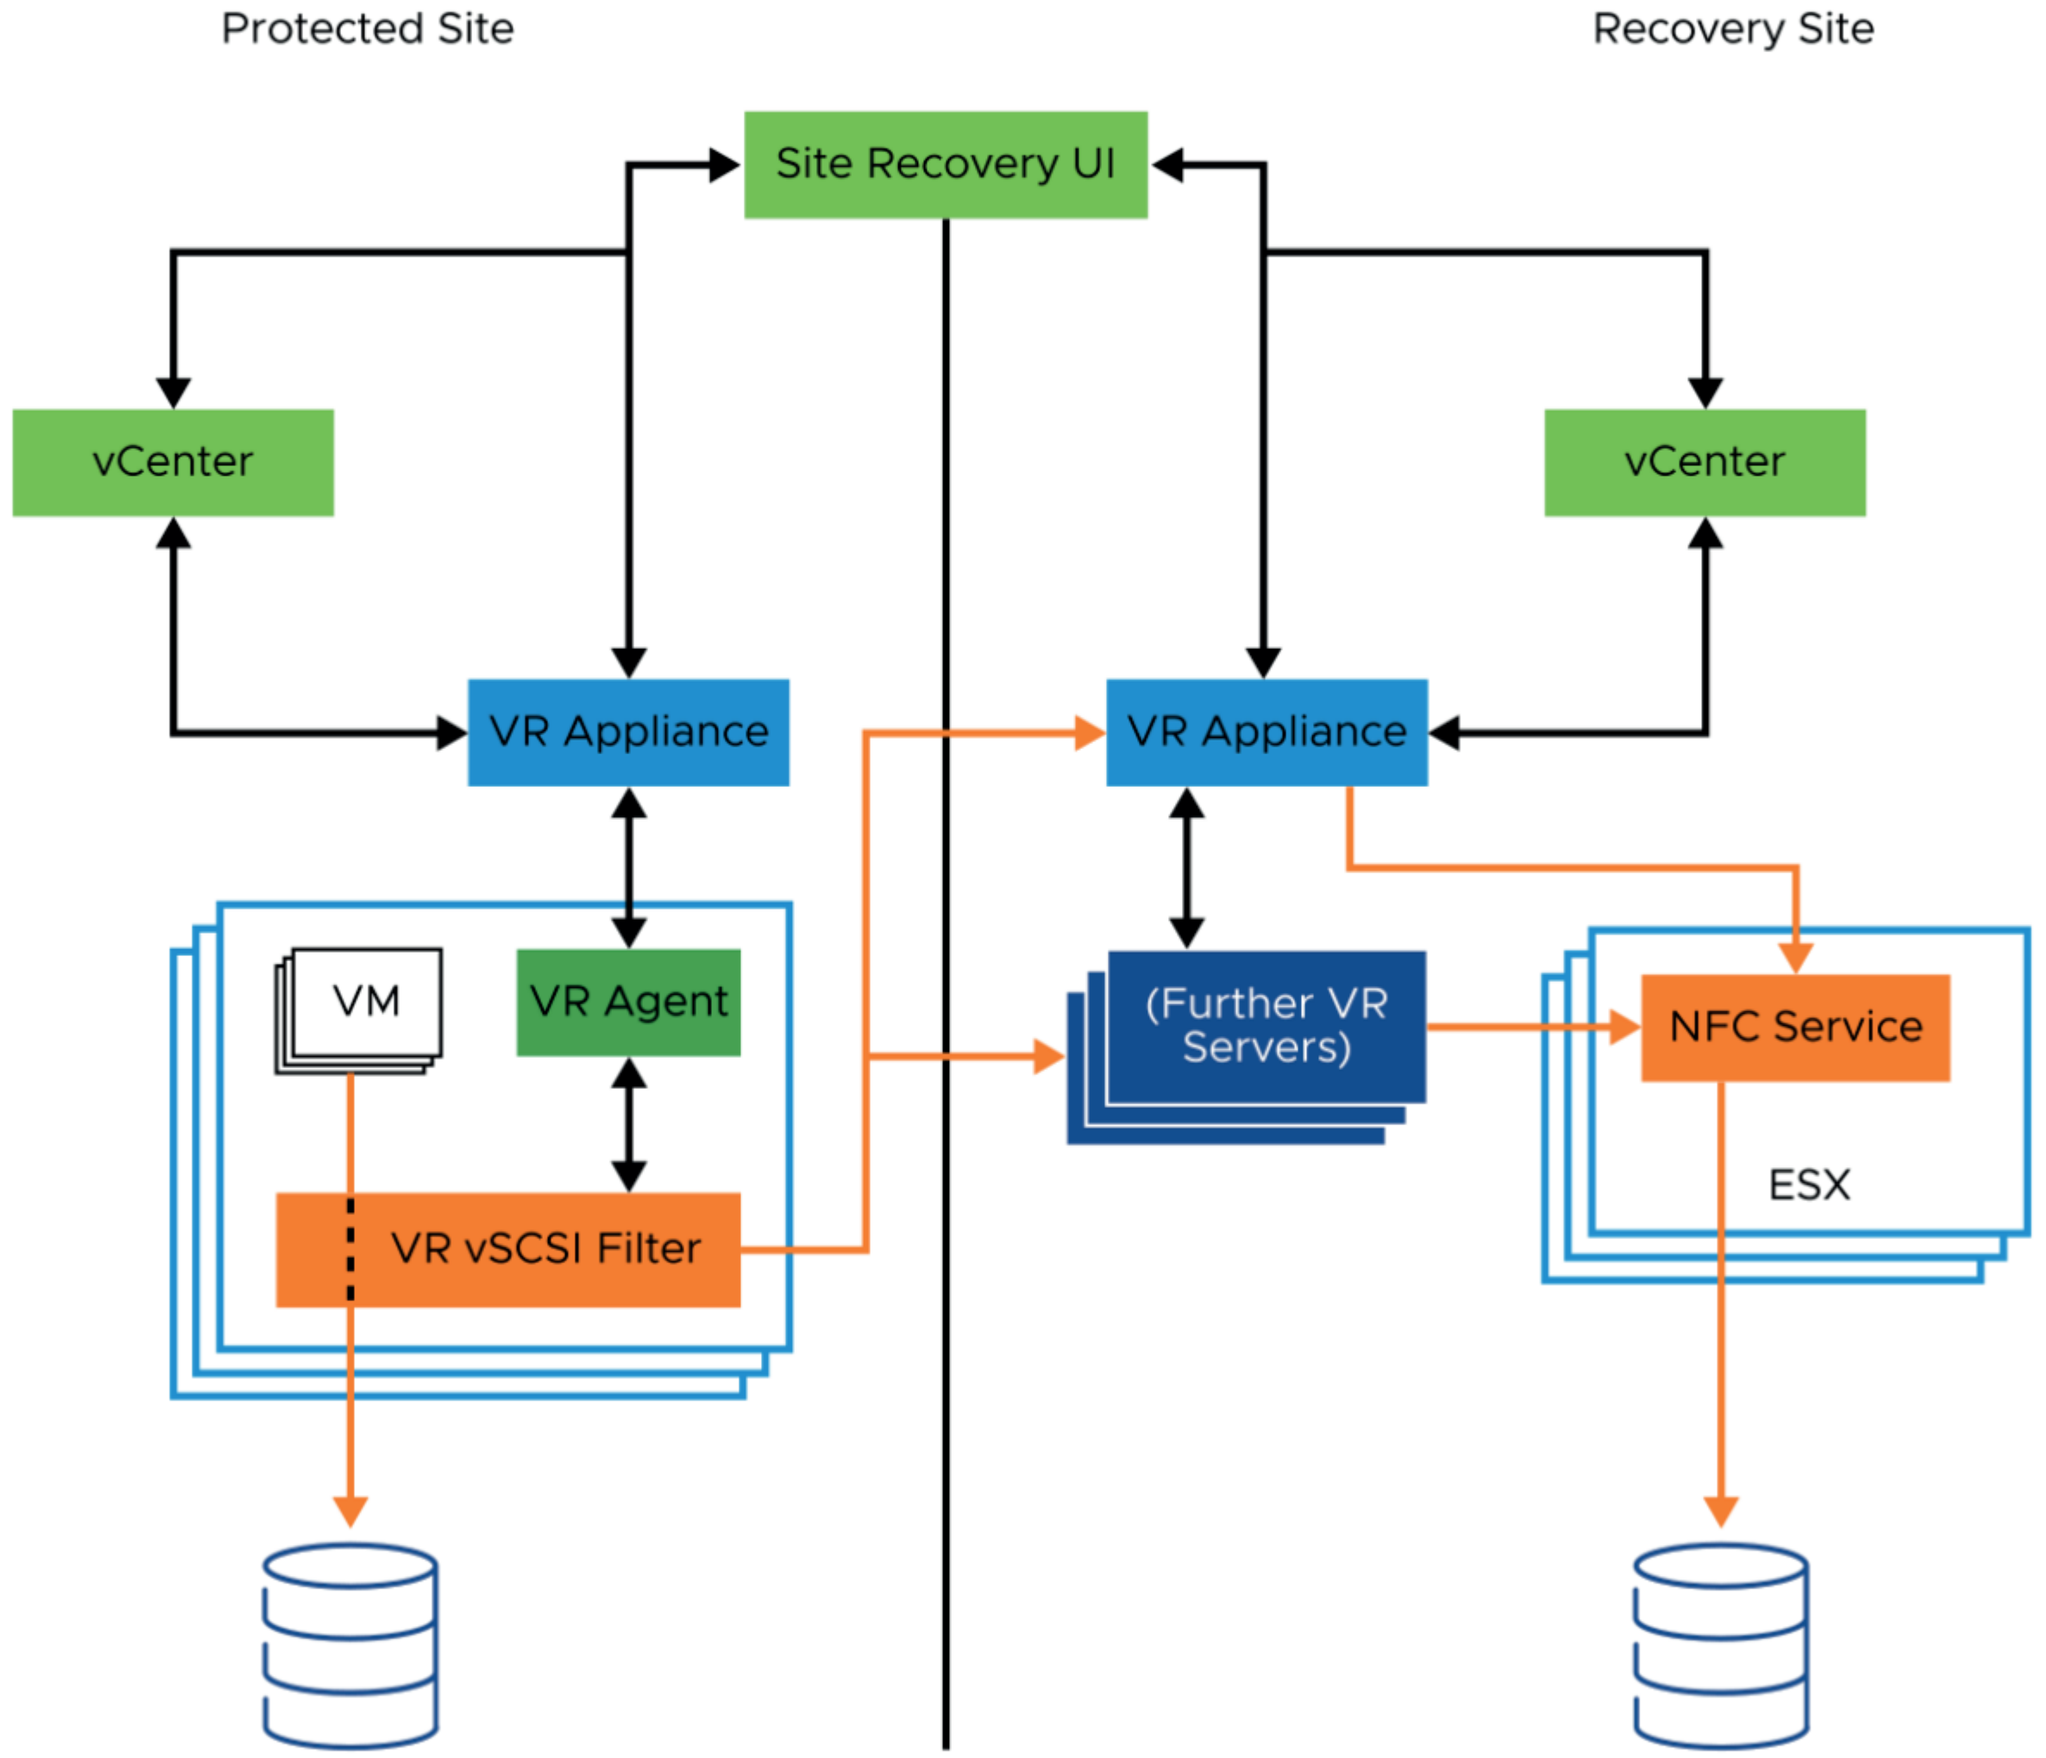

Technet24

2-37 Replication in a Single vCenter

You can also replicate a virtual machine between datastores at the same vCenter.

Only one vSphere Replication appliance is deployed on the single vCenter.

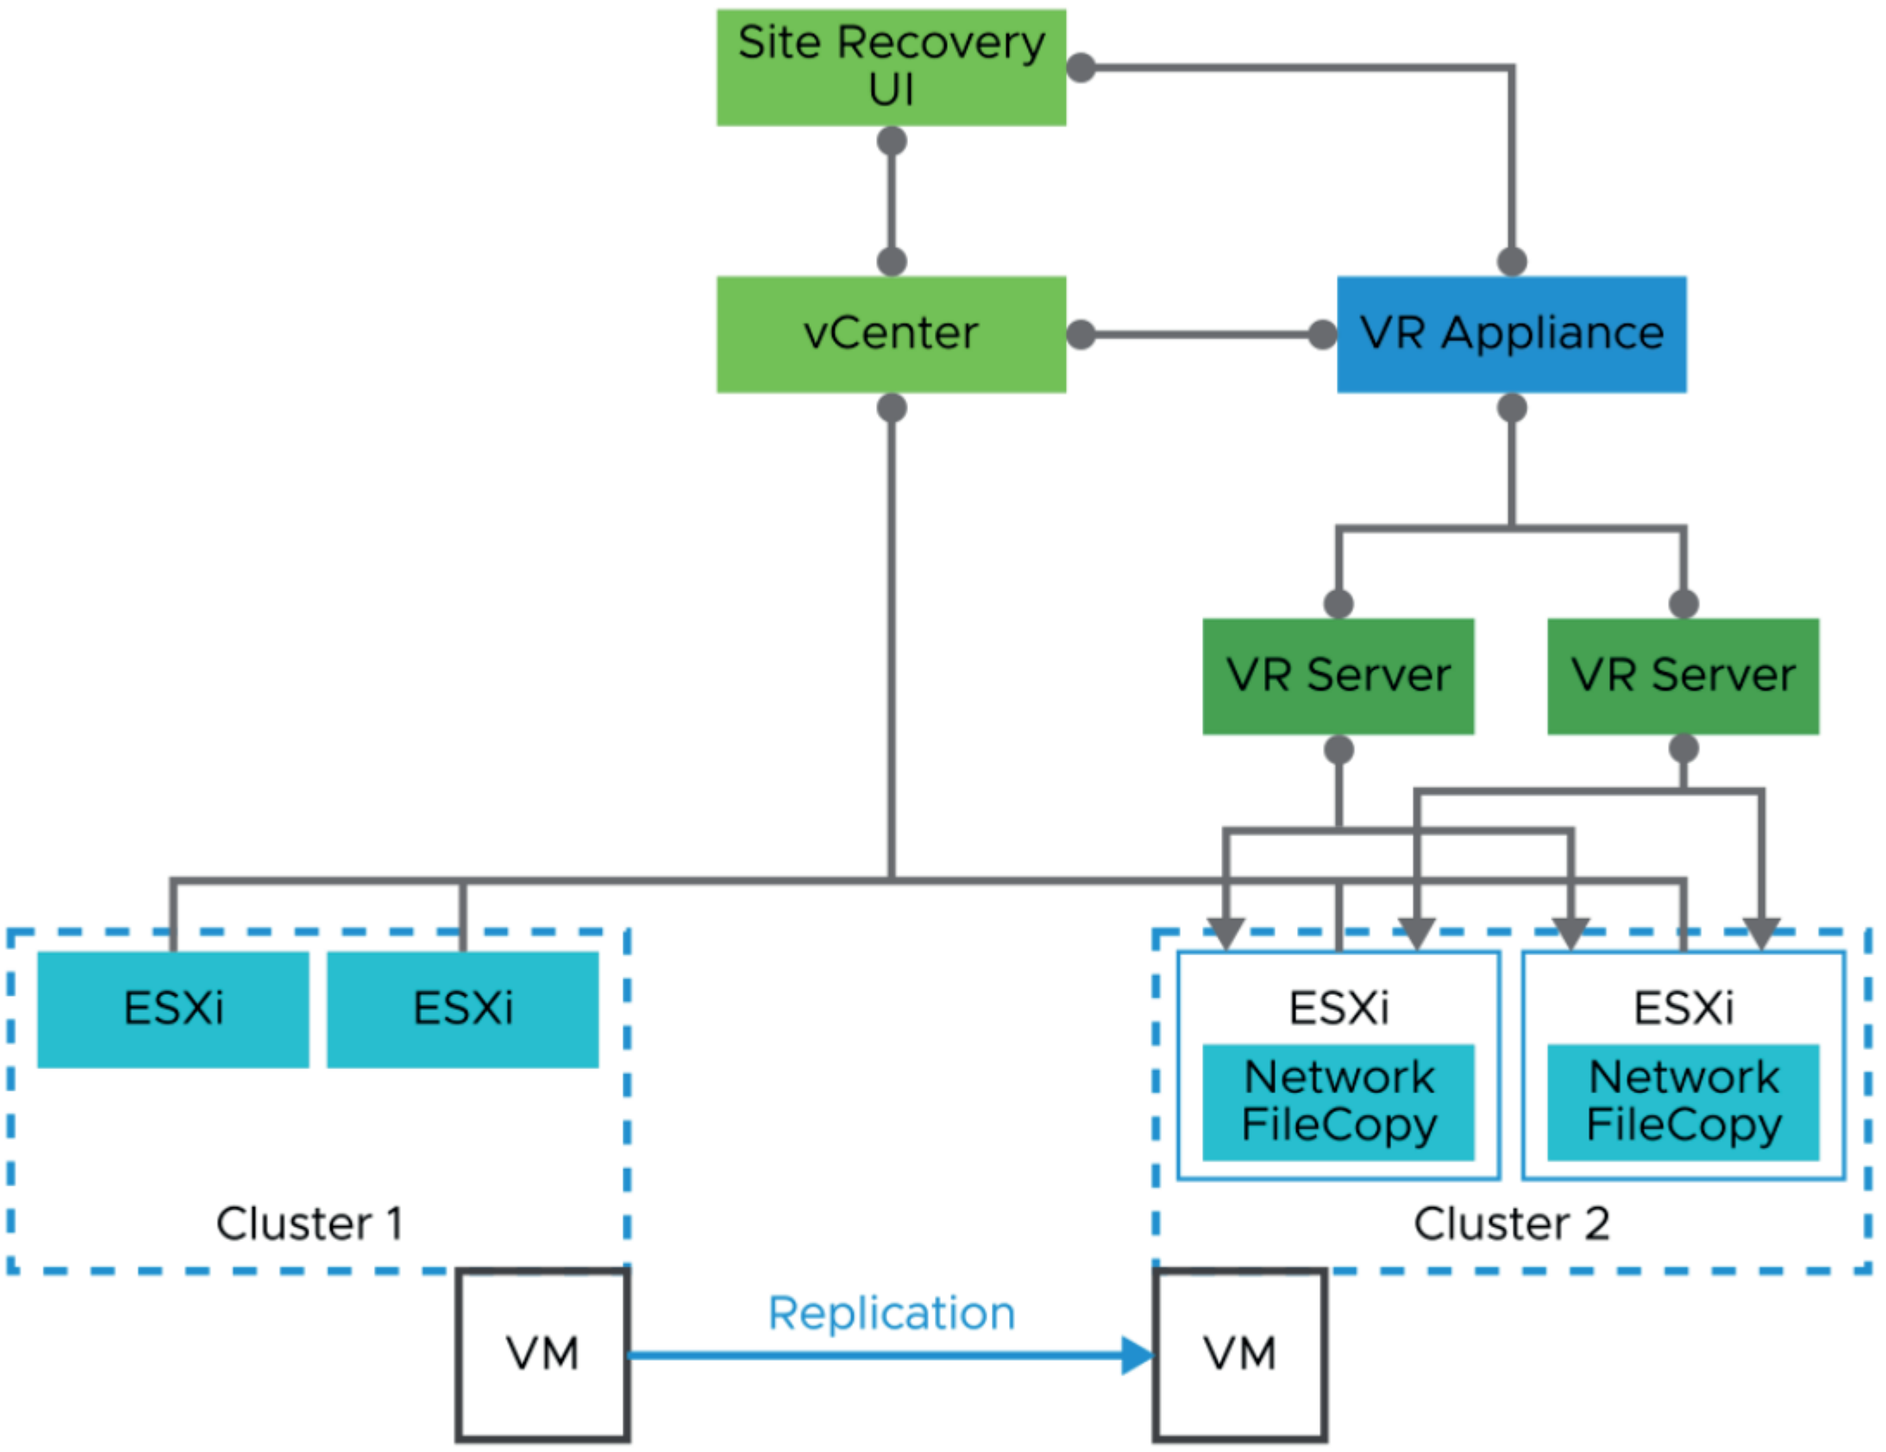

2-38 Replication to a Shared Target Site

You can replicate virtual machines to a shared target site.

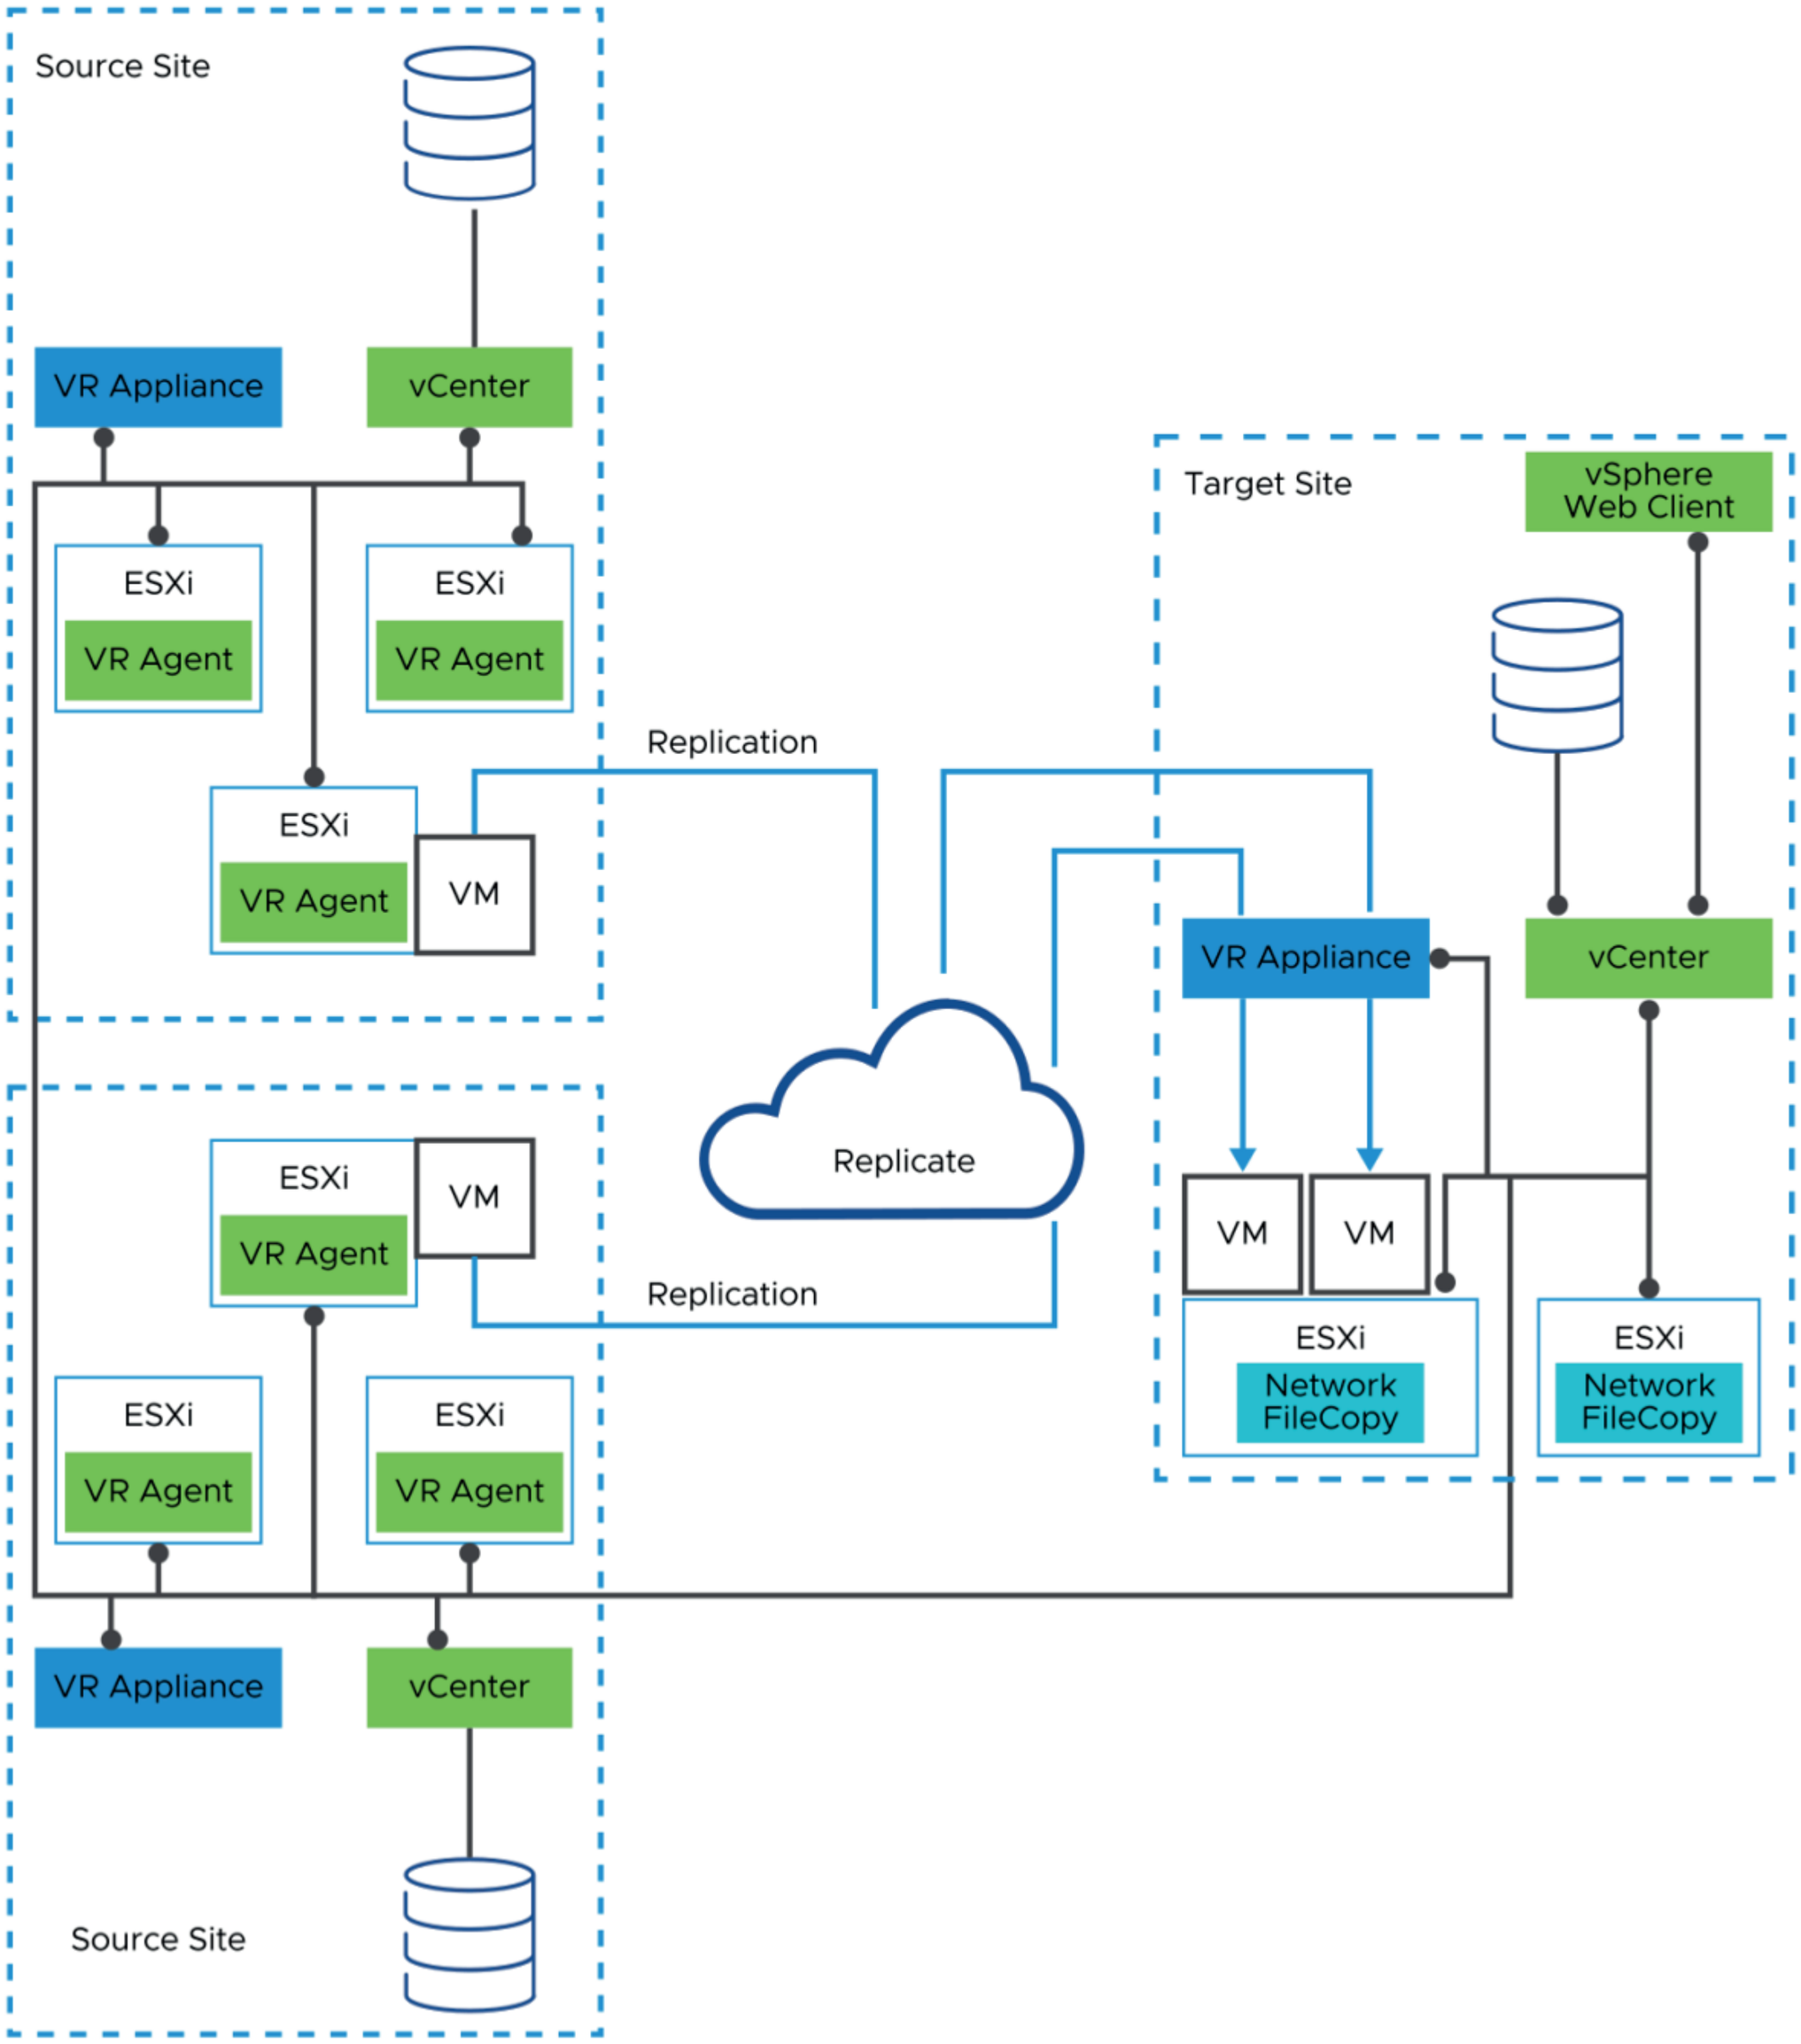

Technet24

2-39 Deploying the vSphere Replication Appliance

You use the vSphere Client to deploy the vSphere Replication appliance on an ESXi host:

1. Download the installation package to a local directory or obtain its URL.
2. Use the standard OVF deployment wizard to deploy the appliance.
   a.. From the main menu, select **Inventory** and click the **Hosts and Clusters** icon
   b.. Right-click an ESXi host and select **Deploy OVF template.**

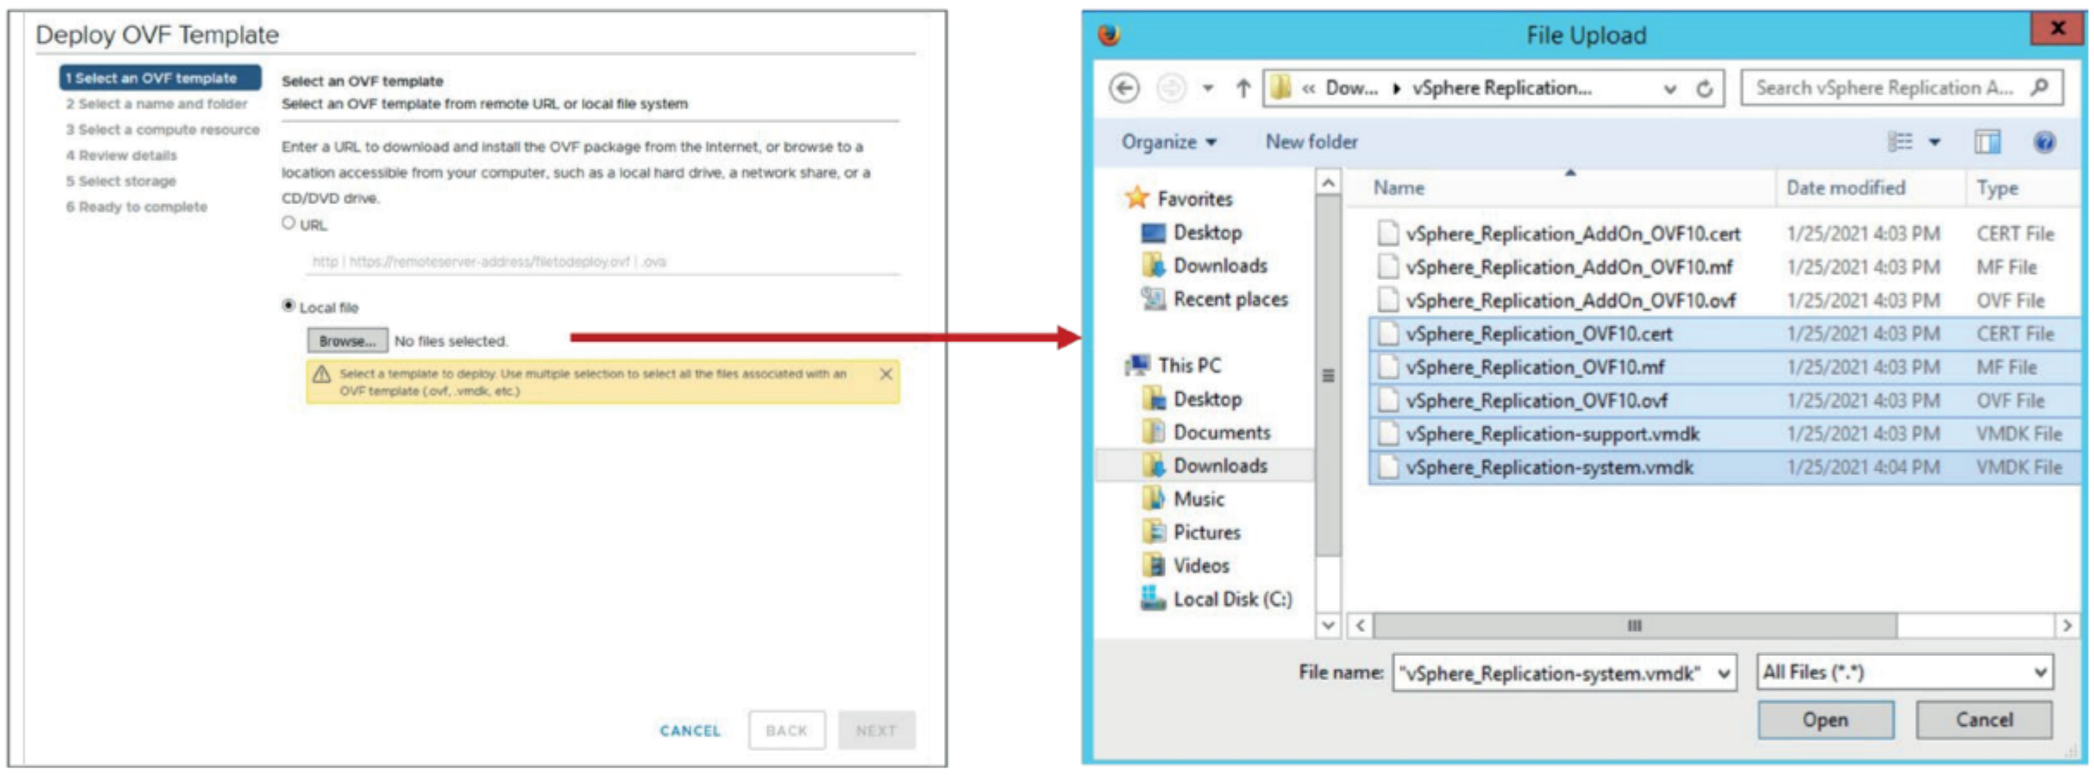

You can deploy vSphere Replication with either an IPv4 or IPv6 address. Mixing IP addresses, for example having a single appliance with an IPv4 and an IPv6 address, is not supported.

After you deploy the vSphere Replication appliance, you use the VRMS appliance management interface to register the endpoint and the certificate of the vSphere Replication management server with the vCenter Lookup Service. You also use the VRMS appliance management interface to register the vSphere Replication solution user with the vCenter Single Sign-On administration server.

Once the vSphere Replication appliances have been deployed and registered with vCenter at both sites, the vSphere Replication appliances need to be paired with the vCenter on the other site using the vSphere client.

For more details on deploying the vSphere Replication appliance, see *VMware vSphere Replication Documentation* at https://docs.vmware.com/en/vSphere-Replication/index.html.

2-40 Cloud Provider

In addition to installing and configuring the vSphere Replication appliance, you must configure the connection to your cloud provider.

You can configure a connection to the cloud provider before you start the Configure Replication wizard or while you configure a replication task.

2-41 Lab 3: (Simulation) Deploying vSphere Replication and Site Recovery Instances

Deploy vSphere Replication and Site Recovery Instances:

1. Deploy and Register the On-Premises vSphere Replication Instance with vCenter
2. Deploy and Register the On-Premises Site Recovery Instance with vCenter

2-42 Review of Learner Objectives

- Identify the components in the vSphere Replication architecture
- Deploy and configure vSphere Replication and Site Recover instances

Technet24

2-43 Lesson 4: VM Replication

2-44 Learner Objectives

- Recover replicated virtual machines

2-45 Configuring vSphere Replication for a Single VM

vSphere Replication can protect individual VMs and their virtual disks by replicating them to another location.

To configure vSphere Replication for a VM in the vSphere Client, right-click the VM in the inventory and select **All Site Recovery actions > Configure Replication**.

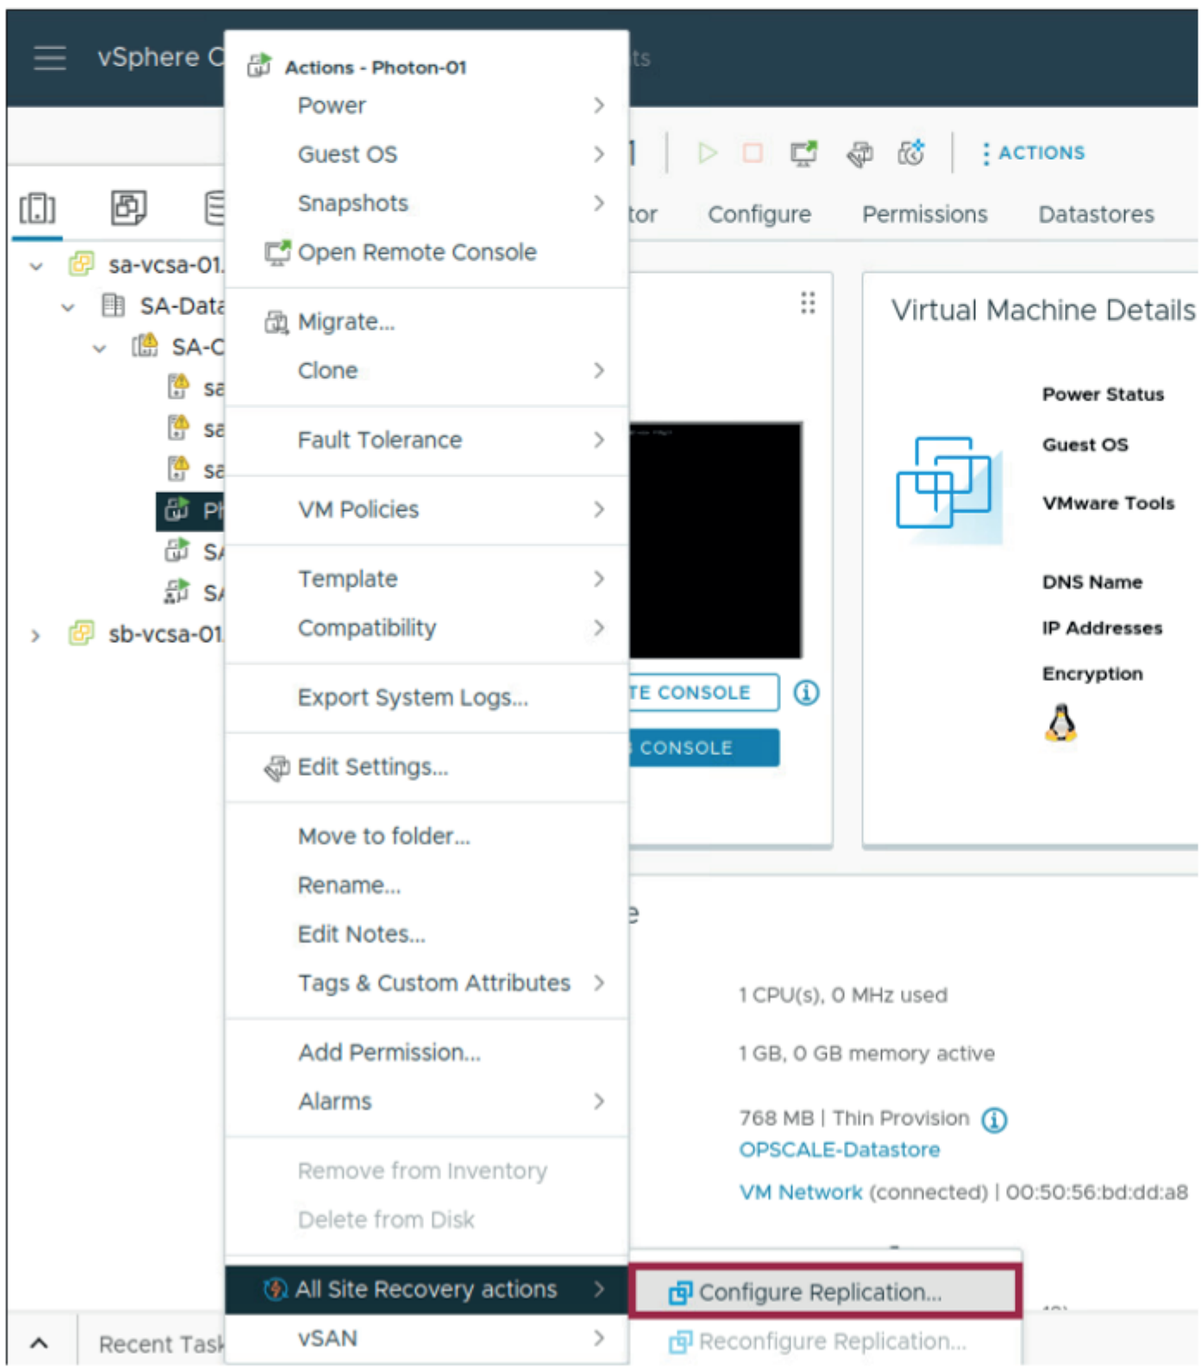

2-46 Configuring Recovery Point Objective and Point in Time Instances

During replication configuration, you can set an RPO and configure retention of instances from multiple points in time.

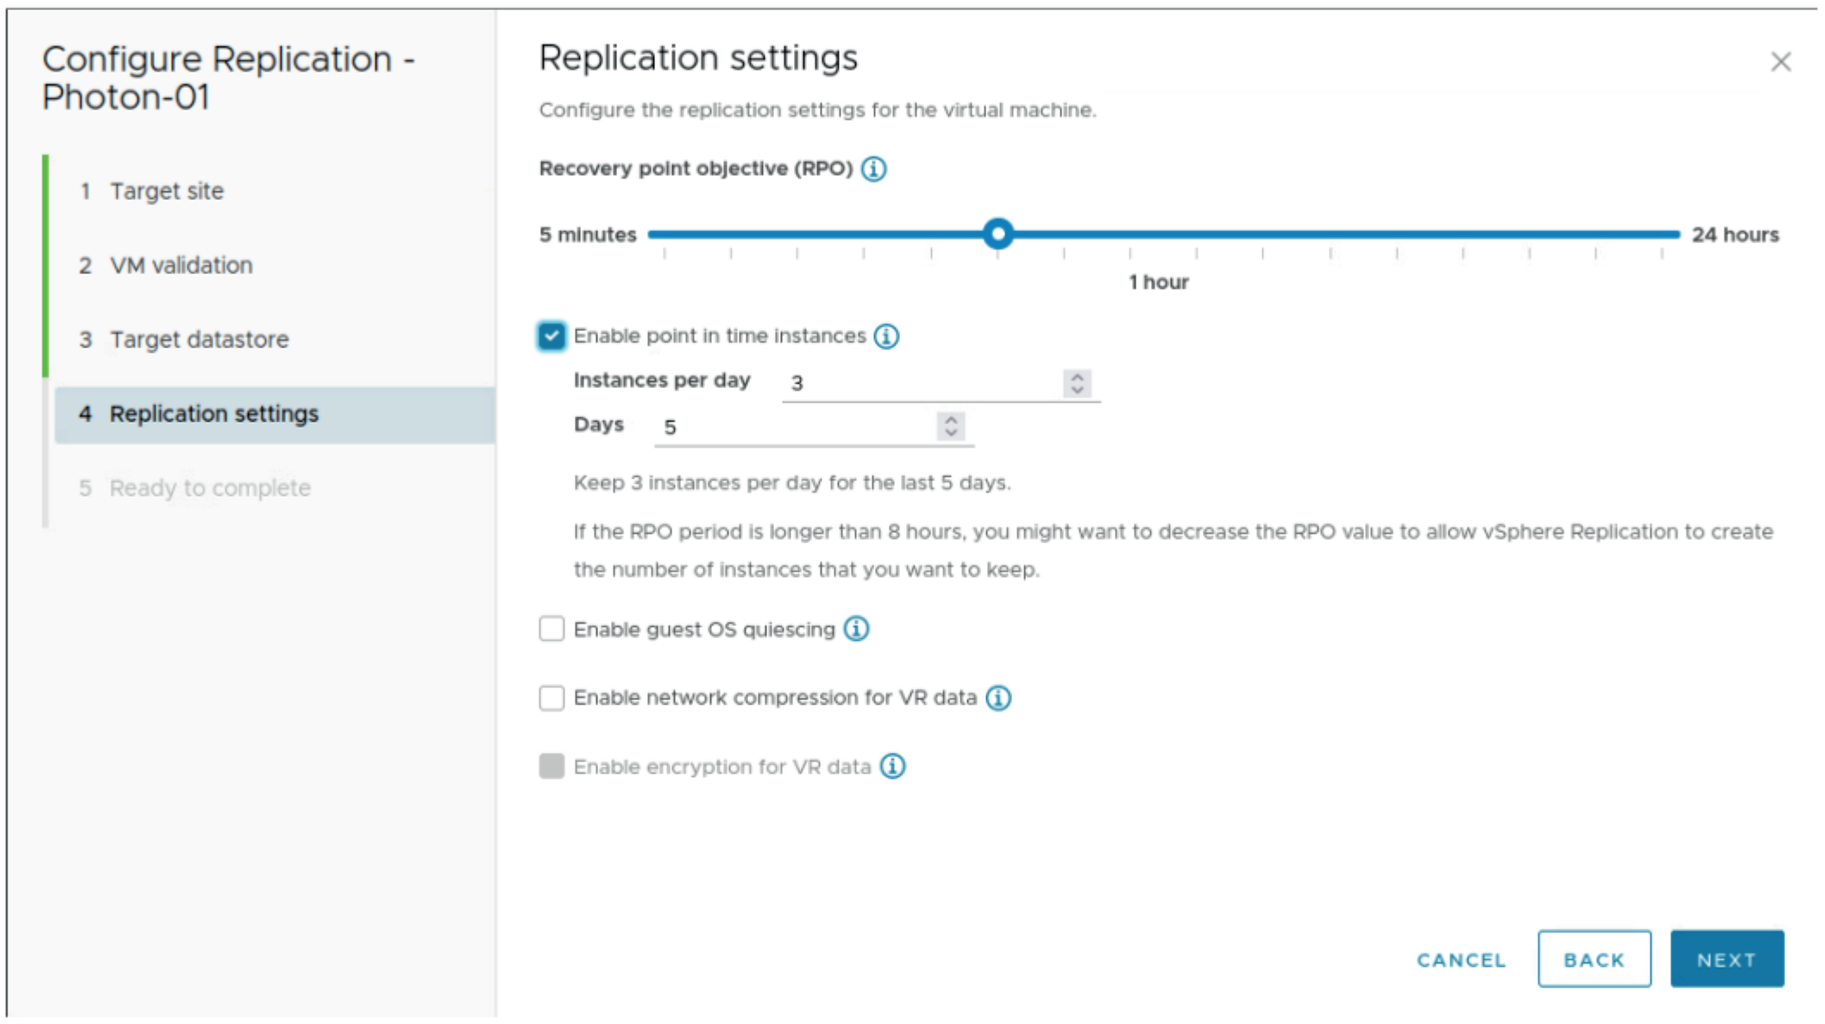

The value that you set for the recovery point objective (RPO) affects replication scheduling.

When you configure replication, you set an RPO to determine the time between replications. For example, an RPO of 1 hour aims to ensure that a VM loses no more than 1 hour of data during the recovery. For smaller RPOs, less data is lost in a recovery, but more network bandwidth is consumed to keep the replica up to date.

For a discussion about how the RPO affects replication scheduling, see *vSphere Replication Administration* at https://docs.vmware.com/en/vSphere-Replication/8.6/vsphere-replication-86-admin.pdf.

Technet24

2-47 Recovering Replicated VMs (1)

To perform the recovery, you use the Recover virtual machine wizard in the vSphere Client at the target site.

With vSphere Replication, you can recover VMs that were successfully replicated at the target site.

You can recover one VM at a time on the **Replications** tab.

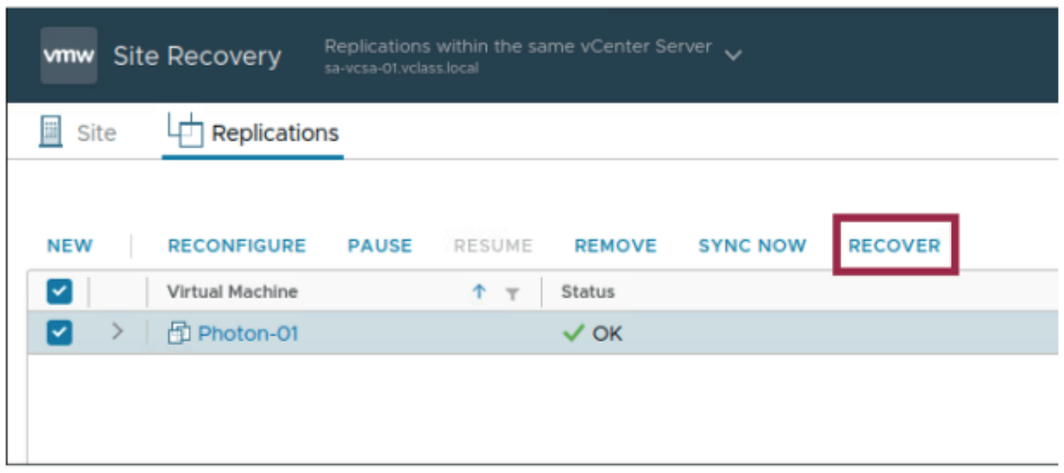

2-48 Recovering Replicated VMs (2)

When recovering the VM. You are asked to select either to recover the VM with all the latest data or to recover the VM with the most recent data available on the target site.

Recover with recent changes to avoid data loss:

- vSphere Replication performs a full synchronization of the VM from the source site to the target site before recovering the VM.
- This option requires that the data of the source VM be accessible.
- You can select this option only if the VM is powered off.

Recover with latest available data:

- vSphere Replication recovers the VM by using the data from the most recent replication on the target site, without performing synchronization.
- Selecting this option results in the loss of any data that changed since the most recent replication.
- Select this option if the source VM is inaccessible or if its disks are corrupted.
- Select this option only if the VM is powered off

2-49 Lab 4: (Simulation) Configuring Replication for a Single VM

Configure replication for a single VM:

1. Configure replication for the virtual machine
2. Verify the replication status

2-50 Review of Learner Objectives

- Recover replicated virtual machines

2-51 Key Points

- With a centralized repository, you have a single location where VMware Tools are located
- To protect your VM's data, you can use a backup solution based on vSphere Storage APIs - Data Protection
- You can use vSphere Replication to protect VMs as part of a disaster recovery strategy

Questions?

Technet24

Module 3

vSphere Cluster Operations

3-2 Importance

Resource pools allow you to delegate control over resources of a host or a cluster.

vSphere Cluster Services (vCLS) ensures that if vCenter becomes unavailable, cluster services remain available to maintain the resources and health of the workloads that run in the clusters.

3-3 Module Lessons

1. Resource Pools
2. vSphere Cluster Services (vCLS)

Technet24

3-4 Lesson 1: Resource Pools

3-5 Learner Objectives

- Create a resource pool
- Set resource pool attributes
- Describe expandable reservations
- Describe scalable shares

3-6 Reservations, Limits, and Shares (1)

Beyond the CPU and memory configured for a VM, you can apply resource allocation settings to a VM to control the amount of resources granted.

- A reservation specifies the guaranteed minimum allocation for a VM.
- A limit specifies an upper bound for CPU or memory that can be allocated to a VM.
- A share is a value that specifies the relative priority or importance of a VM's access to a given resource.

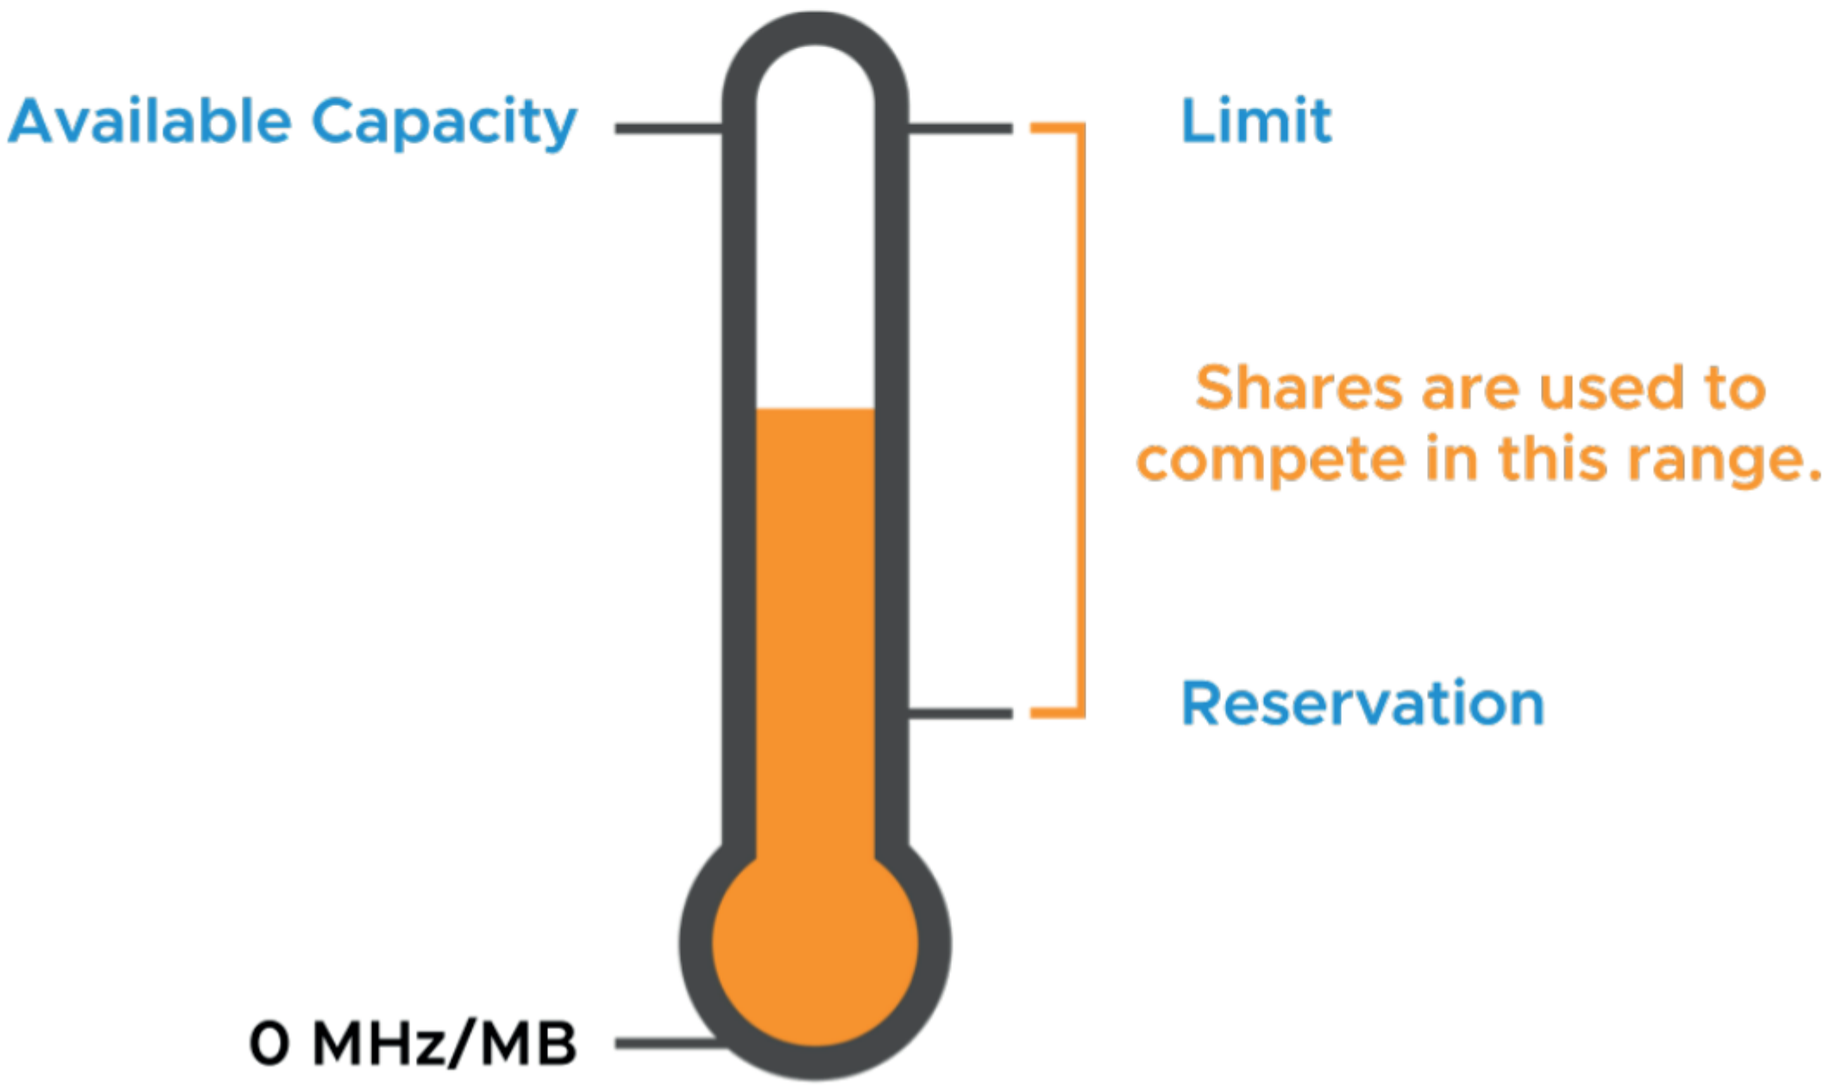

If host resources are over-committed and VM workloads are high, resource contention can occur.

To manage resources efficiently, vSphere provides mechanisms to allow less, more, or an equal amount of access to a defined resource. vSphere also prevents a VM from consuming large amounts of a resource. vSphere grants a guaranteed amount of a resource to a VM whose performance is not adequate or that requires a certain amount of a resource to run properly.

When host memory or CPU contention begins, a VM's allocation target is somewhere between its specified reservation and specified limit, depending on the VM's shares and the system load. vSphere uses a share-based allocation algorithm to achieve efficient resource use for all VMs and to guarantee a given resource to the VMs that need it most.

3-7 Reservations, Limits, and Shares (2)

Resource Allocation: RAM

Reservations

- Memory reserved to a VM is guaranteed never to swap or balloon.
- If an ESXi host does not have enough unreserved RAM to support a VM with a reservation, the VM does not power on.
- Reservations are measured in MB, GB, or TB.
- The default is 0 MB.

Limits:

- VMs never consume more physical RAM than is specified by the memory allocation unit.
- VMs might use the VM swap mechanism (.vswp) if the guest OS attempts to consume more RAM than is specified by the limit.

Resource Allocation: CPU

Reservations

- CPU that is reserved for a VM is guaranteed to be immediately scheduled on physical cores. The VM is never placed in a CPU ready state.
- If an ESXi host does not have enough unreserved CPU to support a VM with a reservation, the VM does not power on.
- Reservations are measured in MHz or GHz.
- The default is 0 MHz.

Limits:

- VMs never consume more physical CPU than is specified by the CPU allocation unit.
- CPU threads are placed in a ready state if the guest OS attempts to schedule threads faster than the limit allows.

Technet24

3-8 About Resource Pools

A resource pool is a logical abstraction of hierarchically managed CPU and memory resources.

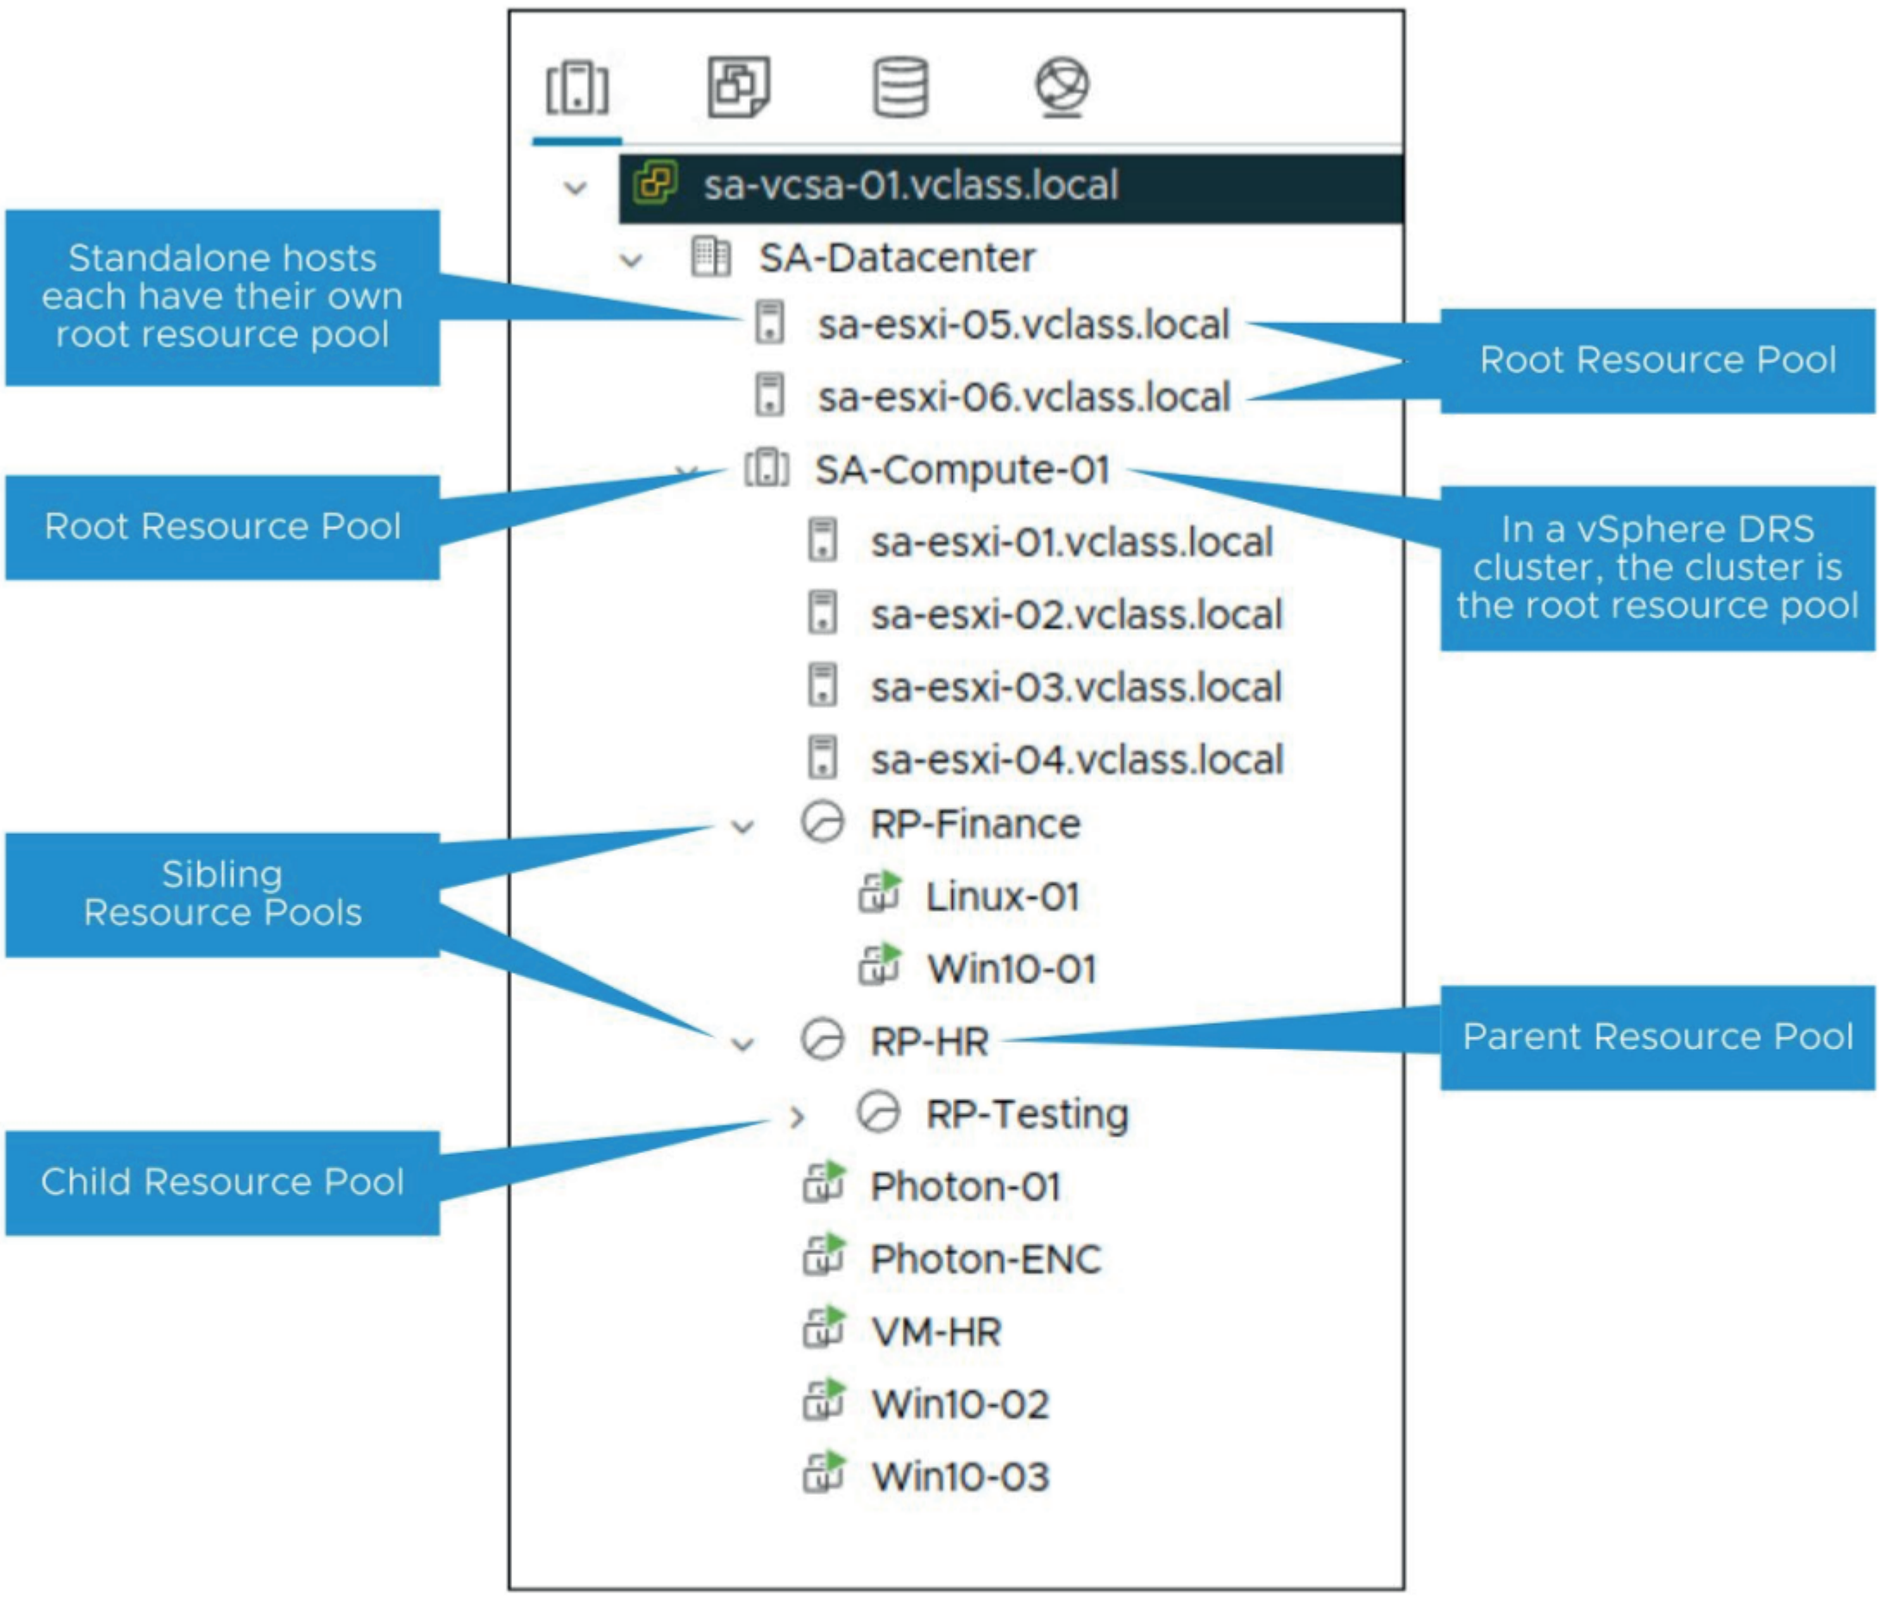

A resource pool is a logical abstraction for managing resources. Resource pools can be grouped into hierarchies and used to hierarchically partition available CPU and memory resources. In the example on the slide, RP-HR is the parent resource pool for RP-Testing-UI. RP-Finance and RP-HR are siblings.

With a resource pool, you can divide and allocate resources to VMs and other resource pools. You can control the aggregate CPU and memory resources of the compute resource. The compute resource can be either a standalone host or a vSphere DRS cluster. Resource pools are also used to delegate privileges to other users and groups.

The topmost resource pool is called the root resource pool. Each standalone host and each vSphere DRS cluster has an (invisible) root resource pool that groups the resources of that host or cluster. The root resource pool does not appear because the resources of the host (or cluster) and the root resource pool are always the same.

3-9 Reasons to Use Resource Pools

With resource pools, you can delegate control over resources of a standalone host or a vSphere DRS cluster. Using resource pools can result in the following benefits:

- Flexible hierarchical organization: Add, remove, or reorganize resource pools or change resource allocations as needed.
- Isolation between pools, sharing in pools: Top-level administrators can make a pool of resources available to a department-level administrator.
- Access control and delegation: VM creation and management are performed within the boundaries of the resources to which the resource pool is entitled. Delegation is usually done with permission settings.
- Separation of resources from hardware: If you use VMware vSphere® Distributed Resource Scheduler™ (DRS) clusters, the resources of all hosts are always assigned to the cluster.
- Management of sets of VMs running a multitier service: Group VMs for a multitier service in a resource pool.

Technet24

3-10 Resource Pool Attributes

You can create a child resource pool of any ESXi host, resource pool, or vSphere DRS cluster:

- Shares: Low, Normal, High, Custom
- Reservation: In MHz, GHz, MB, or GB
- Reservation Type:
  - **Expandable** check box selected: VMs and sub-pools can draw from this pool's parent.
  - **Expandable** check box deselected: VMs and sub-pools can draw only from this pool, even if its parent has free resources.
- Limit:
  - In MHz, GHz, MB, or GB
  - Unlimited access, by default, up to maximum amount of accessible resources

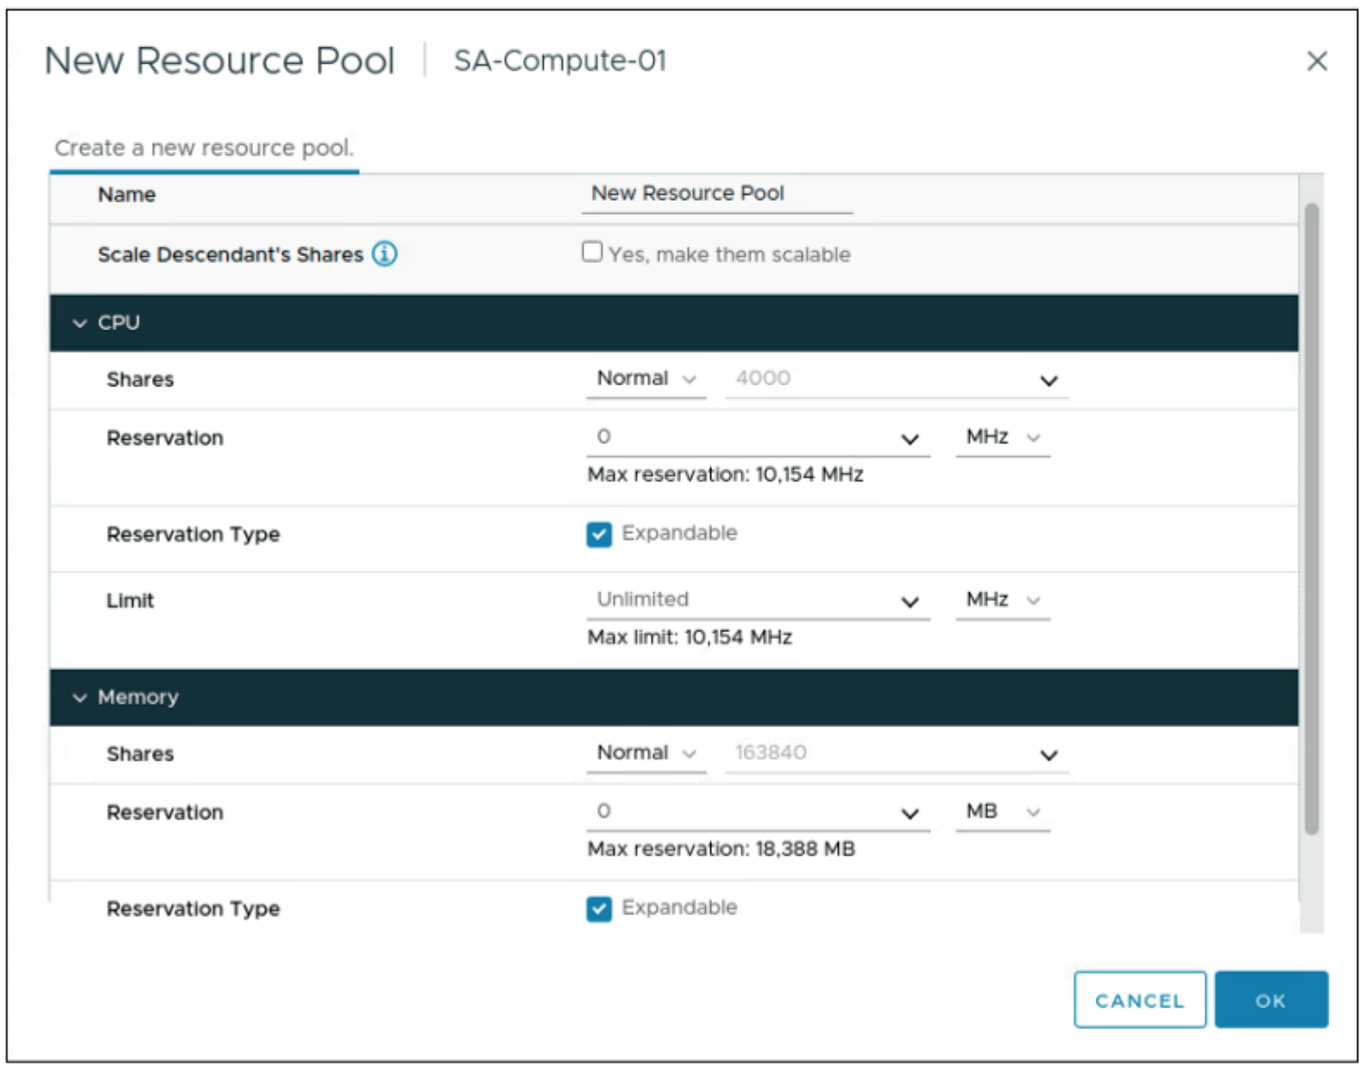

Like VMs, resource pools have reservation, limit, and share values for CPU and memory resources.

High Shares receives 8000 CPU, Normal Shares receives 4000 CPU, Low Shares receives 2000 CPU based on 16GB RAM and 4 vCPUs.

Resource pools also have a setting that enables reservations to be expanded. With this expandable reservation attribute, a resource pool that cannot satisfy a reservation request can search through its hierarchy to find unreserved capacity to satisfy the reservation request.

3-11 Resource Pool Example

Resource attributes are set on a resource pool.

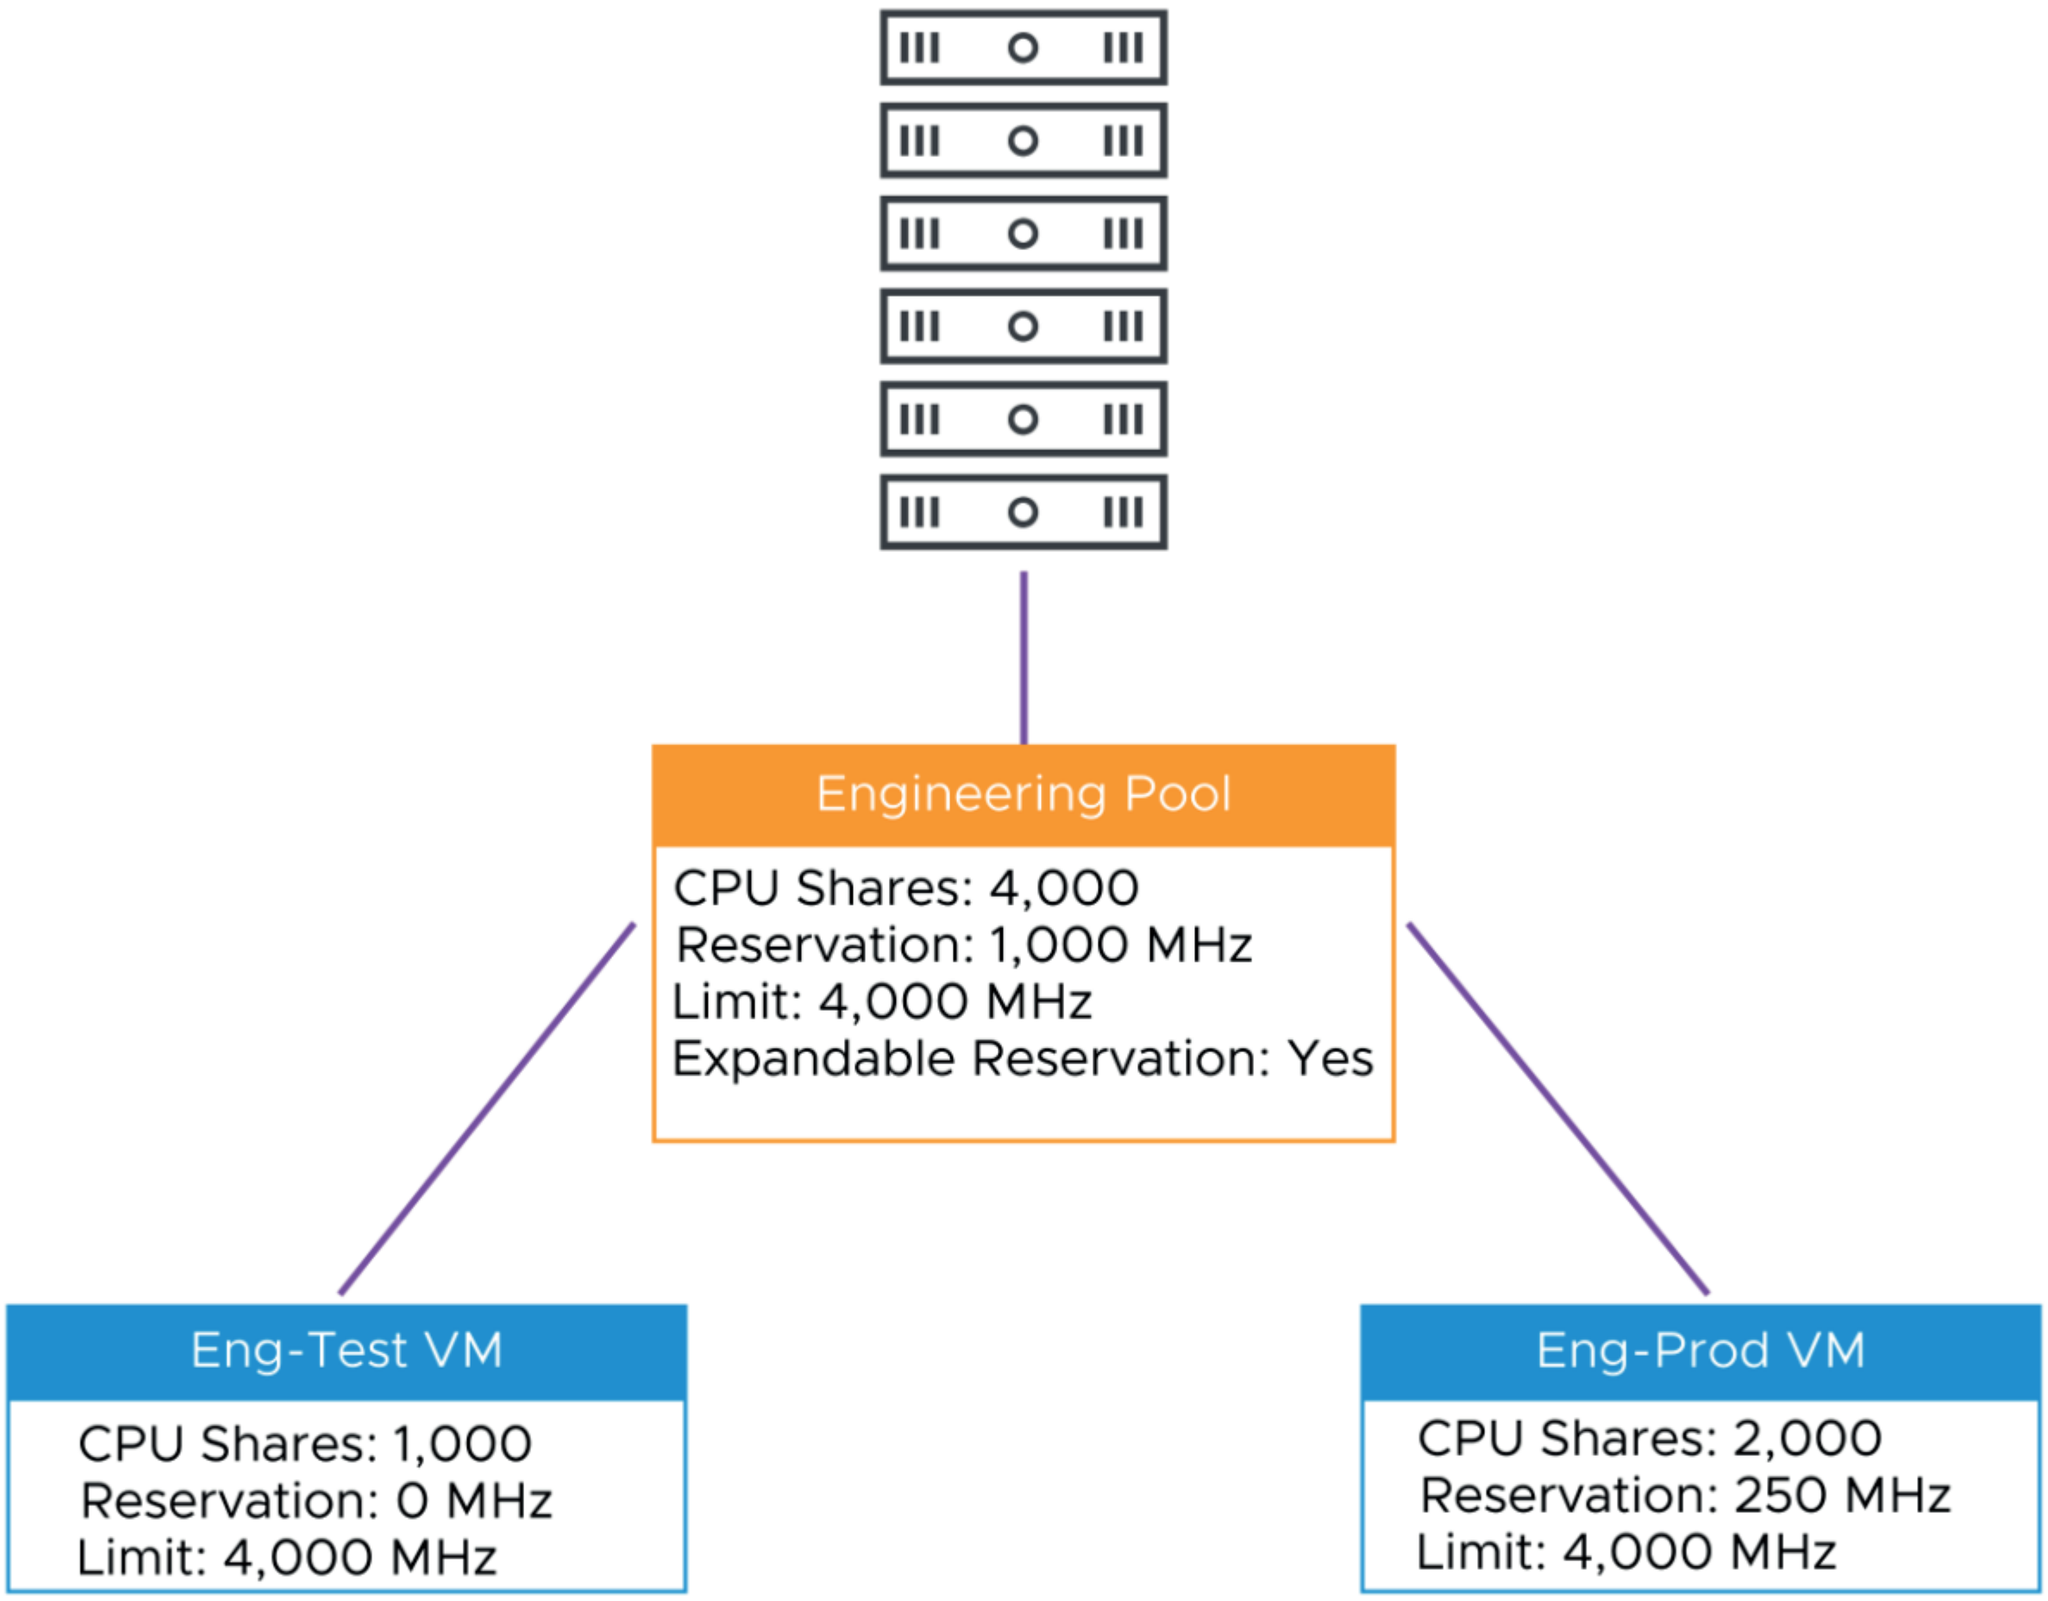

Technet24

3-12 Resource Pool Example: CPU Shares

The Finance resource pool has twice as many CPU shares as the Engineering resource pool. The Finance resource pool is entitled to twice as many CPU resources as the Engineering resource pool.

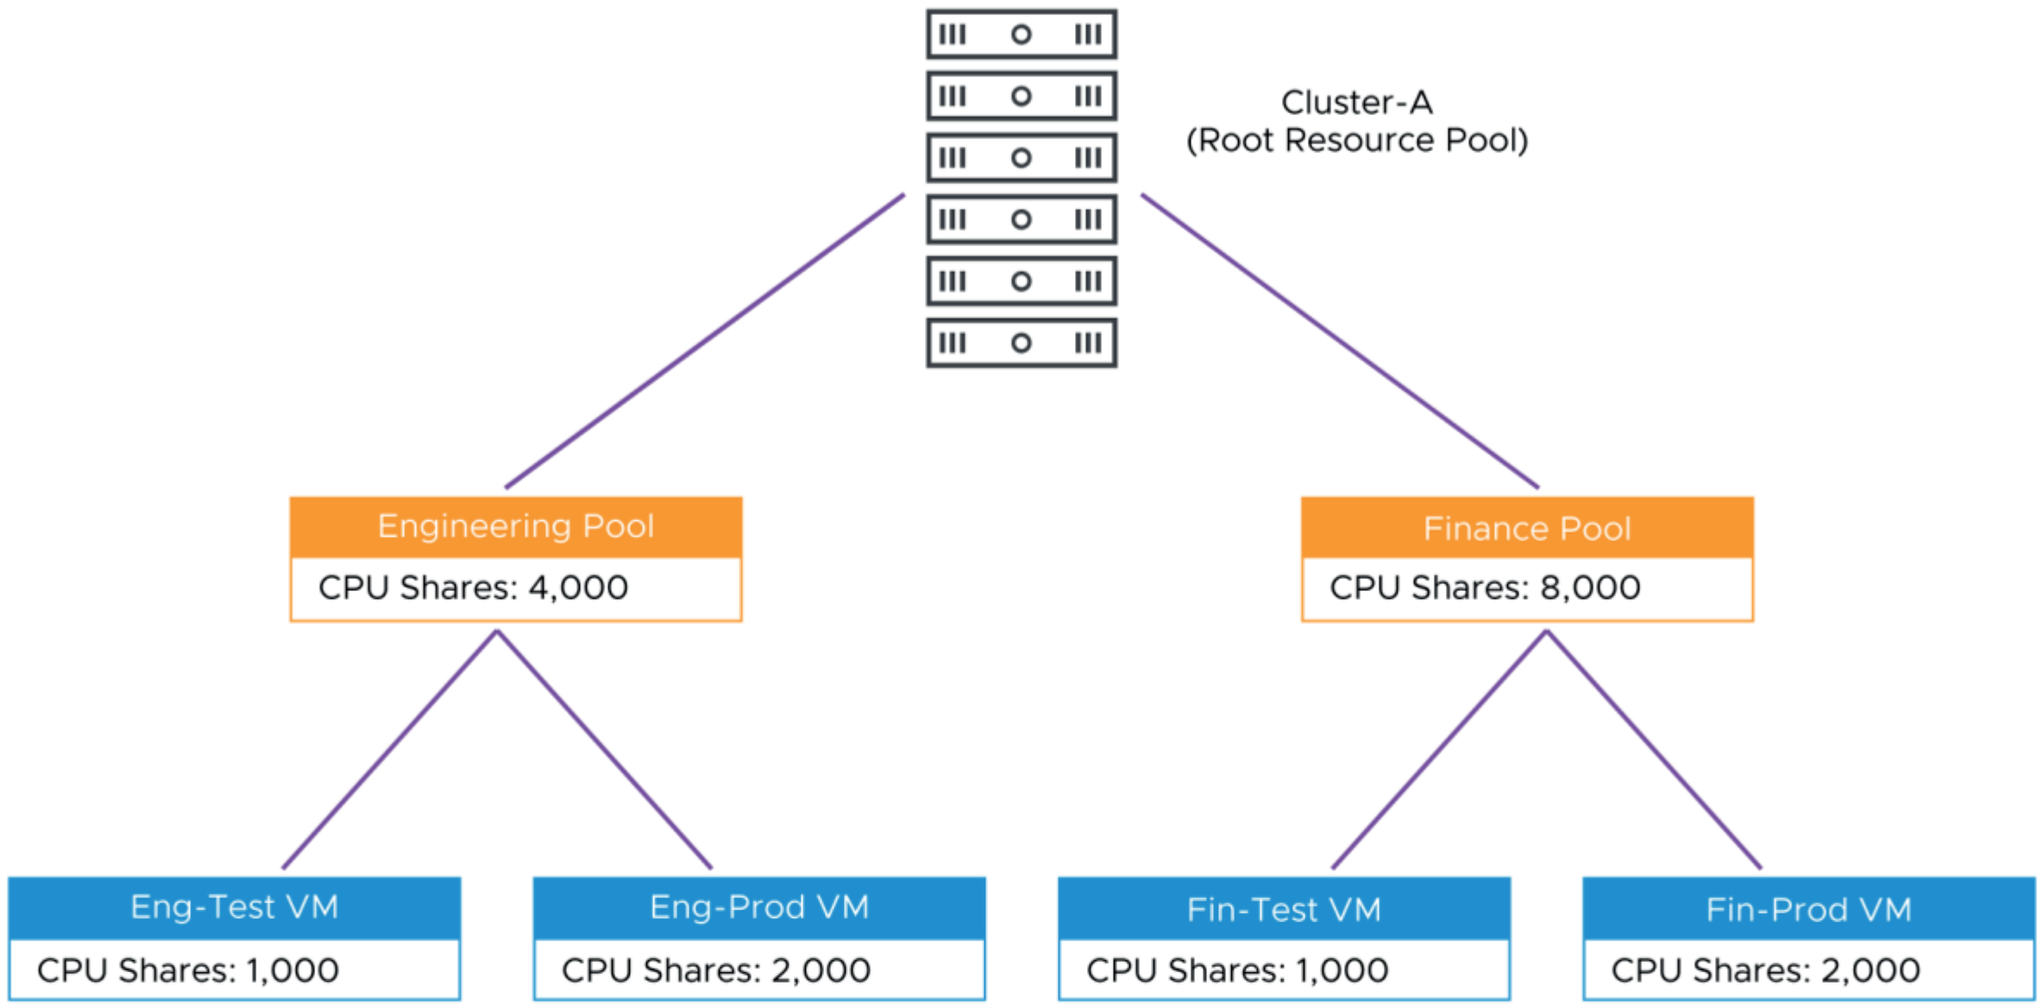

Resource pools can be organized hierarchically. The root resource pool is the topmost resource pool. The root resource pool includes the sum for all CPUs (in megahertz) and the sum of all the installed RAM (in megabytes) available in the compute environment (standalone host or cluster).

Except for the root resource pool, every resource pool has a parent resource pool. A resource pool might contain child resource pools or only VMs that are powered on in the pool.

A child resource pool is used to allocate resources from the parent resource pool for the child's consumers. Administrative control can also be delegated to individuals or organizations. A child resource pool cannot exceed the capacity of the parent resource pool. By creating a child pool, you reserve resources from the parent pool, whether VMs in the child pool are powered on.

Shares specify the relative priority or importance of either a resource pool or VM. If a resource pool has twice as many shares of a resource as another resource pool, it is entitled to consume twice as much of that resource. The same logic applies to VMs.

3-13 Resource Pool Example: CPU Contention

The VMkernel scheduled all the VMs onto the same physical CPU. These VMs are in direct competition with one another.

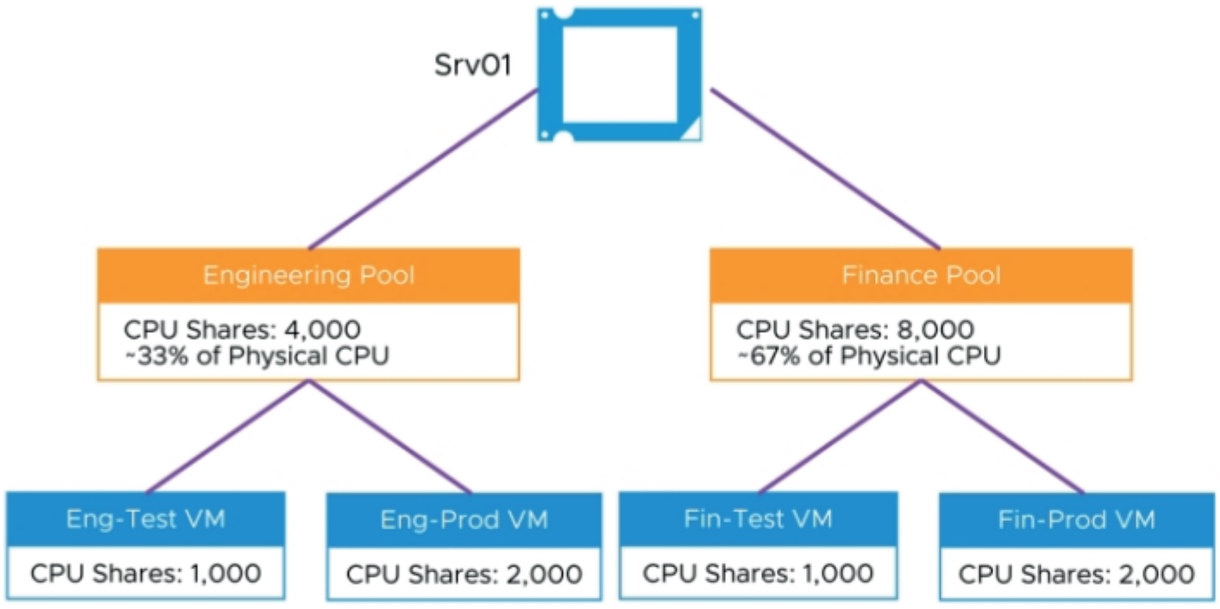

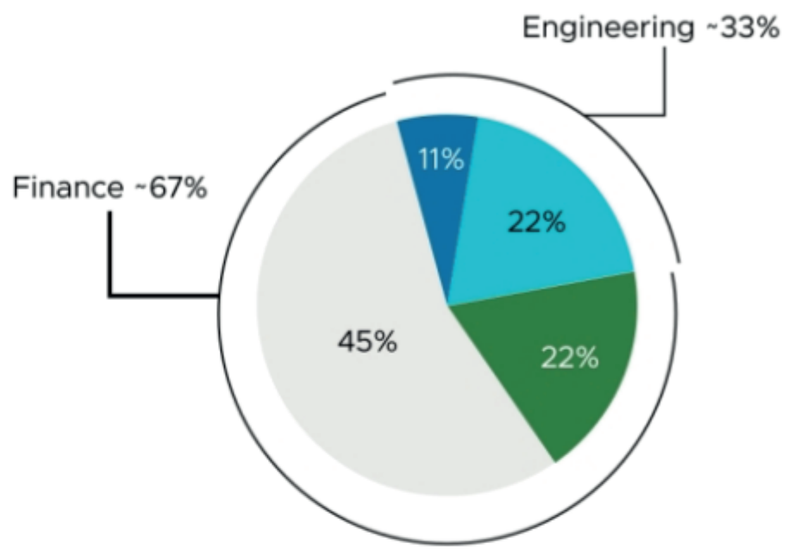

The example on the slide uses approximations to explain how the number of shares affects the amount of CPU allocated to a VM.

The host named Srv01 has only one CPU. The CPU is not named Srv01, that's the name of the ESXi host

The Engineering pool gets approximately 33 percent of the CPU and splits its allotment between VMs Eng-Test and Eng-Prod. Likewise, the Finance pool gets approximately 67 percent of the CPU and splits its allotment between VMs Fin-Test and Fin-Prod. A VM's resource settings are constrained by the resources of the resource pool to which the VM belongs.

The VM Eng-Test gets approximately 33 percent of the CPU allocation of the Engineering resource pool [1,000/(1,000 + 2,000)]. This figure is equal to about 11 percent of the physical CPU (33 percent of 33 percent equals approximately 11 percent). Each of the VMs gets a percentage of the physical CPU allocated to its resource pool based on its individual share allocation.

Technet24

3-14 About Expandable Reservations

- ◊ Expanded reservations are not released until the VM that caused the expansion is shut down or its reservation is reduced.
- ◊ Borrowing resources occurs recursively from the ancestors of the current resource pool:
  - The Expandable Reservation option must be enabled.
  - This option offers more flexibility but less protection.
- ◊ **NOTE: A mismanaged or badly sized expandable reservation might claim all unreserved capacity.**

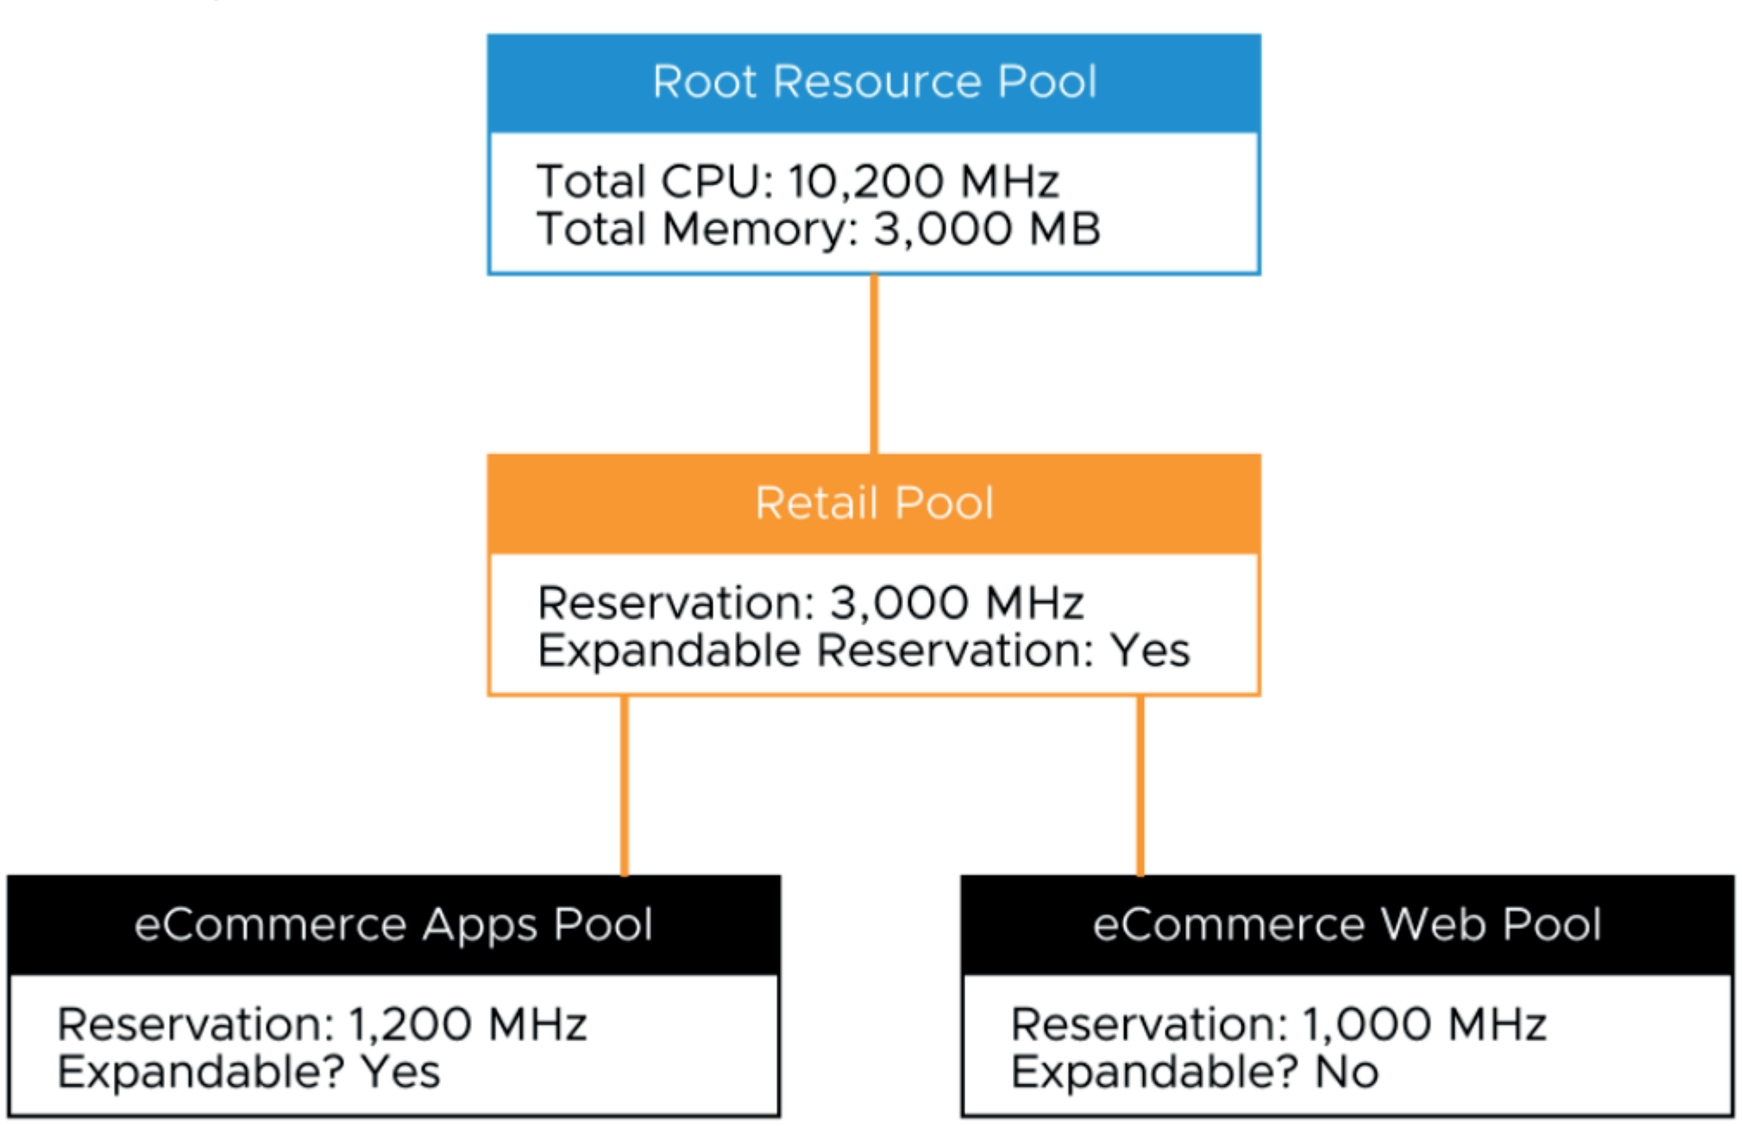

In the example, eCommerce Apps is an expandable child resource pool.

The reservation of a child resource pool cannot exceed that of its parent. The search for unused resources goes through the ancestry of the root resource pool or to the first resource pool where the Expandable Reservation Type is not enabled.

Use expandable reservation carefully. A single child resource pool can use all its parent's available resources, leaving nothing directly available for other child resource pools. This issue can occur if VMs are improperly moved around to get them to start based on resource reservations.

You might want to deactivate the **Expandable** Reservation Type when you configure a fixed amount of resources to a group. For example, an IT administrator whose customers are different organizations in the company that have paid for a fixed amount of CPU and memory resources might want to deactivate the **Expandable** Reservation Type.

3-15 Example of Expandable Reservation (1)

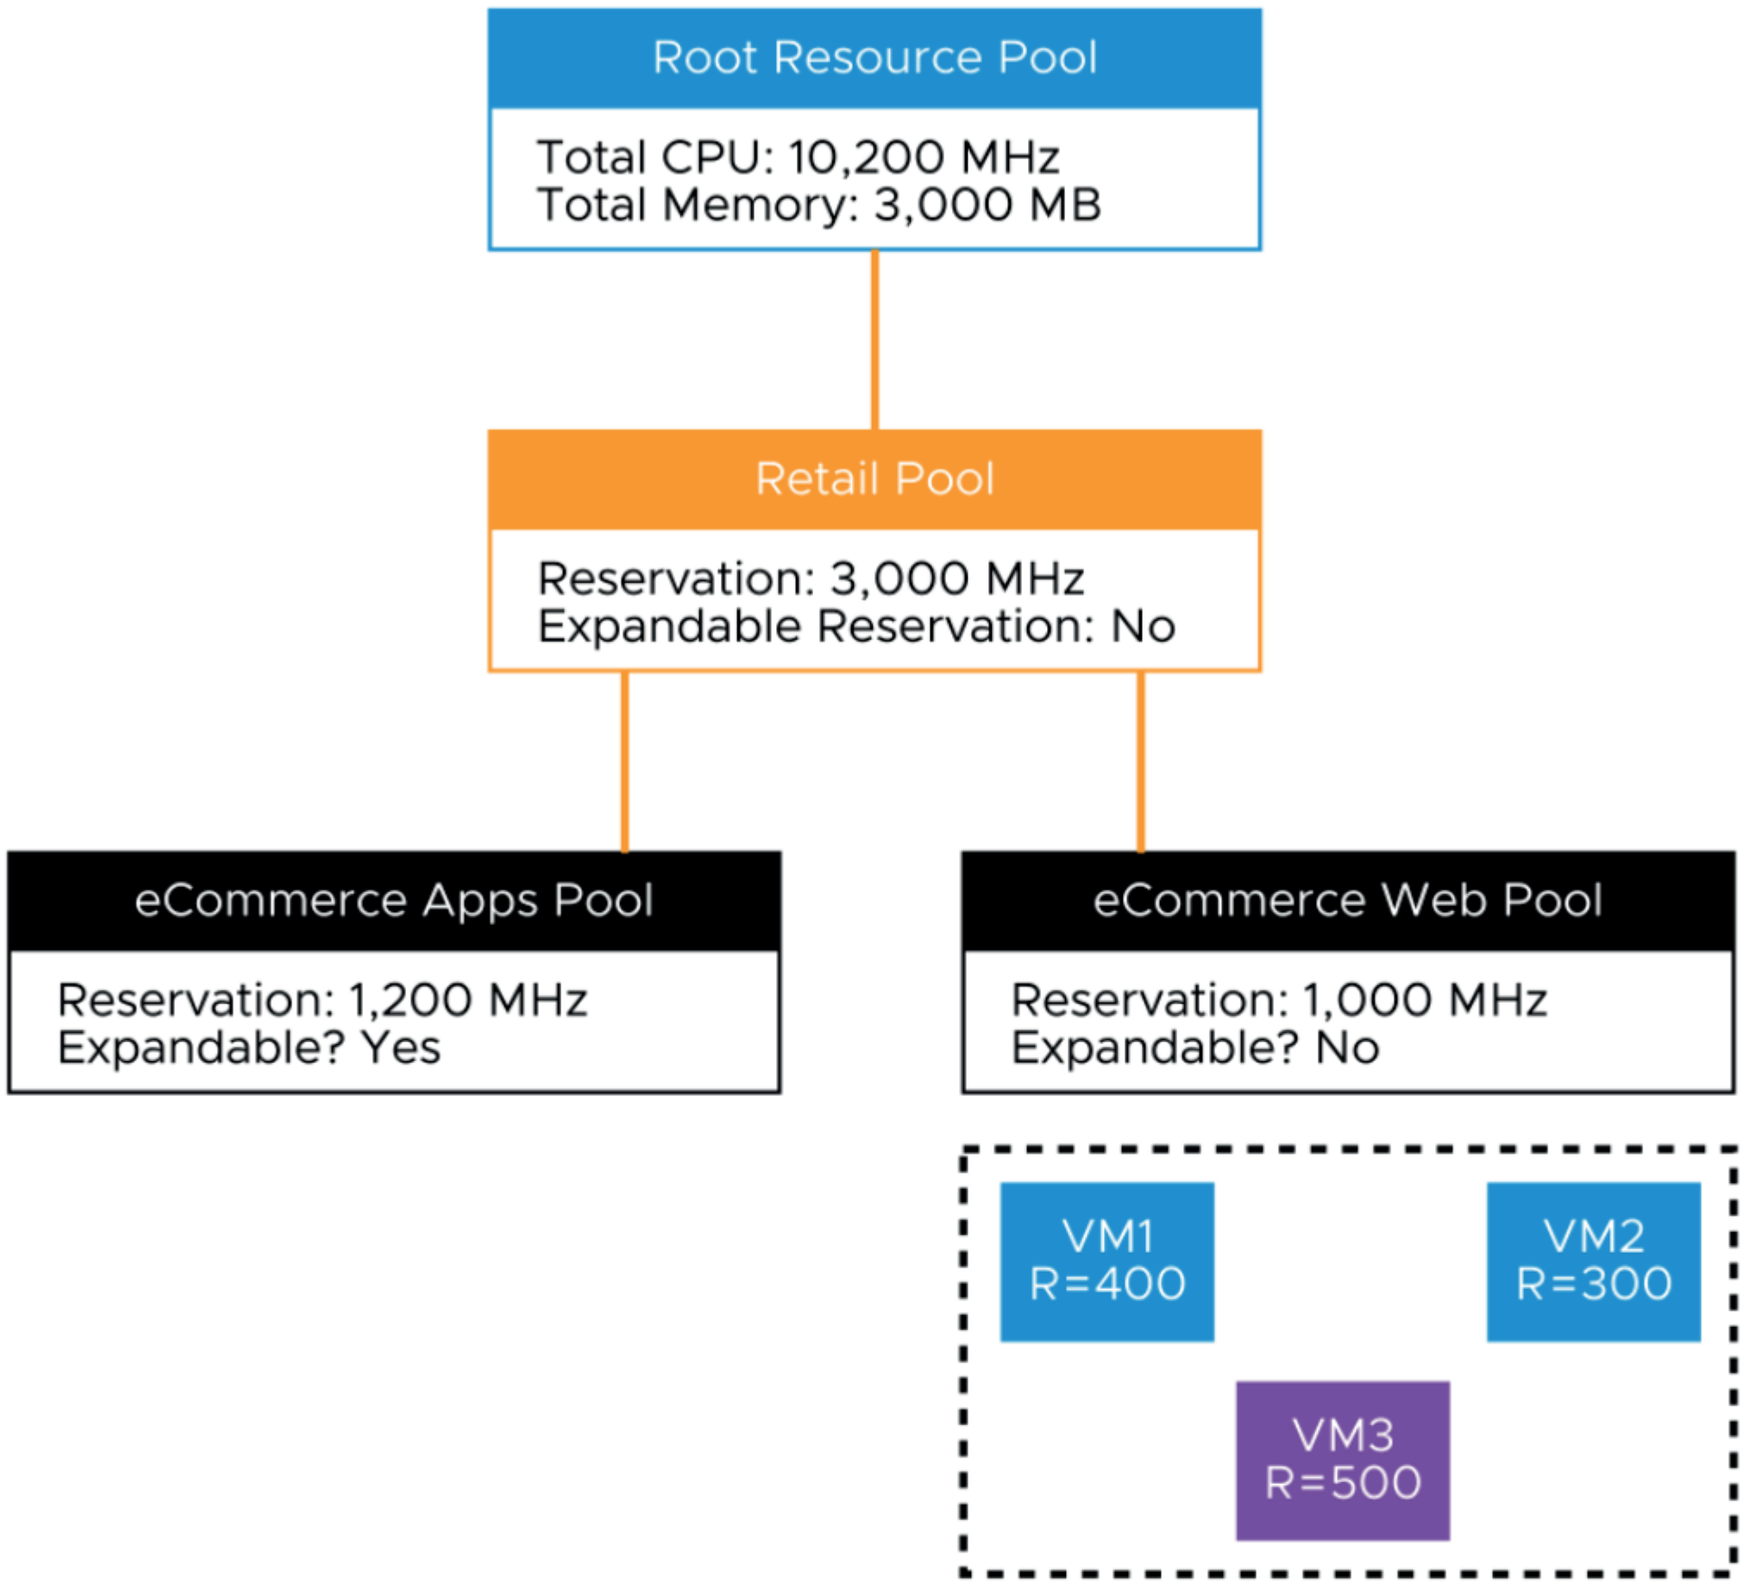

eCommerce resource pools reserve 2,200 MHz of the 3,000 MHz that the Retail pool has reserved.

You power on VMs in the eCommerce Web pool.

With the Expandable Reservation option deactivated on the eCommerce Web pool, VM3 cannot start with a reservation of 500 MHz:

- Lower the VM reservation.
- Enable **Expandable** as the Reservation Type.
- Increase the eCommerce Web pool's reservation.

Technet24

3-16 Example of Expandable Reservation (2)

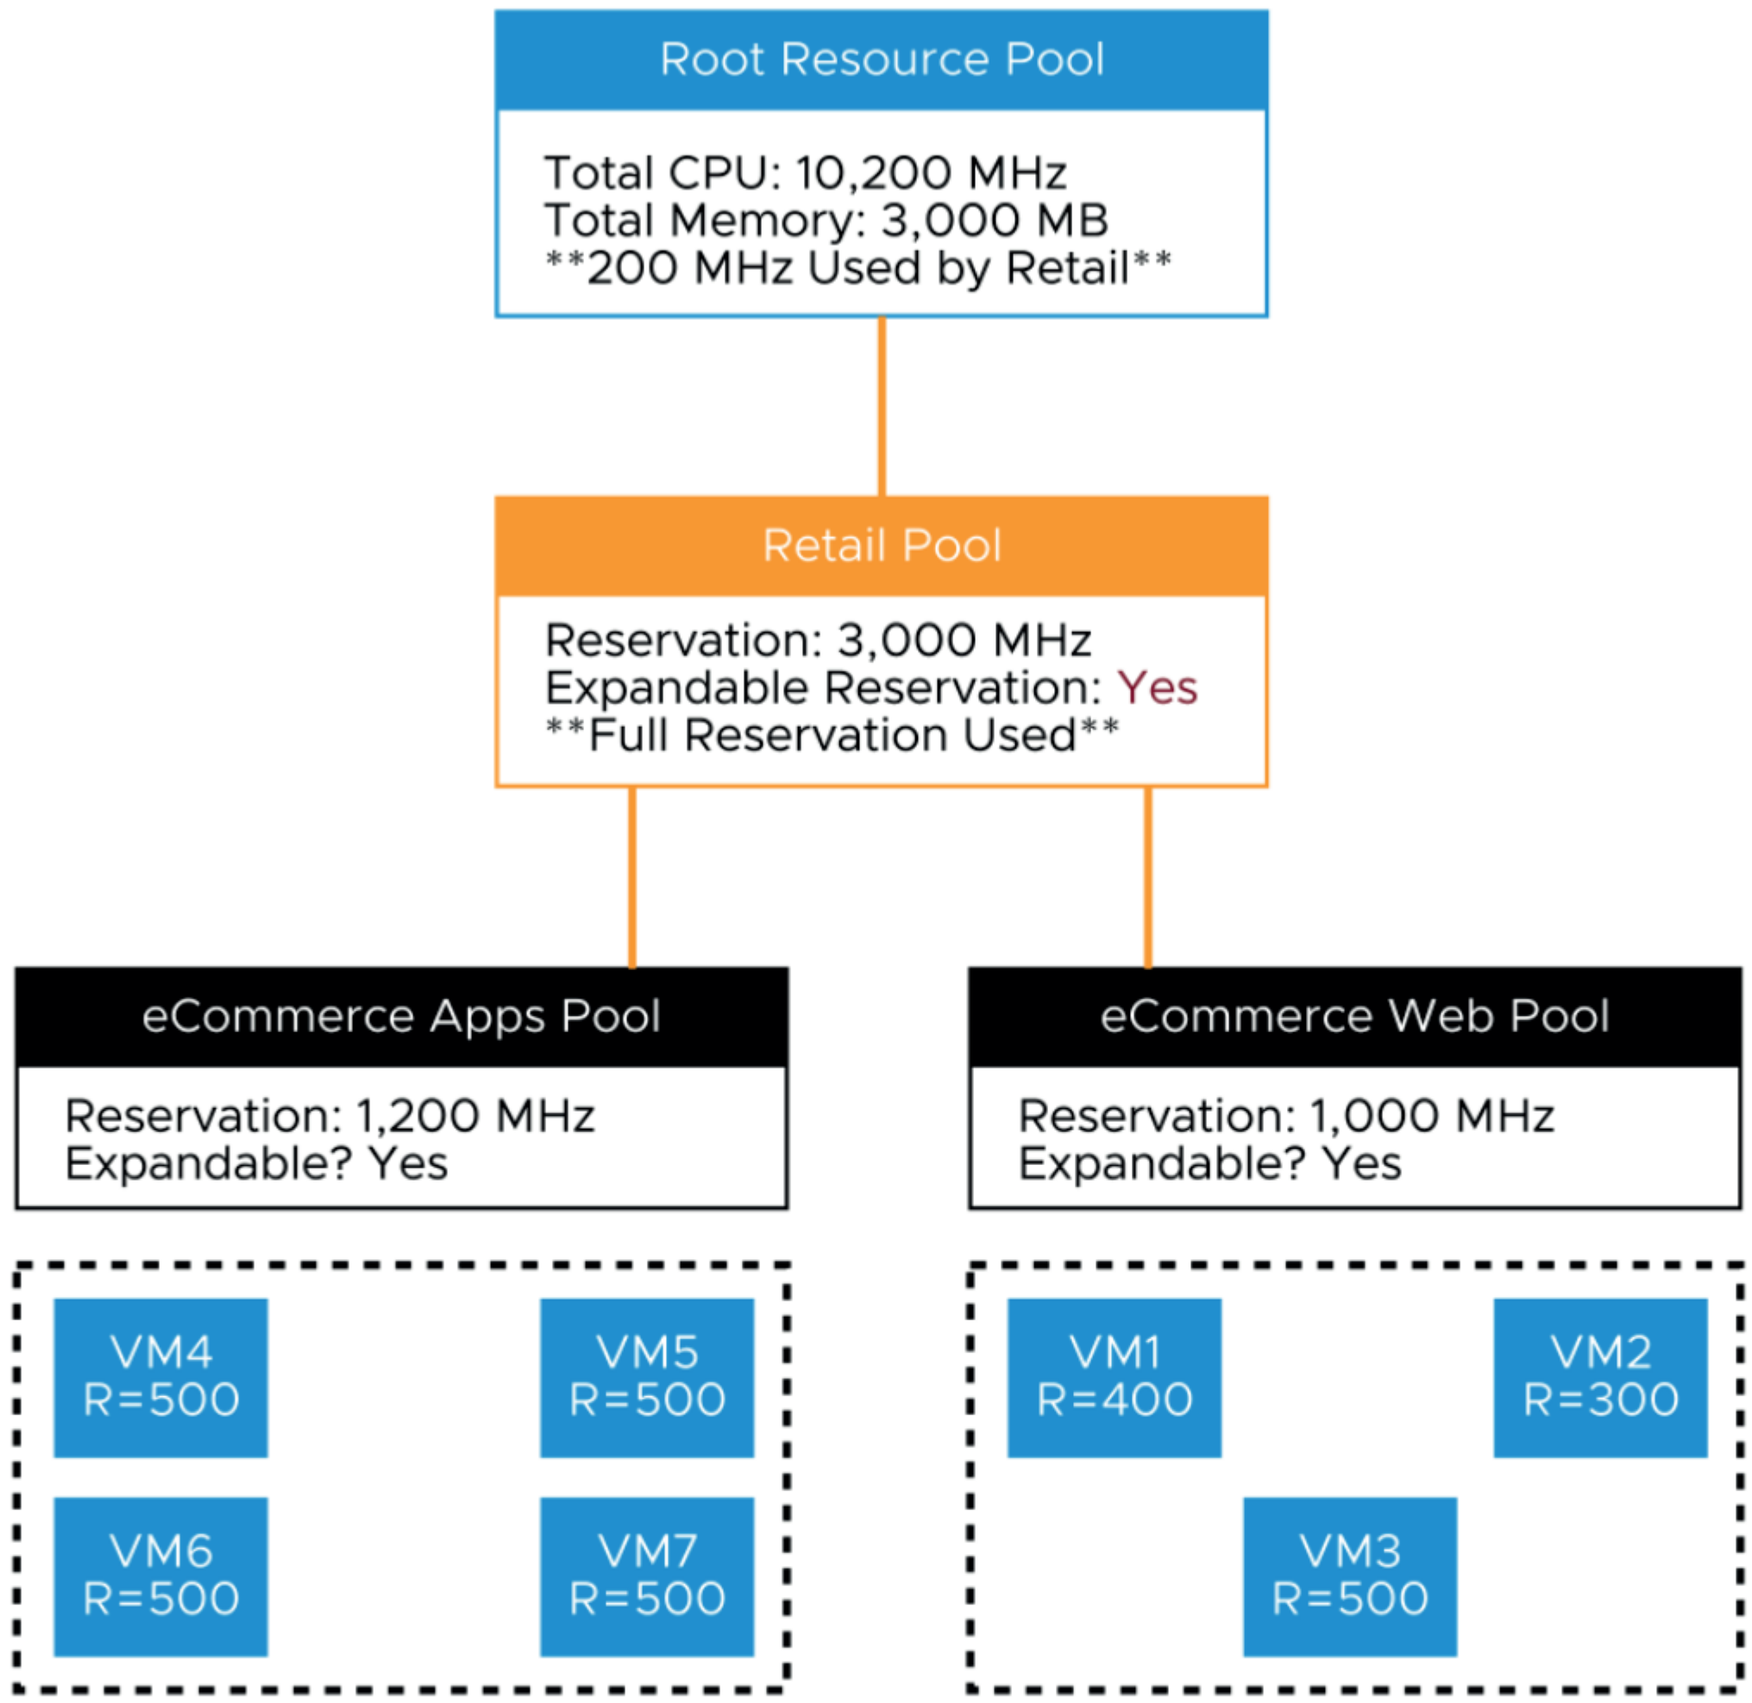

The **Expandable** Reservation Type is enabled on the eCommerce Web pool.

The system considers the resources available in the child resource pool and its direct parent resource pool.

The VM's reservation is charged against the reservation for eCommerce Web pool.

eCommerce Web pool's reservation is charged against the reservation for Retail.

3-17 Resource Pool Summary Tab

The resource pool's **Summary** tab displays information that applies to the ESXi and its resources.

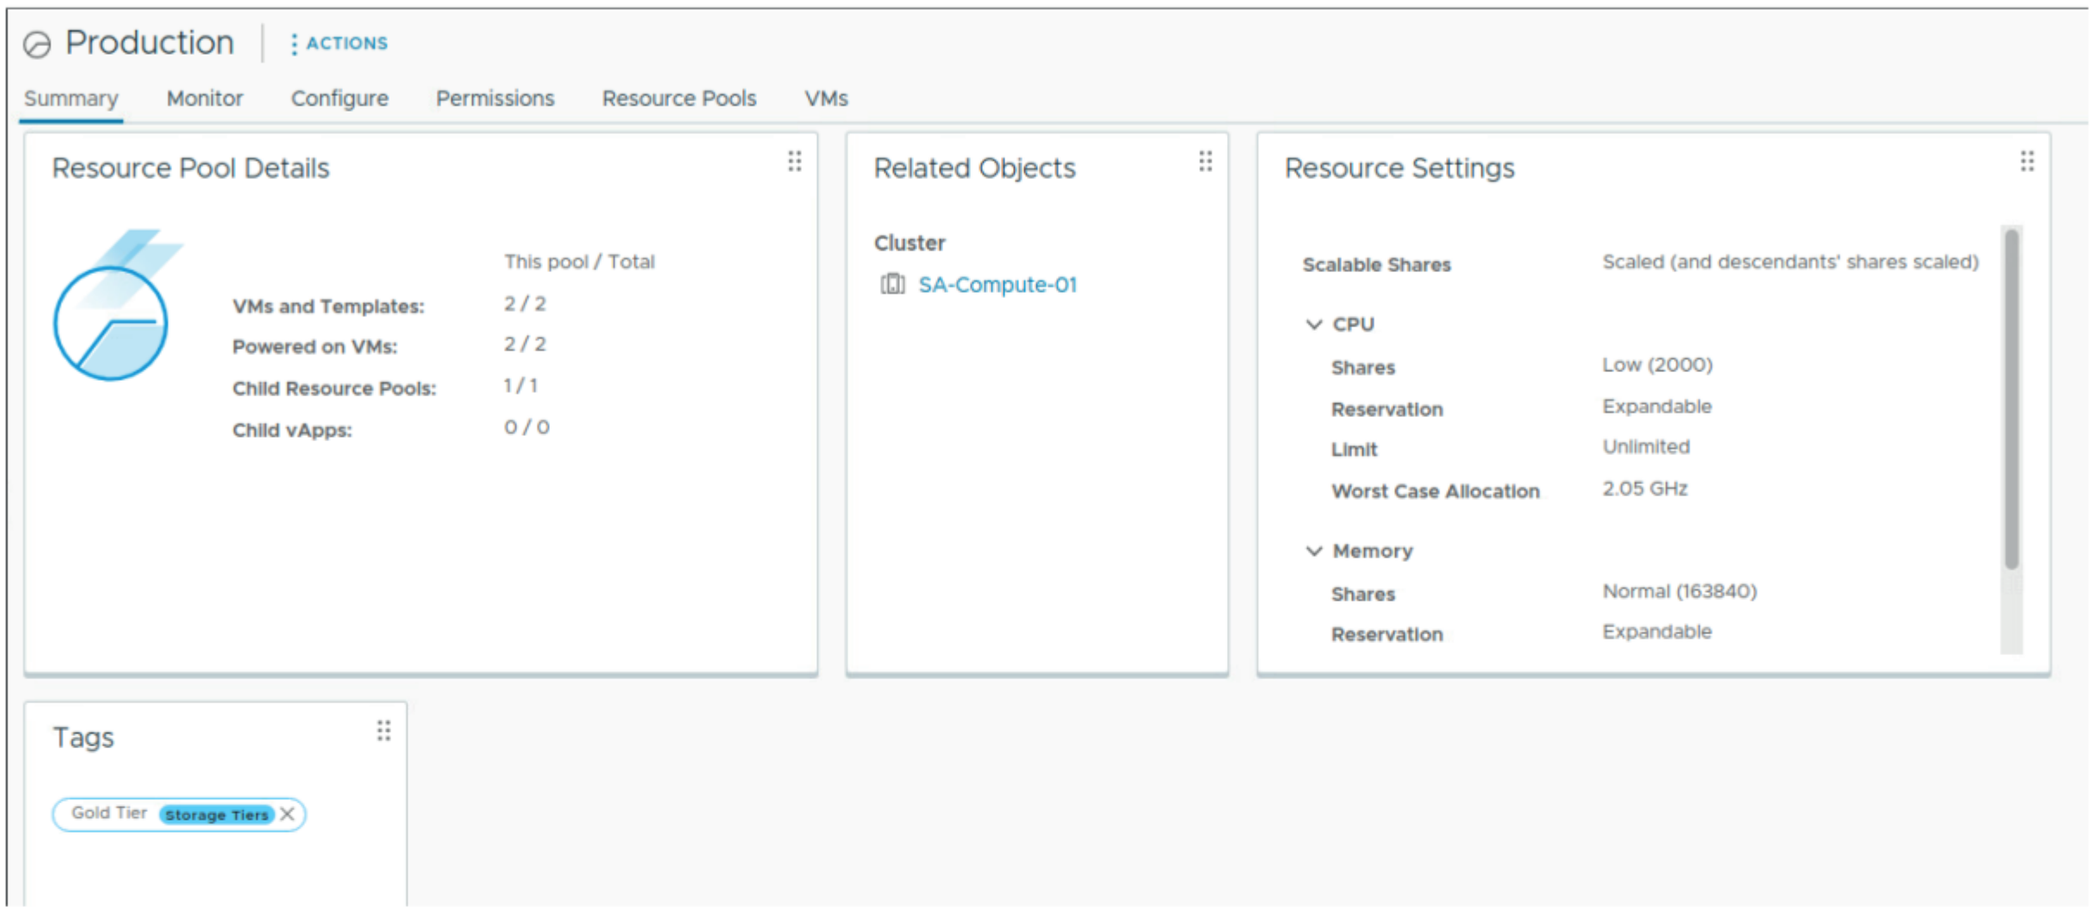

The **Summary** tab provides important information:

- The Resource Settings pane displays CPU and Memory share settings.
- The Resource Consumers pane displays the number of VMs, number of powered-on VMs, and child resource pools that are in the selected resource pool.
- The Tags pane displays tags assigned to objects that reside in the resource pool.

Technet24

3-18 Resource Reservation Details

In the vSphere Client, you can view information about a resource pool's CPU, memory, and storage resources.

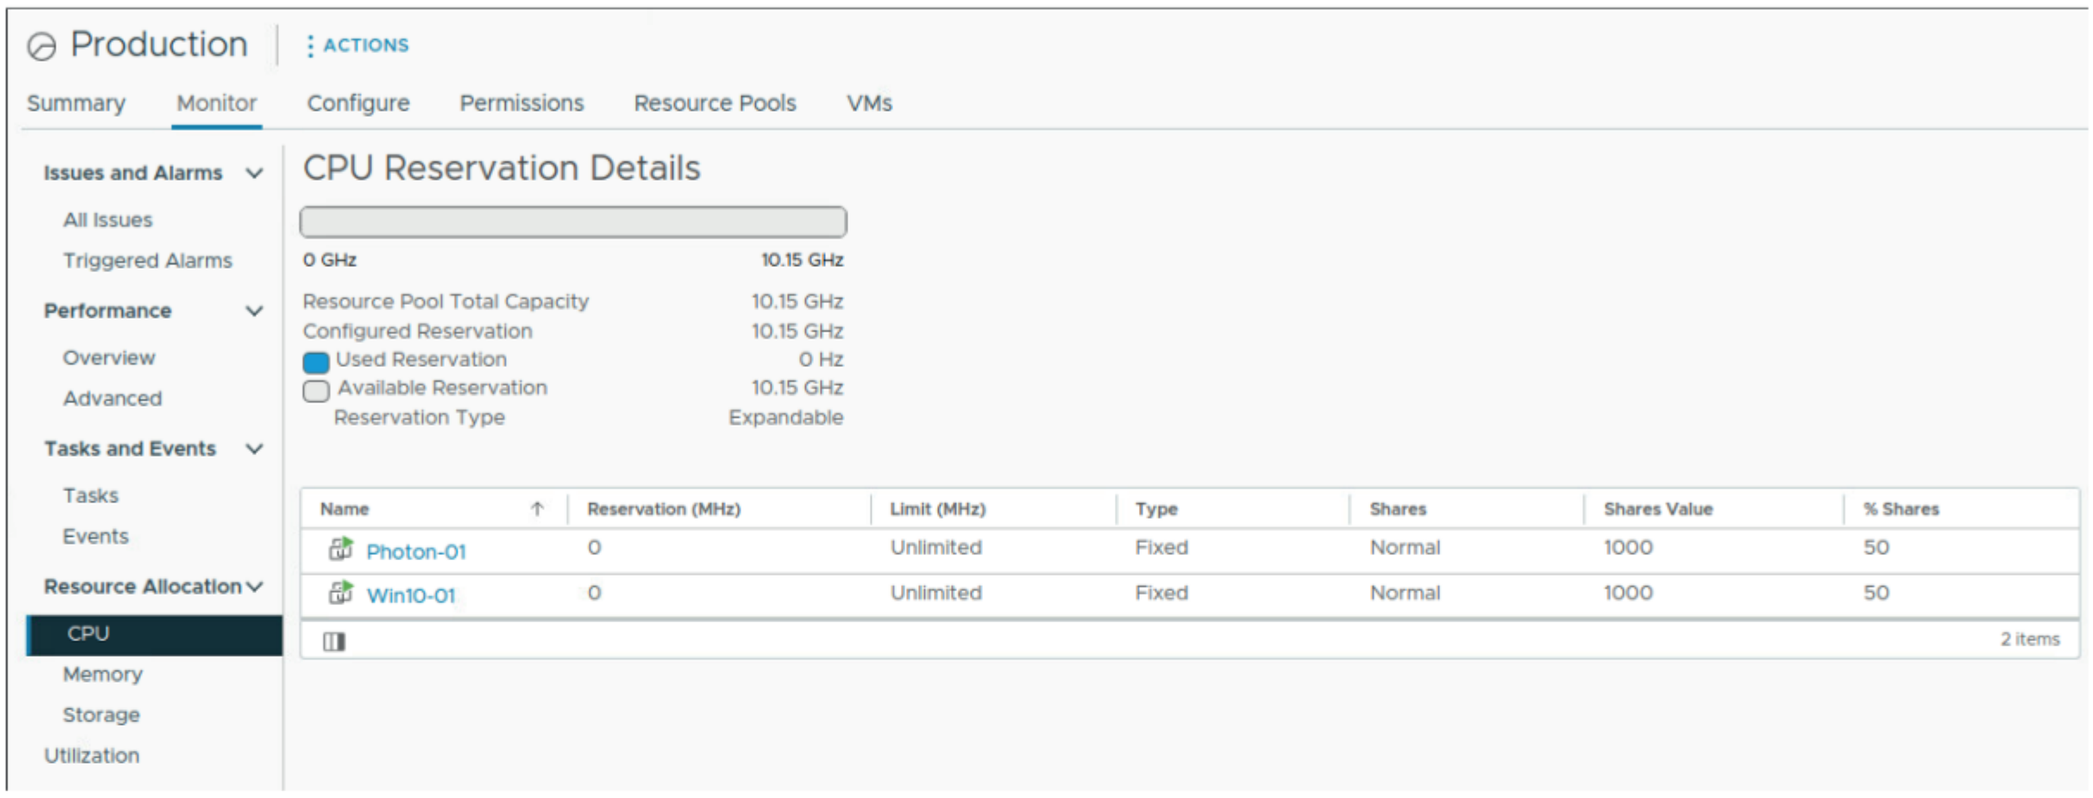

For the resource pool's CPU and memory resources, the following information is provided:

- Amount of CPU and memory reservation that is configured
- Type of reservation used (expandable or fixed)
- Amount of reservation in use by VMs and child pools
- Amount of CPU and memory that is available to be reserved

3-19 Resource Pool Shares

A resource pool is a logical abstraction of hierarchically managed CPU and memory resources. Shares specify the resource pool priority relative to other resource pools.

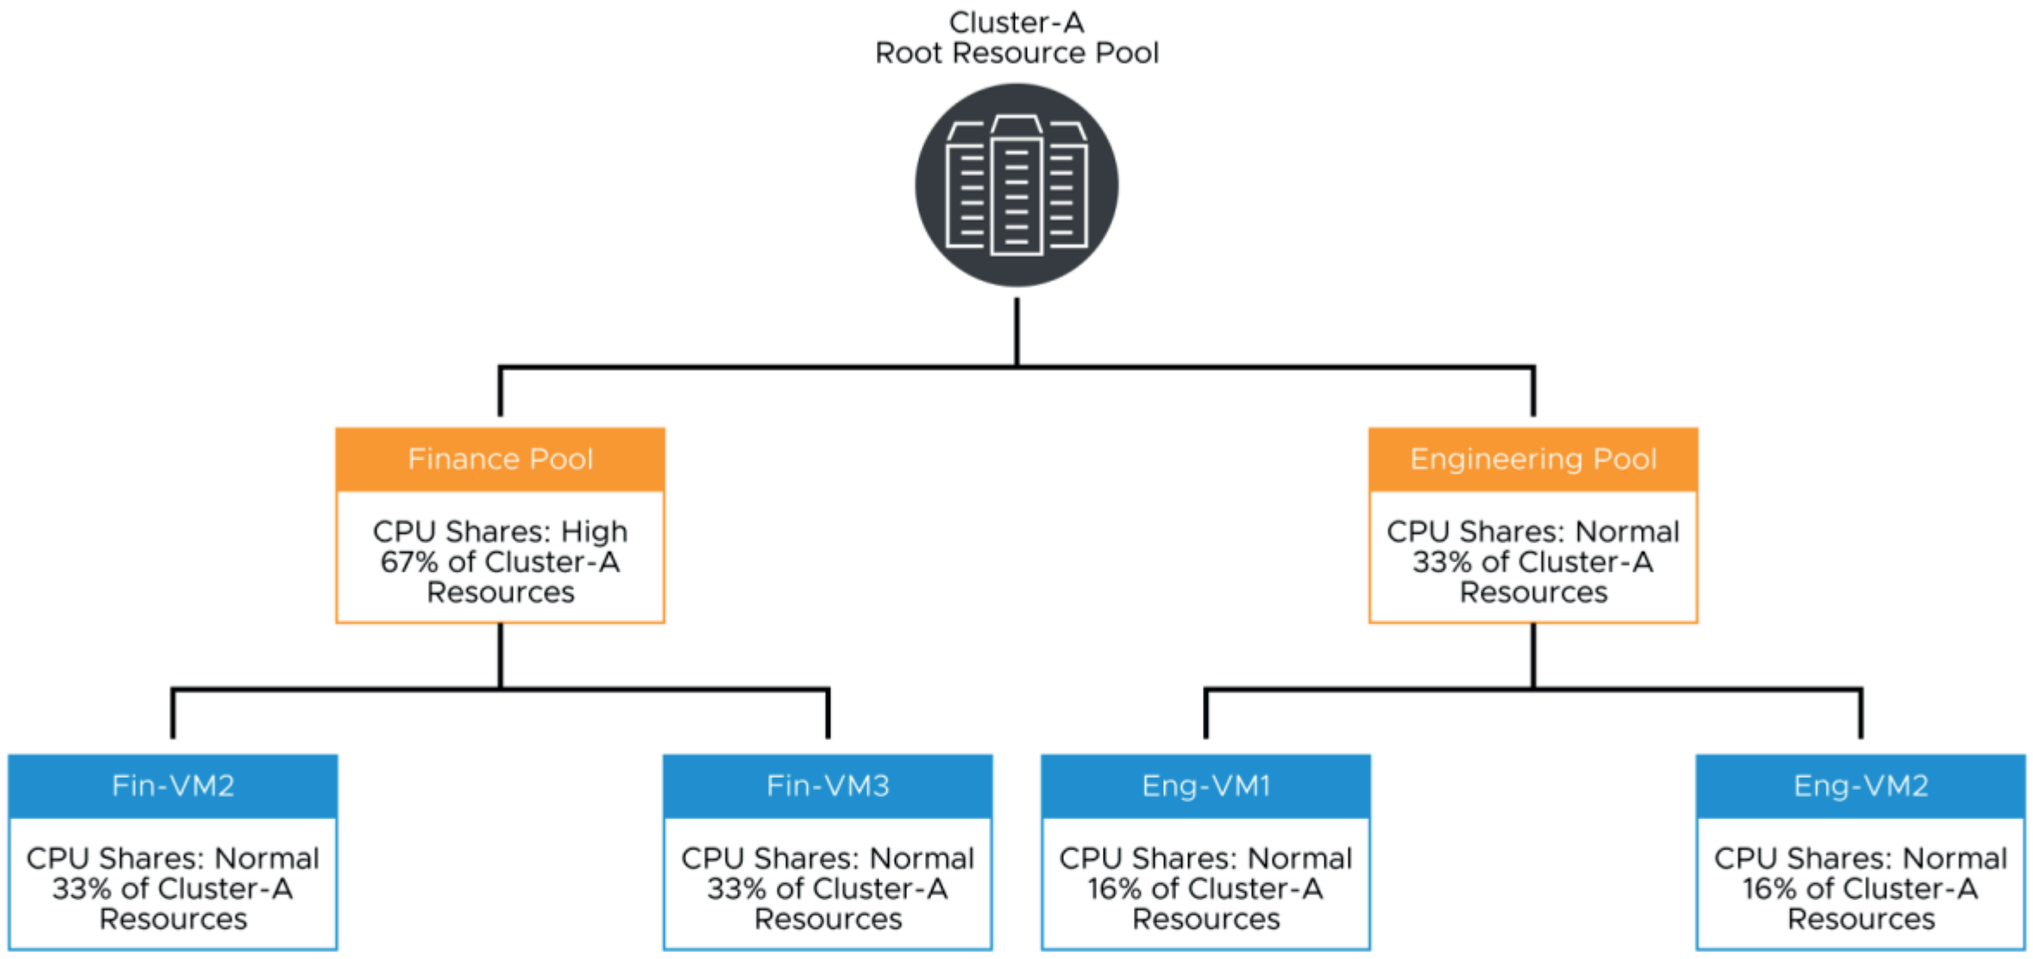

Resource pools can be organized hierarchically. The root resource pool is the topmost resource pool. The root resource pool includes the sum for all CPUs (in megahertz) and the sum of all the installed RAM (in megabytes) available in the compute environment (standalone host or cluster).

A child resource pool is used to allocate resources from the parent resource pool for the child's consumers. A child resource pool cannot exceed the capacity of the parent resource pool. By creating a child pool, you reserve resources from the parent pool. In the example, the root resource pool is Cluster-A. Cluster-A has two child resource pools: the Engineering pool and the Finance pool.

Shares specify the relative priority or importance of a resource pool. In the example, you want the finance VMs to have twice as much processor time as the engineering VMs. You can configure the Finance pool with twice as many CPU shares as the Engineering pool.

You can also specify shares per VM and distribute the resource allocation within the resource pool. VM shares are relative only to other powered-on VMs within the same resource pool. In the example, all VMs have an equal number of shares. However, the Finance VMs are entitled to about 67% of the cluster processor time. Each Finance VM gets about 33% of the entire cluster CPU resources, or twice as much as for engineering VMs.

Technet24

3-20 Diluted Shares

When you power on more VMs in a resource pool, you can dilute the resource allocation for each VM in the pool.

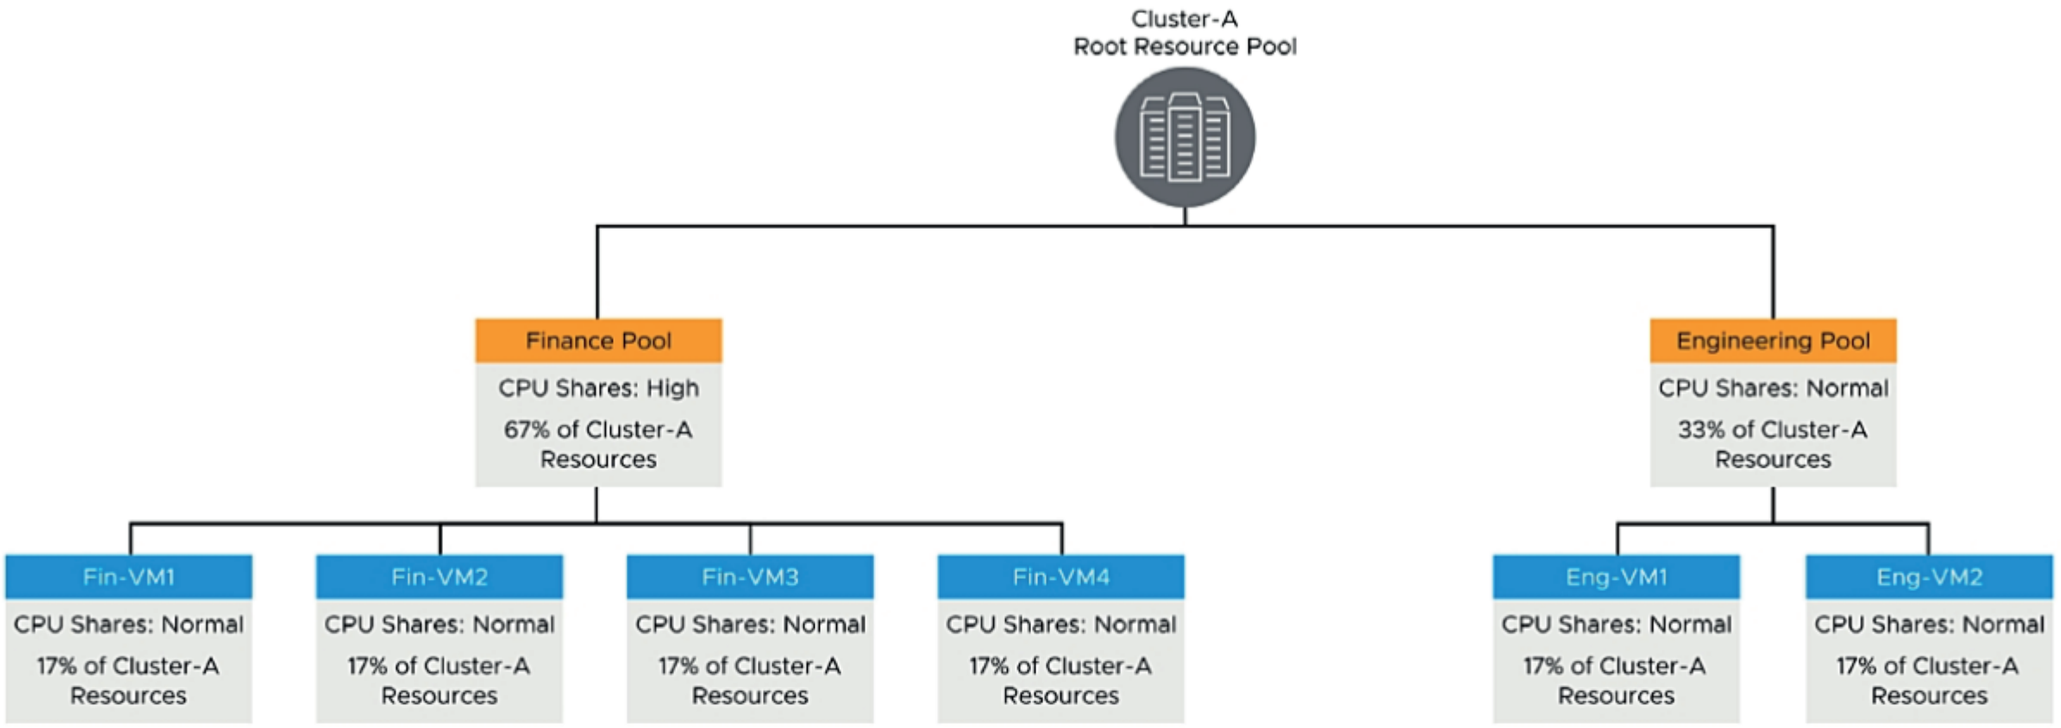

When you power on additional VMs in a resource pool, the resource allocation per VM is recalculated based on the value configured for shares. In the example, two more VMs are powered on in the Finance pool. All VMs in the cluster have an equal number of shares. Therefore, the resources are divided by four.

Because the finance VMs are entitled to about 67% of the cluster processor time, each finance VM gets about 17% of the entire cluster CPU resources. The value of shares within the Finance pool is diluted and is no longer twice as much as the engineering VMs.

3-21 About Scalable Shares (1)

Introduced in vSphere 7, you can enable scalable shares to dynamically scale up and scale down the shares of sibling resource pools and maintain the resource allocation ratio.

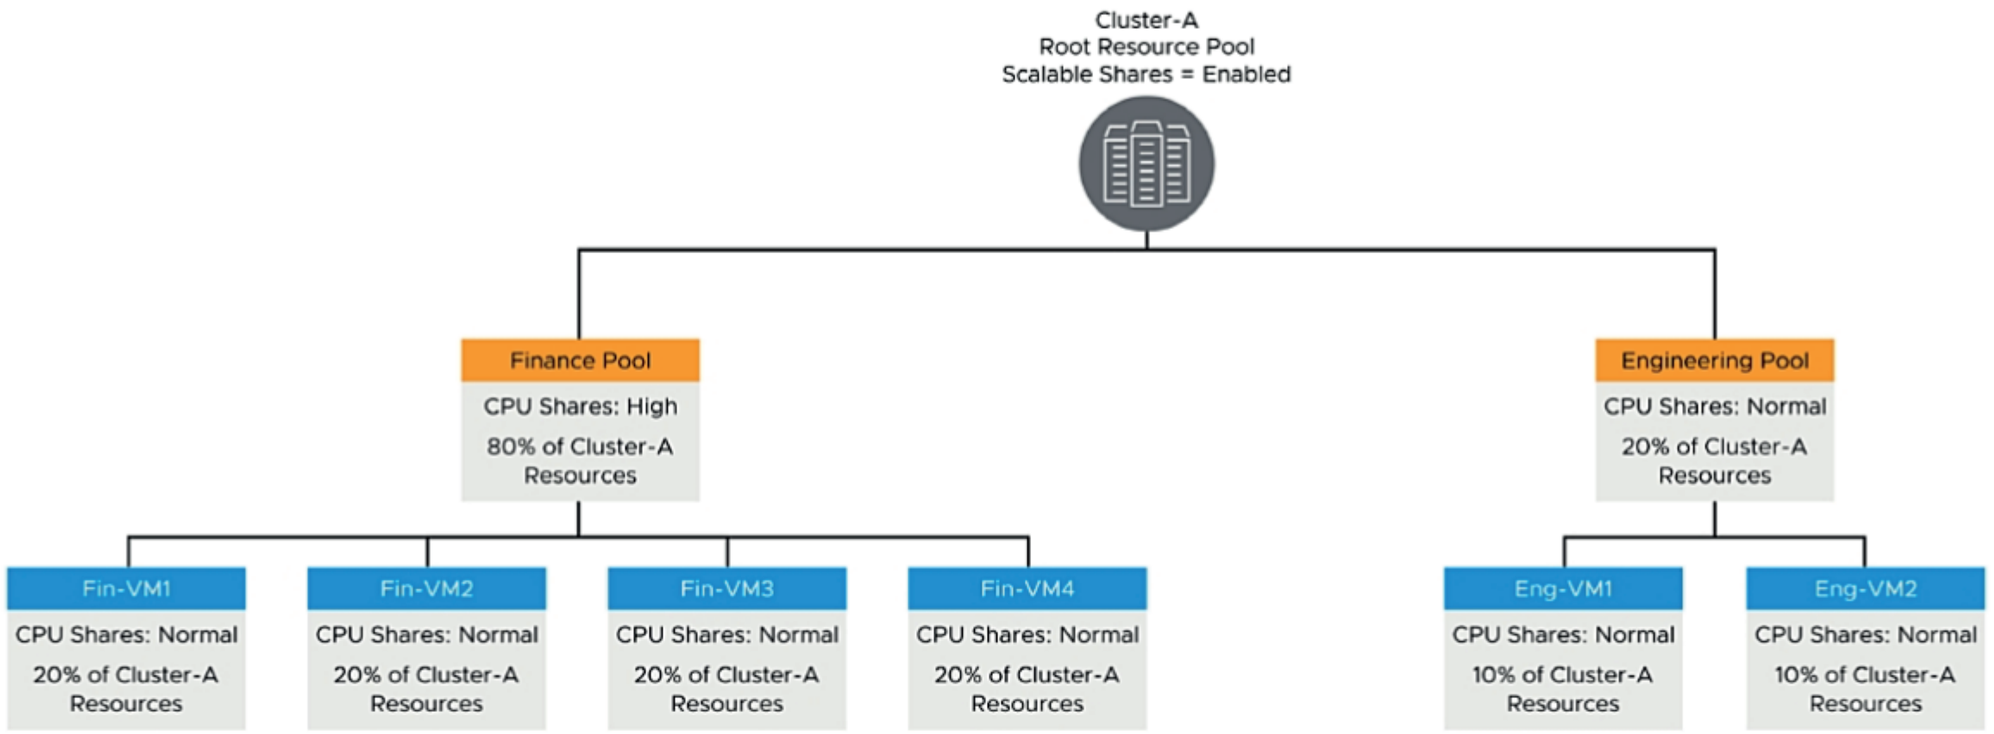

When you enable scalable shares on a resource pool, all child resource pools maintain the original ratio of shares configured in the parent pool.

In the example, you enable scalable shares on Cluster-A to dynamically adjust the resource allocation of child resource pools. All VMs have an equal number of shares, so the shares configuration of the child resource pools adjusts to give every finance VM twice as much as the engineering VMs. The shares configuration of the child resource pools adjusts to 80% for the Finance pool, and 20% for the Engineering pool. Each finance VM gets 20% of the Cluster-A processor time, and the engineering VMs get 10%.

Technet24

3-22 About Scalable Shares (2)

The scalable shares feature makes the child resource pools scalable:

- DRS dynamically scales up and scales down the shares of the resource pools to maintain a resource allocation ratio.
- The scalable shares algorithm maintains the resource allocation ratio by summing the shares of all entities in the resource pool and multiplying by the user-specified initial shares value.
- Scalable shares make the child resource pool scalable. It does not make the parent resource pool scalable.

3-23 Enabling Scalable Shares

You can enable scalable shares at the cluster or resource pool level. The resource allocation of child resource pools becomes scalable.

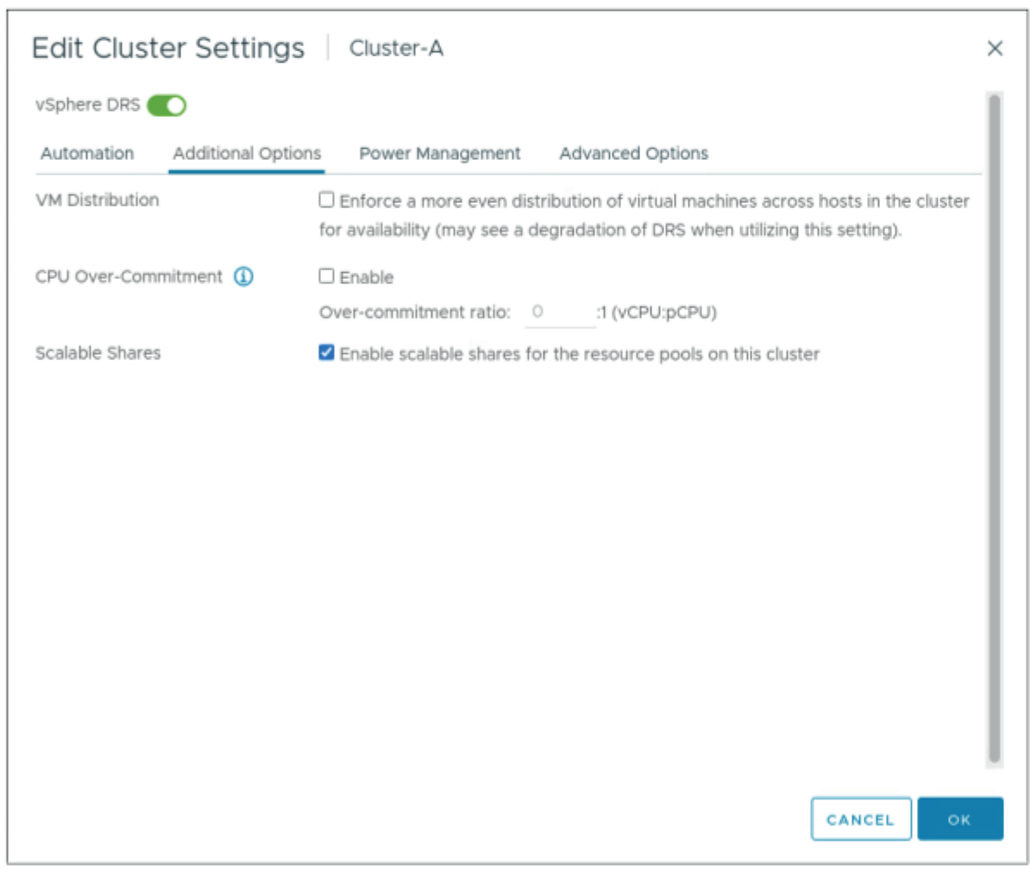

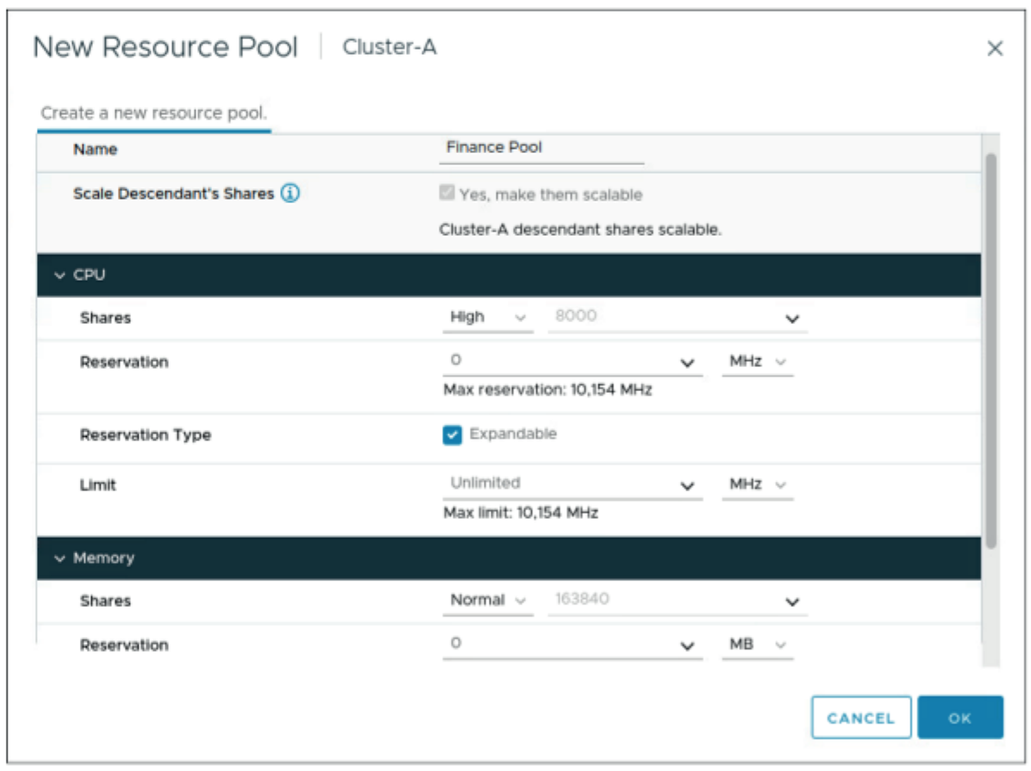

3-24 Lab 5: Managing Resource Pools

Create and use resource pools:

1. Re-configure VMs
2. Create CPU Contention
3. Create Resource Pools
4. Verify Resource Pool Functionality
5. Knowledge Check

3-25 Review of Learner Objectives

- Create a resource pool
- Set resource pool attributes
- Describe expandable reservations
- Describe scalable shares

Technet24

3-26 Lesson 2: vSphere Cluster Service (vCLS)

3-27 Learner Objectives

- Describe vCLS VMs
- Identify vCLS workflows
- Describe the function of the vCLS
- Recognize operations that might disrupt the healthy functioning of vCLS VMs

3-28 About vCLS (1)

The vCLS deploys vCLS virtual machines (vCLS VMs) to each vSphere cluster.

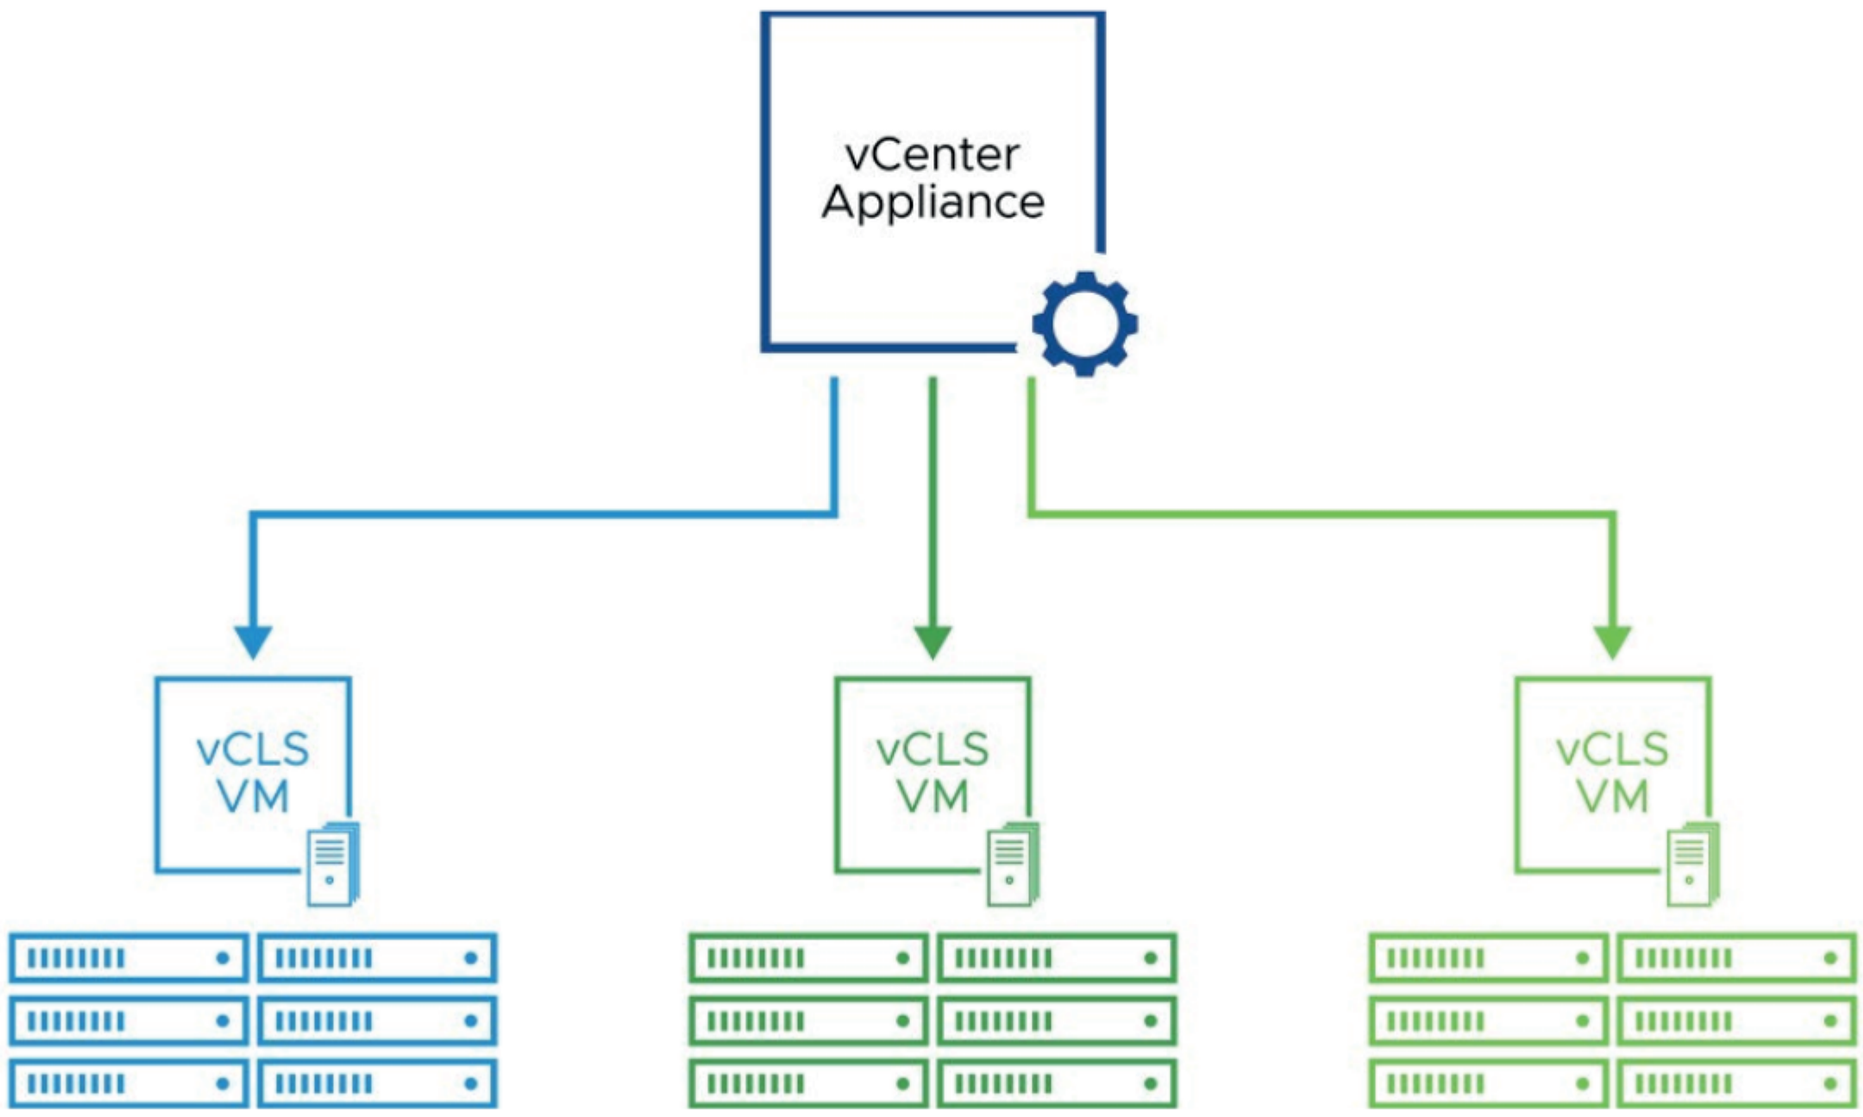

vCLS VMs are deployed to a cluster at creation and after hosts are added to the cluster. In a future release, vCLS VMs will provide vCLS (vSphere HA and vSphere DRS) to workloads even if vCenter is offline.

vCLS VMs are deployed to existing vSphere clusters after vCenter is updated to vCenter 7.0 Update 1.

vSphere Distributed Resource Scheduler (vSphere DRS) cannot function if vCLS VMs are not present in the vSphere cluster.

3-29 About vCLS (2)

vCLS VMs run in every cluster, even if cluster services like vSphere DRS or vSphere HA are not activated on the cluster.

The number of vCLS VMs depends on the number of ESXi hosts in a cluster.

Number of Hosts in a Cluster	Number of vCLS VMs
1	1
2	2
3 or more	3

3-30 About vCLS (3)

The vSphere cluster shows an alert message if healthy vCLS VMs are not available in the cluster.

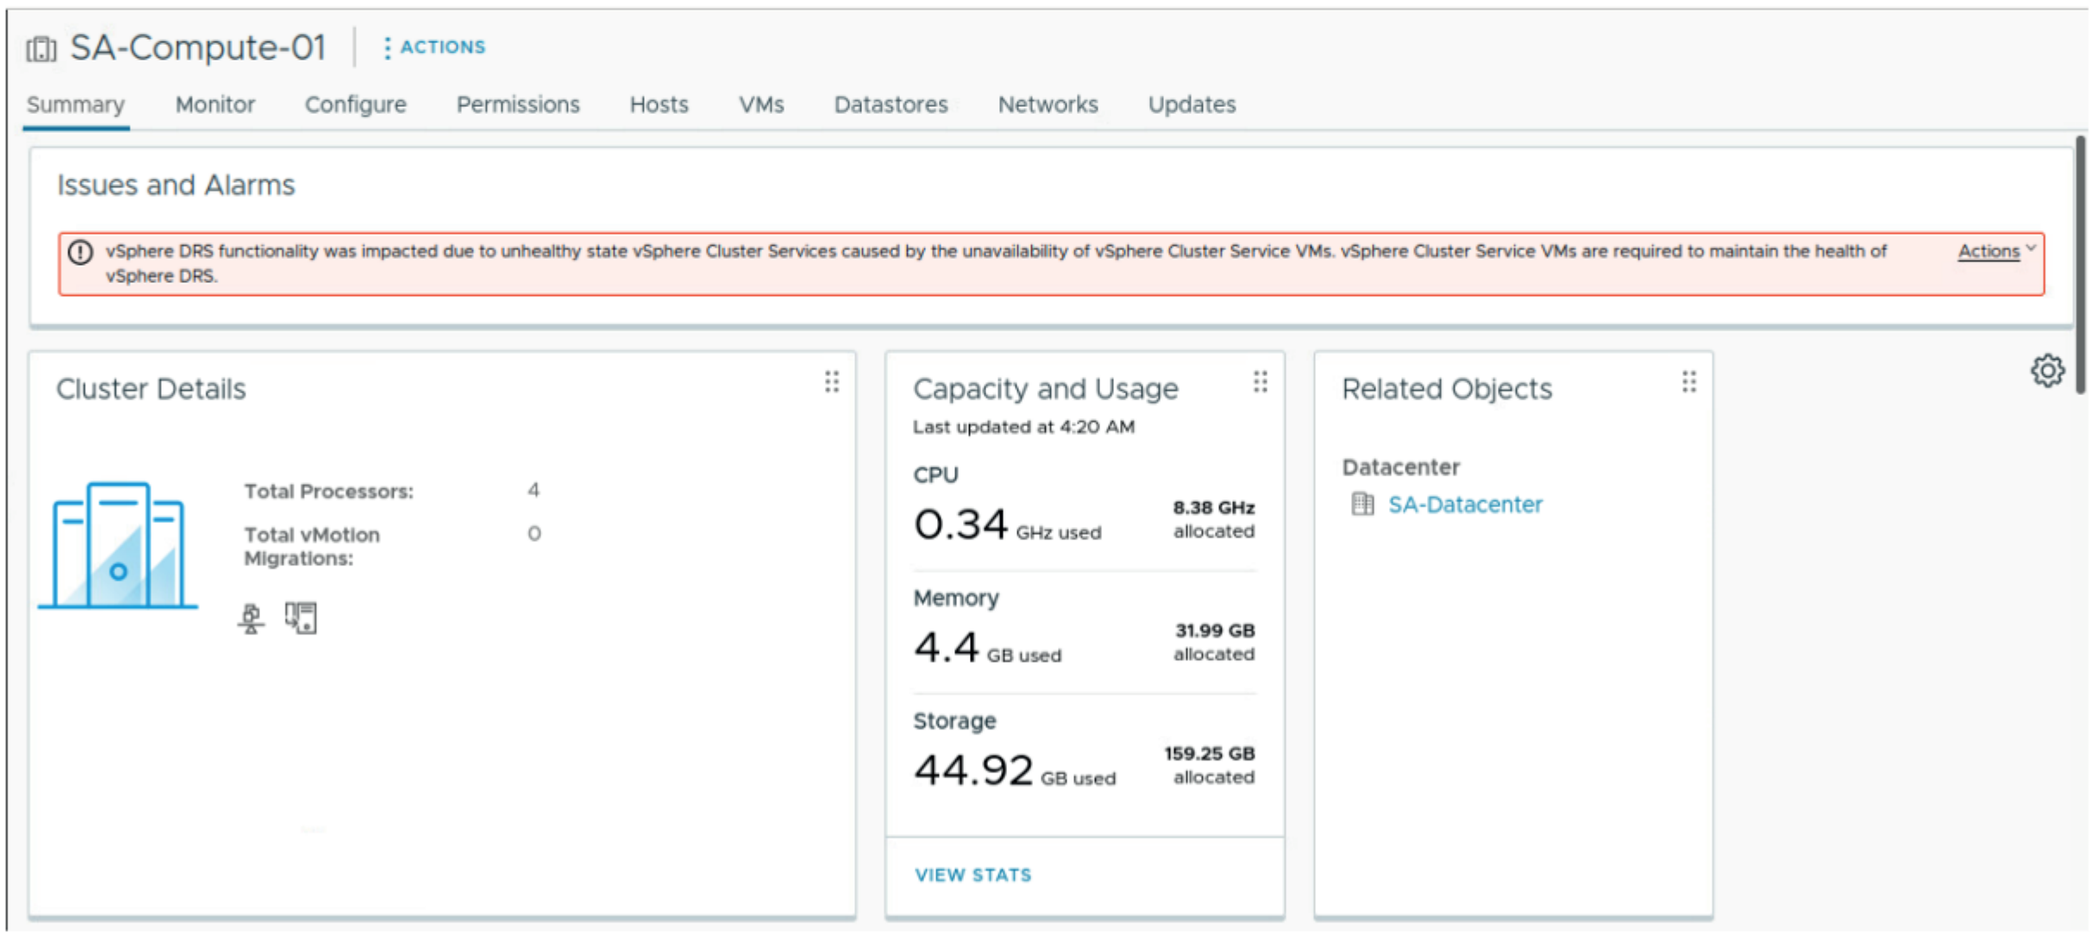

Technet24

3-31 vCLS Components

The vCLS introduces vCLS Manager and vCLS Resource Manager.

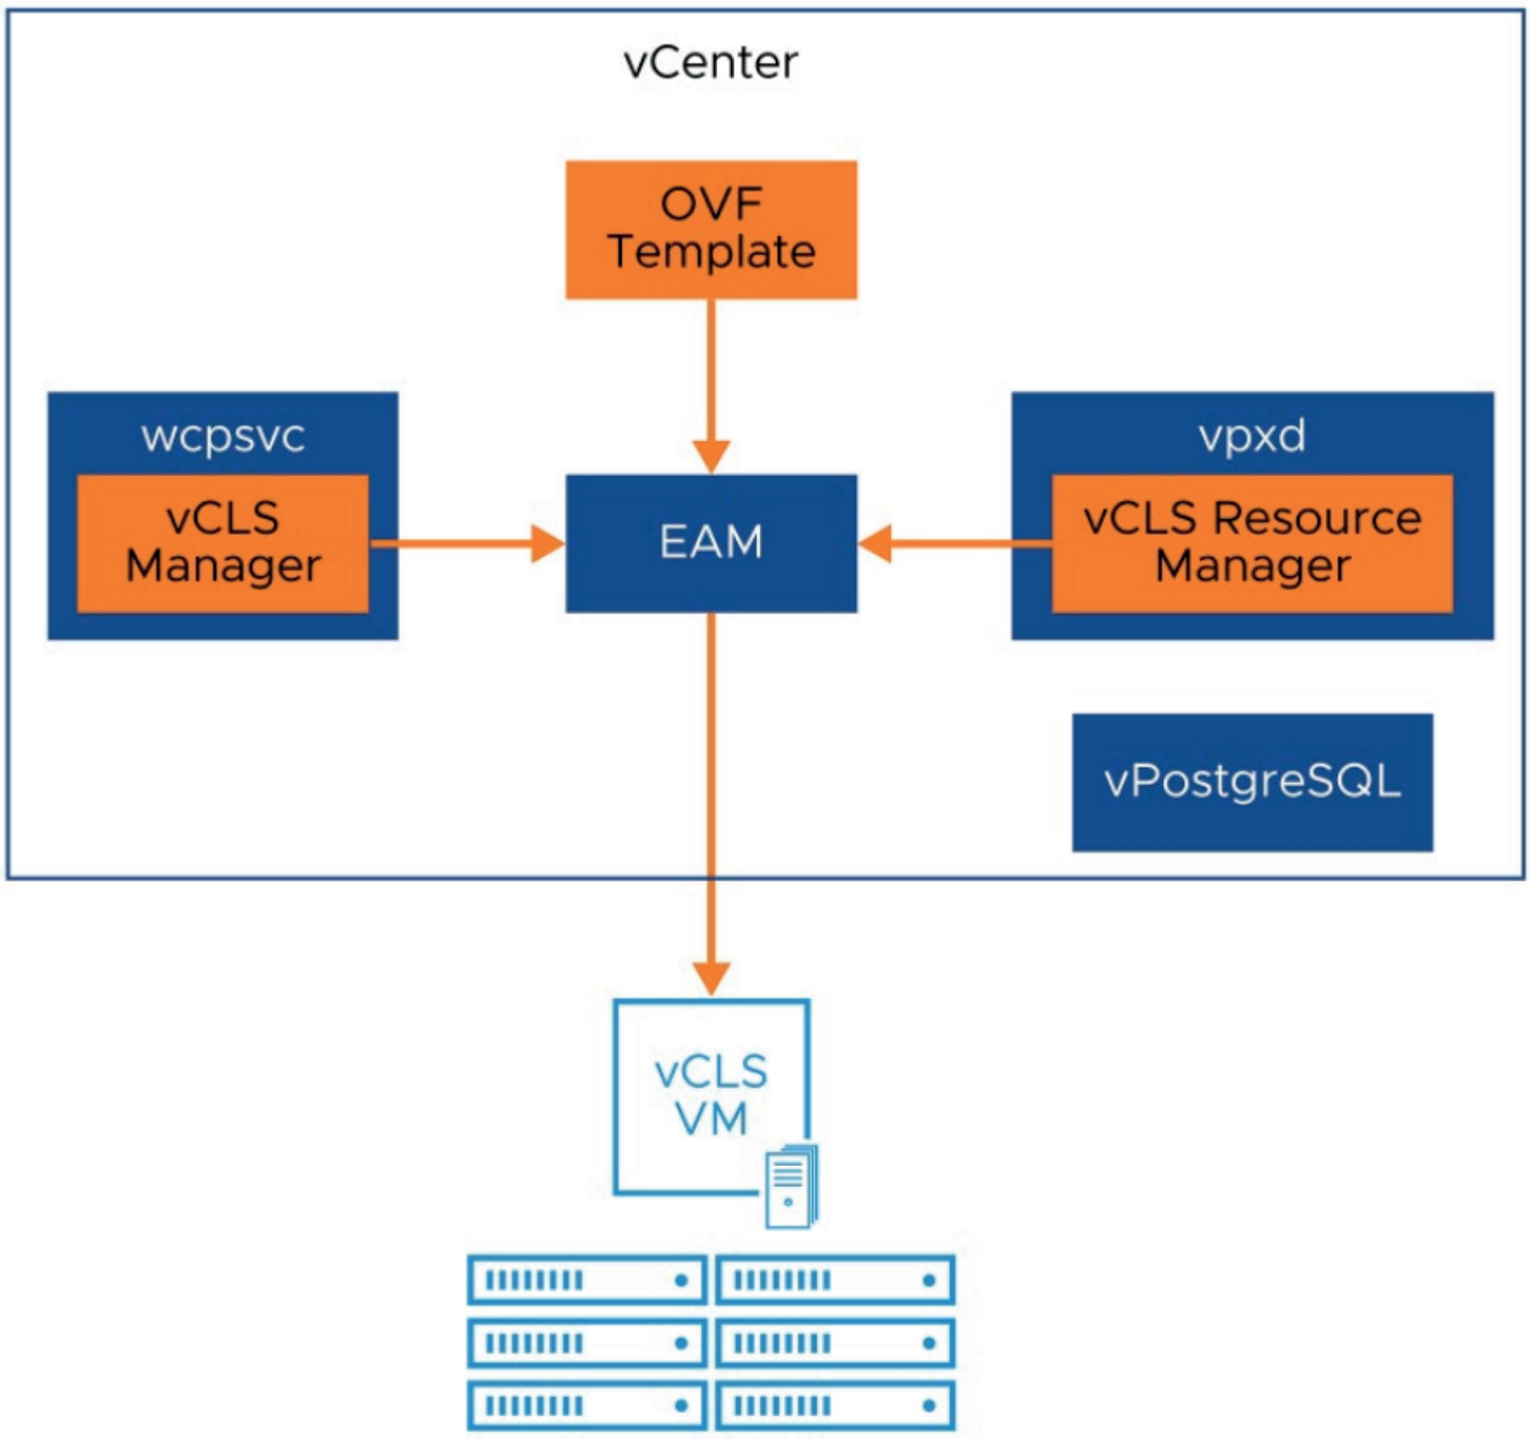

Each vCLS component performs a unique function:

- ESX Agent Manager (EAM):
  - Deploys vCLS VMs to the ESXi hosts in the vSphere cluster
  - Receives VM placement information from vCLS Resource Manager

- vCLS Manager:
  - New module running in the wcpsvc service
  - Manages and monitors a vSphere ESX Agent Manager agency for each set of vCLS VMs
  - Persists the EAM agency information in the vCenter database
  - Customizes each vCLS VM during deployment
  - Performs password rotation for vCLS VMs
  - Operates a desired state model
- vCLS Resource Manager:
  - New module running in the vmware-vpxd service
  - Manages vCLS initial placement and failover placement
- vCLS OVF:
  - Virtual machine OVF template for vCLS VMs.
  - Stored at /storage/lifecycle/vmware-hdcs/ in vCenter.
  - Patches and updates to vCenter may contain updated versions of the OVF. EAM will push the updated vCLS VMs 85out to the cluster and then remove the outdated vCLS VMs

Technet24

3-32 About vCLS VMs (1)

vCLS VMs are present in each vSphere cluster.

vCLS VMs have the following characteristics:

- Managed by vCenter
- Do not have an assigned network interface card (NIC) or IP address
- Deployed to a shared datastore when possible
- Configured with 1 virtual CPU, 128 MB memory, and 2 GB thin-provisioned disk
- Run a minimal installation of Photon OS

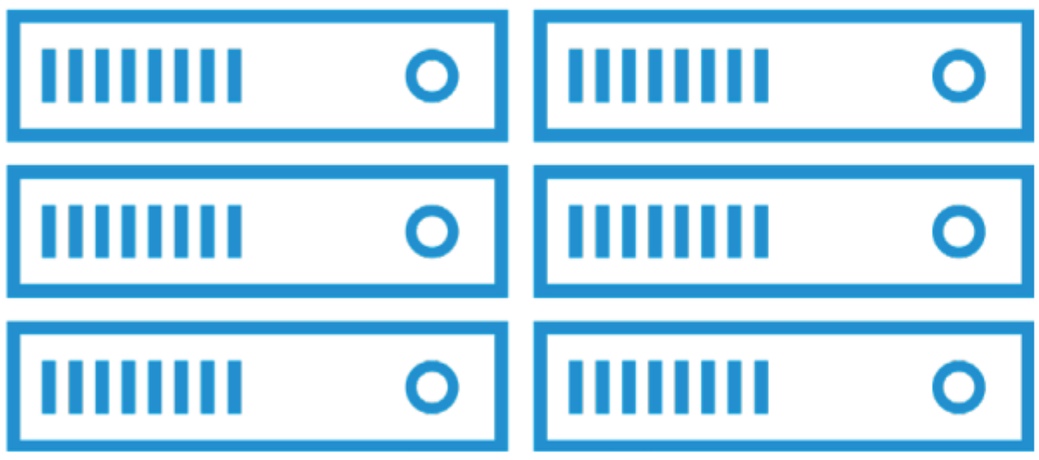

A vCLS is deployed from an OVF with a minimal installed profile of Photon OS. vCLS Manager manage the resources, power state, and availability of these VMs. vCLS VMs are required for maintaining the health and availability of vSphere HA and DRS. Any impact on the power state or resources of these VMs might degrade the health of vCLS and cause vSphere DRS to stop working in the cluster.

Modifying vCLS VMs is not supported.

vCLS VMs are visible in the VMs tab of the host or cluster, from the VM and Templates tab and when connected directly to an ESXi host using VMware Host Client.

When a shared datastore is not available, vCLS VMs are deployed to local datastores. With vSphere 7 Update 3, vCLS agent VM's that are currently running on local storage or any datastore not in the vCLS allowed list, are migrated using Storage vMotion to a vCLS allowed datastore option, as configured by the user.

vCLS VMs are automatically powered on by vCenter, if they are manually powered off.

3-33 About vCLS VMs (2)

In the vSphere Client, vCLS VMs are not visible in the inventory tree of the Hosts and Clusters view.

You view these VMs from the VMs tab in the Hosts and Clusters view. Alternatively, you view vCLS VMs from the VMs and Templates view.

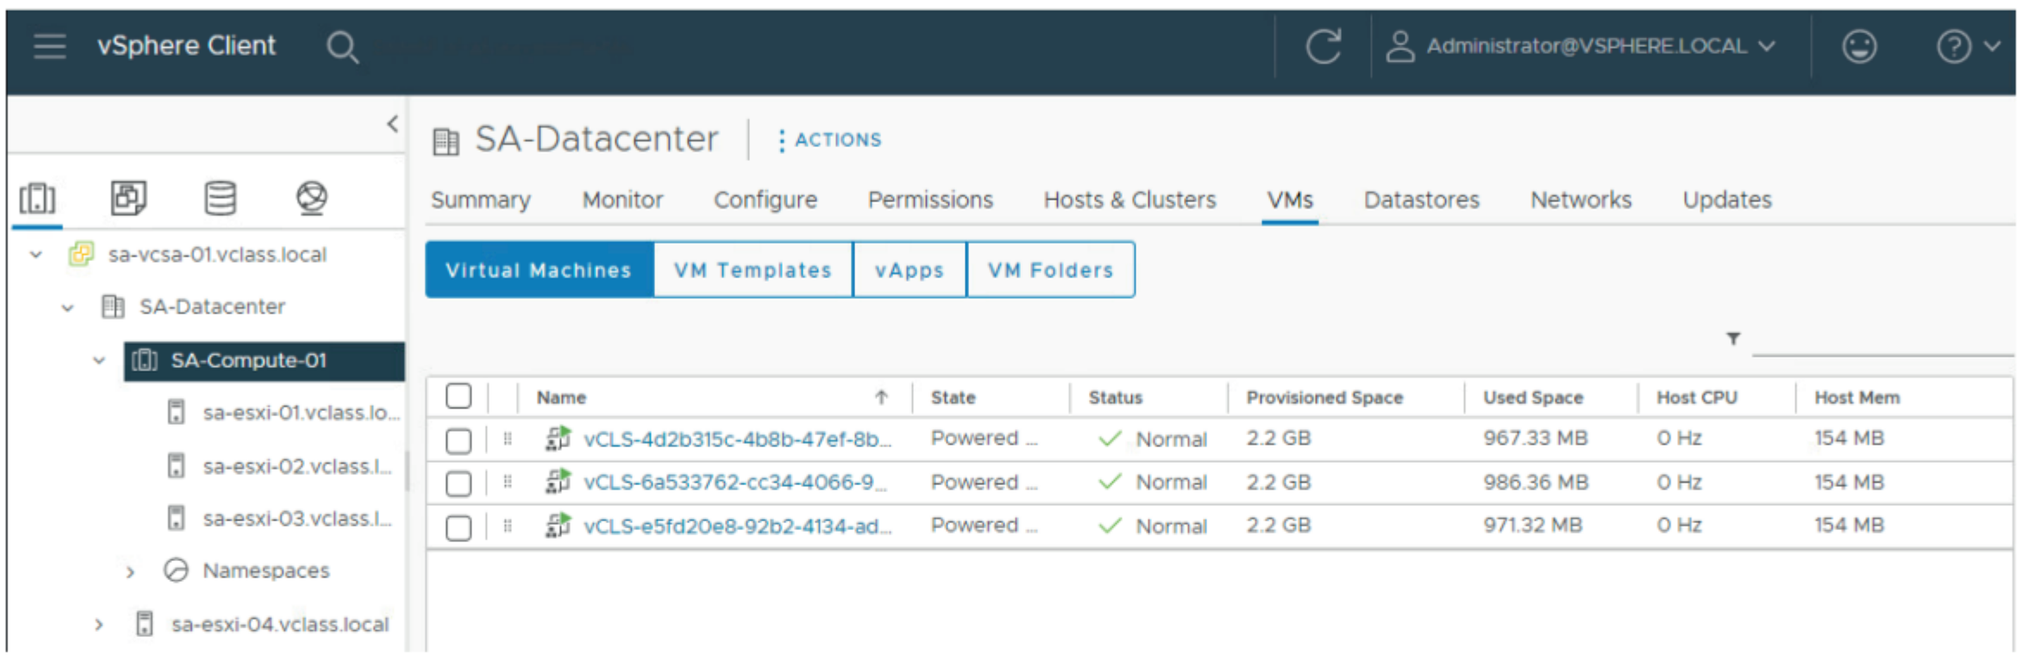

Introduced in vSphere 7 Update 3, the vCLS VMs now have a unique name based up their UUIDs.

3-34 About EAM Agency

The lifecycle for vCLS agent VMs is maintained by the vSphere ESX Agent Manager (EAM).

In the vSphere Client, you select **Administration** > **vCenter Server Extensions** > **vSphere ESX Agent Manager** (EAM) to view the ESX agencies.

The vCLS Agency is automatically added and does not require any configuration.

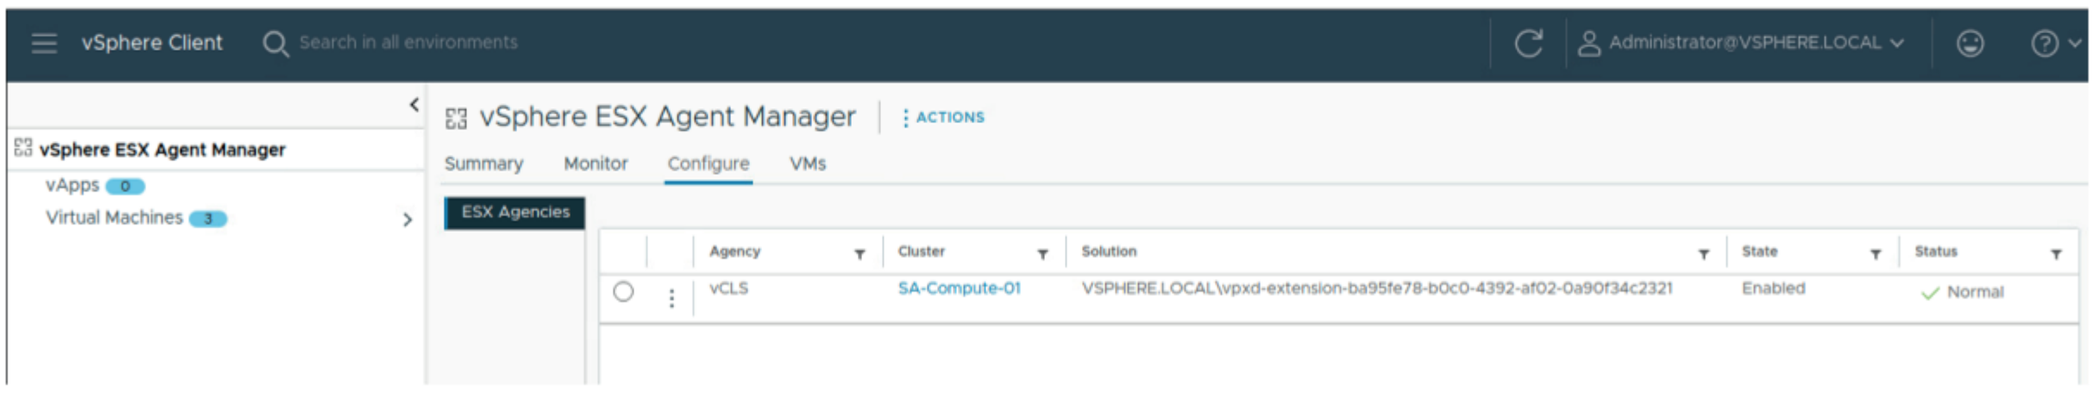

Technet24

3-35 vCLS Cluster Creation Workflow

vCLS Manager deploys vCLS VMs when a new vSphere cluster is created.

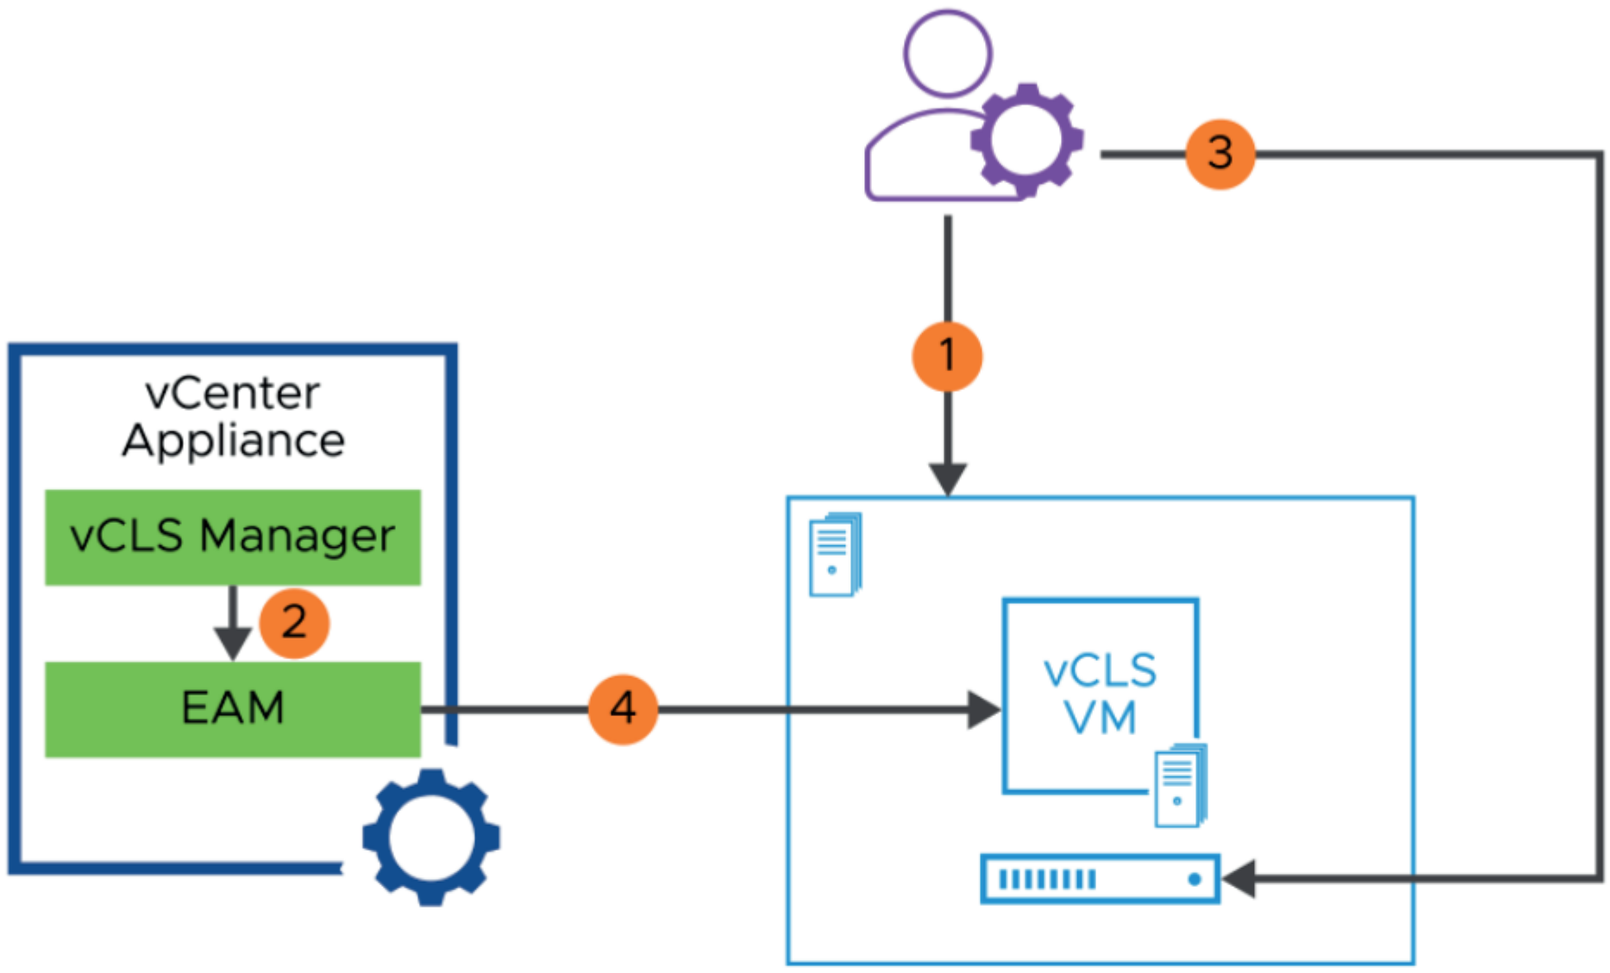

vCLS Manager deploys vCLS VMs when a new vSphere cluster is created:

1. A vSphere administrator creates a vSphere cluster.
2. The vCLS Manager creates an vSphere ESX Agent Manager (EAM) agency associated with the vSphere cluster.
3. The vSphere administrator adds an ESXi host to the cluster.
4. EAM deploys a vCLS to the ESXi host.

Additional vCLS VMs are deployed as additional ESXi hosts are added to the vSphere cluster. A maximum of three vCLS VMs are deployed and maintained in a single vSphere cluster.

3-36 vCLS Cluster Upgrade Workflow

The vCLS automatically updates vCLS VMs when new updates to the vCLS OVF template become available.

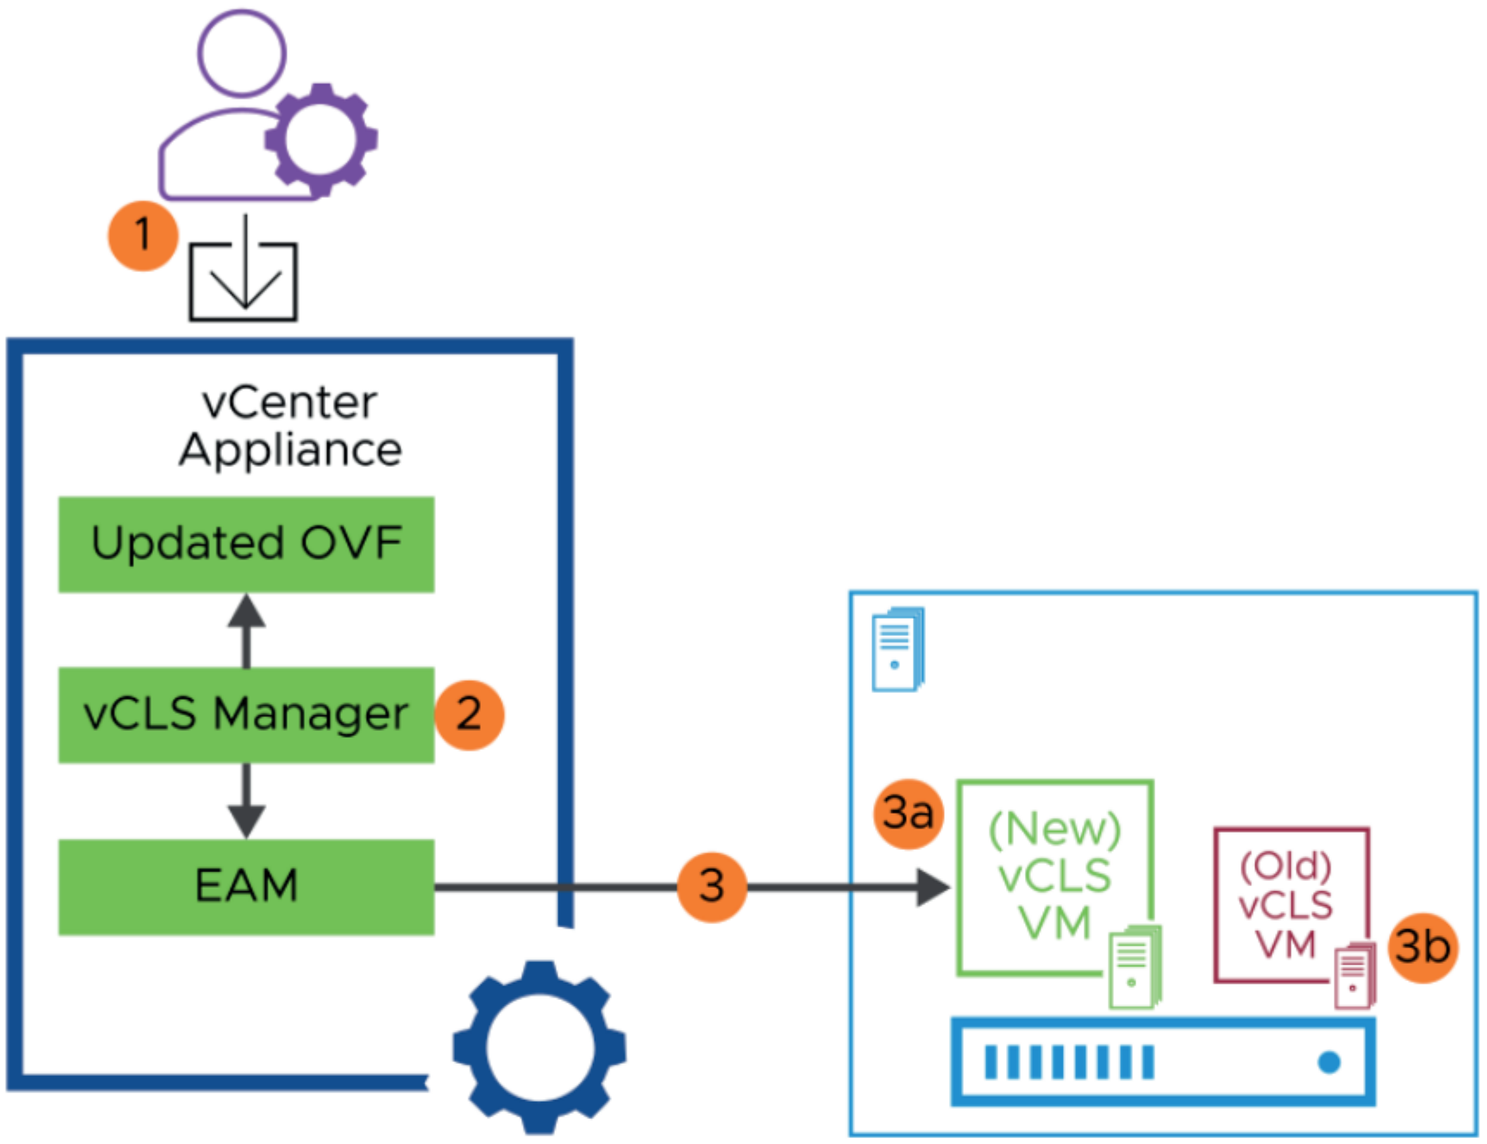

When a new update to the vCLS OVF template becomes available:

1. A vSphere administrator updates or patches vCenter. The update includes a new version of the vCLS OVF template.
2. vCLS Manager observes a new version of the vCLS OVF template and updates the EAM agency.
3. EAM initiates an upgrade of vCLS VMs:
   a.. A new vCLS is deployed and powered on.
   b.. An old vCLS is shut down and deleted.

   This process repeats until all vCLS VMs are replaced with new versions.

Technet24

3-37 Moving ESXi Hosts Between Clusters

The vCLS automatically ensures the desired state of each EAM agency in a vSphere Cluster.

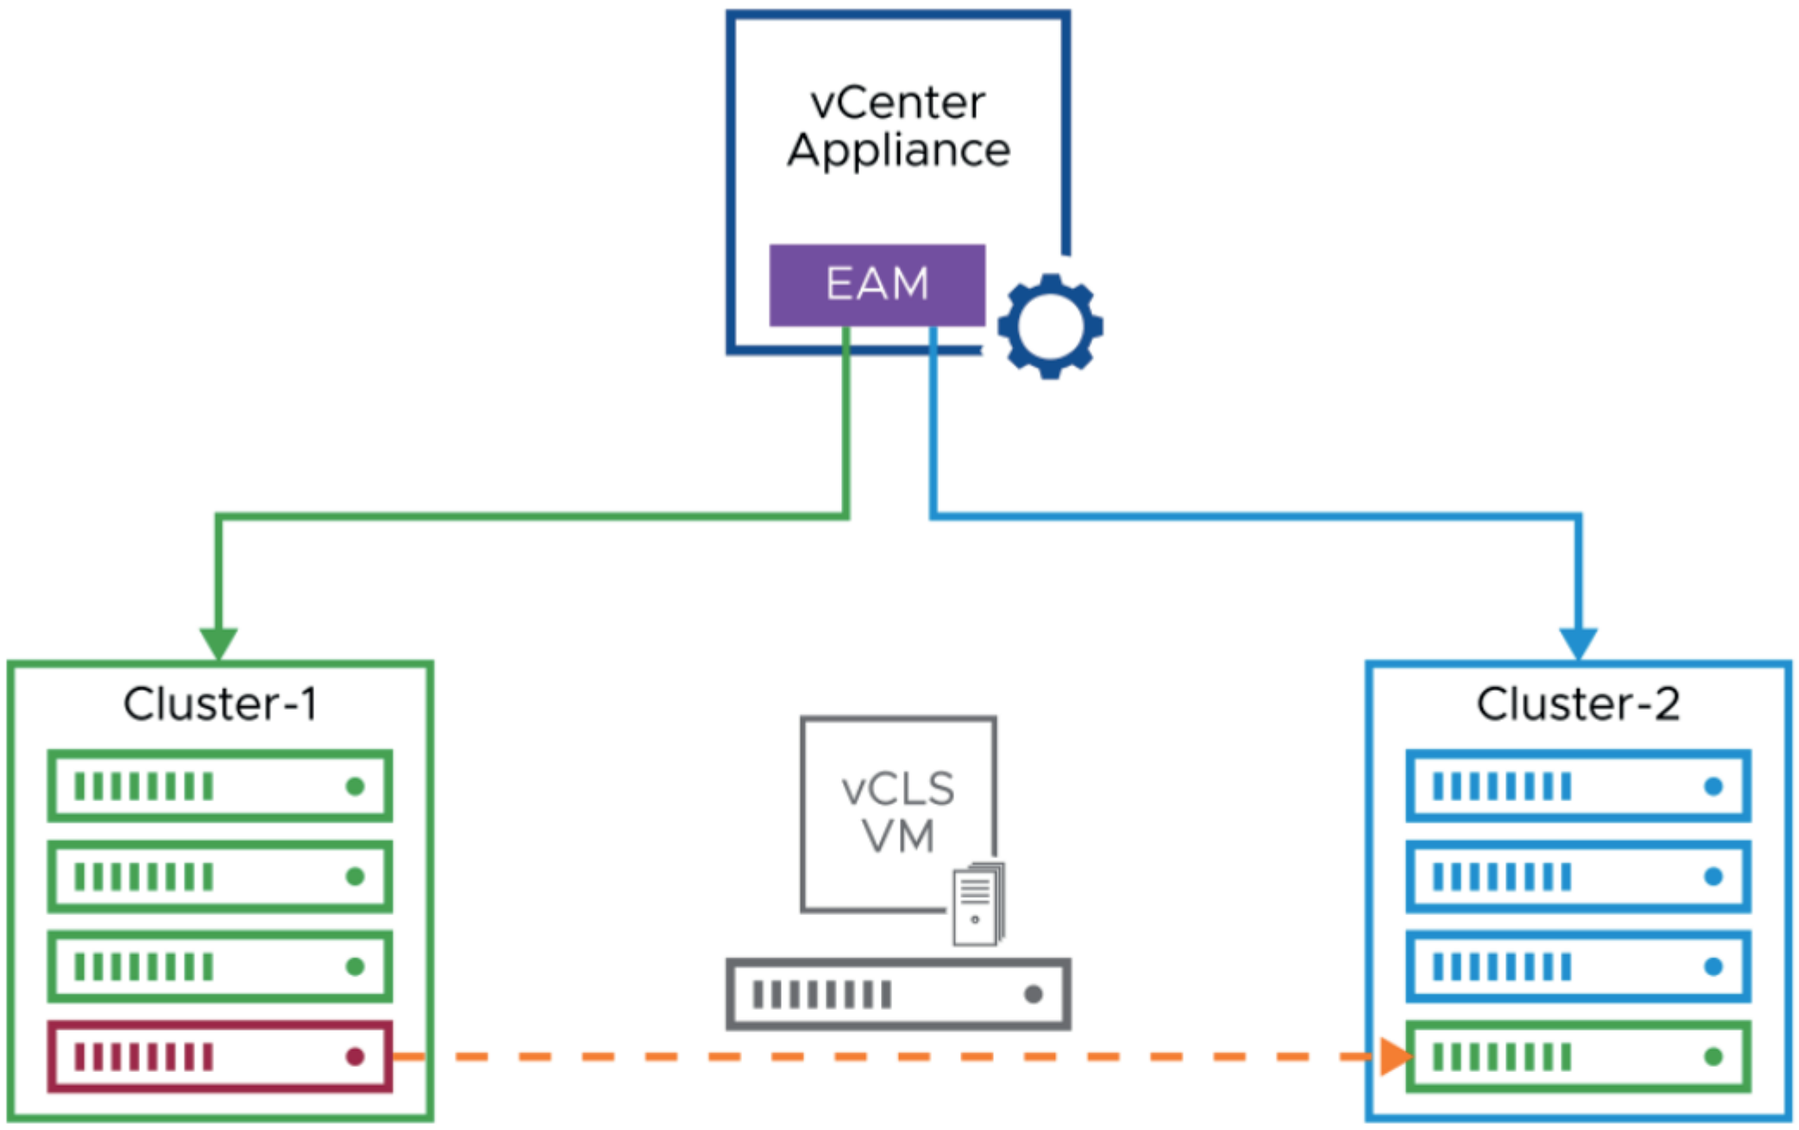

For example:

1. A vSphere administrator moves a host from one vSphere cluster to another vSphere cluster.
2. If vCLS VMs are on the host, they are deleted when the host is added to the destination vSphere cluster.
3. The EAM agency for the source and destination vSphere cluster is updated.
4. New vCLS VMs are deployed to the host when it is added to the destination vSphere cluster, if necessary to satisfy the EAM agency.

If an ESXi host, running vCLS VMs, is moved from a vCenter 7.0 Update 1 environment to a vCenter 7.0 or earlier environment, vCLS VMs must be manually powered down and deleted.

3-38 vCLS Functions

vCLS ensures cluster services such as vSphere DRS and vSphere HA are all available to maintain the resources and health of the workloads running in the clusters, independent of the vCenter instance availability.

Introduced in vSphere 7.0 Update 3, you can now configure:

- Preferred Datastores for vCLS VMs
- Anti-Affinity for vCLS VMs with specific other VMs

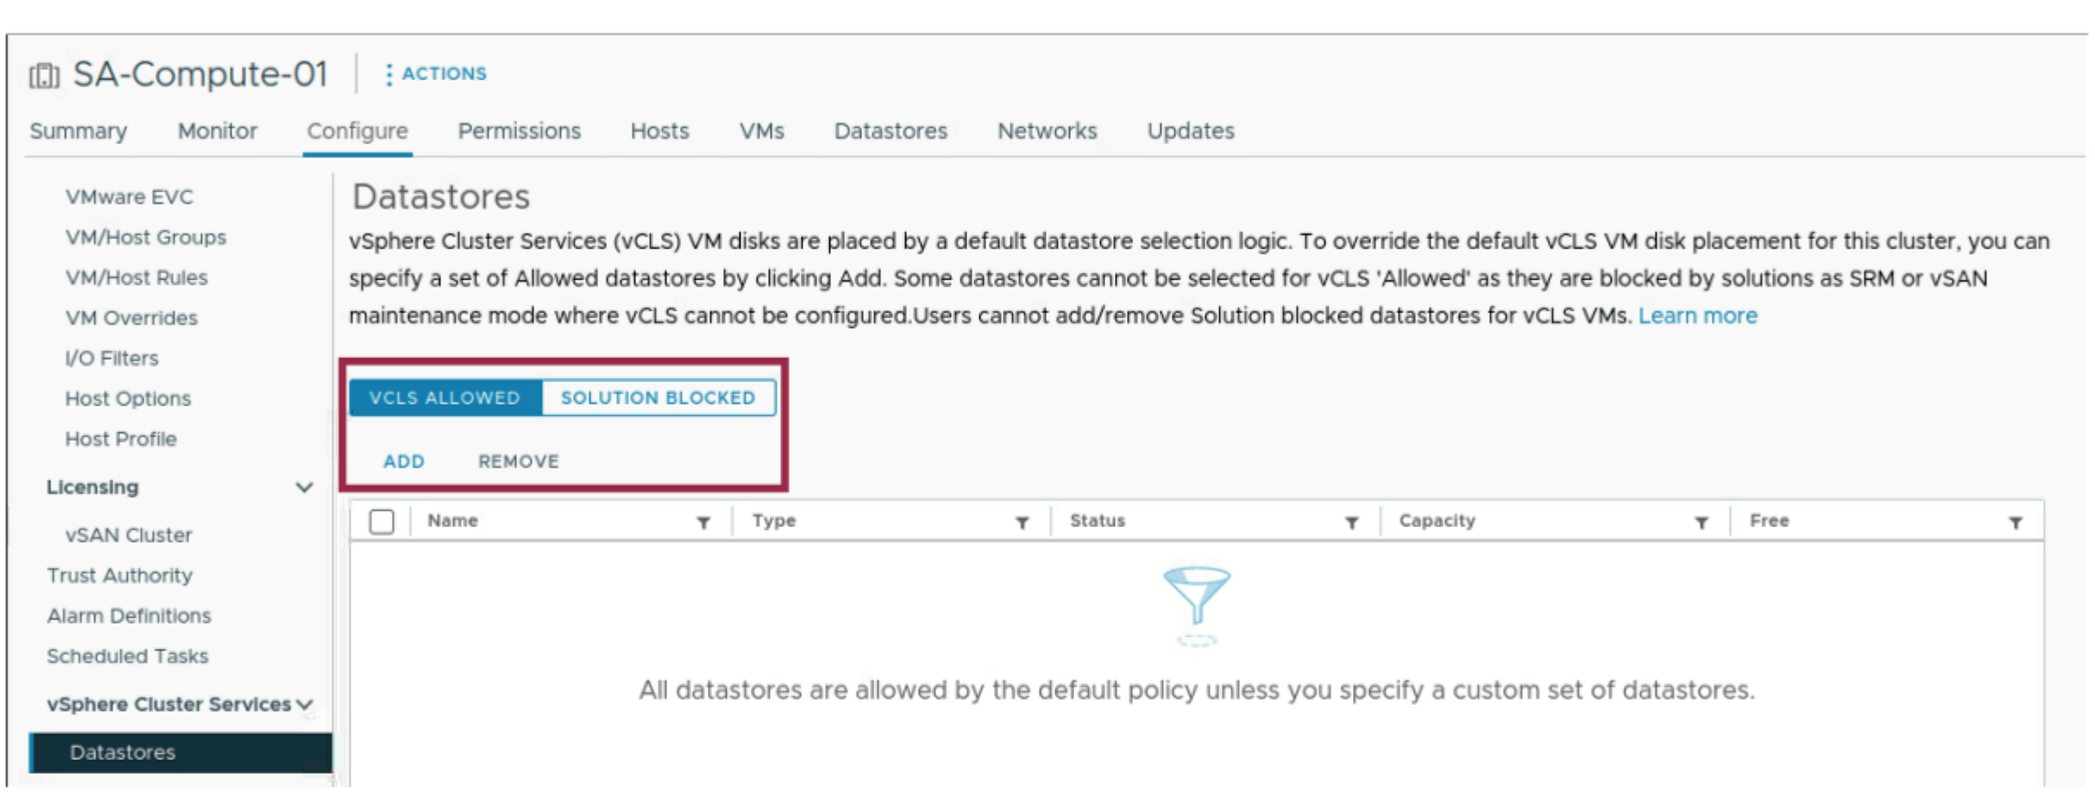

Technet24

3-39 Maintaining Health of vCLS (1)

The vCLS VMs are always powered-on because vSphere DRS depends on the availability of these VMs.

Only administrators can perform selective operations on vCLS VMs.

To avoid failure of cluster services, avoid performing any configuration or operations on the vCLS VMs.

The vCLS VMs are protected from accidental deletion

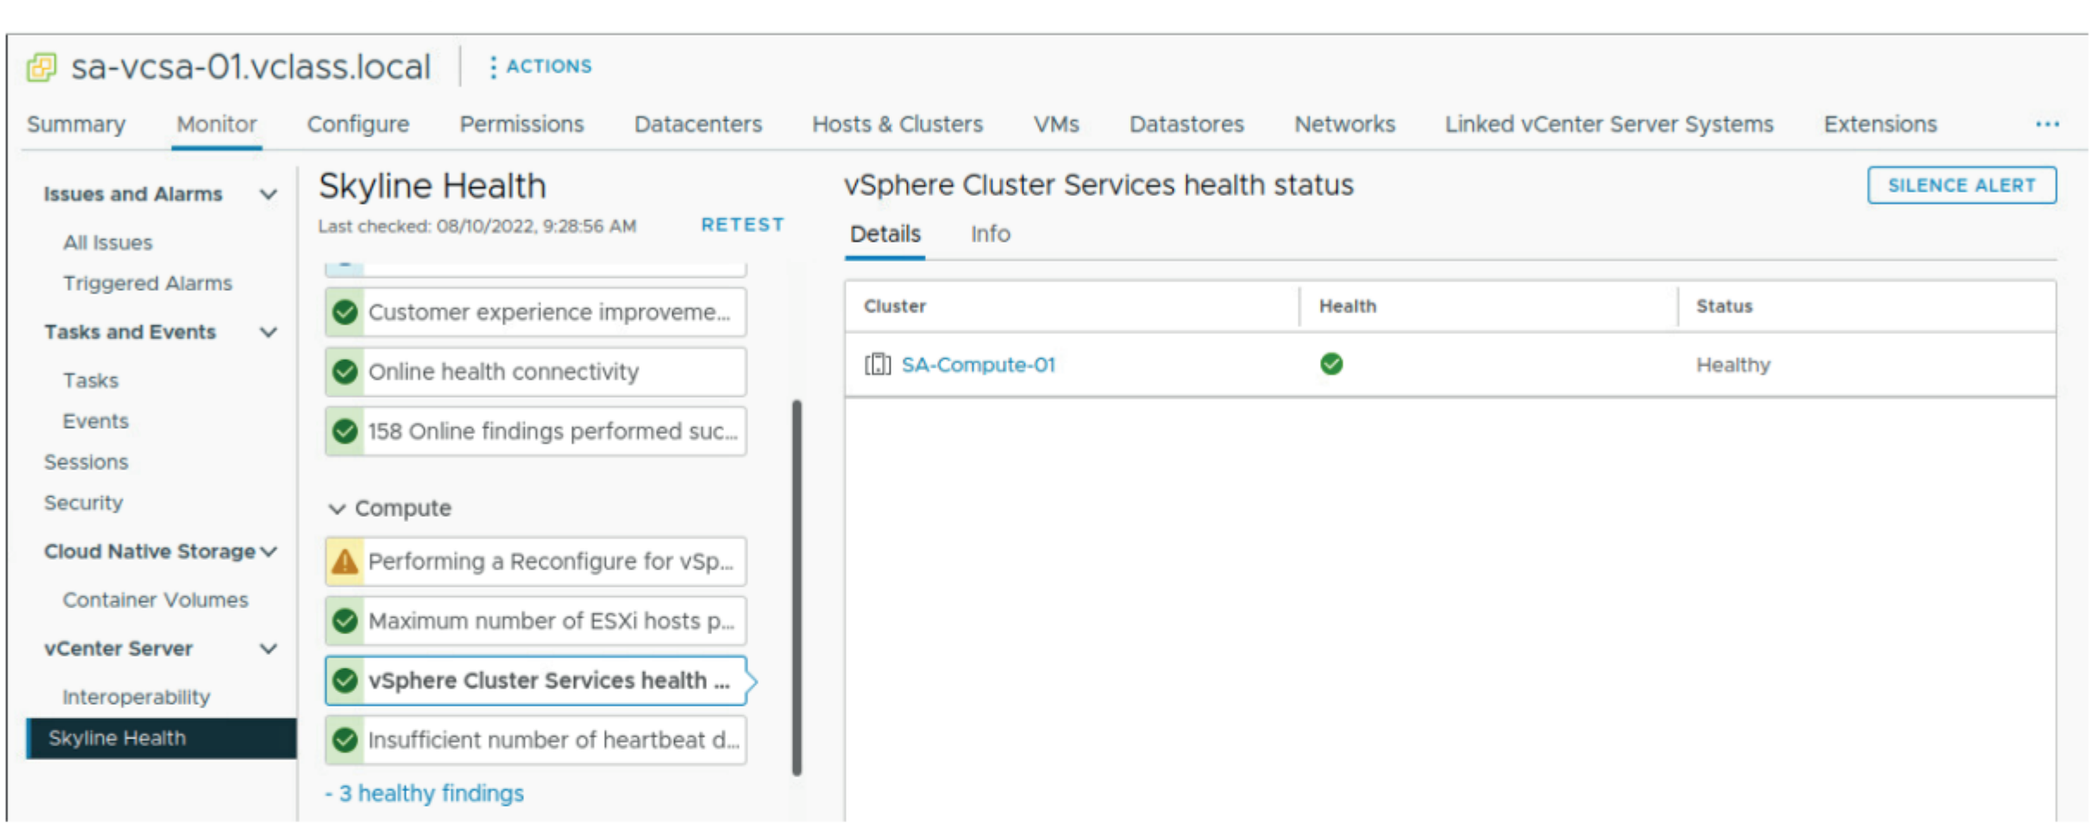

3-40 Maintaining Health of vCLS (2)

Operations that might disrupt the healthy functioning of vCLS VMs:

- Changing the power state of the vCLS VMs
- Resource reconfiguration of the vCLS VMs
- Deleting the vCLS VMs from disk
- VM encryption
- Moving vCLS VM to a different folder
- Renaming the vCLS folders
- Enabling HA admission control policy on vCLS VMs
- Moving vCLS VMs to a resource pool
- Recovering vCLS VMs from a snapshot
- Changing the BIOS
- Removing the vCLS VMs from the inventory
- Configuring PMem
- Cloning vCLS VMs
- Renaming the vCLS VMs
- Enabling DRS rules and overrides on vCLS VMs
- Enabling HA overrides on vCLS VMs

3-41 Retreat Mode

There is no way to disable vCLS on a vSphere cluster and still have vSphere DRS function on that cluster.

However, should it be necessary, you can disable vCLS on a cluster by activating Retreat Mode on the cluster.

For example, when a datastore is placed in maintenance mode, if the datastore hosts vCLS VMs, you must manually vMotion the vCLS VMs to a new location or put the cluster in retreat mode.

Impacts on your vSphere cluster:

- vSphere DRS will not function on that cluster if DRS is enabled for that cluster.
- vSphere HA will not perform optimal placement during a host failure scenario because HA depends on DRS for placement recommendations.

Technet24

3-42 Activating Retreat Mode

To put a cluster in Retreat Mode

1 Navigate to the cluster on which vCLS must be deactivated.

2 Copy the cluster domain ID from the URL of the browser. It should be similar to `domain-c<number>`.

3 Navigate to the vCenter **Configure** tab.

4 Under **Advanced Settings**, click the **Edit Settings** button.

5 Add a new entry `config.vcls.clusters.domain-c<number>.enabled`. Use the domain ID copied in step 2.

6 Set the Value to `False`.

7 Click Save.

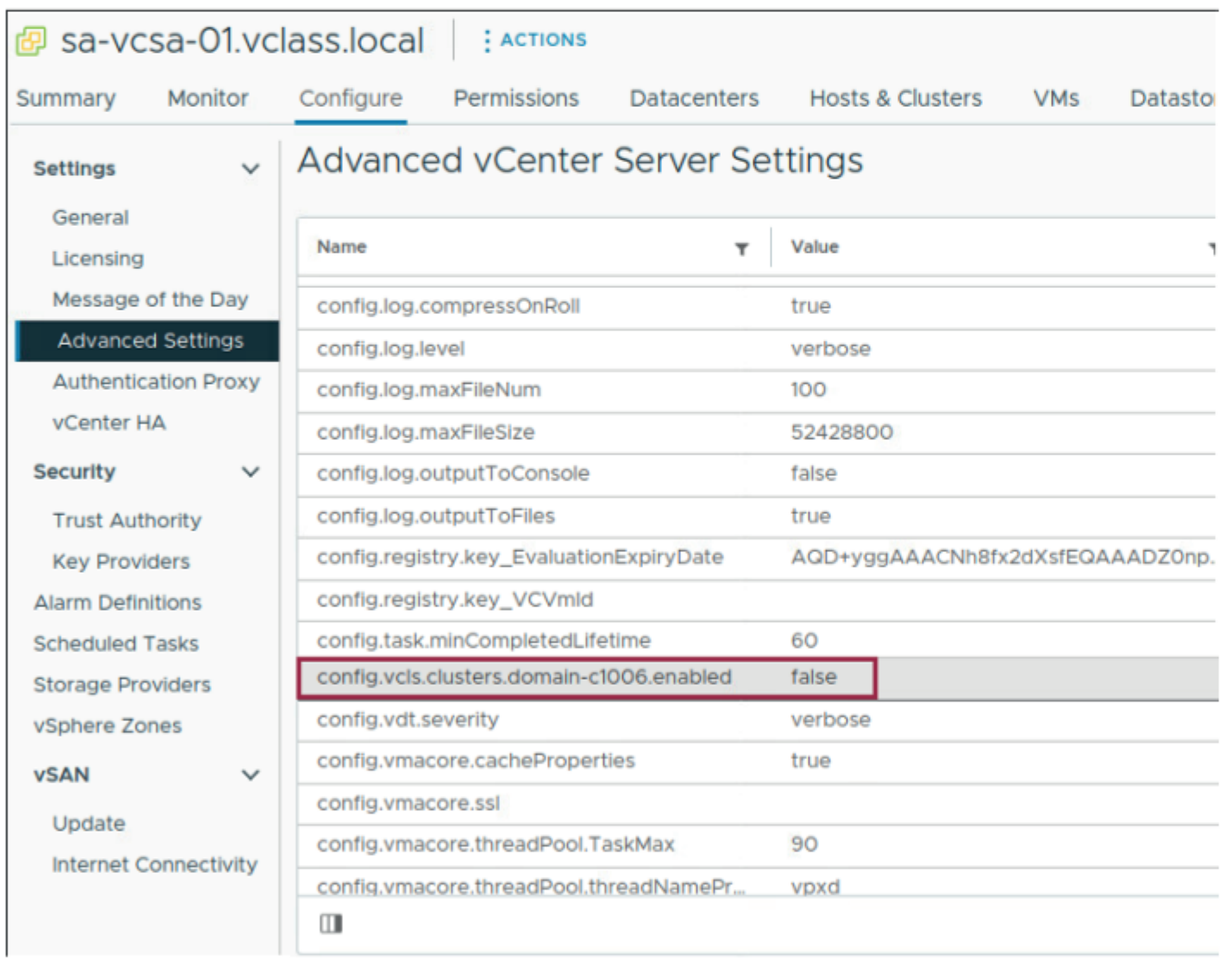

3-43 Troubleshooting Log Files

Log files related to vCLS tasks can be found in different locations on the vCenter.

ESX Agent Manager	Location
`eam.log`	`/var/log/vmware/eam/`

vCLS Manager	Location
`wcpsvc.log`	`/var/log/vmware/wcp/`

The root password for vSphere Cluster Service VMs can be extracted by running the `/usr/lib/vmware-wcp/decrypt_clustervm_pw.py` script from a root SSH session on vCenter. The VM console interface is used to access vSphere Cluster Service VMs.

To ensure cluster services health, avoid accessing the vCLS VMs. This document is intended for explicit diagnostics on vCLS VMs.

3-44 Lab 6: Enabling vCLS Retreat Mode

Enable vCLS retreat mode on a vSphere luster:

1. Verify the Cluster Domain ID
2. Enable vCLS Retreat Mode and Monitor vCLS VMs
3. Revert the Changes

Technet24

3-45 Review of Learner Objectives

- Describe vCLS VMs
- Identify vCLS workflows
- Describe the function of the vCLS
- Recognize operations that might disrupt the healthy functioning of vCLS VMs

3-46 Key Points

- A resource pool is a logical abstraction of hierarchically managed CPU and memory resources.
- vSphere Cluster Service (vCLS) VMs are deployed to a cluster at creation and after hosts are added to the cluster.

Questions?

Module 4

Network Operations

4-2 Importance

As you scale your vSphere environment, the features and functions of vSphere Distributed Switch can help you manage networking.

4-3 Module Lessons

1. Managing vSphere Distributed Switches
2. Using Network I/O Control in Distributed Switches
3. NetFlow and Port Mirroring in Distributed Switches
4. vSphere Distributed Services Engine

Technet24

4-4 Lesson 1: Managing vSphere Distributed Switches

4-5 Learner Objectives

- Manage network traffic with the traffic filtering and marking policy
- Recognize NSX virtual distributed switches in the vSphere Client
- Perform a distributed switch health check
- Back up and restore a distributed switch configuration
- Describe the automatic rollback and recovery of networking configurations

4-6 About Distributed Switches

A distributed switch functions as a single virtual switch across all associated hosts. Distributed switches have several benefits over standard switches:

- Distributed switches centralize the virtual network administration, and simplifies the data center administration.
- Distributed switch ports are statically assigned by vCenter and offer more granular control over network statistics and policies.

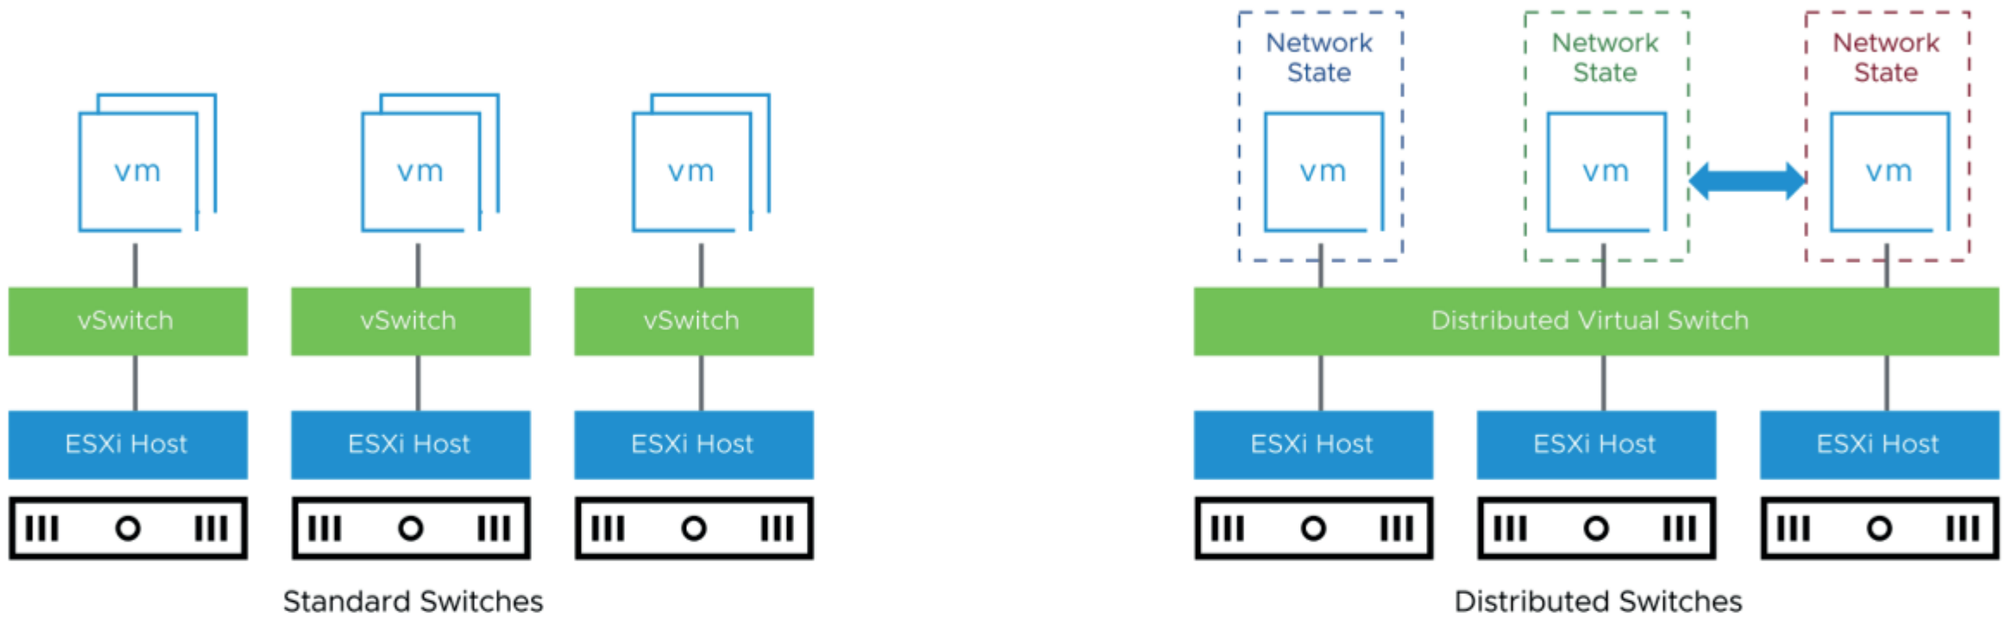

Having the network configuration at the data center level (vSphere Distributed Switch), not at the host level (standard switch), gives distributed switches the following advantages:

- Data center setup and administration are simplified through this centralized network configuration. For example, adding a host to a cluster and making it compatible with vSphere vMotion is much easier than with a standard switch.
- Distributed ports migrate when VMs are migrated. For example, when you migrate a VM with vSphere vMotion, the distributed port statistics and policies move with the VM, which simplifies debugging and troubleshooting.

Technet24

4-7 About the Traffic Filtering and Marking Policy

With the traffic filtering and marking policy for vSphere distributed switches, you can protect your virtual network from unwanted traffic and security attacks and manage network traffic to meet specified service level agreement (SLA) targets.

The traffic filtering and marking policy has the following features:

- It can permit or deny specific types of traffic.
- It can apply a QoS tag to mark a certain type of traffic.
- It is equivalent to the access control list feature available on physical switches.

The traffic filtering and marking policy consists of one or more network traffic rules, defined at the distributed port group or uplink port group level.

Use the traffic filtering and marking policy to create a set of rules for security and QoS tagging of packets flowing through distributed switch ports.

The distributed switch applies rules on traffic at different places in the data stream. It can apply network traffic rules on the data path between the VM network adapter and the distributed port. It can also apply network traffic rules between the uplink port and the physical network adapter.

The traffic filtering and marking policy is supported only on distributed switches.

4-8 Creating Network Traffic Rules

You can define network traffic rules for processing traffic related to VMs or to physical adapters.

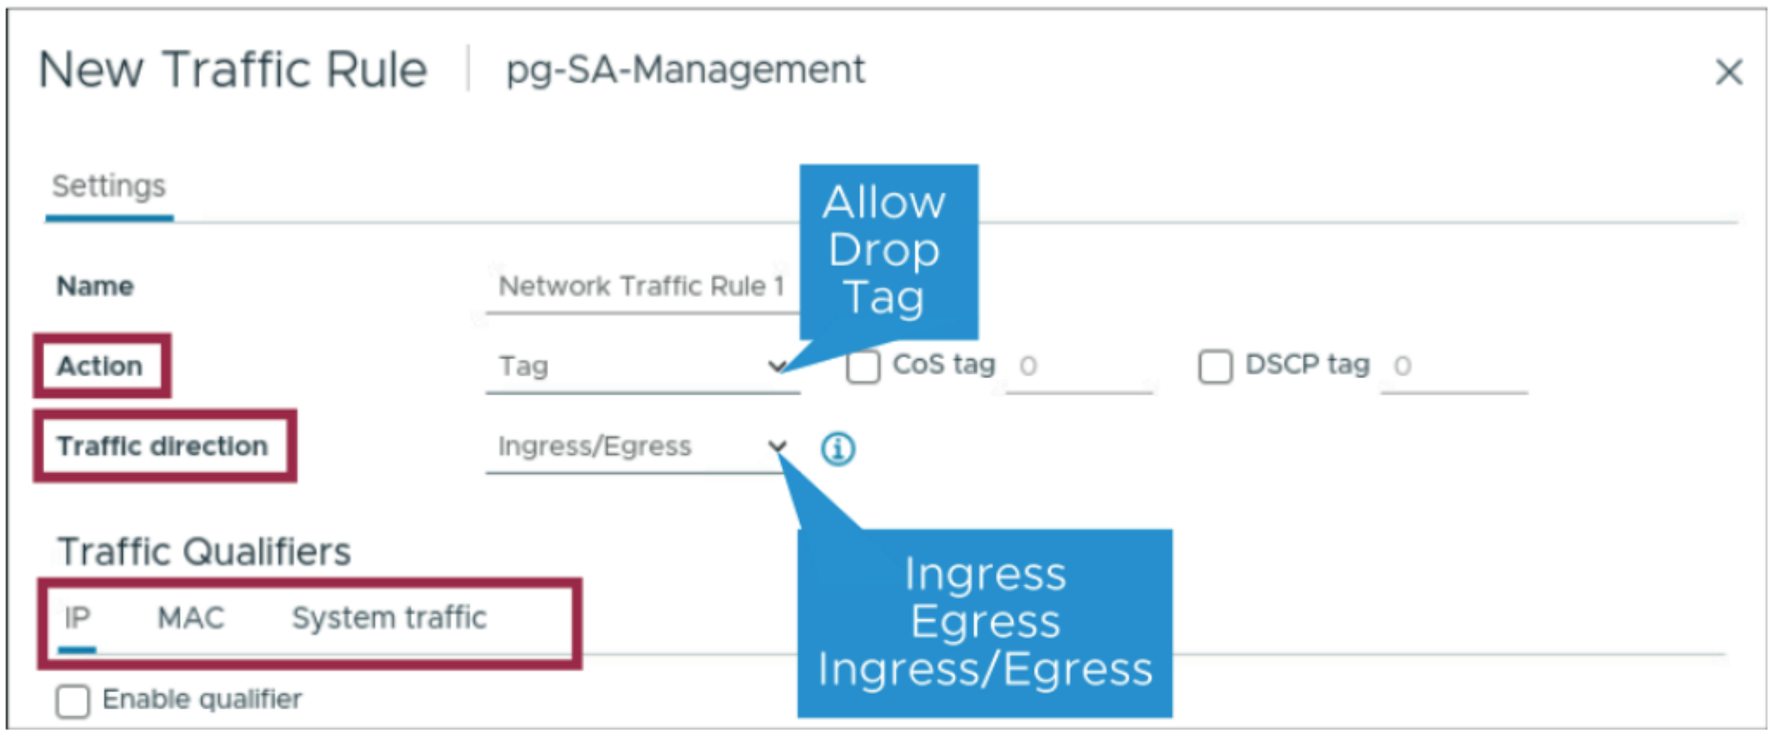

In general, a network traffic rule consists of a qualifier for traffic and an action for restricting or prioritizing the matching traffic.

In a network traffic rule, the Allow action allows traffic to pass through the distributed switch port or uplink port, and the Drop action blocks the traffic. The Tag action marks (or tags) traffic passing through the distributed switch port or uplink port.

From the point of view of the Distributed Switch

- Ingress means in-bound from the VM, vmkernel port or from and uplink.
- Egress is out-bound to the VM, to vmkernel port, to an uplink.

A qualifier represents a set of matching criteria related to a networking layer. You can match traffic based on the system traffic type, layer 2 traffic properties, or layer 3 traffic properties. You can use the qualifier for a specific networking layer, or you can combine qualifiers to match packets more precisely.

Technet24

4-9 Network Traffic Rule: Example 1

This rule allows incoming traffic from systems on VLAN 32 that have a MAC address starting with 00:50:56 and denies all other traffic.

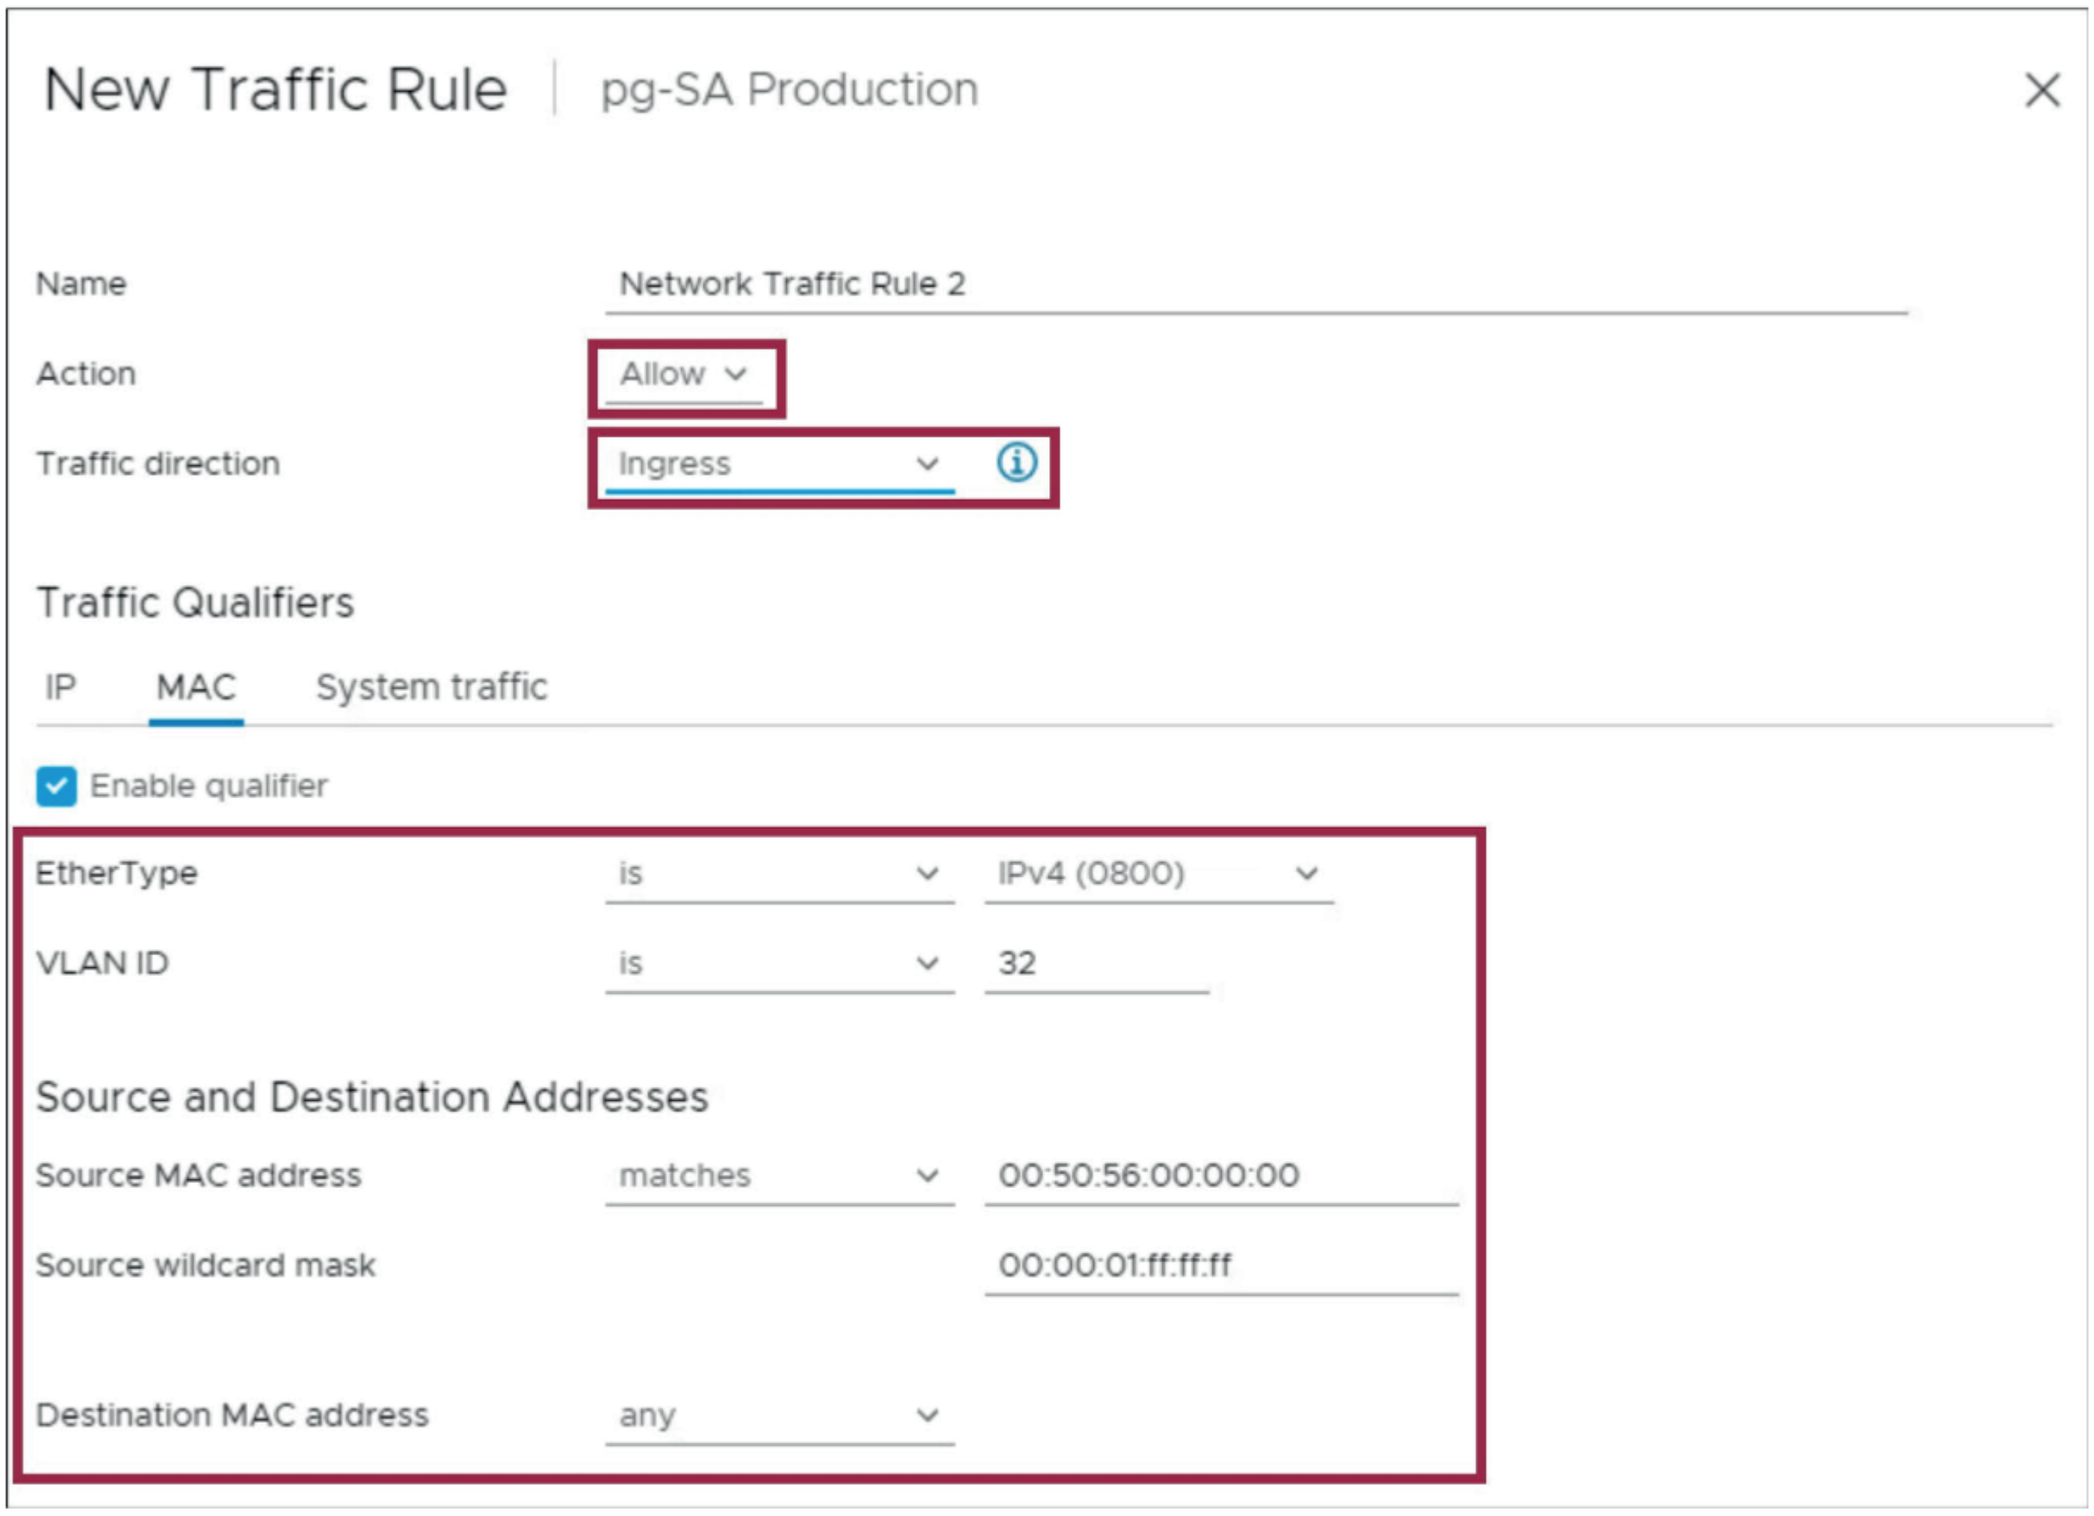

Use a MAC traffic qualifier to define matching criteria for the layer 2 (data link layer) properties of packets. These properties include MAC address, VLAN ID, and a protocol that consumes the frame payload (IPv4, IPv6, or Address Resolution Protocol).

Locating traffic with a VLAN ID on a distributed port group works with Virtual Guest Tagging. To match traffic to VLAN ID if Virtual Switch Tagging is active, use a rule on an uplink port group or uplink port.

4-10 Network Traffic Rule: Example 2

This rule allows incoming and outgoing VM traffic, known as Ingress/Egress

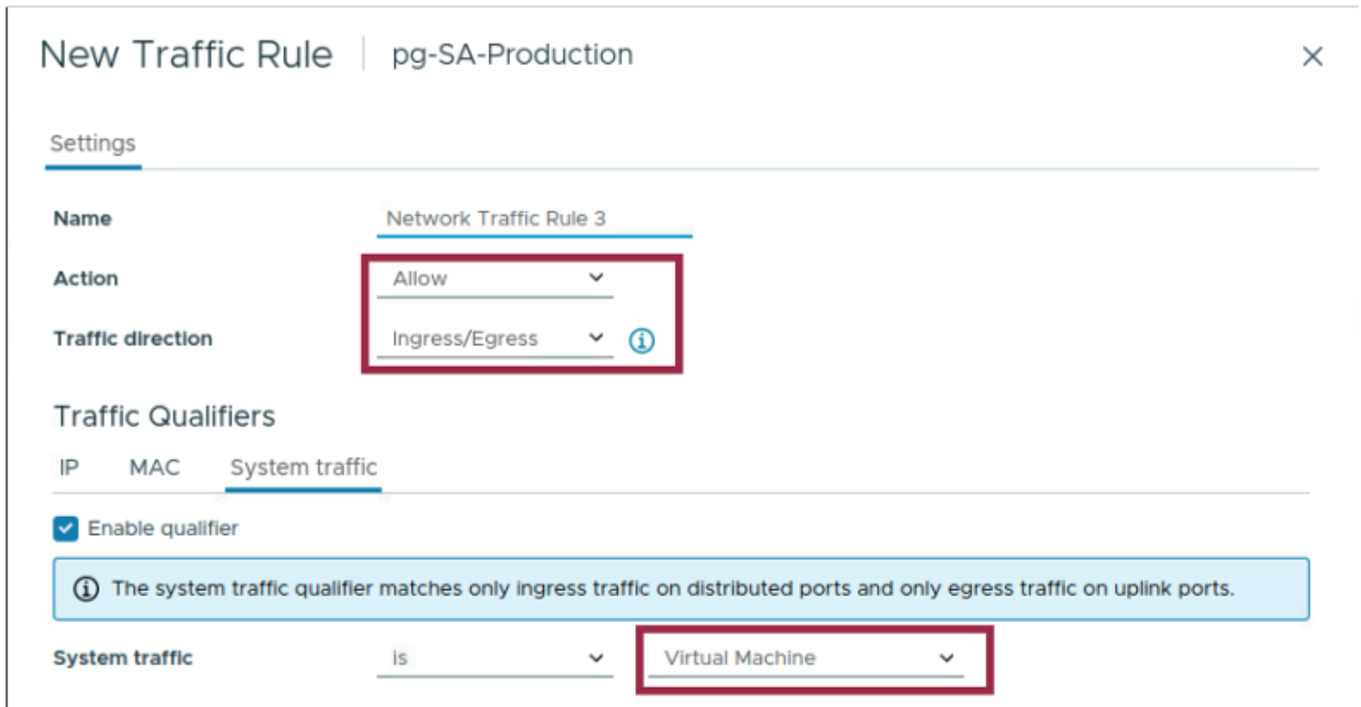

Use the system traffic qualifier to match packets to the type of virtual infrastructure data that is flowing through the ports of the distributed port group.

In this example, the rule enforces that the traffic through the pg-SA Production port group must be VM traffic. The distributed switch drops any other types of traffic that are sent to this port group.

You can select the type of traffic through the ports of the group that carries system data, such as traffic for management from vCenter, iSCSI storage, and vSphere vMotion.

When you use qualifiers for the system data type, both layer 2 and layer 3 packet attributes set the properties that packets must have to match the rule.

Ingress - Traffic moving out of a physical interface.

Egress - Traffic moving into a physical interface.

Technet24

4-11 Marking Network Traffic (1)

You can assign priority tags to traffic that has higher networking requirements for bandwidth, low latency, and so on.

You can mark the traffic with a Class of Service (CoS) tag in layer 2 or a Differentiated Serviced Code Point (DSCP) tag in layer 3.

Marking traffic has the following benefits:

- You can tag important traffic so that the physical network does not drop these packets during congestion.
- You can provide end-to-end Quality of Service (QoS) and service level agreement (SLA) requirements.

Marking, or priority tagging, is a way to mark traffic that has higher Quality of Service (QoS) demands. With priority tagging, the network can recognize different classes of traffic. Network devices can handle the traffic from each class according to its priority and requirements.

Similar to traffic filtering, marking also requires that you classify the traffic first, based on these qualifiers: system, MAC, and IP. After you define your traffic qualifiers, you can decide how to mark your traffic.

You can prioritize traffic by using a Class of Service (CoS) priority tag, a Differentiated Serviced Code Point (DSCP) tag, or both.

You can mark traffic with a CoS priority tag in the layer 2 packet header. Accepted values are 0 through 7.

You can mark traffic with a DSCP tag in the layer 3 packet header. Accepted values are 0 through 63.

You can assign priority tags to traffic, such as VoIP and streaming video, that has higher networking requirements for bandwidth, low latency, and so on.

4-12 Marking Network Traffic (2)

This rule marks incoming SIP UDP packets from subnet 192.168.2.0/24.

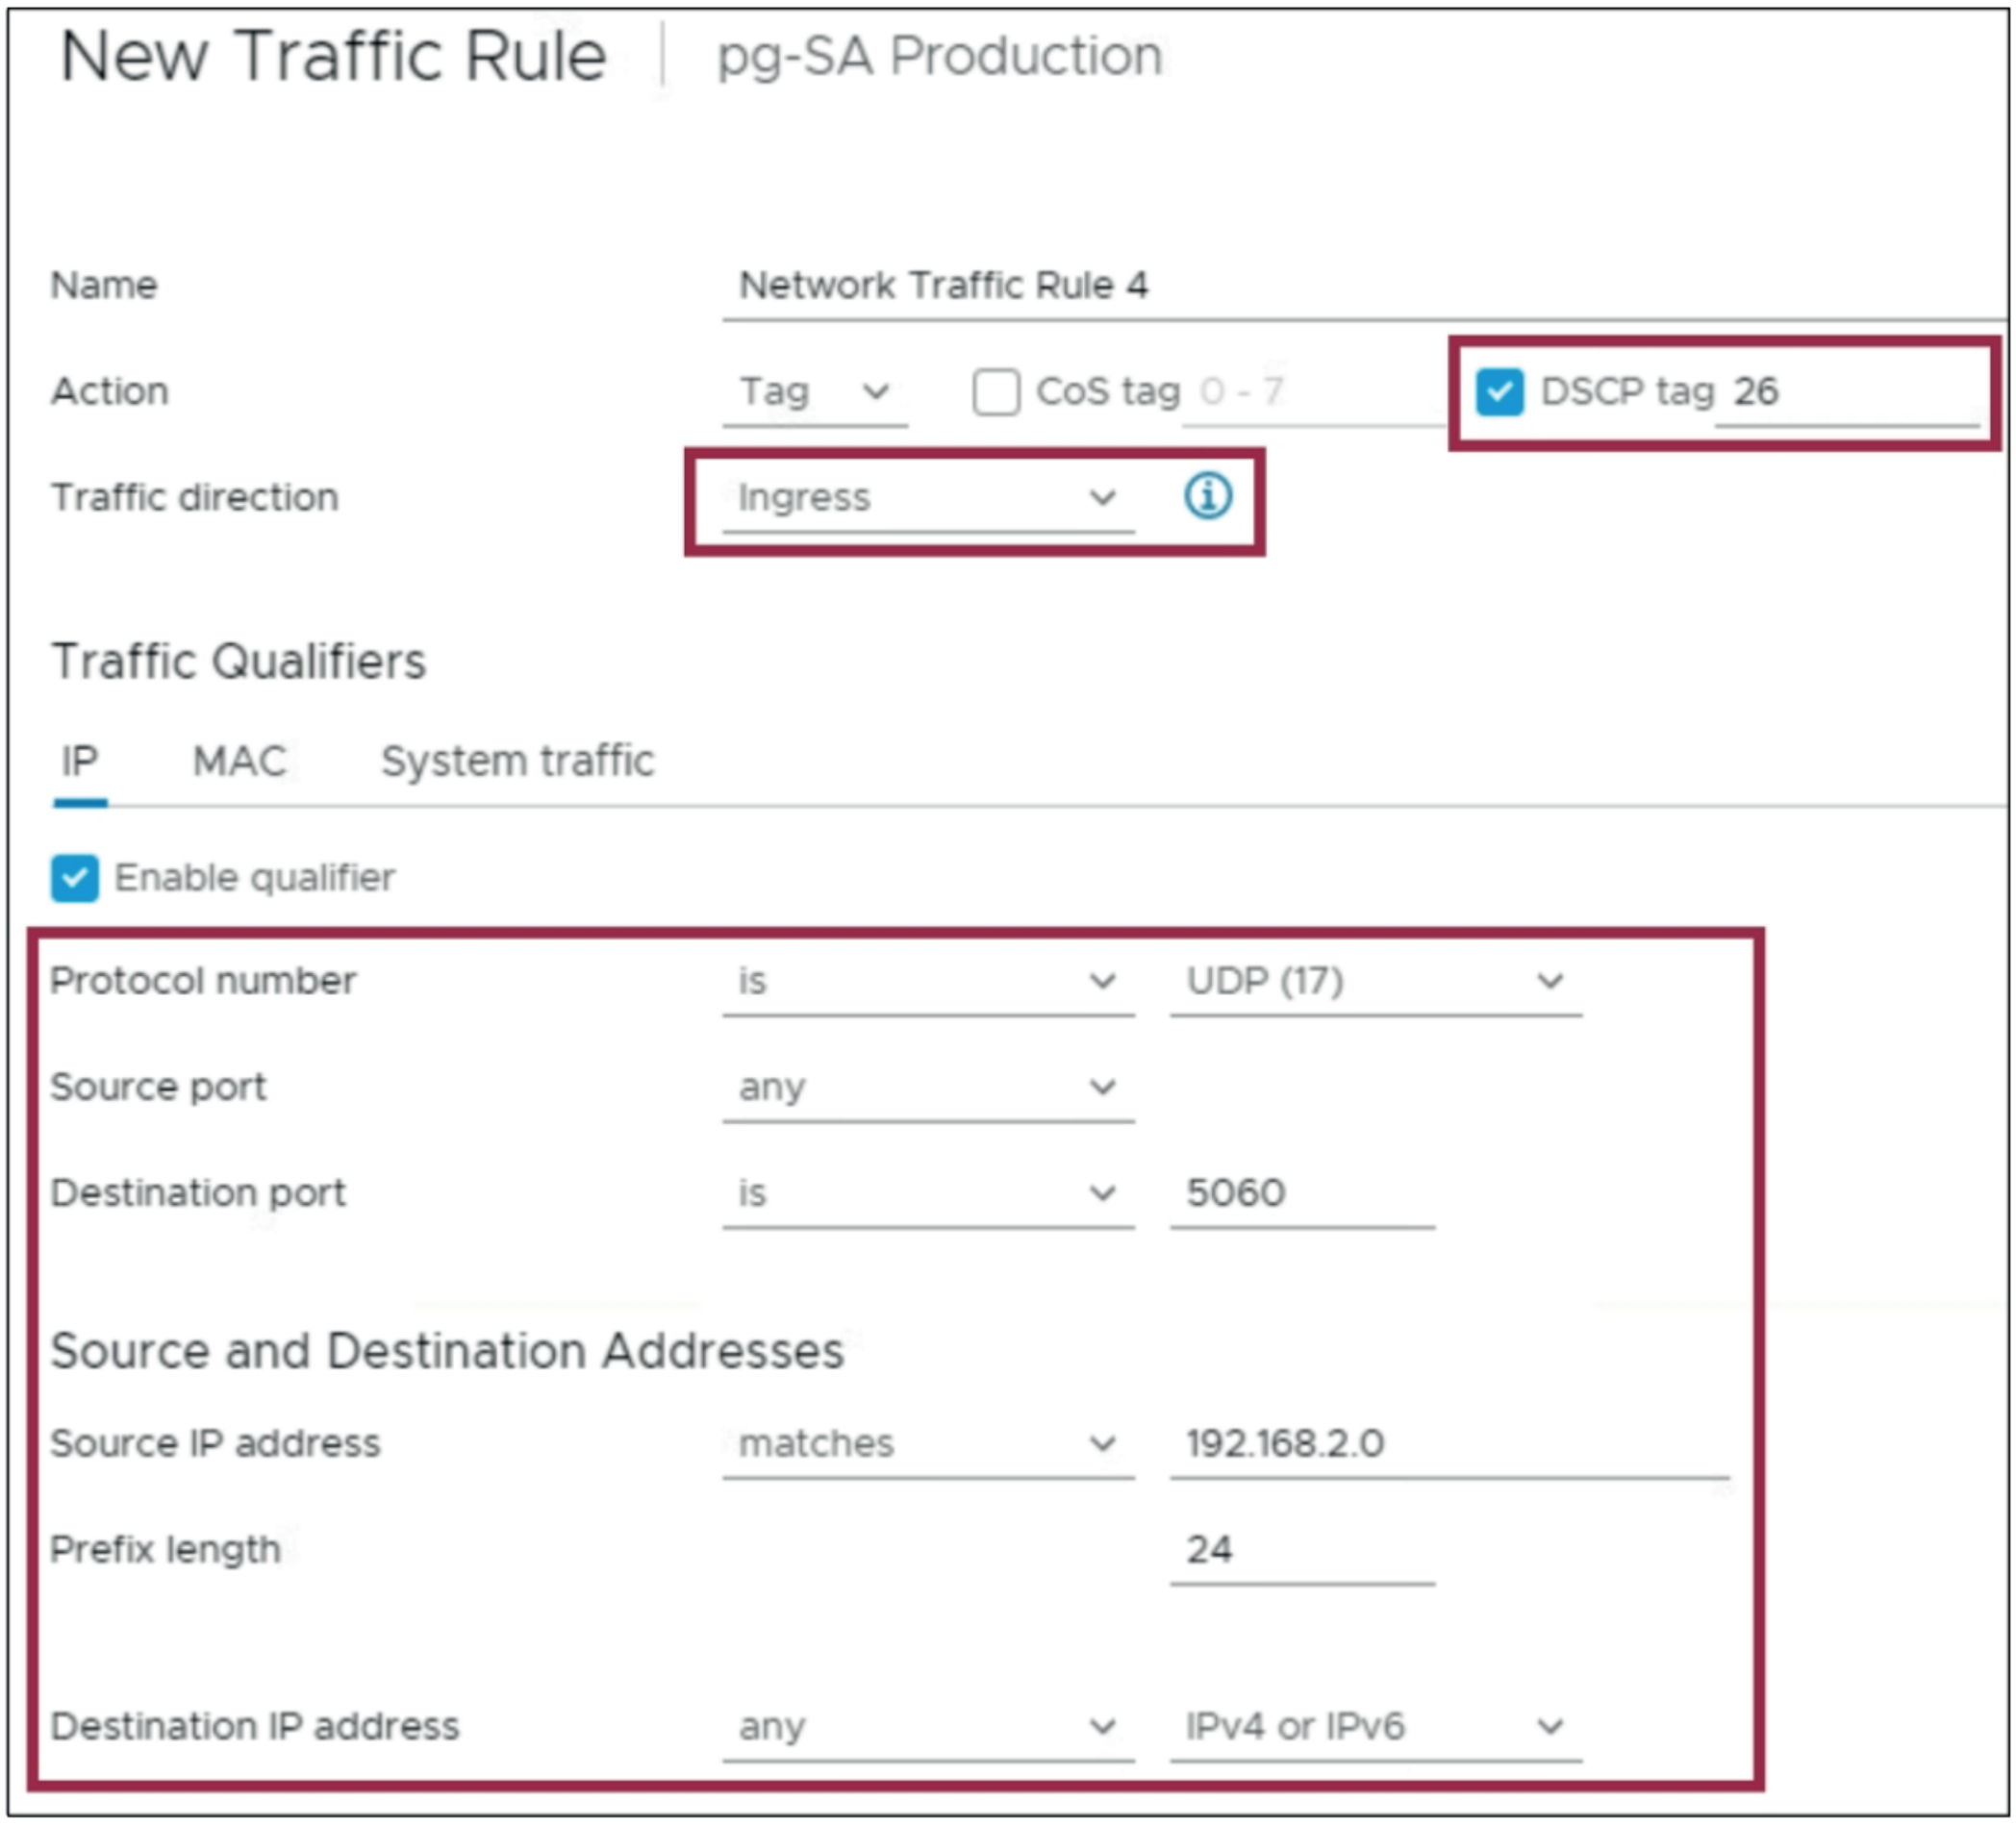

This example marks VoIP traffic. VoIP flows have special requirements for QoS in terms of low loss and delay. The traffic related to Session Initiation Protocol (SIP) for VoIP usually has a DSCP tag equal to 26.

Technet24

4-13 Access to NSX Virtual Distributed Switch (1)

vSphere 7 and later supports both VDS and the NSX virtual distributed switch (N-VDS).

- NSX logical switches can be presented in vCenter as a form of virtual port groups.
- vSphere 7 and VMware NSX 3.0 must be configured with the following minimum versions:
  - vCenter 7.0
  - ESXi 7.0
- vSphere 8 and VMware NSX 4.0 must be configured with the following minimum versions:
  - vCenter 8.0
  - ESXi 7.0
- vCenter must be registered to the NSX Manager as a compute manager.
- For logical switches to be visible in the vSphere Client, you must configure VDS with the NSX port group from the transport zone.

VMware now supports DPU-based Acceleration for NSX in this VMware NSX 4.0.1.1.

4-14 Access to NSX Virtual Distributed Switch (2)

NSX Manager provides full access for managing NSX-T Data Center objects:

- Manage the configuration of VDS
- Manage the configuration of logical switches (NSX port groups)
- Manage the configuration of NSX ports

vCenter provides limited access to NSX-T Data Center objects:

- Manage the configuration of VDS
- Limited access to NSX port groups and NSX ports:
  - You can view some NSX configuration data.
  - You can assign VM network adapters or vNICs to NSX port groups.
  - You can use vSphere vMotion to migrate a VM to and from NSX port groups.
  - You can add VMkernel adapters to NSX port groups.

Starting with vSphere 7, NSX port groups behave as regular distributed port groups so VMkernel adapters can be added to NSX port groups from the vSphere Client.

The vSphere Client does not support the following NSX port group actions:

- Delete
- Modify settings
- Use in port mirroring session
- Include configuration in VDS exports

Technet24

4-15 About Health Check for Distributed Switches

With the health check support feature, you can identify and troubleshoot configuration errors or inconsistencies between the virtual and physical switch configurations in a vSphere distributed switch.

By regularly examining settings on distributed and physical switches, health check identifies common configuration errors:

- Mismatched VLAN trunks between the distributed switch and physical switch
- Mismatched MTU settings between the distributed switch, physical adapter, and physical switch ports
- Mismatched virtual switch teaming policies for the physical switch port-channel settings

Health check is a feature that detects certain inconsistencies between the physical and the virtual networks. Key parameters, such as VLAN tags, MTUs, and the NIC teaming configuration, must be configured consistently on the physical and the virtual switches. An inconsistent configuration can lead to network connectivity problems.

Health check searches for configuration inconsistencies and reports them to the administrator. The default check interval is one minute.

For example, a small change to a portgroup's settings on a VDS could affect VMs and vmkernel ports on up to 2000(?) hosts

As a best practice, enable health check for brief time periods because this feature creates many MAC addresses and can affect production traffic.

4-16 Health Check Example

Two ESXi hosts use one distributed switch with two distributed port groups each connected to two physical switches.

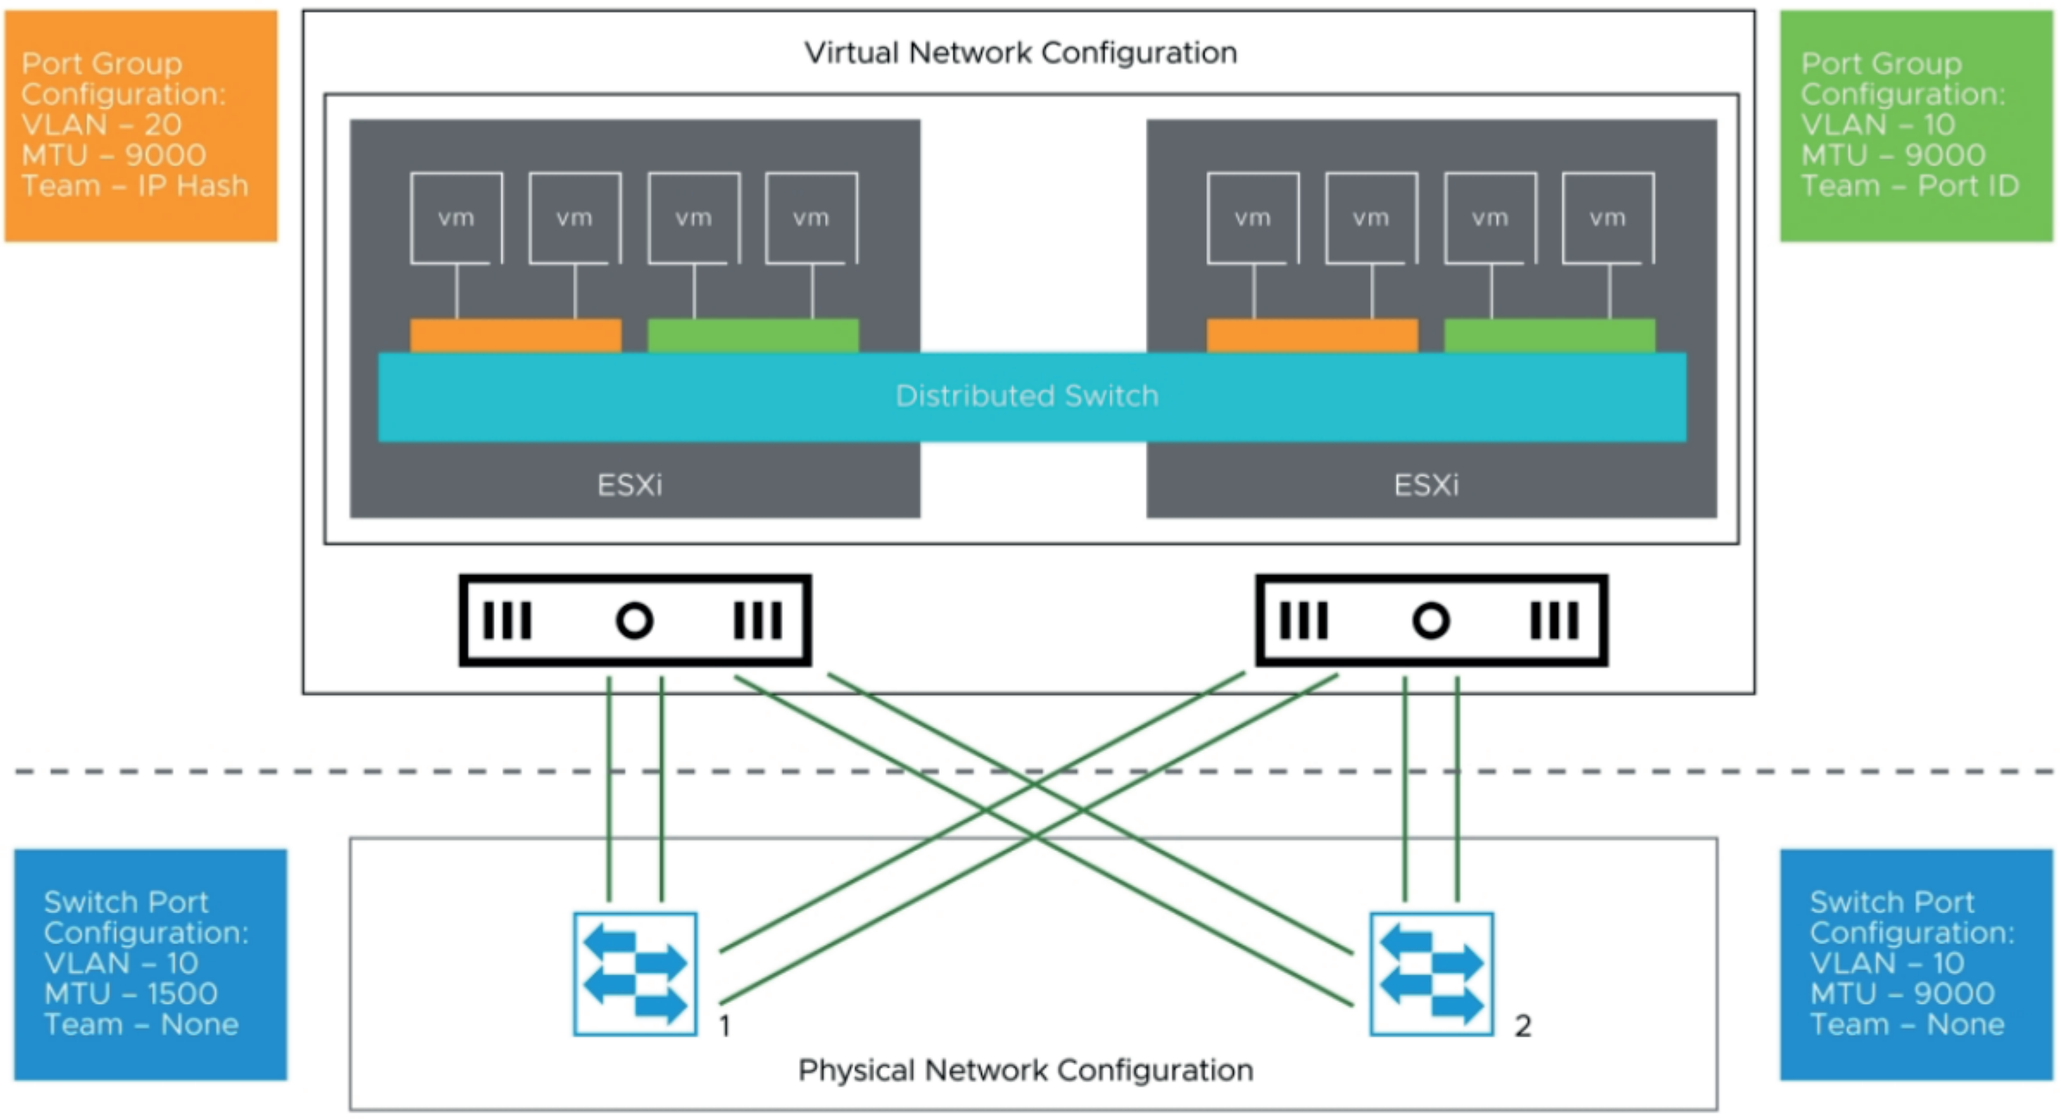

In addition to the distributed switch, two physical switches with their switch port configurations are shown.

Compare Green virtual port group (VLAN 10) with the settings on physical switch 2. The VLAN ID and the MTU settings are identical. The virtual port group is configured with the port ID teaming selection. But because the configuration is internal to the virtual switch and requires no configuration on the physical switch, the virtual and the physical settings are consistent and should not experience connectivity issues.

Compare Orange virtual port group (VLAN 20) to the settings on physical switch 1. The VLAN IDs and the MTU settings are different. The virtual port group teaming setting is set to IP hash. For IP hash to operate properly, EtherChannel or Link Aggregation Control Protocol (LACP) must be configured on the physical switch.

Health check can detect and report the configuration differences between the port group and the switch port configuration by using layer 2 Ethernet packets. At one-minute intervals (by default), request and acknowledge packets are sent back and forth between the virtual interface and the switch. When packets are dropped, a configuration warning appears in the vSphere Client.

Technet24

4-17 Enabling Health Check

Health check can be performed on the following network components:

- VLAN and MTU
- Teaming and failover

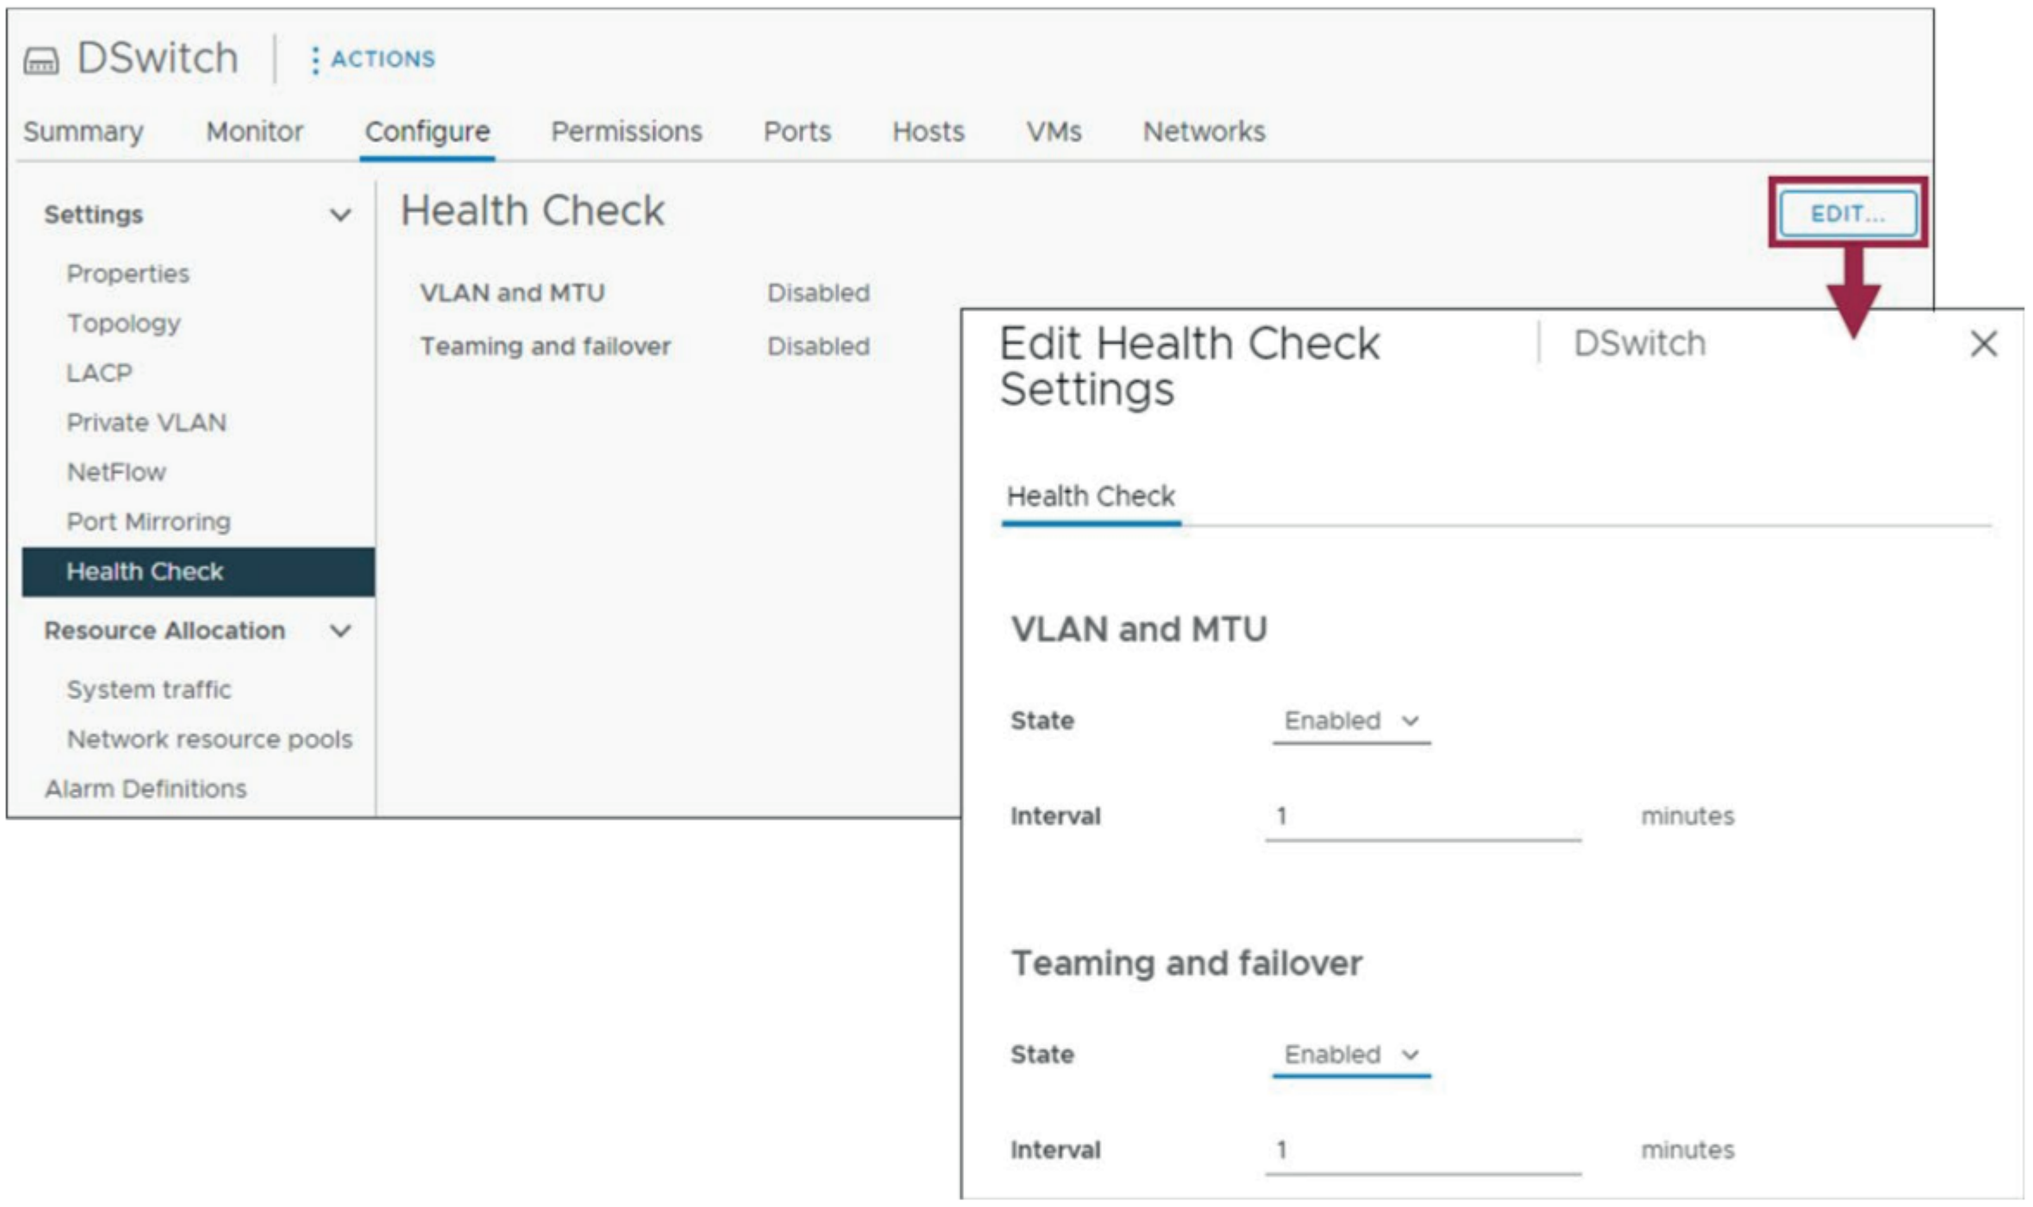

Health check has the following requirements:

- For VLAN and MTU checks, at least two physical uplinks must be connected to the distributed virtual switch.
- For the teaming policy check, at least two active uplinks must be in the team, and at least two ESXi hosts must be connected to the virtual distributed switch.

See *vSphere Networking* at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-networking-guide.pdf.

4-18 Monitoring Health Check Results

After health check runs for a few minutes, you can monitor the results in the **Distributed switch** > **Monitor** > **Health** pane in the vSphere Client.

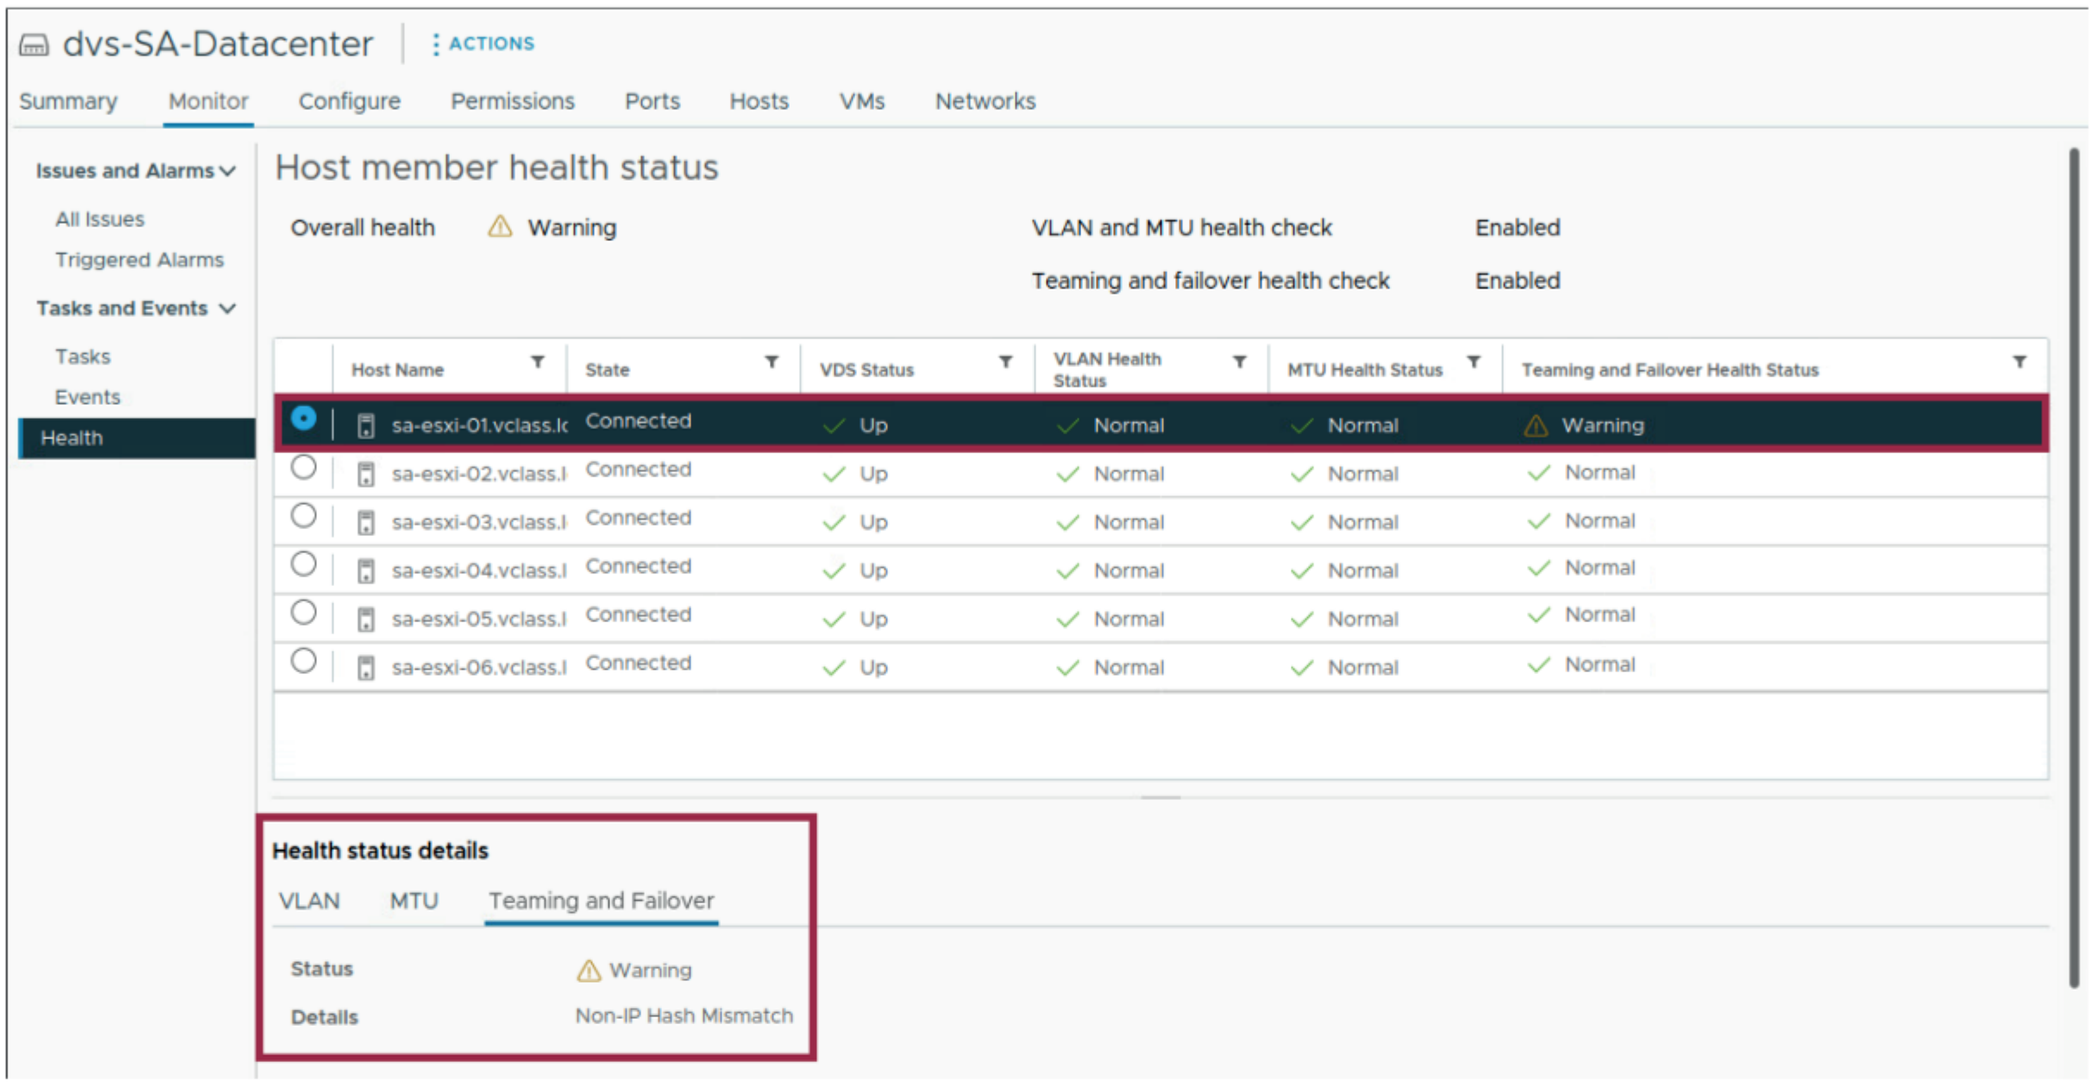

The health status details are organized into three categories: VLAN, MTU, and Teaming and Failover.

Technet24

4-19 Backing Up and Restoring Configurations

You can back up and restore the configuration of your distributed switch, distributed port groups, and uplink port group for deployment, rollback, and sharing purposes.

The following operations are supported:

- Back up the configuration to a file.
- Restore the switch and port group configuration from a backup.
- Create a switch or port group from the backup.
- Revert to a previous port group configuration after changes are made.

You perform these operations by using the export, import, and restore functions available for distributed switches.

You have other available restore options if a switch configuration is lost, for example, when vCenter database corruption occurs, or if the virtual switch or port group settings are misconfigured.

Other ways to restore the virtual switch include restoring the database completely or rebuilding the switch. Although both measures restore the switch, they can be time-consuming.

4-20 Exporting Distributed Switch Configurations

You can export distributed switch and distributed port group configurations to a file.

When exporting, you can perform the following tasks:

- Make a backup of your distributed switch configuration.
- Create a template of a distributed switch configuration.
- Create a revision control system for your distributed switch configuration.

You export the distributed switch and distributed port group configuration to a file on the system that is running the vSphere Client. The file preserves valid network configurations, which can then be distributed to other deployments.

You can use the backup file created from the export function as a template for creating similar distributed switch configurations on other vCenters.

You can keep revisions by saving the distributed switch configuration after each change. By keeping revisions, you can restore the current configuration to an older configuration if necessary.

You can automate this task with the following PowerCLI cmdlets:

- `Export-VDSwitch`: Exports the configuration of a specified vSphere distributed switch to a ZIP file.
- `Export-VDPortGroup`: Exports the configuration of a specified distributed port group to a specified ZIP file.

Technet24

4-21 Restoring and Importing Distributed Switch Configurations

After you export a distributed switch configuration, you can restore or import a configuration:

- Restoring resets the configuration of an existing distributed switch from an exported configuration file.
- Importing creates distributed switch from an exported configuration file.

You can use the restore function to reset a distributed switch configuration that has become corrupted.

You can use the import function to create distributed switch, for example, on a different vCenter system.

With the restore and import functions, you can recreate a distributed switch configuration.

Restoring a distributed switch configuration overwrites the current settings of the distributed switch and its port groups. Full functionality is restored if network settings fail or the vCenter database is corrupted. The restore function does not delete port groups that are not part of the configuration file.

When you import a distributed switch configuration a new distributed switch is created, so you can create multiple copies of an existing deployment.

4-22 Rollback and Recovery of the Host Management Network

Rollback prevents the accidental misconfiguration and loss of host connectivity to vCenter by rolling back to the previous valid management network configuration.

Rollback provides the following options to recover from management network misconfigurations:

- Automatic rollback if a misconfiguration is detected
- Direct Console User Interface (DCUI) to recover the management network manually in the event rollback is disabled

Automatic rollback and recovery is a vSphere feature that helps prevent management network outages. You protect the management network in the following ways:

- The automatic rollback feature detects any configuration changes on the management network. If the host cannot reach vCenter, the changes do not take effect.
- You can also reconfigure the management network on the distributed virtual switch by using the Direct Console User Interface (DCUI). Using the DCUI to change the management network parameters of the switch changes those parameters for all hosts connected to the switch.

Technet24

4-23 Rollback Details

By rolling back configuration changes, vSphere protects hosts from losing connection to vCenter because of misconfiguration of the management network.

Updates that trigger a rollback:

- Host-level rollback: Triggered by a change in the host networking configurations, such as a physical NIC speed change, a change in the MTU configuration, or a change in IP settings.
- Distributed switch-level rollback: Occurs after the user updates distributed switch-related objects, such as a port group or distributed ports.

If you make a change using the vSphere client, vCenter can trigger rollback on a vDS. If you make a change using CLI or host client, then the host will trigger rollback.

Rollback is enabled by default.

The following types of updates can also trigger a rollback:

- If a change in the host networking configuration causes connectivity problems between the host and the vCenter system, a rollback occurs. Examples of misconfigurations include a mismatch in the speed and duplex of the physical NIC, a mismatch in the jumbo frame's configuration, and a misconfiguration in the IP settings of the management network.
- If misconfiguration of the distributed virtual switch causes connectivity issues between the hosts and the vCenter system, a rollback can occur. Port group and distributed virtual port misconfigurations can cause network downtime caused by host management outages.

4-24 Recovery Through the DCUI

If rollback is deactivated, an alternative recovery method is using the DCUI. This provides an easy way for the user to connect directly to the host and fix the distributed switch properties.

DCUI recovery must be performed per host.

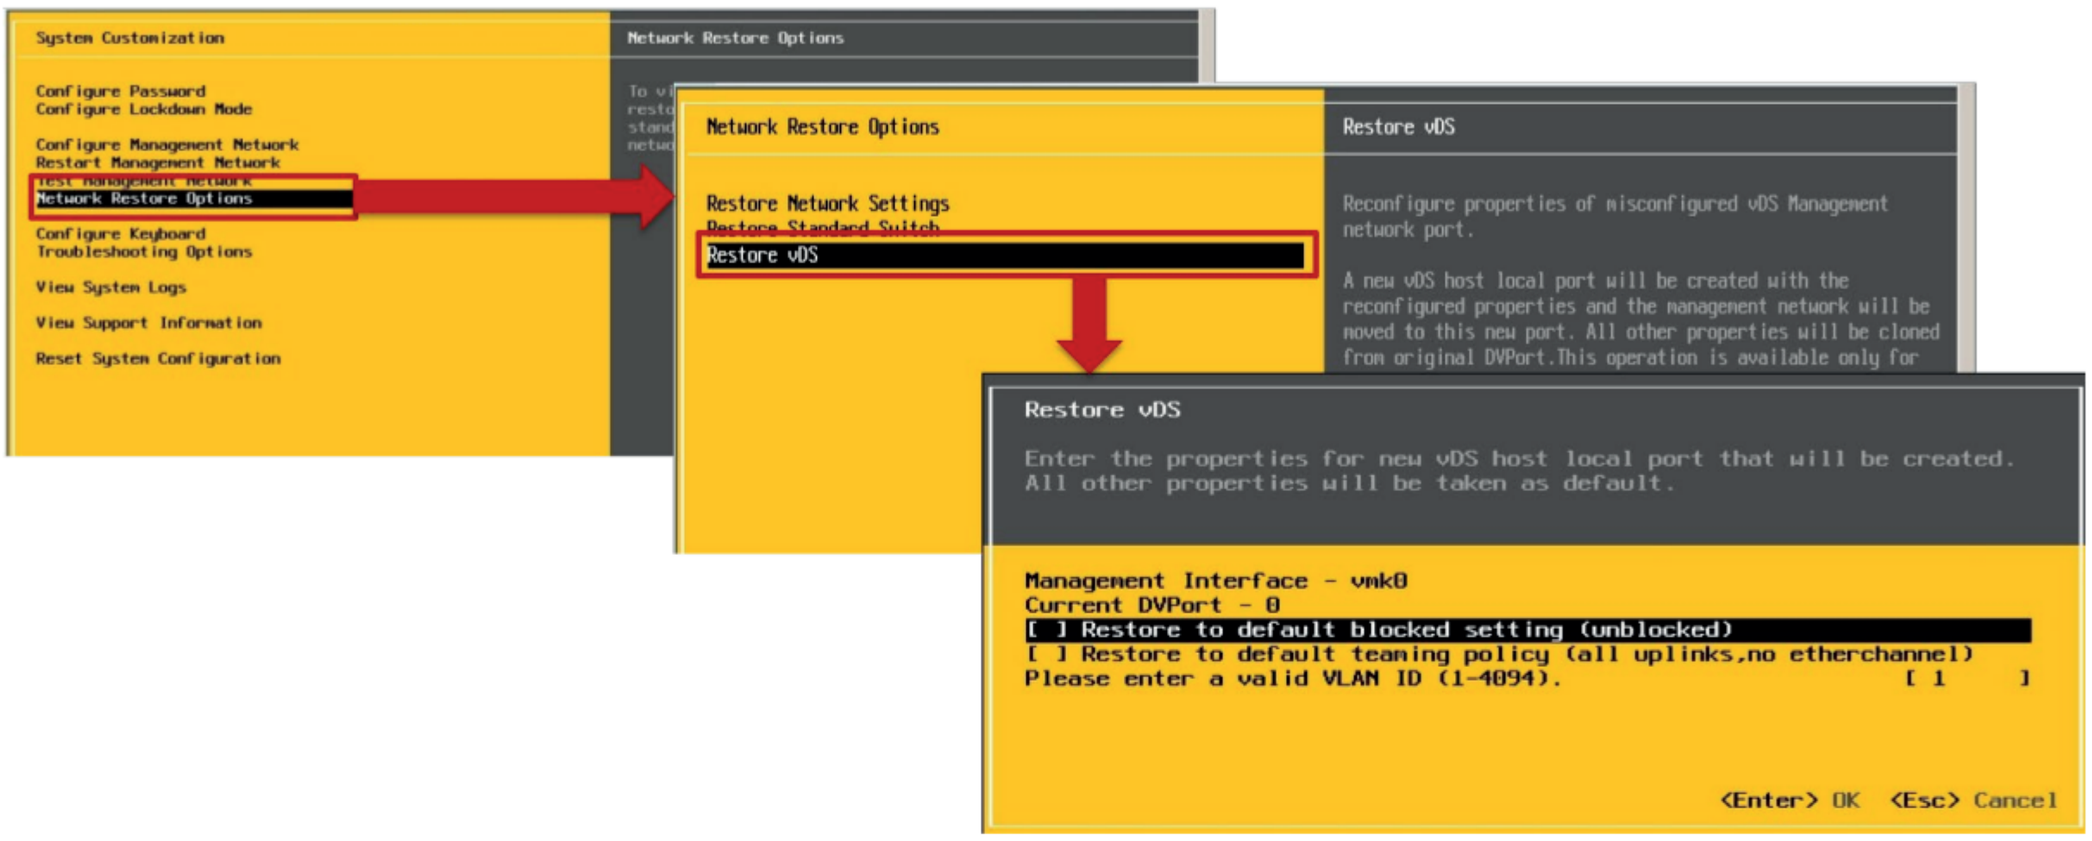

To roll back the configuration manually from the DCUI, take these steps:

1. In the DCUI, press the F2 key.
2. When prompted, log in to the DCUI with the appropriate credentials.
3. Use the down arrow key to scroll to **Network Restore Options** and press Enter.
4. Under Network Restore Options, select **Restore vDS** to roll back the distributed network settings.

   You can either restore the host settings to the default blocked setting or restore the default teaming policy of the host.
5. Select the appropriate option and press Enter.

Technet24

4-25 Lab 7: Configuring vSphere Distributed Switches

Create and configure a distributed switch:

1. Create a Distributed Switch
2. Add ESXi Hosts to the Distributed Switch
3. Examine Your Distributed Switch Configuration
4. Migrate VMs to Another Distributed Switch Port Group

4-26 Lab 8: Managing vSphere Distributed Switches

Use the vSphere Client to create and maintain a vSphere Distributed Switch (VDS) at the data center level:

1. Add a New Port Group to VDS
2. Enable the VDS Health Check
3. Investigate the VDS Health Check Status
4. Remediate the VDS Issue
5. Deactivate the VDS Health Check Service
6. Back Up the VDS Configuration
7. Knowledge Check

4-27 Review of Learner Objectives

- Manage network traffic with the traffic filtering and marking policy
- Recognize NSX virtual distributed switches in the vSphere Client
- Perform a distributed switch health check
- Back up and restore a distributed switch configuration
- Describe the automatic rollback and recovery of networking configurations

4-28 Lesson 2: Using Network I/O Control in Distributed Switches

4-29 Learner Objectives

- Describe how Network I/O Control enhances performance
- Describe how to use Network I/O Control to allocate bandwidth to different types of system traffic
- Describe how to use Network I/O Control to allocate bandwidth to VMs

Technet24

4-30 About Network I/O Control (1)

You can use Network I/O Control to allocate network bandwidth to business-critical applications and to resolve situations when there is insufficient network bandwidth to satisfy the needs of multiple traffic types.

Network I/O Control allocates network bandwidth over distributed switches by using network resource pools for VM and system traffic.

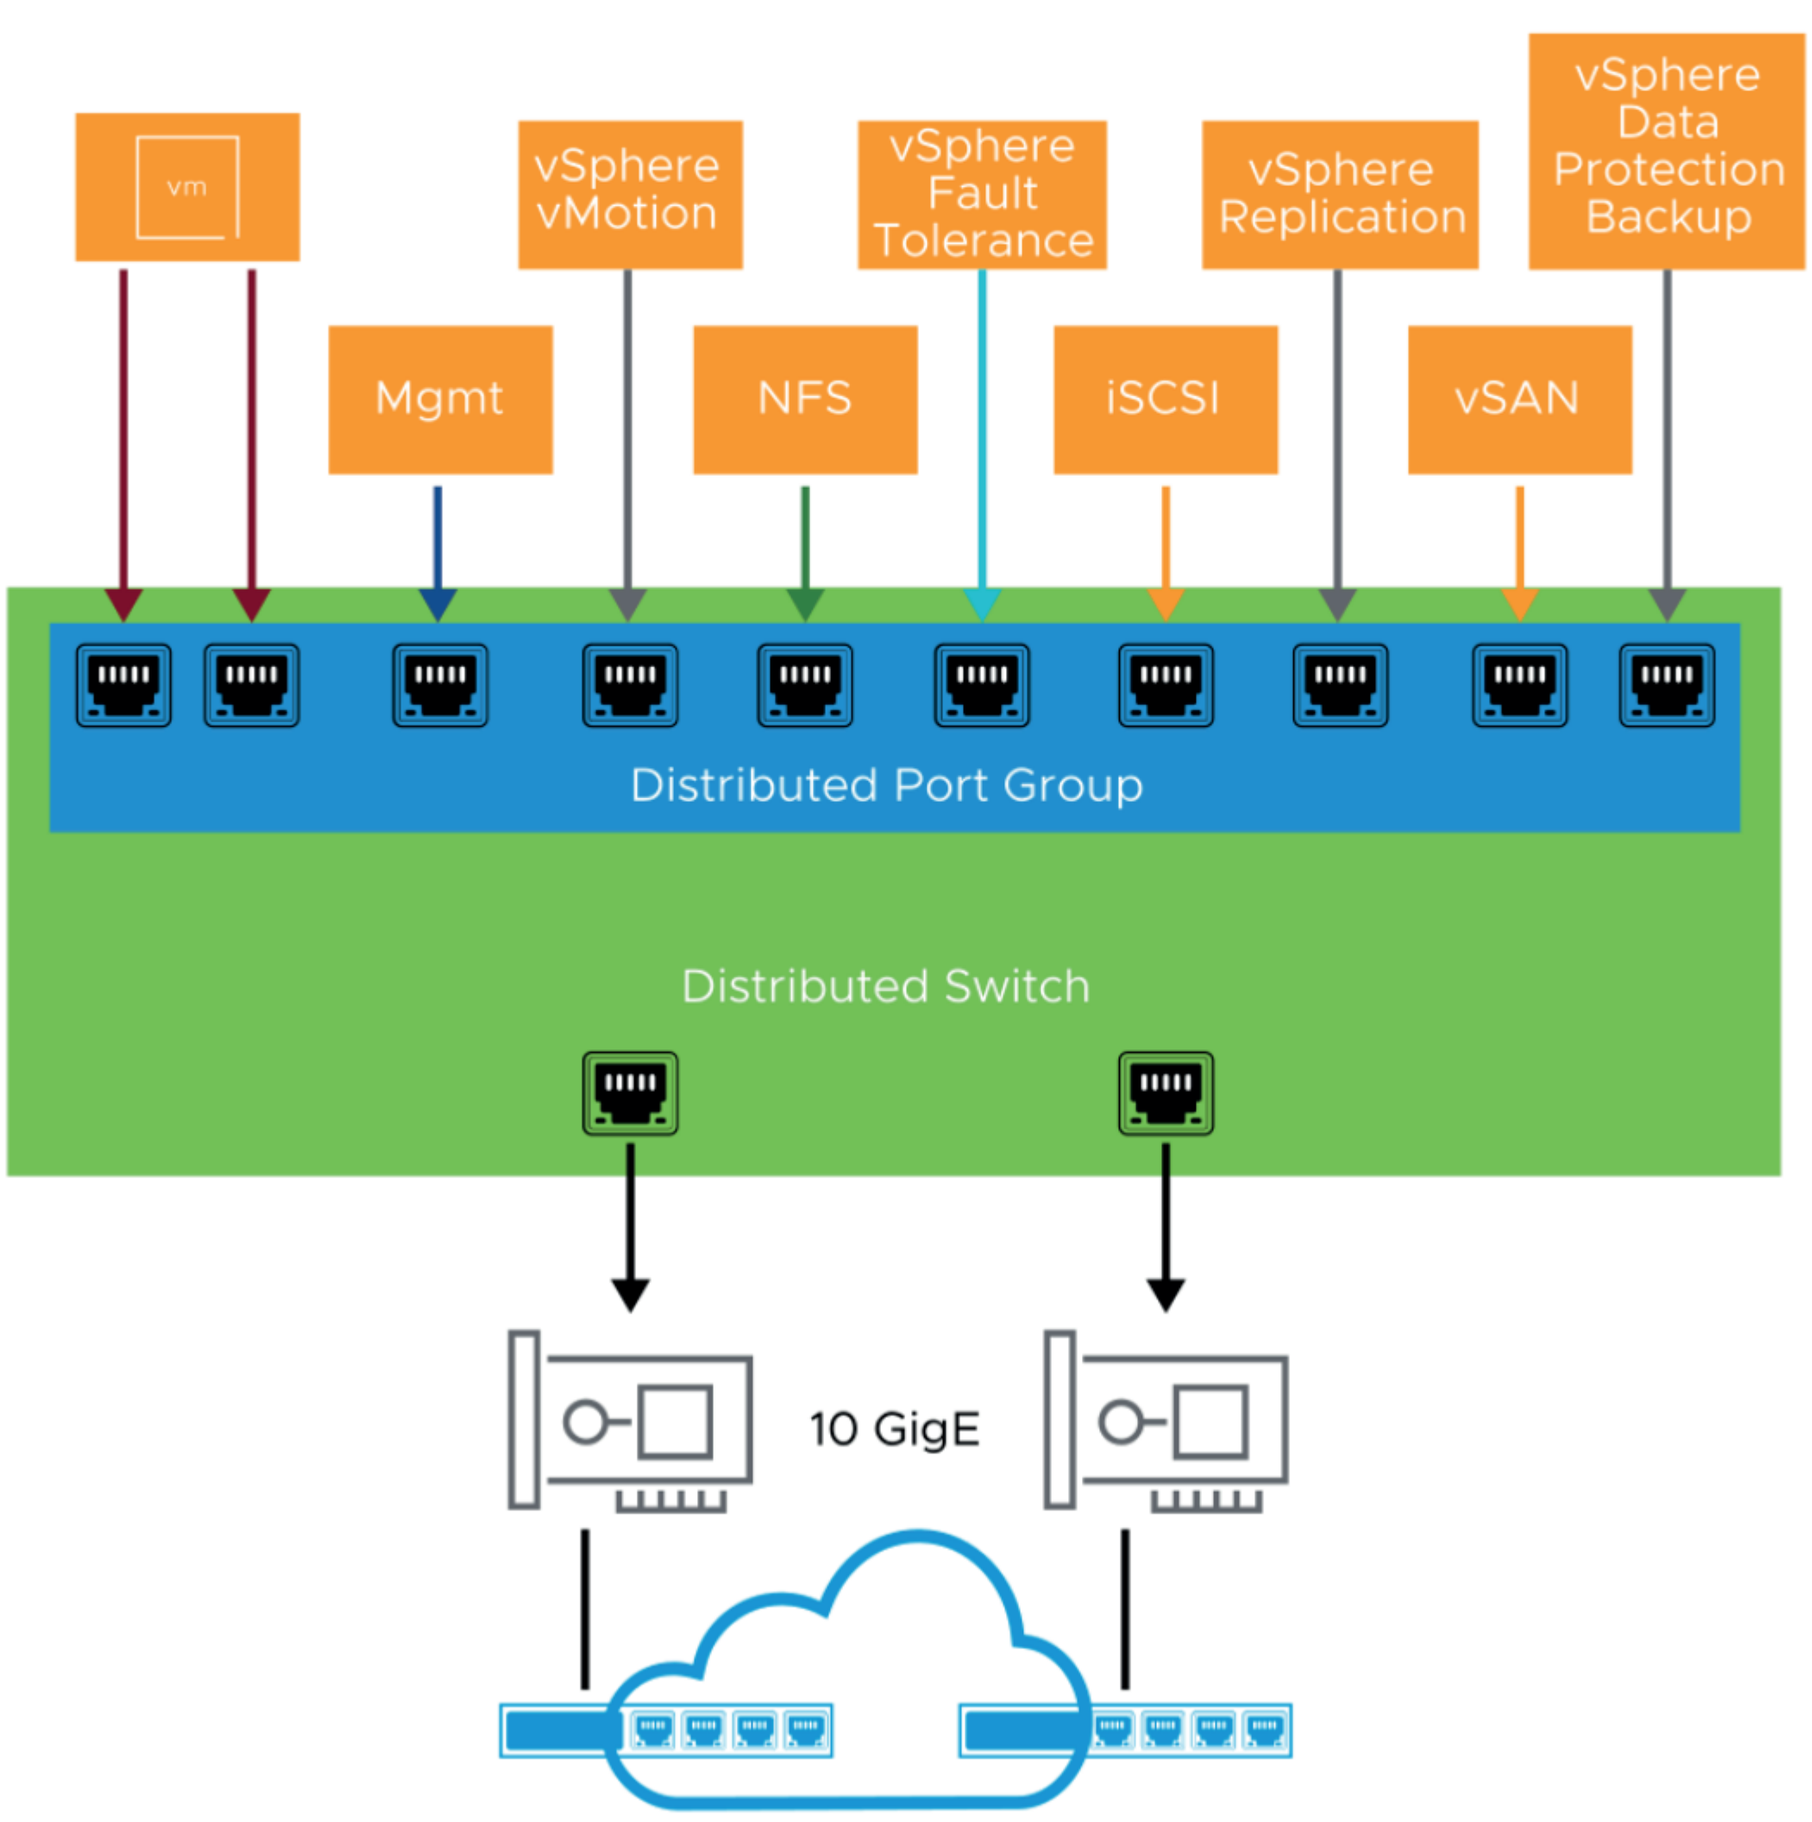

4-31 About Network I/O Control (2)

vSphere 6.x added support for Network I/O Control 3.0, which has the following features:

- You can reserve bandwidth for system traffic and VM traffic based on the capacity of the physical adapters on a host
- You can set limits to define the maximum bandwidth that a certain traffic type can consume on an adapter. If no one else is using the additional bandwidth, the traffic type with the limit also cannot consume it.
- You have detailed resource control at the VM network adapter level.

  This feature is similar to the model that you use for allocating CPU and memory resources. You use shares, reservations, and limits to control bandwidth.

Technet24

4-32 About the Bandwidth Allocation Model for System Traffic

Network I/O Control allows for Shares, Limits and Reservations to be configured to control how bandwidth is allocated to both system traffic (such as management, iSCSI, vSAN etc.) and VM traffic.

Bandwidth Parameter	Description
Shares	The relative priority of a system traffic type against other system traffic types that are active on the same physical adapter. Use the following values to define the number of shares: • Low: 25 • Normal: 50 • High: 100 • Custom: A user-defined value (from 1 to 100)
Reservations	The minimum bandwidth, in Mbps, that must be guaranteed on a single physical adapter.
Limits	The maximum bandwidth, in Mbps or Gbps, that a system traffic type can consume on a single physical adapter.

System Traffic is the management of all the traffic types that use the physical network adapters.

VM Traffic is the management of a specific VM's traffic through its virtual NIC, depending on the system traffic setting for VM traffic.

The amount of bandwidth that is available to a system traffic type is determined by its relative shares and the amount of data that the other system features transmit.

Reserved bandwidth that is unused becomes available to other types of system traffic. However, Network I/O Control does not redistribute unused system traffic capacity to VM placement.

For example, you configure a reservation of 2 Gbps for iSCSI. Because iSCSI uses a single path, the distributed switch might never impose this reservation on a physical adapter. The vacant bandwidth is not allocated to VM traffic.

Network I/O Control cannot safely meet a potential need for bandwidth for system traffic. This failure to reallocate unused reserved bandwidth to VM system traffic is especially true of a new iSCSI path where you must provide bandwidth to a new VMkernel adapter.

4-33 Configuring Bandwidth Allocations for System Traffic

To use Network I/O Control, you can configure shares, bandwidth reservation, and limits for each type of system traffic.

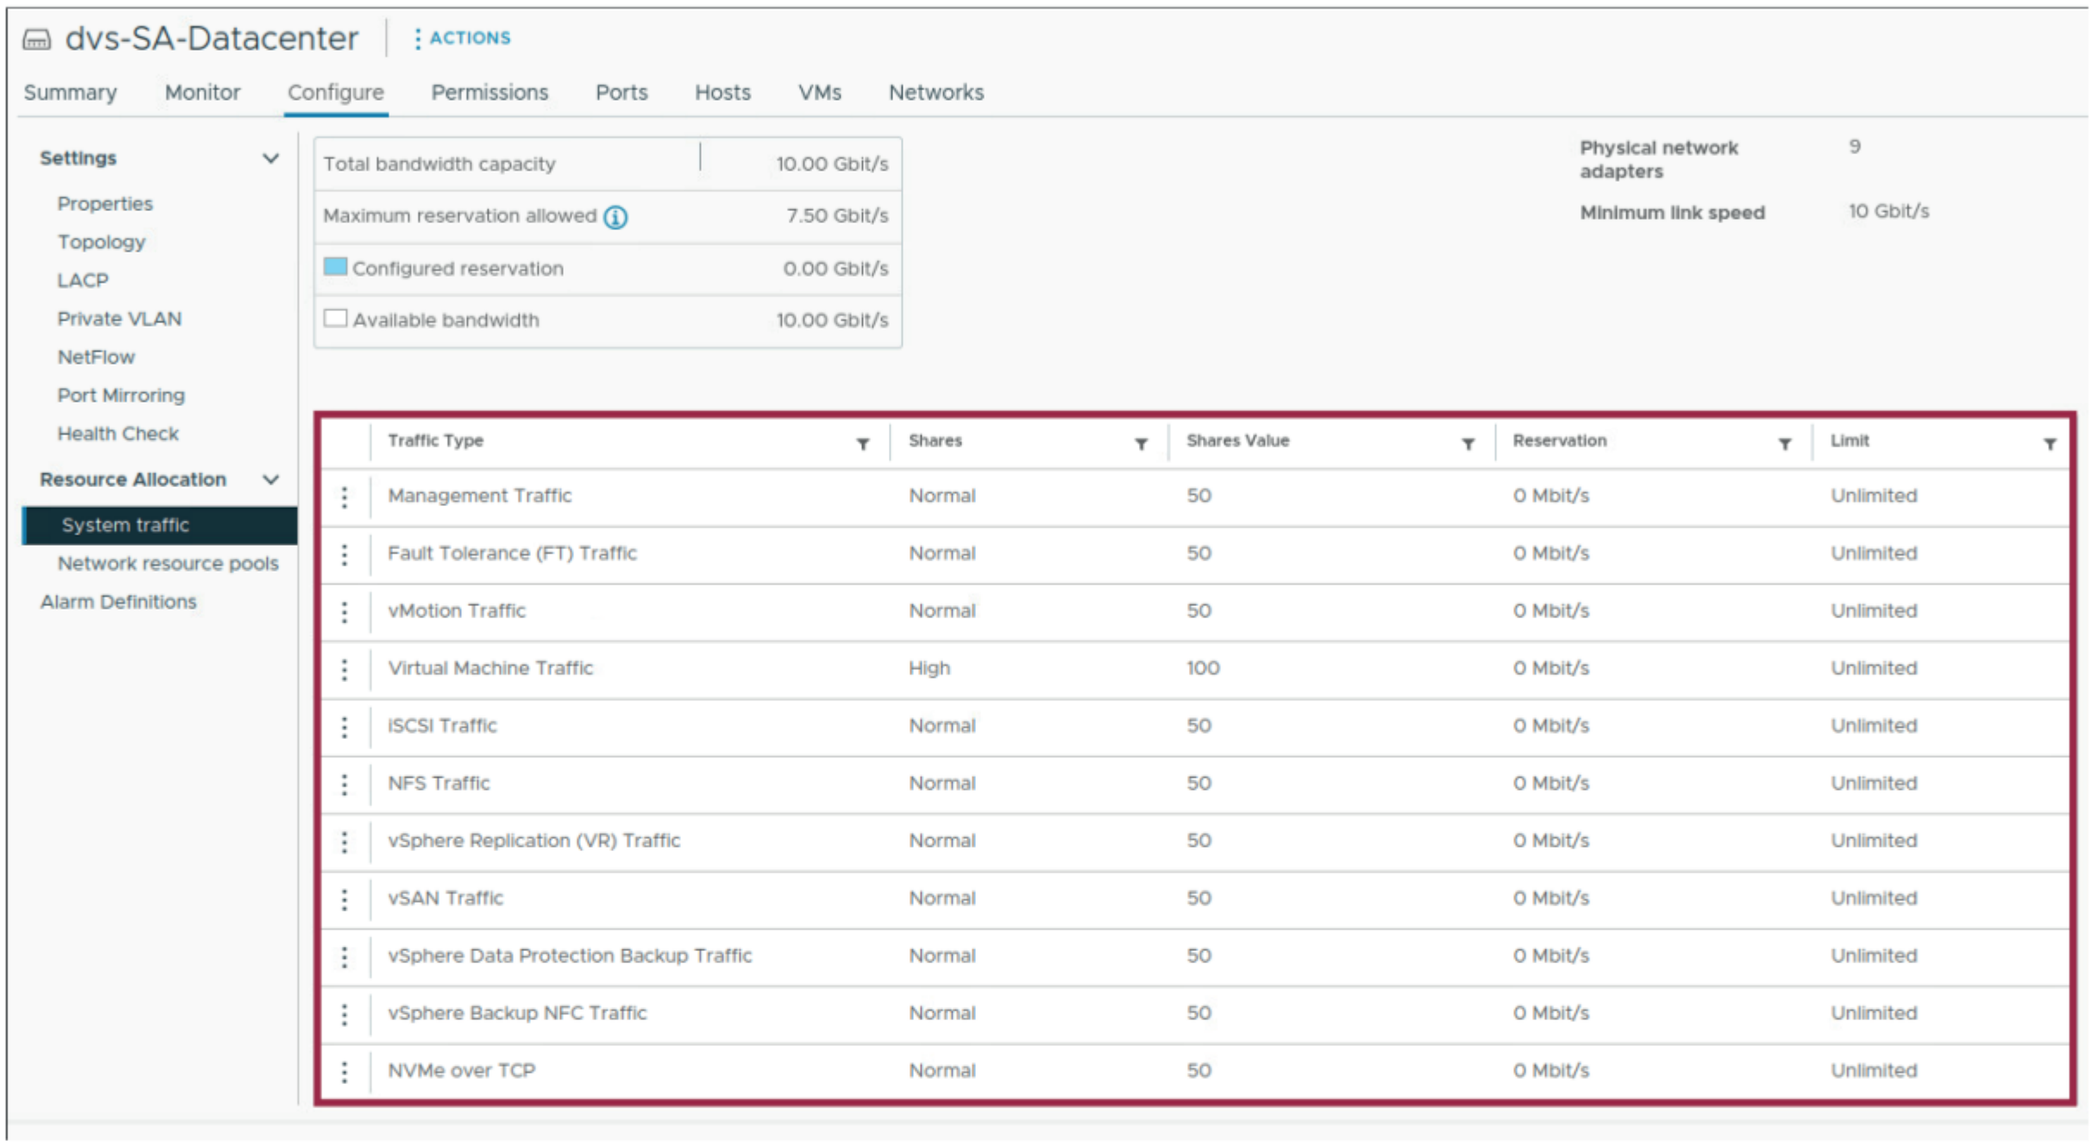

You can use Network I/O Control to configure bandwidth allocation for the traffic that is related to the main features of vSphere:

- VMs
- Management
- vSphere vMotion
- NFS
- vSphere Fault Tolerance
- iSCSI
- vSphere Replication
- vSAN
- vSphere Data Protection

Technet24

4-34 Reserving Bandwidth for System Traffic (1)

Network I/O Control allocates the requested bandwidth on each physical network adapter. You can reserve no more than 75 percent of the bandwidth of a physical network adapter.

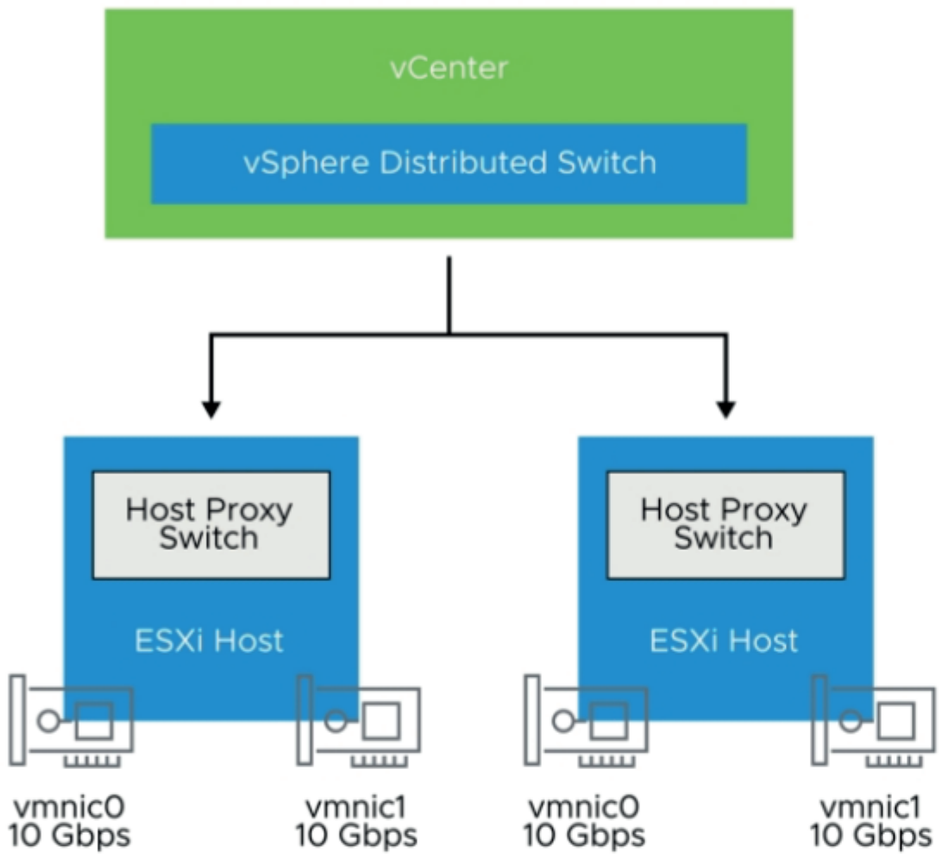

Bandwidth Reservation for System Traffic on Physical NIC

System Traffic	Bandwidth Reservation Gbps
Management	0.5
iSCSI	1.0
vSphere Fault Tolerance	1.0
vSphere vMotion	1.0
Virtual Machine	0.5
Free Capacity	6.0

Total: Less than or equal to 7.5 Gbps (75% of 10 Gbps)

You can guarantee a minimum bandwidth to a system feature for optimal operation according to the capacity of the physical adapters.

For a distributed switch that is connected to ESXi hosts with 10 GbE network adapters, you can configure a reservation to guarantee 1 Gbps for vSphere vMotion traffic, 1 Gbps for vSphere Fault Tolerance, 0.5 Gbps for VM traffic, and so on.

You might leave more capacity unreserved so that the host can dynamically allocate bandwidth according to shares, limits, and use, and to reserve only enough bandwidth for the operation of a system feature.

You can reserve up to 75 percent of the bandwidth of a physical network adapter. For example, if your ESXi host uses a 40 GbE network adapter, 30 Gbps (.75 x 40) can be reserved.

4-35 Reserving Bandwidth for System Traffic (2)

The vSphere Client shows bandwidth use information and the share, reservation, and limit values for each traffic type.

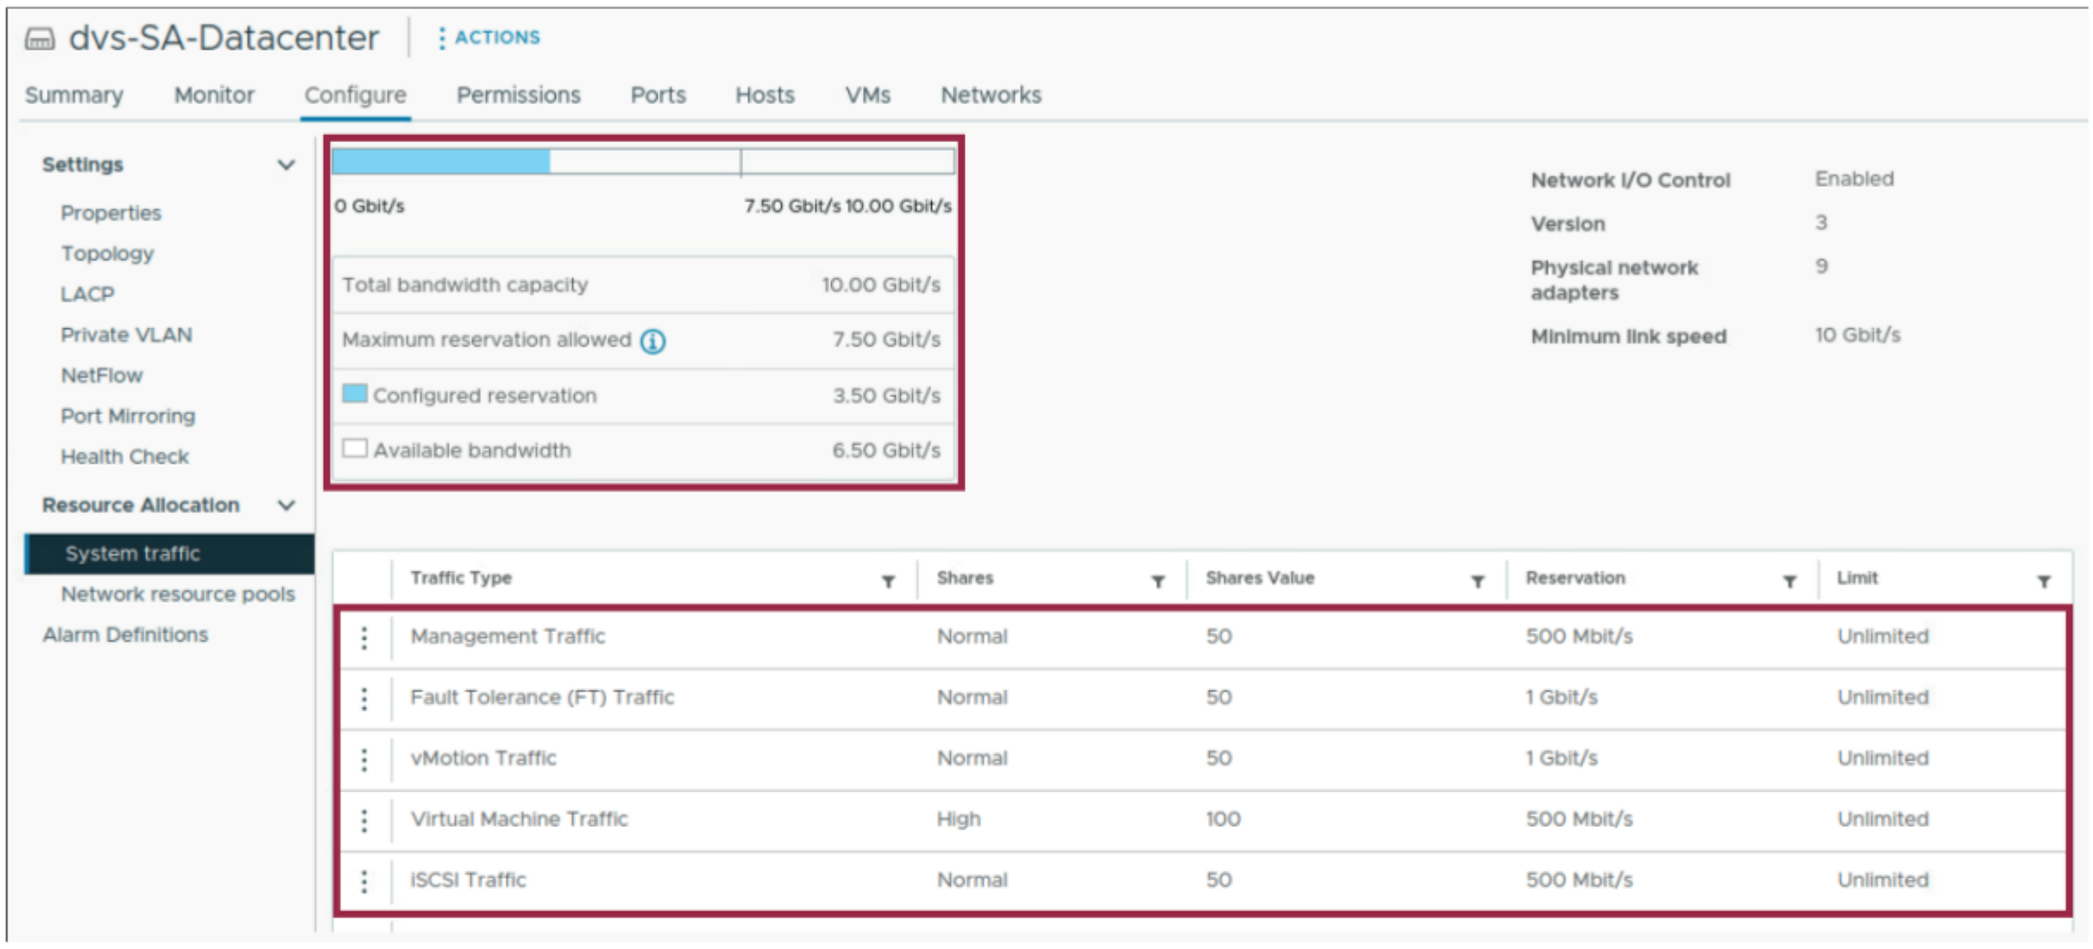

Technet24

4-36 Bandwidth Aggregation for Network Resource Pools

A network resource pool represents a part of the aggregated bandwidth that is reserved for the virtual machine system traffic on all physical adapters connected to the distributed switch.

By default, distributed port groups are assigned to the network resource pool named default, where no quota is configured.

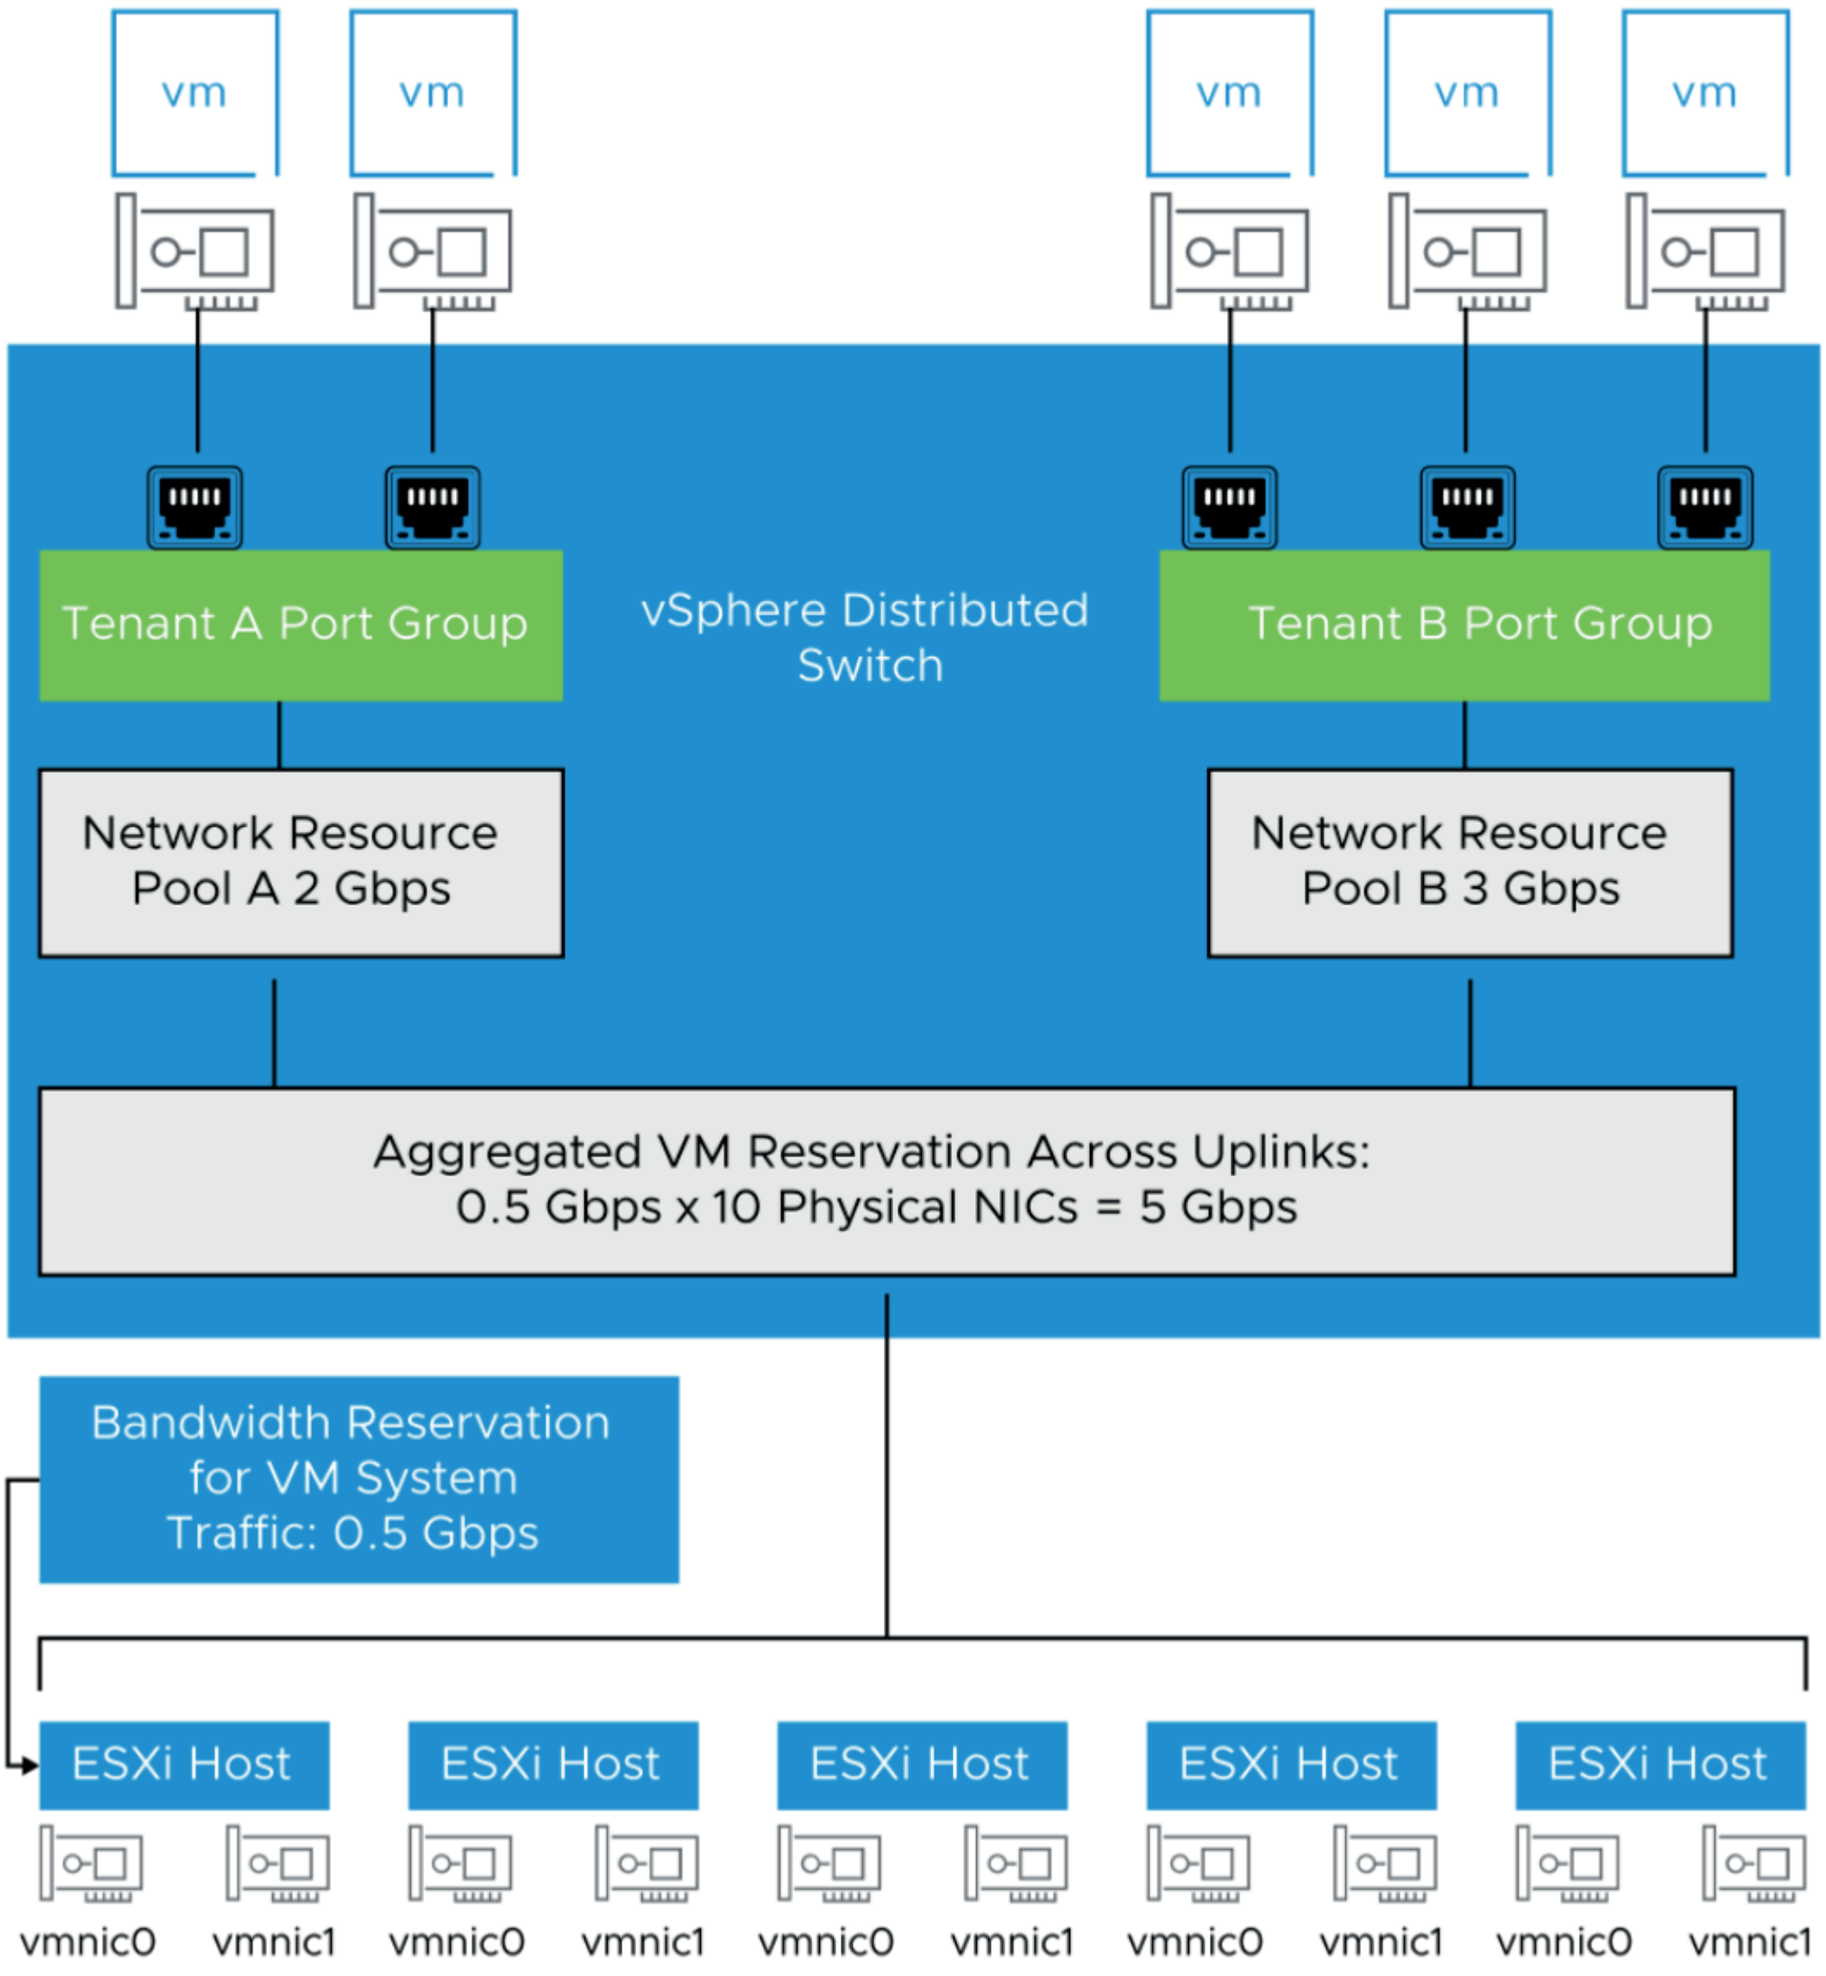

(vmnic0 and vmnic1 are 10 Gbps adapters.)

Network I/O Control allocates bandwidth for VMs across the entire distributed switch and on the physical adapter carrying the VM traffic.

To use Network I/O Control to enable bandwidth allocation for VMs, configure the VM system traffic. The bandwidth reservation for VM traffic is also used in admission control. When you power on a VM, admission control verifies that enough bandwidth is available.

4-37 Creating Network Resource Pools

Once a reservation has been configured for the VM traffic type, you can create a network resource pool to reserve bandwidth for a set of VMs.

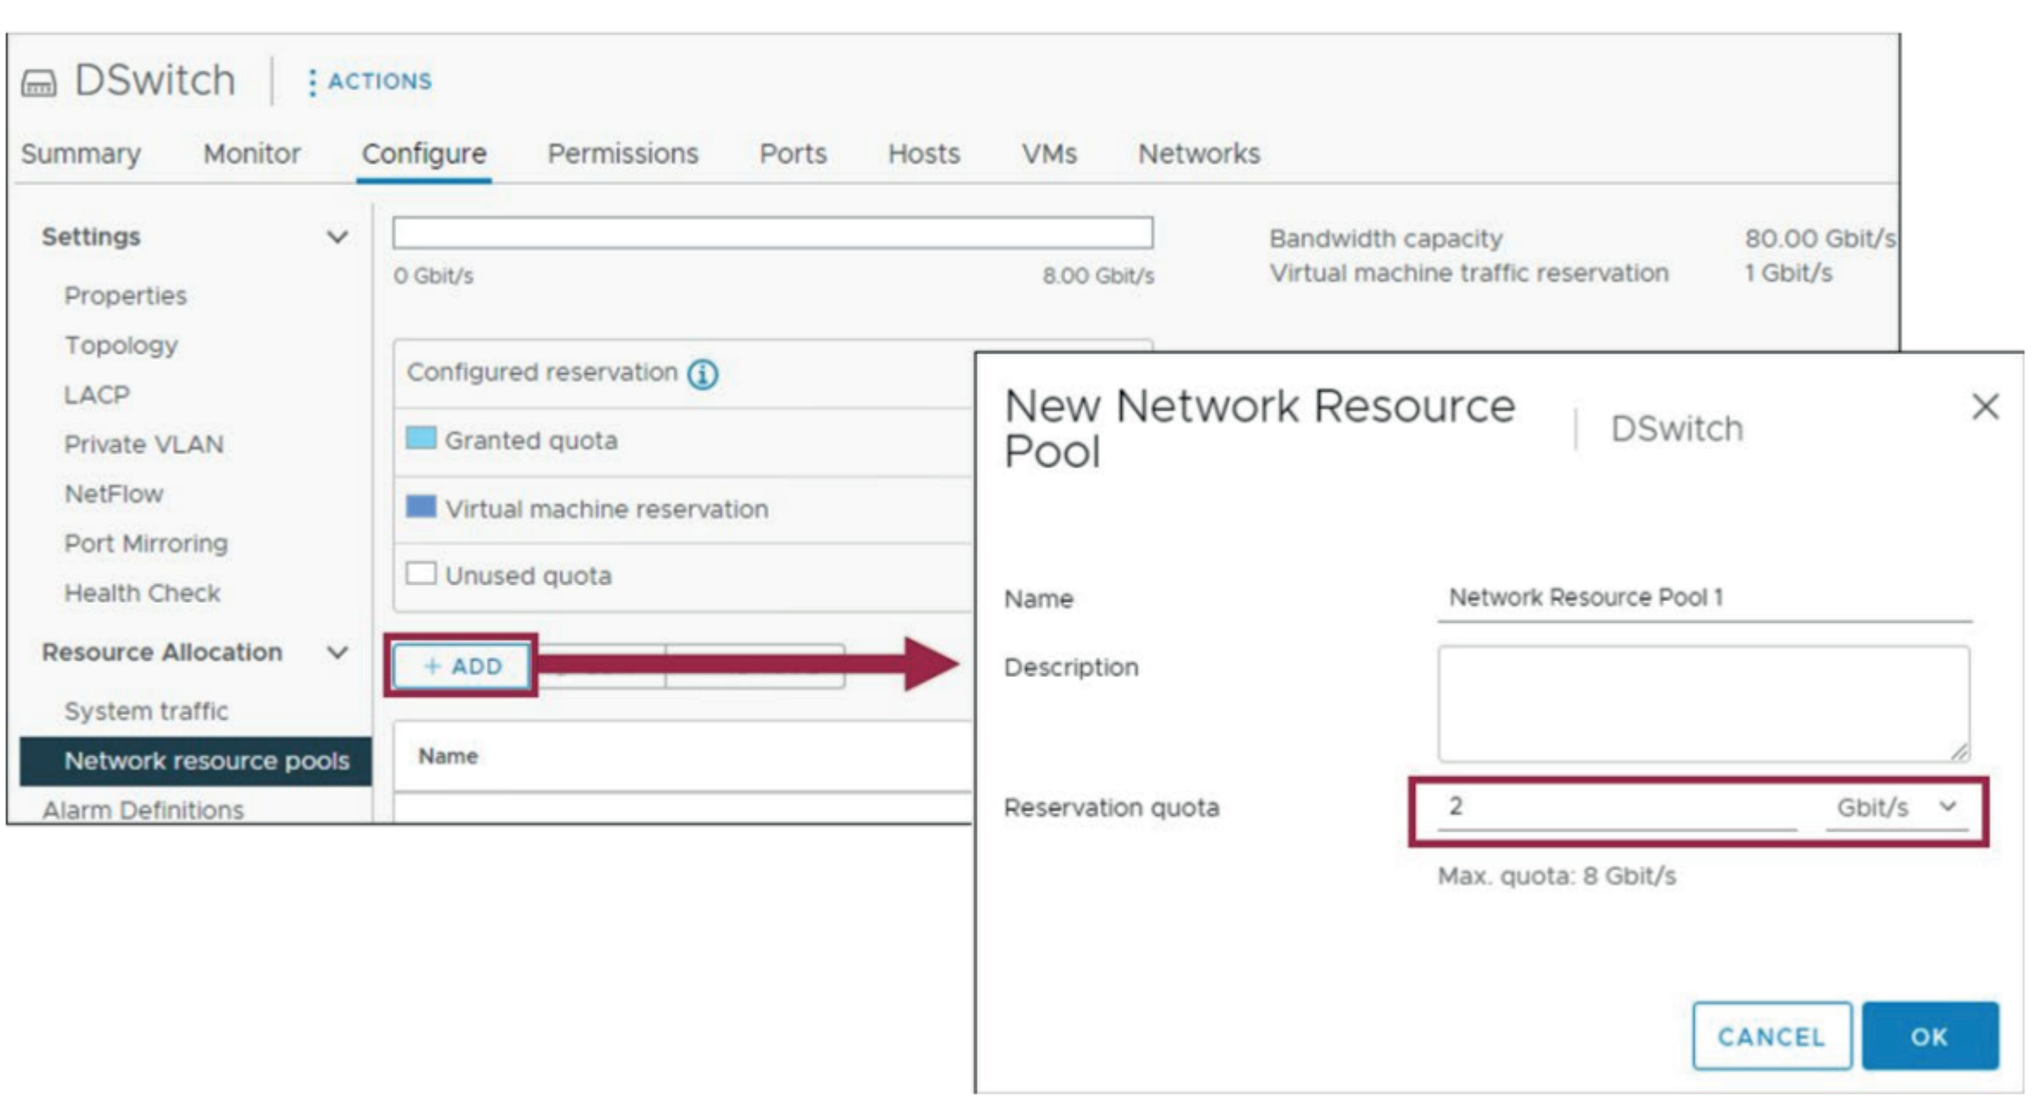

For example, if the VM system traffic has 0.5 Gbps reserved for each 10 GbE uplink on a distributed switch that has 10 uplinks, the total aggregated bandwidth that is available for VM reservation on this switch is 5 Gbps. Each network resource pool can reserve a quota of this 5 Gbps capacity.

The bandwidth reservation that is dedicated to a network resource pool is shared among the distributed port groups associated with the pool. A VM receives bandwidth from the pool through the distributed port group that the VM is connected to.

Technet24

4-38 Bandwidth Allocation for VM Traffic

Bandwidth is allocated to individual VMs based on configured shares, reservations, and limits for the VM network adapters.

Bandwidth Parameter	Description
Shares	The relative priority of the traffic type through the VM network adapter against the capacity of the physical adapter that is carrying the VM traffic to the network. Use the following values to define the number of shares: • Low: 25 • Normal: 50 • High: 100 • Custom: A user-defined value (from 1 to 100)
Reservations	The minimum bandwidth, in Mbps, that the VM network adapter must receive on the physical adapter.
Limits	The maximum bandwidth on the VM network adapter for traffic to other VMs on the same or on another host.

A network resource pool provides a reservation quota to VMs. The quota represents a portion of the bandwidth that is reserved for VM system traffic on the physical adapters connected to the distributed switch.

You can set aside bandwidth from the quota for the VMs that are associated with the pool. The reservation from the network adapters of powered-on VMs that are associated with the pool must not exceed the quota of the pool.

4-39 Bandwidth Allocation for Individual VMs

The total bandwidth reservation of the VMs on a host cannot exceed the reserved bandwidth that is configured for the VM system traffic.

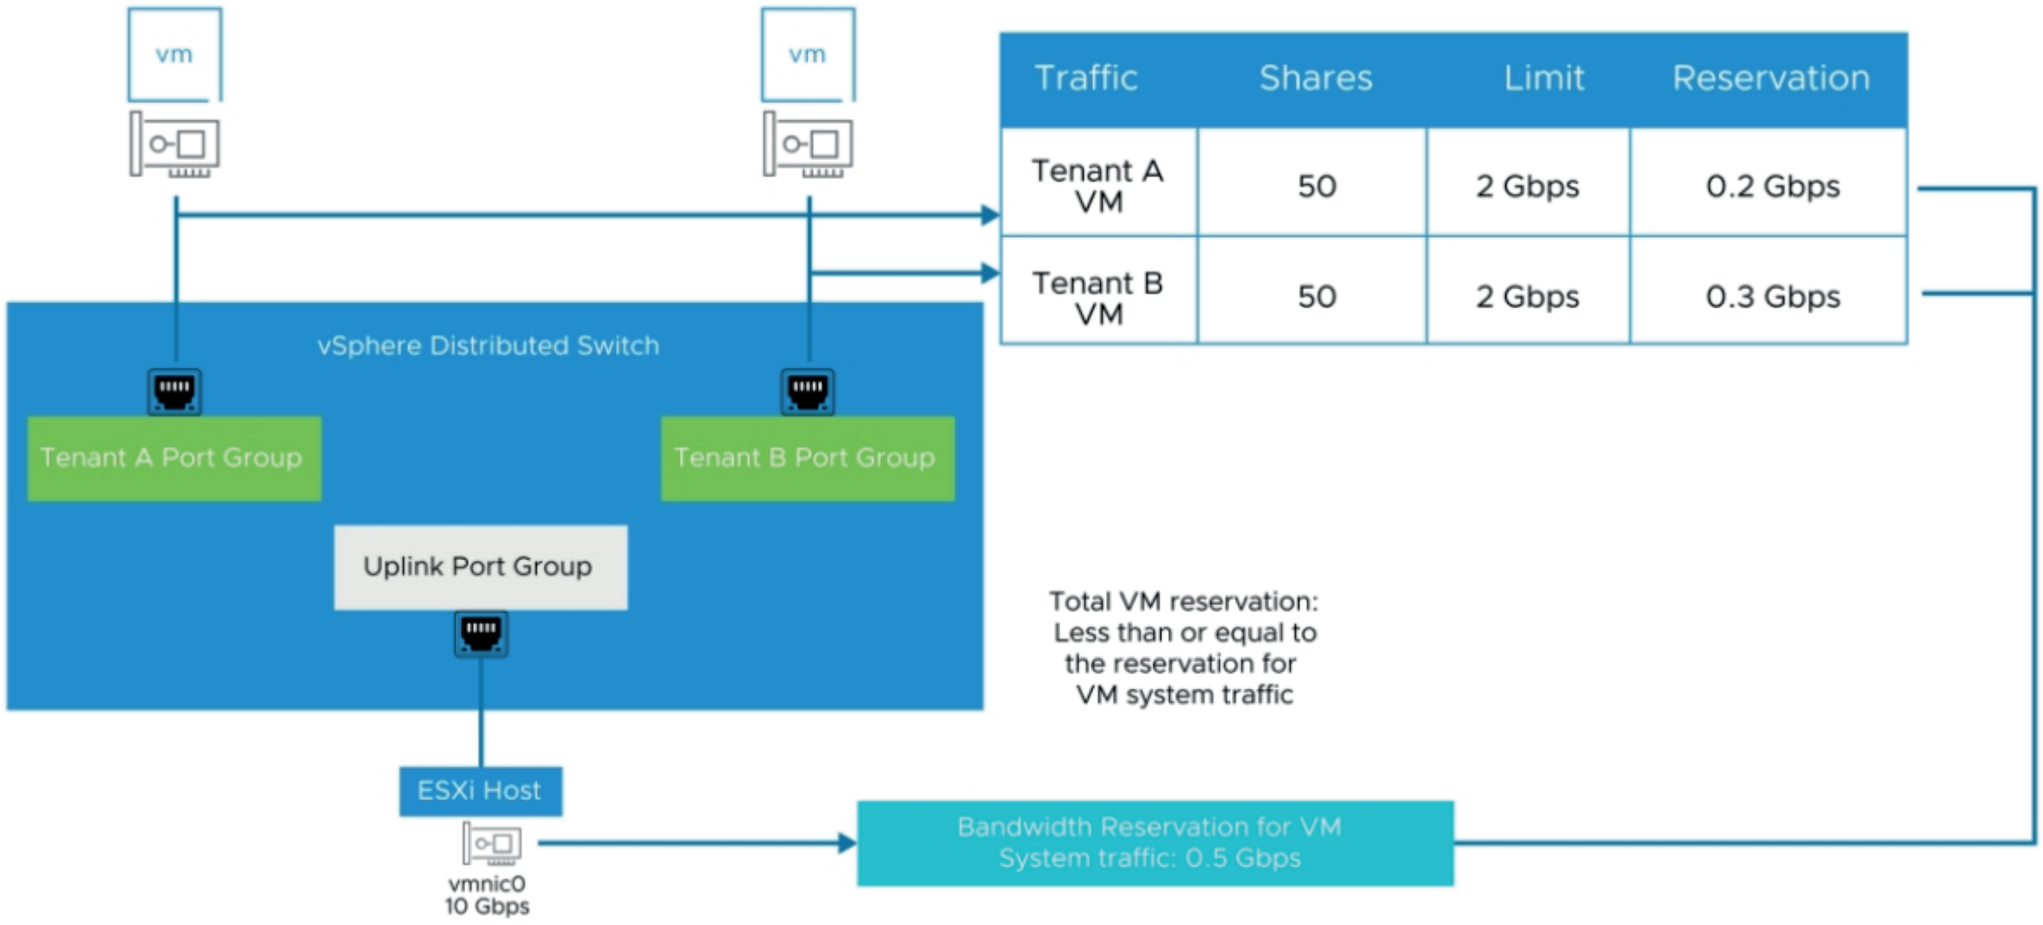

To guarantee bandwidth, Network I/O Control implements a traffic placement engine that becomes active if a VM has bandwidth reservation configured. The distributed switch tries to direct the traffic from a VM network adapter to a physical adapter that can supply the required bandwidth and is in the scope of the active teaming policy.

The real limit and reservation also depend on the traffic-shaping policy on the distributed port group that the adapter is connected to. For example, if a VM network adapter asks for a limit of 200 Mbps and the average bandwidth configured in the traffic-shaping policy is 100 Mbps, the effective limit becomes 100 Mbps.

Technet24

4-40 Bandwidth Admission Control in vSphere DRS (1)

Bandwidth admission control verifies that the VM reservation can be met.

When a VM is powered on in a DRS-enabled cluster, vSphere DRS places the VM on a host that has the capacity to guarantee the bandwidth that is reserved for the VM.

If you power on a VM that is in a cluster, vSphere DRS places that VM on a host with the capacity to guarantee the bandwidth that is reserved for the VM, according to the active teaming policy.

To use admission control in vSphere DRS, you must perform the following tasks:

- Configure bandwidth allocation for the VM system traffic on the distributed switch.
- Configure the bandwidth requirements of a VM that is connected to the distributed switch.

If a VM reservation cannot be met, the VM is not powered on.

4-41 Bandwidth Admission Control in vSphere DRS (2)

vSphere DRS migrates a VM to another host to satisfy the bandwidth reservation of a VM in the following situations:

- The reservation is changed to a value that the current host can no longer satisfy.
- The physical adapter carrying the VM's traffic goes offline.

In this example, Host1 loses an uplink, which leaves Host1 with one working uplink. Because Host2's uplinks still work, vSphere DRS migrates either VM1 or VM2 to Host2 to meet the VM's network reservation requirement.

4-42 Bandwidth Admission Control in vSphere HA

When a host fails, vSphere HA powers on the failed VMs on another host in the cluster, that satisfies the bandwidth reservation and teaming policy.

Bandwidth admission control prevents a VM from starting if the bandwidth reservation for that VM cannot be met.

To use admission control in vSphere HA, you must perform the following tasks:

- Allocate bandwidth for VM system traffic.
- Configure the bandwidth requirements of a VM that is connected to the distributed switch.

4-43 Activity: Network I/O Control (1)

You want to guarantee that the VMs connected to a distributed port group called Prod 1 can access 25 percent of the available reserved bandwidth for VM traffic.

Use the bandwidth reservation information in the example to answer the following question.

What must you configure to provide the appropriate bandwidth to the VMs on Prod 1?

System Traffic	Bandwidth Reservation Gbps
Management	4
iSCSI	12
Virtual Machine	8

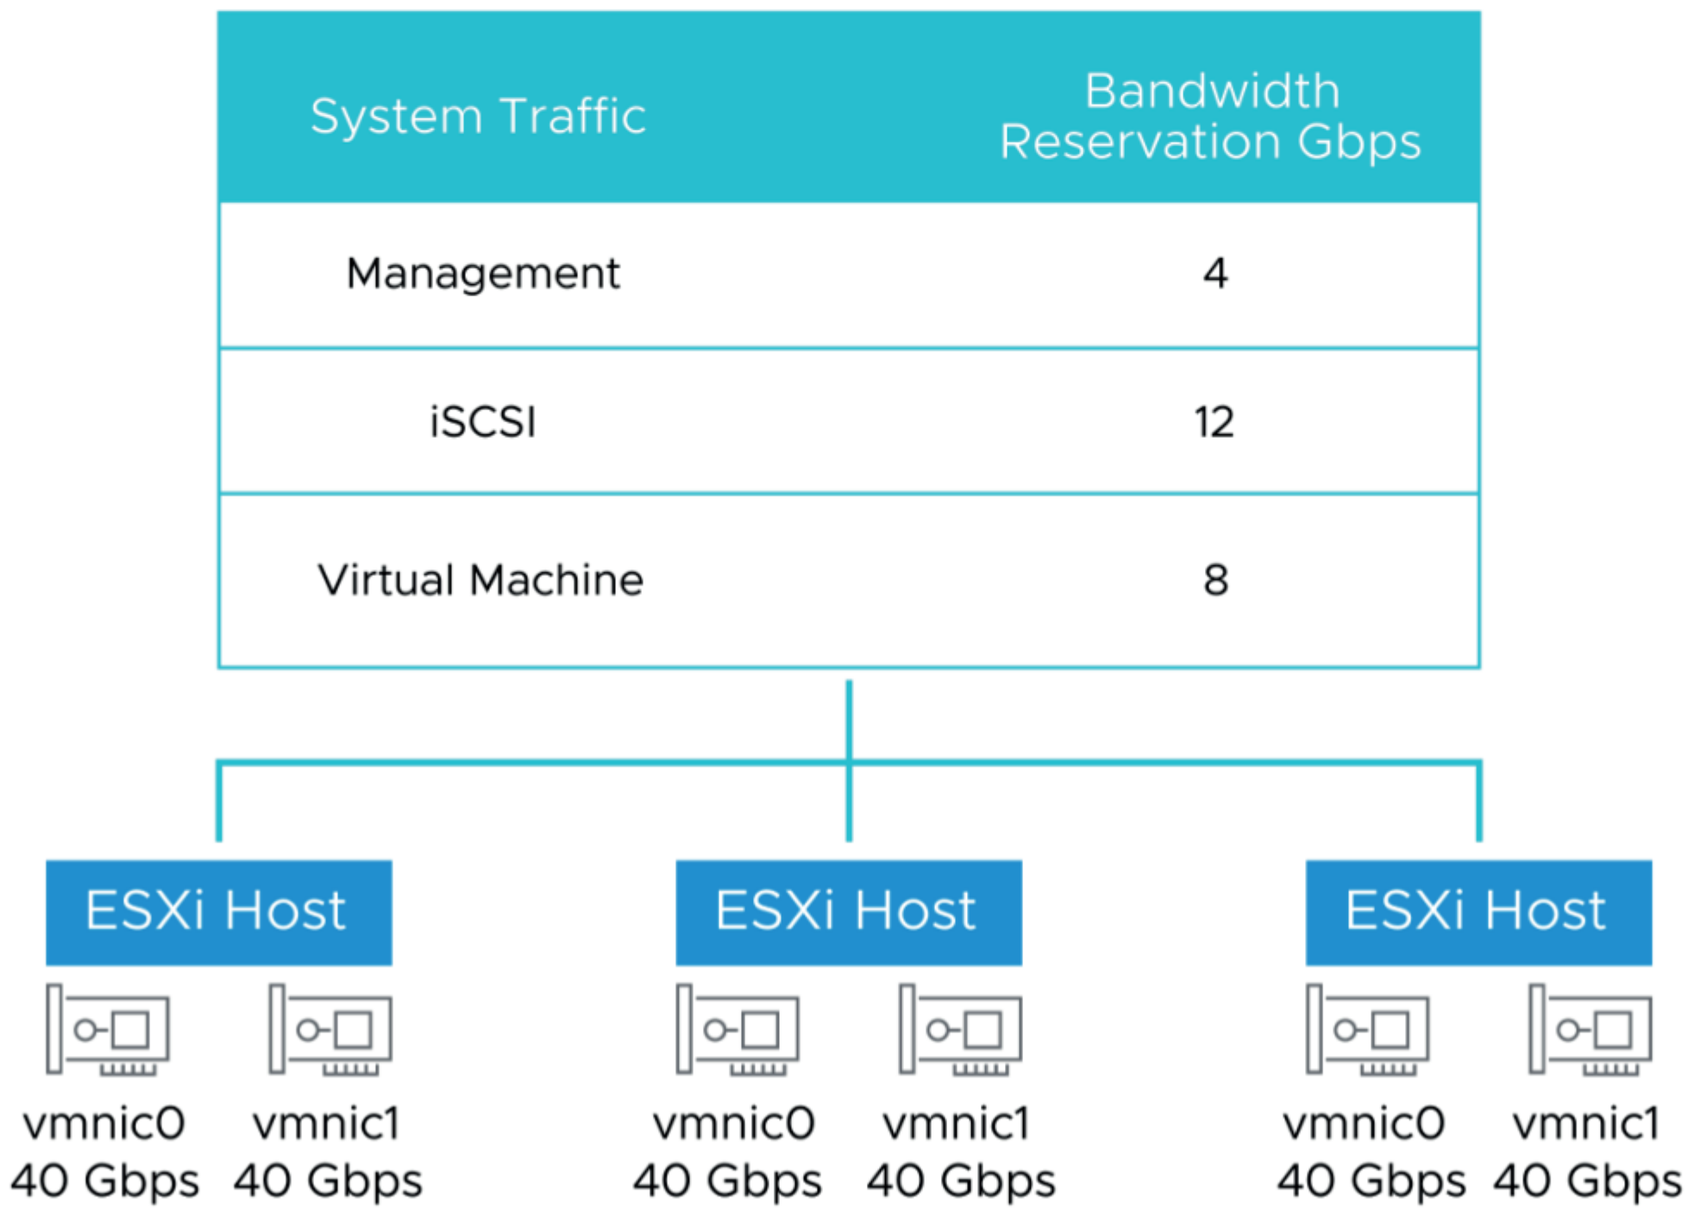

Technet24

4-44 Activity: Network I/O Control (2)

The aggregated amount of reserved bandwidth for VM traffic is 48 Gbps.

Twenty-five percent of the aggregated amount is 12 Gbps.

Therefore, you must configure a network resource pool with a 12 Gbps reservation for Prod 1.

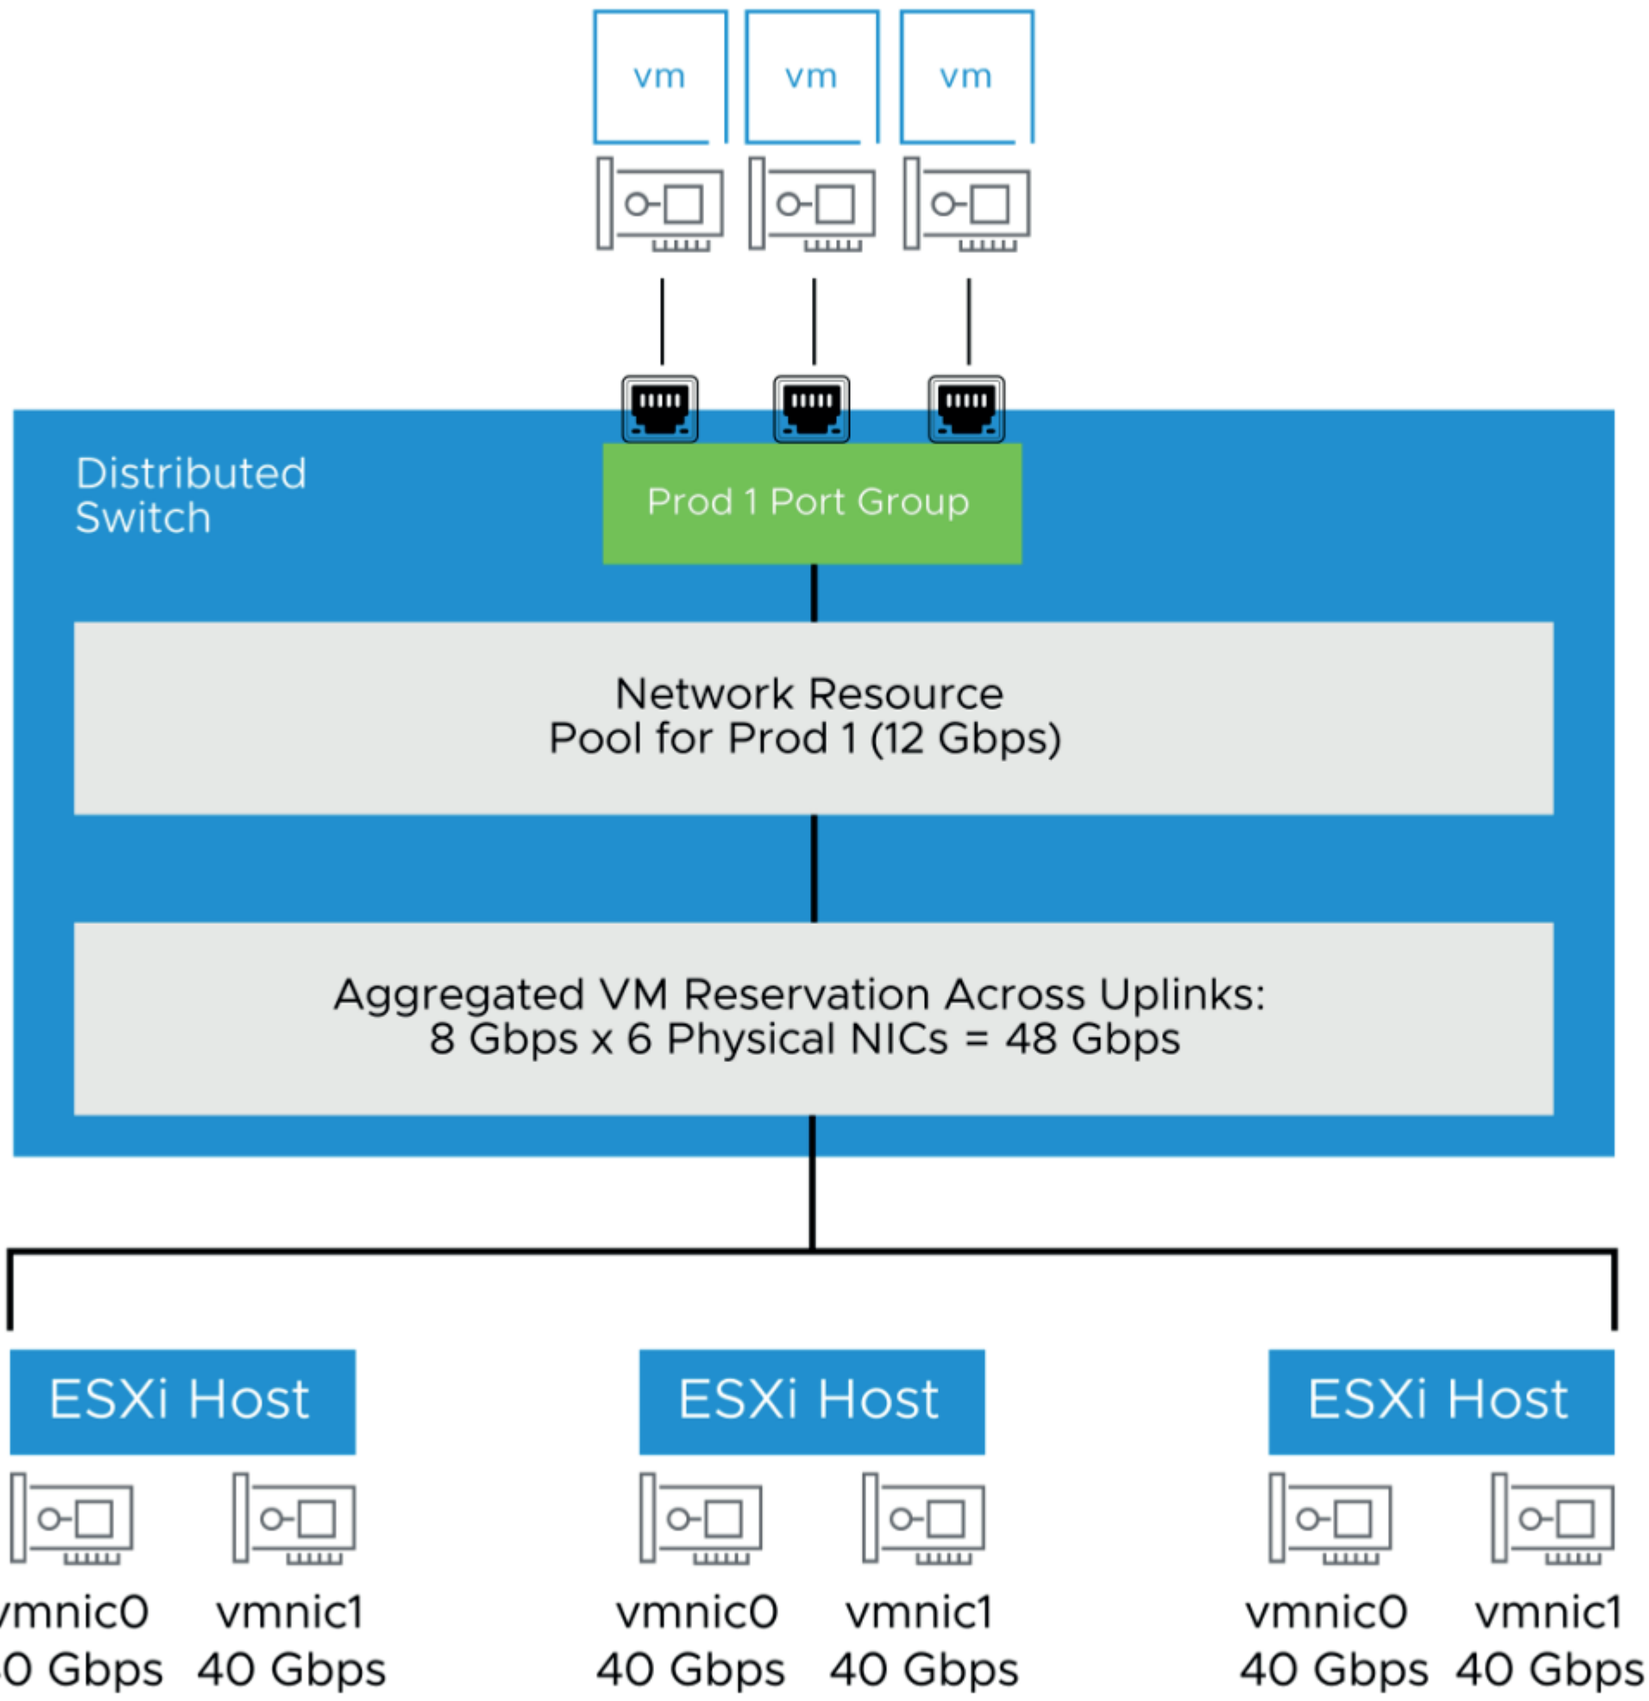

4-45 Review of Learner Objectives

- Describe how Network I/O Control enhances performance
- Describe how to use Network I/O Control to allocate bandwidth to different types of system traffic
- Describe how to use Network I/O Control to allocate bandwidth to VMs

4-46 Lesson 3: NetFlow and Port Mirroring in Distributed Switches

4-47 Learner Objectives

- Describe how to use NetFlow for monitoring network security and performance
- Configure port mirroring on a distributed switch

4-48 About NetFlow

NetFlow is a network analysis tool for monitoring the network and viewing the VM traffic that flows through a distributed switch.

You can use NetFlow for profiling, intrusion detection and prevention, networking forensics, and compliance.

vSphere distributed switches support IPFIX (NetFlow version 10).

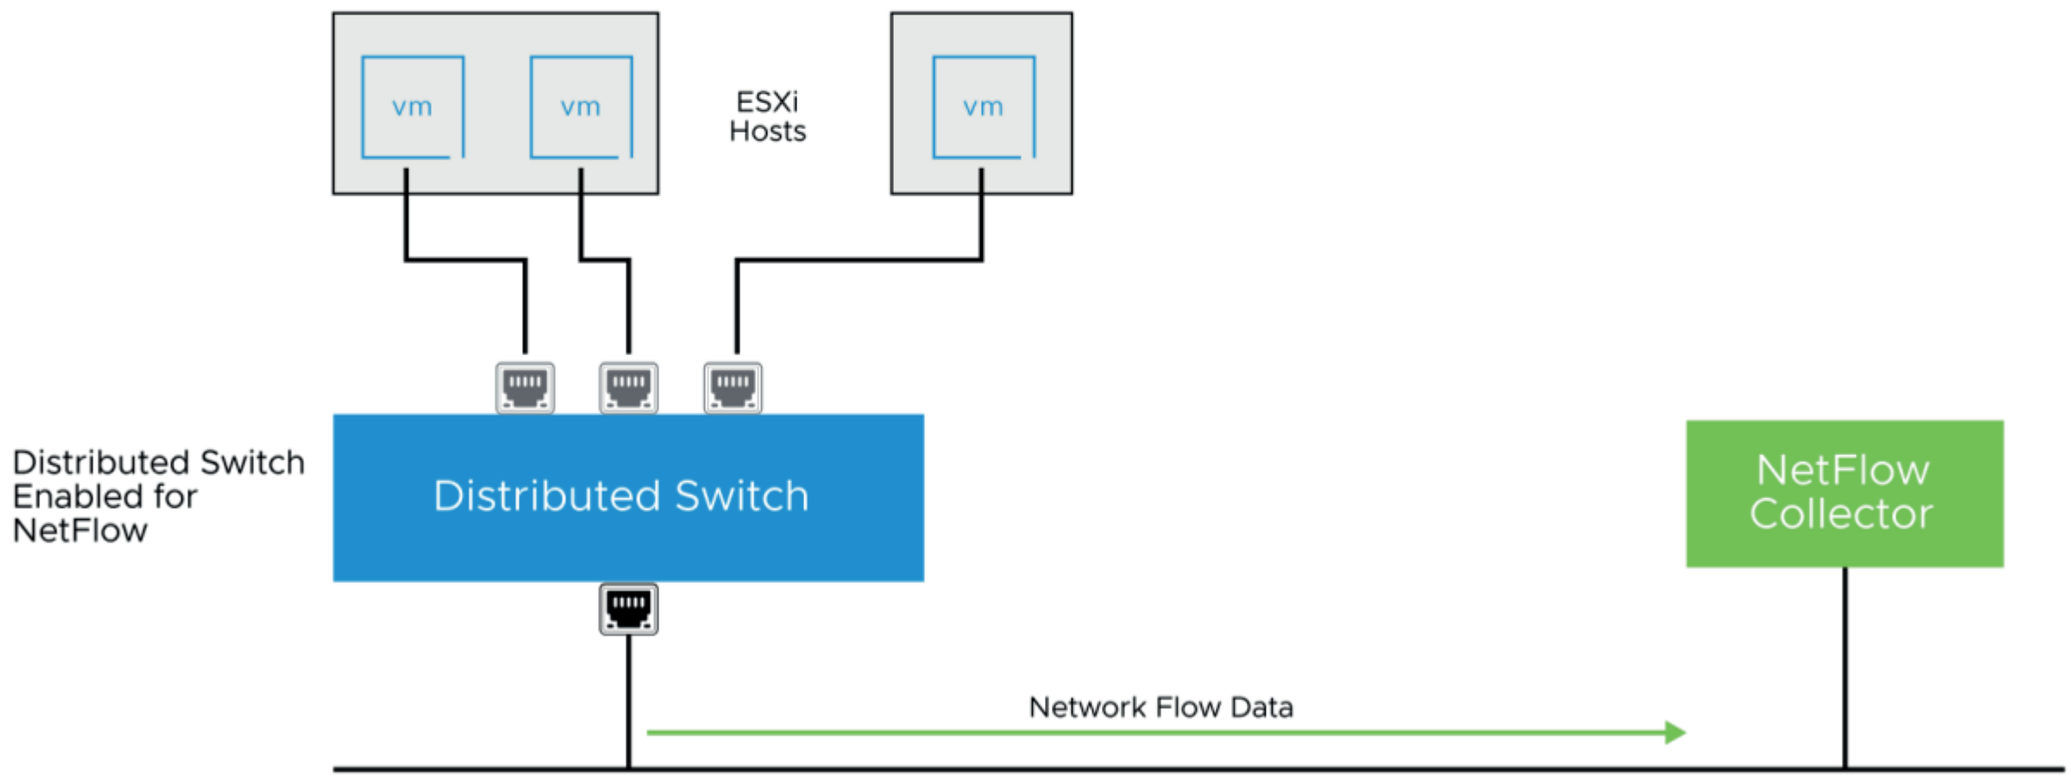

NetFlow is a protocol that Cisco Systems developed for analyzing network traffic. NetFlow is an industry-standard specification for collecting network data for monitoring and reporting. The data sources are network devices, such as switches and routers. For ESXi deployments, NetFlow supports detailed monitoring and analysis of VM network traffic.

Standard switches do not support NetFlow.

NetFlow collectors are available from third-party providers.

Technet24

4-49 About Network Flows

A network flow is a unidirectional sequence of packets. Each packet shares a common set of properties.

NetFlow captures the following types of flows:

- Internal flow: Represents intrahost VM traffic
- External flow: Represents interhost VM traffic and physical machine-to-VM traffic

Flow records are sent to a NetFlow collector for analysis.

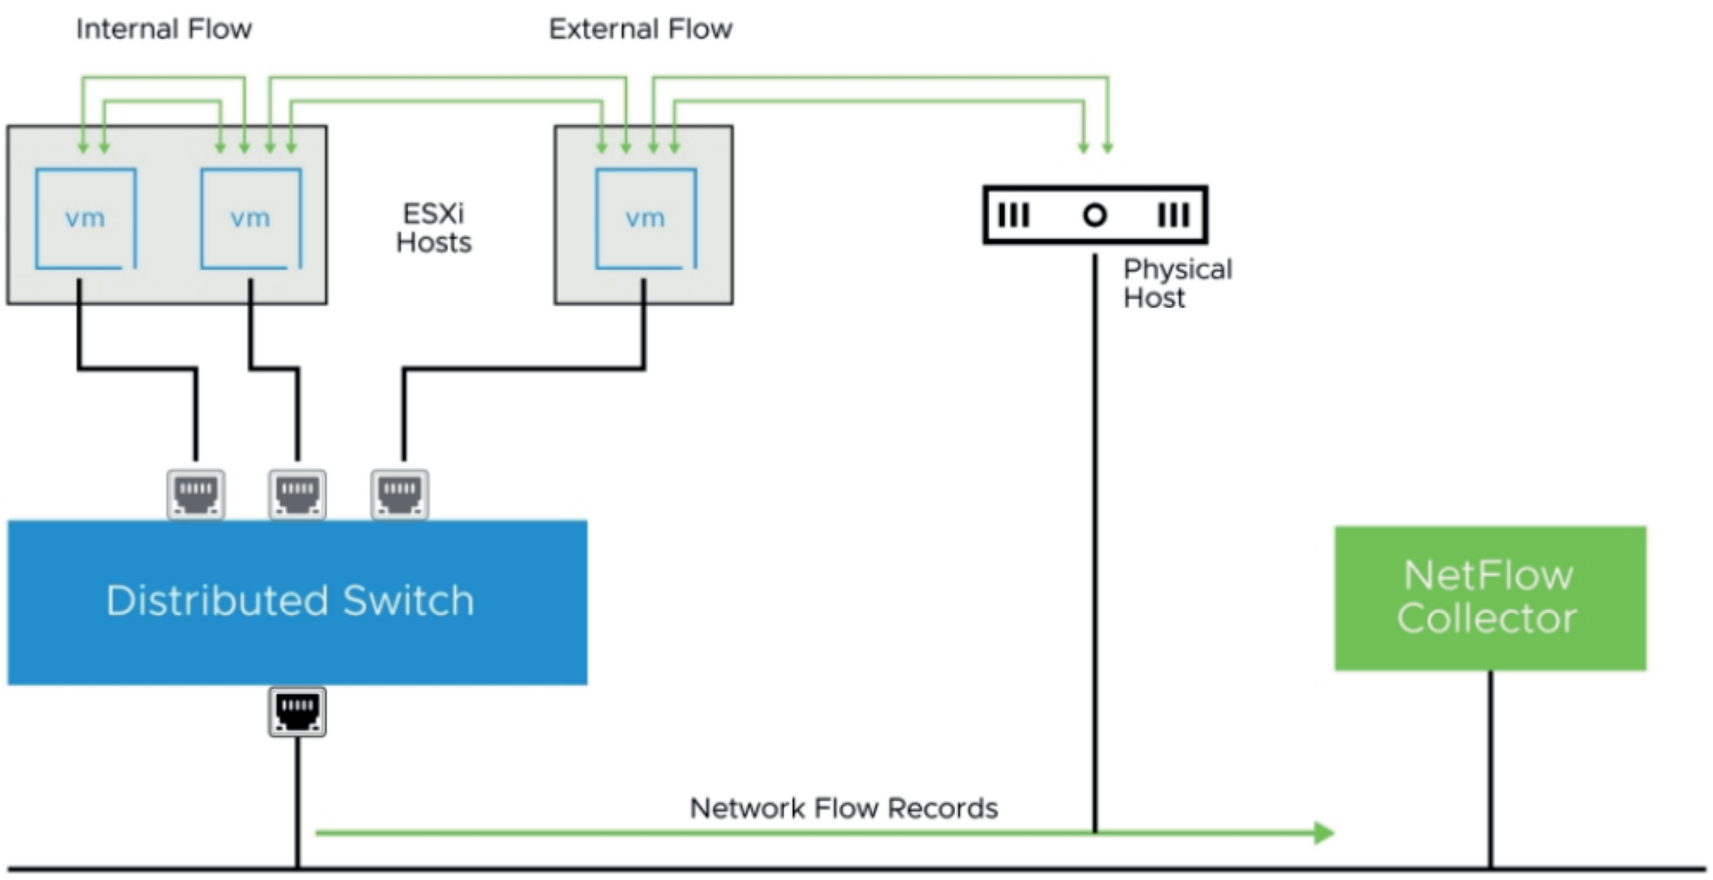

Network flows give you a complete view of VM traffic, which can be collected for historical views and used for multiple purposes.

Internal flows are generated from intrahost VM traffic, that is, traffic between VMs on the same hosts.

External flows are generated from traffic between VMs on different hosts or VMs on different distributed switches. External flows are also generated from traffic between physical machines and VMs.

A flow is a sequence of packets that share properties such as source and destination IP addresses, source and destination ports, input and output interface IDs, and protocol.

A flow is unidirectional. Flows are processed and stored as flow records by supported network devices, such as a distributed switch. The flow records are sent to a NetFlow collector for additional analysis.

Although flow processing is an efficient method, NetFlow can strain the distributed switch. NetFlow requires both additional processing and storage on the host for the flow records to be processed and exported.

4-50 Network Flow Analysis (1)

Network flow data is sent to a third-party NetFlow collector, which accepts and stores network flow records.

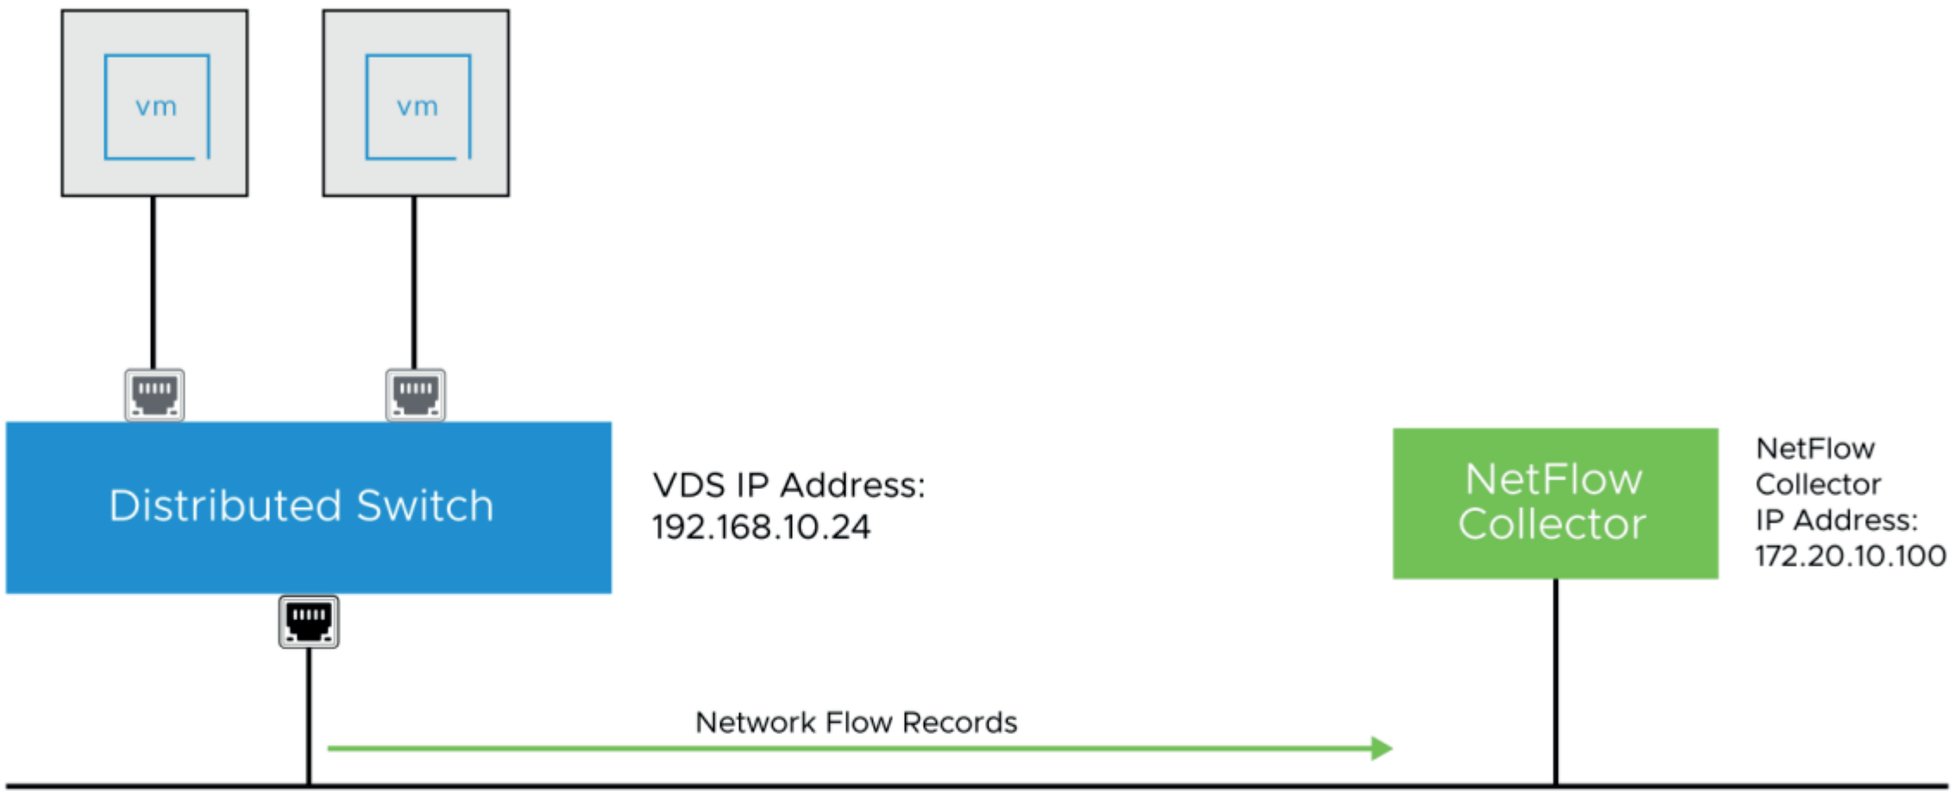

NetFlow sends aggregated network flow data to a NetFlow collector. Third-party vendors have NetFlow collector products.

A NetFlow collector accepts and stores the completed network flow records. NetFlow collectors vary in functionality by vendor. A NetFlow collector provides the following features:

- Analysis software to mine, aggregate, and report on the collected data
- A storage system for long-term storage so that you can archive the network flow data
- A customized user interface, often based on a web browser

The NetFlow collector can report on various kinds of networking information:

- The current top network flows consuming the most bandwidth in a particular (virtual) switch
- The IP addresses that are behaving irregularly
- The number of bytes that a particular VM sent and received in the past 24 hours

Technet24

4-51 Network Flow Analysis (2)

A NetFlow collector stores flow-based data and analyzes the collected data:

- Includes a database for storing the collected data:
  - You can investigate and isolate excessive network bandwidth use, bottlenecks, and unexpected application traffic.
  - You can view historical records to diagnose the cause of outages or breaches.
- Mines, aggregates, and reports on the collected data:
  - You can analyze network traffic by rate, volume, and utilization.
  - You can analyze trends in VM and host traffic.

With NetFlow data, you can investigate the causes of excessive use of network bandwidth, bottlenecks, and unexpected application traffic. The historical records that you put in long-term storage can help you diagnose what caused these outages or breaches.

Because NetFlow data comes from NetFlow-enabled network devices, additional network probes to collect the flow-based data are not required. NetFlow collectors and analyzers can provide a detailed set of network performance data. Given enough storage on the NetFlow collector, flow data can be archived for a long time, providing a long-term record of network behavior.

4-52 Configuring NetFlow on Distributed Switches

You configure NetFlow settings on a distributed switch.

To configure NetFlow on a distributed switch, specify the following NetFlow settings:

- Collector IP address and Collector port
- Observation Domain ID
- Switch IP address

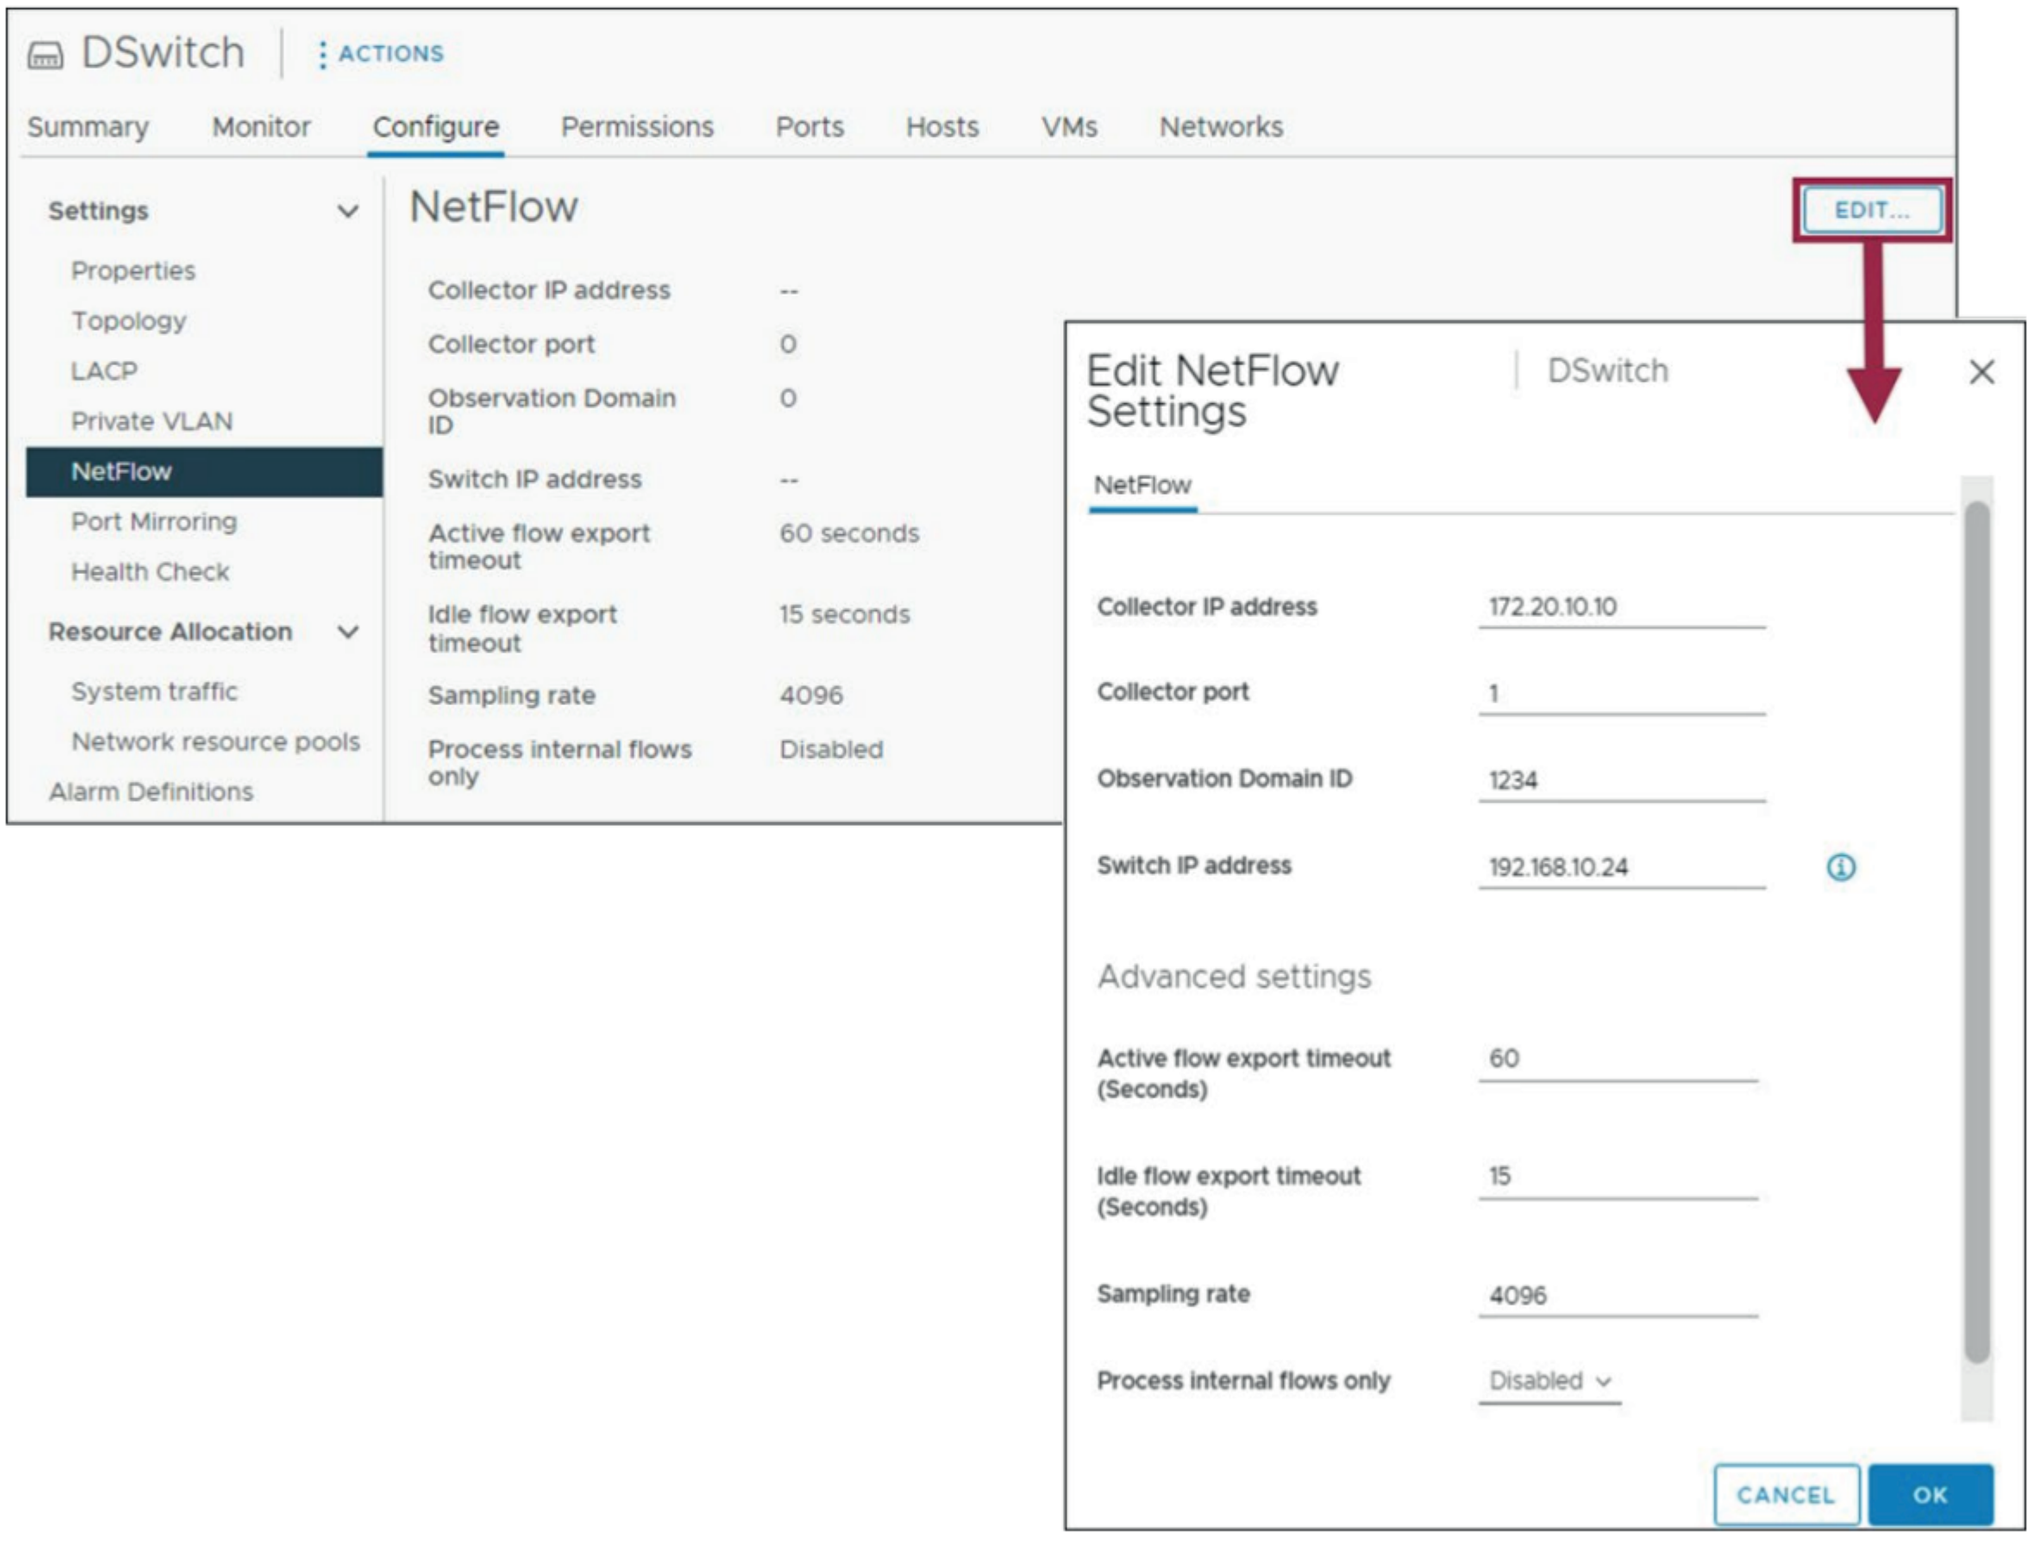

To configure NetFlow on a distributed switch, specify the following NetFlow settings:

- **Collector IP address** and **Collector port**: The IP address and port number that are used to communicate with the NetFlow collector system. The values are assigned by the collector.
- **Observation Domain ID**: A value that identifies the information related to the switch during the configuration of NetFlow Settings of a vSphere Distributed Switch.

Technet24

- **Switch IP address**: An optional IP address that is used to identify the source of the network flow to the NetFlow collector. The IP address is not associated with a network port, and it does not have to be pingable.

  This IP address is used to fill the source IP field of NetFlow packets. With the address, the NetFlow collector can interact with the distributed switch as a single switch, rather than as a separate, unrelated switch for each associated host. If this IP address is not configured, the host's management IP address is used instead.

For a description of the advanced NetFlow settings, see vSphere Networking at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-networking-guide.pdf.

4-53 Enabling NetFlow on a Distributed Port Group

After configuring NetFlow on the distributed switch, you either activate or deactivate NetFlow on a distributed port group, a specific port, or at the uplink.

4-54 About Port Mirroring

Port mirroring duplicates network packets from a source to a destination. It has the following uses:

- Helps with troubleshooting
- Provides input for network analysis appliances

Many network switch vendors implement port mirroring in their products.

vSphere port mirroring supports the following switch functionalities:

- Switched Port Analyzer (SPAN)
- Remote SPAN (RSPAN)
- Encapsulated RSPAN

With RSPAN, mirrored traffic can be directed to a remote monitor.

The RSPAN session can span multiple switch hops on a network.

With ERSPAN, the session can span an IP network.

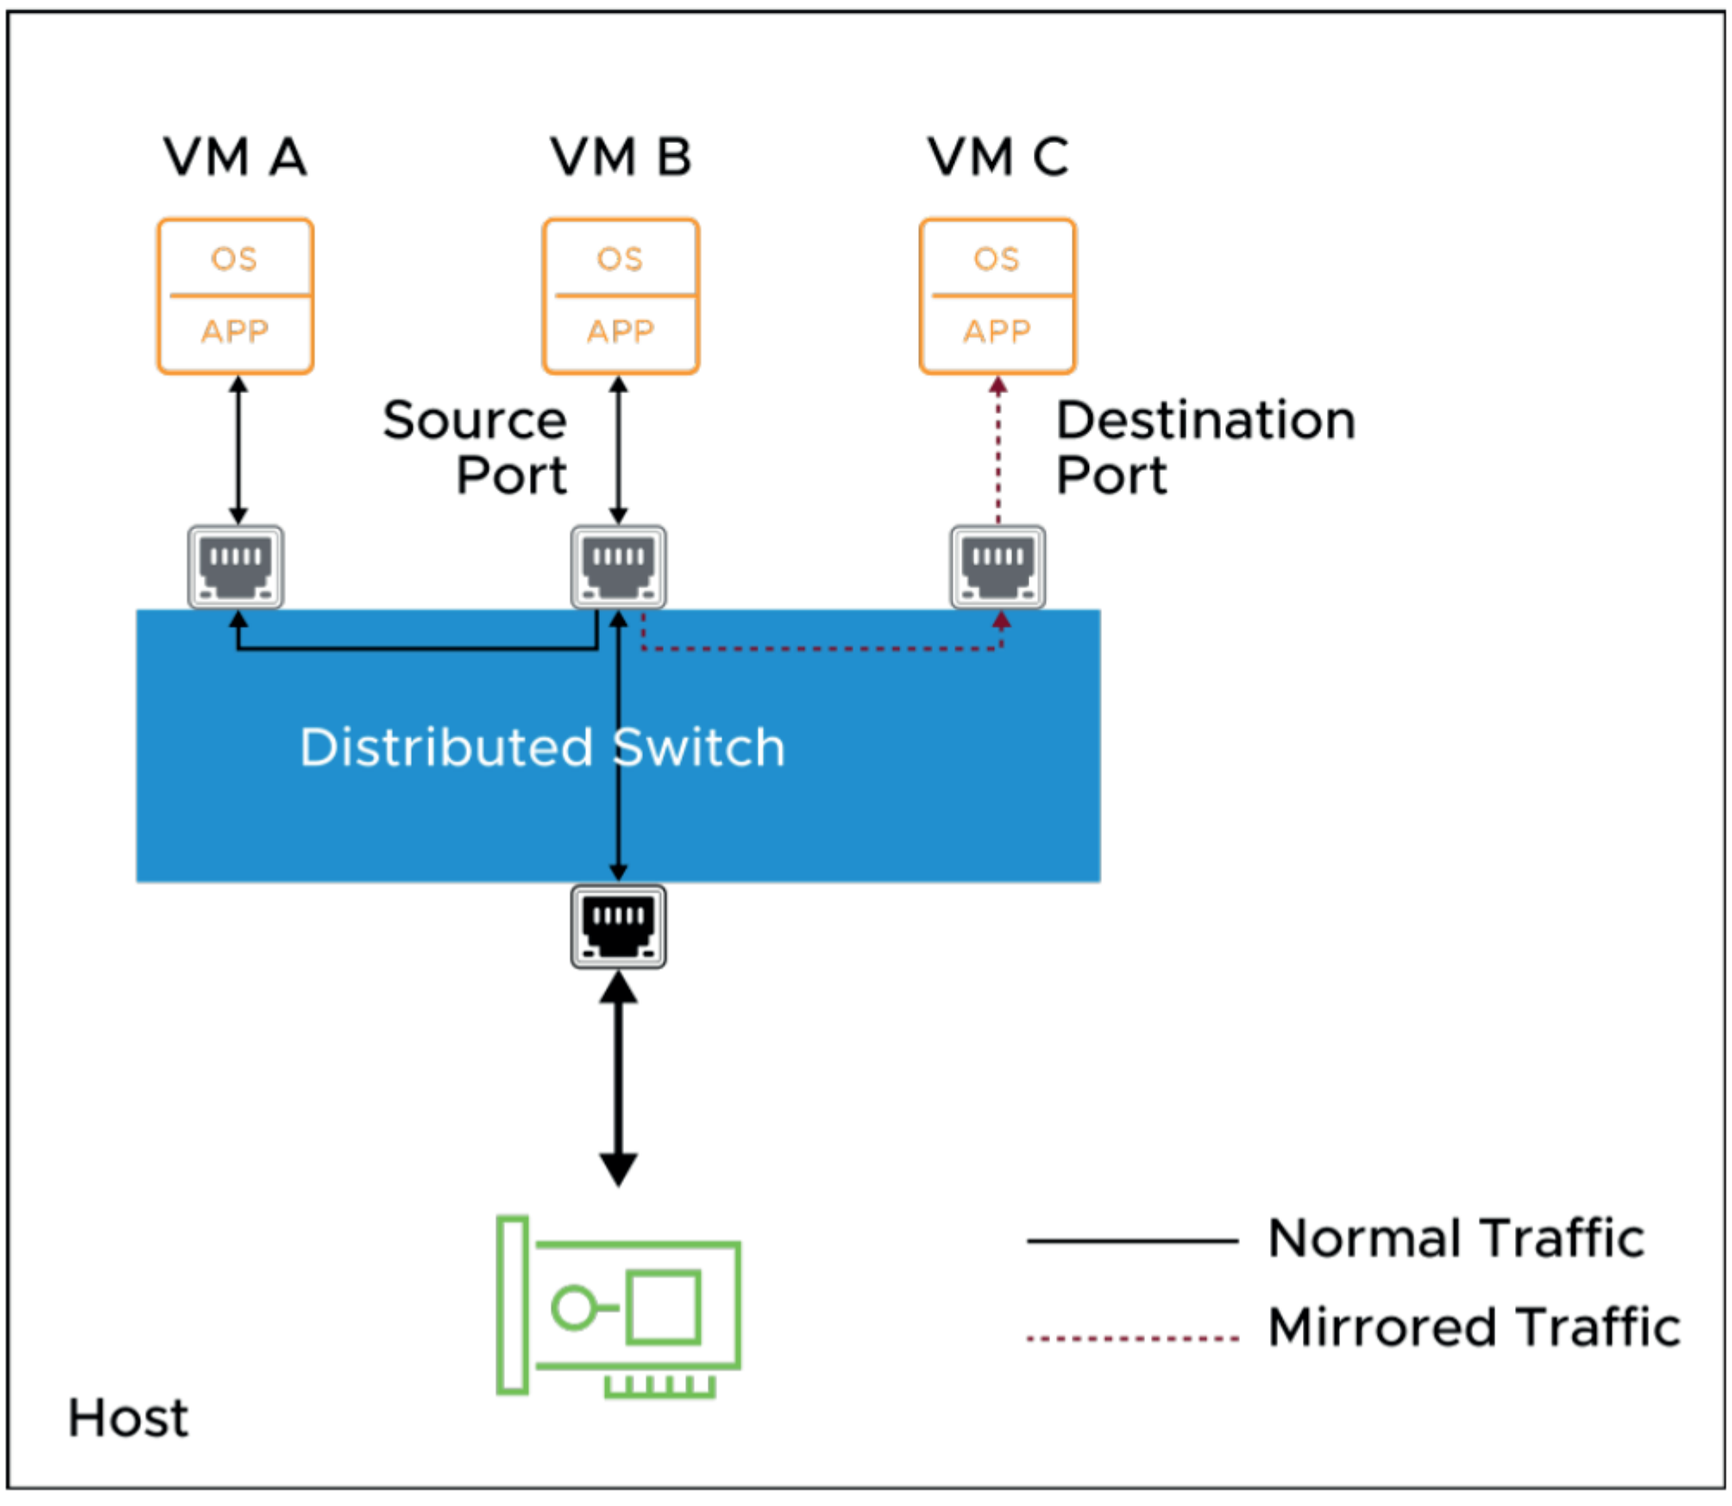

Technet24

On physical switches, administrators often must mirror traffic to special ports when troubleshooting network-related problems. Port mirroring is commonly used for network appliances that require the monitoring of network traffic, for example, intrusion detection systems.

Many network switch vendors implement port mirroring in their products. For example, port mirroring on a Cisco Systems switch is called Switched Port Analyzer (SPAN).

With port mirroring, you do not need to enable promiscuous mode on a distributed switch to troubleshoot network issues. If you enable promiscuous mode on a distributed port, the port sees all the network traffic passing through the distributed switch. You cannot select the port or port group that a promiscuous port is allowed to see. The promiscuous port sees all traffic that is on the broadcast domain.

Port mirroring supports Cisco Remote Switch Port Analyzer (RSPAN) and Encapsulated Remote Switch Port Analyzer (ERSPAN).

4-55 Port Mirroring Session Types (1)

You create a port mirroring session to mirror distributed switch traffic to ports, uplinks, and remote IP addresses.

You must select a port mirroring session type.

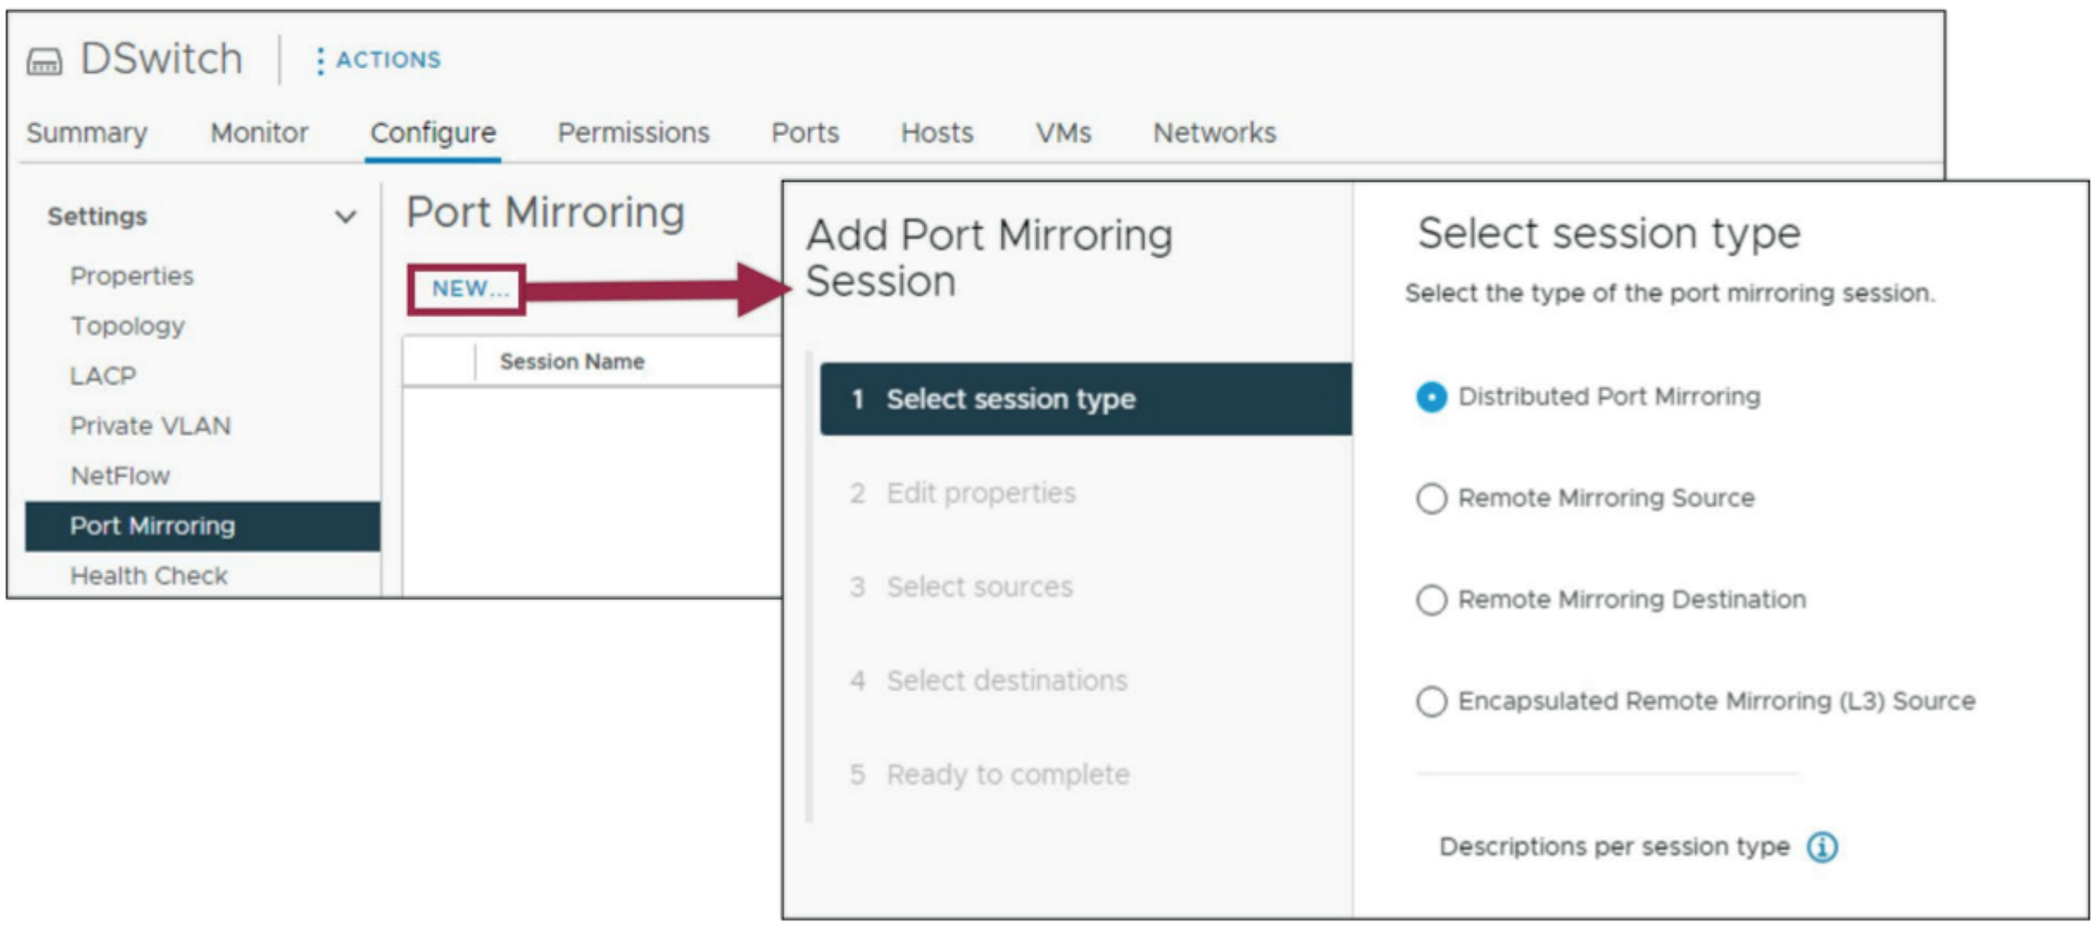

4-56 Port Mirroring Session Types (2)

You can select from the following port mirroring session options:

- **Distributed Port Mirroring**: Port mirroring between distributed ports can only be local. If the source and the destination ports are on different hosts because of vSphere vMotion migration, mirroring between them does not work. However, if the source and destination ports move to the same host, port mirroring works.
- **Remote Mirroring Source**: When a source-distributed port is moved from host A to host B, the original mirroring path from the source port to A's destination uplink is removed on A. A new mirroring path from the source port to B's destination uplink is created on B. The uplink that is used is determined by the uplink name specified in the session.
- **Remote Mirroring Destination**: When a destination-distributed port is moved from host A to host B, all original mirroring paths from source VLANs to the destination port are moved from A to B.
- **Encapsulated Remote Mirroring (L3) Source**: When a source-distributed port is moved from host A to host B, all original mirroring paths from the source port to the destination IP addresses are moved from A to B.

Distributed Port Mirroring - RSPAN

Remote Mirroring Source - RSPAN

Remote Mirroring Destination - RSPAN

Encapsulated Remote Mirroring (L3) Source - ERSPAN

Technet24

4-57 Port Mirroring Session Properties

You can configure one or more advanced properties based on the port mirroring session type that you select.

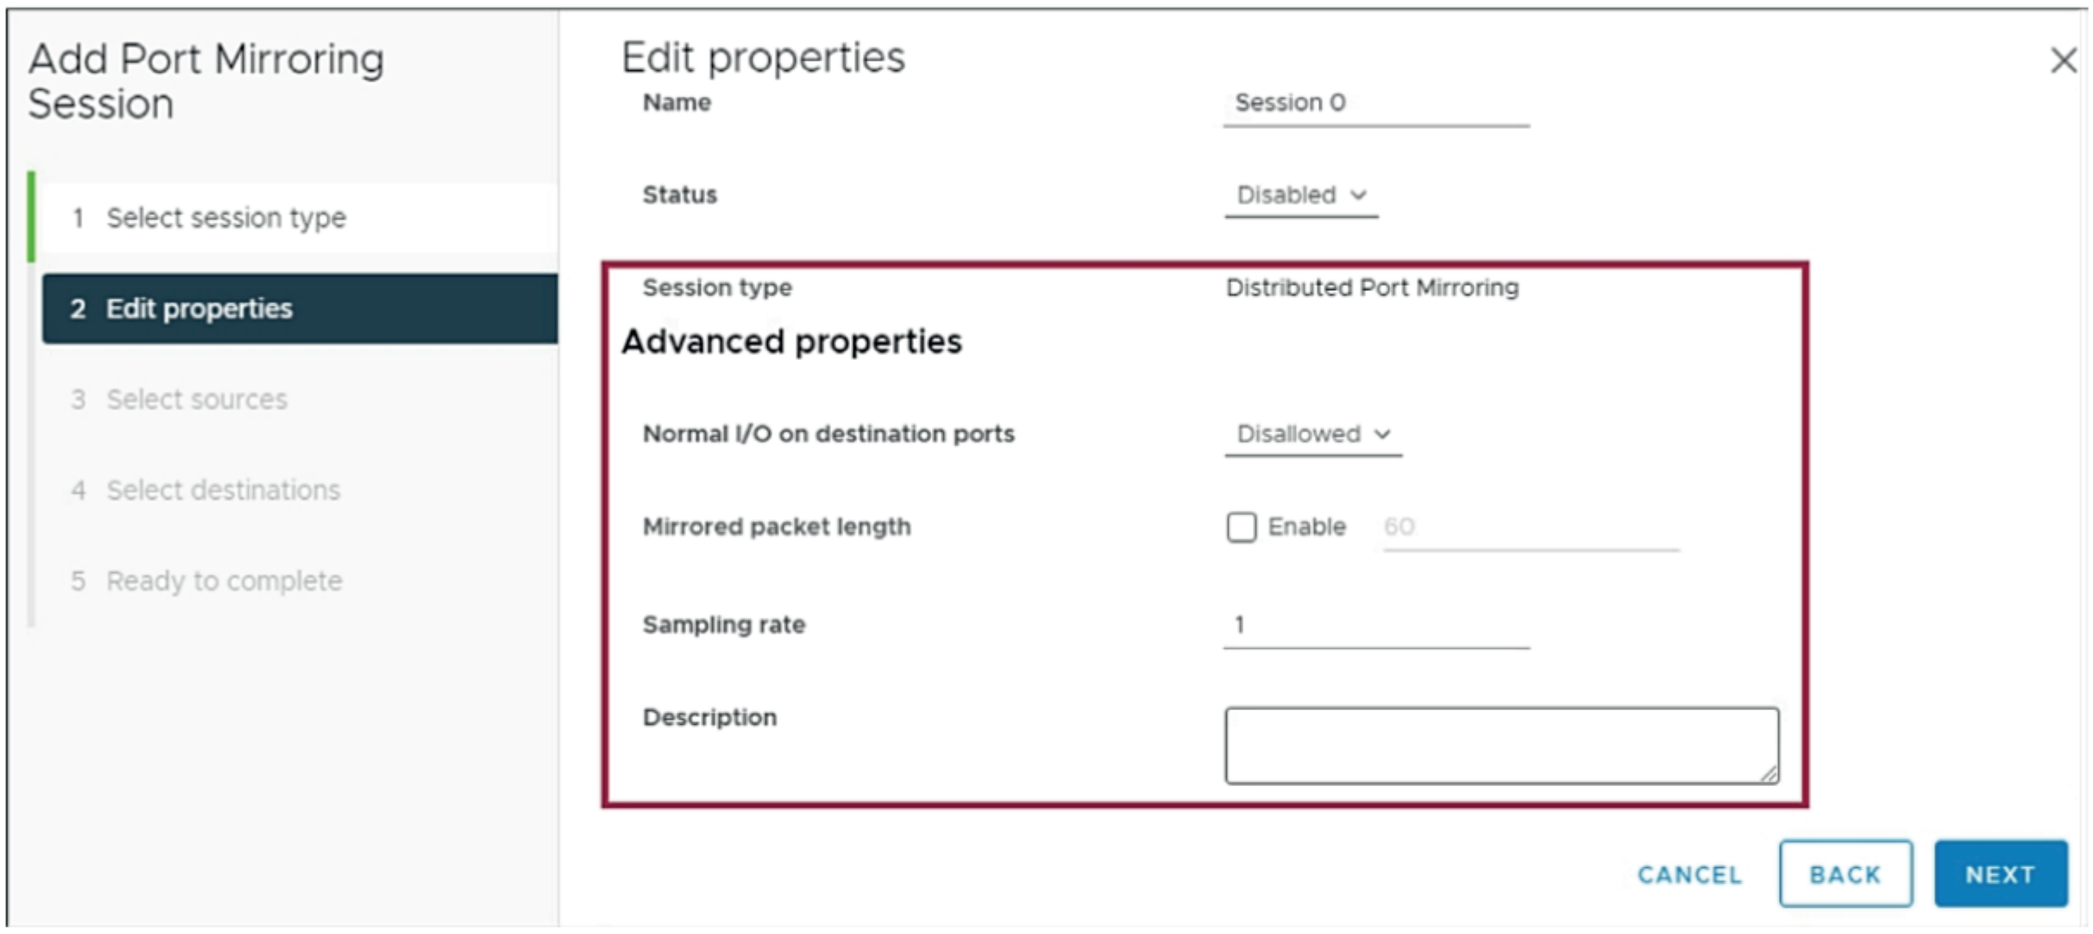

A distributed port mirroring type session can have the following properties:

- TCP/IP stack type (default or mirror)
- Normal I/O on destination ports (can be allowed or disallowed)
- Mirrored packet length (in bytes)
- Rate at which packets are sampled
- Description

Every port mirroring session is uniquely identified by its name. A session can also have a description.

For a description of the advanced session properties, see *vSphere Networking* at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-networking-guide.pdf.

4-58 Source and Destination in a Port Mirroring Session

A port mirroring session requires details about the source and the destination.

Source details include:

- Traffic direction
- One or more port IDs

Destination details include:

- One or more port IDs
- One or more uplinks
- IP address

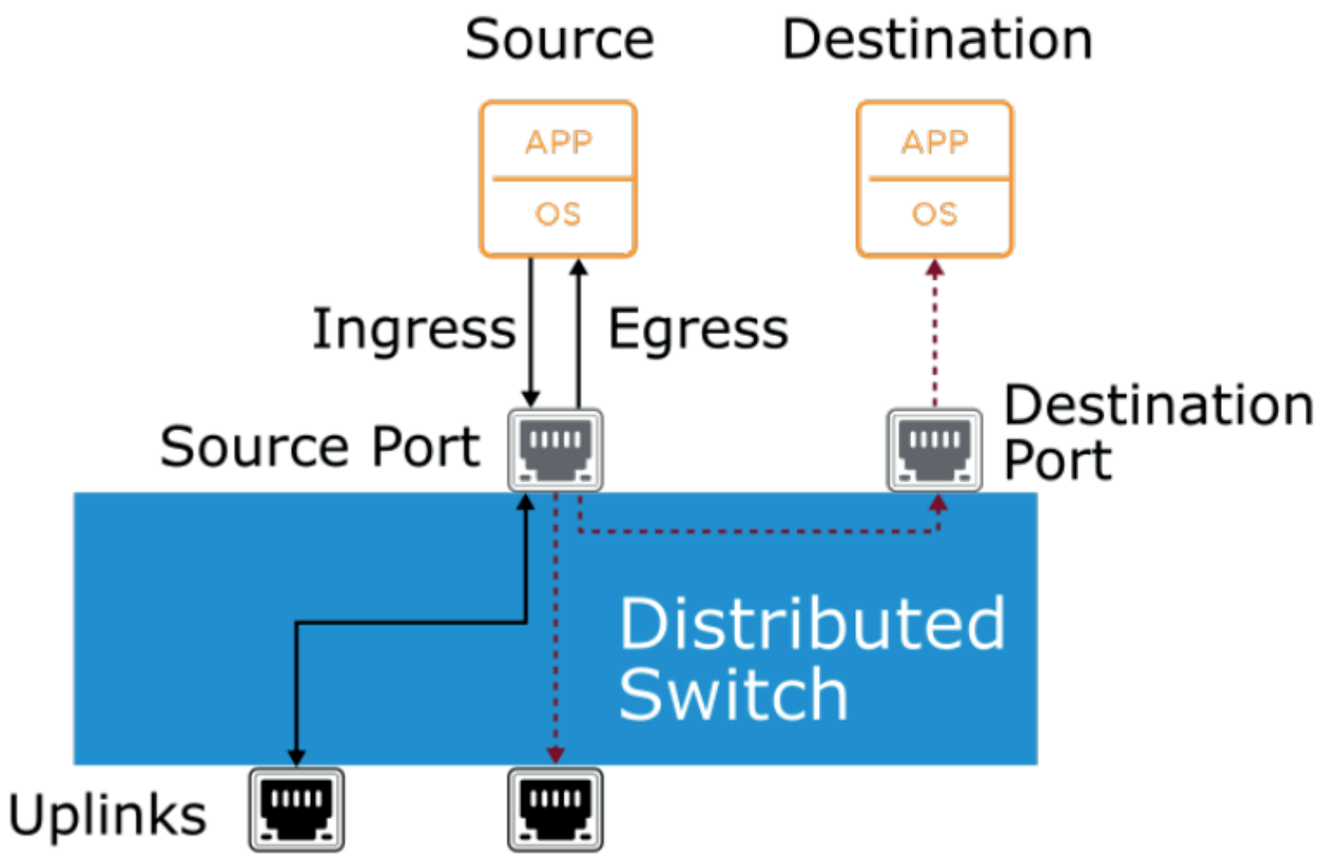

Normal Traffic

Mirrored Traffic

When creating a port mirroring session, you must configure the source and the destination. When configuring the source, you must specify the traffic direction.

Traffic direction is categorized as ingress or egress. Ingress traffic direction is traffic flowing from the source VM into the distributed switch. Egress traffic direction is traffic flowing from the distributed switch into the source VM.

Technet24

To avoid flooding the network with mirrored traffic, port mirroring has the following restrictions:

- In a session, a port cannot be a source and a destination.
- A port cannot be a destination for more than one session.
- A promiscuous port cannot be an egress source or destination.
- An egress source cannot be a destination for sessions, to avoid cycles of mirroring paths.

4-59 Lab 9: Using Port Mirroring

Configure port mirroring and capture network traffic on a distributed switch:

1. Prepare to Capture Mirrored Network Traffic
2. Configure Port Mirroring on the Distributed Switch
3. Verify That Port Mirroring Is Capturing Traffic
4. Restore the Distributed Switch Configuration

4-60 Review of Learner Objectives

- Describe how to use NetFlow for monitoring network security and performance
- Configure port mirroring on a distributed switch

4-61 Lesson 4: vSphere Distributed Services Engine

4-62 Learner Objectives

- Define vSphere Distributed Services Engine
- Describe the use cases and benefits of vSphere Distributed Services Engine
- Explain the architectural components of vSphere Distributed Services Engine

4-63 SmartNICs

A SmartNIC is a NIC with a general-purpose CPU, out-of-band management, and virtualized device functionality.

- **General-purpose CPU:** Having a general-purpose CPU allows one to run arbitrary code and applications directly on the NIC, such as networking and storage services, which both improves performance (because of fast access to the network I/O path) and saves core CPU cycles.
- **Out-of-band management:** The CPU complex on the SmartNIC can be managed independently from the server's CPU.
- **Virtualized device functionality:** SmartNICs can expose "virtual" devices on the PCI bus that appear to the core CPU OS and apps as if they are actual hardware devices.

SmartNICs are also referred as Data Processing Units (DPU).

Technet24

4-64 Traditional Infrastructure Challenges

Some of the challenges in today's environments are:

- Meeting the modern apps new requirements
- Shared Hypervisor CPU resources between apps and infrastructure (network, storage and hypervisor management tasks)
- Network and security services performance impact on applications
- Inadequate isolation between provider and tenant in multitenant environments based on bare metal clouds

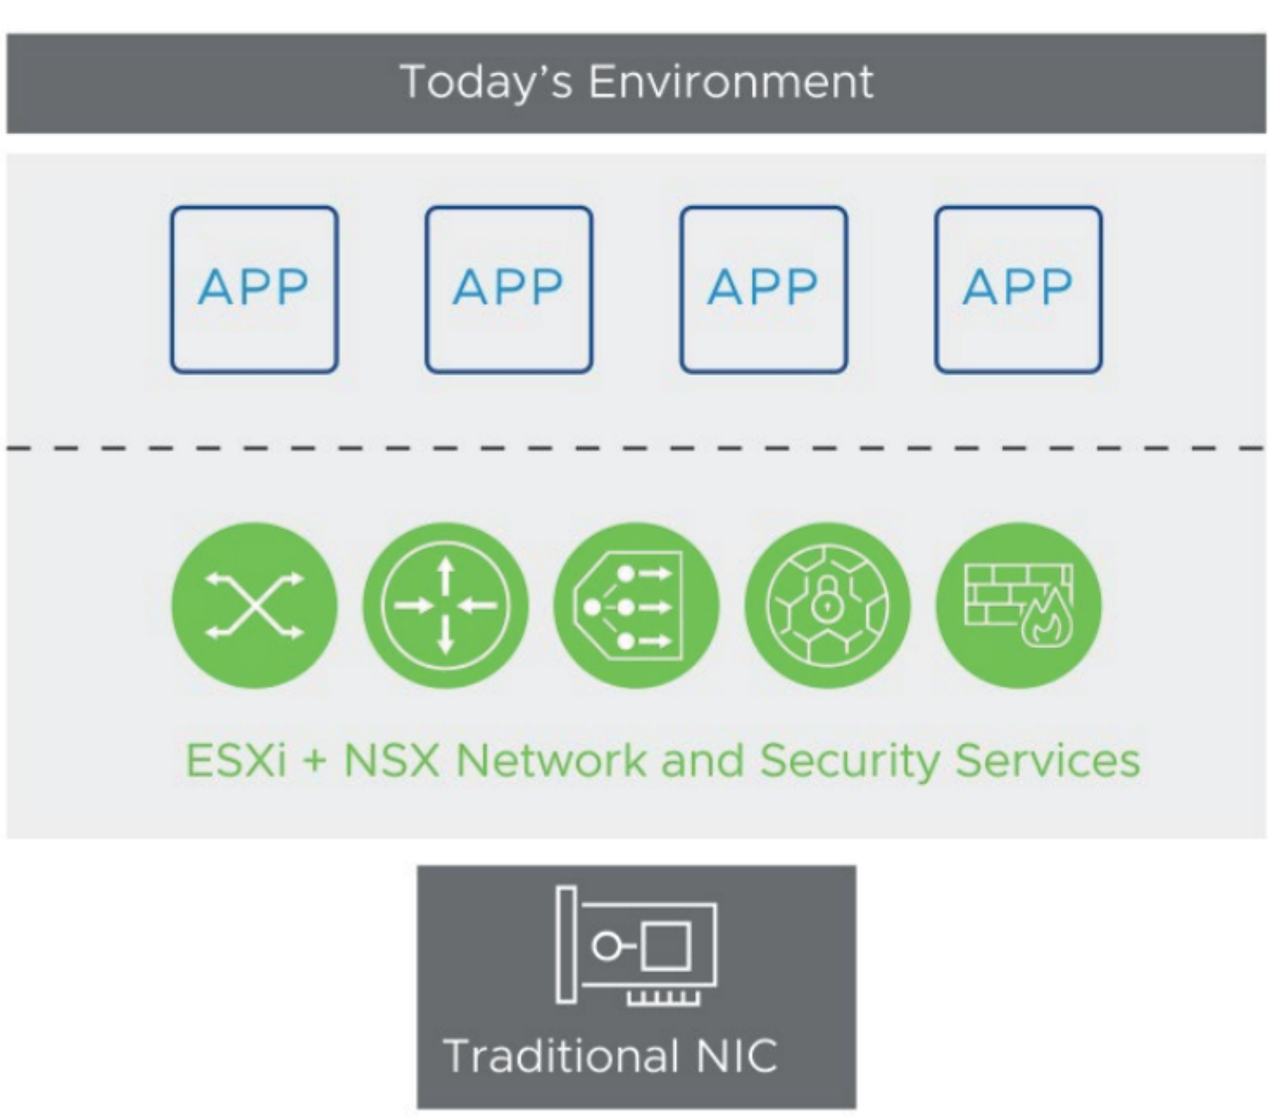

- Modern apps require: Speed of delivery, Application Availability, Security of applications and better performance
- Hypervisor CPU resources not dedicated for modern apps. This creates CPU and I/O bottlenecks.
- Bare Metal clouds are seeing interest among customers with demands for high performance workloads, direct hardware access and custom multi-tenant environments but they are lacking the isolation between the provider network and tenant network typical of SDDC environments.

4-65 About vSphere Distributed Services Engine

In vSphere 8, vSphere Distributed Services Engine unlocks the power of SmartNICs/Data Processing Units (DPUs) for hardware accelerated data processing to boost infrastructure performance.

vSphere Distributed Services Engine will leverage the SmartNIC/DPU to:

- **Boost** infrastructure performance
- **Strengthen** infrastructure security with an enhanced zero-trust model
- **Simplify** infrastructure and workload management

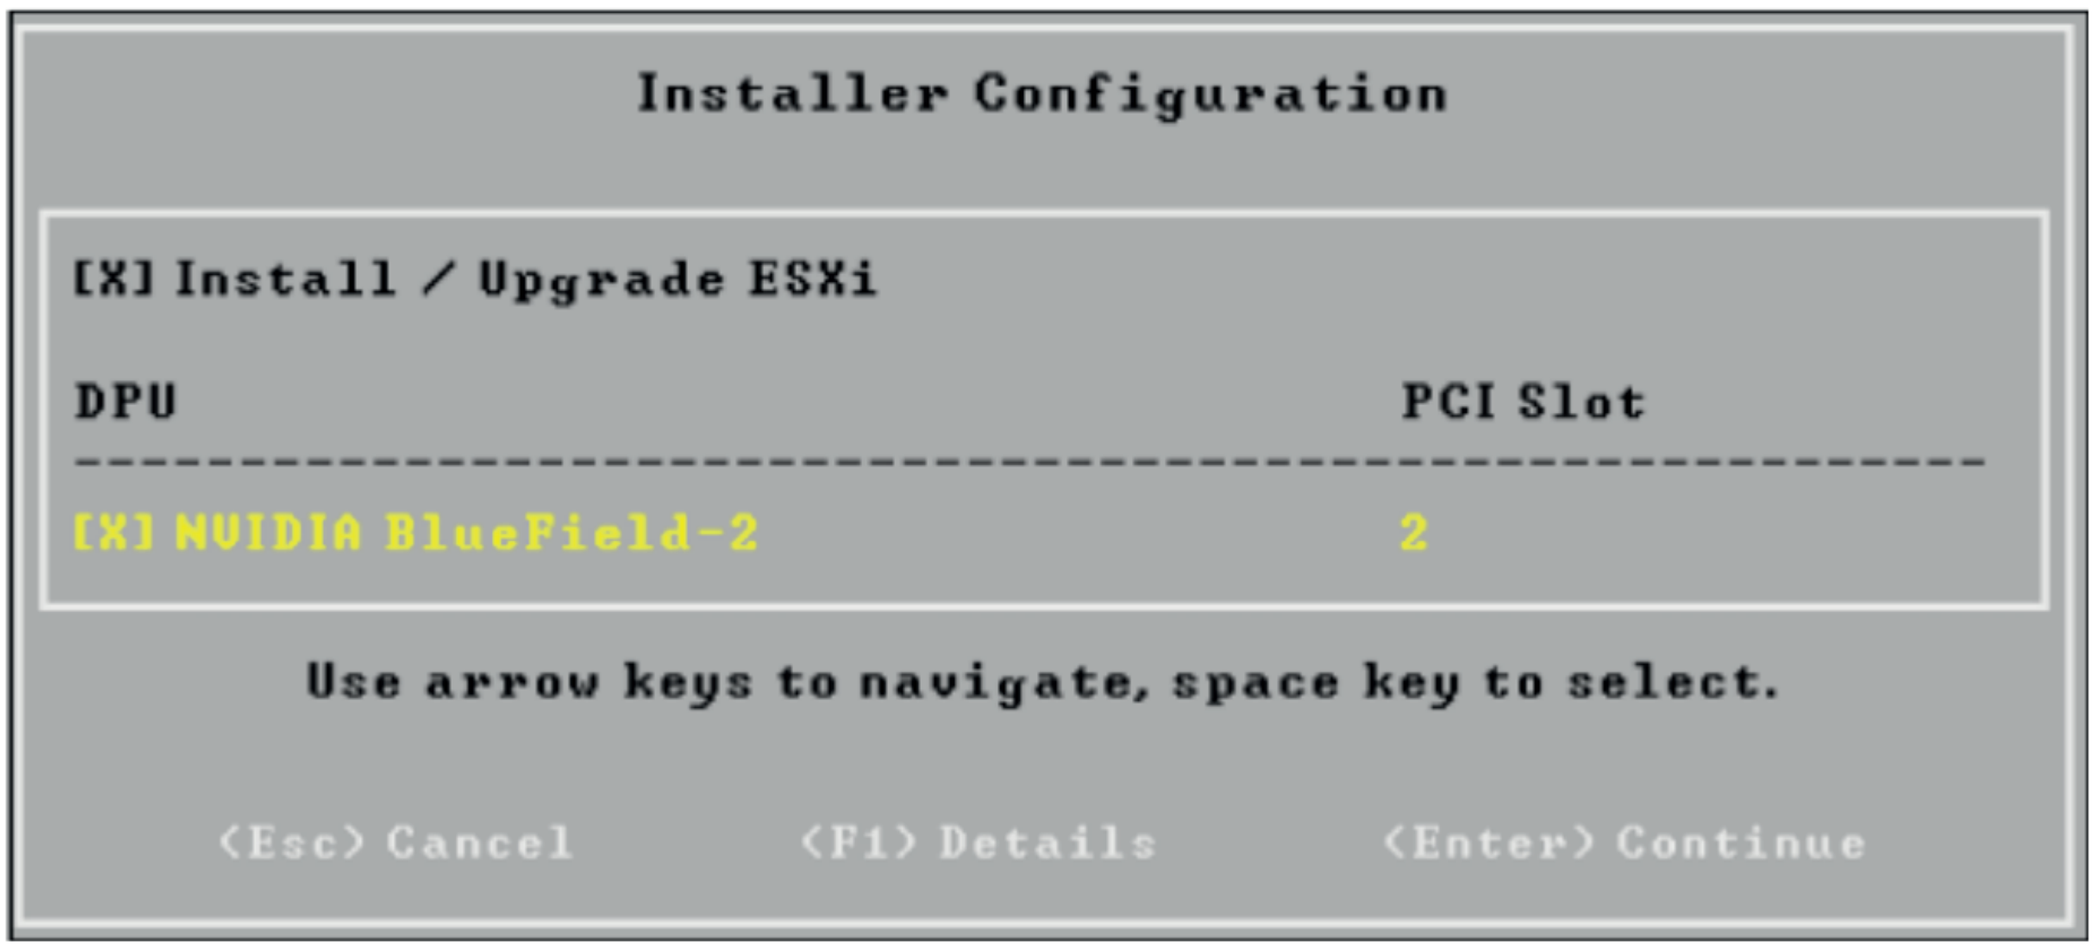

Technet24

4-66 Next Generation Infrastructure with vSphere Distributed Services Engine

SmartNICs offer the following advantages:

- Dedicated CPU resources and HW acceleration.
- Full datapath offloading to achieve very high throughput and low latency.
- Security services can be enabled in the SmartNIC/DPU without performance impact
- Enhanced observability and operations capabilities for monitoring, troubleshooting, logging and compliance.
- Isolation between tenant and provider both on ESXi and bare metal platforms (available in future releases)

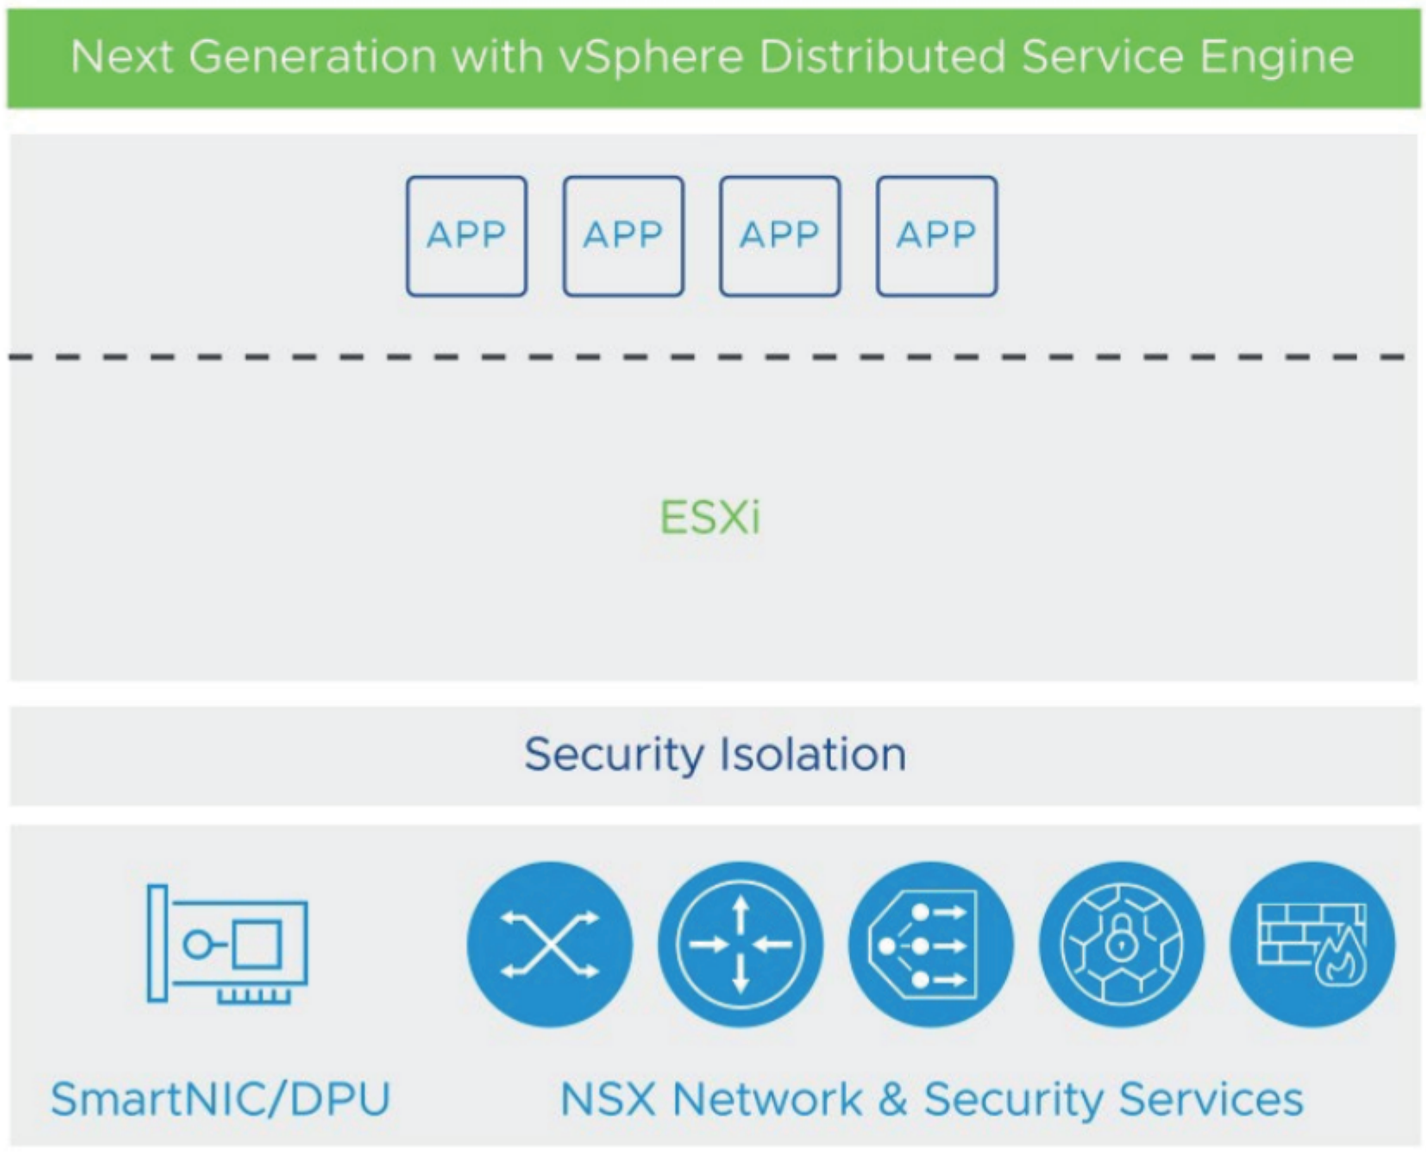

Next Generation architecture runs NSX services directly in the SmartNIC.

4-67 vSphere Distributed Services Engine Benefits

These are the main use cases of vSphere Distributed Services Engine:

- Applications with high Network bandwidth demands
- Security services offloading for a better performance
- Enhanced observability requirements

- Networking offloads for high bandwidth demanding applications like video streaming apps.
- Security offloads for a better performance for applications with high security/encryption demands

4-68 vSphere Distributed Services Engine Ecosystem

Leading system OEMs like Dell Technologies, HPE or Lenovo include vSphere Distributed Services Engine in their integrated platforms.

Nvidia

AMD/Pensando

Technet24

4-69 Architecture Changes with vSphere Distributed Services Engine

Key changes with the vSphere Distributed Services Engine architecture are:

- NSX and Infrastructure services like storage and I/O control run on the vSphere Distributed Services Engine processor
- New modules are required to run ESXi and NSX functionalities on the new architecture with vSphere Distributed Services Engine:
  - ESXiO
  - NSXiO

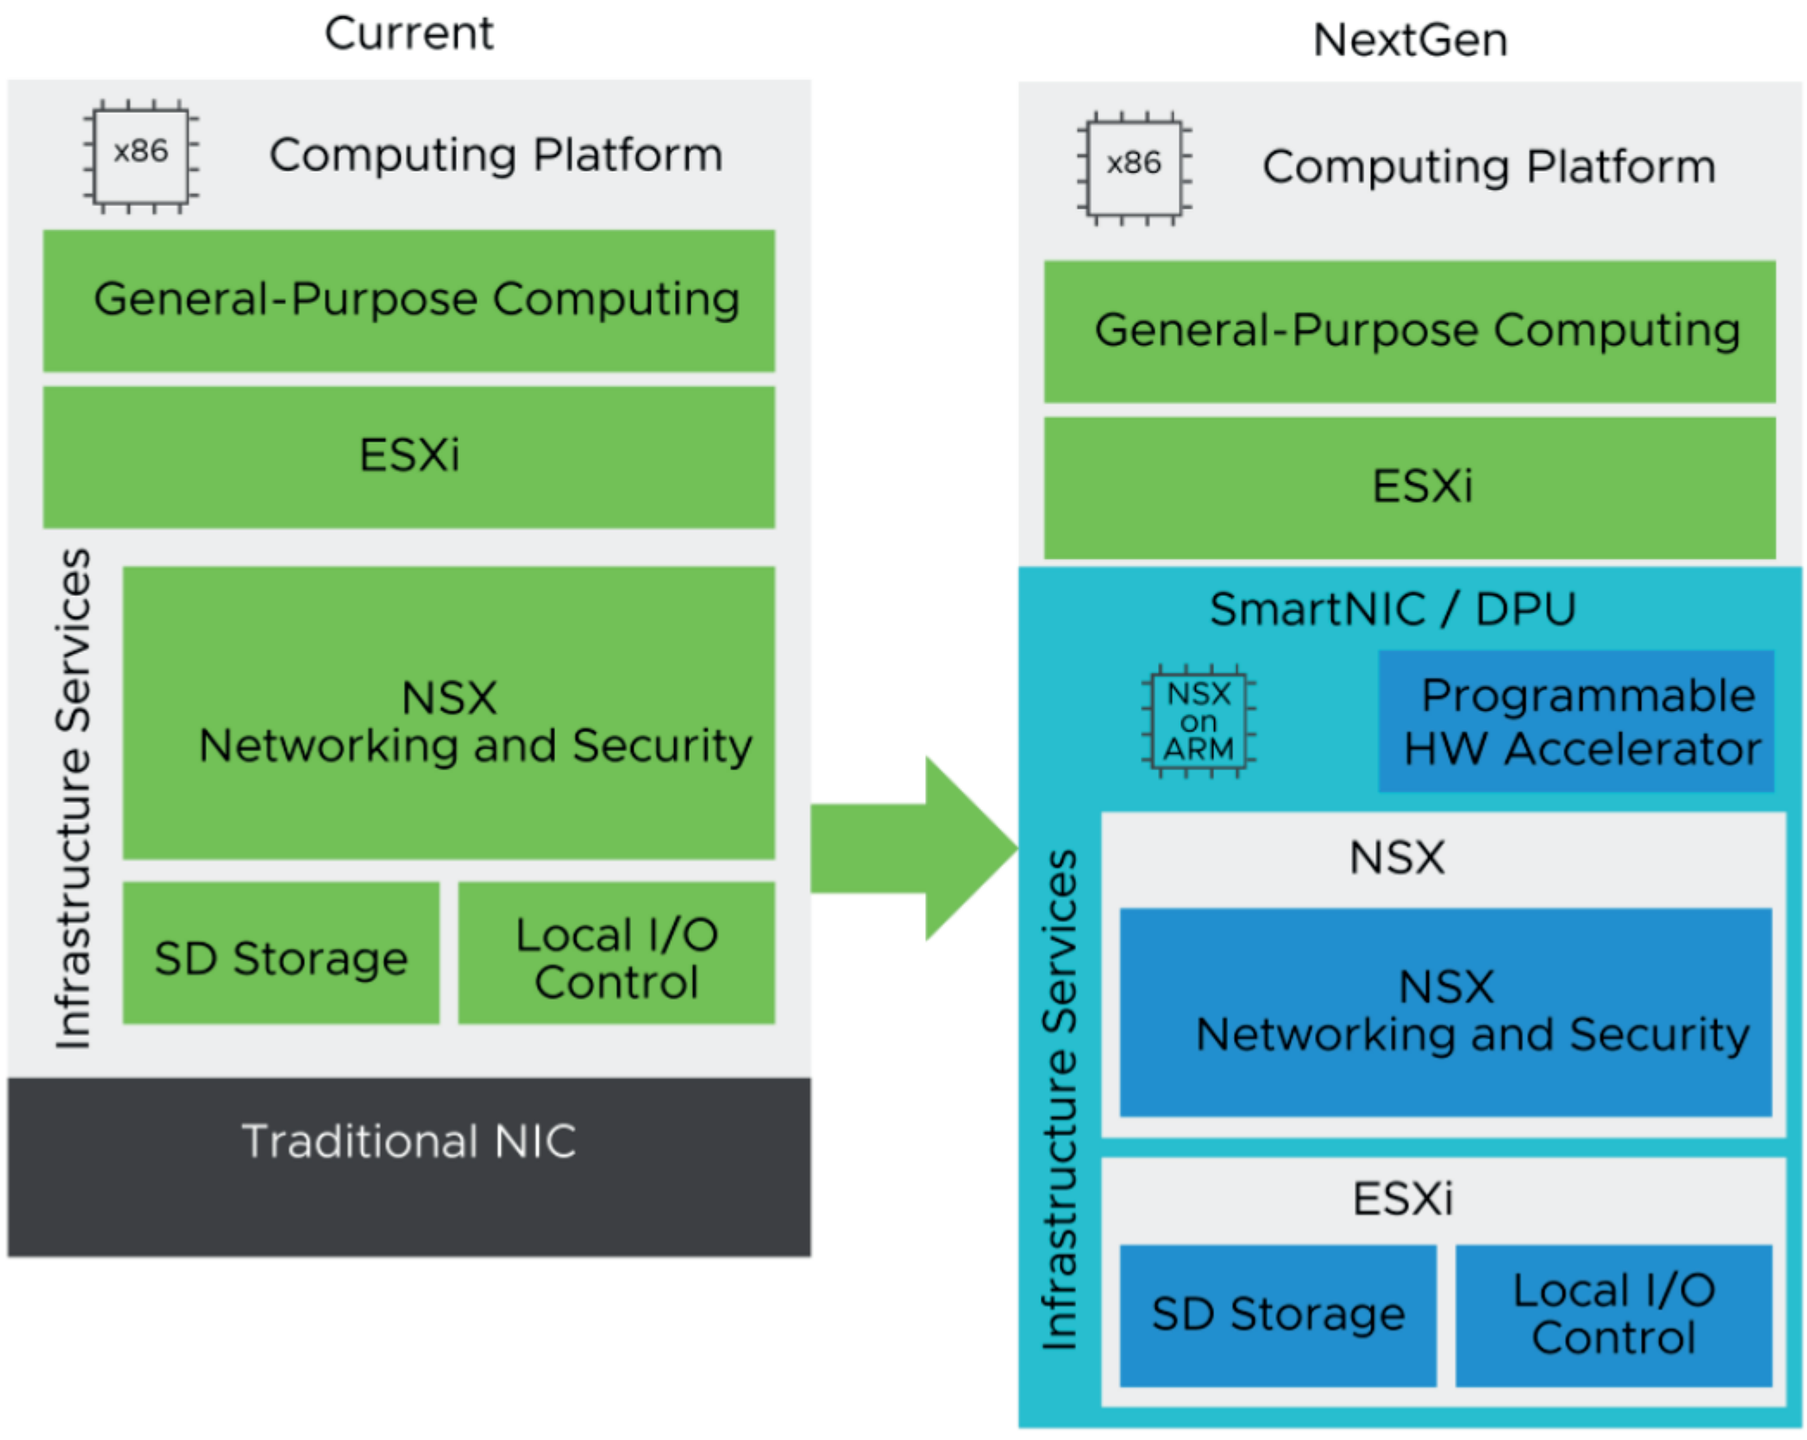

The new architecture runs the infrastructure services on the SmartNIC providing the necessary separation between application workloads that run on the x86 computing platform and the infrastructure services.

SmartNIC hardware is independent from the kernel and has programable pipeline with dedicated CPUs. This allows network and storage virtualization to run directly on the NIC saving x86 CPU cycles and improving performance via offload.

These modules are required to run ESXi and NSX functionalities on the new architecture with smartNICs:

- ESXio: Lightweight ESXi component running on the ARM processor in a SmartNIC
- NSXio: NSX components (transport node components) that are running on the SmartNIC.

4-70 vSphere Distributed Services Engine Hardware Components

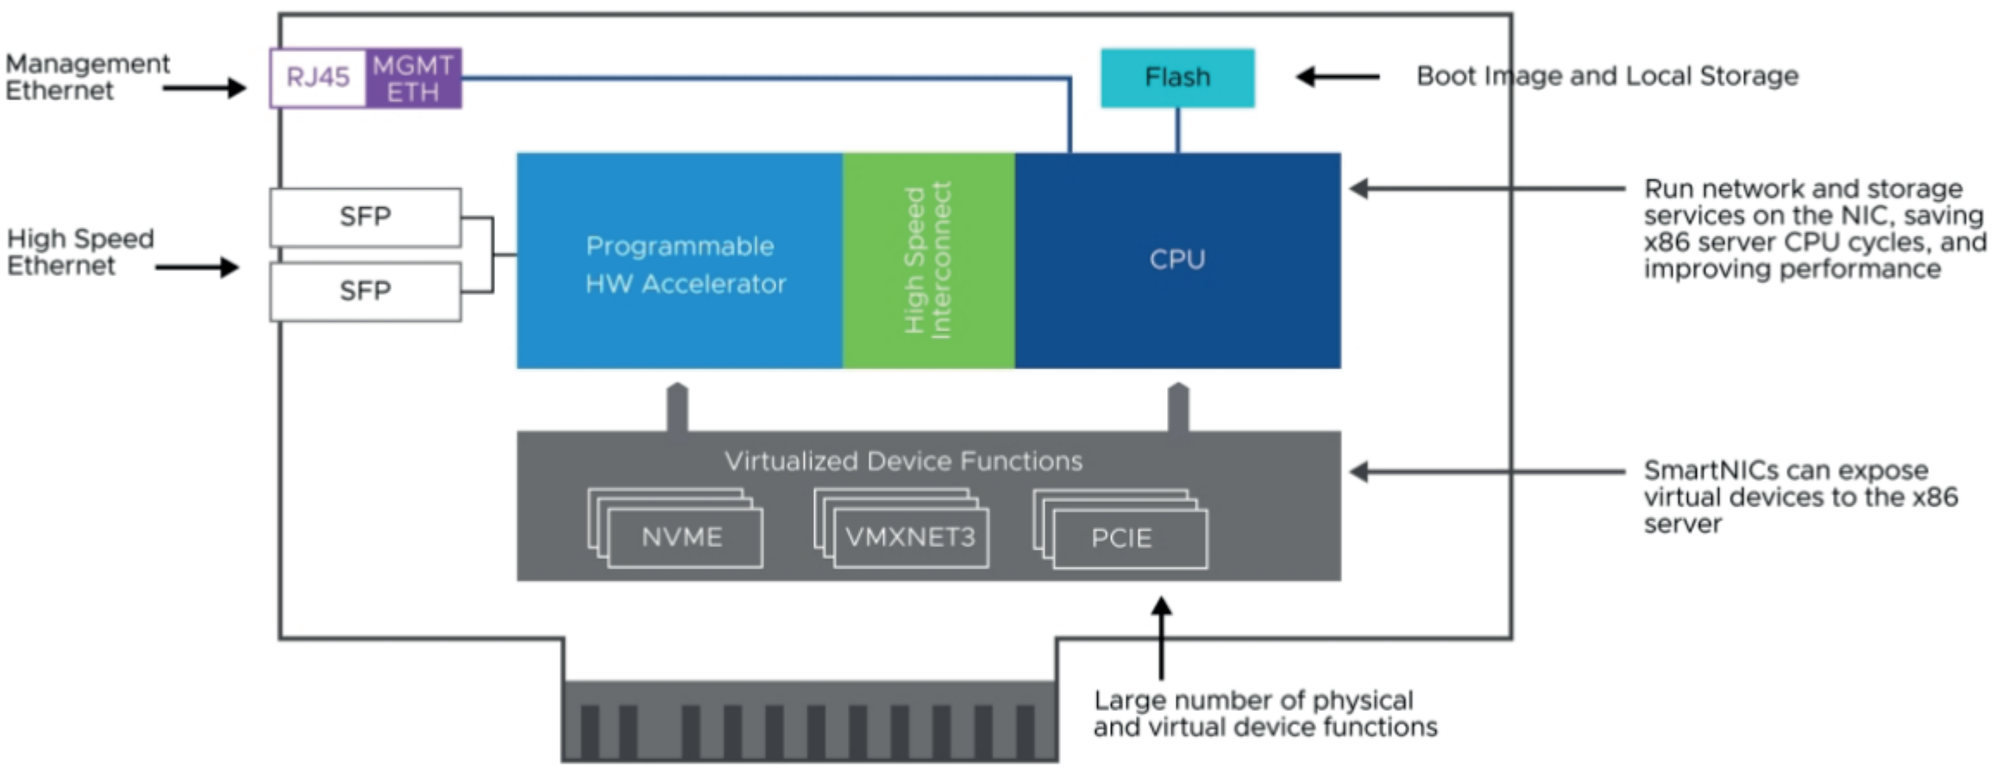

vSphere Distributed Services Engine components:

- Flash memory to store the boot image and local storage
- Ethernet port for management purposes
- Two or more SFP ports for High Speed Ethernet.
- Programmable hardware accelerator packet processing pipeline for datapath offloading.
- Embedded ARM processors with its own memory to save server x86 CPU cycles and improve network performance.
- Virtualized Device Functions to expose virtual device functions to the ESXi host. Large number of physical and virtual devices are included (NVMe, VMXNET3 and PCIe)

Technet24

4-71 Review of Learner Objectives

- Define vSphere Distributed Services Engine
- Describe the use cases and benefits of vSphere Distributed Services Engine
- Explain the architectural components of vSphere Distributed Services Engine

4-72 Key Points

- A distributed switch provides functions that are similar to a standard switch. But the distributed switch defines a single configuration that is shared across all associated hosts.
- Network I/O Control allocates bandwidth to each type of system traffic by using shares, reservations, and limits.
- Distributed switch and distributed port group configurations can be backed up and restored.
- Distributed switches support the use of network analysis and troubleshooting tools, specifically NetFlow and port mirroring.
- vSphere Distributed Services Engine unlocks the power of SmartNICs/Data Processing Units (DPUs) for hardware accelerated data processing to boost infrastructure performance.

Questions?

Module 5

Storage Operations

5-2 Importance

The scalability features and software-defined storage configurations in vSphere can help you handle growth efficiently in your environment. For datastore growth and balancing issues, you can use vSphere Storage DRS. With the knowledge and skills in storage policy based management, you can align storage with the application demands of your VMs.

5-3 Module Lessons

1. NVMe Storage Systems
2. iSCSI Extension for RDMA (iSER)
3. VMFS Datastores
4. vSAN Datastores
5. vSphere Storage APIs - Array Integration (VAAI)
6. vSphere Storage API for Storage Awareness and vSphere APIs for I/O Filtering (VASA and VAIO)
7. Storage Policy Based Management Overview
8. Using VM Storage Policies
9. Storage Policies for vSAN
10. Storage Policies for vSphere Virtual Volumes
11. Storage I/O Control
12. Datastore Clusters and vSphere Storage DRS

Technet24

5-4 Lesson 1: NVMe Storage Systems

5-5 Learner Objectives

- Recognize the support for NVMe
- Configure ESXi for NVMe over RDMA (RoCE v2)
- Describe the VMware Pluggable Storage Architecture
- Describe the purpose of the High-Performance Plug-in

5-6 vSphere Support for NVMe

Non-Volatile Memory Express (NVMe) is a standardized protocol for high-performance multi-queue communication with NVM devices. ESXi supports the NVMe protocol to connect to local and networked storage devices.

Introduced in vSphere 7 the following features were added to enhance NVMe support

- Pluggable Storage Architecture (PSA) stack using a High-Performance Plug-In (HPP) with multipathing
- Hot-add and hot-remove of local NVM devices, if the server supports it
- Support for networked NVMe devices connected via Fibre Channel or RDMA

ESXi uses Pluggable Storage Architecture (PSA) as a VMkernel layer for managing storage multipathing.

The PSA is an open and modular framework. It coordinates software modules that are responsible for multipathing operations.

Modules include generic multipathing modules (MPPs) that VMware provides, such as Native Multipathing Plug-In (NMP) and High Performance Plug-in (HPP), plus third-party MPPs as well.

5-7 Supported NVMe Topologies

The NVMe storage can be directly attached to a host indirectly through different fabric transports, which is known as Shared NVMe over Fabrics, or simply NVMe-oF.

NVMe Transport	ESXi 7 and Later Support
NVMe over PCIe	Local storage
NVMe over RDMA	Shared NVMe-oF storage
NVMe over Fibre Channel (FC-NVMe)	Shared NVMe-oF storage
NVMe over TCP	Shared NVMe-oF storage

5-8 FC-NVMe and NVMe over RDMA (RoCE v2)

vSphere 7 and later supports remote access to NVMe-oF. This includes both FC-NVMe and NVMe over RDMA topologies.

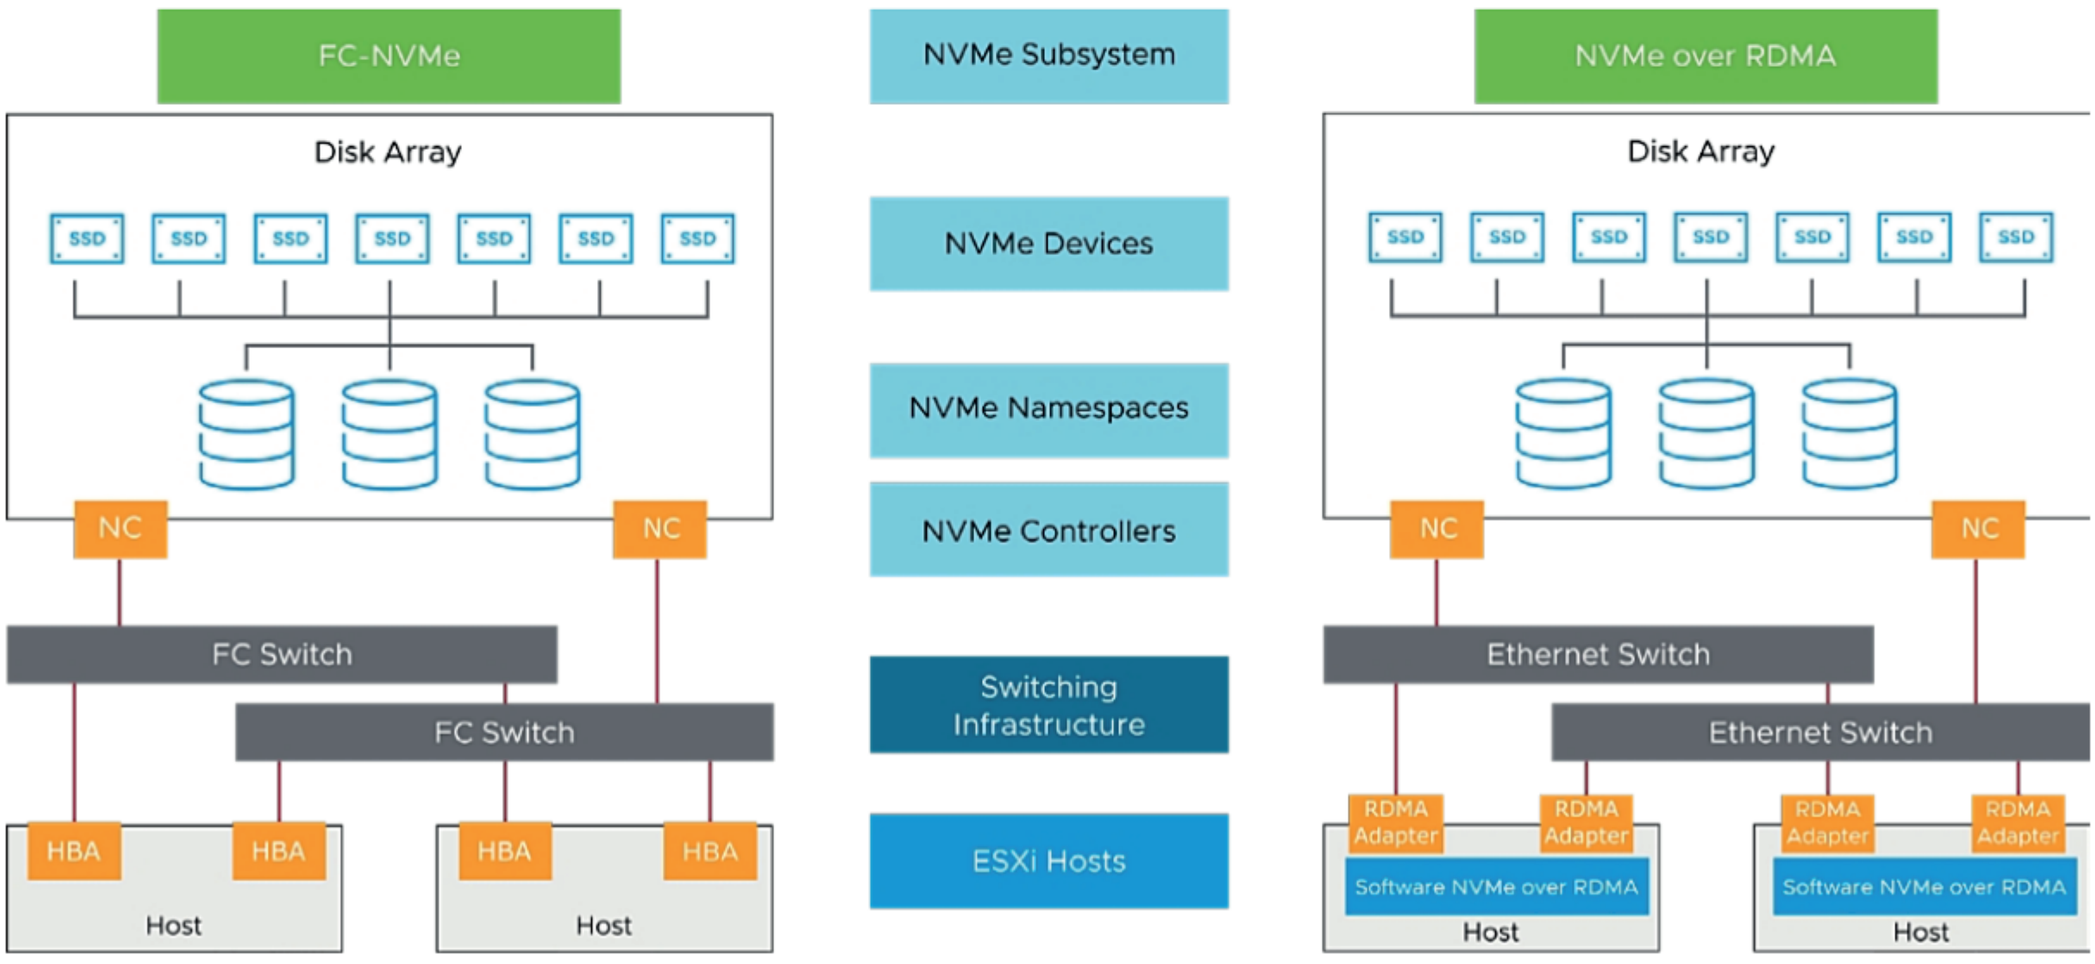

Technet24

5-9 Support for FC-NVMe

FC-NVMe is a technology that can map NVMe onto the Fibre Channel protocol.

With FC-NVMe, data and commands can be transferred between a host computer and a target storage device.

To access the FC-NVMe storage, you must install a Fibre Channel storage adapter that supports NVMe on your ESXi host.

You do not need to configure the controller. After you install the required hardware NVMe adapter, it automatically connects to all targets and controllers that are reachable.

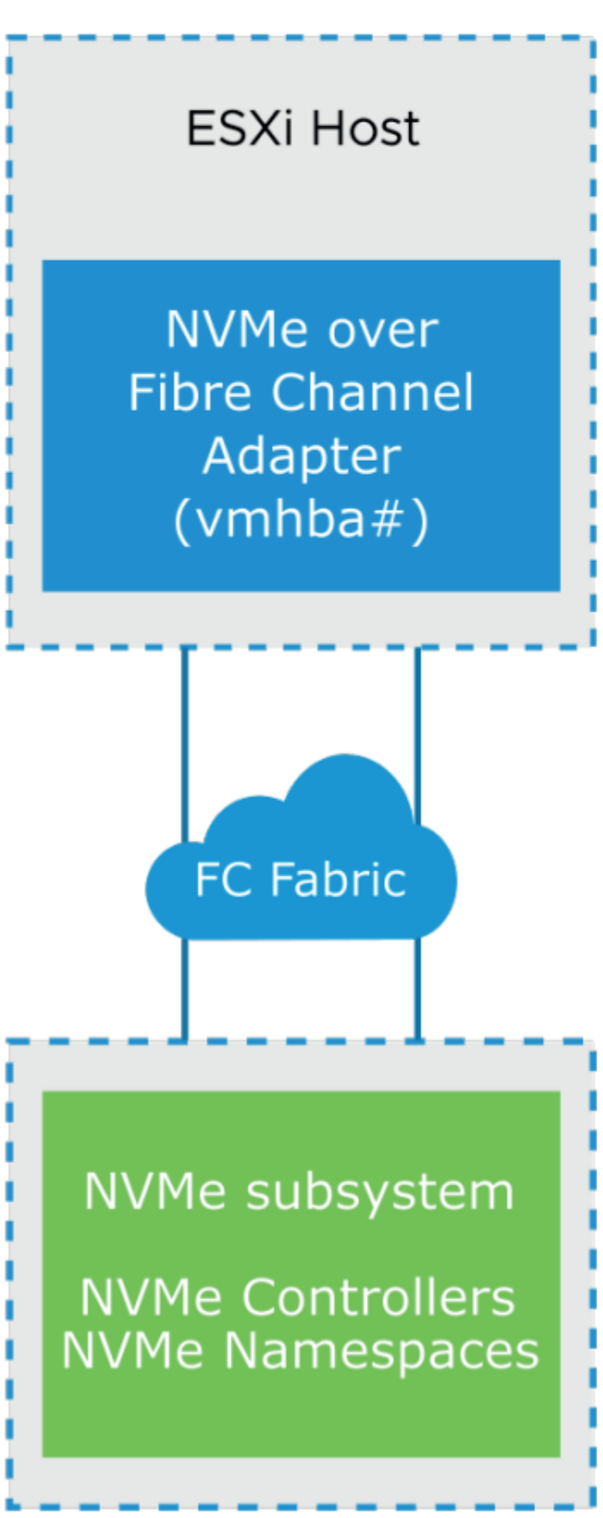

Your ESXi host detects the device and displays it in the vSphere Client as a standard Fibre Channel adapter (vmhba) with the storage protocol indicated as NVMe.

You do not need to configure the controller. After you install the required hardware NVMe adapter, it automatically connects to all targets and controllers that it can reach.

You can later reconfigure the adapter and disconnect its controllers or connect other controllers that were not available during the host boot.

5-10 Requirements for FC-NVMe

When you use FC-NVMe, the following components are required:

- Fibre Channel storage array that supports NVMe.
- NVMe controller on your storage array.
- ESXi 7 or later.
- Hardware NVMe adapter on your ESXi host: This adapter is typically a Fibre Channel host bus adapter (HBA) that supports NVMe.

Configuration maximums for NVMe-oF:

- 256 namespaces
- 2048 paths
- 4 paths per namespace per host

For more information about NVMe concepts, see the chapter called About VMware NVMe Storage in the *vSphere Storage* documentation at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-storage-guide.pdf.

Technet24

5-11 Support for NVMe over RDMA (RoCE v2)

The NVMe over RDMA (RoCE v2) technology uses RDMA for transport between two systems on the network.

Data exchange occurs in the main memory, bypassing the OS and processor of both systems.

ESXi 7 and later supports RoCE v2 technology, which enables RDMA over an Ethernet network.

To access storage, the ESXi host uses an RDMA network adapter installed on your host and a software NVMe over RDMA storage adapter.

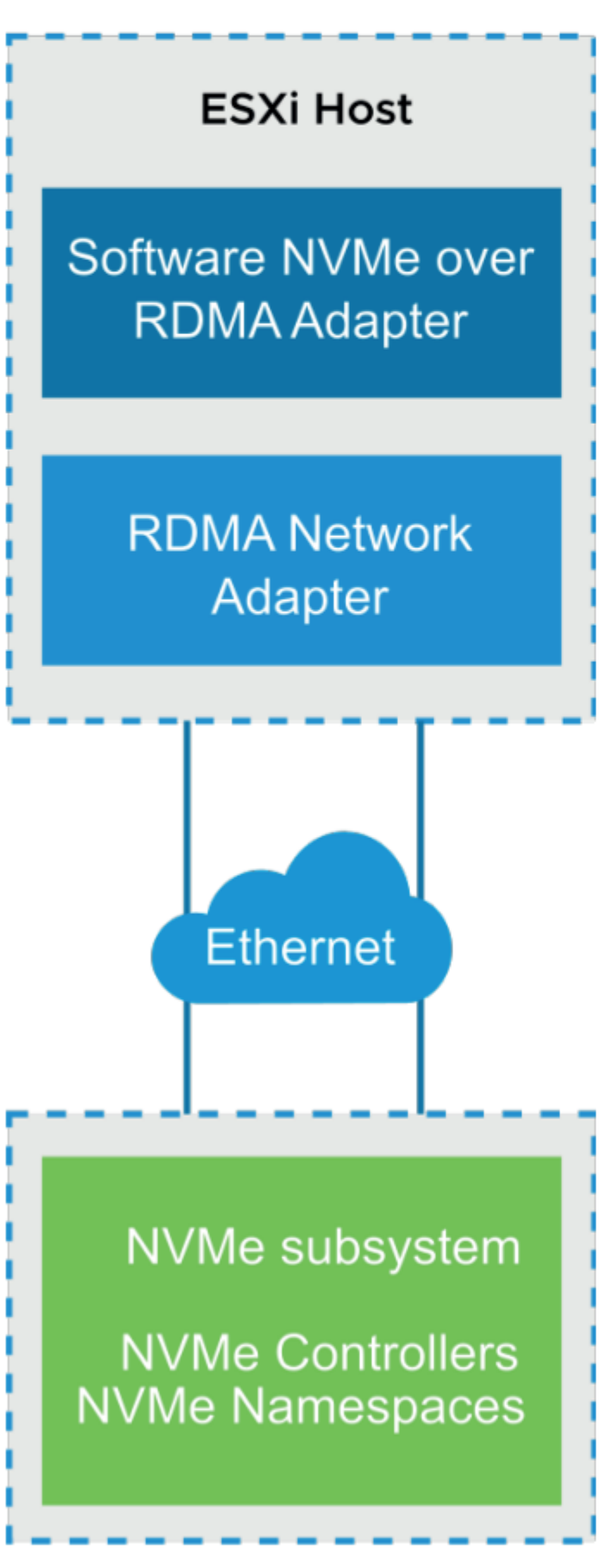

NVMe over RDMA with RoCE v2 technology provides ultra-efficient storage to optimize performance-sensitive workloads over your existing Ethernet fabric.

You must configure both adapters to be able use storage discovery. The adapter configuration process involves setting up VMkernel binding for an RDMA network adapter and enabling a software NVMe over RDMA adapter.

5-12 Requirements for NVMe over RDMA (RoCE v2)

NVMe over RDMA (RoCE v2) requires the following components:

- NVMe storage array with NVMe over RDMA (RoCE v2) transport support
- NVMe controller on your storage array
- ESXi 7 or later
- Ethernet switches supporting a lossless network
- Network adapter on your ESXi host that supports RoCE v2
- Software NVMe over RDMA adapter enabled on your ESXi hosts

Configuration maximums for NVMe-oF:

- 32 namespaces
- 128 paths
- 4 paths per namespace per host

5-13 Configuring ESXi for NVMe over RDMA (RoCE v2) (1)

To configure your ESXi host to use NVMe over RDMA (RoCE v2):

1. On your ESXi host, install a network interface that supports RDMA (RoCE v2).
2. In the vSphere Client, verify that the RDMA adapter is discovered by your host.
3. Configure VMkernel binding for the RDMA adapter.
4. Verify the VMkernel binding configuration for the RDMA adapter.
5. Enable the software NVMe over RDMA adapter.
6. Discover NVMe controllers on your array (automatic or manual).
7. Rescan storage to discover NVMe namespaces.
8. Create your datastores.

Technet24

5-14 Configuring ESXi for NVMe over RDMA (RoCE v2) (2)

Identify the physical network adapter that corresponds to the RDMA adapter.

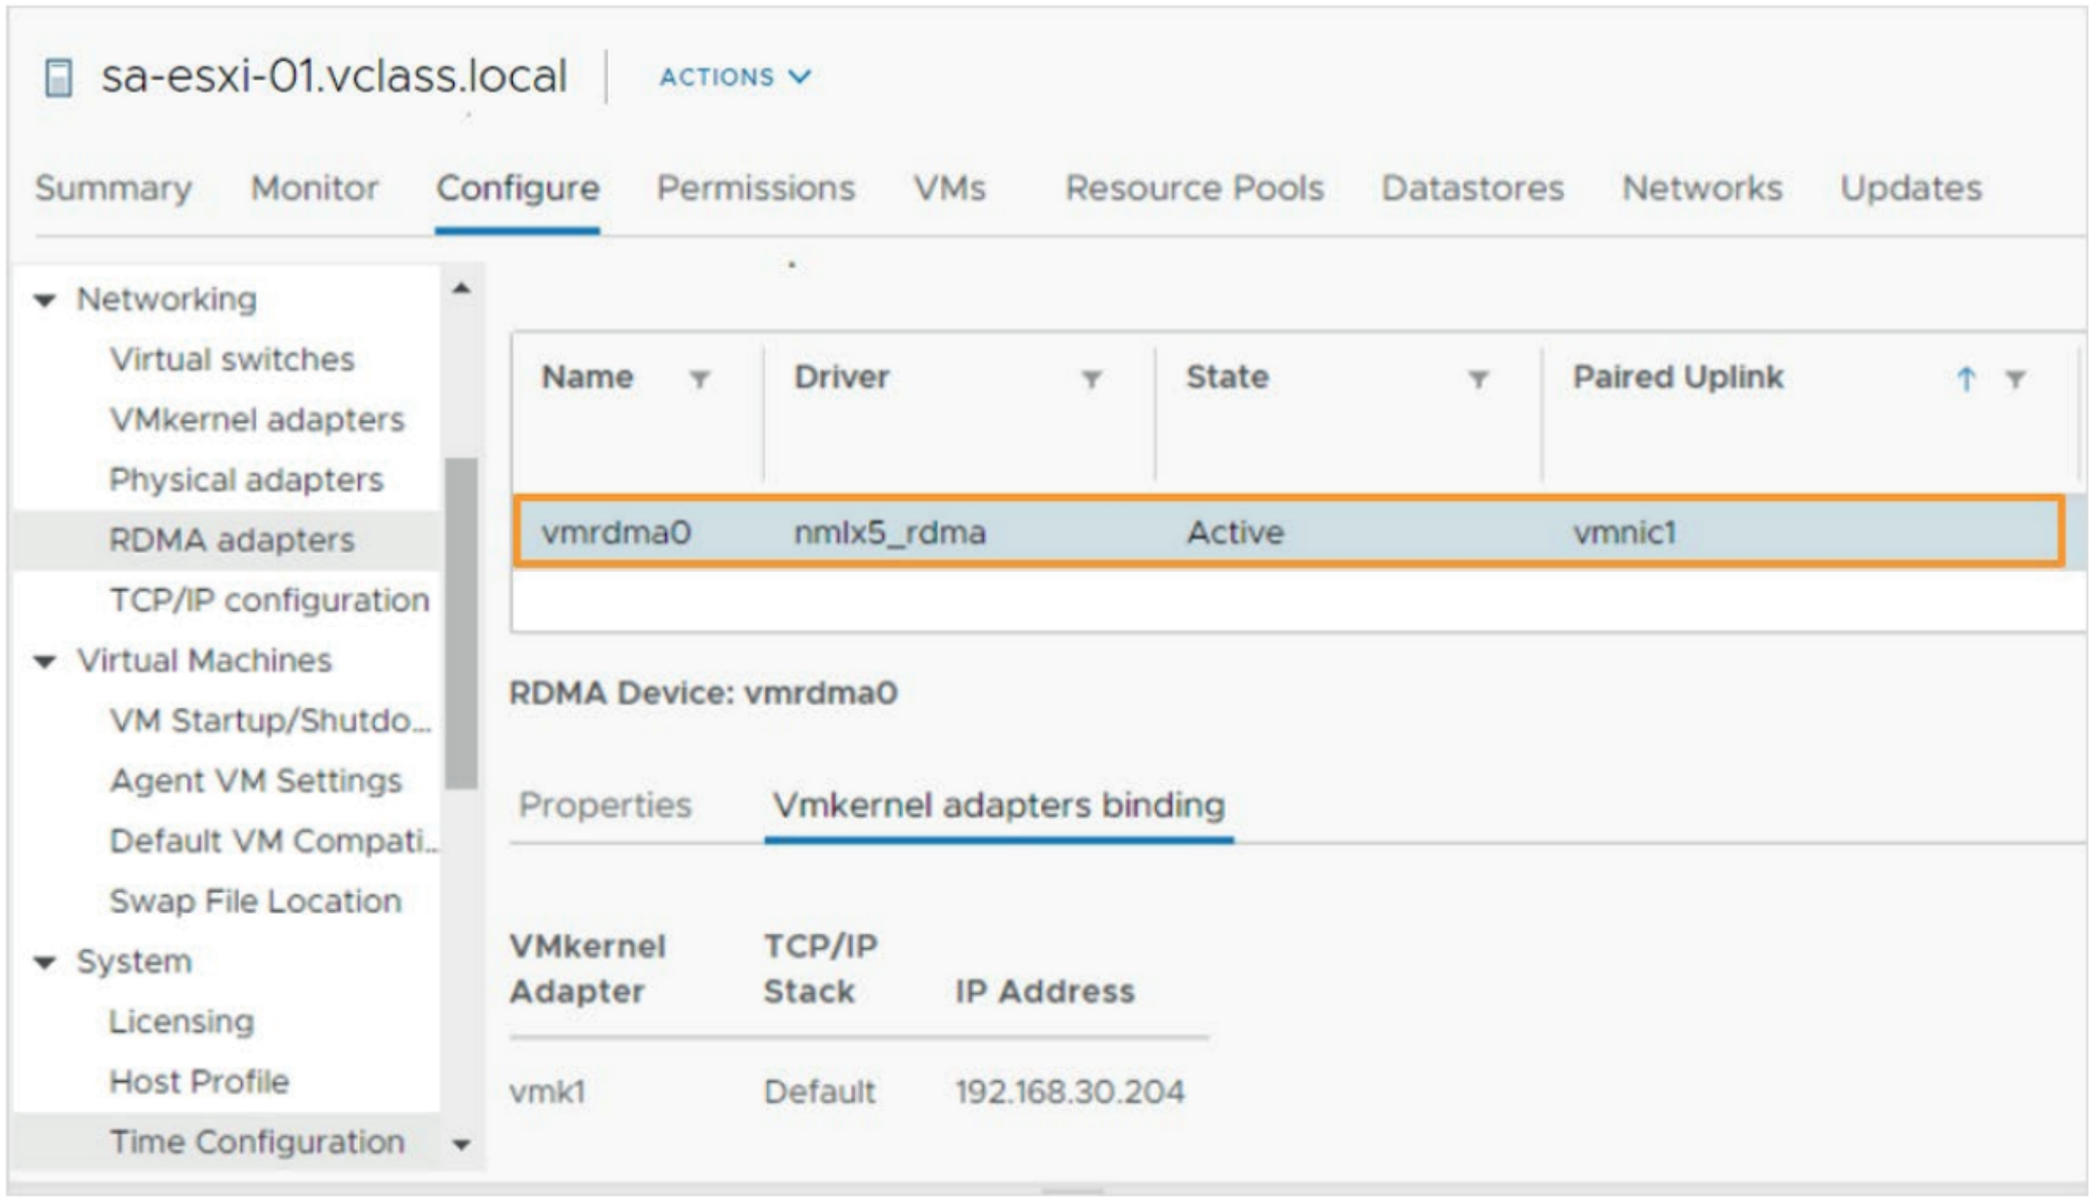

5-15 Configuring ESXi for NVMe over RDMA (RoCE v2) (3)

Configuring the Port binding for NVMe over RDMA involves creating a switch and connecting the physical network adapter and the VMkernel adapter to the switch.

Through this connection, the RDMA adapter becomes bound to the VMkernel adapter.

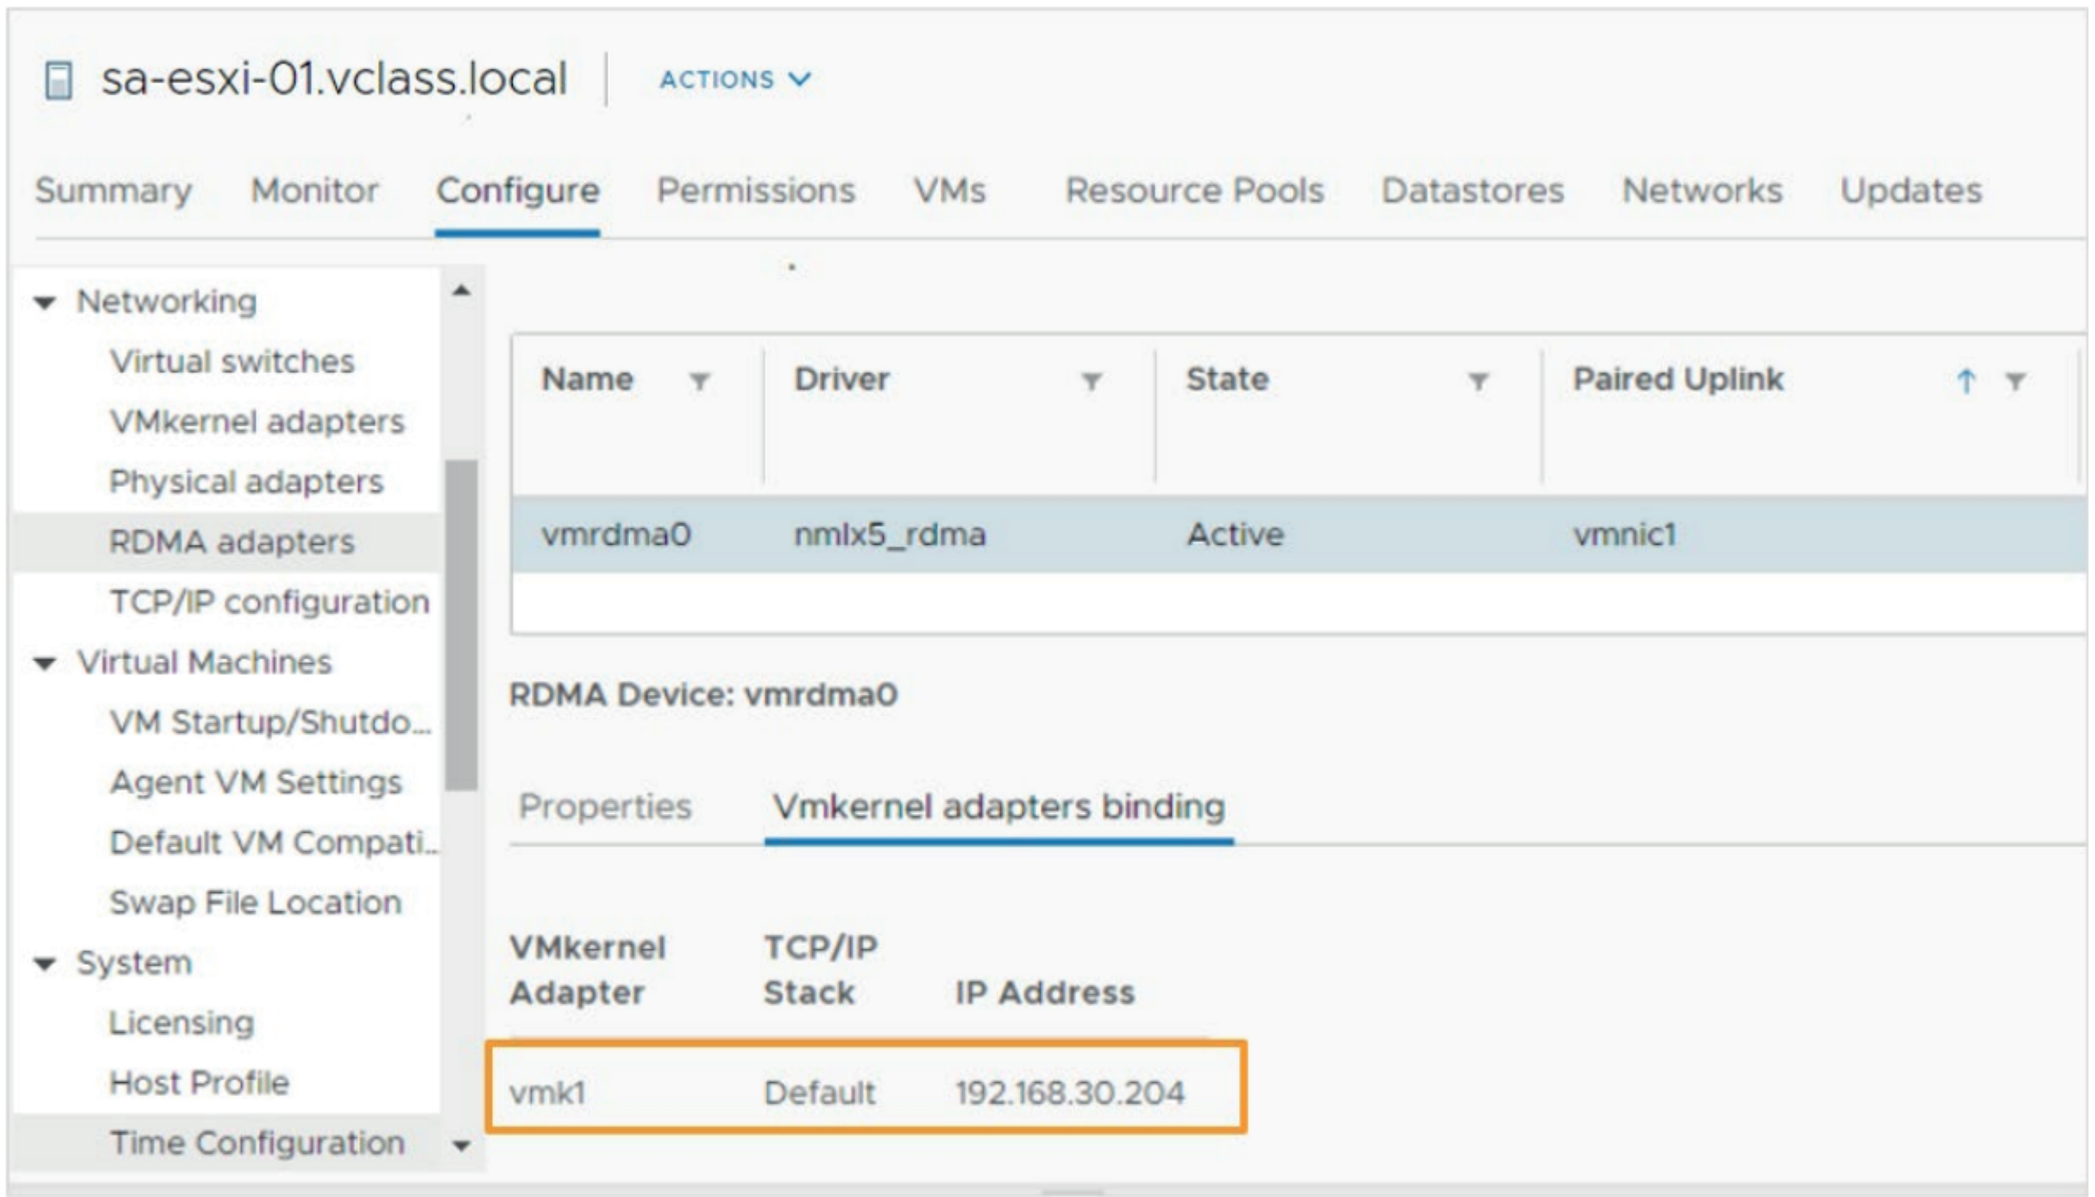

Technet24

5-16 Configuring ESXi for NVMe over RDMA (RoCE v2) (4)

Select **Add software NVMe over RDMA adapter** and select an appropriate RDMA adapter (vmrdma) from the drop-down menu.

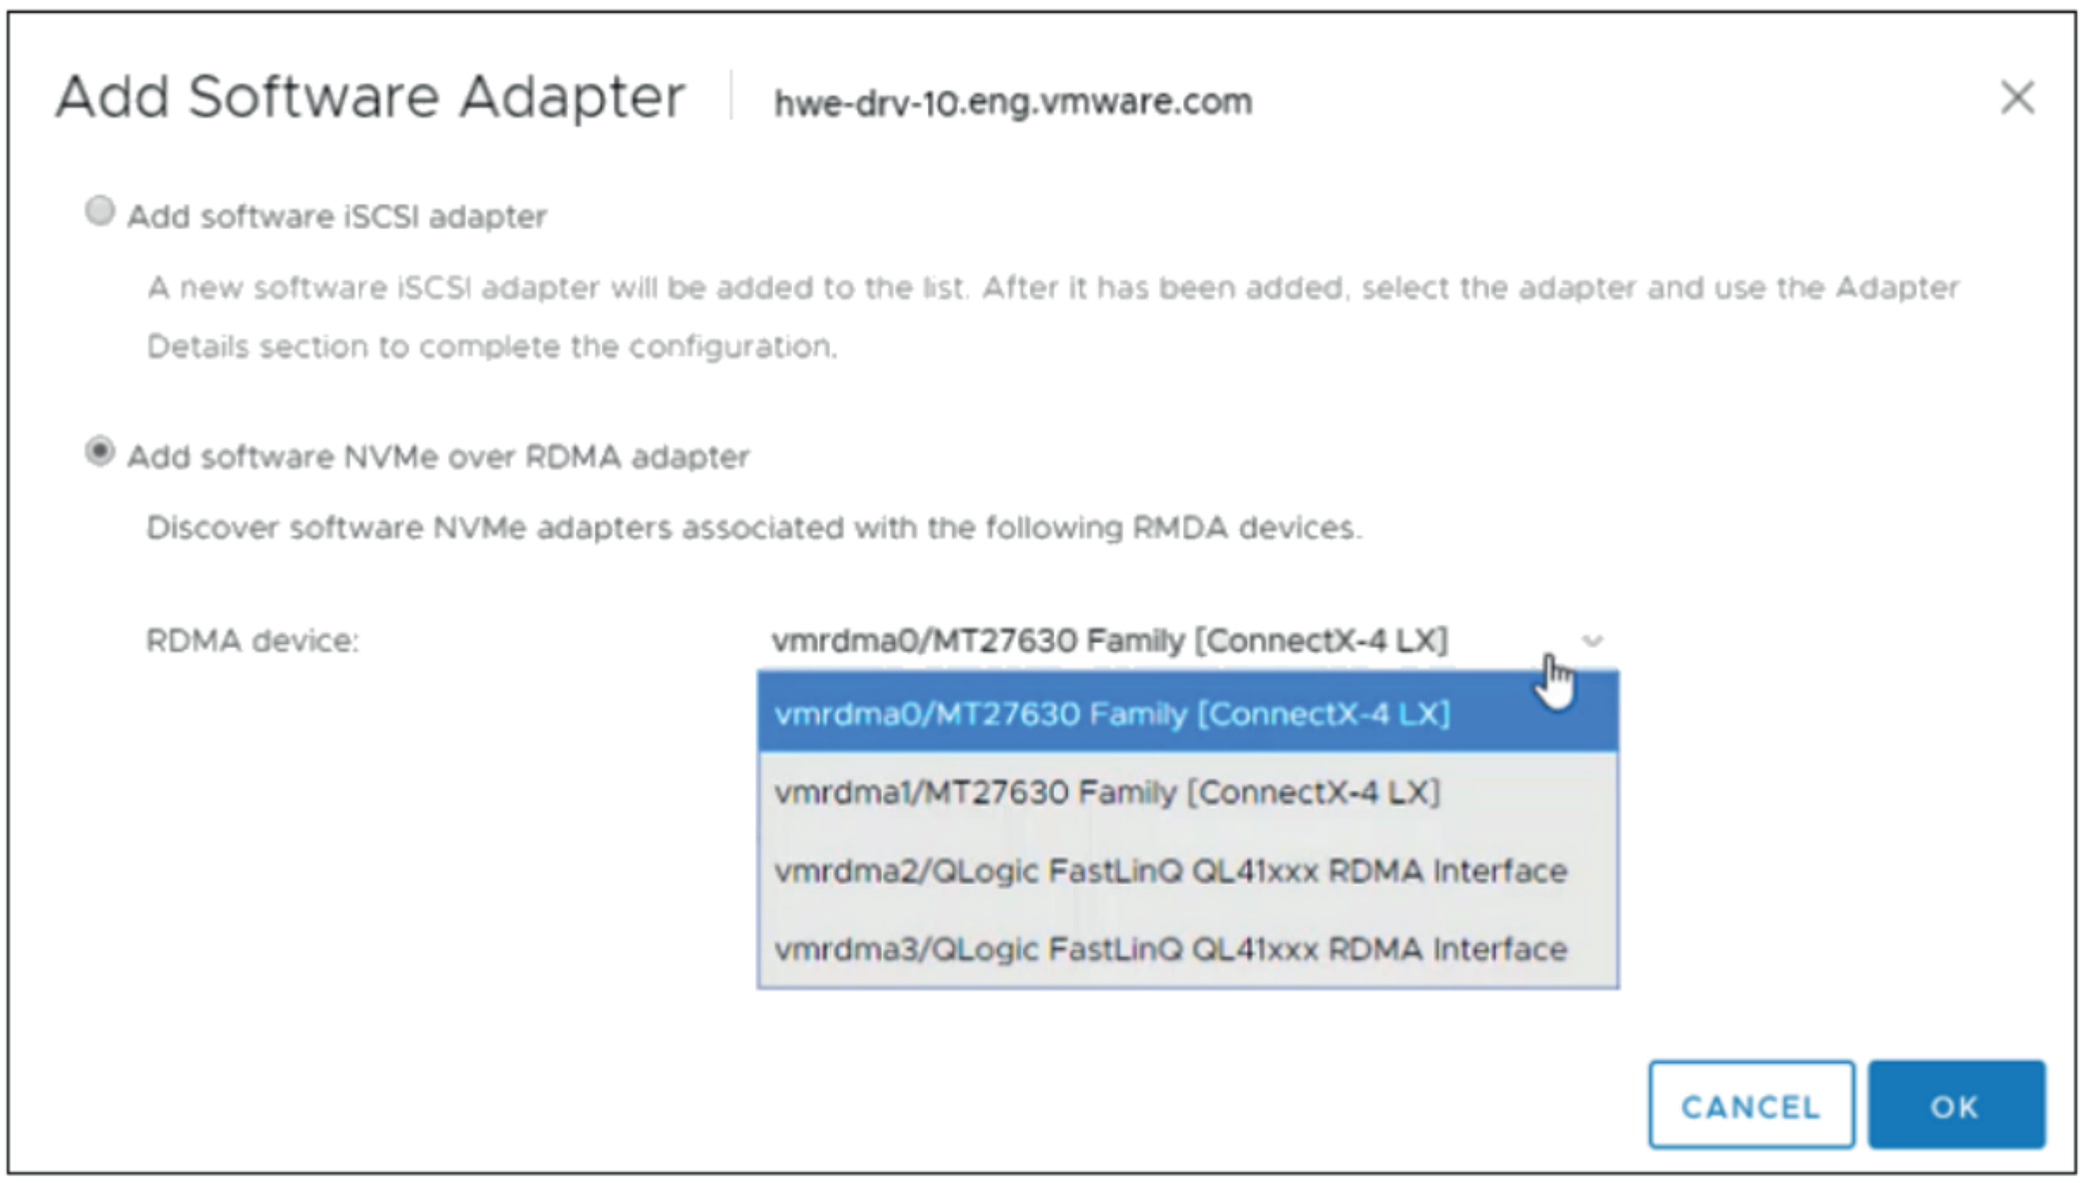

The software NVMe over RDMA adapter appears on the list as a vmhba storage adapter. You can remove the adapter if you must free the underlying RDMA network adapter for other purposes.

For more information about configuring adapters for NVMe over RDMA (RoCE v2) storage, see *Configuring Adapters for NVMe over RDMA (RoCE v2) Storage* at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-storage/GUID-F4B42510-9E6D-4446-816A-5012866E0038.html.

5-17 Pluggable Storage Architecture

Pluggable Storage Architecture is an open and modular framework that coordinates various software modules responsible for multipathing operations.

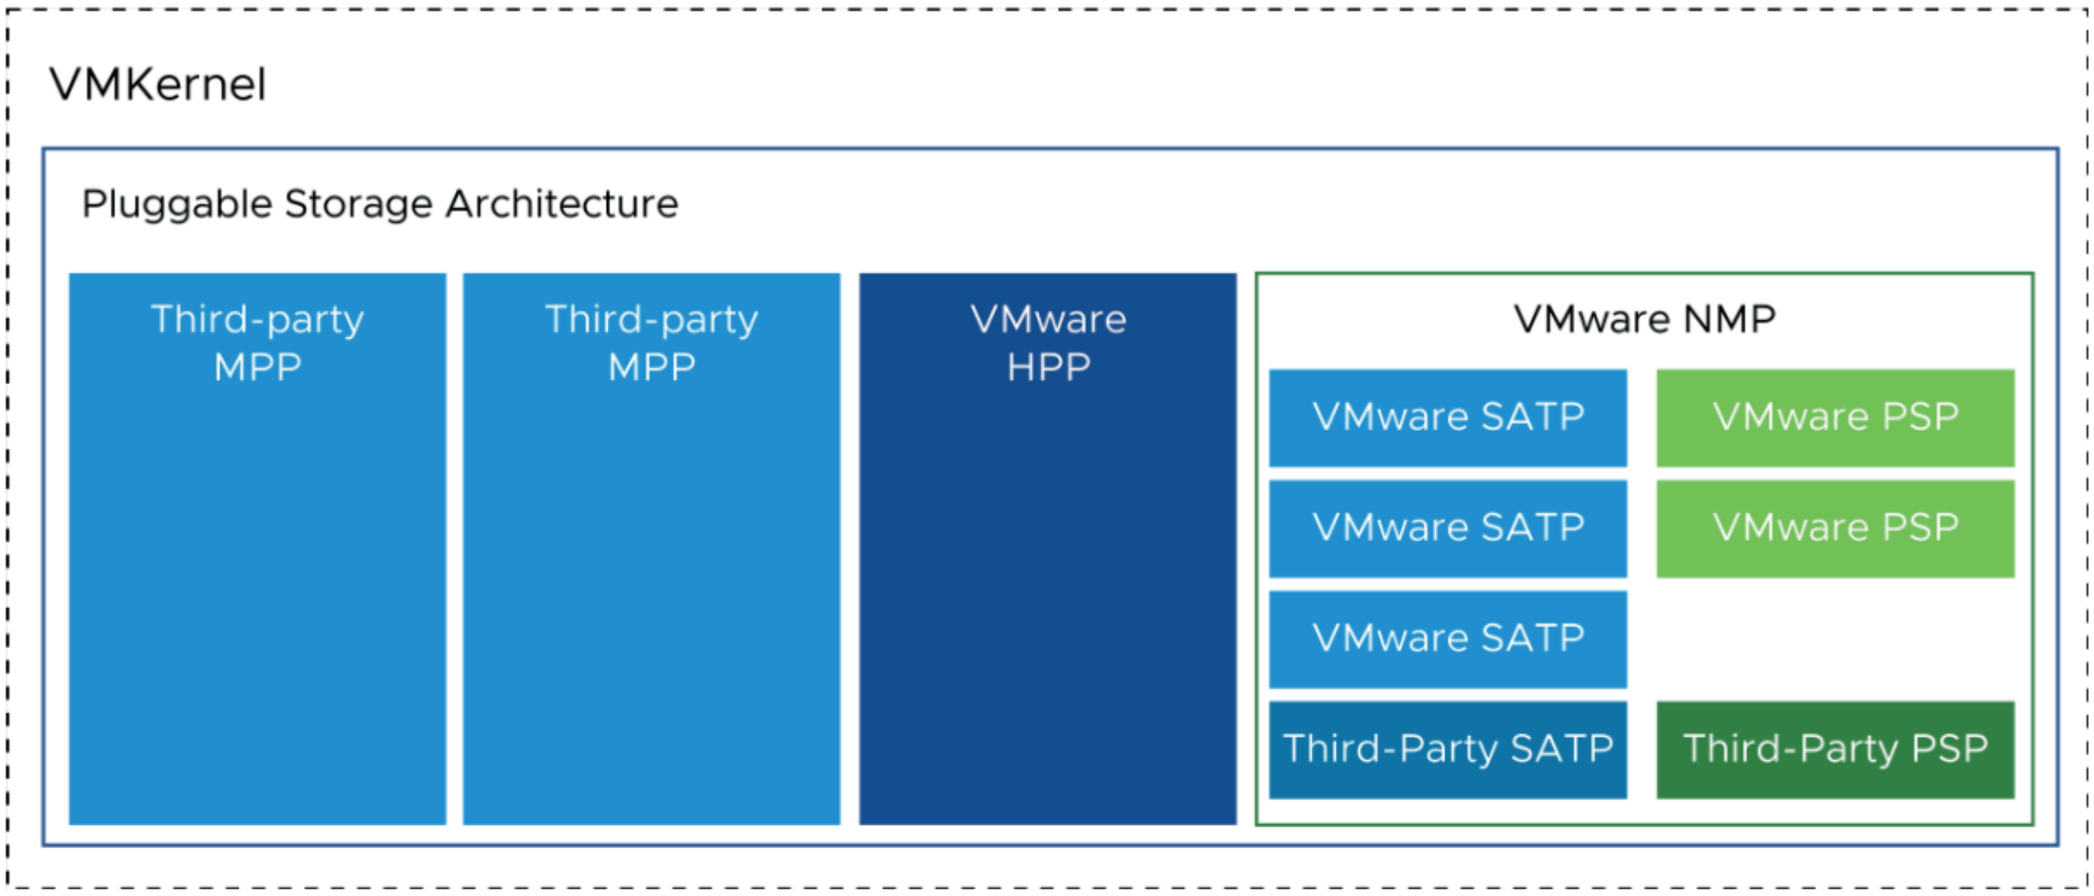

Acronym	Definition
PSA	Pluggable Storage Architecture
NMP	Native Multipathing Plug-in: Generic VMware multipathing module that is used with SCSI storage devices
PSP	Path Selection Plug-in: Handles path selection for a SCSI storage device
SATP	Storage Array Type Plug-in: Handles path failover for a given SCSI storage array

Acronym	Definition
MPP (third party)	Multipathing Plug-in: A multipathing module developed and provided by a third party
HPP	High-Performance Plug-in provided by VMware: Used with ultra-fast local and networked flash devices
PSS	Path Selection Scheme: Handles multipathing for NVMe storage devices

Technet24

VMware provides generic native multipathing modules, called VMware NMP and VMware HPP.

In addition, the PSA offers a collection of VMkernel APIs that third-party developers can use.

Software developers can create their own load balancing and failover modules for a particular storage array. These third-party multipathing modules (MPPs) can be installed on the ESXi host and run in addition to the VMware native modules, or as their replacement.

As the diagram on the slide shows, multiple third-party MPPs can run in parallel with the VMware NMP or HPP.

When installed, the third-party MPPs can replace the behavior of the native modules. The MPPs can take control of the path failover and the load-balancing operations for the specified storage devices.

5-18 High-Performance Plug-In for NVMe

The HPP performs the following functions:

- Acts as the default plug-in that claims NVMe-oF targets
- Supports only active-active and implicit Asymmetric Logical Unit Access (ALUA) targets
- Improves the performance of ultra-fast flash devices installed on ESXi
- Can replace the Native Multipathing Plug-in for local NVMe devices
- Supports multipathing for NVMe-oF

To support multipathing, HPP uses a Path Selection Scheme (PSS) when selecting physical paths for I/O requests. You can use vSphere ESXi Shell or the vSphere CLI (Prior to vSphere 7.0) to configure the HPP and PSS.

In vSphere 7 and later, HPP integrates path selection and failover logic within a single module, reducing cross-module calls.

For local NVMe devices, NMP remains the default plug-in, but you can replace it with HPP.

5-19 Support for High-Performance Plug-In Multipathing

The HPP improves the performance of NVMe devices on your ESXi host.

HPP Support	vSphere 7 and Later
Storage devices	Local NVMe PCIe Shared NVMe-oF (active-active and implicit ALUA targets only)
Multipathing	Yes
Second-level plug-ins	No Path Selection Scheme (PSS)
SCSI-3 persistent reservations	No
4Kn devices with software emulation	No
vSAN	Yes

Technet24

5-20 High-Performance Plug-In Best Practices

To achieve the fastest throughput from a high-speed storage device, follow these recommendations:

- Use the HPP for NVMe local or networked devices.
- Do not activate the HPP for HDDs or slower flash devices.
- When using FC-NVMe, follow general recommendations for Fibre Channel storage.
- When using NVMe-oF, do not mix transport types to access the same namespace.
- Before registering NVMe-oF namespaces, ensure that active paths are presented to the host.
- Configure your VMs to use VMware Paravirtual controllers (PVSCSI).
- Configure your VMs to use high-latency sensitivity.
- If a single VM drives a significant share of the device's I/O workload, consider spreading the I/O across multiple virtual disks. Attach the disks to separate virtual controllers in the VM.

Do not activate the HPP for HDDs or slower flash devices. The HPP does not provide any performance benefits with devices incapable of at least 200,000 IOPS.

If you do not attach the disks to separate virtual controllers in the VM, I/O throughput might be limited. This limitation is due to saturation of the CPU core that is responsible for processing I/O operations on a virtual storage controller.

5-21 Review of Learner Objectives

- Recognize the support for NVMe
- Configure ESXi for NVMe over RDMA (RoCE v2)
- Describe the VMware Pluggable Storage Architecture
- Describe the purpose of the High-Performance Plug-in

5-22 Lesson 2: iSCSI Extension for RDMA (iSER)

5-23 Learner Objectives

- Describe support for iSER in vSphere
- Configure iSER on ESXi hosts

5-24 About iSER

iSCSI Extensions for RDMA (iSER) provides high-performance iSCSI connectivity for standard RDMA NICs.

In addition to supporting traditional iSCSI, ESXi supports the iSER protocol. When the iSER protocol is enabled, the iSCSI framework on the ESXi host can use the RDMA transport instead of TCP/IP.

vSphere 7.0 added support for vSphere Virtual Volumes over iSER.

In vSphere 7 and later, support for iSER requires the following:

- RoCE-enabled devices
- A target array that supports iSER on RoCE

Technet24

5-25 Configuring iSER

Configure iSER on your ESXi host so that the iSCSI framework on the host can use Remote Direct Memory Access (RDMA) transport instead of TCP/IP.

Configuration Step	Description
Enable the VMware iSER adapter.	Use the `esxcli` command to enable the VMware iSER adapter.
Modify general properties for iSCSI or iSER adapters.	If needed, change the default name and alias assigned to your adapter.
Configure port binding for iSCSI or iSER.	You must create network connections to bind the iSER engine and the RDMA-capable network adapter.
Configure dynamic or static discovery for iSCSI and iSER.	With the dynamic discovery, each time the initiator contacts a specified iSER storage system, it sends the `SendTargets` request to the system. With static discovery, manually enter information for the targets.
Set up CHAP for iSCSI or iSER.	This step is necessary only if your environment uses the Challenge Handshake Authentication Protocol (CHAP).
Enable jumbo frames for networking.	If your environment supports jumbo frames, enable them for the adapter.

When installed on the host, the RDMA-capable adapter appears in vCenter Server as a network adapter (vmnic).

To make the adapter functional, you must enable the VMware iSER component and connect the iSER adapter to the RDMA-capable vmnic. You can then configure properties, such as targets and CHAP, for the iSER adapter.

NOTE: iSER does not support NIC teaming. When configuring port binding, use only one RDMA adapter per virtual switch.

5-26 Requirements for iSER

To enable the VMware iSER adapter, use the `esxcli` command and ensure that you meet the following requirements:

- Verify that your iSCSI storage supports the iSER protocol.
- Install the RDMA-capable adapter on your ESXi host.
- Use the RDMA-capable switch.
- Enable flow control on the ESXi host. To enable flow control for the host, use the `esxcli system module parameters` command.
- Configure RDMA switch ports to create lossless connections between the iSER initiator and target.

To enable flow control, see the VMware knowledge base article 1013413 at http://kb.vmware.com/kb/1013413.

Technet24

5-27 Running iSER

Run the following command from the ESXi host to enable the iSER adapter.

```
esxcli rdma iser add
```

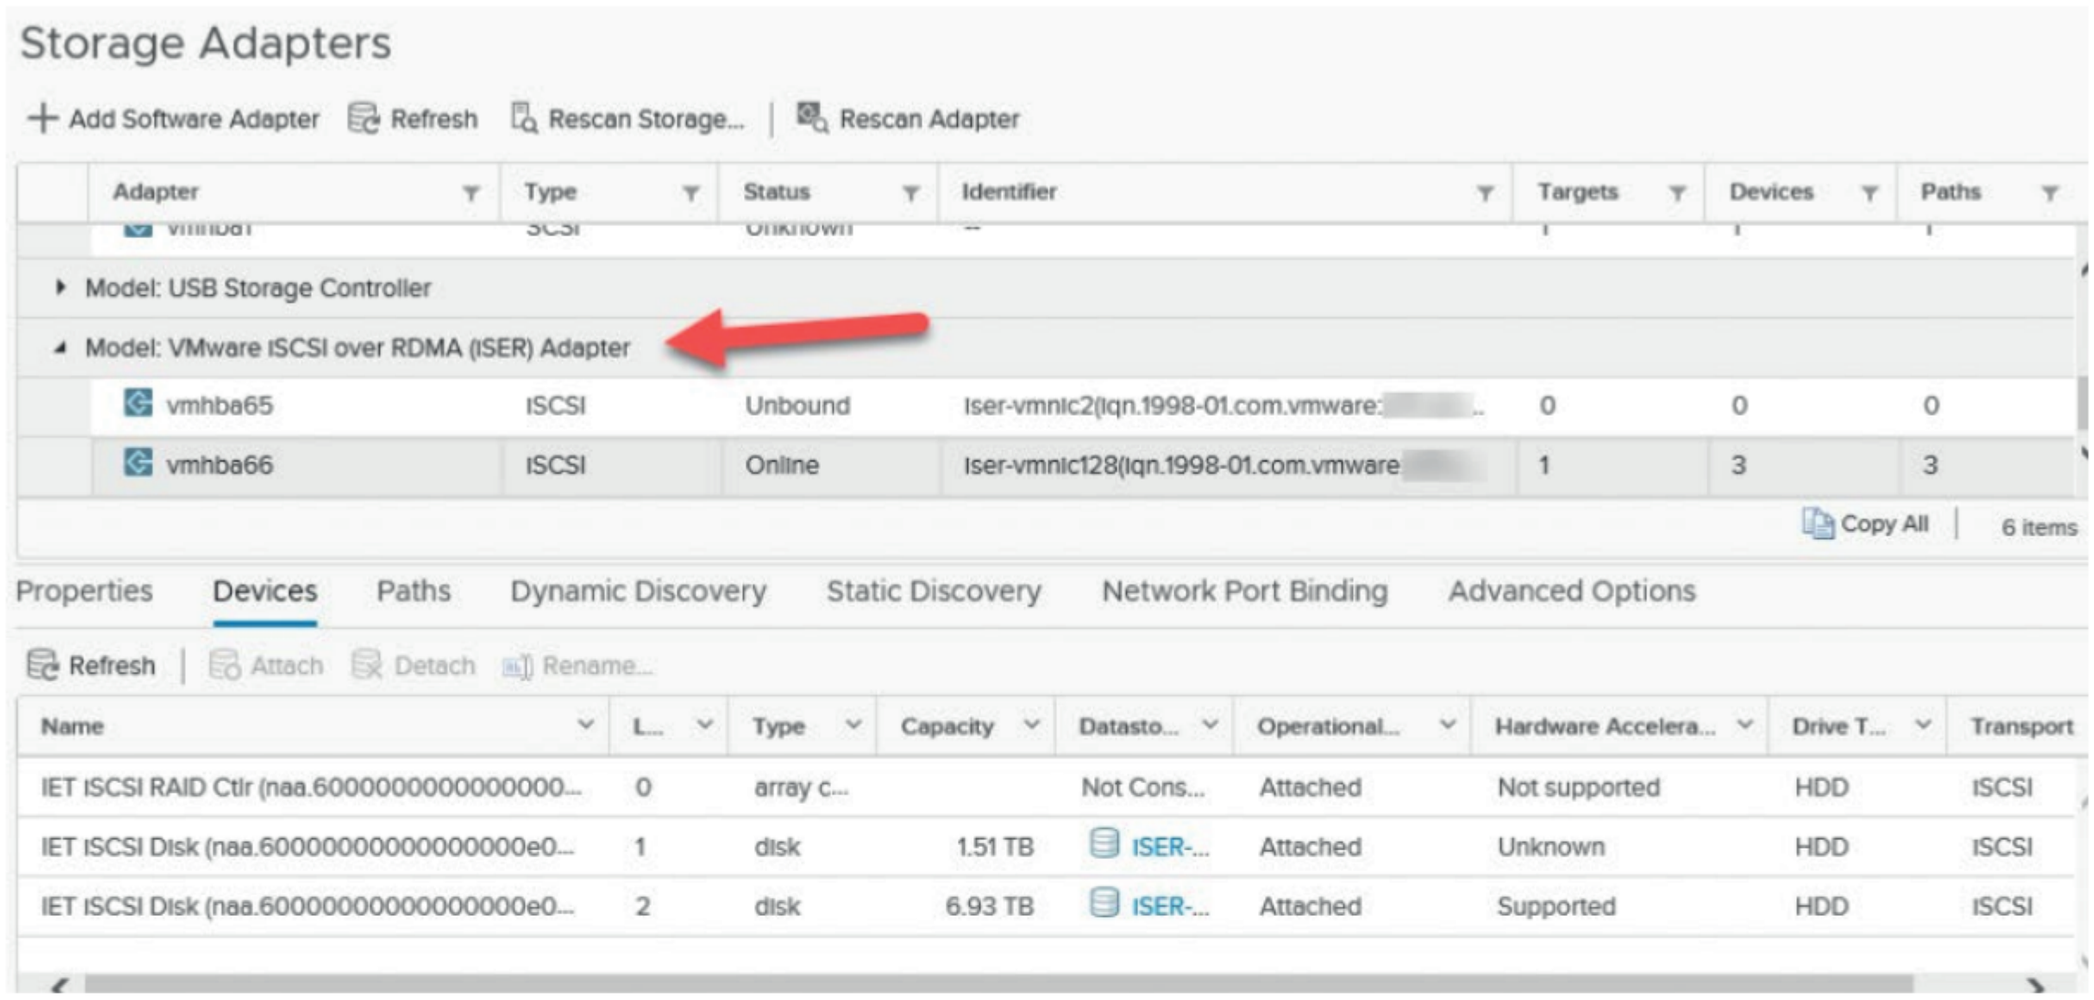

5-28 Review of Learner Objectives

- Describe support for iSER in vSphere
- Configure iSER on ESXi hosts

5-29 Lesson 3: VMFS Datastores

5-30 Learner Objectives

- Explain why VMFS is a high-performance, scalable file system
- Identify the performance improvements of VMFS6 over VMFS5
- Describe the migration procedure from VMFS5 to VMFS6

5-31 About VMFS Datastores

VMFS Datastores are logical containers that hide specifics of physical storage from virtual machines and provide a uniform model for storing the virtual machine files

VMFS is a high-performance, cluster file system that is optimized for VMs.

vSphere 6.5 and later versions support VMFS5 and VMFS6 datastores

VMFS5 was introduced in vSphere 5.0. VMFS6 as introduced in vSphere 6.5.

Technet24

5-32 VMFS5 and VMFS6 Datastores

VMFS5 and VMFS6 datastores are highly scalable:

- The maximum size for both a datastore and a single extent is 64 TB.
- The file system block size used is 1 MB, which supports files up to 62 TB.
- Small files are efficiently stored:
  - File system blocks are divided into subblocks to efficiently store small files.
  - Data of small files (less than or equal to 1 KB) is stored directly in the file descriptor.
- A virtual disk can be hot-extended to any size, up to 62 TB.

For more information about the hot-extend feature, see *vSphere Storage* at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-storage-guide.pdf.

5-33 VMFS6 Datastores

VMFS6 provides improvements in scalability and performance over VMFS5:

- VMFS6 (and NFS) datastores can be mounted onto up to 128 ESXi hosts
- Faster file creation times
- Parallel disk operations and concurrent transactions

Shared VMDK on VMFS6 supports SCSI-3 Persistent Reservations (SCSI-3 PR) at the virtual disk (VMDK) level:

- A Windows Server Failover Cluster (WSFC) can use shared disks on VMFS.
- It can support clustered virtual disks on specific datastores (requires supported hardware).
- You can use the vSphere Client to migrate from RDMs to VMFS.

For virtual-compatibility RDMs, you can perform a Storage vMotion migration.

From physical-compatibility RDMs to virtual disks, downtime is required to migrate.

5-34 Storage Device Formats

ESXi supports storage devices with traditional and advanced sector formats:

- 512n: A traditional sector format.
- 512e (512-byte emulation): A 4K drive that can emulate 512 byte sectors for legacy apps and guest OS's
- 4Kn (4K native): Presents 4KB sectors without any emulation.

Sector Format	Logical Sector Size	Physical Sector Size
512n	512	512
512e	512	4,096
4Kn	4,096	4,096

The storage industry is moving towards advanced format drives to provide large capacity drives to servers and storage arrays.

4K drives have advantages over 512-byte sector drives:

- Improved total cost of ownership: 4K drives require less space for error correction codes (ECCs) than regular 512-byte sector drives. This space saving results in greater data density on 4K drives, which leads to an improved total cost of ownership.
- Improved availability: 4K drives have a larger ECC field, which improves error correction and data integrity.
- Improved performance: 4K drives are expected to have better performance than 512-byte sector drives. However, this improved performance is true only when the guest operating system is configured to issue I/O operations that are aligned to the 4K sector size.

Because legacy applications and operating systems do not support 4K native drives, the drive industry introduced 4K sector size drives in 512-byte emulation mode to support legacy applications. The drives have a physical sector size of 4K, but the logical sector size is 512 bytes.

These drives are called 512e drives, where the letter e stands for emulation mode. 512e drives help drive adoption of 4K drives. The goal is to move to 4K native drives.

Technet24

5-35 Storage Devices Supported by VMFS Datastores

You can deploy VMFS5 and VMFS6 datastores on 512n and 512e storage devices.

With vSphere 6.7 and later, VMFS6 datastores can be deployed on 4Kn direct-attached storage devices.

ESXi continues to expose 512-byte sector virtual disk files to the guest operating system.

Storage Device	Supported by VMFS5	Supported by VMFS6
512n	Yes.	Yes.
512e	Yes. However VMFS5 is not supported on local 512e devices.	Yes.
4Kn	No.	Yes.

For more information about vSphere support for 512e and 4K native drives, see VMware knowledge base article 2091600 at http://kb.vmware.com/kb/2091600.

5-36 Limitations of Using 4Kn Drives

When used in a vSphere environment, 4Kn drives have the following limitations:

- Only local SAS and SATA hard disk drives are supported.
- 4Kn SSD drives, 4Kn NVMe drives, and 4Kn drives as RDM are not supported.
- Booting from 4Kn drives is supported with UEFI only.
- Only the Native Multipathing Plug-In (NMP) can claim the 4Kn devices:
  - The High-Performance Plug-In (HPP) cannot be used to claim these devices.
  - Third-party multipathing plug-ins are not supported.
- For vSAN, 4Kn drives can only be used as capacity HDDS in Hybrid mode.

5-37 Snapshot Formats for VMFS

On the VMFS datastore, the delta disk of a VM snapshot is a sparse disk. Delta disks use different sparse formats depending on the type of VMFS datastore.

The SEsparse format is similar to VMFSsparse, with some enhancements:

- SEsparse is space efficient and supports space reclamation.
- This disk format is recommended for virtual desktop infrastructure workloads.

Sparse Format	Used by VMFS5	Used by VMFS6
VMFSsparse	Yes, for virtual disks smaller than 2 TB.	No.
SEsparse (space-efficient sparse)	Yes, for virtual disks equal to or larger than 2 TB.	Yes, it is the default format.

When you take a snapshot, the state of the virtual disk is preserved, preventing the guest operating system from writing to the disk. A delta or child disk is created. The delta disk represents the difference between the current state of the virtual disk and the state that existed when the previous snapshot was taken.

VMFSsparse is implemented on top of VMFS. The VMFSsparse layer processes the I/O that is issued to a snapshot VM. Technically, VMFSsparse is a redo-log that starts empty, immediately after a VM snapshot is taken. The redo-log expands to the size of its base VM disk (VMDK), when the entire VMDK is rewritten with new data after the VM snapshots are taken. This redo-log is only a file in the VMFS datastore. After the snapshot is taken, the base VMDK that is attached to the VM is changed to the newly created sparse VMDK.

With SEsparse space reclamation, blocks that are deleted by the guest OS are marked, and commands are issued to the SEsparse layer in the hypervisor to unmap those blocks. This operation helps reclaim the space allocated by SEsparse after the guest OS deletes that data.

SEsparse is the recommended disk format for VMware Horizon environments. In these environments, reclamation of storage space is critical because of the many tenants sharing storage.

Technet24

5-38 Migrating Your Data from VMFS5 to VMFS6

Because of the changes in the VMFS6 metadata structures, you cannot directly upgrade from VMFS5 to VMFS6.

A VMFS6 datastore can be created only as a new datastore.

Use vSphere Storage vMotion to migrate data from your VMFS5 datastore to a VMFS6 datastore.

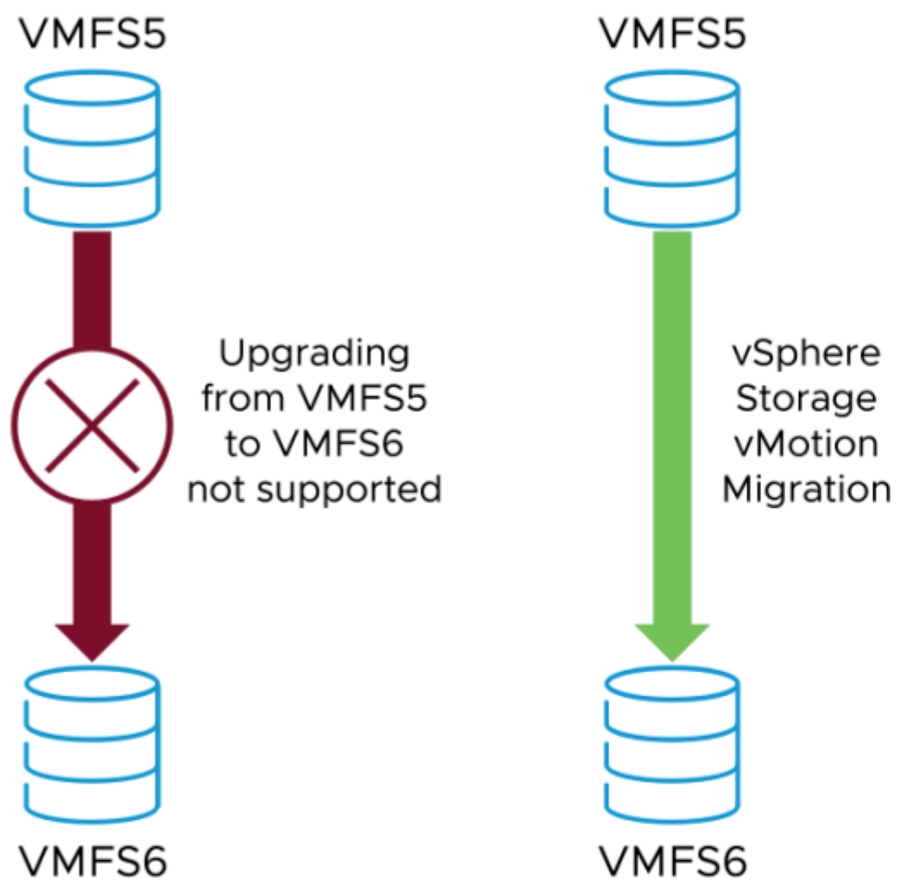

For more information, see VMware knowledge base article 2147824 at https://kb.vmware.com/s/article/2147824.

5-39 Activity: VMFS6 Features (1)

Which scalability and performance features are available only with VMFS6 datastores? Select all that apply.

- ◊ ❒ Support for 4Kn drives
- ◊ ❒ Faster file creation times
- ◊ ❒ SEsparse as a snapshot format
- ◊ ❒ Maximum datastore size of 64 TB
- ◊ ❒ Hot-extends of a virtual disk to any size up to 62 TB
- ◊ ❒ Parallel disk operations and concurrent transactions

5-40 Activity: VMFS6 Features (2)

Which scalability and performance features are available only with VMFS6 datastores? Select all that apply.

◊ ✓ Support for 4Kn drives

◊ ✓ Faster file creation times

◊ ❒ SEsparse as a snapshot format

◊ ❒ Maximum datastore size of 64 TB

◊ ❒ Hot-extends of a virtual disk to any size up to 62 TB

◊ ✓ Parallel disk operations and concurrent transactions

5-41 Review of Learner Objectives

- Explain why VMFS is a high-performance, scalable file system
- Identify the performance improvements of VMFS6 over VMFS5
- Describe the migration procedure from VMFS5 to VMFS6

Technet24

5-42 Lesson 4: vSAN Datastores

5-43 Learner Objectives

- Explain the purpose of a vSAN datastore
- Describe the architecture and requirements of vSAN configuration
- Explain the purpose of vSAN storage policies
- Describe the purpose of fault domains

5-44 vSAN Storage Architecture

Introduced in vSAN 8, vSAN Express Storage Architecture is a new optional storage architecture on the vSAN platform that enables new levels of performance, scalability, resilience, and simplicity.

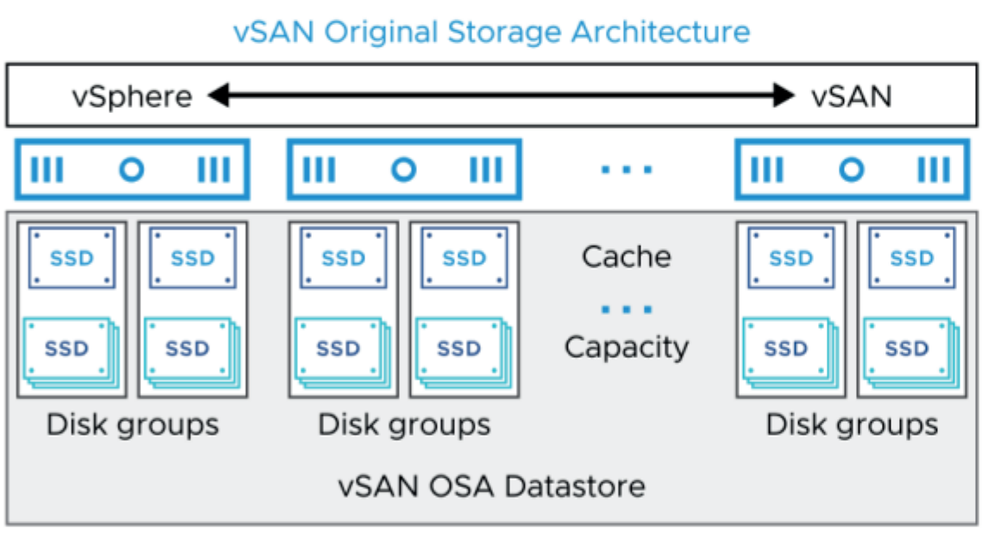

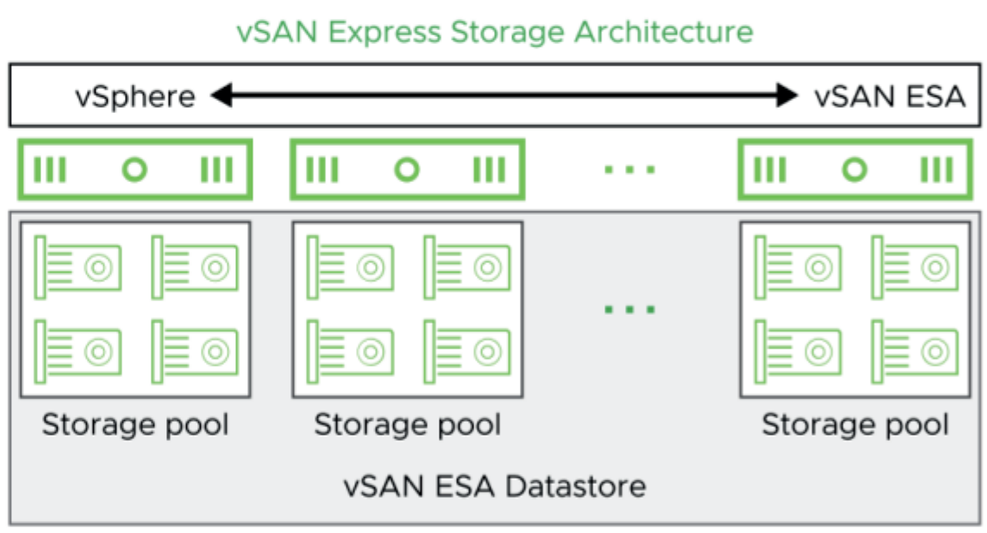

- Dedicated cache and capacity devices
- Accommodates older, wide ranging storage devices
- Uses disk groups to provide performance through caching and capacity

- Flexible architecture using single tier in VSAN 8
- Optimized for high performance NVMe based flash devices
- No disk groups. All devices contribute to performance and capacity

5-45 About vSAN Datastores

vSAN is a software-defined storage solution providing shared storage for vSphere clusters without using traditional external storage.

A vSAN cluster requires:

- A minimum of three hosts to be part of the vSphere cluster and configured for vSAN
- A vSAN network
- Local storage devices on each host that are pooled to create a virtual shared vSAN datastore.

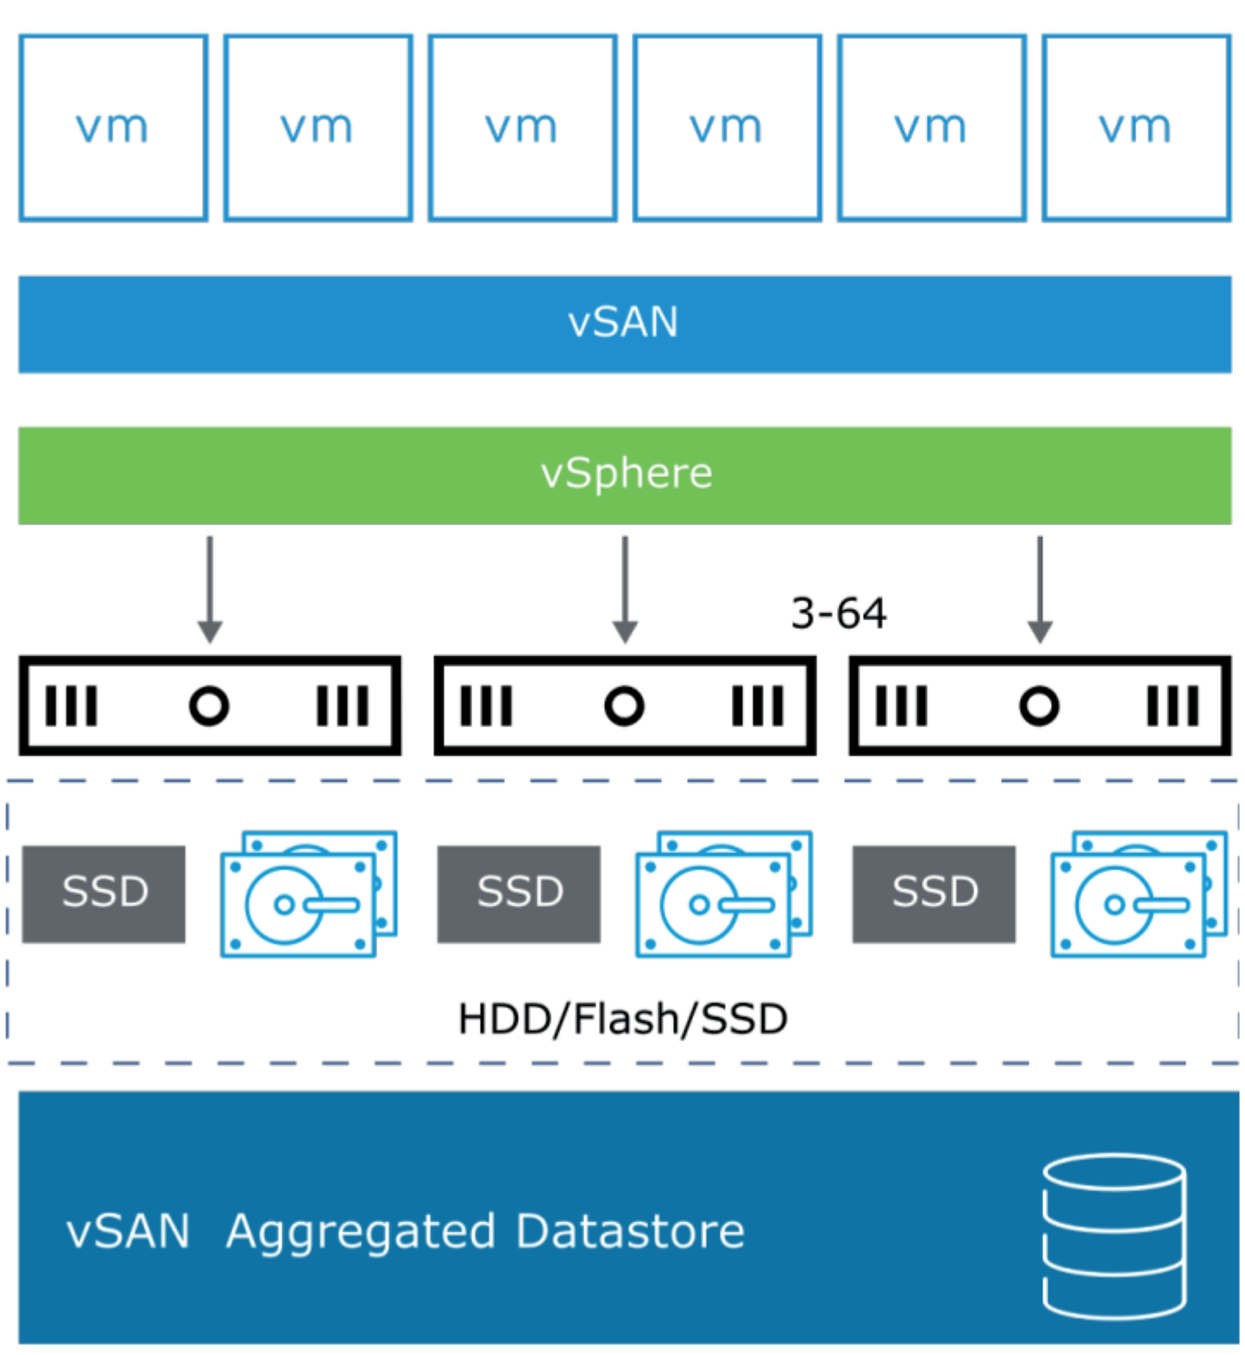

vSAN datastores help administrators use software-defined storage in the following ways:

- Storage policy per VM architecture: With multiple policies per datastore, each VM can have their own unique storage capabilities .
- vSphere and vCenter Server integration: vSAN capability is built in and requires no appliance. You create a vSAN cluster, like vSphere HA or vSphere DRS.
- Scale-out storage: Up to 64 ESXi hosts can be in a cluster. Scale out by populating new nodes in the cluster.
- Built-in resiliency: The default vSAN storage policy establishes RAID 1 redundancy for all VMs.

Technet24

5-46 Disk Groups

Disk groups are vSAN management constructs on all ESXi hosts in a vSAN cluster. A host can include a maximum of five disk groups.

The disk groups are combined to create a single vSAN datastore. A disk group requires:

- One flash device for caching
- One to seven capacity devices for storage

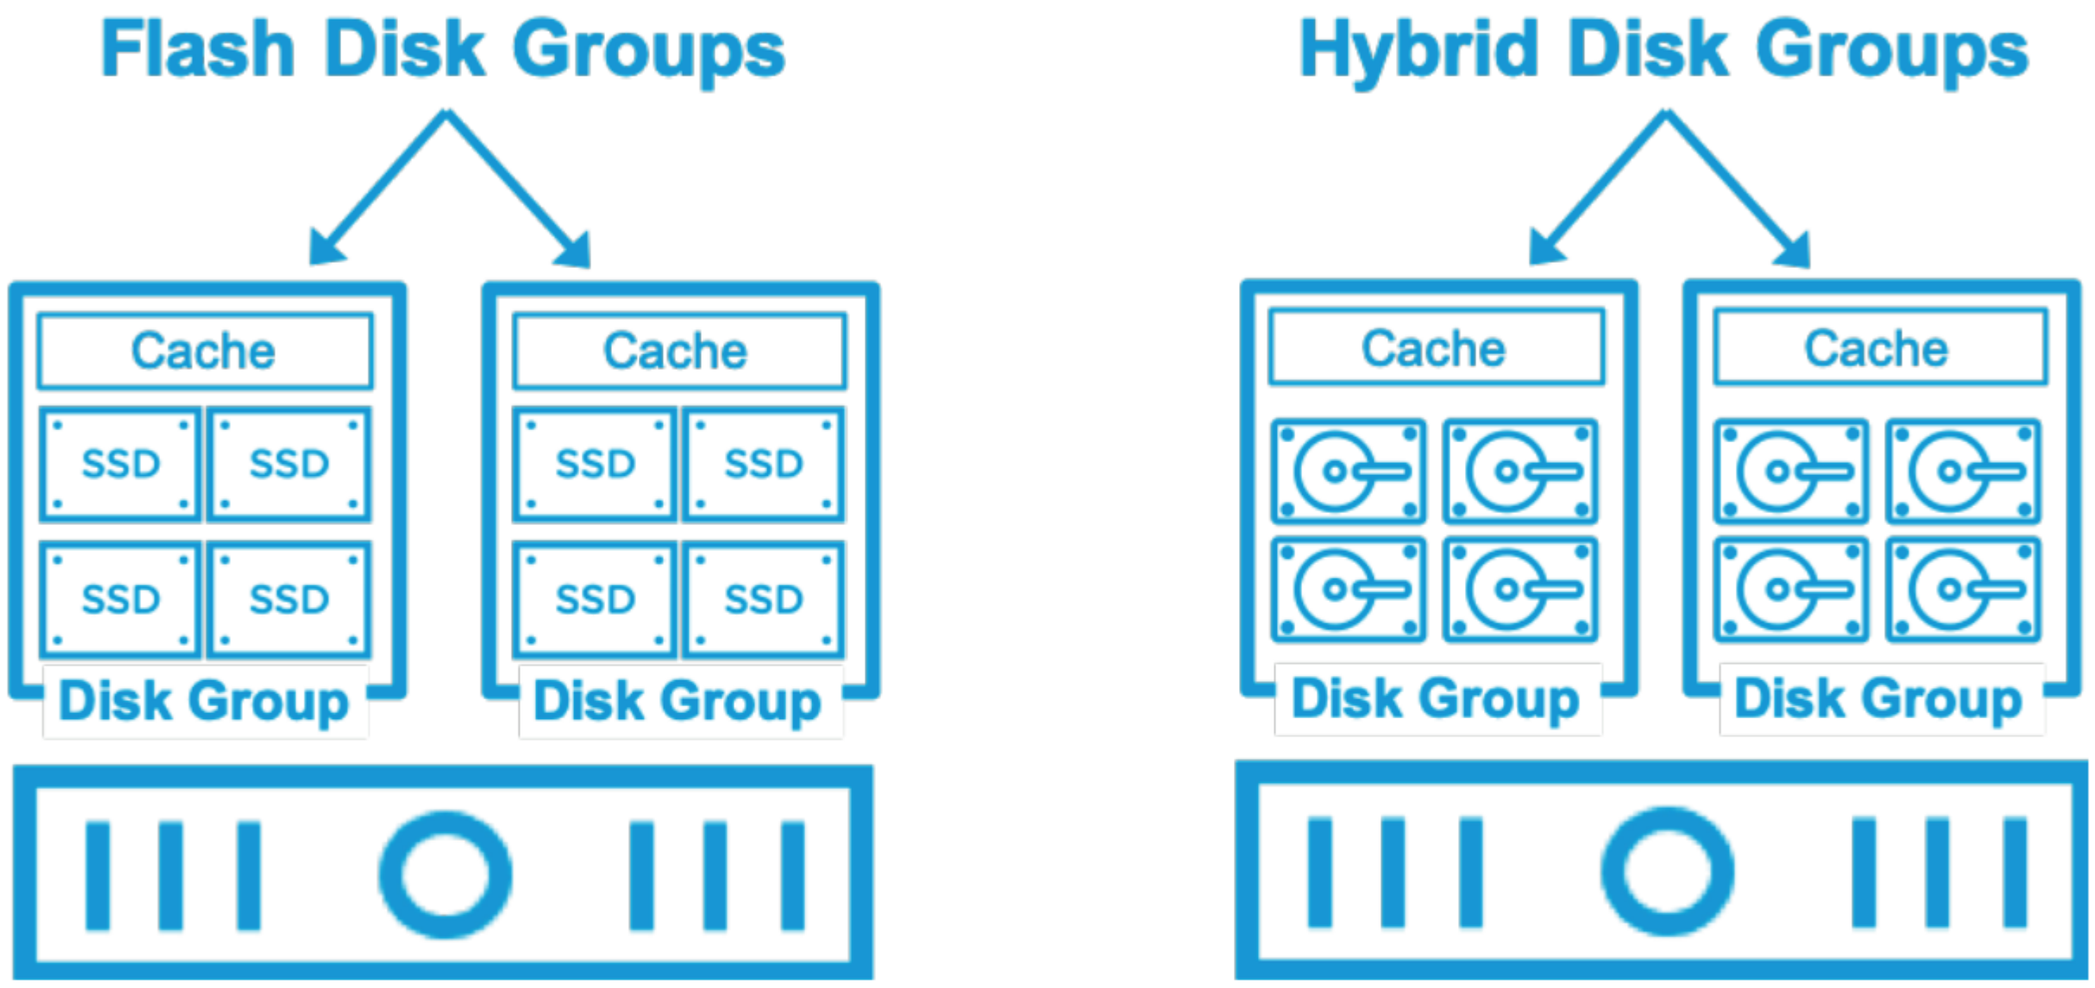

vSAN uses the concept of disk groups to pool together cache devices and capacity devices as single management constructs. A disk group is a pool of one cache device and one to seven capacity devices

5-47 vSAN Hardware Requirements

vSAN capabilities are native to ESXi and require no additional software.

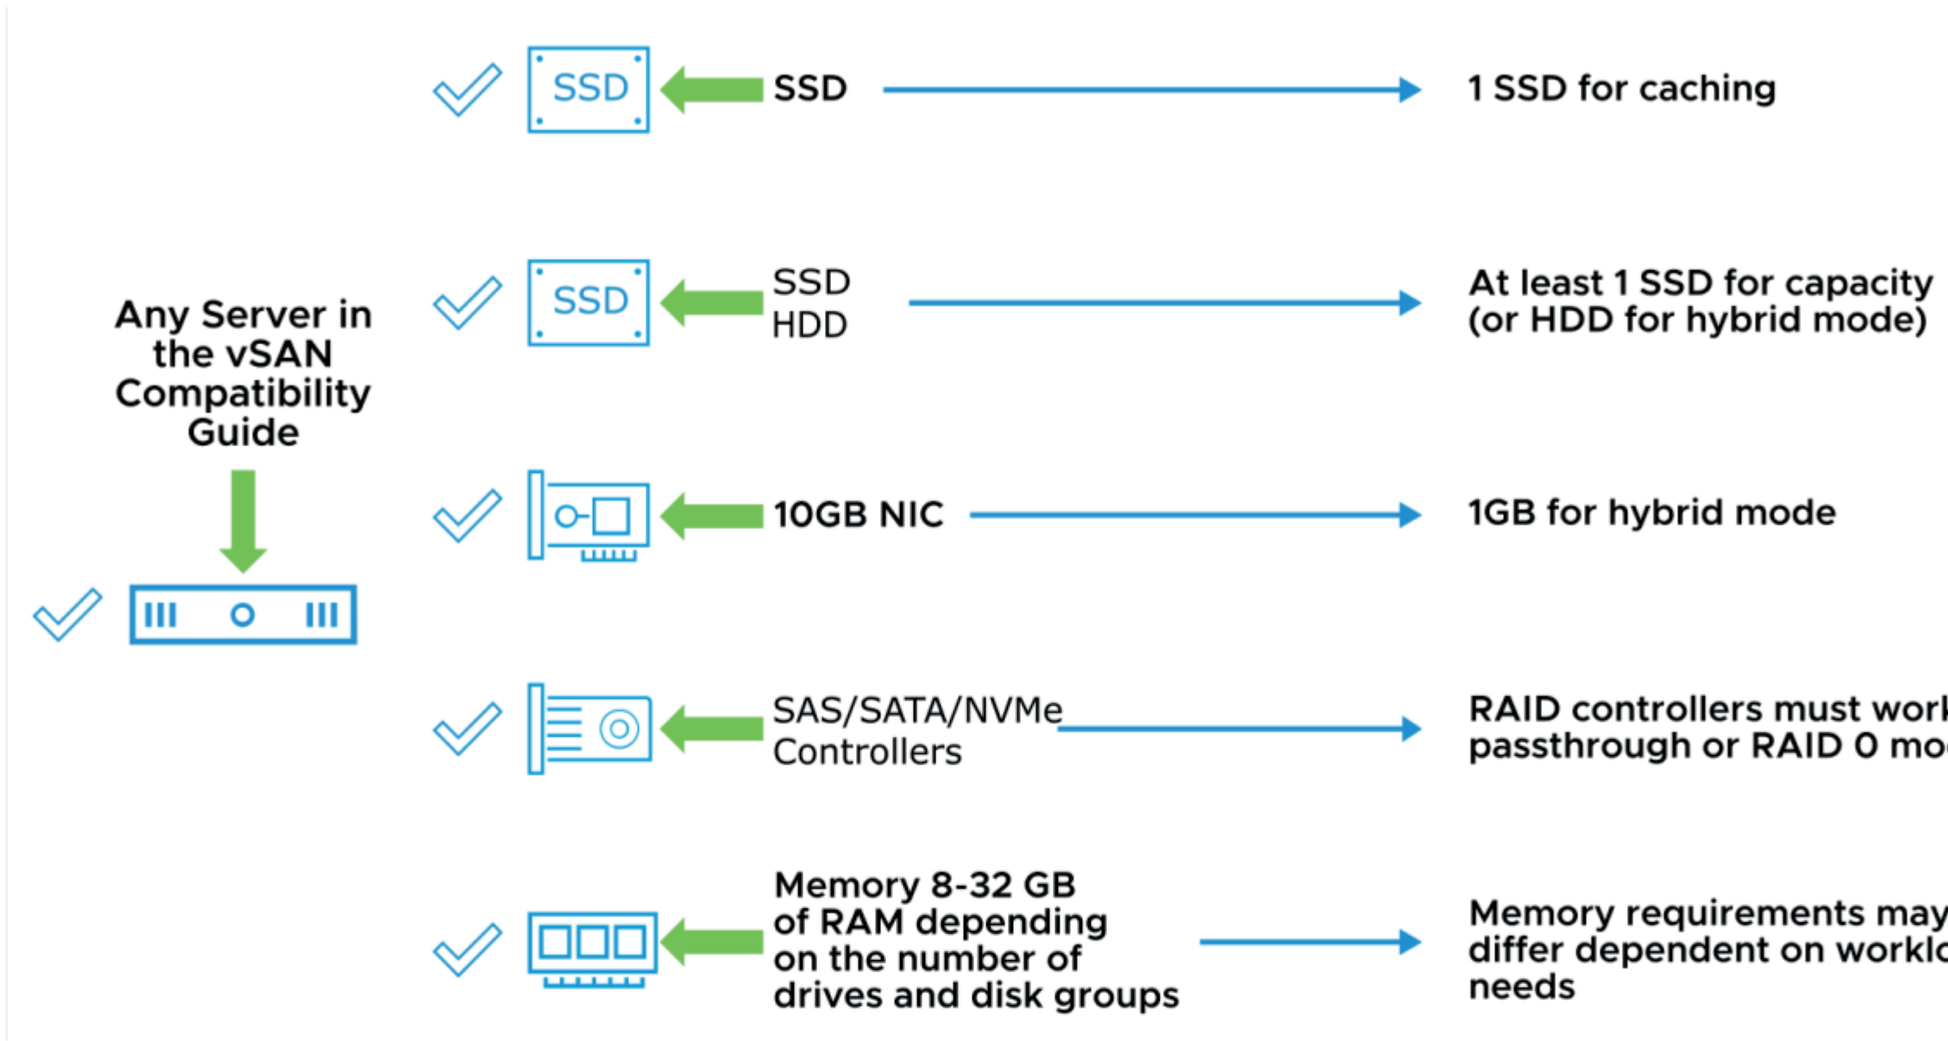

vSAN requires several hardware components that hosts do not normally have:

- One Serial Attached SCSI (SAS), SATA solid-state drive (SSD), or NVMe SSD and one to seven magnetic drives for each hybrid disk group.
- One SAS, or SATA SSD and one to seven flash for all-flash disk groups.
- Dedicated 1 Gbps network (10 Gbps is recommended) for hybrid disk groups.
- Dedicated 10 Gbps network for all-flash disk groups.

  1 Gbps network speeds result in detrimental congestion for an all-flash architecture and are unsupported.
- The vSAN network must be configured for IPv4 or IPv6.

In addition, each host should have a minimum of 32 GB of memory to accommodate a maximum number of five disk groups and a maximum number of seven capacity devices per disk group.

Technet24

5-48 Viewing the vSAN Datastore Summary

The **Summary** tab of the vSAN datastore shows the general vSAN configuration information.

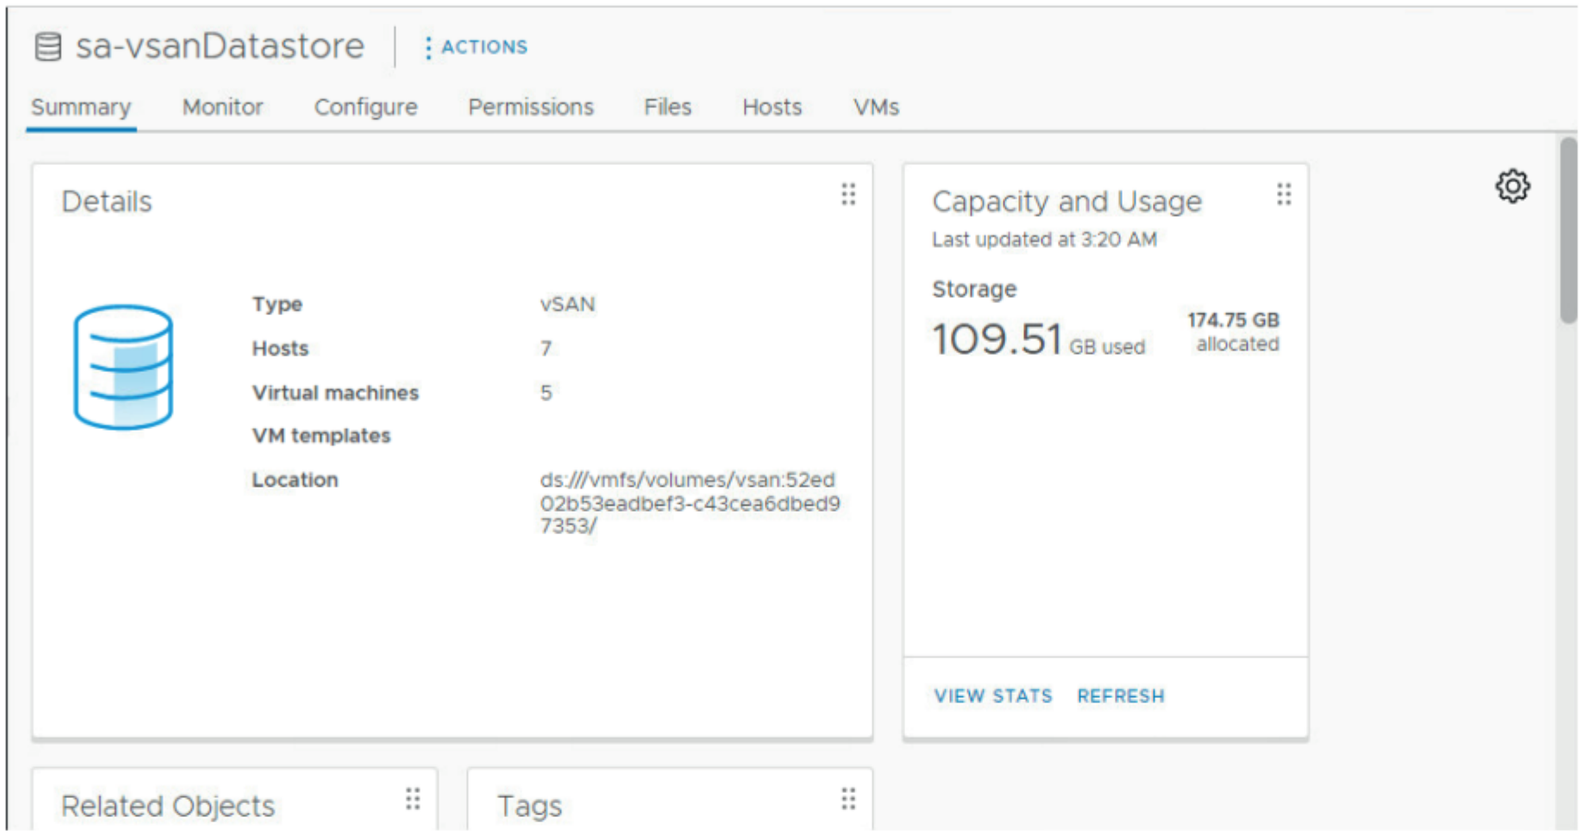

5-49 Objects in vSAN Datastores

vSAN storage is object-based and policy-driven.

VMs created on a vSAN datastore include the following objects:

- One VM home namespace
- One or more VMDK objects
- Thin-provisioned VM swap object
- One or more VM memory objects

Other vSAN objects include:

- vSAN performance data objects
- iSCSI LUN objects

vSAN Object	Traditional VM Files
VM home namespace	`.nvram`, `.vmsd`, `.vmx`, `vmx-*.vswp`, `.log`, `.hlog`
VMDK	`-flat.vmdk`
VM swap	`.vswp`
VM memory	`.vmem`
Snapshot delta	`-00000#-delta.vmdk`, `-00000#-sesparse.vmdk`

A vSAN cluster stores and manages data as flexible data containers called objects. When you provision a VM on a vSAN datastore, a set of objects is created:

- VM home namespace: Stores the virtual machine metadata (configuration files)
- VMDK: Virtual machine disk
- VM swap: Virtual machine swap file, which is created when the VM is powered on
- VM memory: Virtual machine's memory state when a VM is suspended or when a snapshot is taken of a VM and its memory state is preserved
- Snapshot delta: Created when a virtual machine snapshot is taken

Technet24

5-50 VM Storage Policies

Storage policies define how objects that are included in a VM are stored.

Storage policies have the following characteristics:

- Based on storage capabilities
- Defined for a VM or disk at the time of deployment
- Can be applied later
- Can be changed at any time
- Cannot be deleted if they are in use

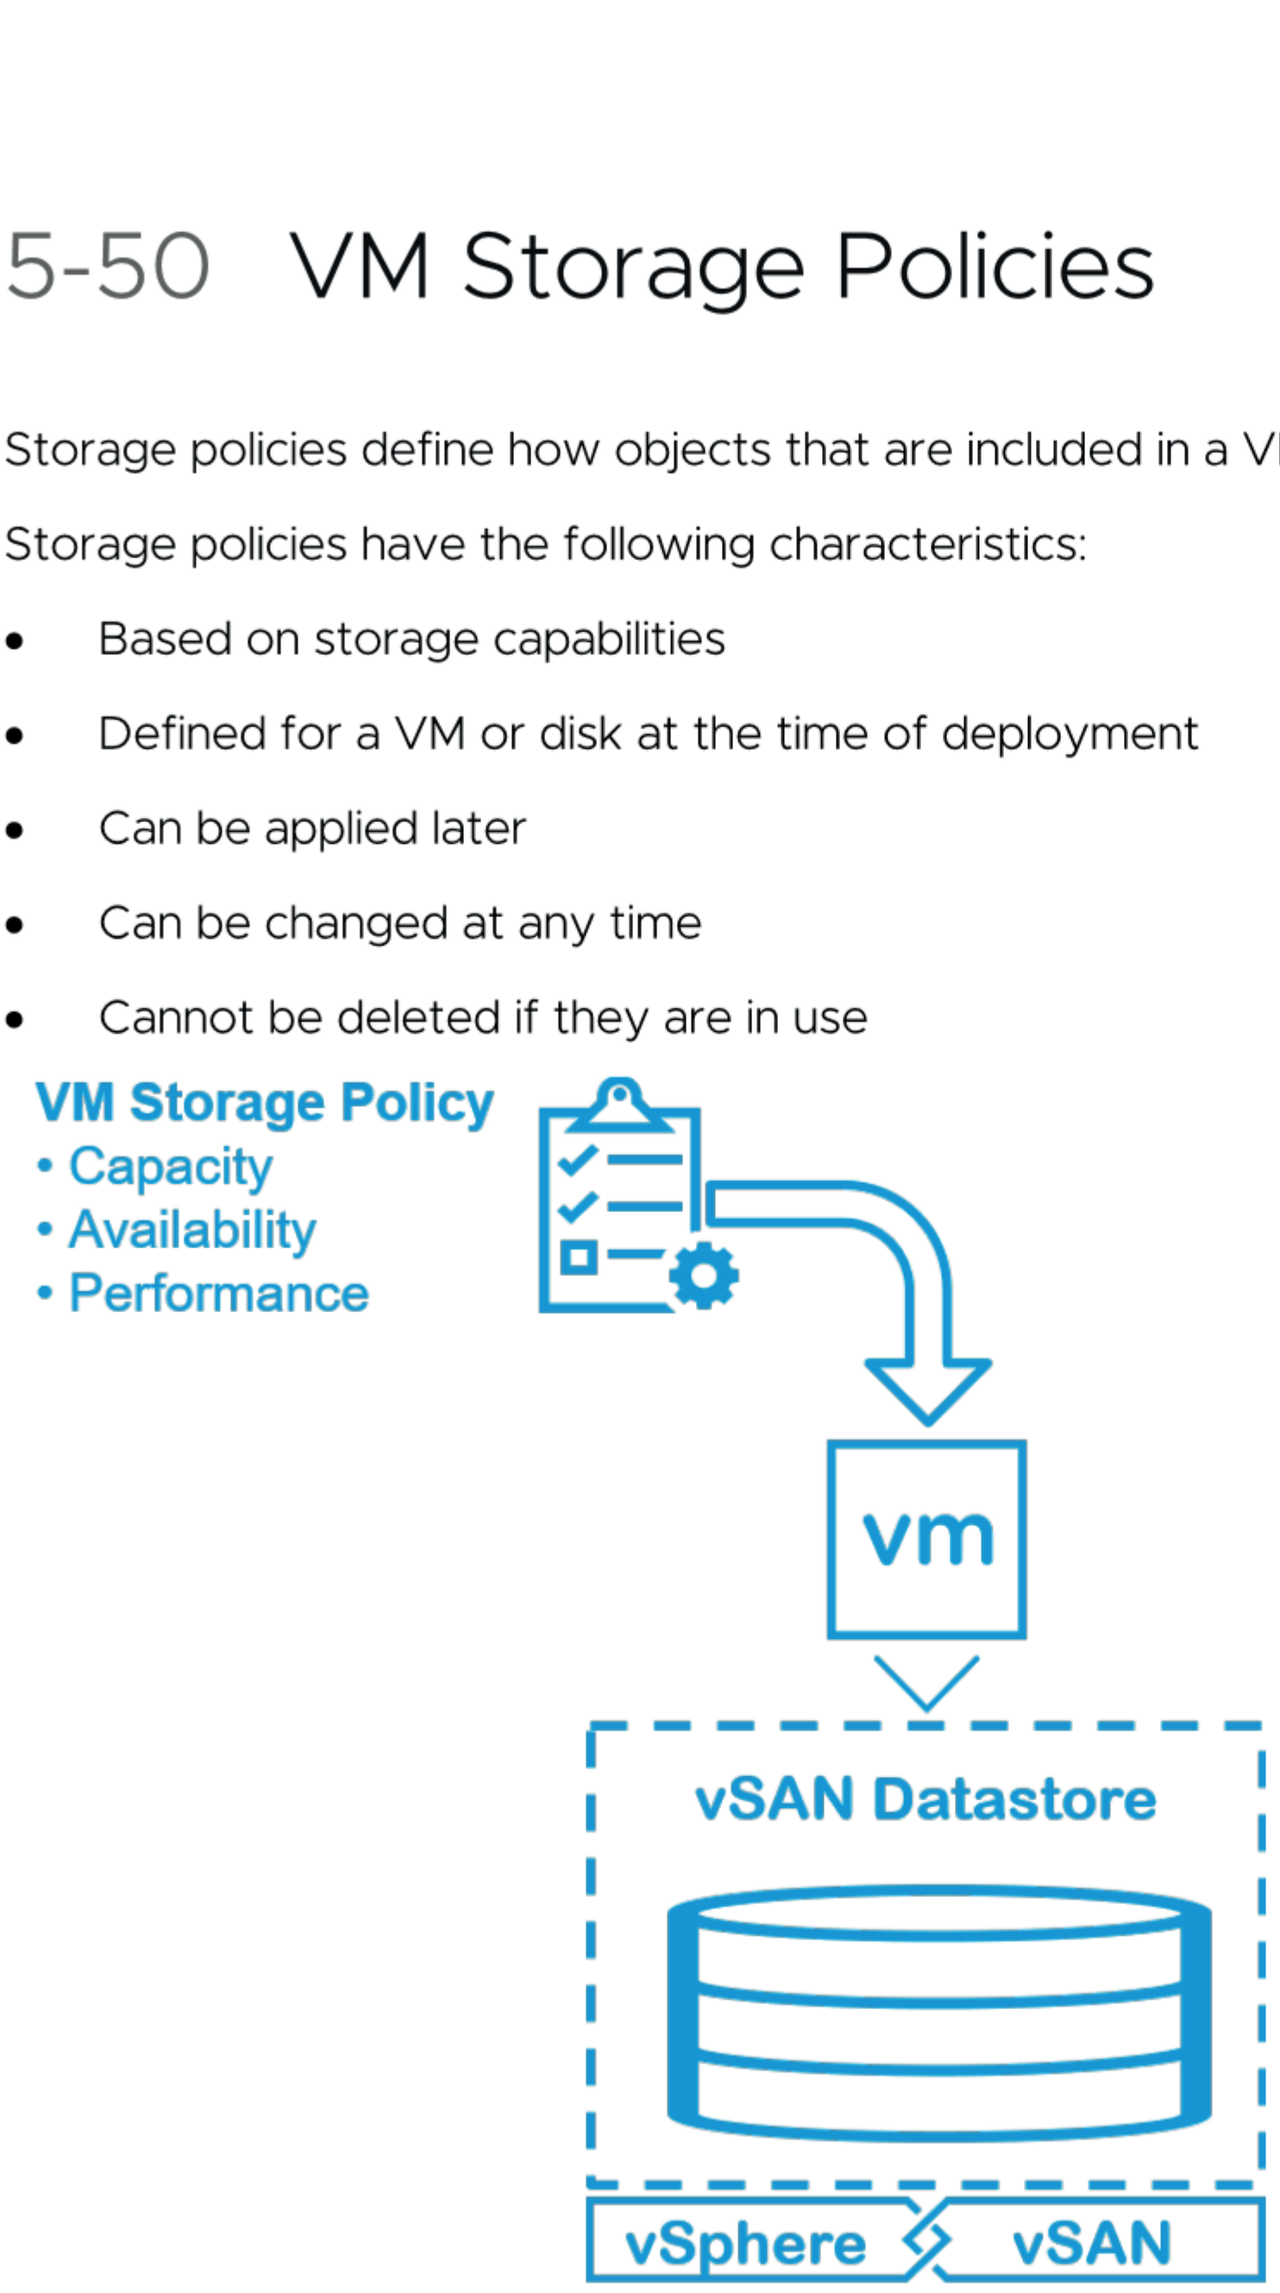

VM storage policies are a set of rules that you configure for VMs. Each storage policy reflects a set of capabilities that meet the availability, performance, and storage requirements of the application or service-level agreement for that VM.

You should create storage policies before deploying the VMs that require these storage policies. You can apply and update storage policies after deployment.

A vSphere administrator who is responsible for the deployment of VMs can select policies that are created based on storage capabilities.

Based on the policy that is selected for the object VM, these capabilities are pushed back to the vSAN datastore. The object is created across ESXi hosts and disk groups to satisfy these policies.

5-51 Viewing VM Settings for vSAN Information

The consumption of vSAN storage is based on the VM's storage policy.

The VM's physical disk placement view provides the following information:

- Component State
- ESXi Host
- The location of disk files on a vSAN datastore

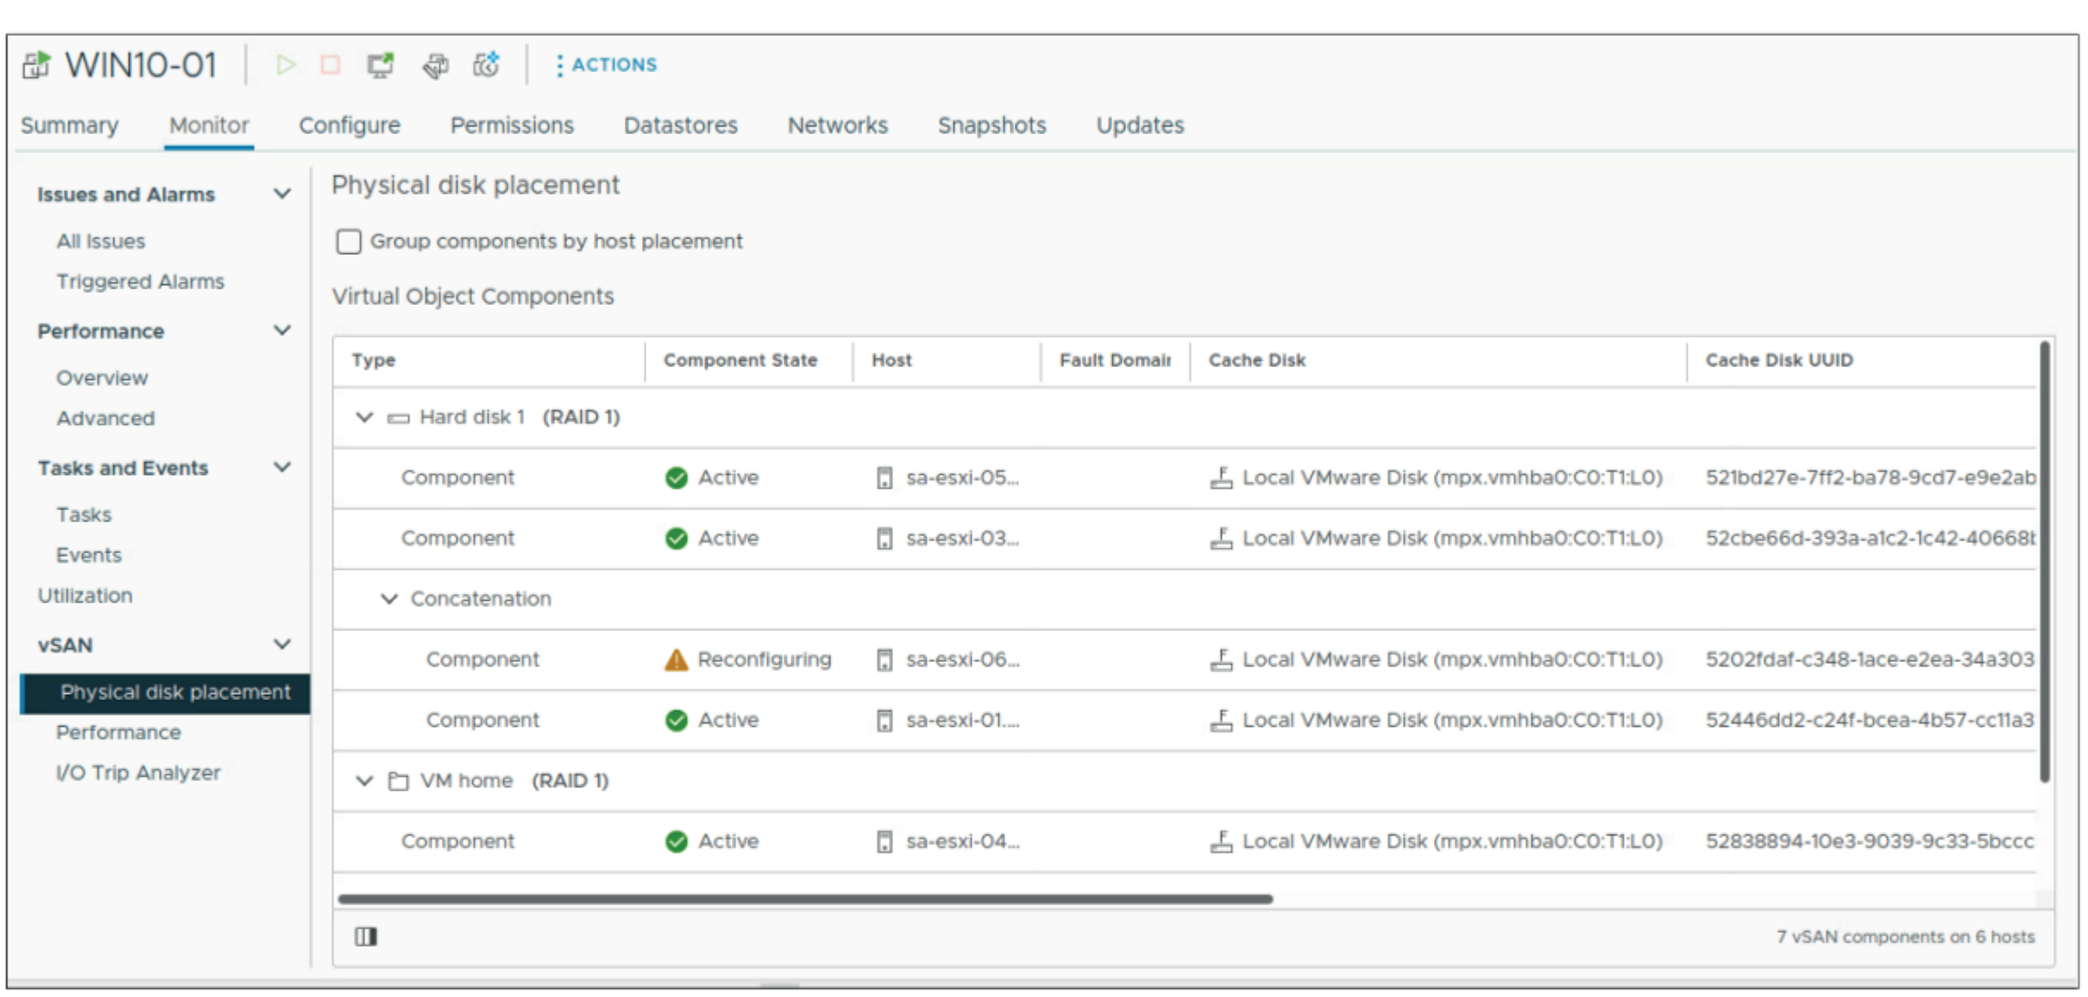

Technet24

5-52 About vSAN Fault Domains

In vSAN, a fault domain represents a set of hardware that shares a single point of failure:

- vSAN can tolerate failures of a single host, entire rack, network switch, and more.
- Protected objects are distributed across different fault domains.

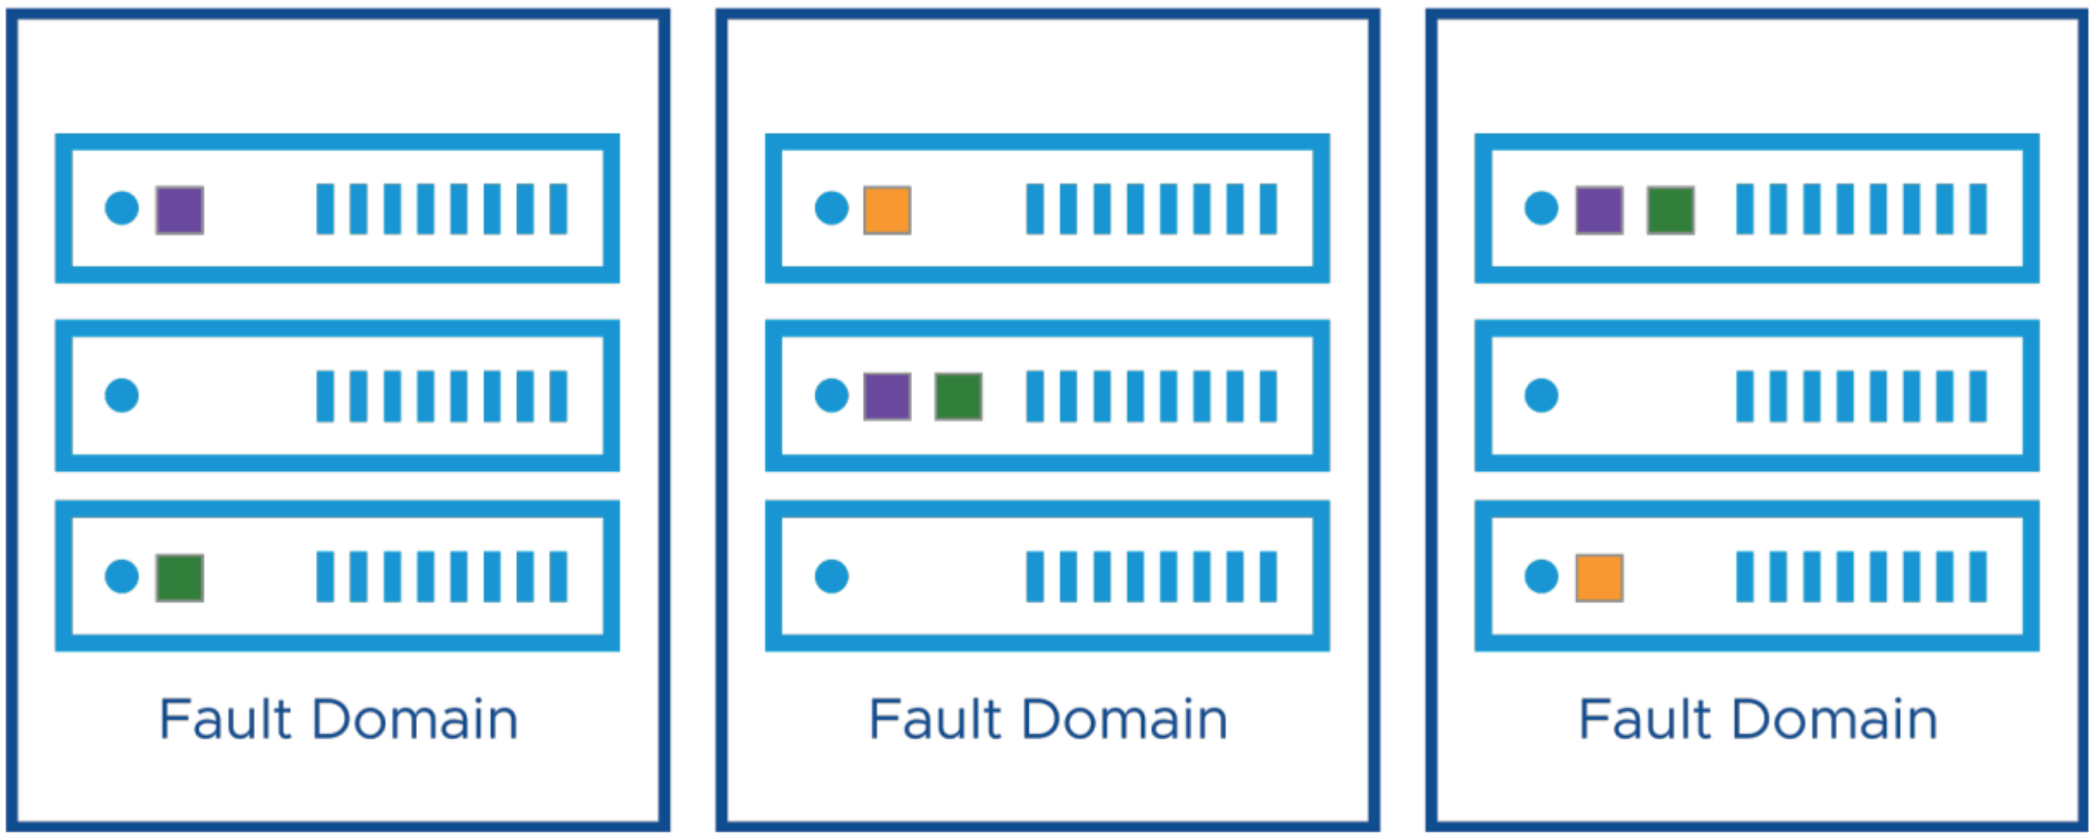

For more information about managing and configuring fault domains in vSAN, see *Administering VMware vSAN* at https://docs.vmware.com/en/VMware-vSphere/8.0/vsan-80-administration-guide.pdf.

5-53 Fault Domain Configurations

vSAN fault domains are configured according to the use case:

- Stretched cluster: vSAN is configured on hosts spanning multiple physical sites.
- Nonstretched cluster: Multiple fault domains are created within a single cluster.
- vSAN two-node configuration: vSAN is configured with two data nodes and a witness node at a separate site.

When performing life cycle operations, vSphere Lifecycle Manager addresses the following user requests:

- Upgrade the vSAN stretched cluster so that hosts from the preferred fault domain are upgraded before hosts from the secondary fault domain.
- Upgrade the vSAN cluster with multiple fault domains so that all the hosts in one fault domain are upgraded first, before moving on to the next fault domain.

5-54 Lab 10: Viewing a vSAN Datastore Configuration

View a vSAN datastore configuration and a virtual machine's components on the vSAN datastore:

1. View a vSAN Datastore Configuration
2. View the vSAN Default Storage Policy
3. View a Virtual Machine on the vSAN Datastore

5-55 Review of Learner Objectives

- Explain the purpose of a vSAN datastore
- Describe the architecture and requirements of vSAN configuration
- Explain the purpose of vSAN storage policies
- Describe the purpose of fault domains

Technet24

5-56 Lesson 5: vSphere Storage APIs - Array Integration

5-57 Learner Objectives

- Describe how vSphere Storage APIs - Array Integration helps storage arrays integrate with vSphere
- Explain how space reclamation works with VMFS datastores

5-58 About vSphere Storage APIs - Array Integration

vSphere Storage APIs - Array Integration helps storage vendors provide hardware assistance to accelerate vSphere I/O operations that are more efficiently accomplished in the storage hardware.

vSphere Storage APIs - Array Integration is a set of protocols and VMkernel functionality that support integration between ESXi and storage arrays:

- With hardware acceleration APIs, arrays can integrate with vSphere to transparently offload certain storage operations to the array.
- Array thin provisioning APIs help prevent out-of-space conditions and perform space reclamation.

vSphere Storage APIs - Array Integration are a family of APIs that are used by third-party hardware, software, and storage vendors to develop components that enhance several vSphere features and solutions.

In a virtualized environment, virtual disks on a VMFS datastore cannot be interpreted by the disk array, so the VMFS datastore cannot use hardware functions per VM or per virtual disk file.

vSphere Storage APIs - Array Integration plug-ins can improve data transfer performance, and they are transparent to the end user.

5-59 Hardware Acceleration APIs for Block Storage

Using hardware acceleration APIs, the ESXi host can offload specific VM and storage management operations to storage hardware:

- Use of these APIs significantly reduces the CPU overhead on the host.
- Block storage devices, such as Fibre Channel and iSCSI devices, support hardware acceleration.
- Hardware acceleration for block storage devices supports several array operations.

Array Operation	Use Cases
Full copy, also called clone blocks, copy offload, or XCOPY	Used by vSphere Storage vMotion, cloning VMs, and deploying a VM from a template
Block zeroing, also called write same	Used when creating eager-zeroed thick virtual disks
Hardware-assisted locking, also called atomic test and set	Improves performance for VMFS metadata changes

With storage hardware assistance, your host performs array operations faster and consumes less CPU, memory, and storage fabric bandwidth.

ESXi hardware acceleration for block storage devices supports the following array operations:

- Full copy: Storage arrays can make full copies of data in the array. The host does not need to read and write the data. This operation reduces the time and network load when cloning VMs, provisioning from a template, or migrating with vSphere Storage vMotion.
- Block zeroing: Storage arrays can zero out several blocks to provide newly allocated storage, free of previously written data. This operation reduces the time and network load when creating VMs and formatting virtual disks. This operation is activated when the Thick Provision Eager Zero option is used for a disk.
- Hardware-assisted locking: This option supports discrete VM locking without use of SCSI reservations. Disk locking is per sector, unlike SCSI reservations, which locks the entire logical unit number (LUN).

The hardware acceleration features are enabled by default and are used on storage arrays that support these features. You can deactivate these features, but it is not recommended. Older arrays, or arrays with outdated firmware, might not support hardware acceleration. In such cases, the host falls back on older and more CPU-intensive techniques to perform the same tasks.

For more information about vSphere Storage APIs - Array Integration, see VMware knowledge base article 1021976 at http://kb.vmware.com/kb/1021976.

Technet24

5-60 Hardware Acceleration for NAS

With hardware acceleration for network-attached storage (NAS), NAS arrays can integrate with vSphere to offload certain storage operations, such as offline cloning, to the array. This integration reduces CPU overhead on the host.

Hardware acceleration for NAS supports the following NAS operations:

- Full file clone: The entire file, instead of file segments, is cloned.
- Reserve space: Space for a virtual disk is allocated in thick format.
- Native snapshot support: VMware Horizon can offload the creation of linked clones to a NAS array.
- Extended file statistics

With hardware acceleration, your host can integrate with NAS devices and use several hardware operations that NAS storage provides:

- Full file clone: This operation is similar to VMFS block cloning, except that NAS devices clone entire files instead of file segments.
- Reserve space: Storage arrays allocate space for a virtual disk file in thick format.

  Typically, when you create a virtual disk on an NFS datastore, the NAS server determines the allocation policy. The default allocation policy on most NAS servers is thin and does not guarantee backing storage to the file. However, the reserve space operation can instruct the NAS device to use vendor-specific mechanisms to reserve space for a virtual disk. As a result, you can create thick virtual disks on the NFS datastore.
- Native snapshot support: This operation enables View to offload the creation of linked clones to a NAS array.
- Extended file statistics: Storage arrays can accurately report space utilization.

Hardware acceleration for NAS is implemented through vendor-specific NAS plug-ins. Vendors typically create and support these plug-ins and distribute them as vSphere installation bundles (VIBs) through a webpage.

5-61 Array Thin Provisioning APIs

With array thin provisioning APIs, a host can integrate with physical storage and detect space use in thin-provisioned LUNs:

- A VMFS datastore that you deploy on a thin-provisioned LUN can detect only the logical size of the LUN.

  For example, if an array reports 2 TB of storage but provides only 1 TB, the datastore considers 2 TB to be the LUN's size.

- Using thin-provisioning integration, the host can perform these tasks:
  - Monitor the use of space on thin-provisioned LUNs to avoid running out of physical space.
  - Inform the array about datastore space that is freed when the following actions occur:
    - Virtual disks and VM files are deleted from the datastore.
    - Virtual disks and VM files are migrated off the datastore by vSphere Storage vMotion.
    - VM snapshots are deleted and consolidated.

Traditional LUNs that arrays present to the ESXi host are thick-provisioned. The entire physical space that is required to back each LUN is allocated in advance.

ESXi also supports thin-provisioned LUNs. When a LUN is thin-provisioned, the storage array reports the LUN logical size, which might be larger than the real physical capacity backing that LUN.

When the datastore expands, or when vSphere Storage vMotion is used to migrate VMs to a thin-provisioned LUN, the array thin provisioning APIs enable the host to communicate with the LUN. The host can warn about breaches in physical space or about out-of-space conditions

No installation steps are required for the array thin provisioning extensions. Array thin provisioning works on all VMFS5 and VMFS6 volumes. Device firmware that is enabled for this API must use the array thin provisioning features. ESXi continuously checks for firmware that is compatible with array thin provisioning. After firmware is upgraded, ESXi starts using the array thin provisioning features.

Technet24

5-62 About Space Reclamation (1)

In a virtualized environment, different levels of abstraction exist at the storage layer:

- User files reside at the guest operating system layer. The guest operating system manages these files.
- The guest OS resides in the VMDK file, which the VMFS datastore manages.
- The VMFS datastore resides on the LUN, which is presented as the storage device to the ESXi host.

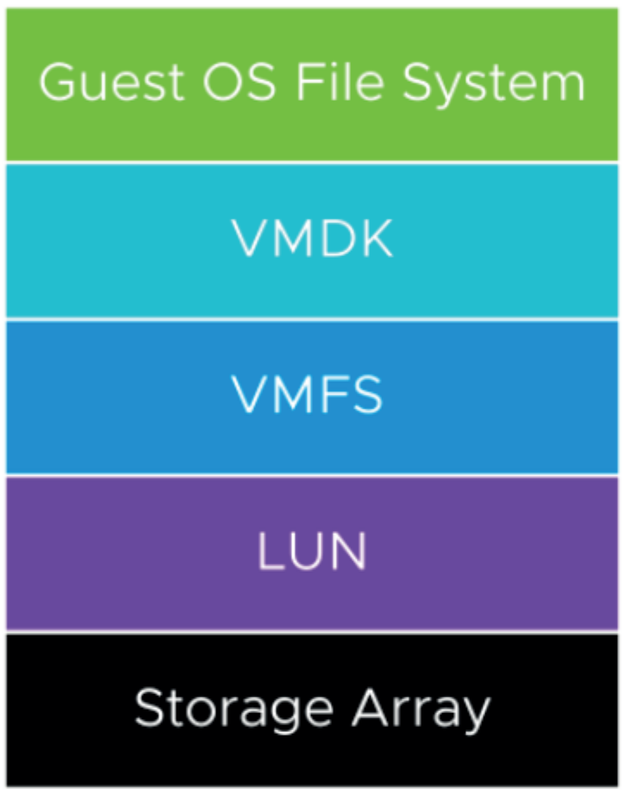

5-63 About Space Reclamation (2)

Data is deleted from the VMFS datastore in these ways:

- Data is deleted within the guest OS.
- A VMDK file is deleted from the datastore.

In either case, the space on the array is not freed.

This space is called dead space or dirty space.

The array cannot reuse this space unless it is told to do so, such as by the `SCSI UNMAP` command.

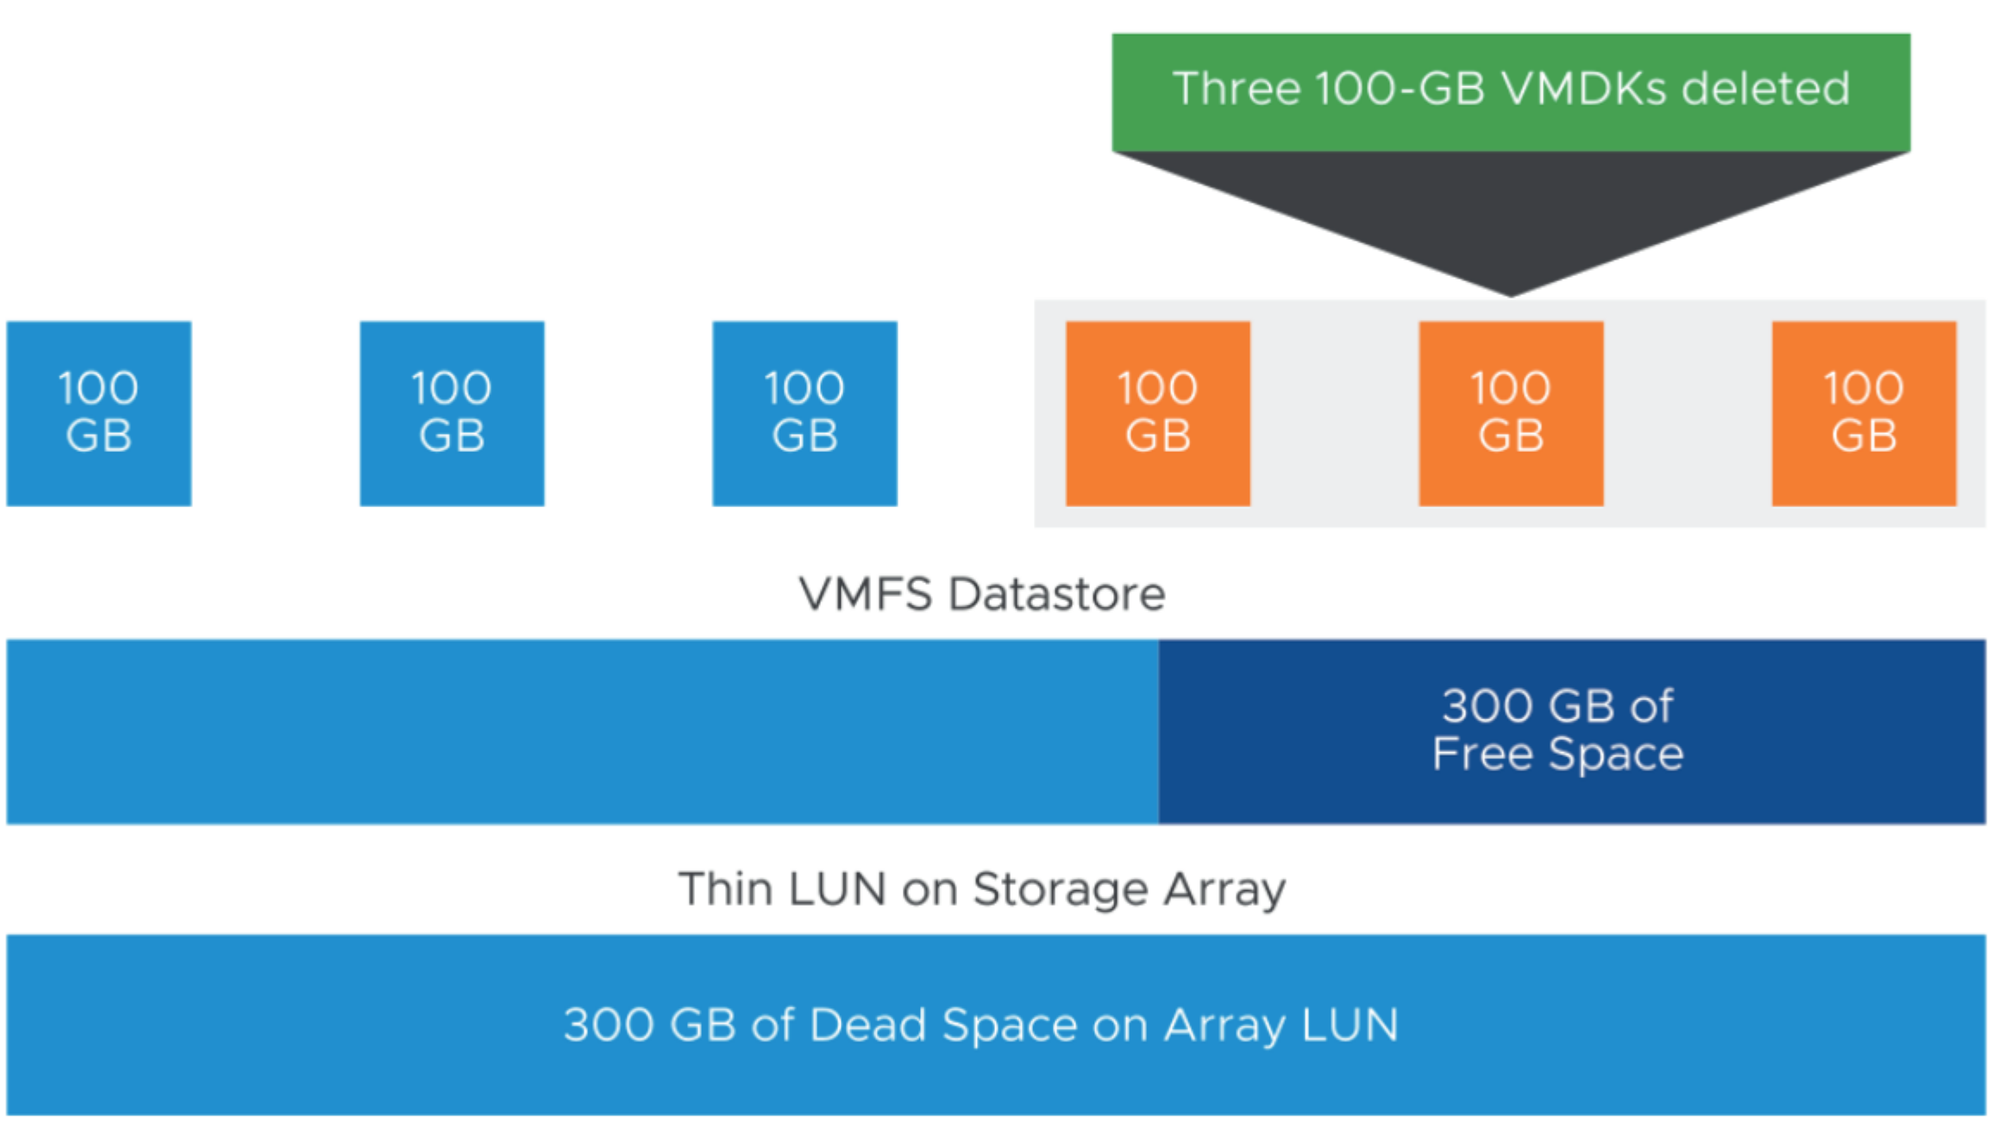

Deleting or removing files from a VMFS datastore frees space within the file system. This free space is mapped to a storage device until the file system releases or unmaps it.

Technet24

5-64 Space Reclamation with VMFS

ESXi supports the reclamation of free space. When you run the SCSI `unmap` command, an ESXi host informs the storage array to reclaim free space:

- With a VMFS5 datastore, you can manually send the `unmap` command with the `esxcli` command.
- With a VMFS6 datastore, the ESXi host automatically sends the `unmap` command to the storage array.

Feature	VMFS5	VMFS6
Automatic space reclamation	No	Yes
Manual space reclamation through ESXCLI	Yes	Yes

When you run the unmap command, an ESXi host informs the storage array that the VM files are removed or deleted from a thin-provisioned VMFS datastore. At this notification, the array can reclaim the freed blocks.

5-65 Performing Manual Space Reclamation on a VMFS5 Datastore

To manually reclaim space on a VMFS5 datastore, you can run the `esxcli` command:

◊ **`esxcli storage vmfs unmap`**

You can automate the process by using this command in a PowerCLI script and scheduling the script to run during off-hours.

Command Option	Description
-l | --volume-label=*volume_label*	The label of the VMFS volume to unmap. A mandatory argument. Use either the –l or the –u option, not both.
-u | --volume-uuid=*volume_uuid*	The UUID of the VMFS volume to unmap. A mandatory argument. Use either the –l or the –u option, not both.
-n | --reclaim-unit=*number*	Number of VMFS blocks to unmap per iteration. An optional argument. The default value is 200.

On VMFS5 datastores, you must run the `esxcli storage vmfs unmap` command manually. The target server prompts you for a user name and password. Other connection options, such as a configuration file and a session file, are supported.

This command can be run without a maintenance window:

- You can specify the reclaim size in blocks, instead of a percentage value, to make the reclaim size more intuitive to calculate.
- Dead space is reclaimed in increments, instead of all at once, to avoid possible performance issues.

When you run the command, it might send many `unmap` requests at a time, which can lock some of the resources during this operation.

For more information about this command, see vSphere Storage at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-storage-guide.pdf.

Technet24

5-66 Enabling Automatic Space Reclamation on a VMFS6 Datastore

When you create a VMFS6 datastore, you can modify the default settings for automatic asynchronous space reclamation:

- **Space Reclamation Granularity**: Defines the unit size for the `UNMAP` command. The default is 1 MB.
- **Space Reclamation Priority**: Controls the rate at which deleted or unmapped blocks are reclaimed on the LUNs backing the datastore. The default priority is Low, which is an unmap rate of 25 MBps.

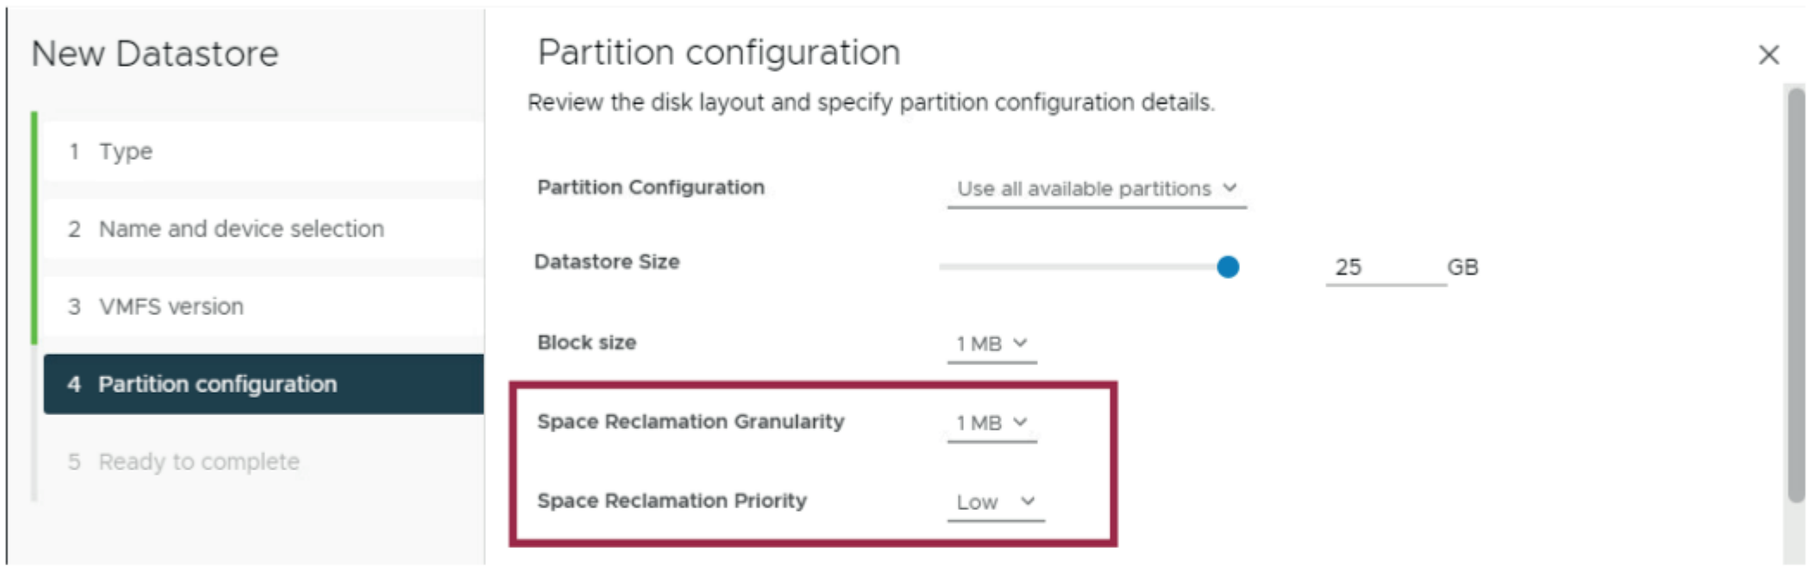

By default, the space reclamation granularity equals the block size, which is 1 MB. Storage sectors smaller than 1 MB are not reclaimed.

By default, the LUN performs the space reclamation operation at a low rate. You can also set the space reclamation priority to None to deactivate the operation for the datastore.

5-67 Configuring the Automatic Unmap Rate

Introduced in vSphere 6.7, you can set a fixed unmap rate for the datastore. The unmap rate is the rate at which automatic `unmap` commands are processed.

You can change the unmap rate to match the storage array's capabilities.

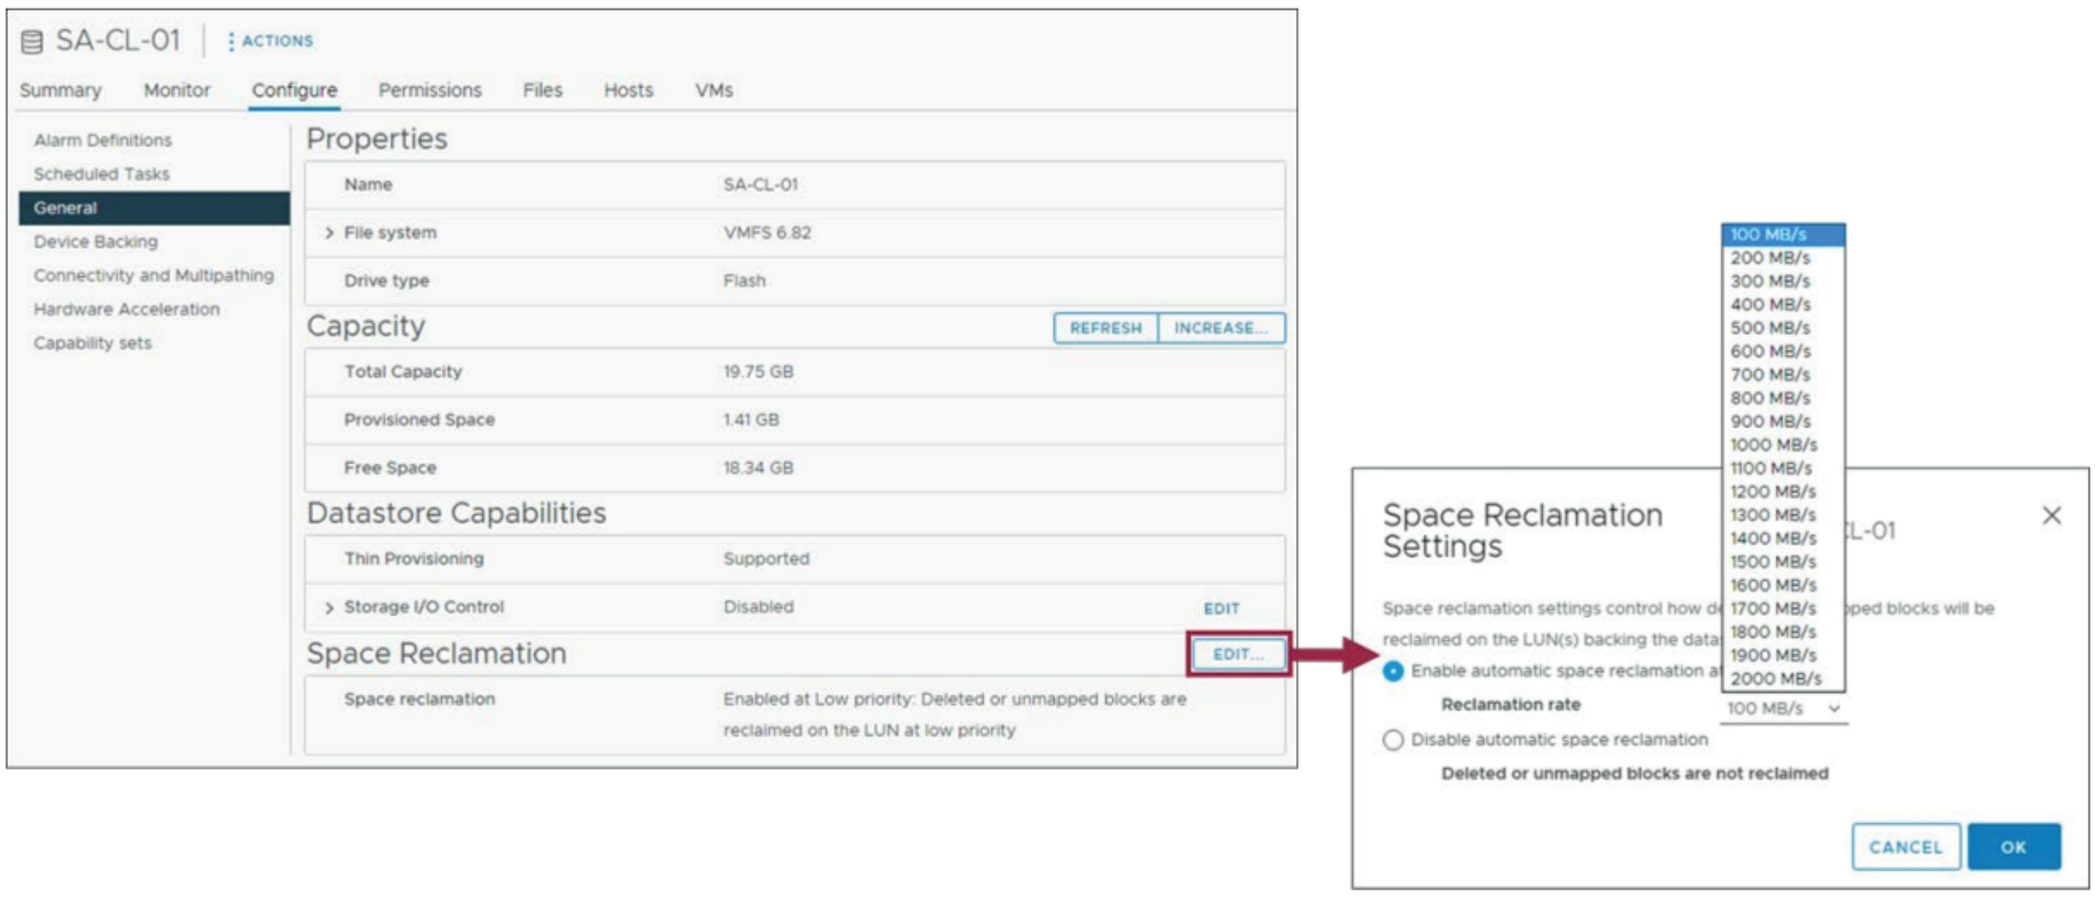

5-68 Space Reclamation Support for Guest Operating Systems (1)

Both VMFS5 and VMFS6 datastores provide support to the `unmap` command that originates from the guest operating system.

Some guest operating systems, such as Windows Server 2012 R2, can send `unmap` commands to the array.

The guest `unmap` commands are translated down to the array and space that can be reclaimed.

For thin-provisioned VMDK files, the virtual disk space decreases by the amount of space reclaimed.

ESXi supports the `unmap` commands that are issued directly from a guest operating system to reclaim storage space.

Technet24

5-69 Space Reclamation Support for Guest Operating Systems (2)

VMFS5 supports the automatic space reclamation requests for a limited number of guest operating systems.

To send unmap requests from the guest OS to the array, the VM must meet the following prerequisites:

- The virtual disk must be thin-provisioned.
- The guest OS must be able to identify the virtual disk as thin-provisioned.
- The EnableBlockDelete advanced ESXi host setting must be set to 1.

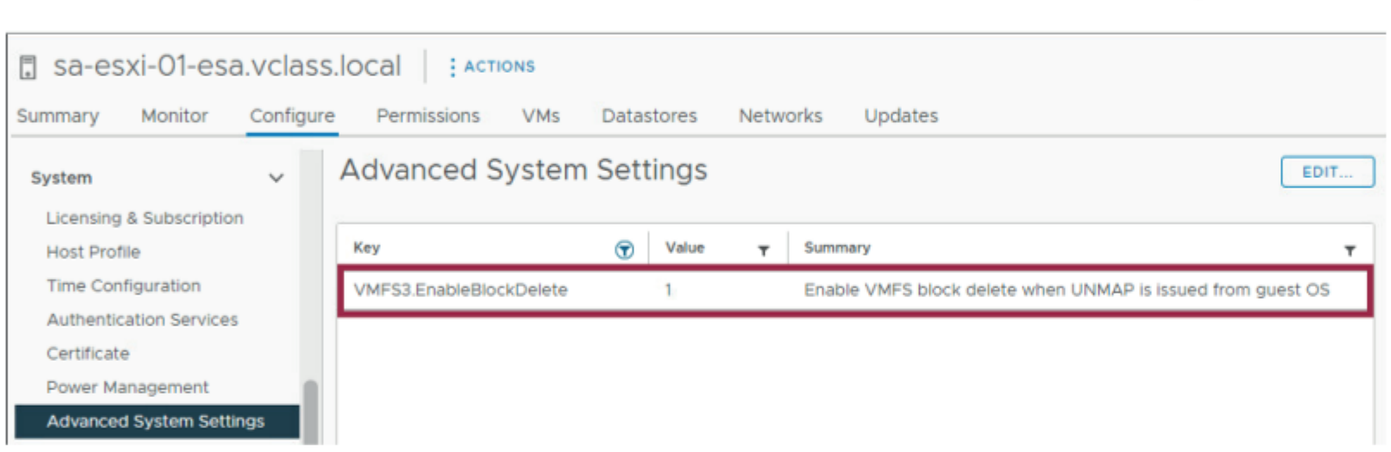

5-70 Space Reclamation Support for Guest Operating Systems (3)

With VMFS6, use the **Space Reclamation Priority** setting on the datastore. The **EnableBlockDelete** option is ignored.

When you use space reclamation with VMFS6 datastores, the following consideration applies:

- VMFS6 processes the unmap request from the guest OS only when the space to reclaim equals 1 MB or is a multiple of 1 MB.

VMFS6 generally supports automatic space reclamation requests from the guest operating systems and passes these requests to the array.

Many guest operating systems can send the `unmap` command and do not require any additional configuration.

The guest operating systems that do not support automatic unmaps might require user intervention.

For information about guest operating systems that support automatic space reclamation on VMFS6, contact your vendor.

5-71 Extended Copy (XCOPY)

Introduced in vSphere 8 Update 1, ESXi now supports Extended Copy to Datastores Across Different Storage Arrays.

XCOPY optimizes the data copy between Datastores across different arrays.

Administrators will now have the ability to offload, migrate and clone workloads to their storage arrays.

While vSphere 8 U1 enables this feature, the actual data migration across the arrays must be supported on the storage array side.

5-72 Activity: vSphere Storage APIs – Array Integration (1)

Which actions apply to an ESXi host that uses storage devices that support vSphere Storage APIs – Array Integration? Select all that apply.

◊ ❒ Inform the storage array that it can reclaim free space in a VMFS datastore.

◊ ❒ Offload certain storage operations, such as offline cloning, to a NAS array.

◊ ❒ Retrieve information about a storage array, such as its storage topology and capabilities.

◊ ❒ Offload specific storage operations to a block storage device, such as Fibre Channel and iSCSI.

Technet24

5-73 Activity: vSphere Storage APIs – Array Integration (2)

Which actions apply to an ESXi host that uses storage devices that support vSphere Storage APIs – Array Integration? Select all that apply.

◊ ✓ Inform the storage array that it can reclaim free space in a VMFS datastore.

◊ ✓ Offload certain storage operations, such as offline cloning, to a NAS array.

◊ ❒ Retrieve information about a storage array, such as its storage topology and capabilities.

◊ ✓ Offload specific storage operations to a block storage device, such as Fibre Channel and iSCSI.

5-74 Review of Learner Objectives

- Describe how vSphere Storage APIs - Array Integration helps storage arrays integrate with vSphere
- Explain how space reclamation works with VMFS datastores

5-75 Lesson 6: vSphere Storage APIs for Storage Awareness and vSphere APIs for I/O Filtering (VASA and VAIO)

5-76 Learner Objectives

- Describe how vSphere API for Storage Awareness ensures that a VM's storage requirements are met
- Explain how vendors can create data services for VMs by using vSphere APIs for I/O Filtering

Technet24

5-77 About vSphere API for Storage Awareness

With vSphere API for Storage Awareness, a storage vendor can develop a software component, called a storage provider, for its storage arrays.

Through this API, a storage provider obtains information about available storage topology, capabilities, and state from the storage array.

vCenter connects to a storage provider. Information from this storage provider appears in the vSphere Client.

Navigate to **vCenter** > **Configure** tab >**Storage Providers**

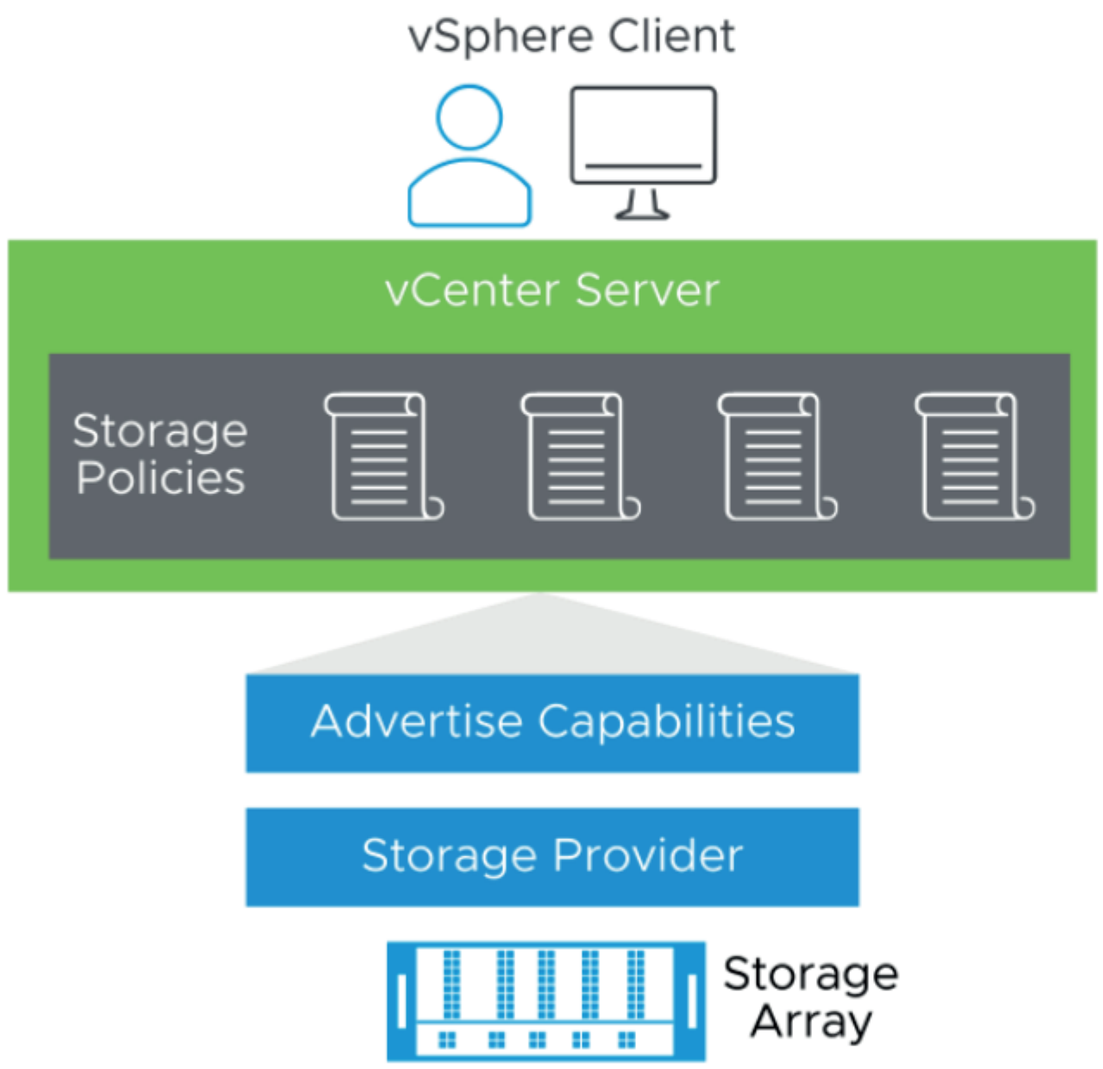

In vCenter, vSphere administrators cannot access the storage capabilities of the LUN on which their VMs are stored. VMs are provisioned to a storage black box. vSphere administrators see only a LUN identifier, such as a Network Address Authority ID (NAA ID) or a T10 identifier.

A storage vendor can use vSphere API for Storage Awareness to provide information about the LUN.

The storage provider is written by the storage array vendor. The storage provider can exist on either the storage array processor or on a virtual appliance. The storage vendor decides.

vCenter connects to the provider to obtain information about the LUN's topology, capabilities, and state. The information is viewed in the vSphere Client.

A storage provider can report information about one or more storage devices. It can support connections to a single vCenter instance or to multiple vCenter instances.

5-78 Benefits of Storage Providers

Storage providers give vSphere administrators information about the topology, capabilities, and state of the LUNs for their VMs. Administrators can use this information in the following ways:

- To monitor the health and use of their LUNs
- To configure VM storage policies with the correct storage in terms of space, performance, and service-level agreement requirements

A storage provider provides capability information about the storage configuration, status, and storage data services in your environment. vSphere Storage DRS uses the information supplied by a storage provider to make VM placement and migration decisions that are compatible with the storage system.

A storage provider provides capability information:

- Storage capabilities: Information about storage characteristics and services, such as the number and type of spindles for a LUN, the I/O operations in megabytes per second, the type of compression that is used, and whether thick-provisioned format is used.
- Storage status: Status of storage entities and alarms and events for configuration changes.

5-79 Storage Provider Categories (1)

In general, vCenter and ESXi use storage providers to obtain information about storage configuration, status, and storage data services in your environment.

Persistence storage providers:

- Manage arrays and storage abstractions
- Can provide data services, such as array-based replication
- Examples:
  - vSAN storage provider
  - vSphere Virtual Volumes storage provider

Data Service Providers

- I/O filter storage providers, or data service providers:
- Examples:
  - Host-based caching
  - Compression
  - Encryption

Technet24

5-80 Storage Provider Categories (2)

Both persistence storage and data service providers can belong to one of these categories.

Built-in Storage Providers:

- Offered by VMware
- Typically, they do not require registration
- Examples:
  - Storage providers that support vSAN or I/O filters are build-in and become registered automatically

Third-party storage providers:

- Offered by a third party and typically require registration
- Examples:
  - vSphere Virtual Volumes storage provider
  - vSphere APIs for Storage Awareness storage providers

5-81 Registering a Storage Provider

In the vSphere Client, select **Configure** > **Storage Providers** to register and manage a storage provider.

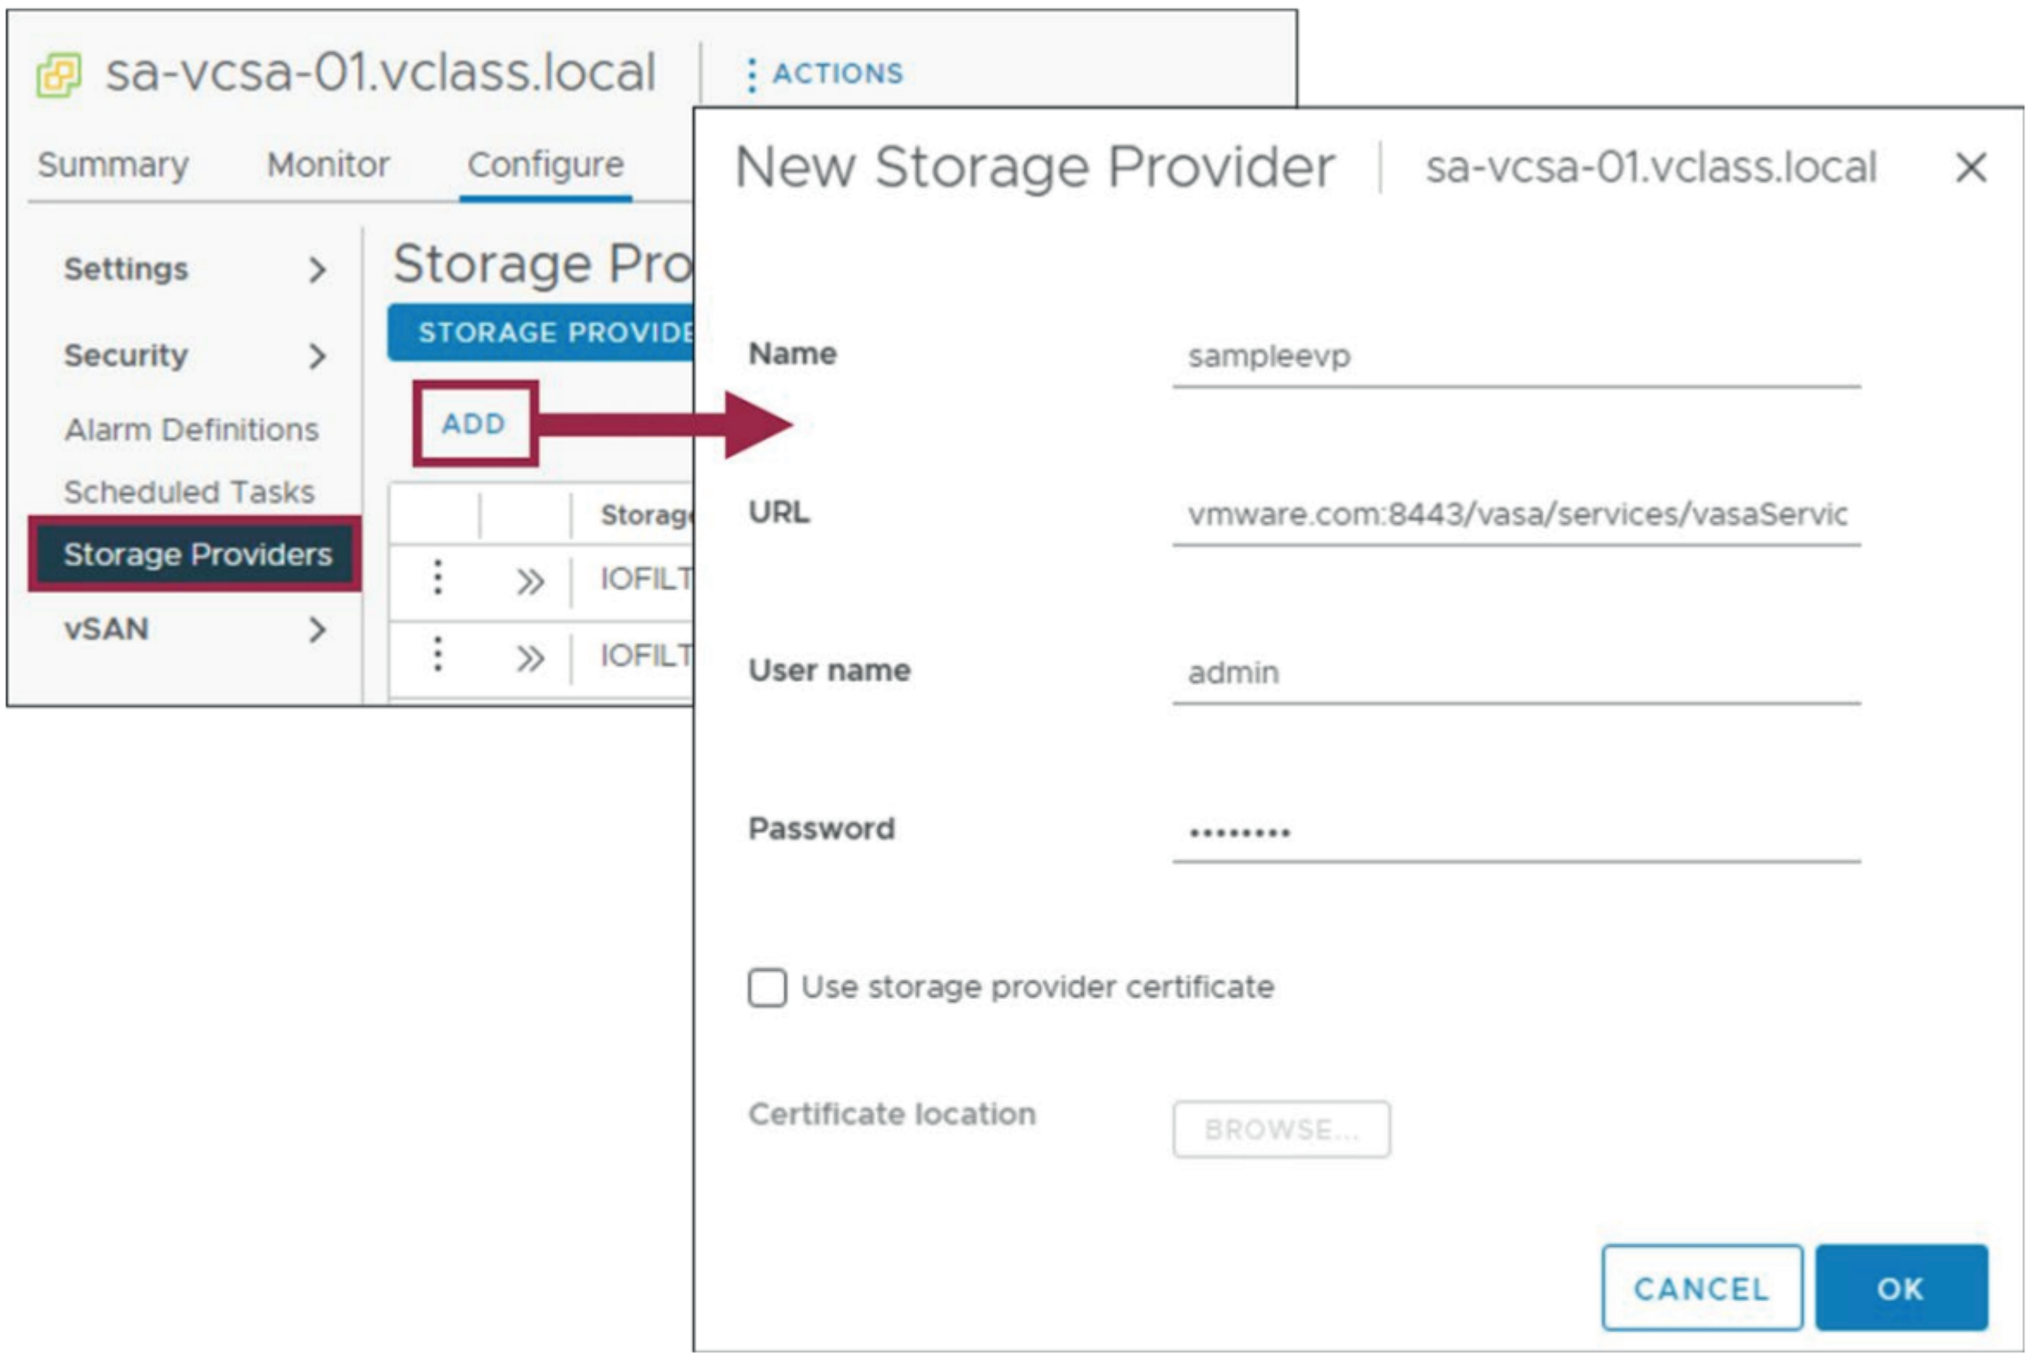

From the Storage Providers pane, you can register a storage provider. All system storage capabilities that are presented by the storage provider appear in the vSphere Client.

To register a storage provider, the storage vendor provides a URL, a login account, and a password.

Users log in to the provider to get array information. vCenter must trust the provider host, so a security certificate from the provider must be installed on the vCenter system.

For information about how to *Register Storage Providers,* see https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-storage/GUID-B50E19B8-C1AD-4FE1-AD90-4DD6D29831FE.html.

Technet24

5-82 About vSphere APIs for I/O Filtering

vSphere APIs for I/O Filtering (VAIO) provide a framework for creating and implementing I/O filters in a VM data stream:

- I/O filters intercept and modify VM disk I/O in the VMkernel before the data is in transit too/from the physical storage.
- I/O filters are independent of storage topology.

With VAIO, VMware and third-party vendors can create data services such as caching and replication.

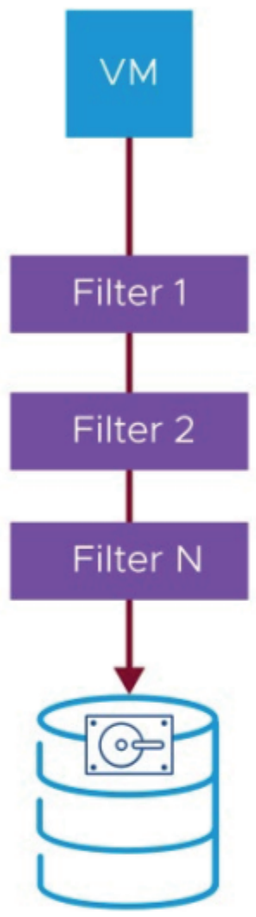

I/O filters can gain direct access to the VM I/O path. You can enable the I/O filter for an individual virtual disk.

VMware offers certain categories of I/O filters. In addition, third-party vendors can create the I/O filters. Typically, these I/O filters are distributed as packages that provide an installer for deploying the filter components on vCenter and ESXi host clusters.

I/O filters can support all datastore types:

- VMFS
- NFS 3
- NFS 4.1
- vSphere Virtual Volumes
- vSAN

5-83 Types of I/O Filters

The types of I/O filters that can be applied to VMs are grouped into filter categories.

Category	Examples
Replication	Replicates all write I/O operations to an external target location, such as another host or cluster
Encryption	Provides encryption mechanisms for VMs
Caching	Implements a cache for virtual disk data
Storage I/O Control	Prioritizes storage I/O that is allocated to VMs during periods of I/O contention

VMware provides certain categories of I/O filters that are installed on your ESXi hosts. In addition, VMware partners can create the I/O filters through the vSphere APIs for I/O Filtering developer program. The I/O filters can serve multiple purposes.

Technet24

5-84 How I/O Filters Work (1)

- The I/O filter framework within the VMkernel is responsible for intercepting I/O and applying the filters associated with a VM.
- I/O filters are run in the VM user world to ensure that any filter problems affect the VM and not the entire host.

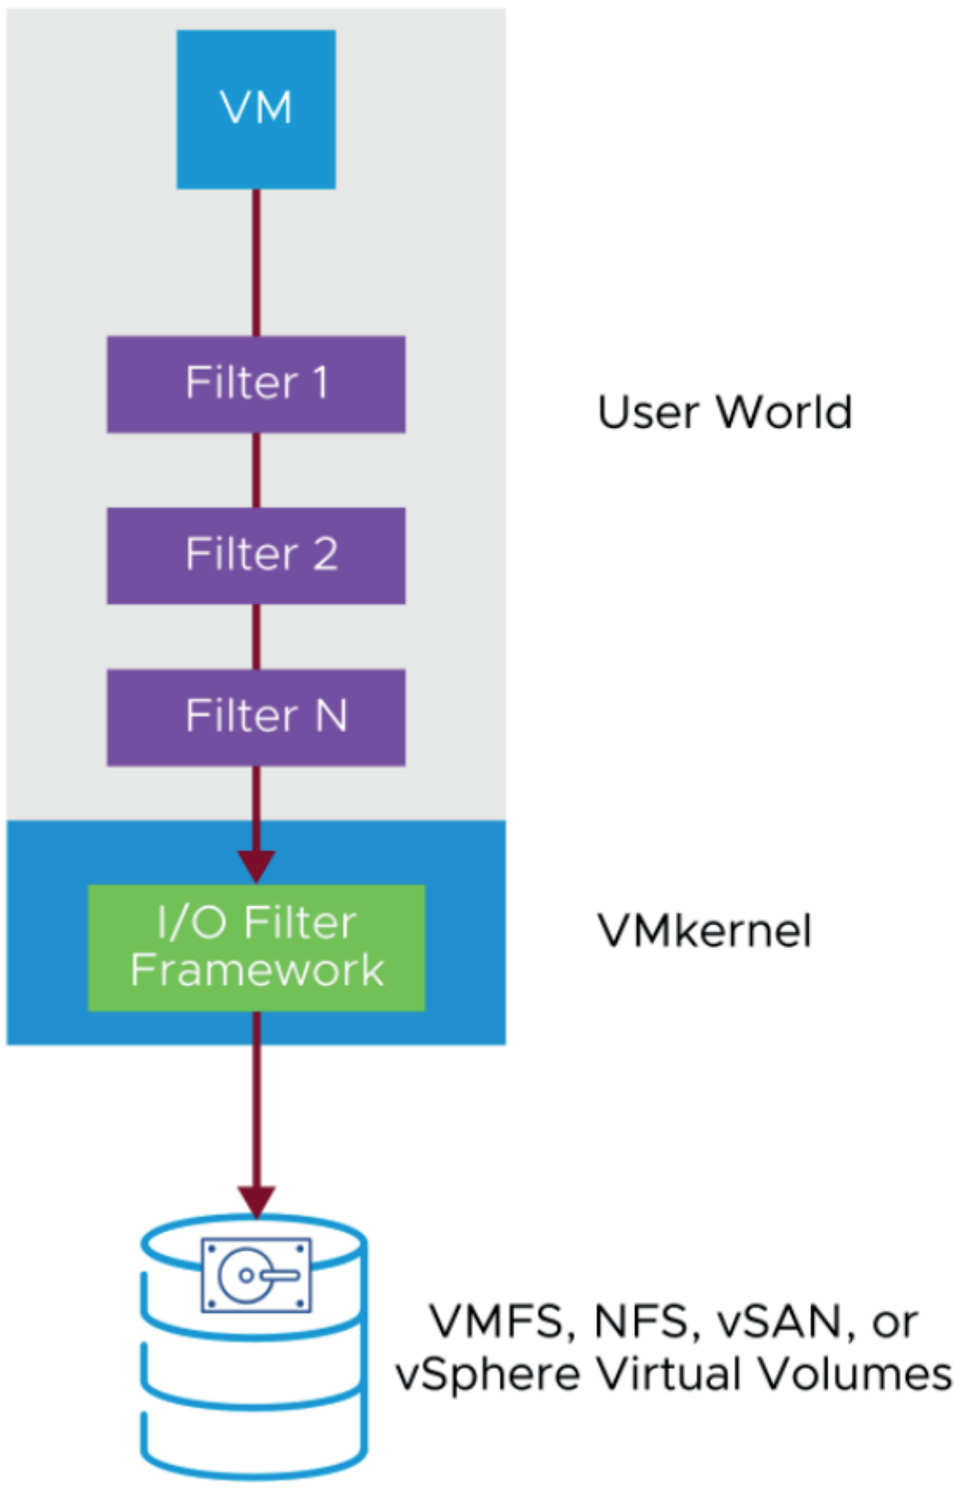

Filters operate by intercepting data in the VMkernel before the data is sent out to physical storage.

The I/O filter framework exists outside the virtual SCSI layer, which means that offline I/O to the disk (that is, modifications when a VM is powered off) can also be intercepted. The framework defines which filters should be applied and in which order.

I/O filters run in the VM's user world. A world is an execution context that is scheduled on a processor. A world is like a process in conventional operating systems.

5-85 How I/O Filters Work (2)

The I/O filter framework applies I/O filters in the following order:

1. Replication
2. Encryption
3. Caching
4. Storage I/O Control

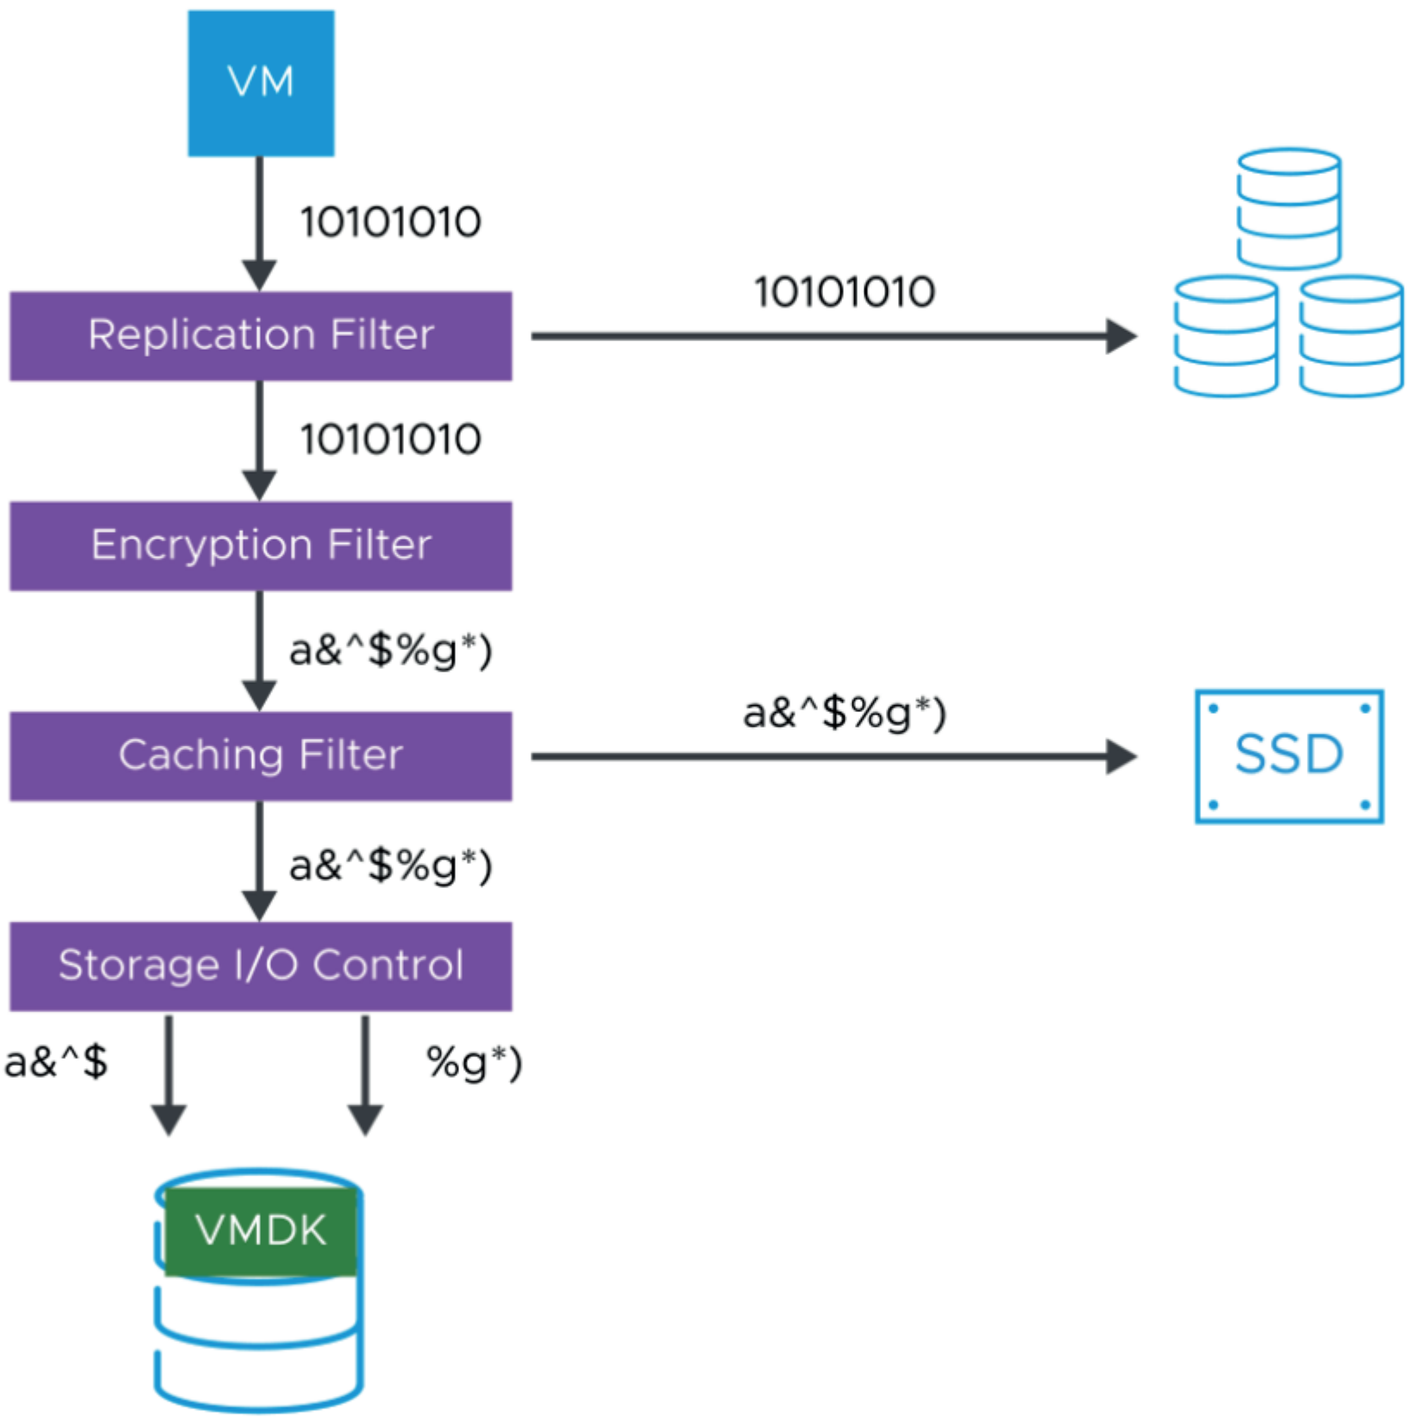

The order in which I/O passes through these filter categories is predefined. This order, as shown on the slide, is important, especially if you plan to implement a replication filter and an encryption filter.

Data passes from the filters in this predefined order before being forwarded to the storage on which the VM files or objects reside.

You can install several filters from the same category on your ESXi host. However, you can have only one filter from the same category per virtual disk.

Because the replication filter comes before the encryption filter, the data that is replicated is in plain text. Additional security measures must be taken to protect these transmissions and the data on the replication target.

Technet24

5-86 I/O Filter Providers

ESXi has preinstalled I/O filters:

- VM encryption
- Storage I/O Control

Third-party I/O filters must be installed according to the vendor instructions.

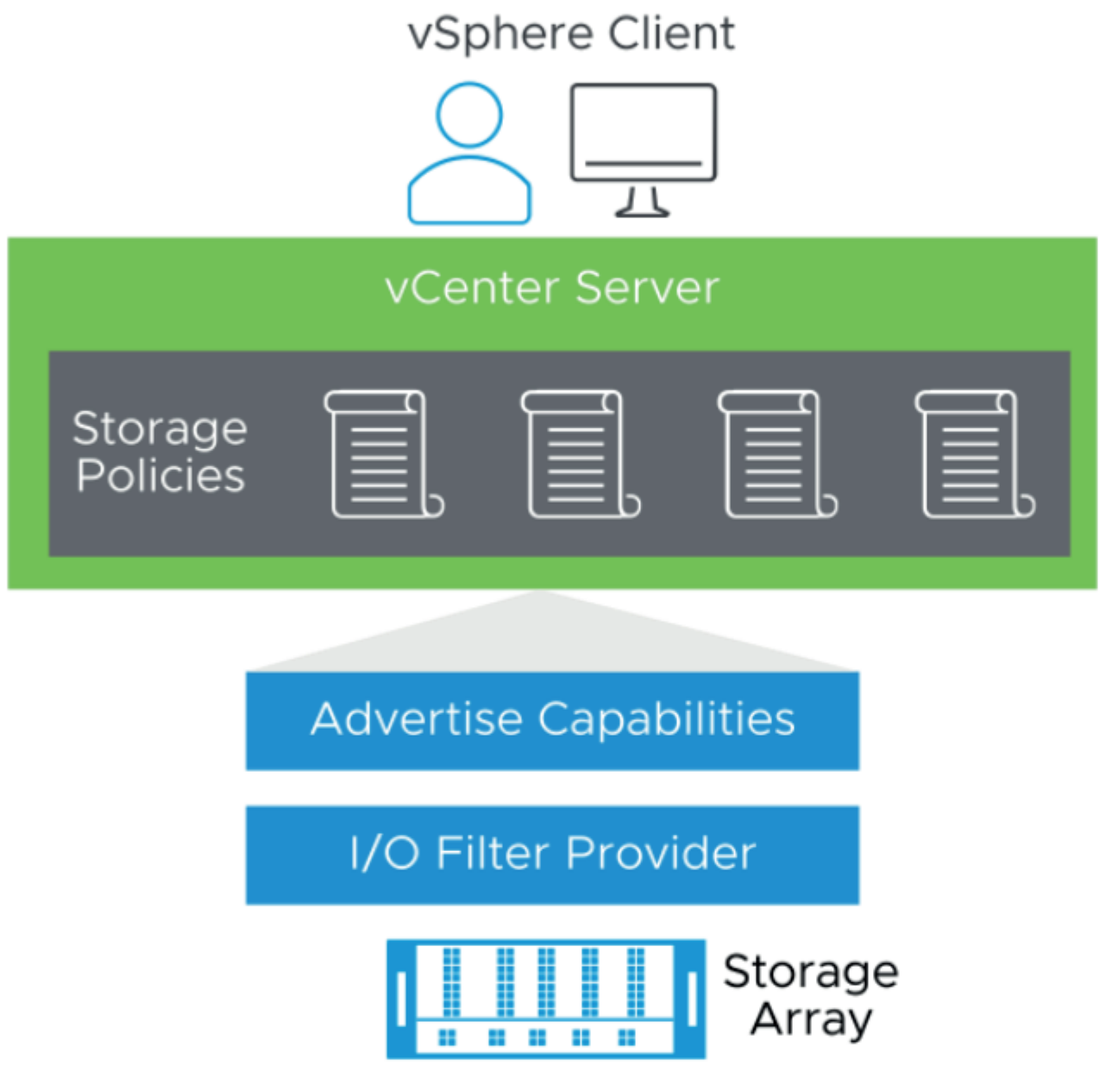

After I/O filters are installed on ESXi hosts, the I/O filter framework configures and registers a storage provider, also called an IOFilter provider, for each host.

The IOFilter provider advertises the new filter capabilities to vCenter. You can use these capabilities to build storage policies.

Storage providers for I/O filtering are software components in vSphere. These providers integrate with I/O filters and report data service capabilities that are supported by I/O filters to vCenter.

If you use I/O filters provided by third parties, you must install the I/O filters in an ESXi host cluster. You cannot install the filters on selected hosts. The filter packages are typically distributed as VIBs, which can include I/O filter daemons, CIM providers, and other associated components.

Typically, to deploy the filters, you run installers provided by vendors.

The capabilities populate the VM Storage Policies pane in the vSphere Client and can be referenced in a VM storage policy. You apply this policy to virtual disks so that the I/O filters can process I/O for the disks.

5-87 Displaying I/O Filter Provider Information

To view the I/O filter providers for each host, select the vCenter object in the inventory and select **Configure > Storage Providers**.

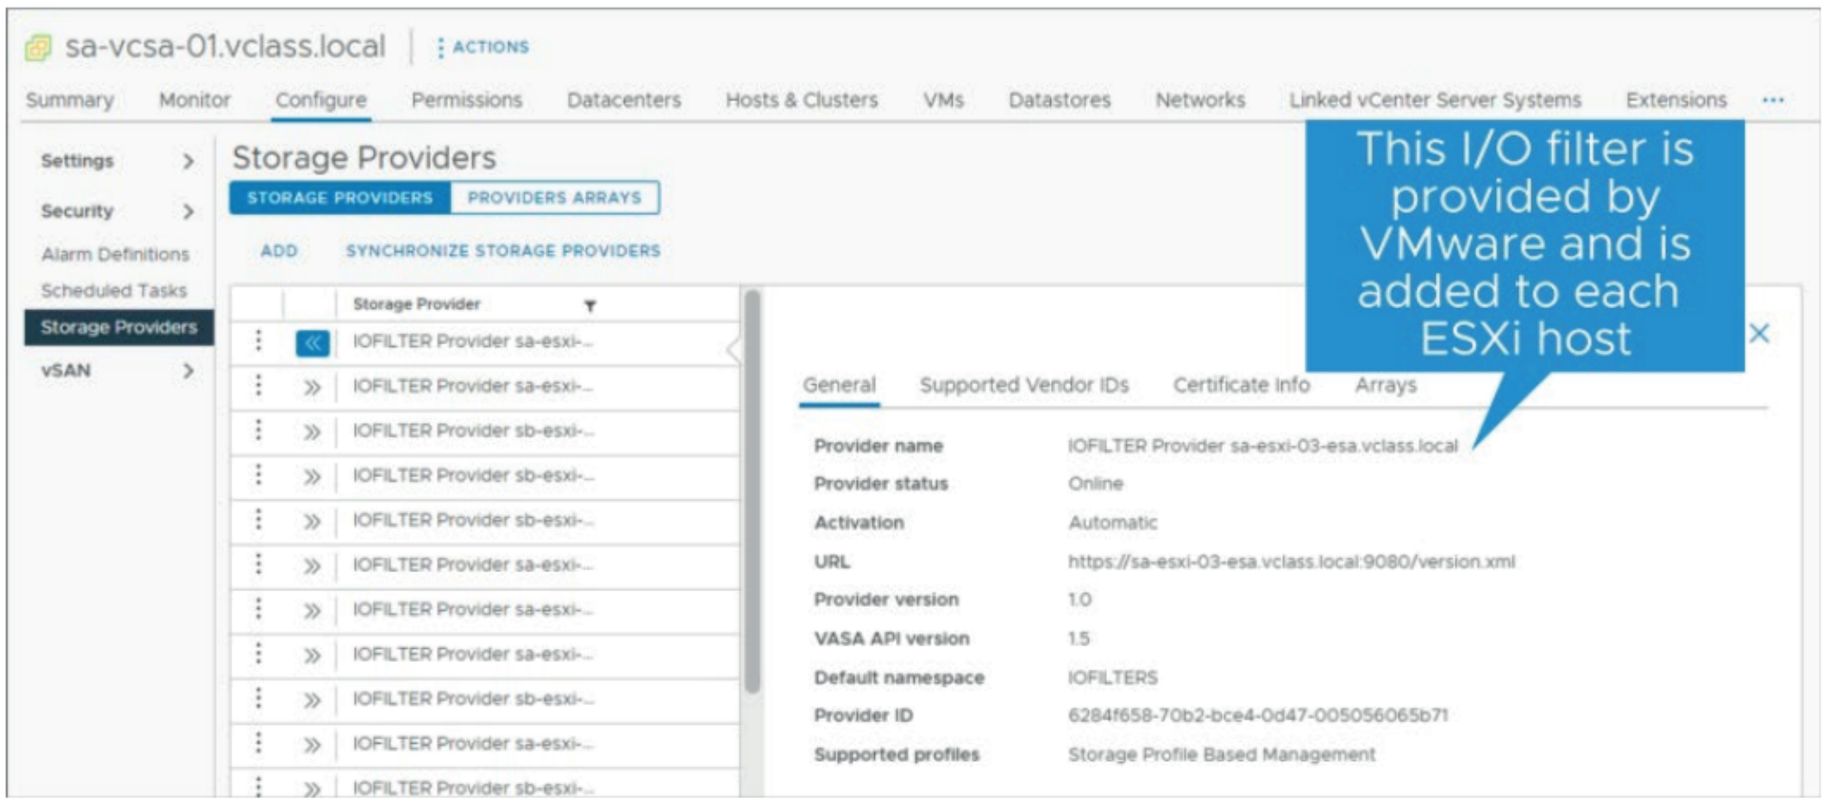

5-88 Review of Learner Objectives

- Describe how vSphere API for Storage Awareness ensures that a VM's storage requirements are met
- Explain how vendors can create data services for VMs by using vSphere APIs for I/O Filtering

Technet24

5-89 Lesson 7: Storage Policy Based Management Overview

5-90 Learner Objectives

- Describe storage policy based management
- Recognize the categories of storage policy rules

5-91 About Storage Policy Based Management (1)

Storage policy based management (SPBM) ensures that VMs use storage that guarantees a specified level of capacity, performance, availability, redundancy, and so on.

SPBM abstracts storage and data services delivered by vSAN, vSphere Virtual Volumes, I/O filters, or traditional storage.

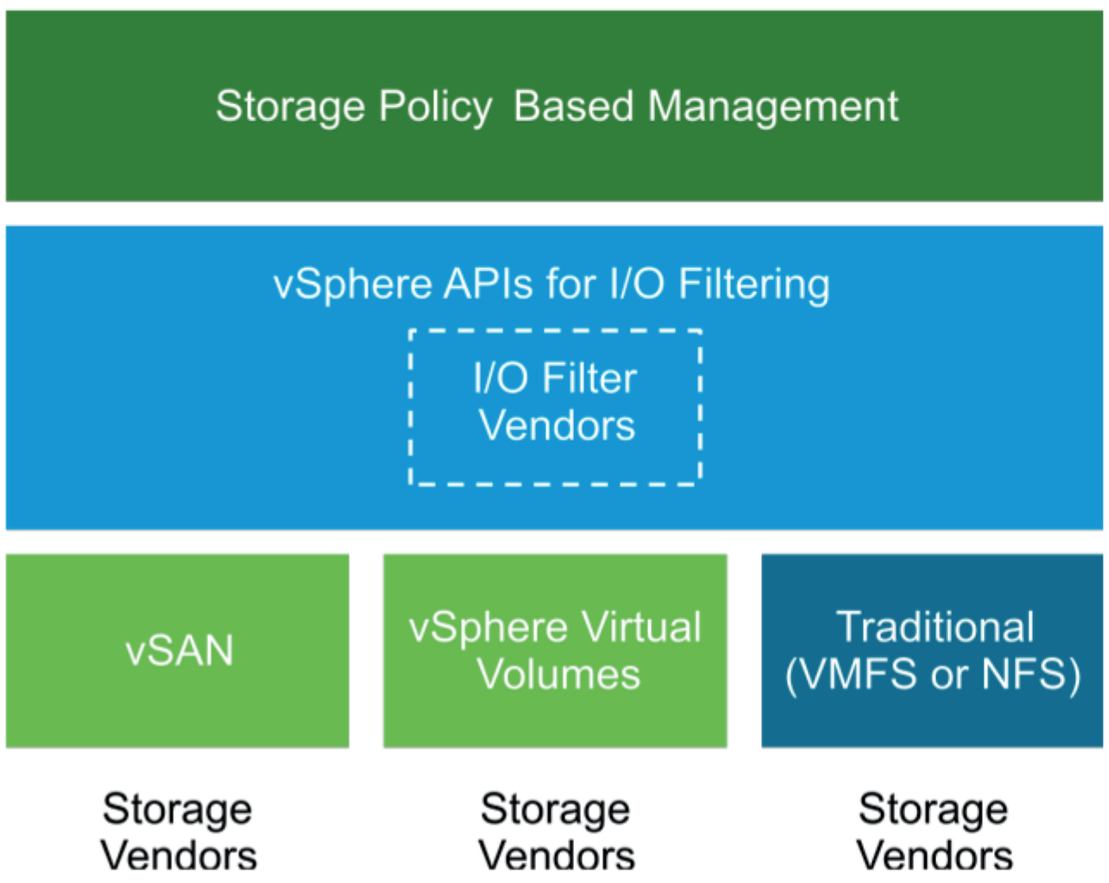

VMware partners and storage vendors can support these storage and data services.

Rather than integrating with each individual vendor or type of storage and data service, storage policy based management (SPBM) provides a universal framework for many types of storage entities.

When an application is to be provisioned, vSphere uses the Storage Policy Based Management module to match the policy requirements against configured storage container capabilities. In this way, vSphere determines where the VMDK files should exist.

5-92 About Storage Policy Based Management (2)

Storage policy based management has the following features:

- Advertisement of storage capabilities and data services that storage arrays and I/O filters
- Bidirectional communication with ESXi and vCenter on one side, and storage arrays and I/O filters on the other
- VM provisioning based on VM storage policies

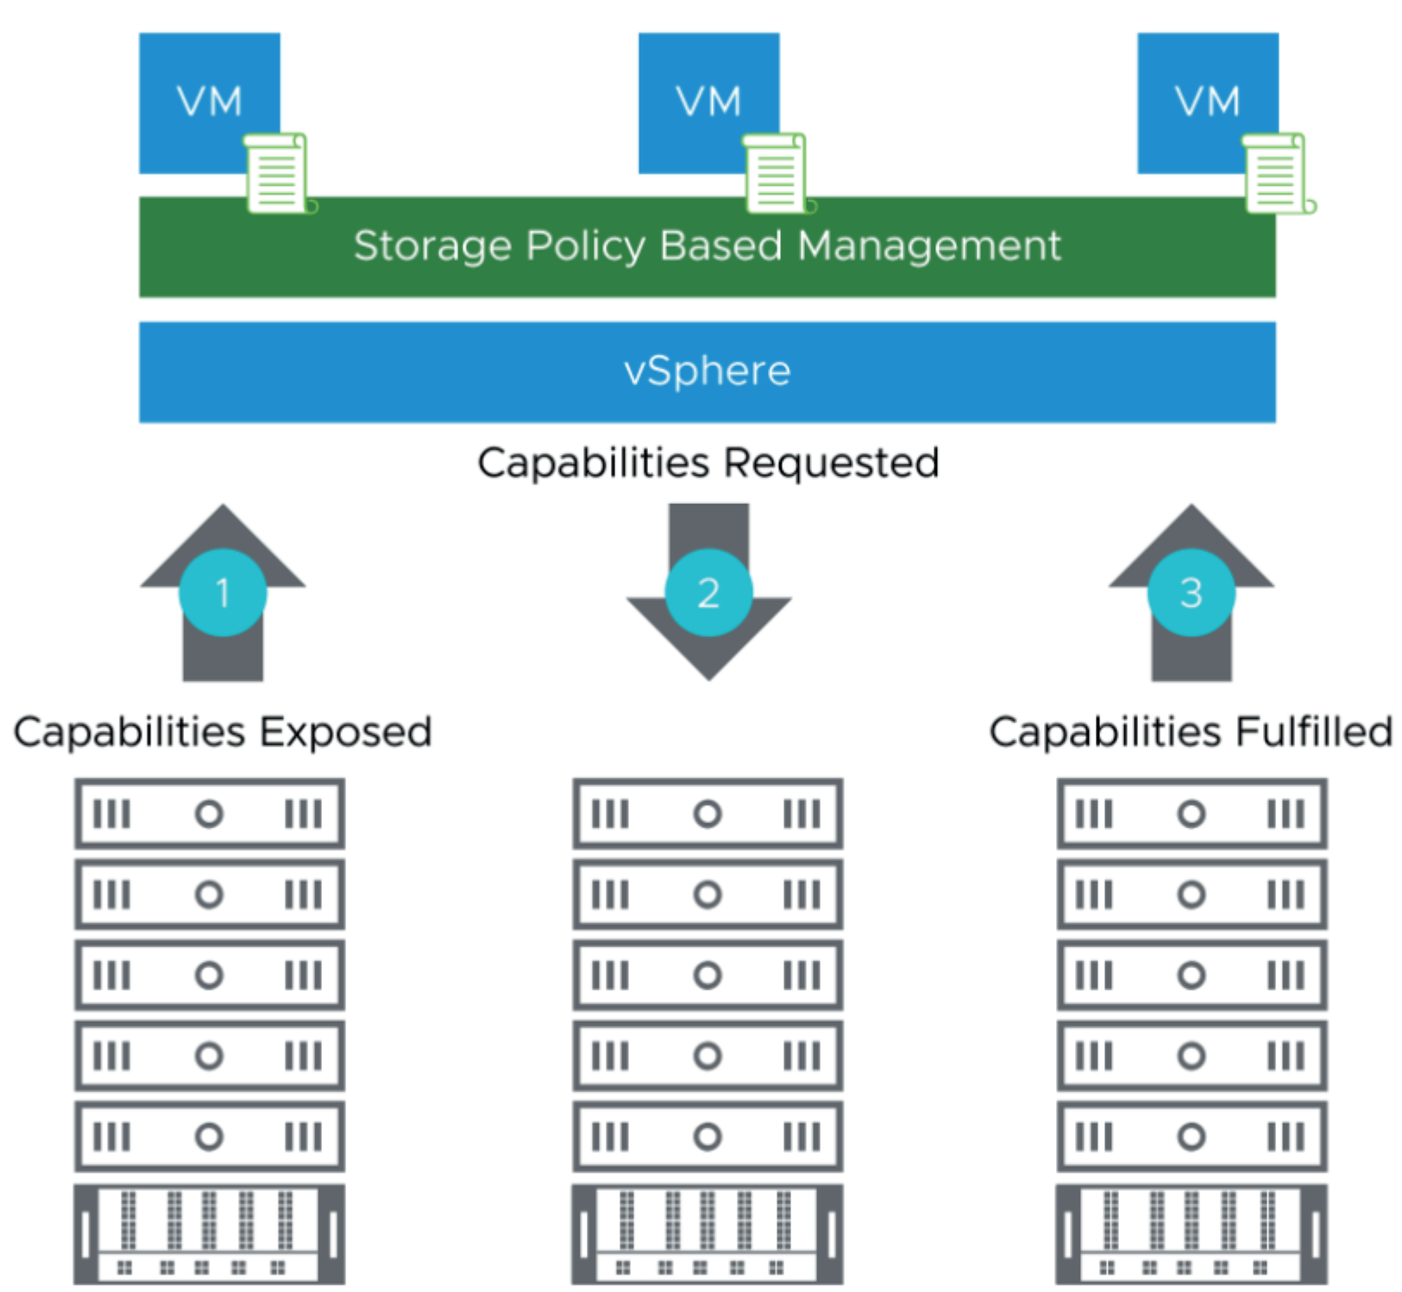

VM storage policies minimize the amount of storage planning that you must do for each VM.

For example, you can use VM storage policies to create basic storage tiers. Datastores with similar capabilities are tagged to form gold, silver, and bronze tiers.

Redundant, high-performance storage might be tagged as the gold tier. Nonredundant, medium performance storage might be tagged as the bronze tier.

The capabilities exposed are performance, capacity, protection, and so on.

The capabilities requested are aligned on VM boundaries.

The capabilities fulfilled are datastore provisioned, and the policy is kept with the datastore.

Technet24

5-93 About VM Storage Policies

A VM storage policy controls which type of storage is provided for the VM and how the VM is placed within storage. It also determines data services that the VM can use.

◊ VM storage policies help you meet the following goals:

- Categorize datastores based on certain levels of service.
- Provision a VM's disks on the correct storage.

VM storage policies capture storage characteristics that VM home files/objects and virtual disks require to run applications on a VM. These storage characteristics are defined in a set of rules.

VM storage policies are essential to VM provisioning.

You use VM storage policies during the provisioning of a VM to ensure that a VM's disks are placed on the storage that is best for its situation.

For example, with VM storage policies, you can ensure that the VM running a critical I/O intensive database is placed in the gold tier.

Ideally, you want to create the best match of predefined VM storage requirements with available physical storage properties.

5-94 About Storage Policy Rules

Rules are the basic elements of a storage policy. Each rule is a statement that describes a single requirement for VM storage and data services.

Category	Description	Examples
Capability-based placement rules	Specify a particular storage requirement for the VM, which guides the decision about VM placement.	• RAID level • Number of failures to tolerate
Tag-based placement rules	Reference datastore tags to define VM placement.	• Gold Tier tag • Silver Tier tag • Bronze Tier tag
Data service rules	Activate specific data services for the VM, but do not define storage placement and VM storage requirement.	• Caching • Replication

Capability-based placement rules describe how VM storage objects are allocated within a datastore to achieve the required level of service. For example, the rules can list virtual volumes as a destination and define the maximum recovery point objective for the virtual volume objects. Storage policy based management finds the virtual volumes datastores that can match the rules and satisfy the storage requirements of the VM.

Tag-based placement rules reference datastore tags. You can use the tag-based rules to fine-tune your VM placement request. For example, you can exclude datastores with a Palo Alto tag from the list of your virtual volumes datastores.

Data service rules activate specific data services. Storage systems or other entities can provide these services. They can also be installed on your hosts and vCenter. You include the data service rules in the storage policy components.

Technet24

5-95 Storage Policy Example (1)

When you create a storage policy, you can select one or more rules that represent service levels demanded by VMs.

These rules define storage characteristics.

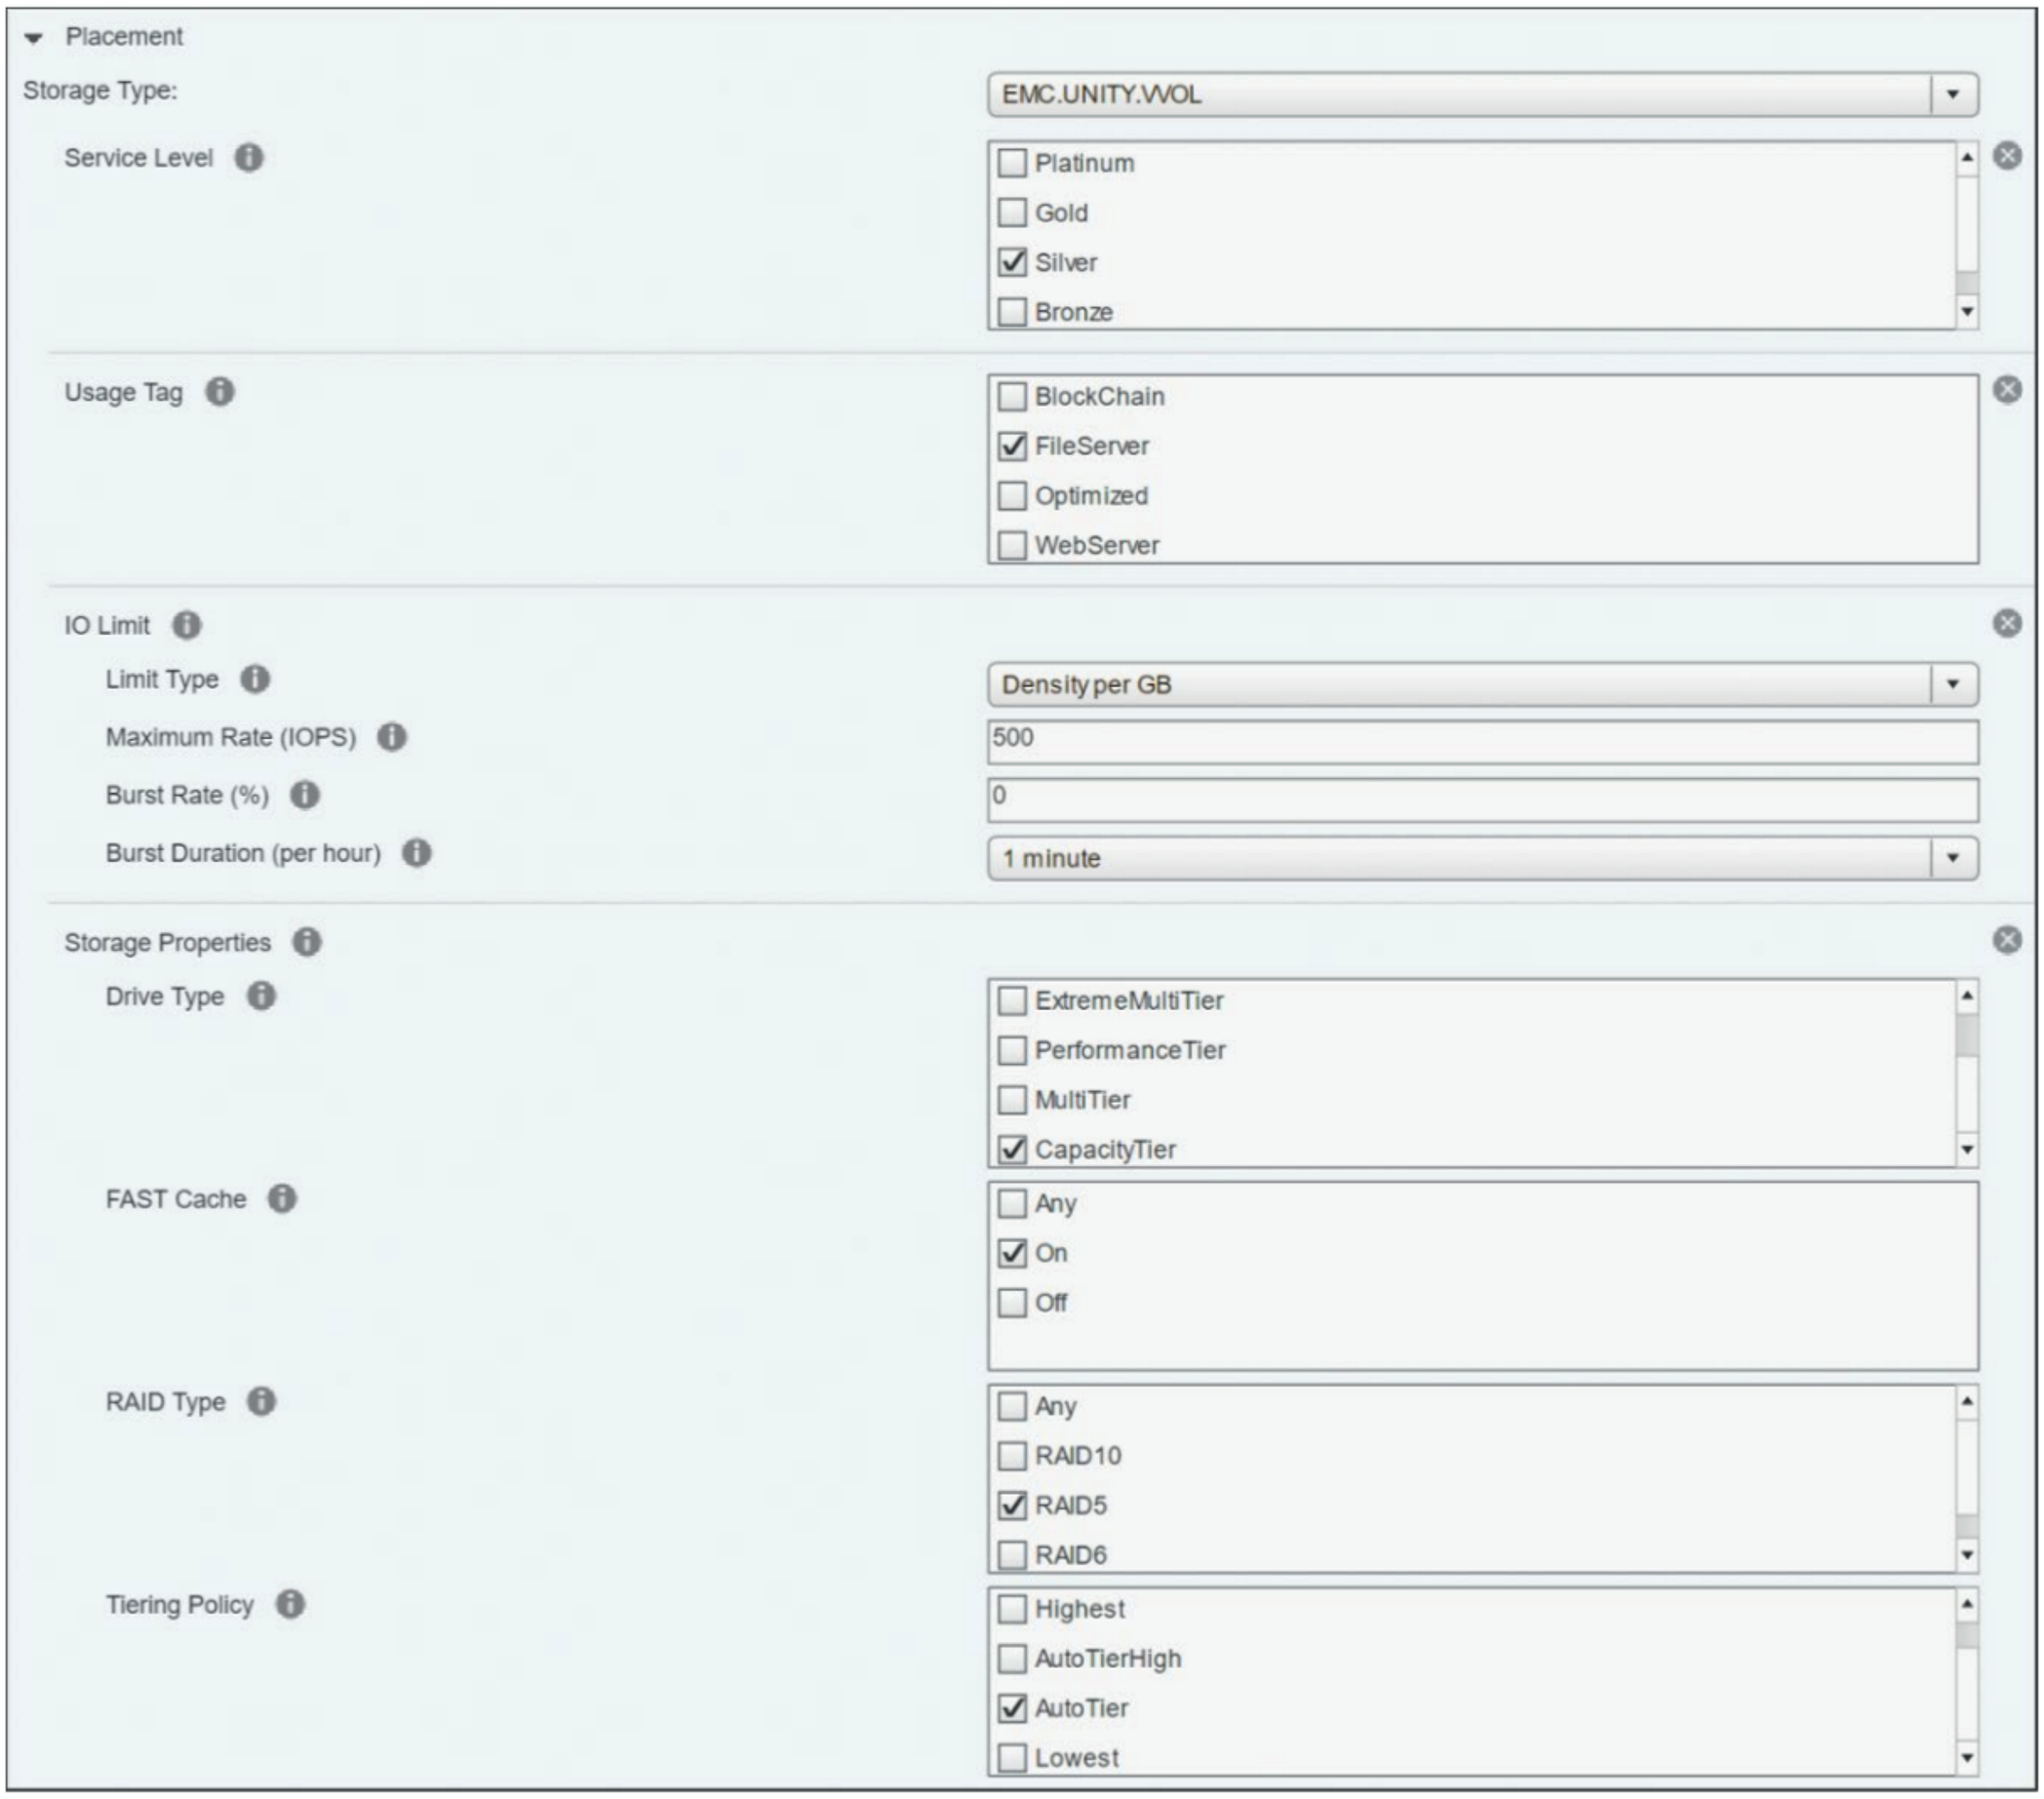

Each array vendor has an identified namespace. The namespace contains the storage container identifiers and assigned capabilities. The capabilities appear as options in the client interface.

5-96 Storage Policy Example (2)

VM storage policies report when VMs run on storage with the correct performance and availability characteristics.

For vSAN and virtual volumes, VM storage policies will ensure that VMs are placed on the correct storage.

For example, to place a database server's disks on appropriate storage:

- Assign the Silver policy to the operating system disk: The Silver policy is set up to use only HDDs.
- Assign the Gold policy to the data disk: The Gold policy uses only SSDs.

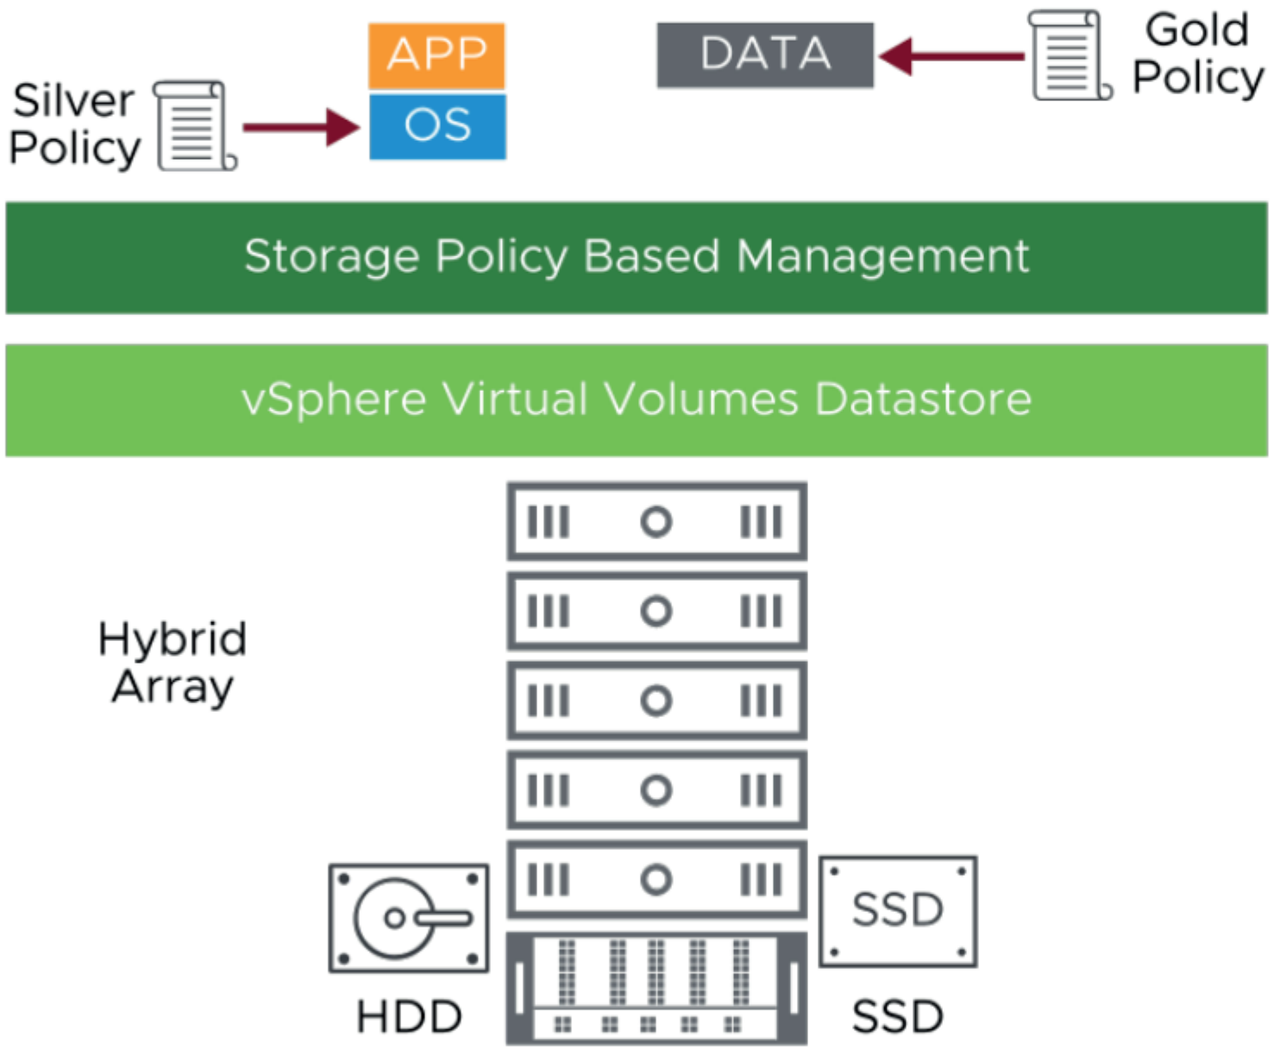

Technet24

5-97 About Storage Policy Components

A VM storage policy can include one or several reusable and interchangeable building blocks, called storage policy components, using VAIO-based storage services.

Each component describes one type of service from a storage provider.

The services vary depending on the storage provider that you use, but they generally belong to one of the following categories:

- Replication
- Encryption
- Caching
- Storage I/O Control

The service provider can be a storage system, an I/O filter, or another entity.

5-98 Storage Policy Component Example

When you create the storage policy component, you define the rules for one specific type and grade of service.

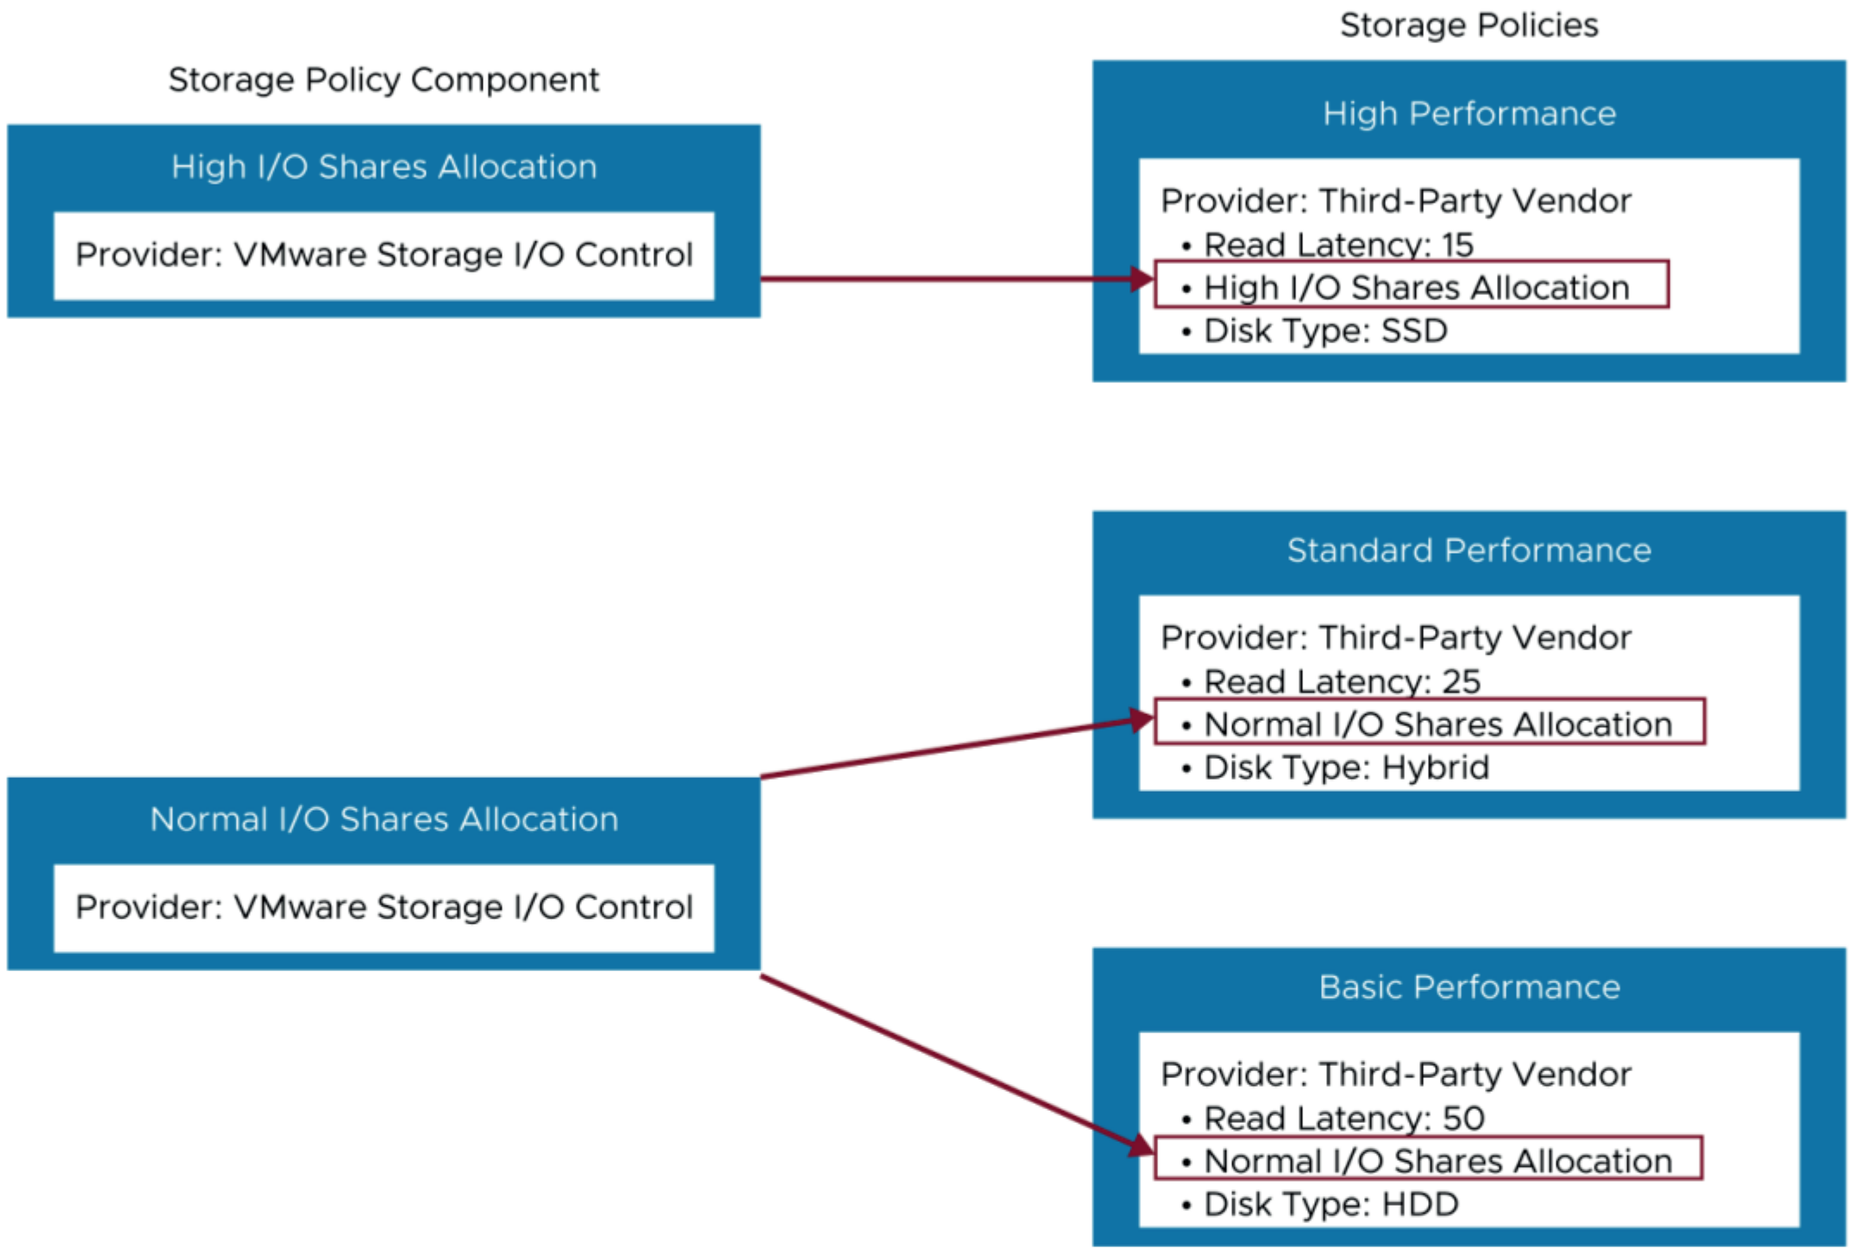

You must add the component to the storage policy and assign the policy to the VM. You cannot assign a component directly to a VM or virtual disk.

You can define the policy components in advance and associate them with multiple VM storage policies.

5-99 Activity: Storage Policy Rules (1)

Which of the following statements best describes Storage Policy Rules on vSphere?

◊ ❒ Storage Policy Rules define the access permissions for vSphere administrators.

◊ ❒ Storage Policy Rules determine the performance characteristics and availability settings for virtual machine storage.

◊ ❒ Storage Policy Rules specify the replication settings for virtual machines.

◊ ❒ Storage Policy Rules define the backup and restore policies for virtual machine data.

Technet24

5-100 Activity: Storage Policy Rules (2)

Which of the following statements best describes Storage Policy Rules on vSphere?

◊ ❒ Storage Policy Rules define the access permissions for vSphere administrators.

◊ ✓ Storage Policy Rules determine the performance characteristics and availability settings for virtual machine storage.

◊ ❒ Storage Policy Rules specify the replication settings for virtual machines.

◊ ❒ Storage Policy Rules define the backup and restore policies for virtual machine data.

5-101 Review of Learner Objectives

- Describe storage policy based management
- Recognize the categories of storage policy rules

5-102 Lesson 8: Using VM Storage Policies

5-103 Learner Objectives

- Create VM storage policies
- Use VM storage policies to create tiered storage

5-104 Creating and Managing Storage Policies

The process of creating and managing storage policies for NFS/VMFS datastores without a VASA provider typically includes several steps:

1. Configure storage to be used for storage policies.
2. Create storage policy components.
3. Create VM storage policies.
4. Apply the storage policy to the VM.
5. Check compliance for the VM storage policy.

5-105 Step 1: Configuring Storage to Be Used for Storage Policies

For entities represented by storage providers, the appropriate provider should be registered:

- Storage provider entities include vSAN, vSphere Virtual Volumes, and I/O filters.
- Some providers are self-registered. Other providers must be manually registered.

For datastores that are not represented by storage providers, create datastore tags:

- Use tags to indicate a property that is not communicated through the storage providers, such as geographical location or administrative group.

When you create storage policies, the vSphere Client shows information from the configuration of datastores and data services in your environment.

This information is obtained from storage providers and datastore tags.

Technet24

5-106 Step 2: Creating Storage Policy Components

You define storage policy components in advance so that you can use them when you create a storage policy.

vCenter provides a few built-in storage policy components, specifically for encryption and Storage I/O Control.

5-107 Step 3: Creating VM Storage Policies

When you create a storage policy, you must select the policy structure:

- **Host based services** define rules to enable and configure data services provided by your host through VAIO.
- **Datastore specific rules** describe requirements or services for a storage type through VASA and tagging methods.

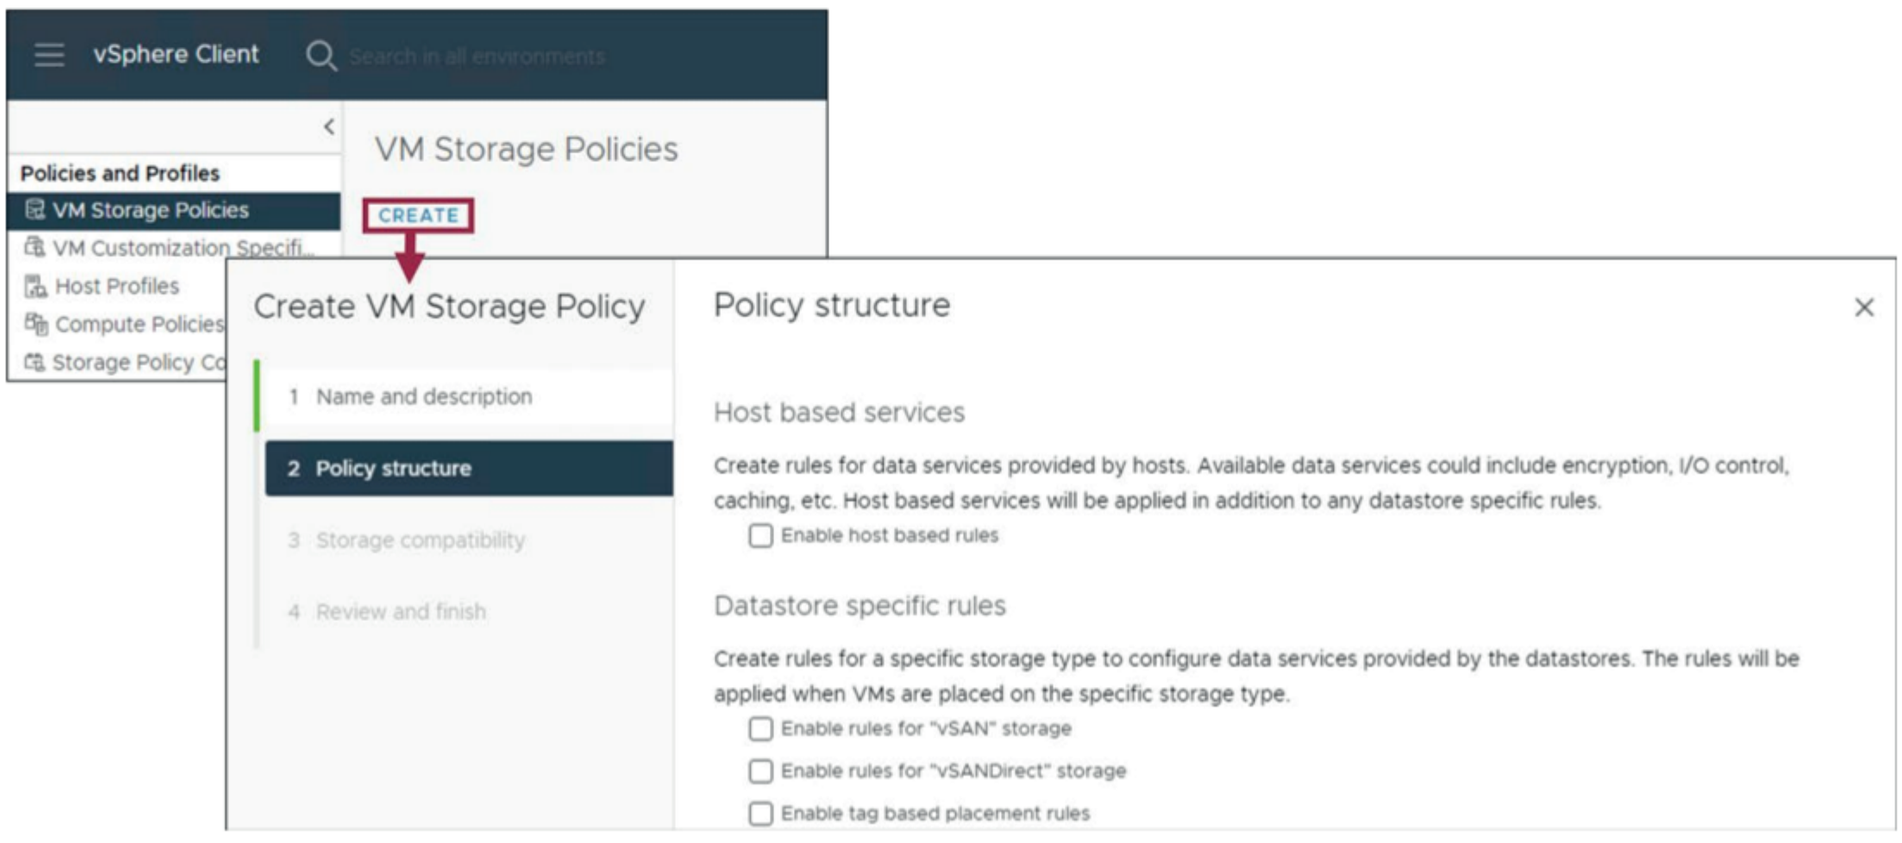

5-108 Example: Creating a Silver Tier Storage Policy (1)

In this example, you create a storage policy named Silver Tier Policy to implement tiered storage in your environment.

Storage tiering allows for differences in storage requirements by creating multiple levels of storage with varying degrees of performance, reliability and cost, depending on the application workload needs.

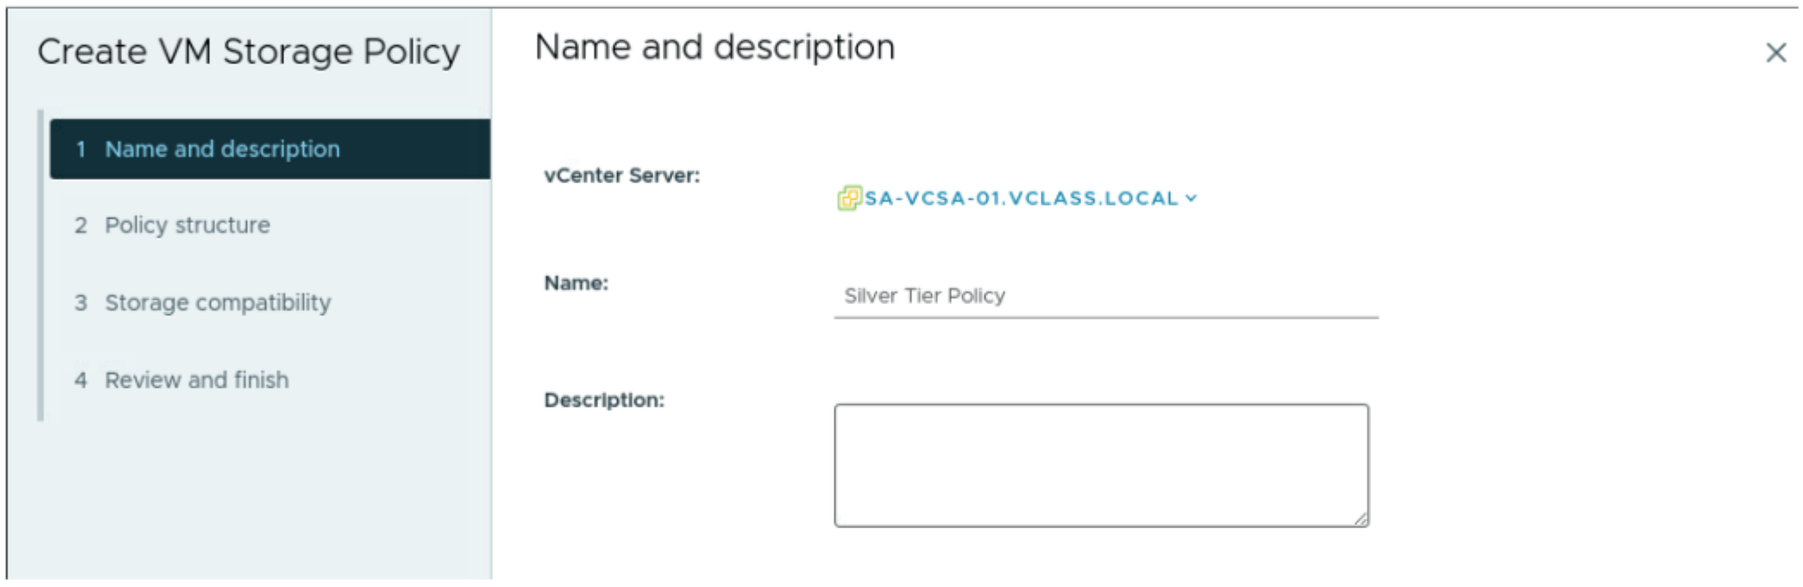

5-109 Example: Creating a Silver Tier Storage Policy (2)

The Silver Tier Policy uses a tag-based placement rule.

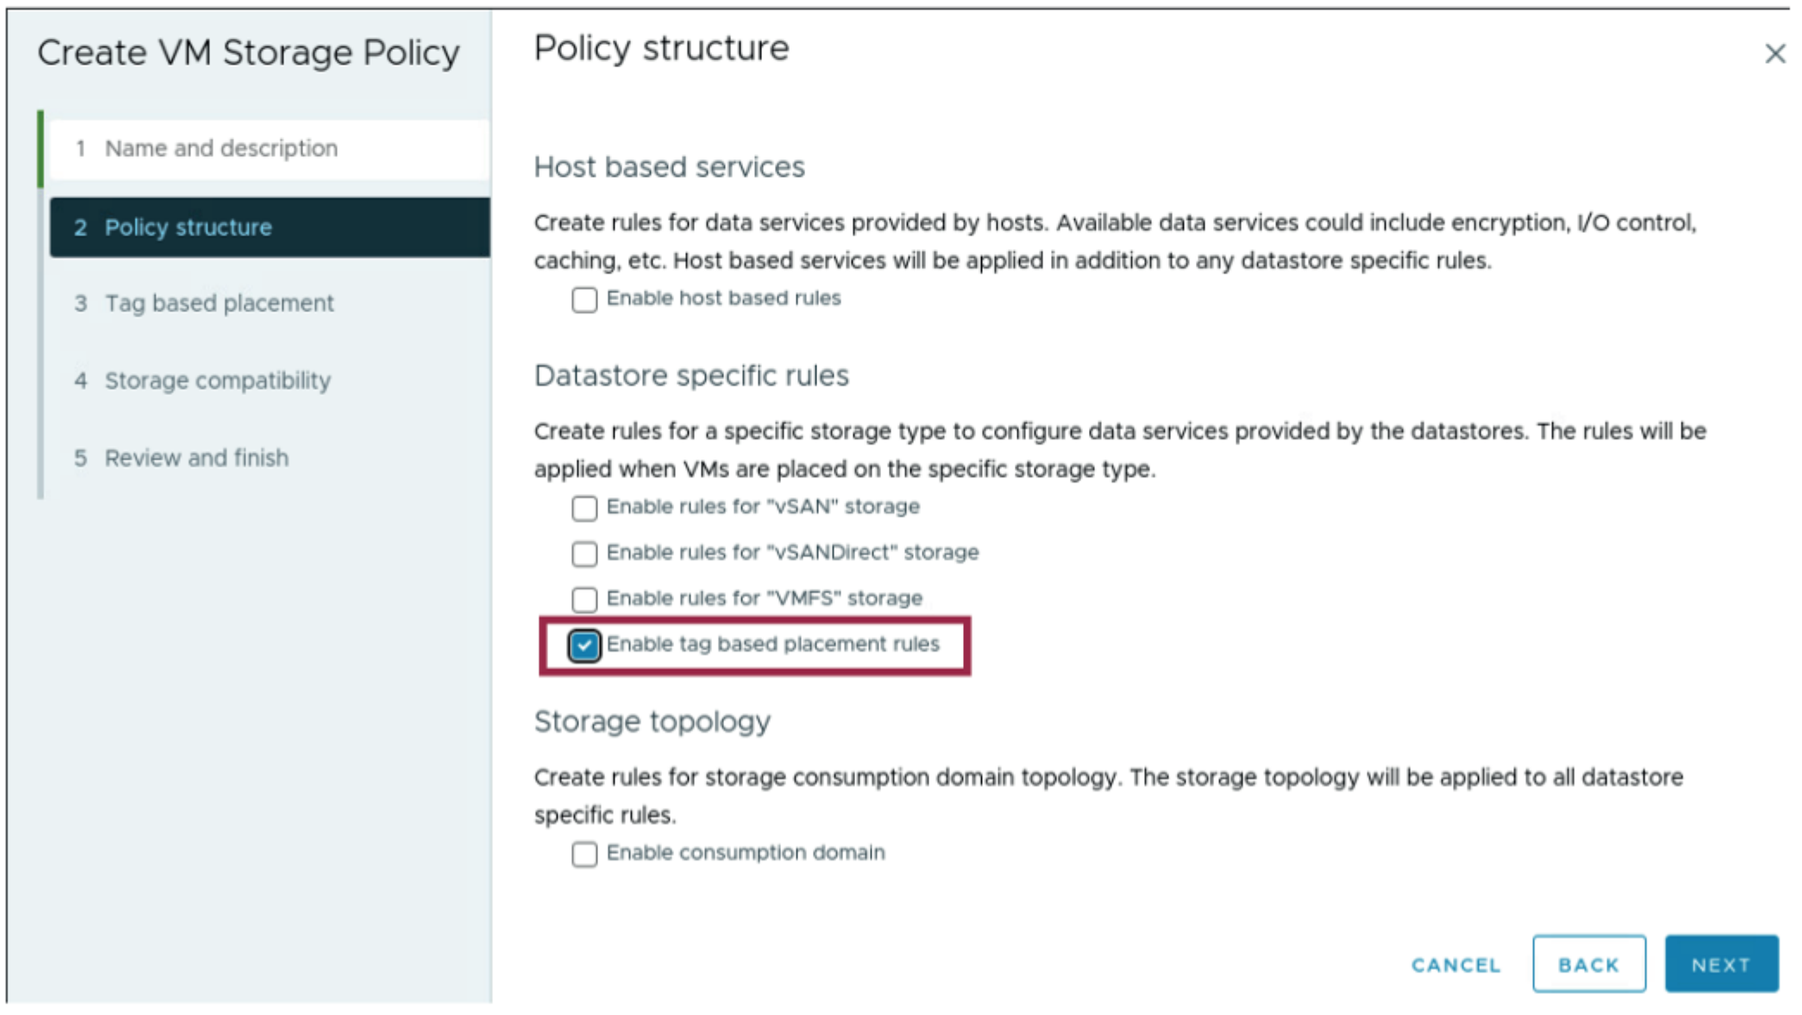

Technet24

5-110 Example: Creating a Silver Tier Storage Policy (3)

The Silver Tier Policy is configured to match any datastores that are tagged as Silver Tier. You select the Storage Tier tag category and then the Silver Tier tag

You select the **Storage Tier** tag category and then the **Silver Tie**r tag.

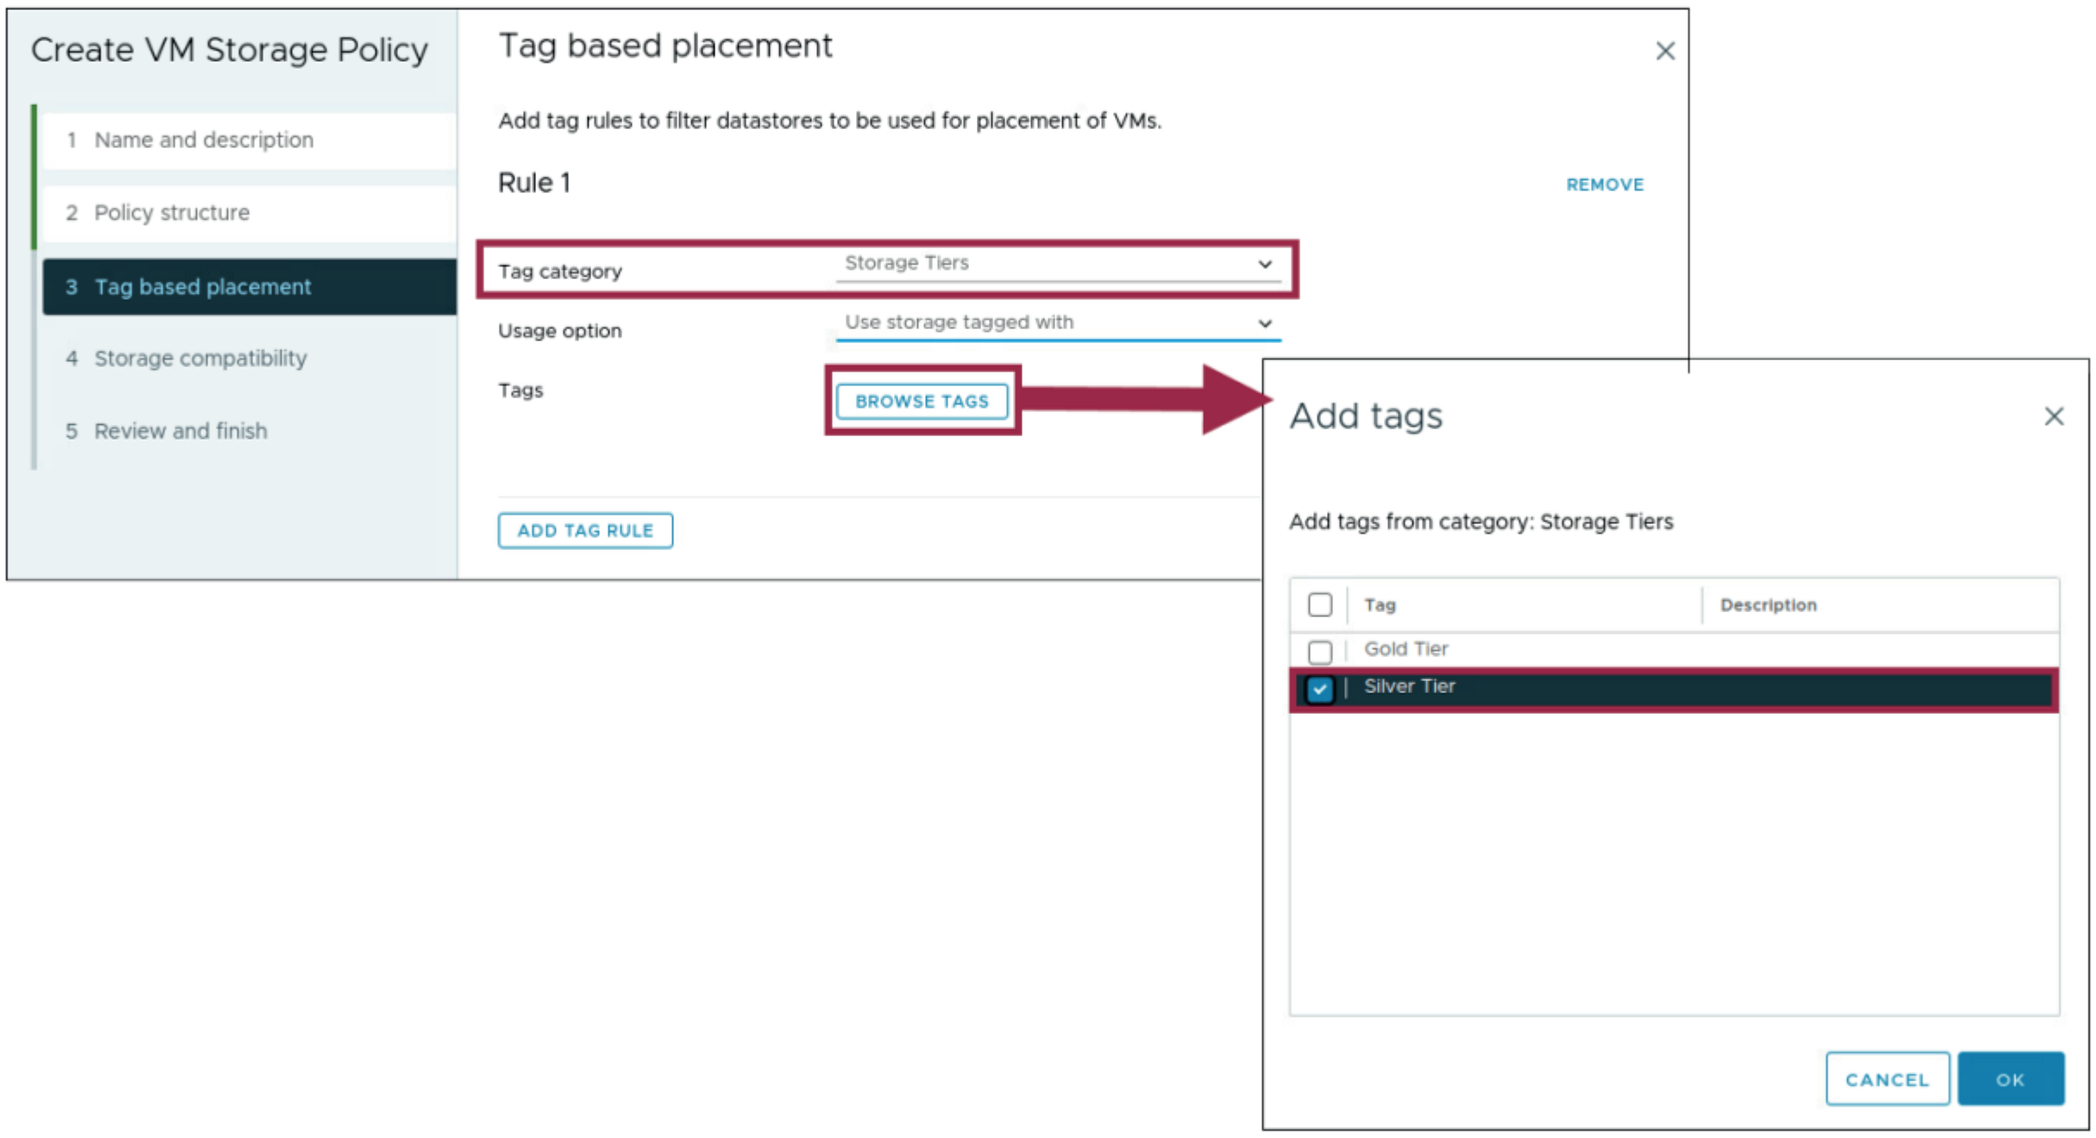

5-111 Example: Creating a Silver Tier Storage Policy (4)

The Storage compatibility pane lists the datastores that are compatible with the Silver Tier Policy.

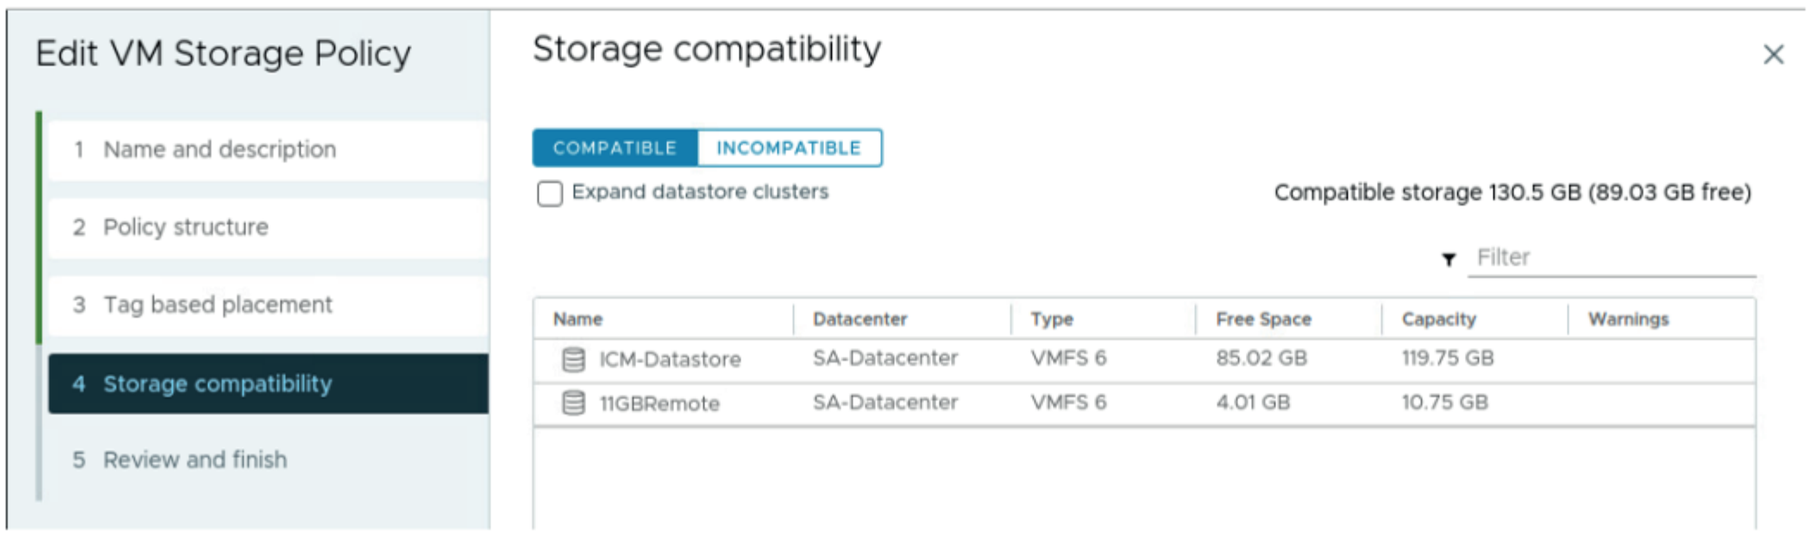

5-112 Step 4: Applying the Storage Policy to the VM

After creating the storage policy, you can assign the policy to a VM when you create, deploy or clone a VM or migrate a VM to a new datastore.

When you assign the Silver Tier Policy to a VM, a list of compatible datastores is shown.

In this example, one datastore is compatible: ICM-Datastore.

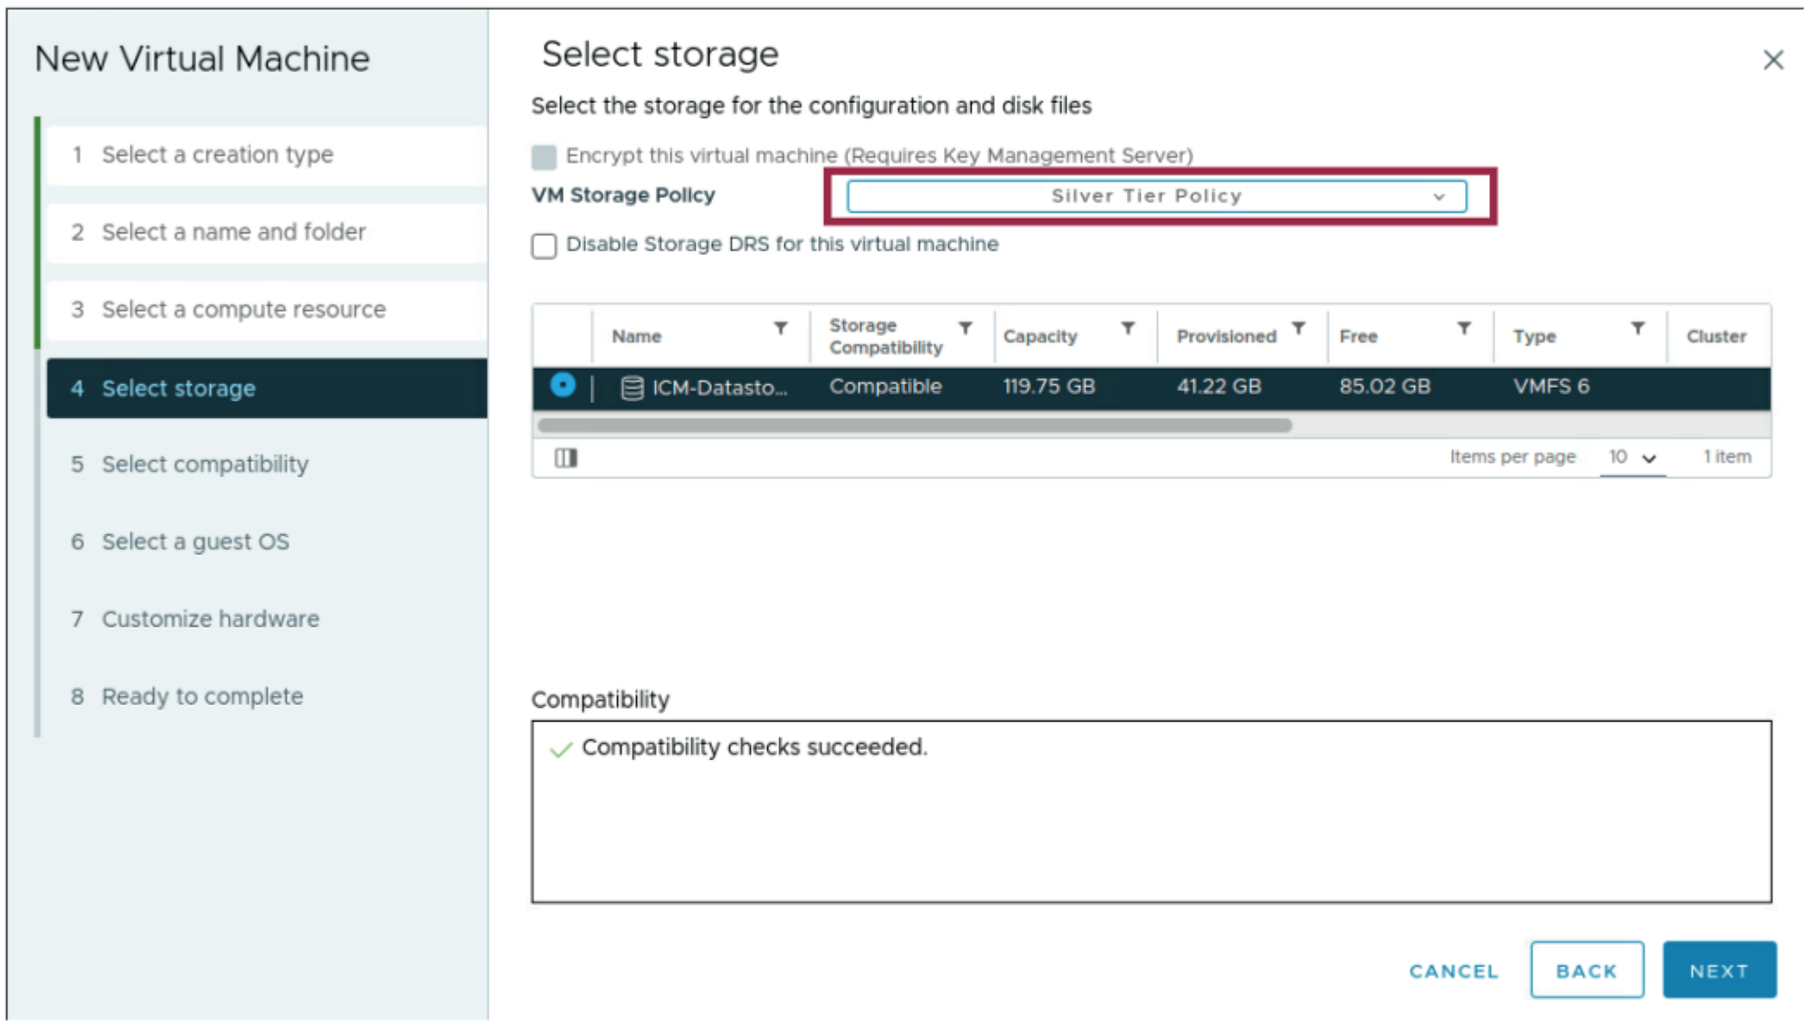

When you select a VM storage policy, the vSphere Client shows the datastores that are compatible with the capabilities of that policy. You can select a datastore or a datastore cluster from the list. If you select a datastore that does not match the storage policy, the vSphere Client shows that the VM is using noncompliant storage.

Technet24

5-113 Step 5: Checking Compliance for the VM Storage Policy

You can check whether VMs use stored on datastores that are compliant with the storage policy.

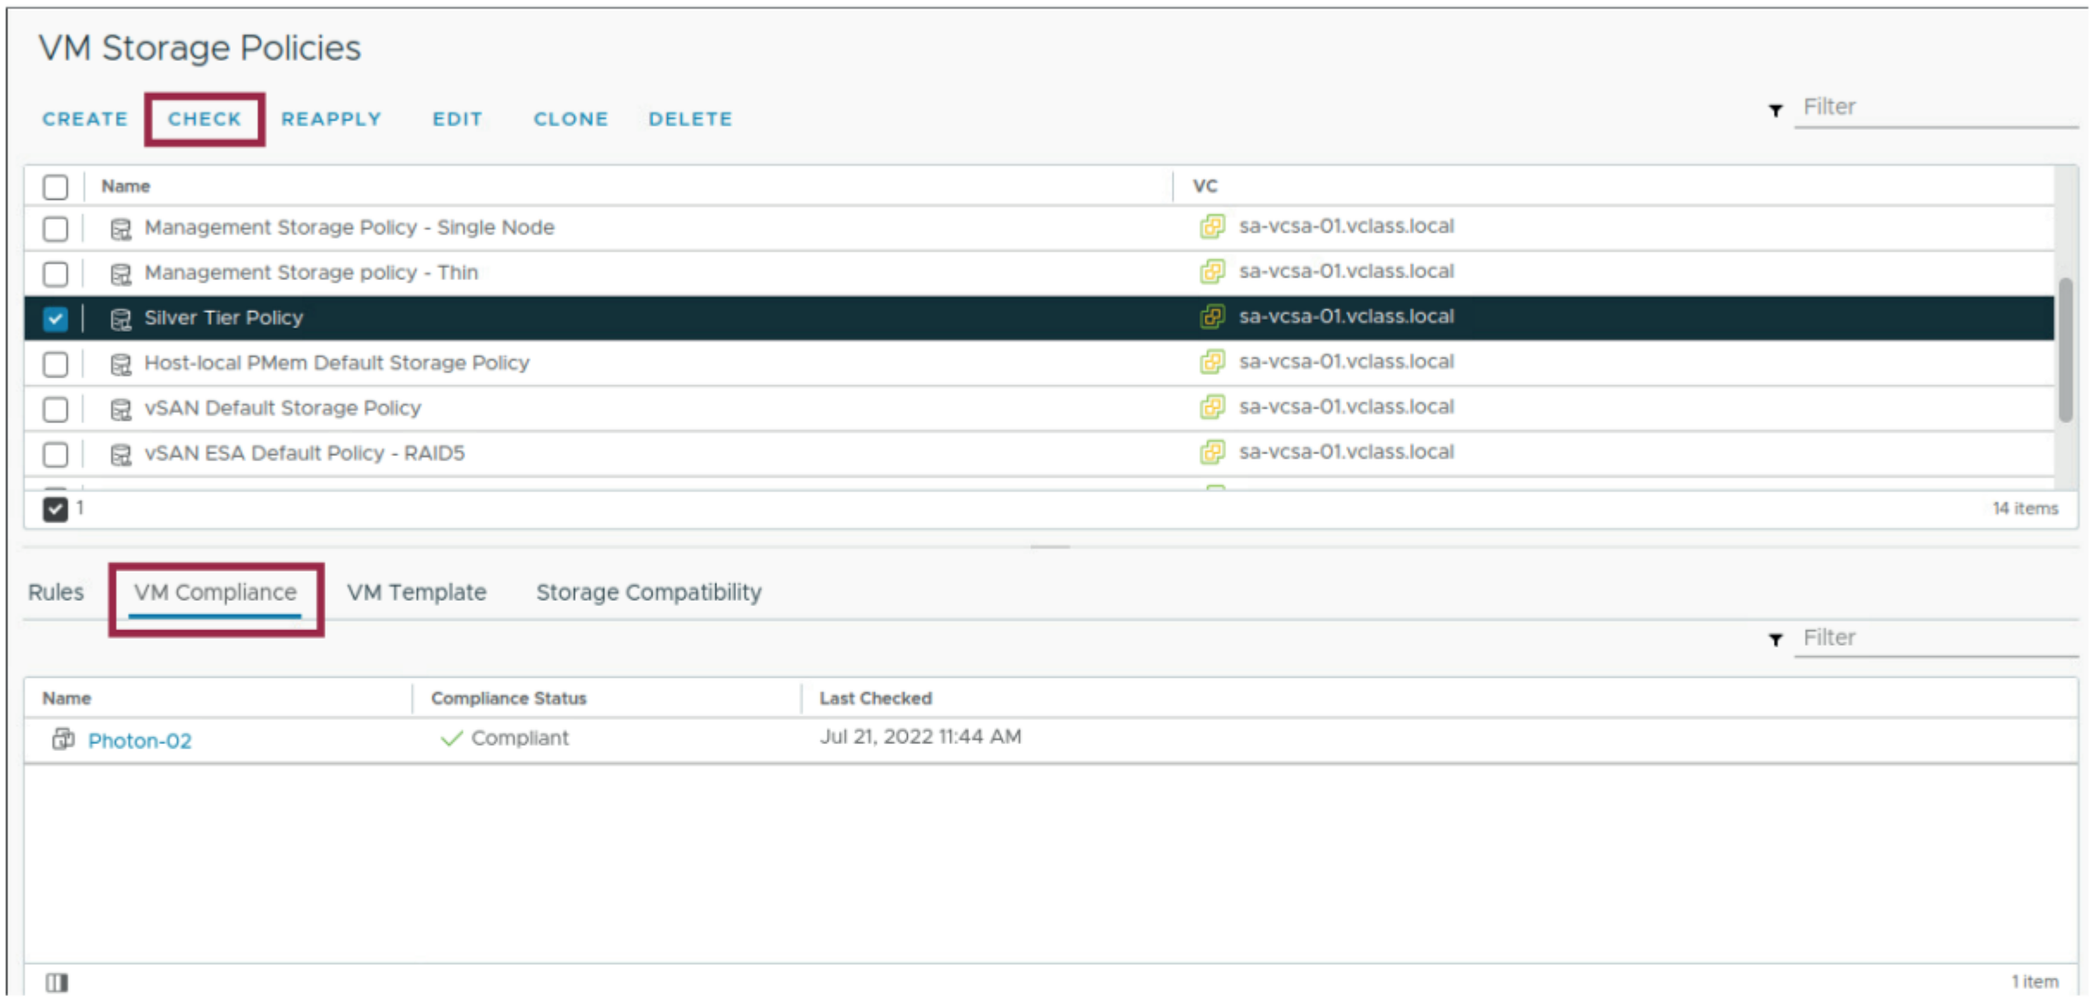

You can use VM storage policies during the ongoing management of VMs.

You can periodically check whether a VM is migrated to or created on inappropriate storage, potentially making it noncompliant.

You can also use storage information to monitor the health and use of storage and report when the VM's storage is noncompliant.

5-114 Lab 11: Using Policy-Based Storage

Use policy-based storage to create tiered storage for VMFS datastore without a VASA provider:

1. Add Datastores for Use by Policy-Based Storage
2. Use vSphere Storage vMotion to Migrate a VMs Storage
3. Configure Storage Tags
4. Create VM Storage Policies
5. Assign Storage Policies to VMs

5-115 Review of Learner Objectives

- Create VM storage policies
- Use VM storage policies to create tiered storage

Technet24

5-116 Lesson 9: Storage Policies for vSAN

5-117 Learner Objectives

- Create vSAN storage policies

5-118 vSAN Storage Providers

The vSAN datastore is represented by a storage provider.

Through the storage provider, the vSAN datastore advertises its capabilities, data services, and other characteristics in the vSphere Client.

You use these characteristics to define capability-based placement rules for your storage policies.

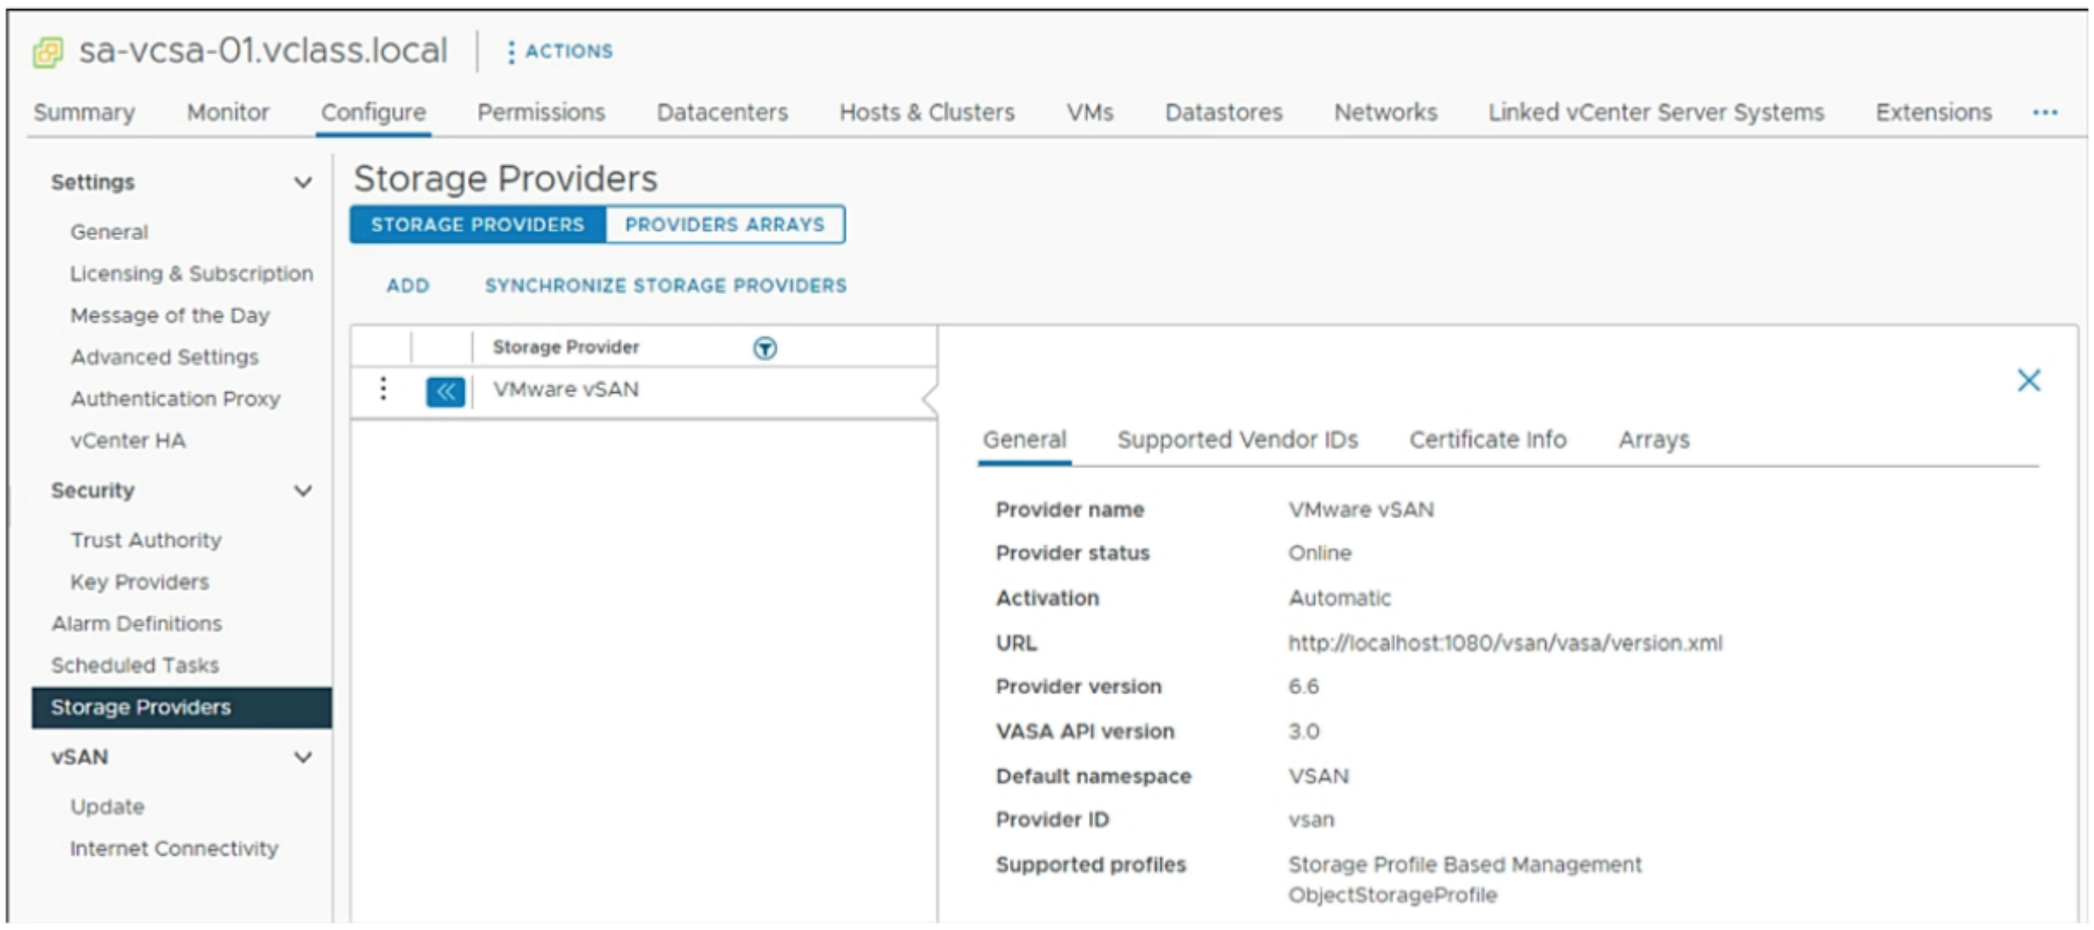

vSAN storage providers report a set of underlying storage capabilities to vCenter. They also communicate with the vSAN layer to report the storage requirements of the VMs.

5-119 vSAN Storage Policies

You can create multiple VM storage policies for a single vSAN datastore:

- vSAN has a default VM storage policy.
- You can create custom storage policies that use vSAN capabilities.
- The default vSAN storage policy is used, unless you select a different storage policy.
- You use the vSphere Client to view, create, and modify policies.

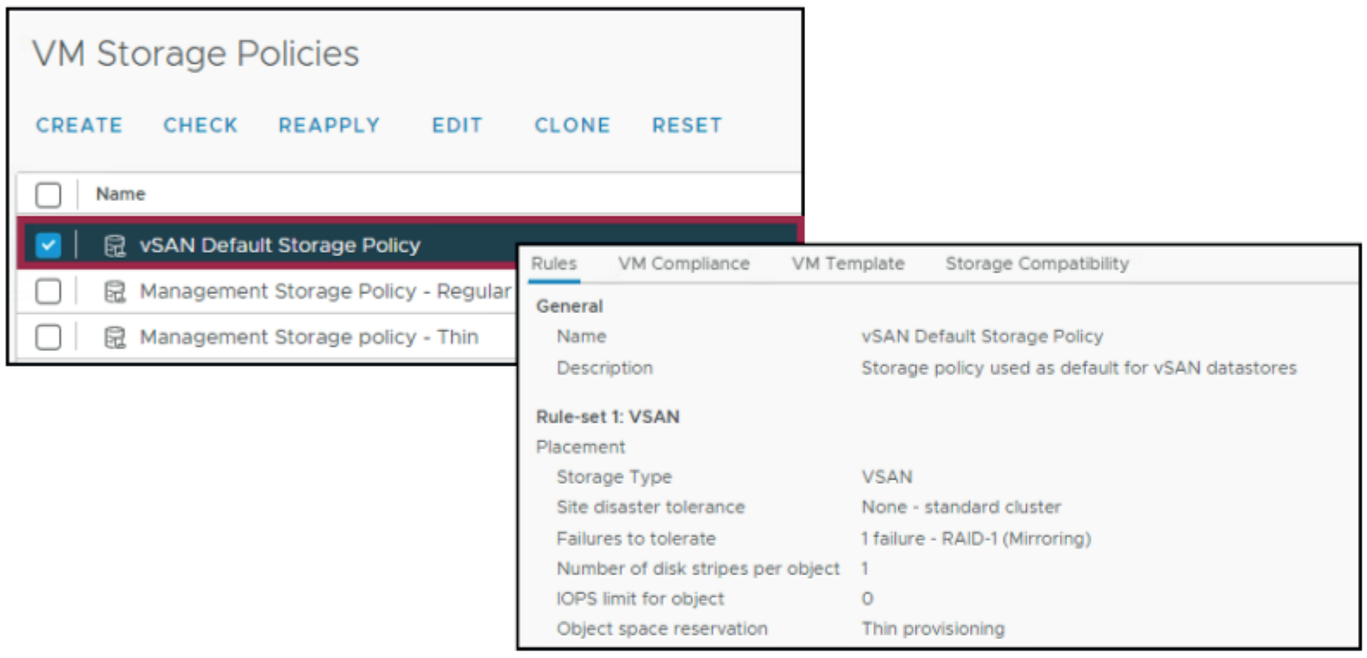

VMs that are deployed to vSAN datastores are assigned at least one VM storage policy.

If a storage policy is not assigned to the VM that is provisioned, a default storage policy is applied to the VM from the datastore.

Technet24

5-120 Creating vSAN Storage Policies

To create a vSAN storage policy, you must select rules for the vSAN storage type.

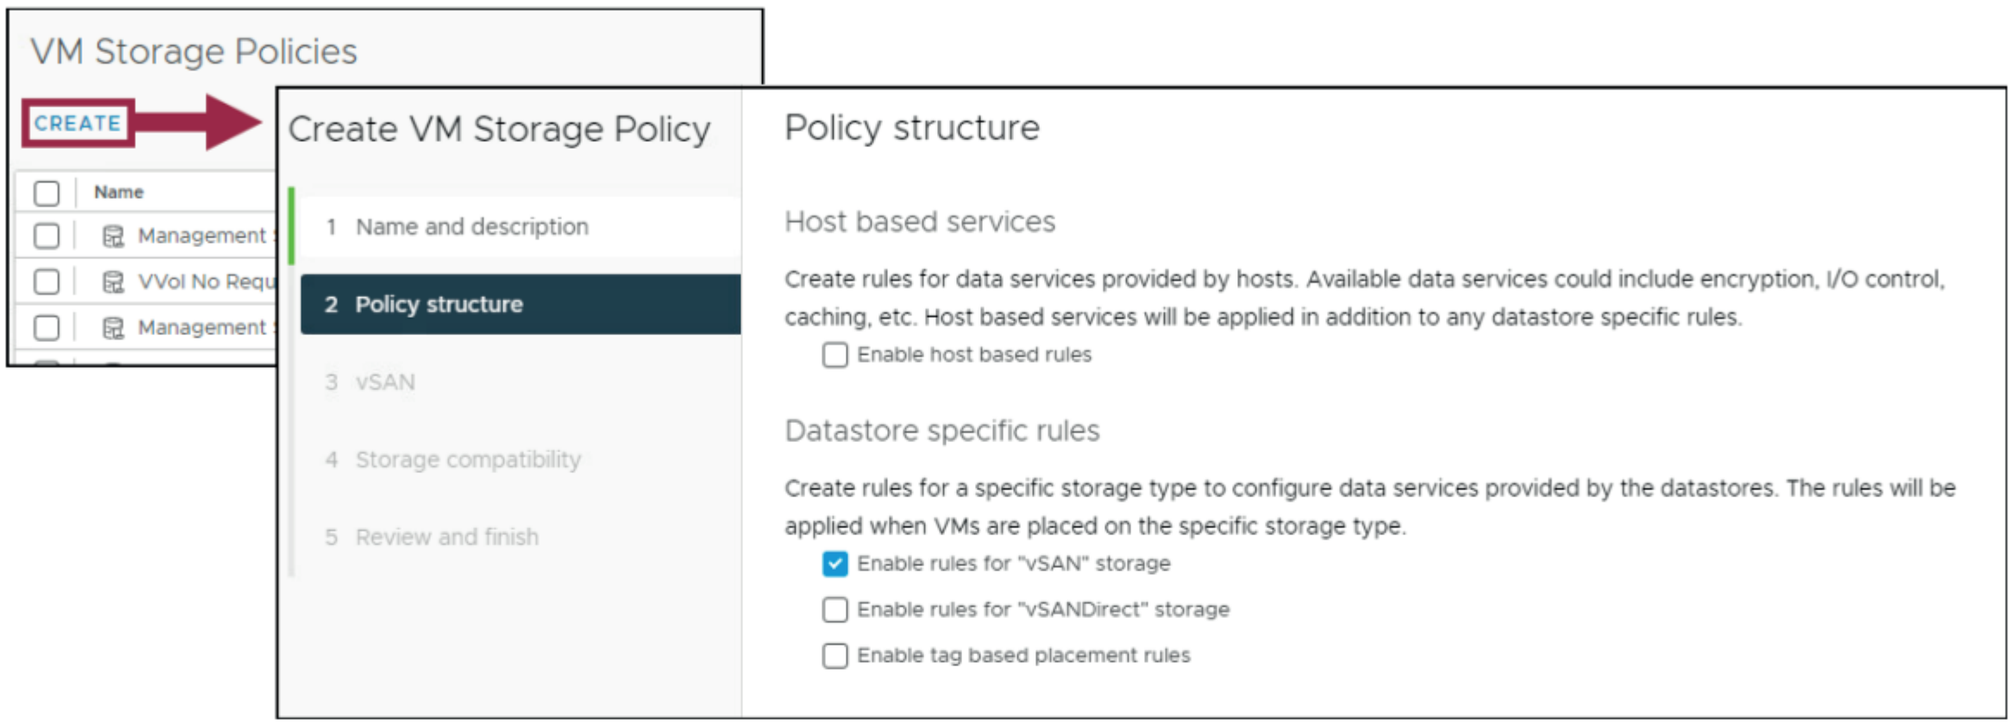

5-121 vSAN Capabilities (1)

You define the rule set or set of capabilities. On the **Availability** tab, you can define how data is protected within a single site cluster and within a stretched cluster.

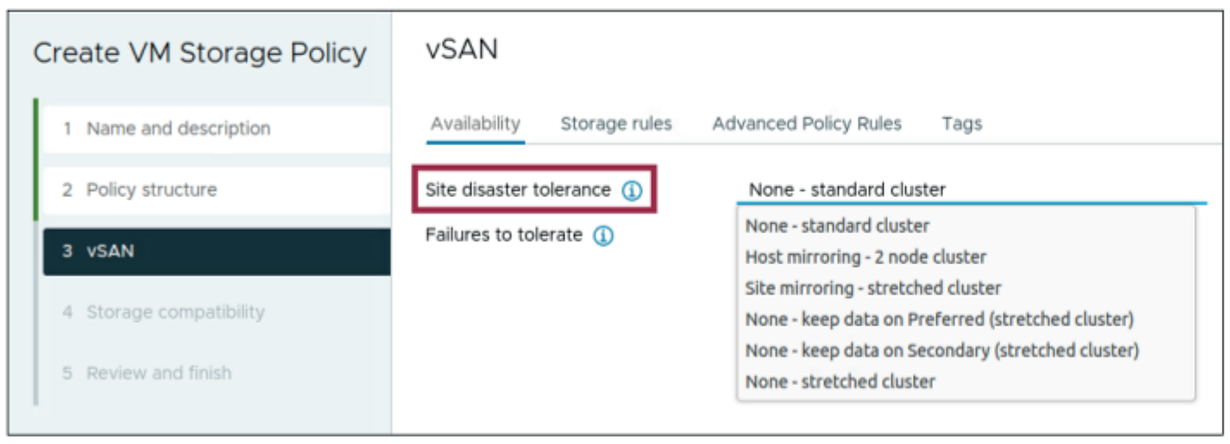

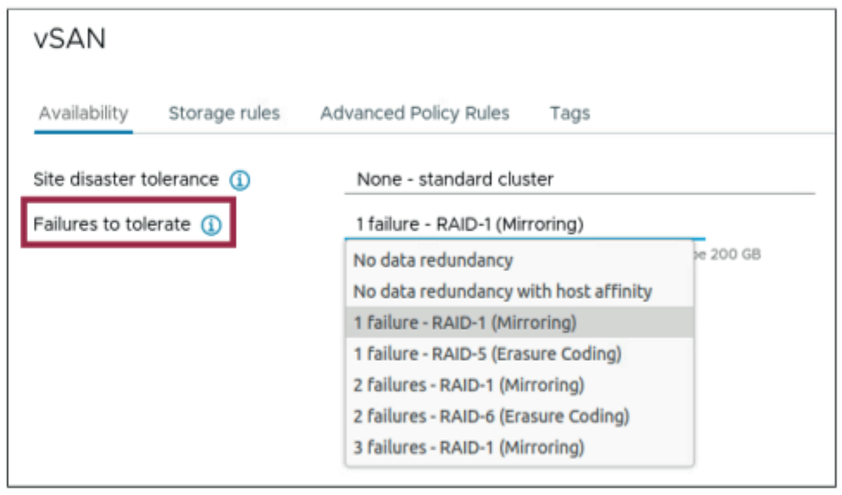

5-122 vSAN Capabilities (2)

On the **Advanced Policy Rules** tab, you can configure one or more advanced vSAN options.

On the **Tags** tab, you can create a tag rule that uses storage associated with a certain tag.

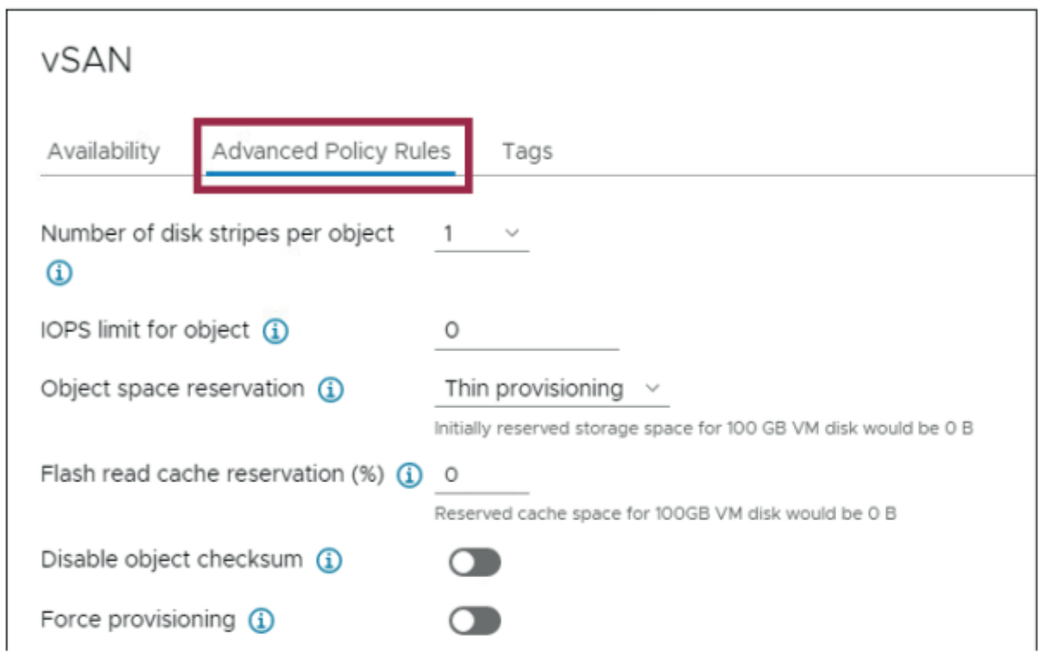

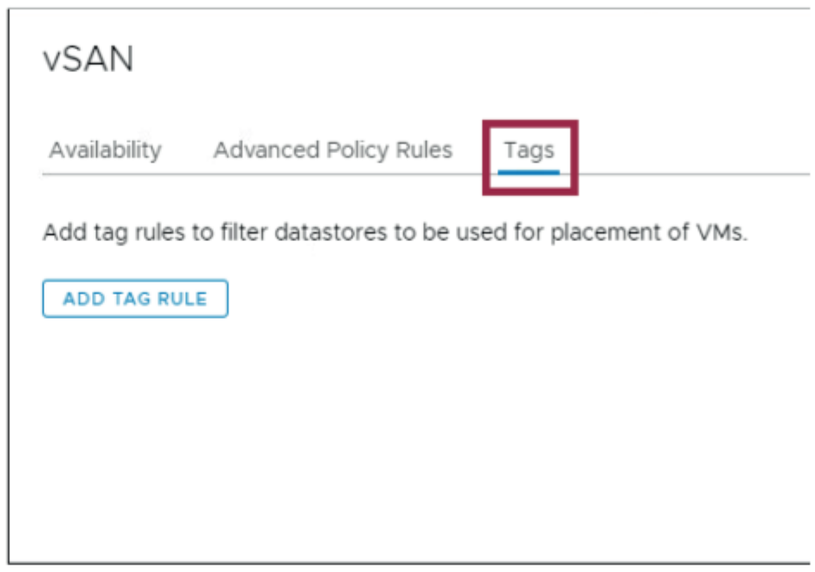

For example, you can use tags to create user-defined storage capabilities and reference them when defining a storage policy for a virtual machine.

For more information about vSAN advanced options, see *Administering VMware vSAN* at https://docs.vmware.com/en/VMware-vSphere/8.0/vsan-80-administration-guide.pdf.

5-123 Lab 12: Creating vSAN Storage Policies

Create and review vSAN storage policies:

1. Examine the Default Storage Policy
2. Create a Custom Policy with No Failure Tolerance
3. Assign the Custom Policy to a VM
4. Make the VM Compliant

5-124 Review of Learner Objectives

- Create vSAN storage policies

Technet24

5-125 Lesson 10: Storage Policies for vSphere Virtual Volumes

5-126 Learner Objectives

- Recognize components in the vSphere Virtual Volumes architecture
- Explain the rule set for creating vSphere Virtual Volumes storage policies

5-127 About vSphere Virtual Volumes

With vSphere Virtual Volumes, storage containers replace traditional storage volumes based on LUNs or NFS shares.

While storage arrays have traditionally provided services like replication, deduplication, and encryption at the LUN or NFS-share level (that is at the VMFS or NFS datastore level), vSphere Virtual Volumes enables array-based features to be applied to individual VMs and virtual disks.

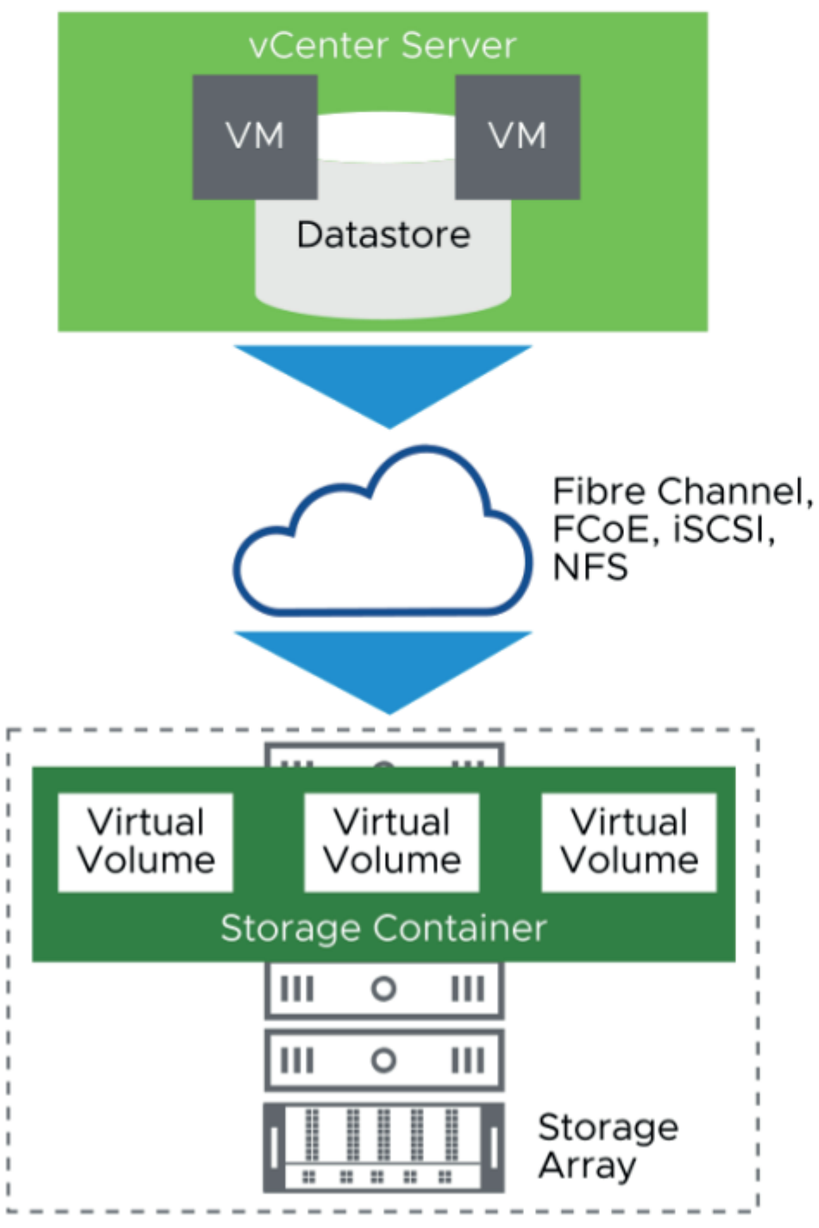

vSphere Virtual Volumes functionality changes the storage management paradigm from managing space inside datastores to managing abstract storage objects handled by storage arrays.

Traditionally, vSphere storage management uses a datastore-centric approach. With this approach, storage administrators and vSphere administrators discuss in advance the underlying storage requirements for VMs. The storage administrator sets up LUNs or NFS shares and presents them to ESXi hosts. The vSphere administrator creates datastores based on LUNs or NFS shares and uses these datastores as VM storage. Typically, the datastore is the lowest level at which data management occurs, from a storage perspective. However, a single datastore contains multiple VMs that might have different requirements.

With the traditional approach, differentiation per VM is difficult.

With vSphere Virtual Volumes, you can differentiate VM services per application by offering a new approach to storage management.

vSphere Virtual Volumes supports Fibre Channel (FC), FCoE, iSCSI, and NFS.

5-128 vSphere Virtual Volumes Architecture

vSphere Virtual Volumes Architecture is an integration and management framework that virtualizes storage arrays, enabling a more efficient operational model that is optimized for virtualized environments and centered on the application instead of the infrastructure.

◊ The following components constitute the vSphere Virtual Volumes architecture:

- vSphere Virtual Volumes storage providers
- Protocol endpoints
- Storage containers
- Virtual volumes datastores
- Virtual volumes

Technet24

5-129 vSphere Virtual Volumes Storage Providers

A vSphere Virtual Volumes storage provider is a software component that mediates communication between vSphere and a storage system.

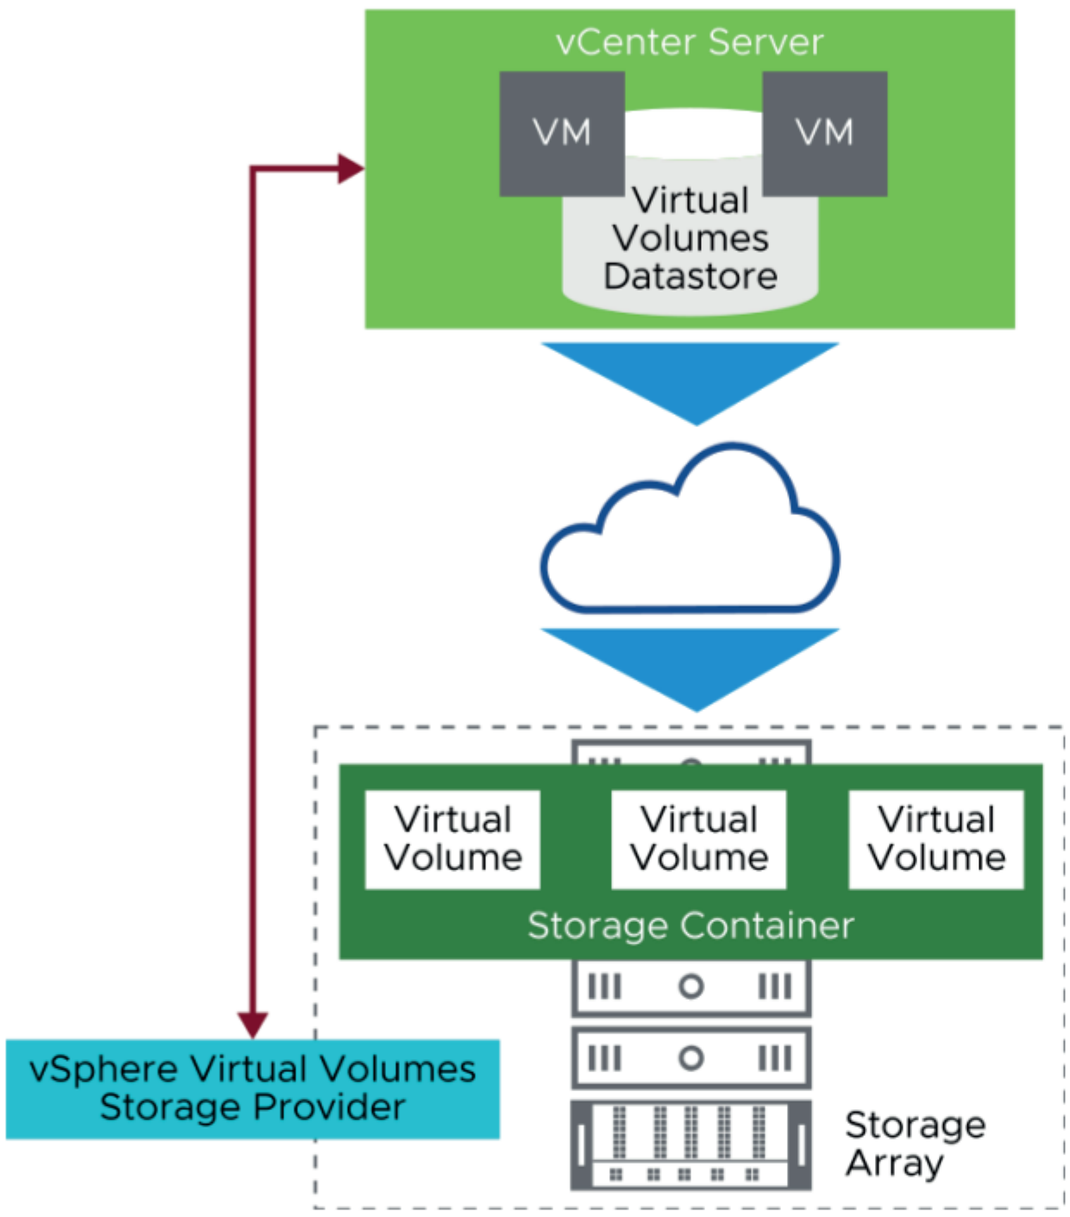

A vSphere Virtual Volumes storage provider is implemented through vSphere API for Storage Awareness and is used to manage all aspects of vSphere Virtual Volumes storage. The storage provider integrates with the Storage Monitoring Service, with vSphere, to communicate with vCenter Server and ESXi hosts.

The storage provider communicates VM storage requirements, which you define in a storage policy, to the storage layer. In this way, a virtual volume created in the storage layer meets the requirements outlined in the policy.

Typically, vendors supply storage providers that can integrate with vSphere and provide support to virtual volumes. Every storage provider must be certified by VMware and properly deployed. For information about deploying the vSphere Virtual Volumes storage provider, contact your storage vendor.

For information about vSphere Virtual Volumes storage partners, see the vSphere Virtual Volumes product page at http://www.vmware.com/products/vsphere/features/virtual-volumes.html.

5-130 Protocol Endpoints

ESXi uses protocol endpoints to establish a data path on demand from VMs to their respective virtual volumes.

Protocol endpoints are configured on the storage system, including one or several per storage container.

Each virtual volume is bound to a specific protocol endpoint. A single protocol endpoint can connect to many virtual volumes.

Protocol Endpoint supports SCSI (iSCSI/FC/FCoE) and NFS).

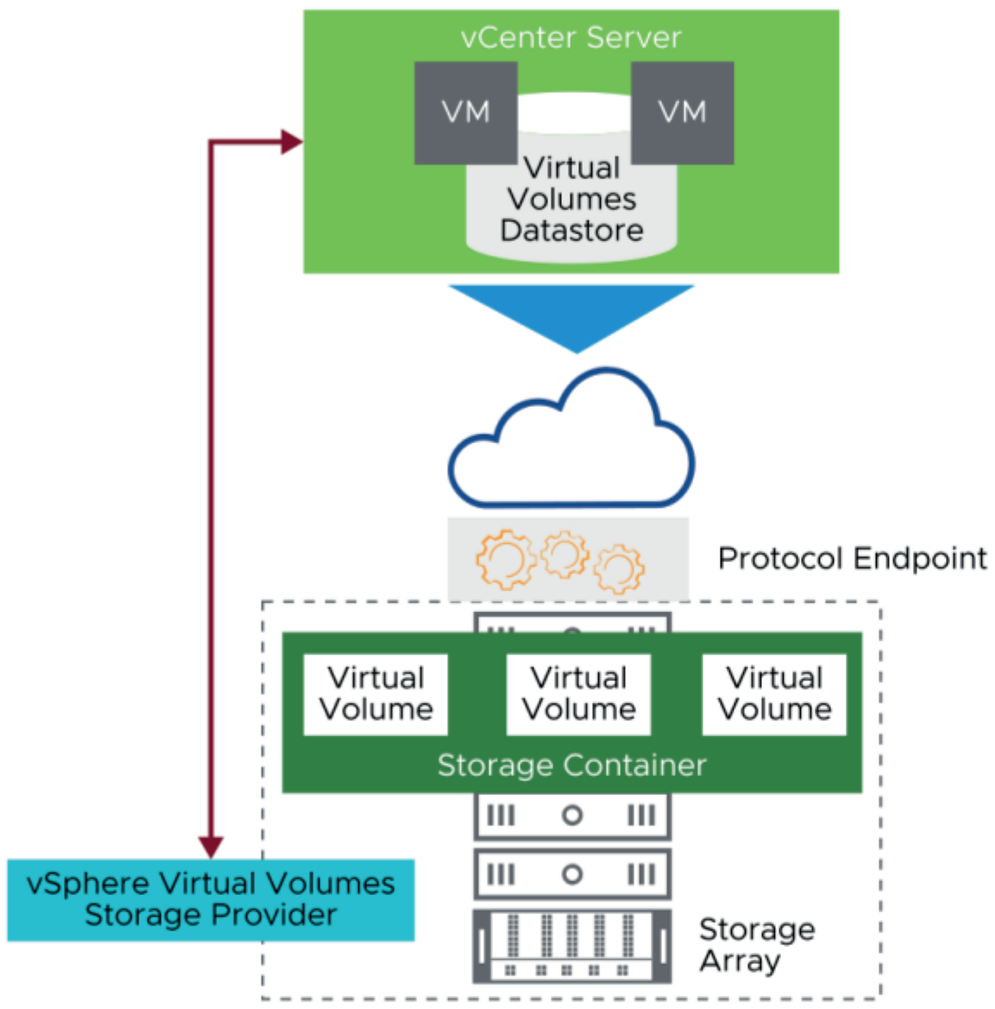

Although storage systems manage all aspects of virtual volumes, ESXi hosts have no direct access to virtual volumes on the storage side. Instead, ESXi hosts use a logical I/O proxy, called the protocol endpoint, to communicate with virtual volumes and virtual disk files that virtual volumes encapsulate.

When the SCSI-based protocol is used, the protocol endpoint represents a LUN defined by a T10-based LUN World Wide Name. For the NFS protocol, the protocol endpoint is a mount point, such as an IP address or a DNS name and a share name.

When a VM on the host performs an I/O operation, the protocol endpoint directs the I/O to the appropriate virtual volume. Typically, a storage system requires a few protocol endpoints.

A storage administrator configures protocol endpoints. Protocol endpoints are part of physical storage and are exported, with associated storage containers, by the storage system through a storage provider. After you map a storage container to a virtual volumes datastore, protocol endpoints are discovered by ESXi and appear in the vSphere Client. Protocol endpoints can also be discovered during a storage rescan.

Technet24

5-131 Storage Containers and vSphere Virtual Volumes Datastores

Virtual volumes reside on a storage container.

A storage container is a logical construct that can span the capacity of an entire storage array, and logically groups virtual volumes based on management and administrative needs.

For example, a storage container can contain all virtual volumes for a particular organization.

A virtual volumes datastore maps to a specific storage container.

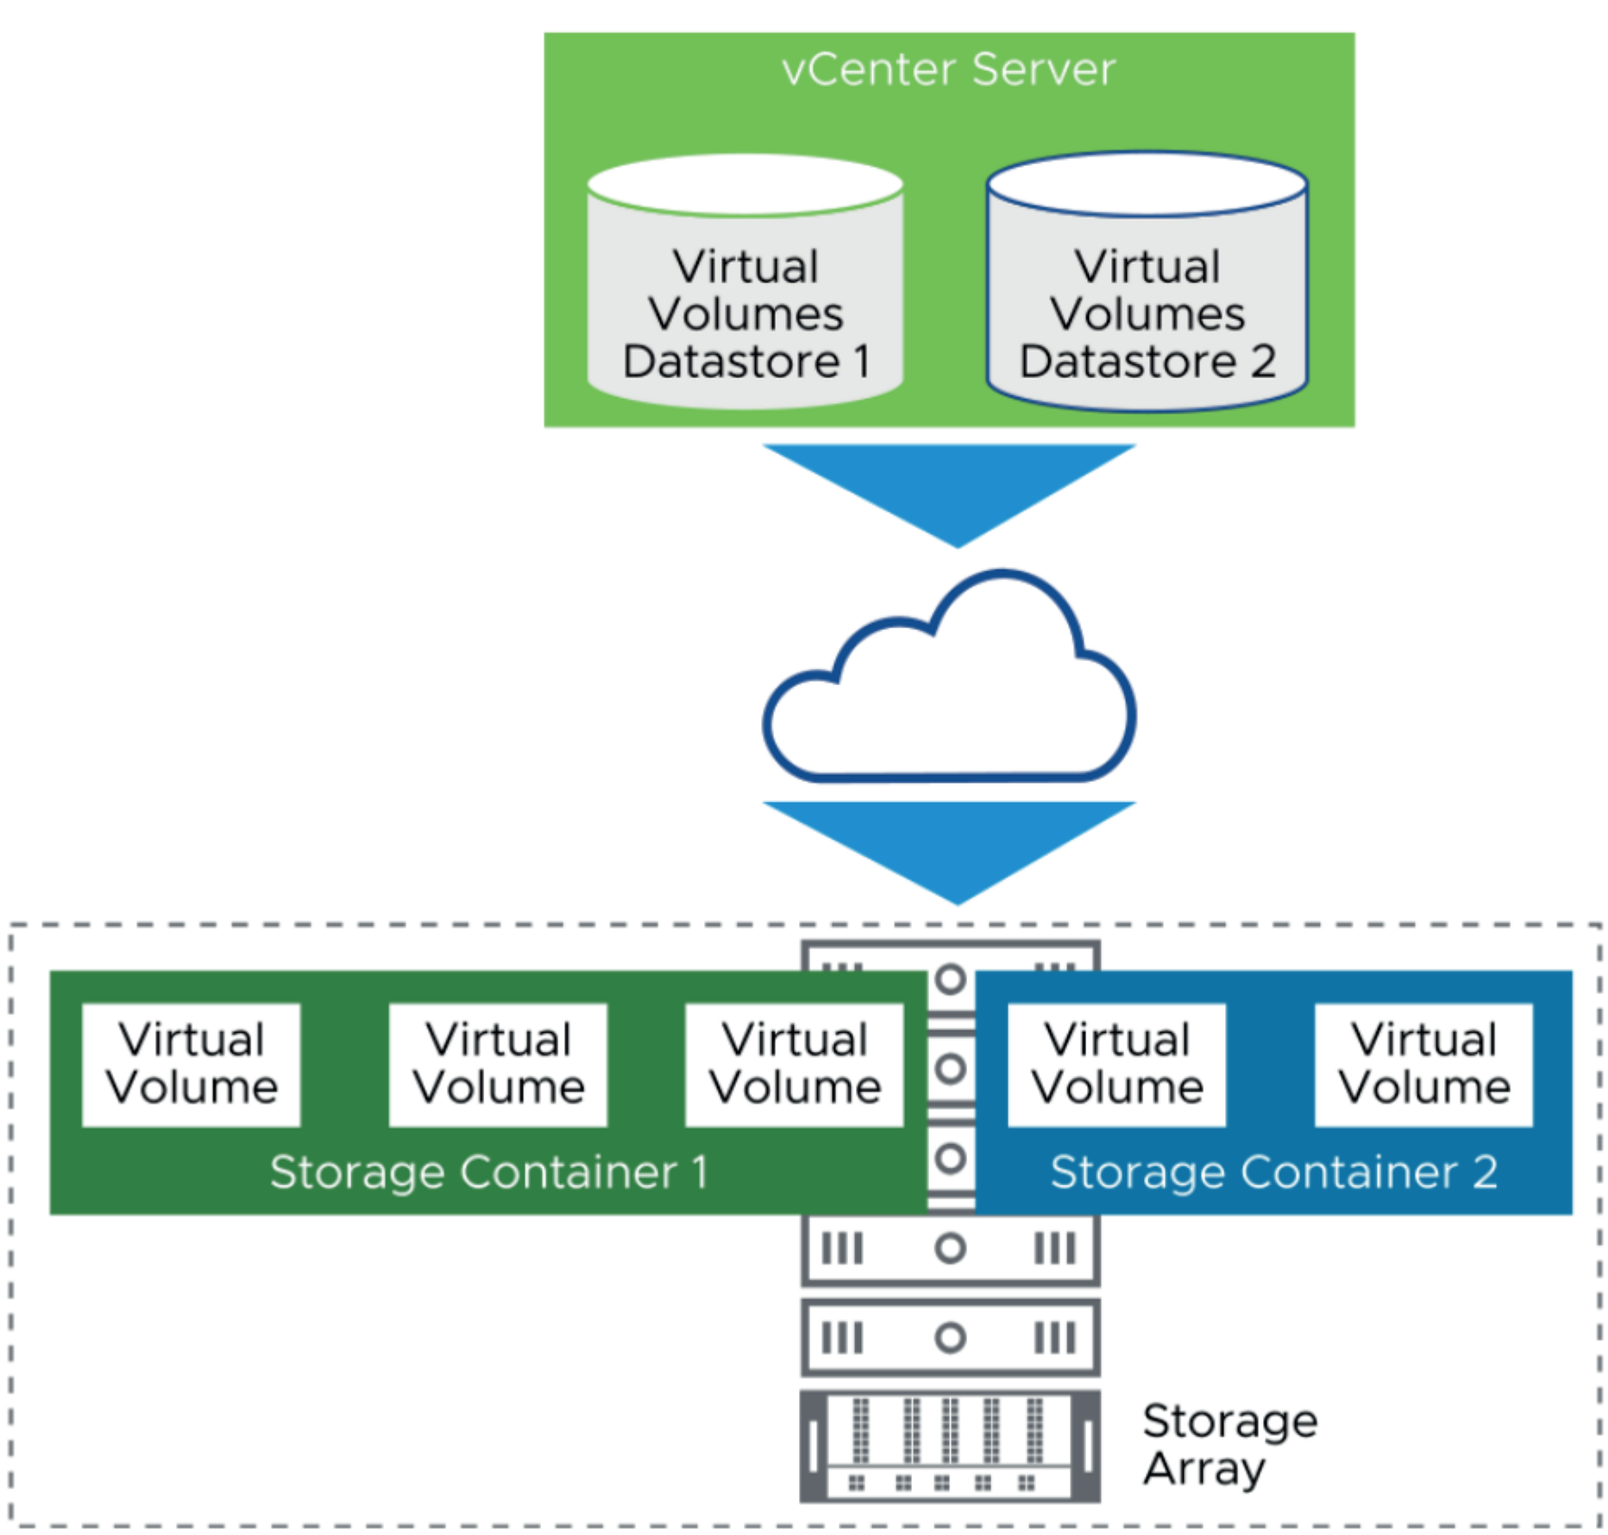

Unlike traditional LUN-based and NFS-based vSphere storage, vSphere Virtual Volumes functionality does not require preconfigured volumes on a storage side. Instead, vSphere Virtual Volumes uses a storage container, which is a pool of raw storage capacity or an aggregation of storage capabilities that a storage system can provide to virtual volumes.

Typically, a storage administrator on the storage side defines storage containers. The number of storage containers and their capacity depend on the vendor-specific implementation, but each storage system requires at least one container.

After vCenter Server discovers storage containers presented by storage systems, you must map a storage container to a virtual volumes datastore. The virtual volumes datastore that you create corresponds directly to the specific storage container and becomes the container's representation in vCenter Server and the vSphere Client.

5-132 Virtual Volumes

Virtual volumes are objects that are stored natively inside a storage system that is connected through an Ethernet or a Fibre Channel network.

Virtual volumes are not preprovisioned but are created automatically by the storage array when you perform virtual management operations.

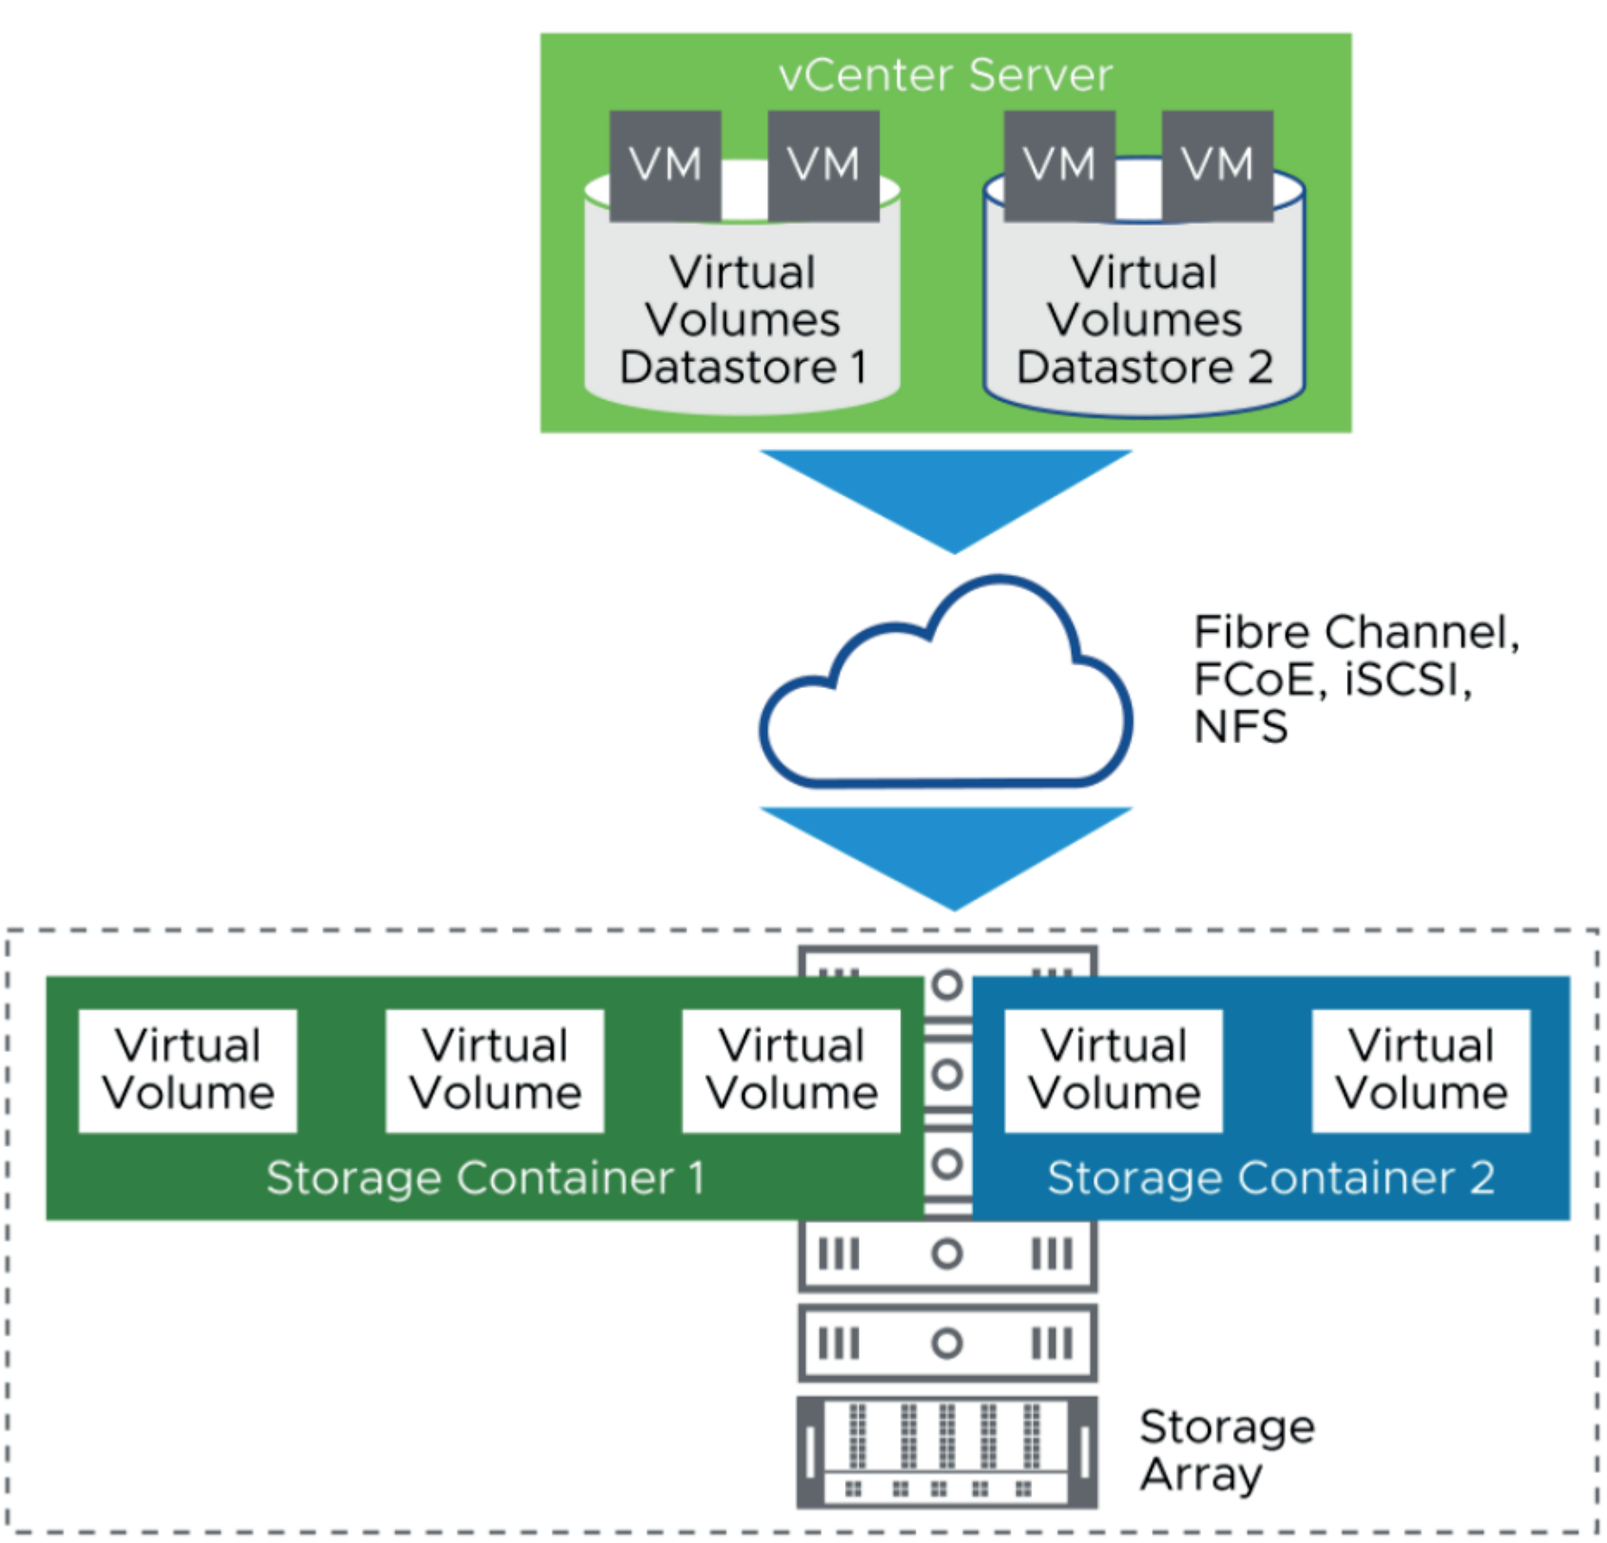

Virtual volumes are encapsulations of VM files, virtual disks, and their derivatives.

Virtual volumes are exported as objects by a compliant storage system and are backed by storage devices in the array. Typically, a unique GUID identifies a virtual volume.

Technet24

5-133 Mapping VM Files to Virtual Volumes

The storage system creates virtual volumes for the core elements that make up the VM:

- Configuration of virtual volumes representing virtual disks
- Data virtual volumes containing a virtual machine's configuration files

Additional virtual volumes can be created for other VM components, such as snapshots and swap files.

vSphere Virtual Volumes maps objects directly to virtual volumes. vSphere can then offload intensive storage operations, such as taking snapshots, cloning, and replication, to the storage system.

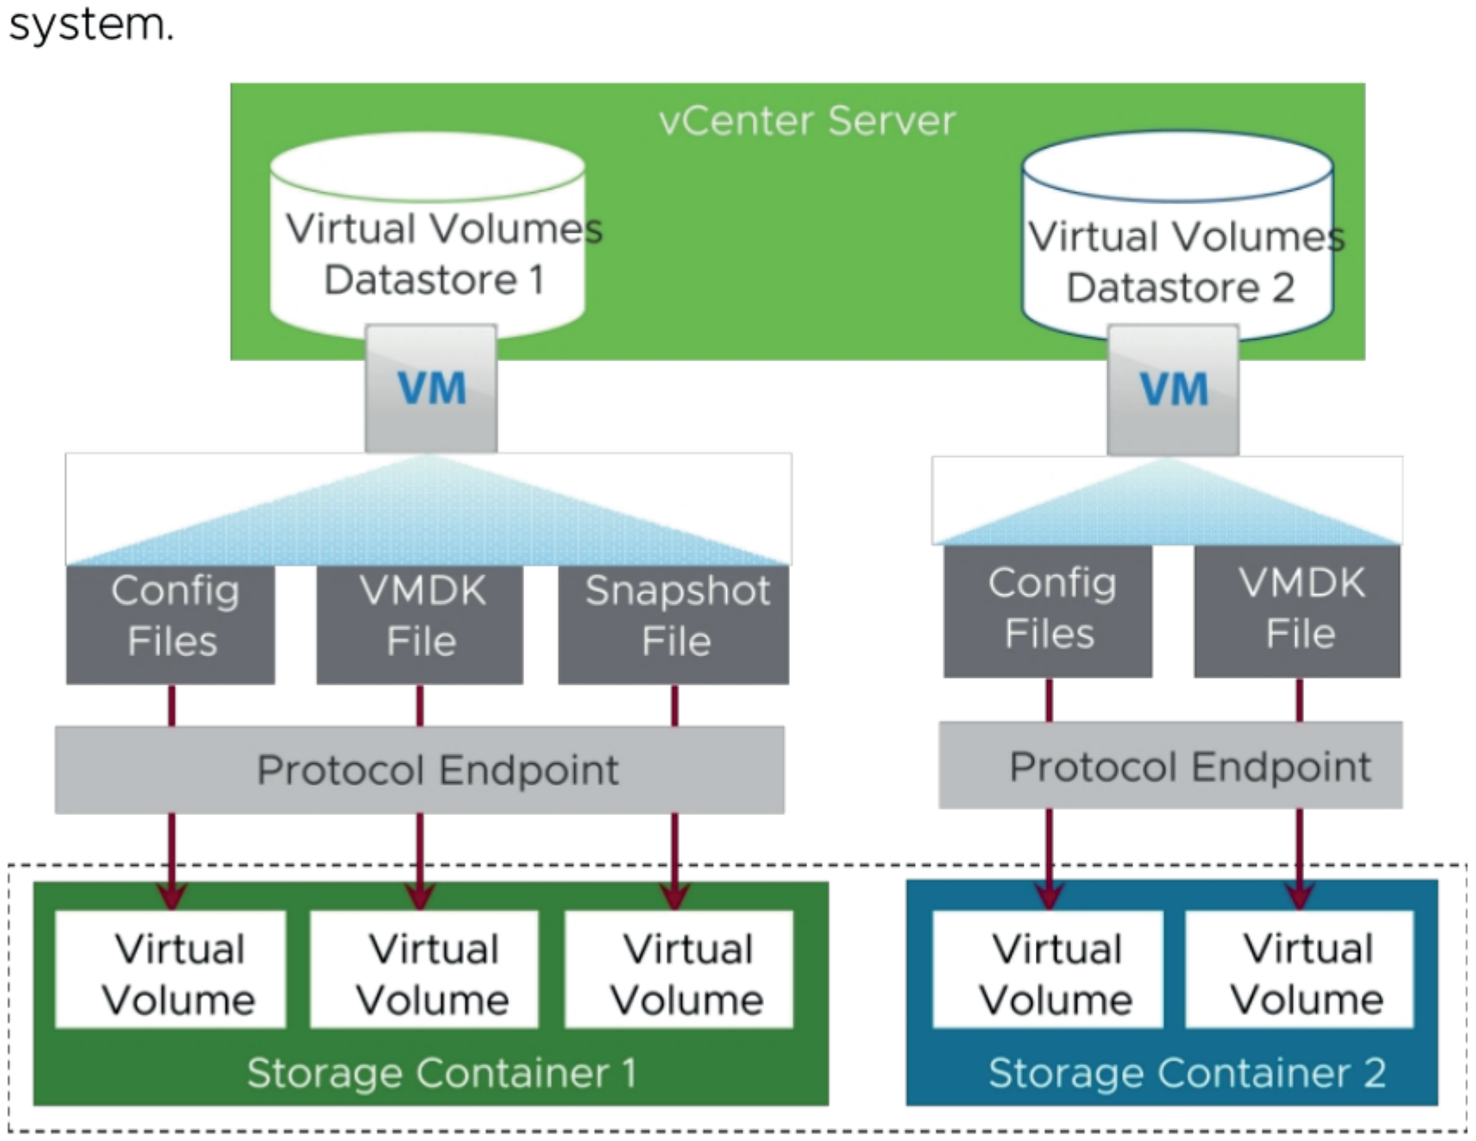

A data virtual volume corresponds directly to each virtual disk (-flat.vmdk) file. Like virtual disk files on traditional datastores, virtual volumes are presented to VMs as SCSI, vSATA, or vNVME disks.

A configuration virtual volume, or a home directory, represents a small directory that contains metadata files for a VM. Metadata files include a VMX file, descriptor files for virtual disks, log files, and so on. The configuration virtual volume is formatted with a file system. When ESXi uses the SCSI protocol to connect to storage, configuration virtual volumes are formatted with VMFS. With the NFS protocol, configuration virtual volumes are presented as an NFS directory.

Additional virtual volumes can be created for other VM components and virtual disk derivatives, such as clones, snapshots, and replicas. These virtual volumes include a swap virtual volume to hold VM swap files and a virtual memory volume to hold the contents of VM memory for a snapsho or when a VM is suspended.

5-134 vSphere Virtual Volumes Replication

vSphere Virtual Volumes supports replication. With this feature, you can offload the replication of VMs to your storage array and use the full replication capabilities of the array.

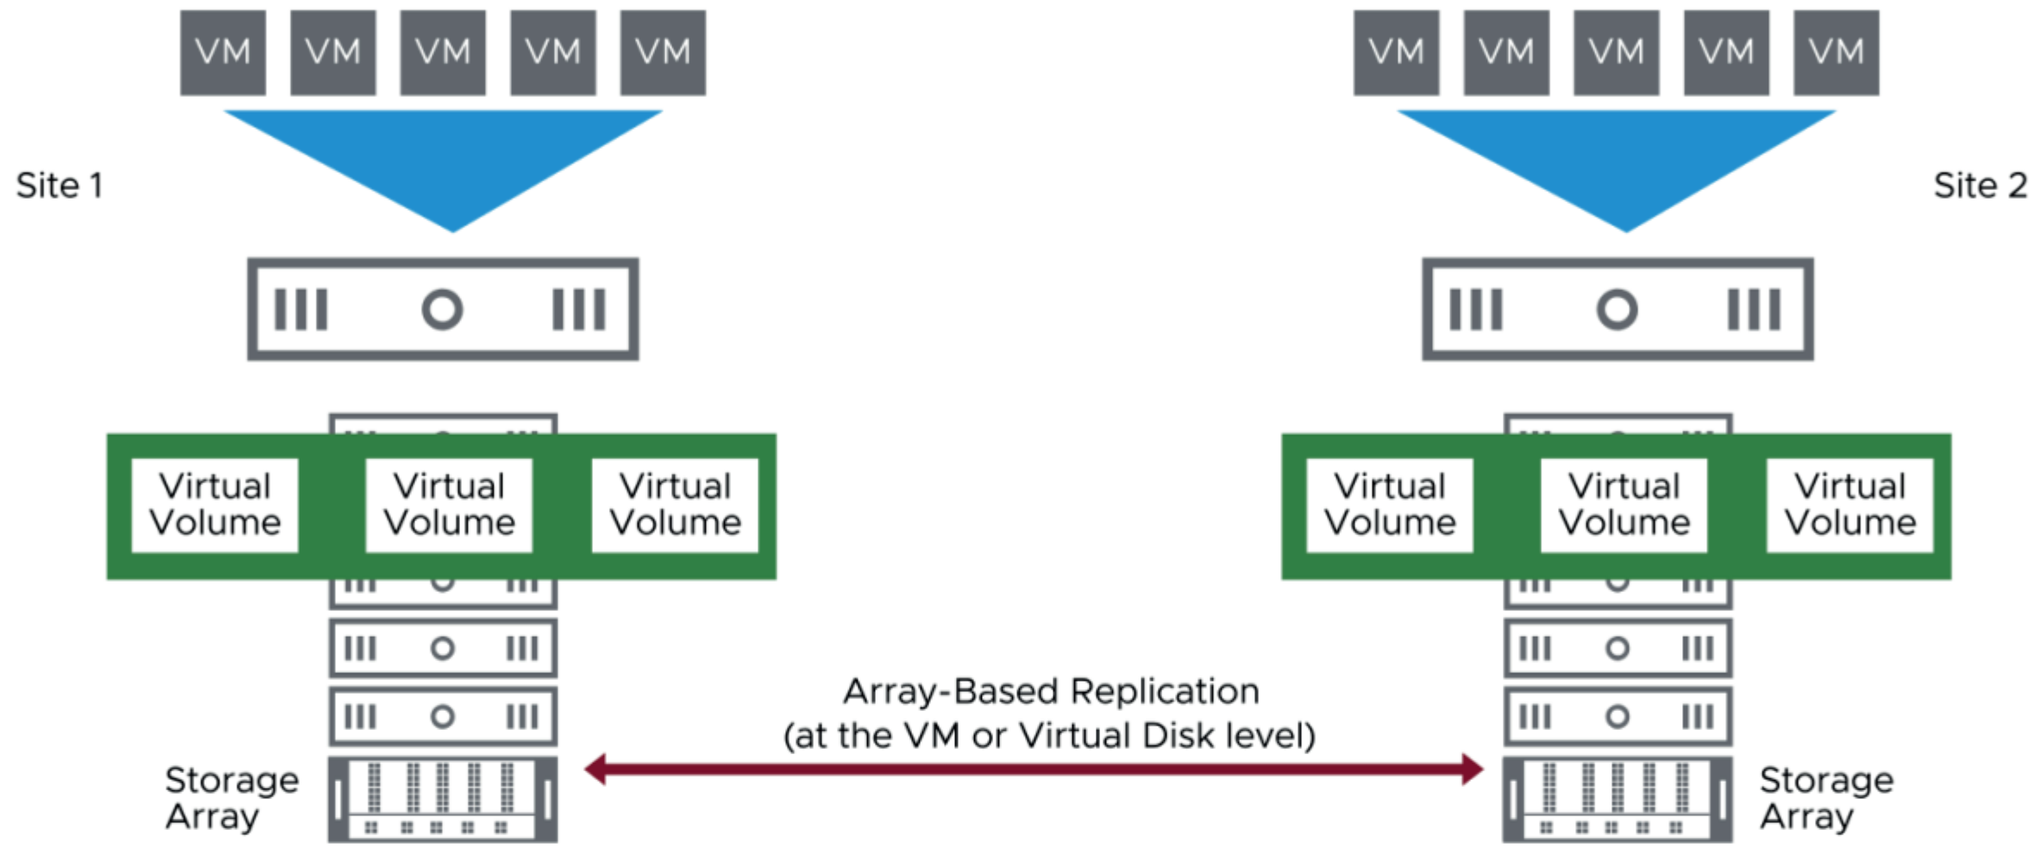

vSphere Virtual Volumes replication has the following properties:

- Available with vSphere API for Storage Awareness version 3.0
- Available for ESXi hosts, version 6.5 and later
- Requires the storage array to be replication-capable

Implementation of vSphere Virtual Volumes replication depends on your array and might vary with different storage vendors. The following requirements generally apply to all vendors:

- The storage arrays that you use to implement replication must be compatible with vSphere Virtual Volumes.
- The arrays must integrate with the version 3.0 or later of the storage (VASA) provider that is compatible with vSphere Virtual Volumes replication.
- The storage arrays must be replication-capable and configured to use vendor-provided replication mechanisms. Typical configurations usually involve one or two replication targets.

  Any required configurations, such as pairing of the replicated site and the target site, must also be performed by the storage system.
- When applicable, replication groups and fault domains for virtual volumes must be preconfigured by the storage system.

Technet24

5-135 Virtual Volume Storage Policies

A VM that runs on a virtual volumes datastore requires a VM storage policy.

If you do not create a VM storage policy compatible with the virtual volumes datastore, the system uses the No Requirements default policy:

- This generic virtual volumes policy contains no rules or storage specifications.
- Storage arrays can determine the optimum placement for the VM objects.

5-136 Virtual Volume Storage Policy: Examples

By using different virtual volumes for different VM components, you can apply storage policies at detailed levels, for example:

- A virtual volume that contains a virtual data disk can have a richer set of services than the virtual volume for the VM boot disk.
- A snapshot virtual volume can use a different storage tier compared to a data virtual volume.
- If information about replication capabilities of the virtual volumes storage array are presented by the storage array's VASA provider, you can activate replication for your business-critical VMs.

5-137 Creating Virtual Volumes Storage Policies (1)

A virtual volumes storage policy can include host-based service rules, VAIO, and datastore-specific rules.

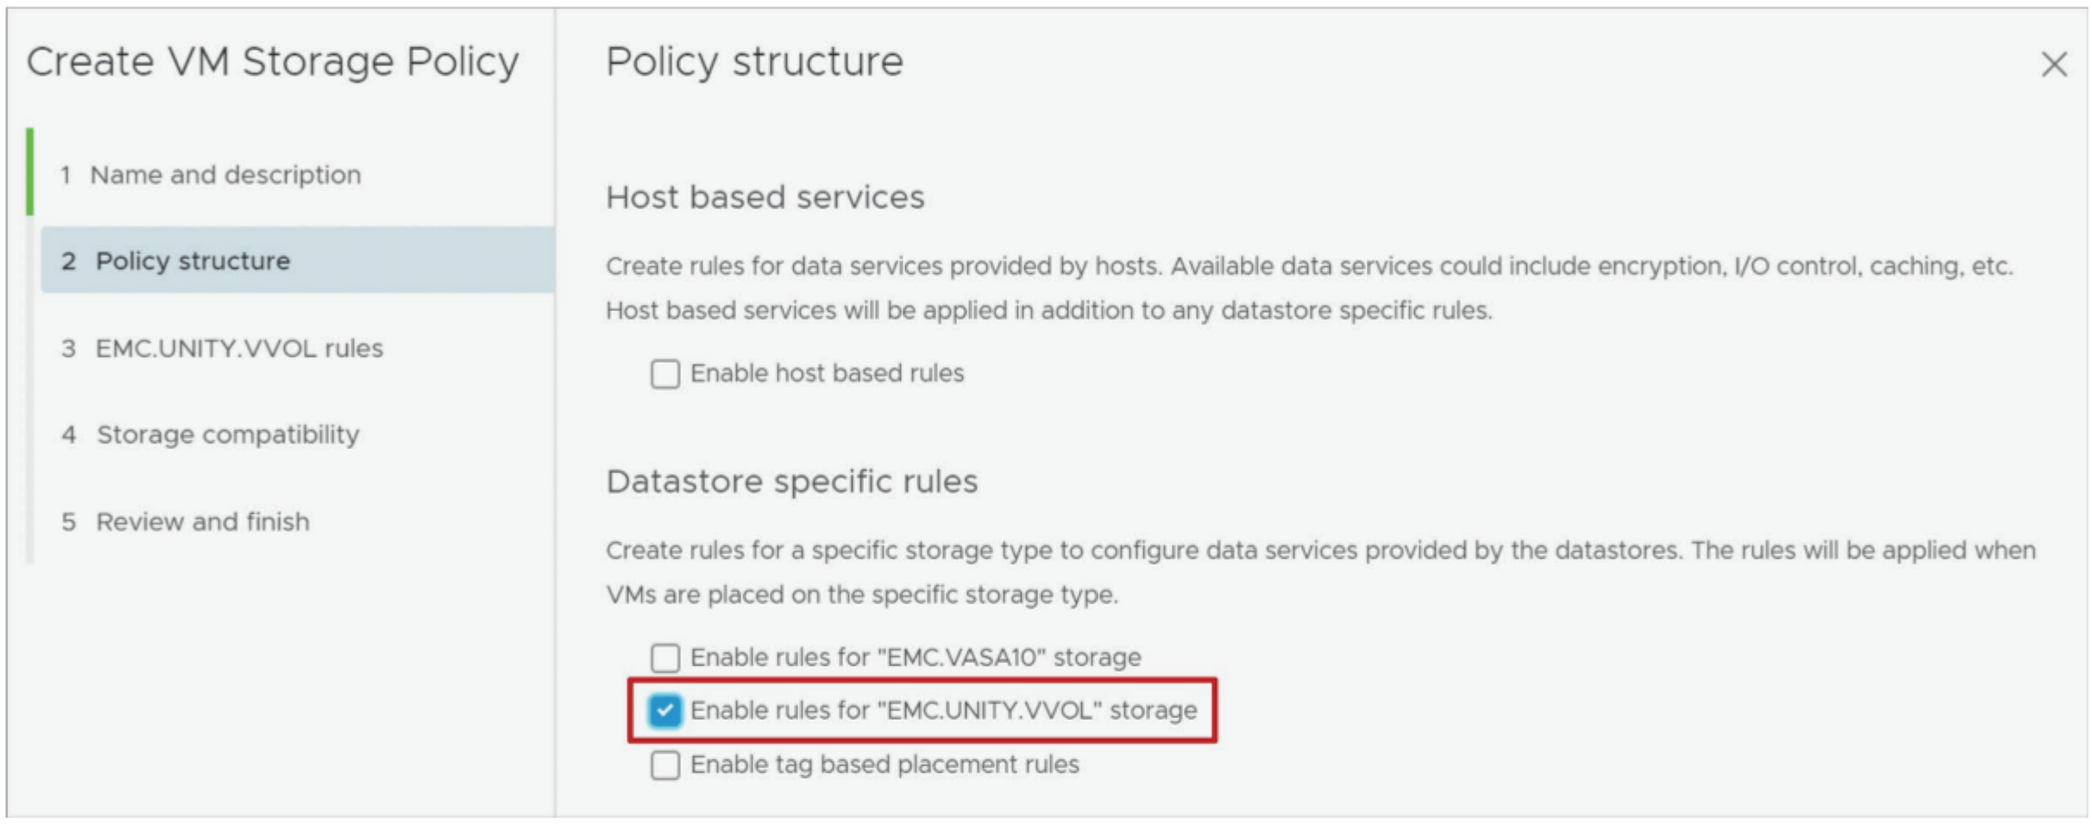

Technet24

5-138 Creating Virtual Volumes Storage Policies (2)

In this example, datastore-specific rules are listed for the EMC Unity virtual volumes storage.

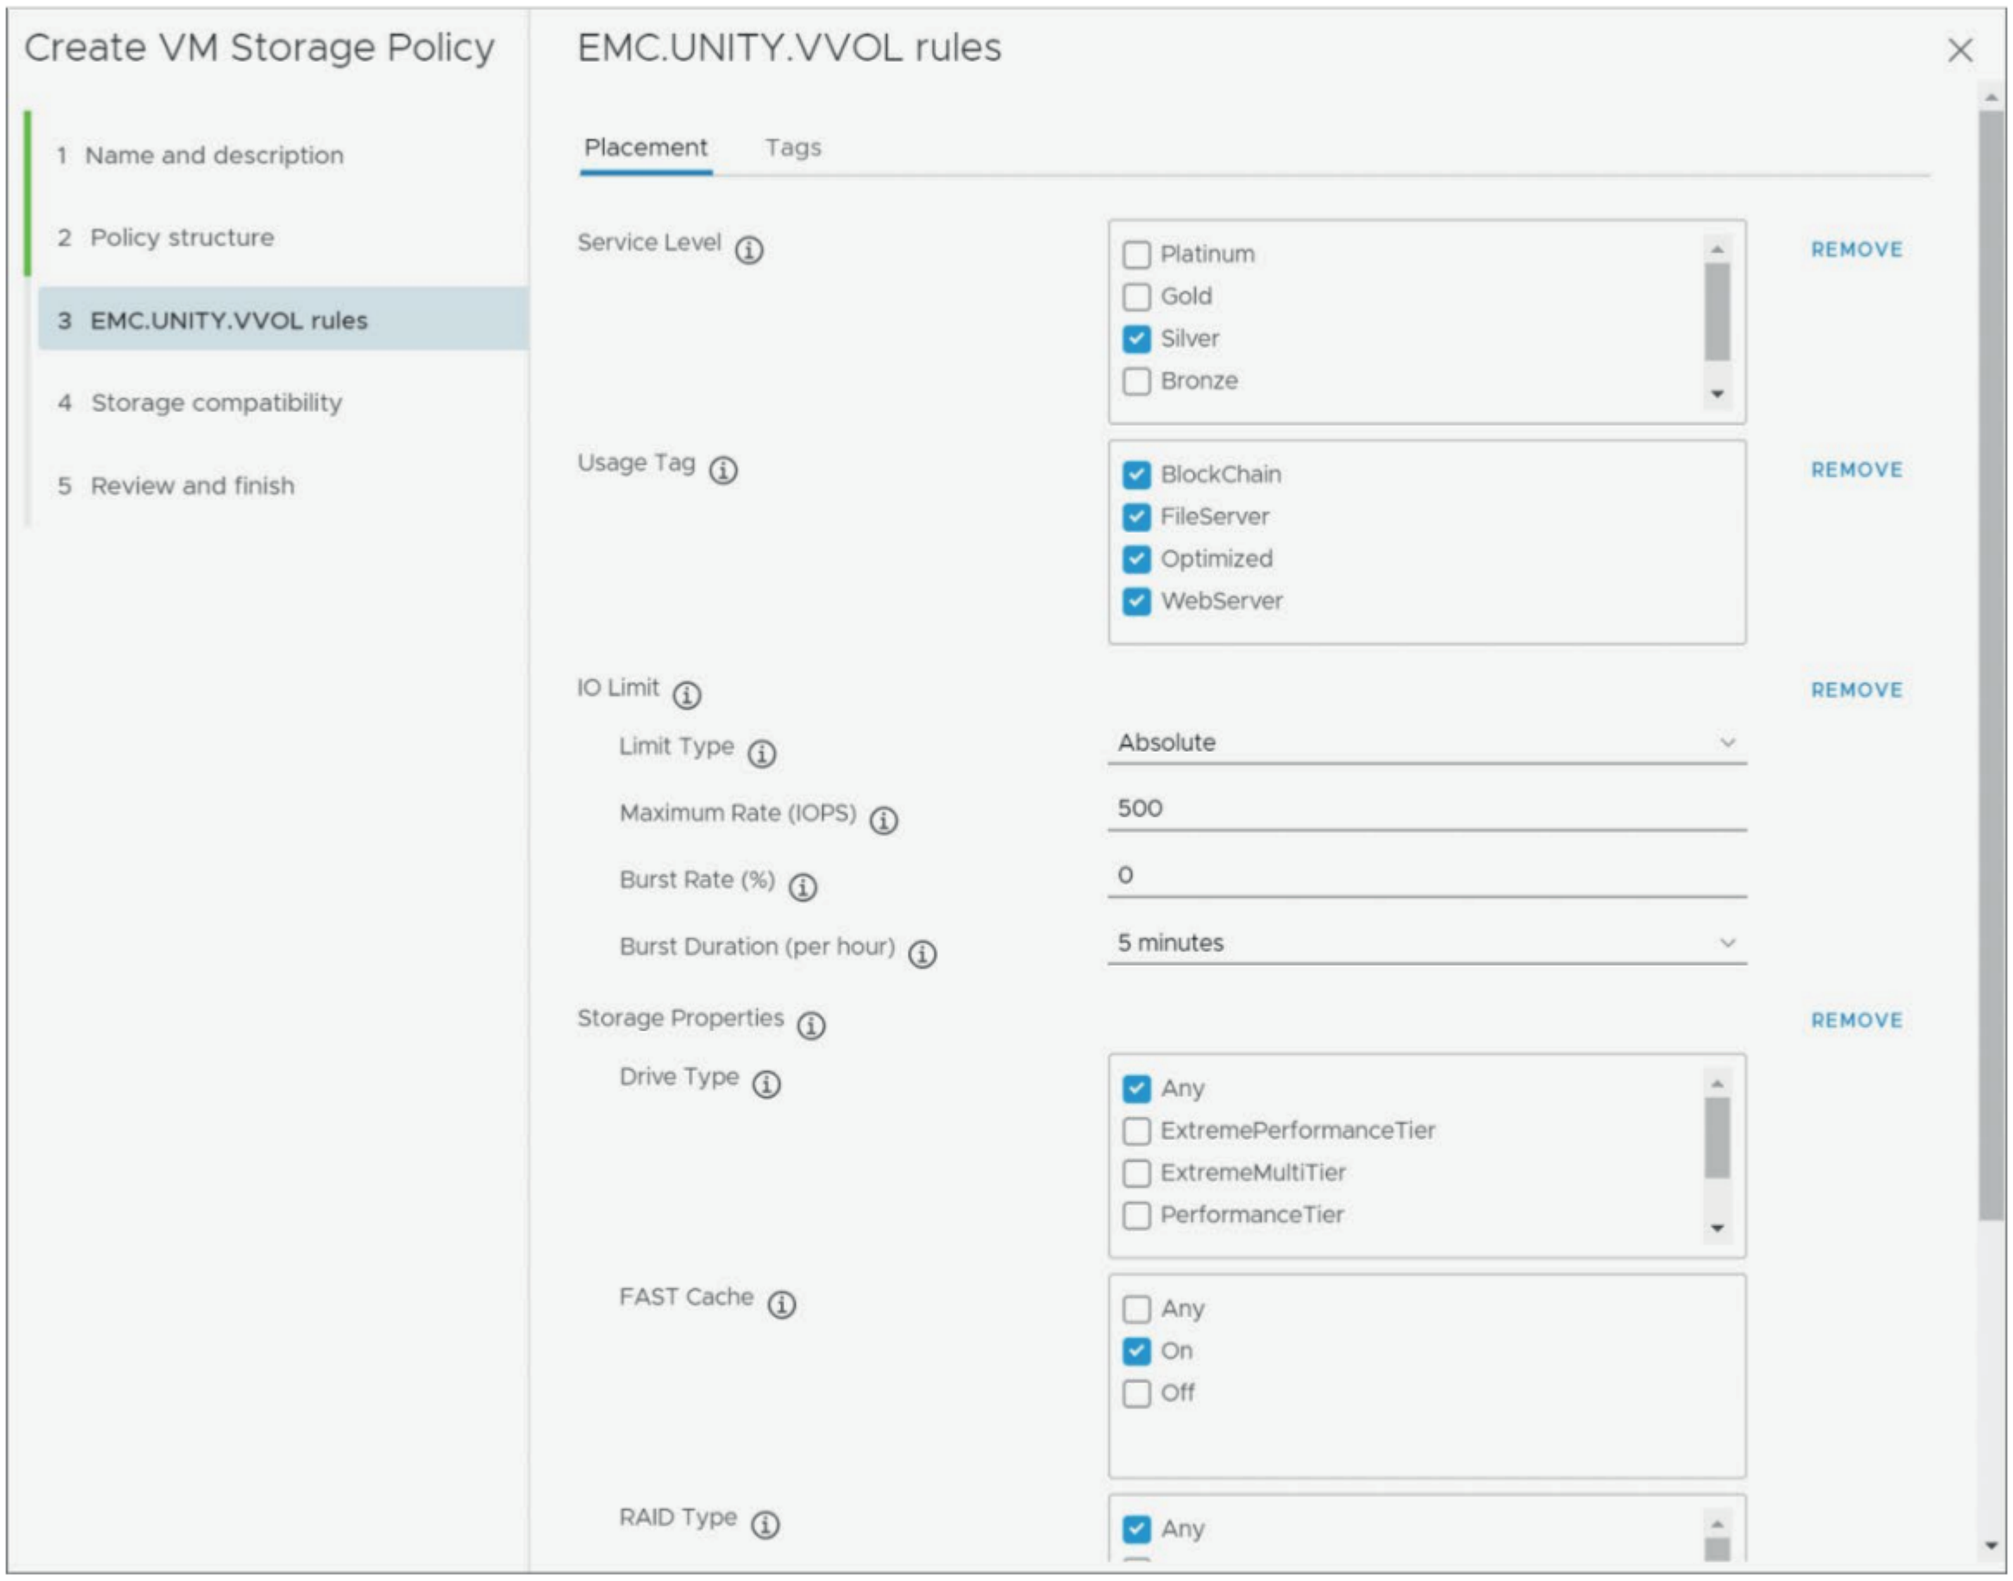

5-139 Activity: Displaying Storage Capabilities (1)

If the capabilities of your virtual volumes storage array do not appear in the storage policy interface in the vSphere Client, what should you check?

◊ ❒ That tags are created for the virtual volumes storage array

◊ ❒ That virtual volumes storage policy components are preconfigured

◊ ❒ That the virtual volumes storage provider is registered with vCenter Server

◊ ❒ That the VVol No Requirements storage policy is not deleted

5-140 Activity: Displaying Storage Capabilities (2)

If the capabilities of your virtual volumes storage array do not appear in the storage policy interface in the vSphere Client, what should you check?

◊ ❒ That tags are created for the virtual volumes storage array

◊ ❒ That the virtual volumes storage policy components are preconfigured

◊ ✓ That the virtual volumes storage provider is registered with vCenter Server

◊ ❒ That the VVol No Requirements storage policy is not deleted

5-141 Review of Learner Objectives

- Recognize components in the vSphere Virtual Volumes architecture
- Explain the rule set for creating vSphere Virtual Volumes storage policies

Technet24

5-142 Lesson 11: Storage I/O Control

5-143 Learner Objectives

- Describe how Storage I/O Control balances the I/O load on a datastore
- Configure Storage I/O Control

5-144 About Storage I/O Control

Storage I/O Control provides cluster-wide storage I/O prioritization, which leads to better workload consolidation and helps reduce extra costs associated with overprovisioning.

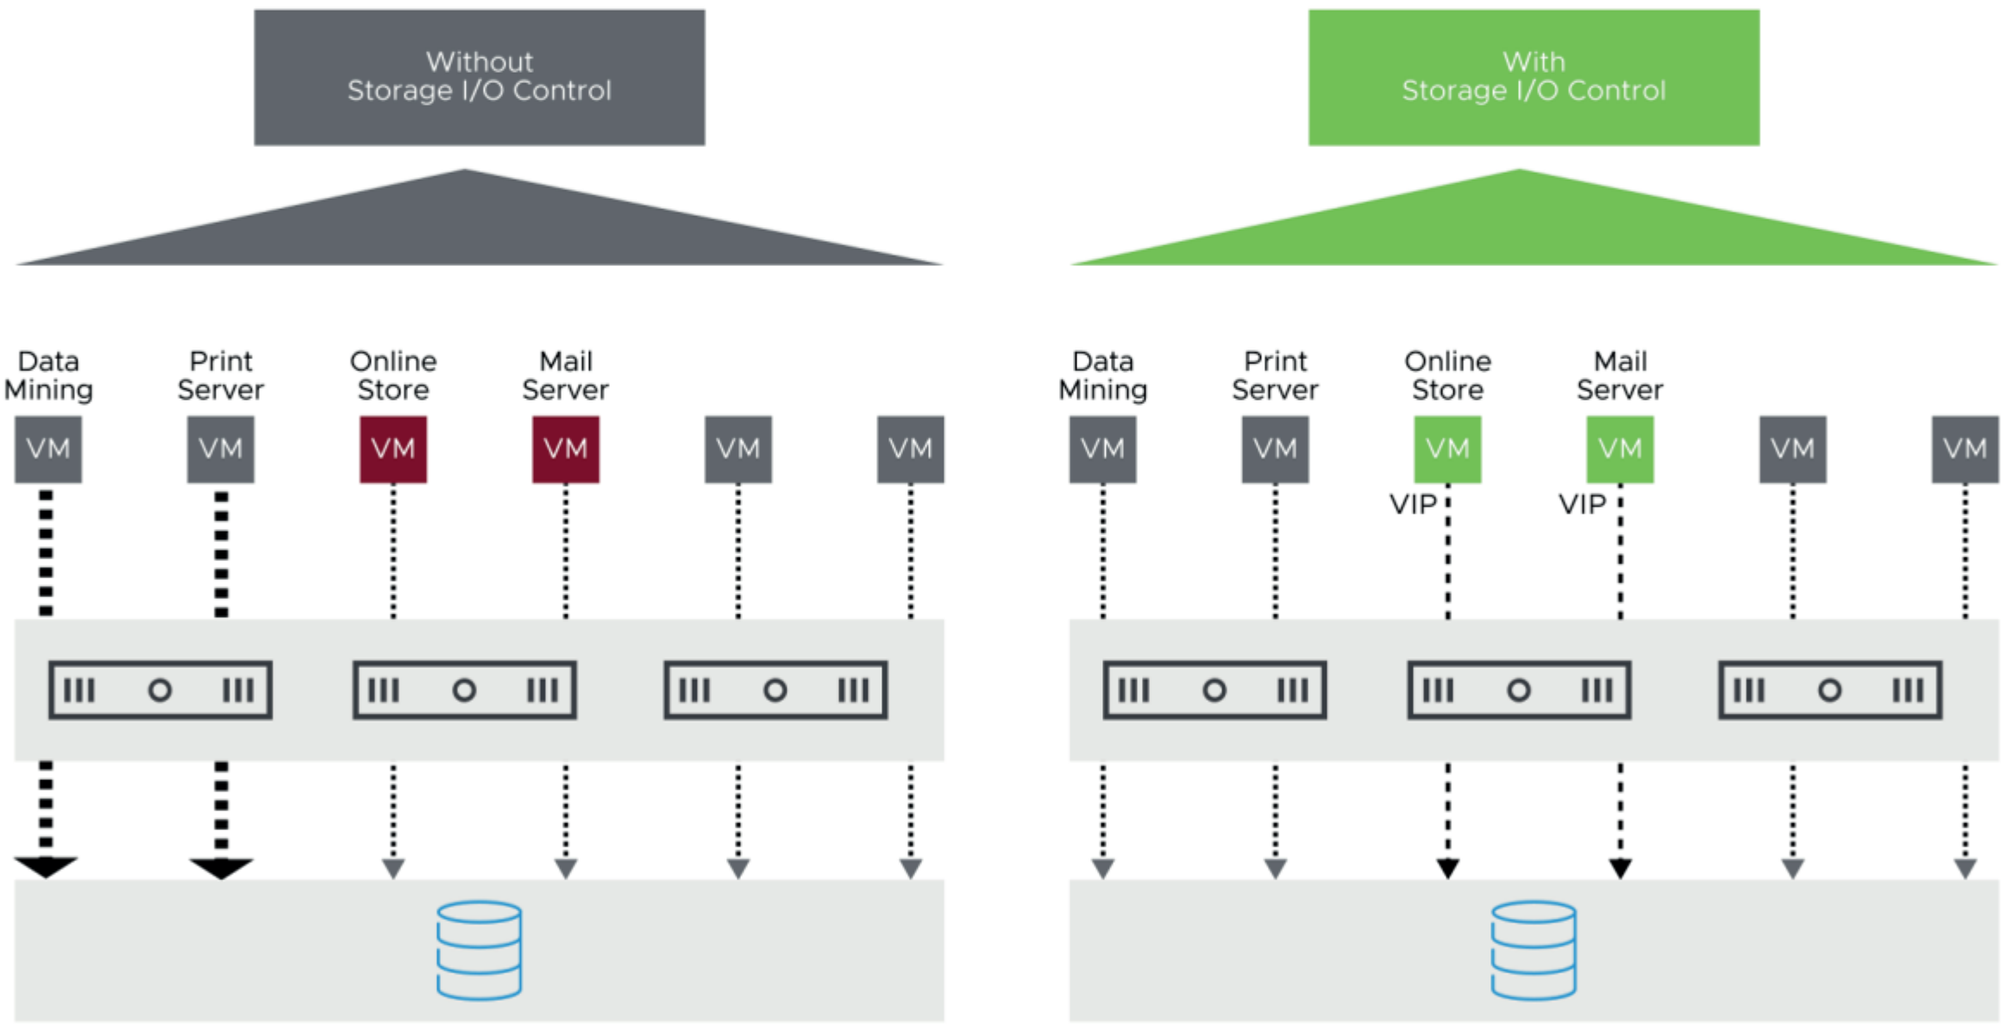

During High I/O from a Noncritical Application

Storage I/O Control extends the constructs of shares, limits, and reservations to handle storage I/O resources. Storage I/O Control is a proportional-share, I/O operations per second (IOPS) scheduler that, under contention, throttles IOPS. You can control the amount of storage I/O that is allocated to VMs during periods of I/O congestion. Controlling storage I/O ensures that more important VMs get preference over less important VMs for I/O resource allocation.

For more information, see *Managing Storage I/O Resources* at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-resource-management/GUID-7686FEC3-1FAC-4DA7-B698-B808C44E5E96.html.

5-145 Storage I/O Control Resource Shares and Limits

Storage I/O shares are similar to shares used for memory and CPU resource allocation.

Shares are typically specified as High, Normal, or Low and these values specify share values with a 4:2:1 ratio, respectively. You can also select Custom to assign a specific number of shares (which expresses a proportional weight) to each virtual machine.

These shares represent the relative importance of a virtual machine regarding the distribution of storage I/O resources. Under resource contention, virtual machines with higher share values have greater access to the storage array.

When you allocate storage I/O resources, you can limit the IOPS allowed for a virtual machine. By default, IOPS are unlimited.

5-146 Enabling Storage I/O Control (1)

With Storage I/O Control, you can ensure that the most important VMs get the appropriate level of I/O resources even during times of congestion.

Storage I/O Control is disabled by default. When Storage I/O Control is enabled on a datastore, all VMs have a default of 1,000 disk shares, no IOPS limit, and no reservation.

You define the number of shares, an upper limit for IOPS, and the number of reserved IOPS for VMs by applying a storage policy.

Storage I/O Control provides quality-of-service capabilities for storage I/O in the form of I/O shares, limits, and reservations that are enforced across all VMs accessing a datastore, regardless of which host they are running on.

When you enable Storage I/O Control on a datastore, ESXi begins to monitor the device latency that hosts observe when communicating with that datastore. When device latency exceeds a defined threshold, the datastore is considered to be congested, and each VM that accesses that datastore is allocated I/O resources in proportion to its shares.

Technet24

5-147 Enabling Storage I/O Control (2)

The following example is of two VMs running Iometer.

	IOPS	Iometer Latency
VM1	1,500	20 ms
VM2	1,500	21 ms

This table compares VM1 and VM2 using the default settings (1000 shares and no IOPS limit). No limits or reservations are set.

	IOPS	Iometer Latency
VM1	1,080	31 ms
VM2	1,900	16 ms

This table compares VM1 and VM2. VM1 still uses 1000 shares and VM2 uses 2000 shares. Shares or limits might be set but reservations are not set.

When you allocate storage I/O resources, you can limit the IOPS that are allowed for a VM. By default, the number of IOPS that are allowed for a VM is unlimited. If the limit that you want to set for a VM is in megabytes per second instead of IOPS, you can convert megabytes per second into IOPS based on the typical I/O size for that VM. For example, a backup application has a typical I/O size of 64 KB. To restrict a backup application to 10 MB per second, set a limit of 160 IOPS: 10 MB per second divided by 64 KB I/O size = 160 IOPS.

VMs VM1 and VM2 are running an I/O load generator called Iometer to generate enough traffic to cause contention .

Each VM is running on a different host, but they are running the same type of workload, which is 16 KB random reads. The shares of VM2 are set to twice as many shares as VM1, which implies that VM2 is more important than VM1. With Storage I/O Control deactivated, the IOPS and the I/O latency that each VM achieves are identical. However, with Storage I/O Control enabled, the IOPS achieved by the VM with more shares (VM2) is greater than the IOPS of VM1. The example assumes that each VM is running enough load to cause a bottleneck on the datastore.

5-148 Monitoring Latency

When you enable Storage I/O Control, ESXi monitors datastore latency and throttles the I/O load if the datastore average latency exceeds the threshold.

By default, Storage I/O Control uses an injector-based model to automatically detect the latency threshold.

The benefit of using the injector-based model is that Storage I/O Control determines the best threshold for a datastore.

When Storage I/O Control is enabled, performance statistics are collected and used to improve vSphere Storage DRS behavior.

Storage I/O Control automatically determines the optimal latency threshold by using injector-based models. However, you can also override this threshold by setting a specific latency value. This default latency setting might be fine for some storage devices, but other devices might reach their latency threshold well before or after the default setting is reached.

Technet24

5-149 Automatic Threshold Detection

With the injector-based model, Storage I/O Control determines the peak throughput of the storage device and sets the threshold to 90 percent of that value:

- Automatic threshold detection works well when a range of disk arrays and datastores are configured.
- The threshold varies according to the performance of each datastore.

You can also manually set the threshold value for a datastore. If this option is selected, the default latency setting is 30 ms.

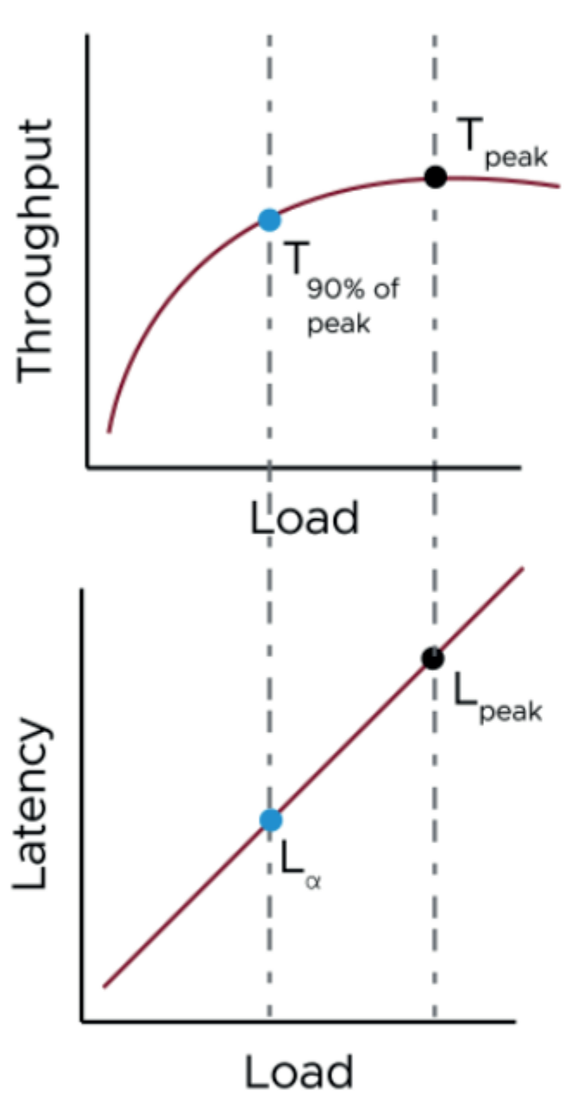

The I/O injector model determines the peak throughput of a datastore. The resulting peak throughput measurement can be used to determine the peak latency of a datastore.

5-150 Storage Policy Components for Storage I/O Control

Introduced in vSphere 6.5, Storage I/O Control is implemented as an I/O filter. The IOFilter provider advertises Storage I/O Control capabilities to vCenter.

Built-in storage policy components define Storage I/O Control capabilities.

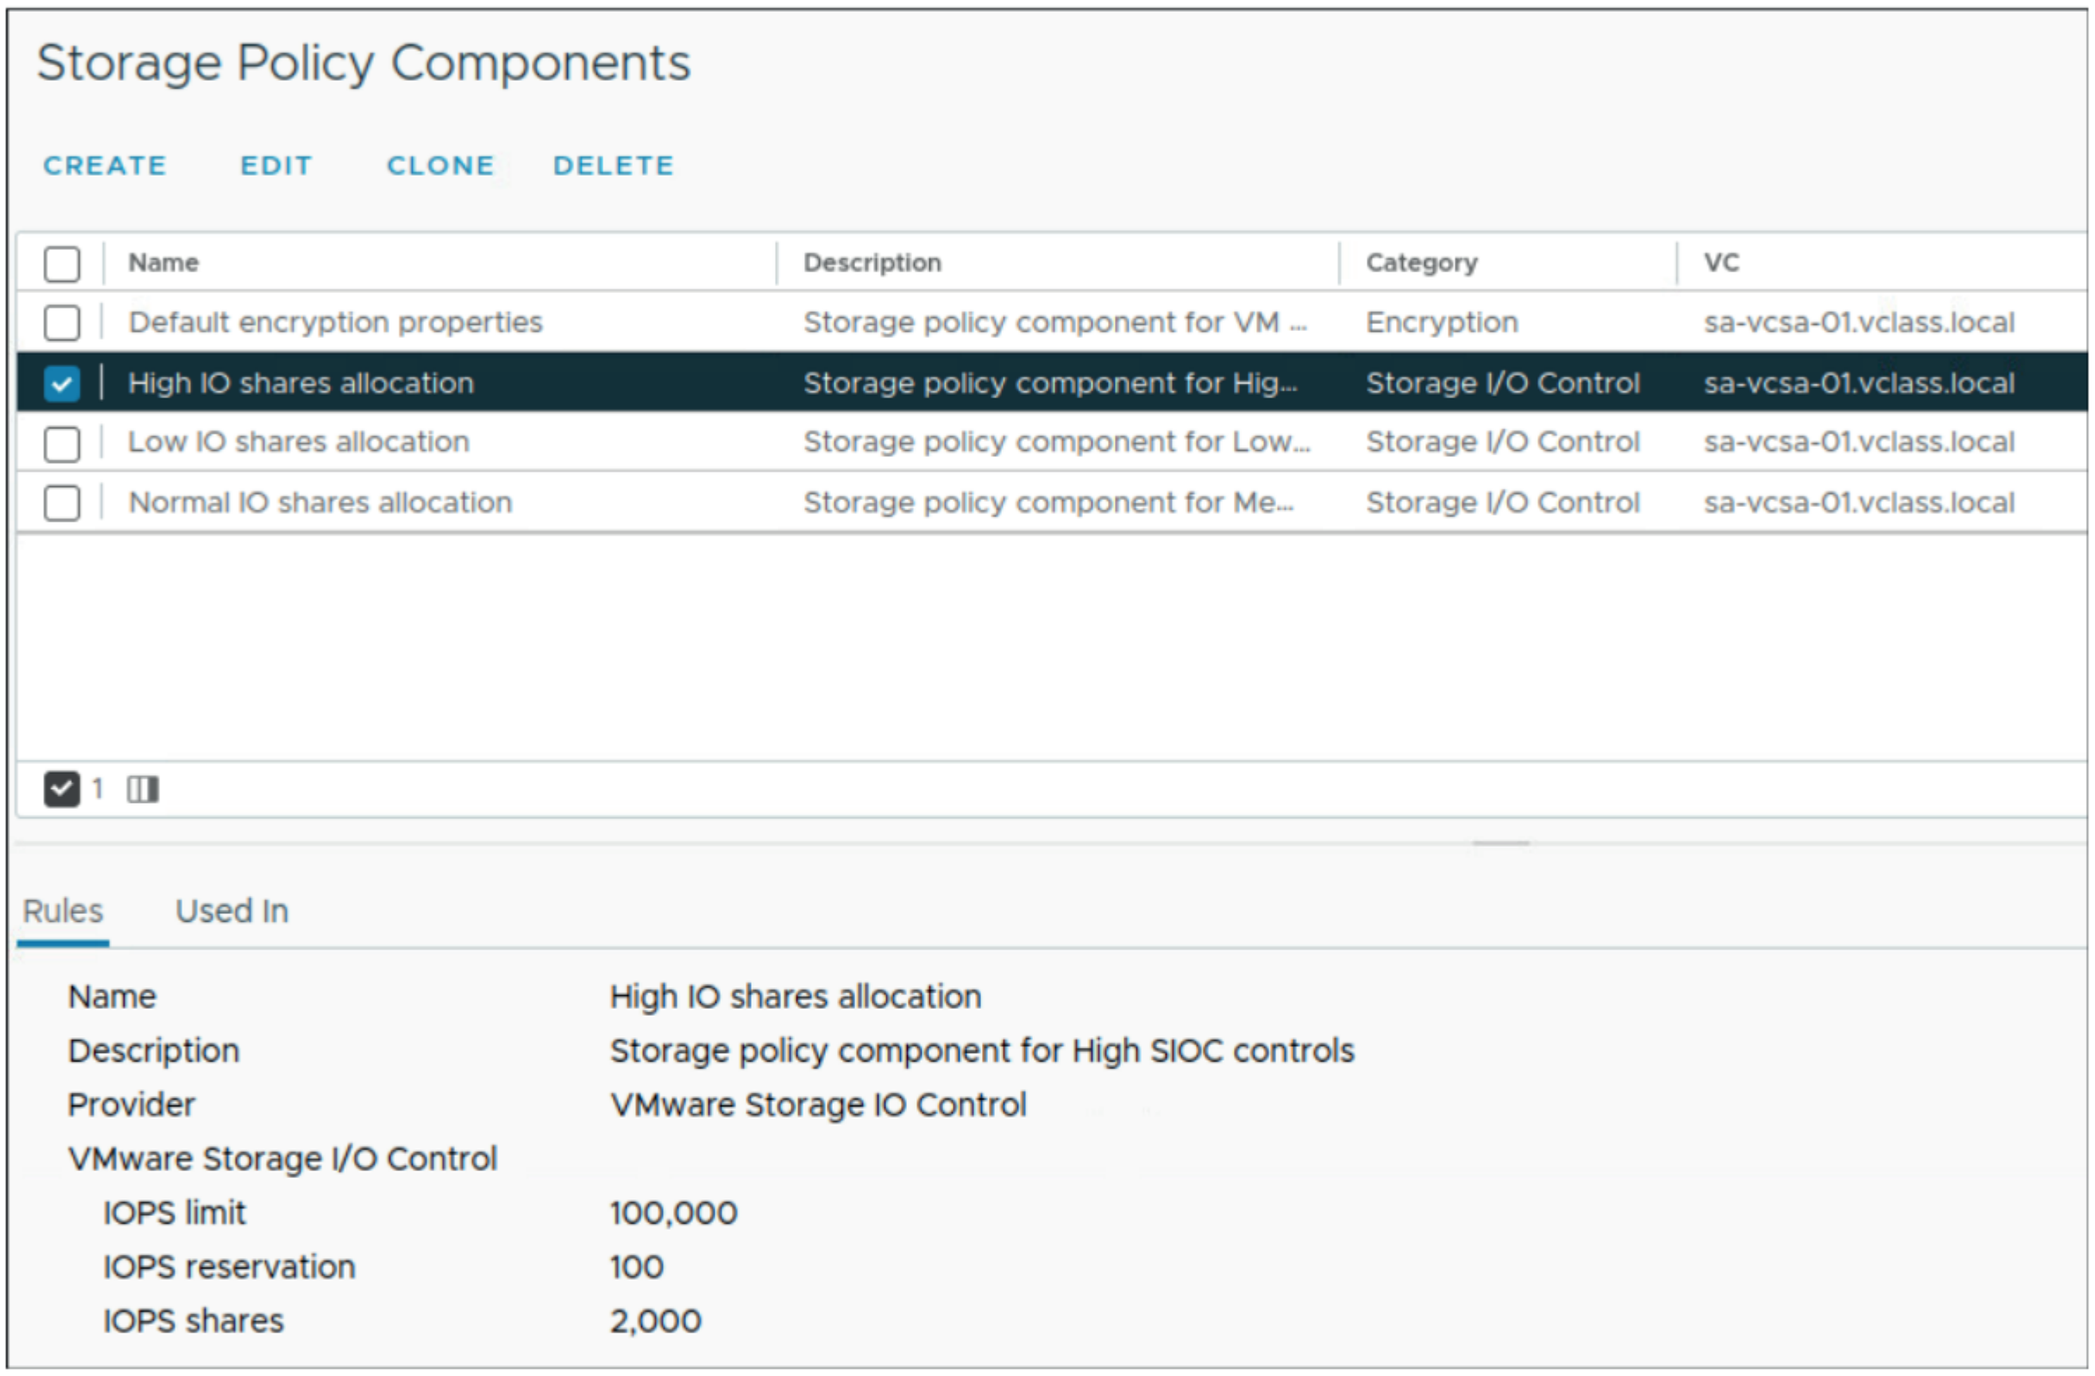

Technet24

5-151 Configuring Storage I/O Control

You can use a storage policy component for Storage I/O Control.

Or, you can create a custom storage policy component, which defines the IOPS limit, IOPS reservation, and IOPS shares. You use the component in the storage policy and assign the policy to a VM.

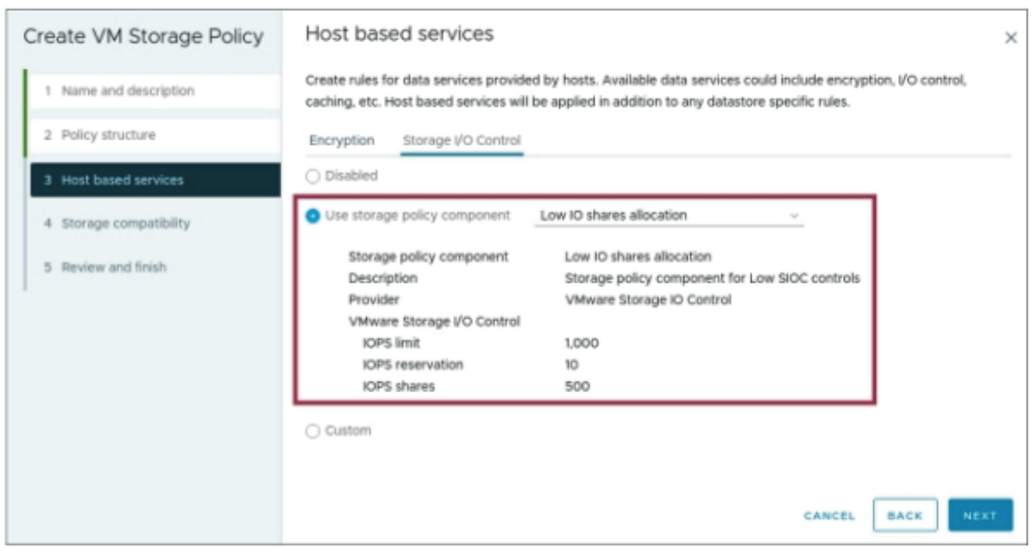

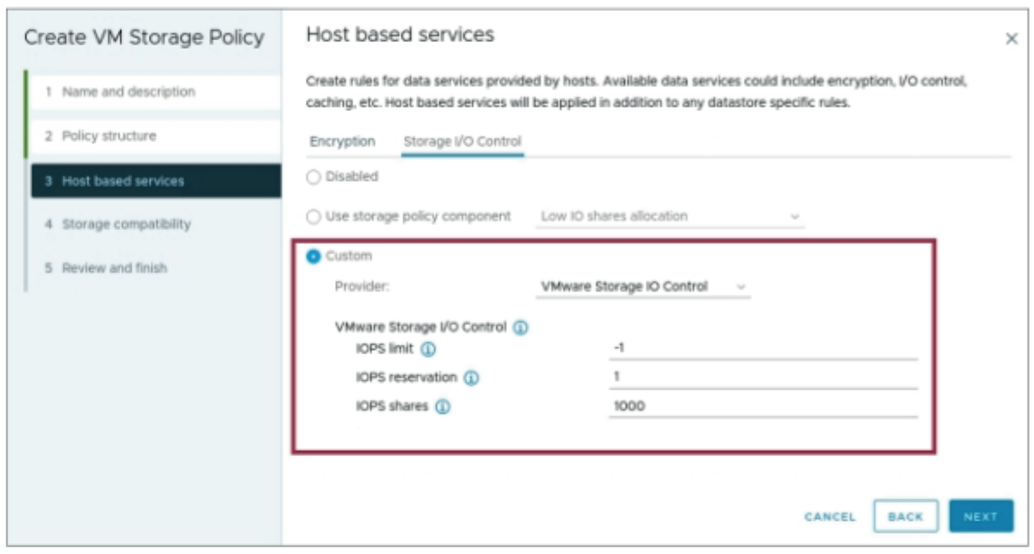

5-152 Storage I/O Control Requirements

Storage I/O Control has several requirements:

- Storage I/O Control is supported only on datastores that are managed by a single vCenter system.
- Storage I/O Control is supported for Fibre Channel, iSCSI, VMFS, and NFS3 storage:
  - Raw device mapping is not supported.
  - vSAN and virtual volumes datastores are not supported.
- Storage I/O Control does not support datastores with multiple extents.
- If your datastores are located on arrays with automated storage tiering capabilities, your array must be certified as compatible with Storage I/O Control.

Automated storage tiering is the ability of an array (or a group of arrays) to migrate LUNs, virtual volumes, or parts of LUNs or virtual volumes to different types of storage media (SSD, Fibre Channel, SAS, SATA, and so on) based on user-set policies and current I/O patterns. No special certification is required for arrays that do not have these automatic migration or tiering features, including those features that allow you to manually migrate data between different types of storage media.

To verify that your automated tiered storage array is certified as compatible with Storage I/O Control, see the *VMware Compatibility Guide* at http://www.vmware.com/resources/compatibility.

Technet24

5-153 Activity: Storage I/O Control (1)

On a datastore with Storage I/O Control enabled, what steps do you take to ensure that the online store VM and the mail server VM have higher IOPS than the other VMs during periods of contention?

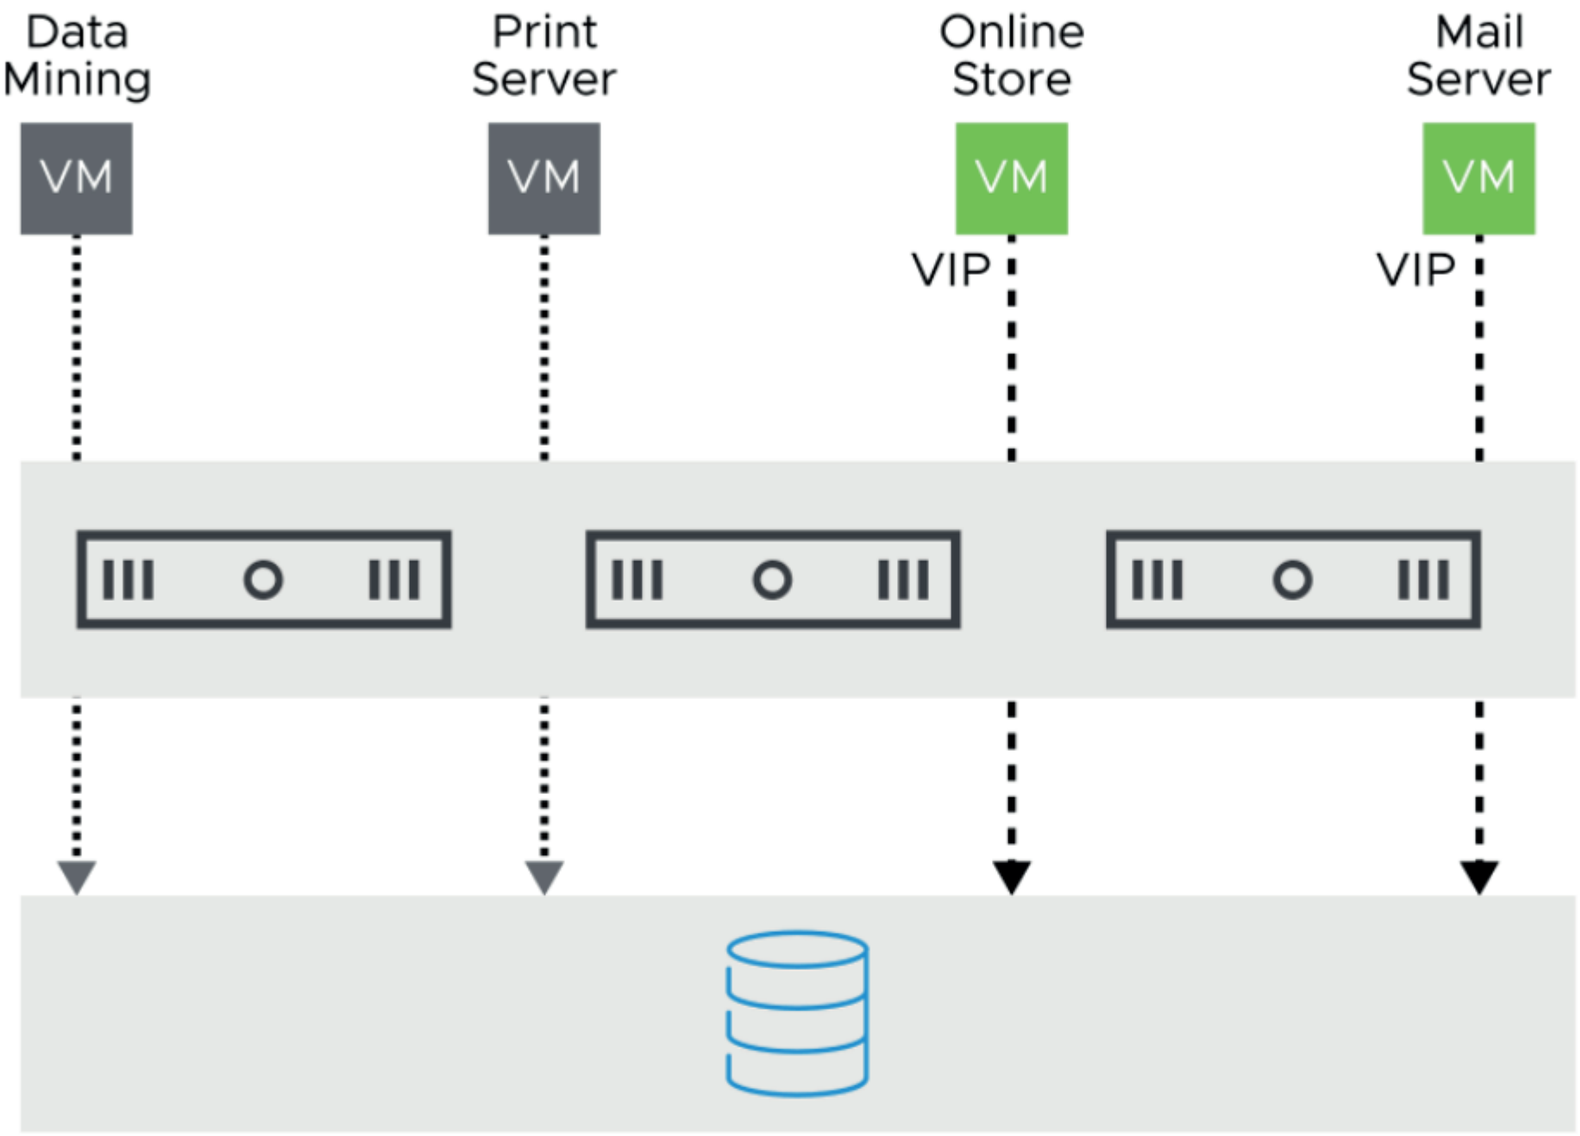

Storage I/O Control Enabled

5-154 Activity: Storage I/O Control (2)

The following steps describe one way to ensure that the online store VM and mail server VM have higher throughput than other VMs during contention:

1. Create a storage policy that uses the predefined storage policy component called High IO Shares Allocation. High IO Shares Allocation has a disk shares value of 2,000.
2. Assign the storage policy to only the online store VM and the mail server VM. All other VMs on the datastore use the default of 1,000 disk shares.

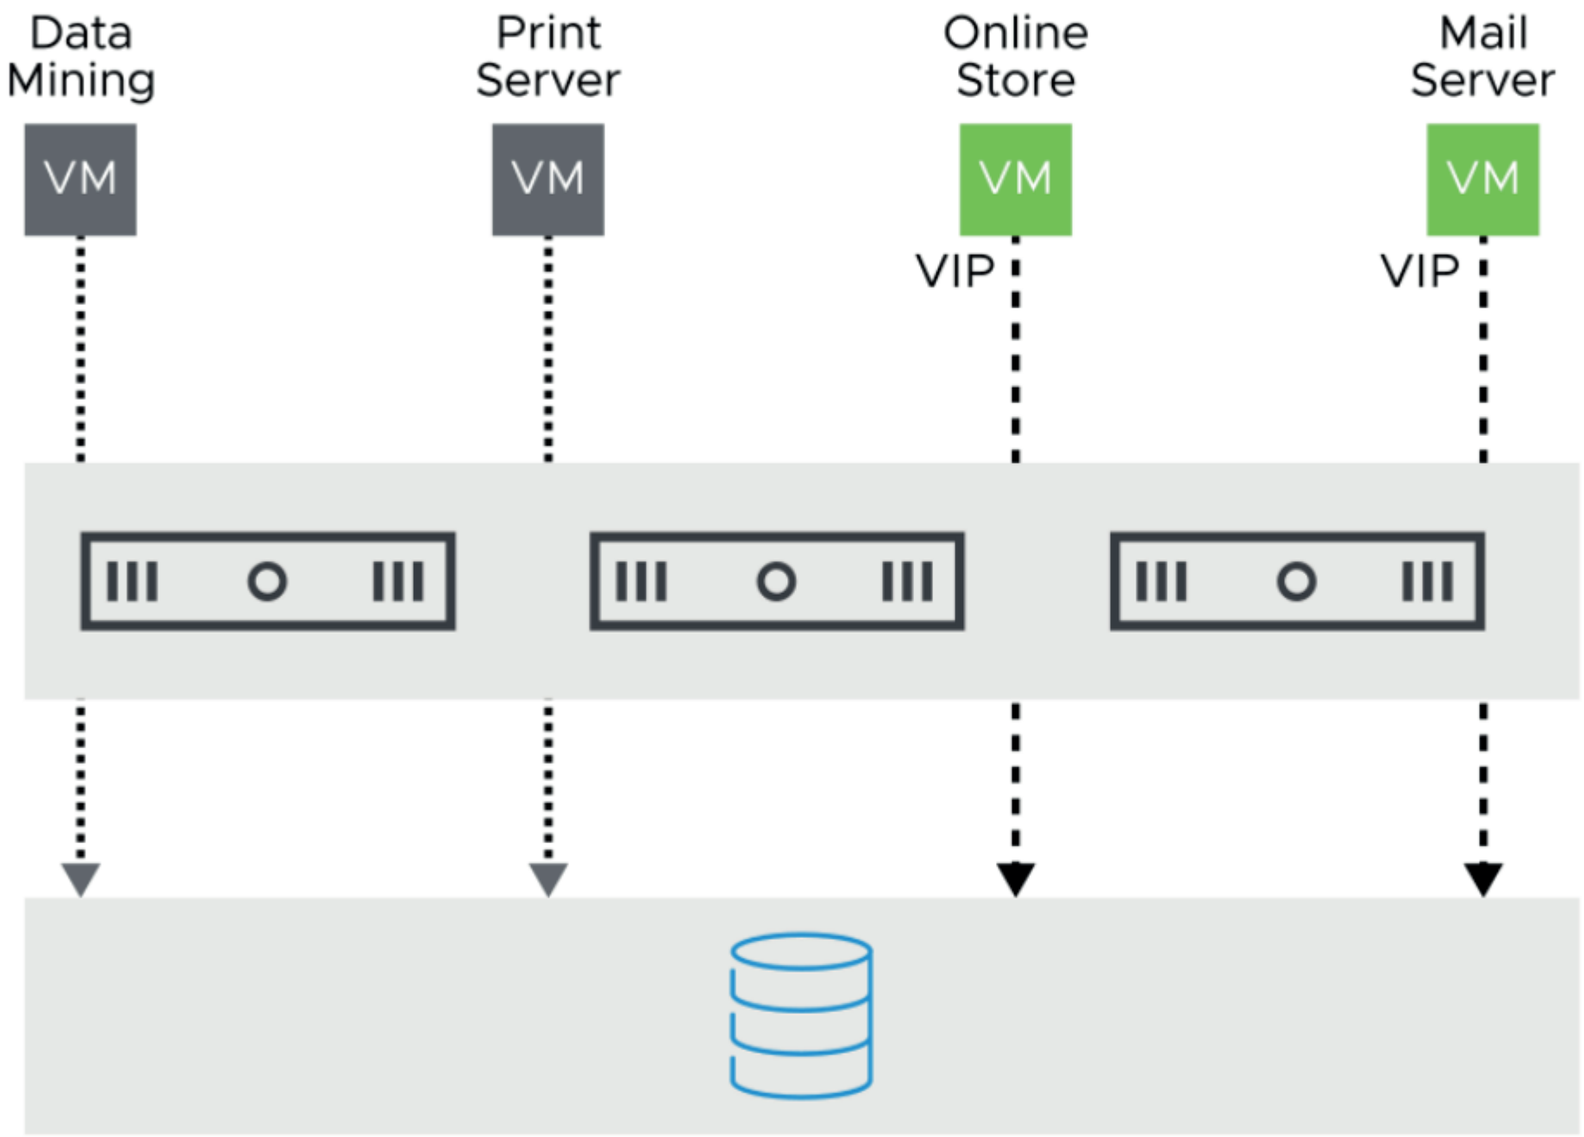

5-155 Review of Learner Objectives

- Describe how Storage I/O Control balances the I/O load on a datastore
- Configure Storage I/O Control

Technet24

5-156 Lesson 12: Datastore Clusters and vSphere Storage DRS

5-157 Learner Objectives

- Create a datastore cluster
- Configure vSphere Storage DRS
- Explain how Storage I/O Control and vSphere Storage DRS complement each other

5-158 About Datastore Clusters

A datastore cluster is a collection of datastores that are grouped together but that work as separate entities.

When you enable a datastore cluster with vSphere Storage DRS, the datastores on the cluster work together to balance capacity and I/O latency.

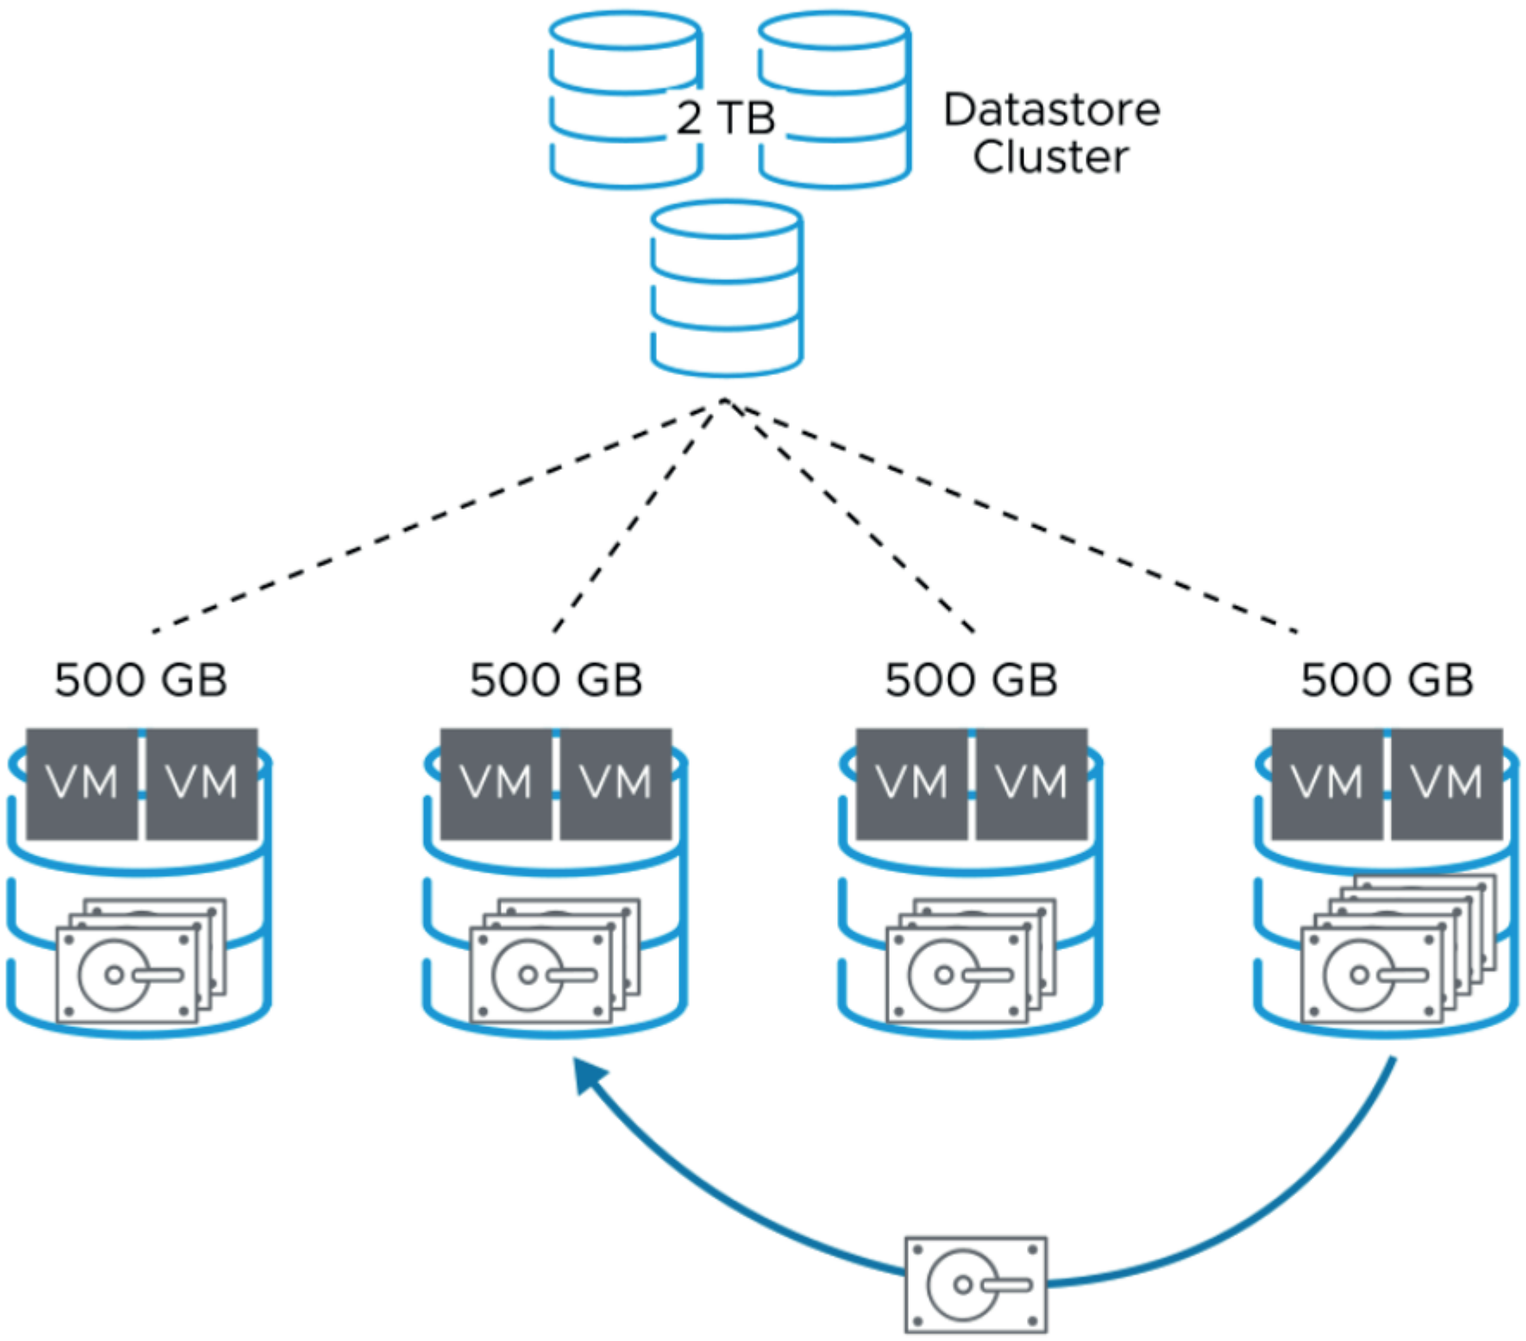

The datastore cluster serves as a container or folder. You can store datastores in the container, but the datastores work as separate entities.

A datastore cluster that is enabled for vSphere Storage DRS is a collection of datastores that work as a single unit.

Technet24

5-159 Datastore Cluster Requirements

Datastores and hosts that are associated with a datastore cluster must meet certain requirements to successfully use datastore cluster features:

- Datastore clusters must contain similar or interchangeable datastores, with the following exceptions:
  - — All datastores in a datastore cluster must be in the same format, either VMFS or NFS.
  - — vSAN and virtual volumes datastores are not supported in a datastore cluster.
  - — Replicated datastores cannot be combined with nonreplicated datastores in the same datastore cluster that is enabled for vSphere Storage DRS.

    This exception applies only to storage arrays that do not support vSphere API for Storage Awareness.
- Datastores shared across multiple data centers cannot be included in a datastore cluster.
- Datastores with hardware acceleration enabled should not be included in the same datastore cluster as datastores that do not have hardware acceleration enabled.

A datastore cluster can contain a mix of datastores of different sizes and I/O capacities, as well as datastores from different arrays and vendors. However, LUNs with different performance characteristics can cause performance problems.

vSphere Storage DRS cannot move VMs between NFS and VMFS datastores.

5-160 Host Cluster to Datastore Cluster Relationships

A vSphere host cluster and a datastore cluster can interact in different relationships.

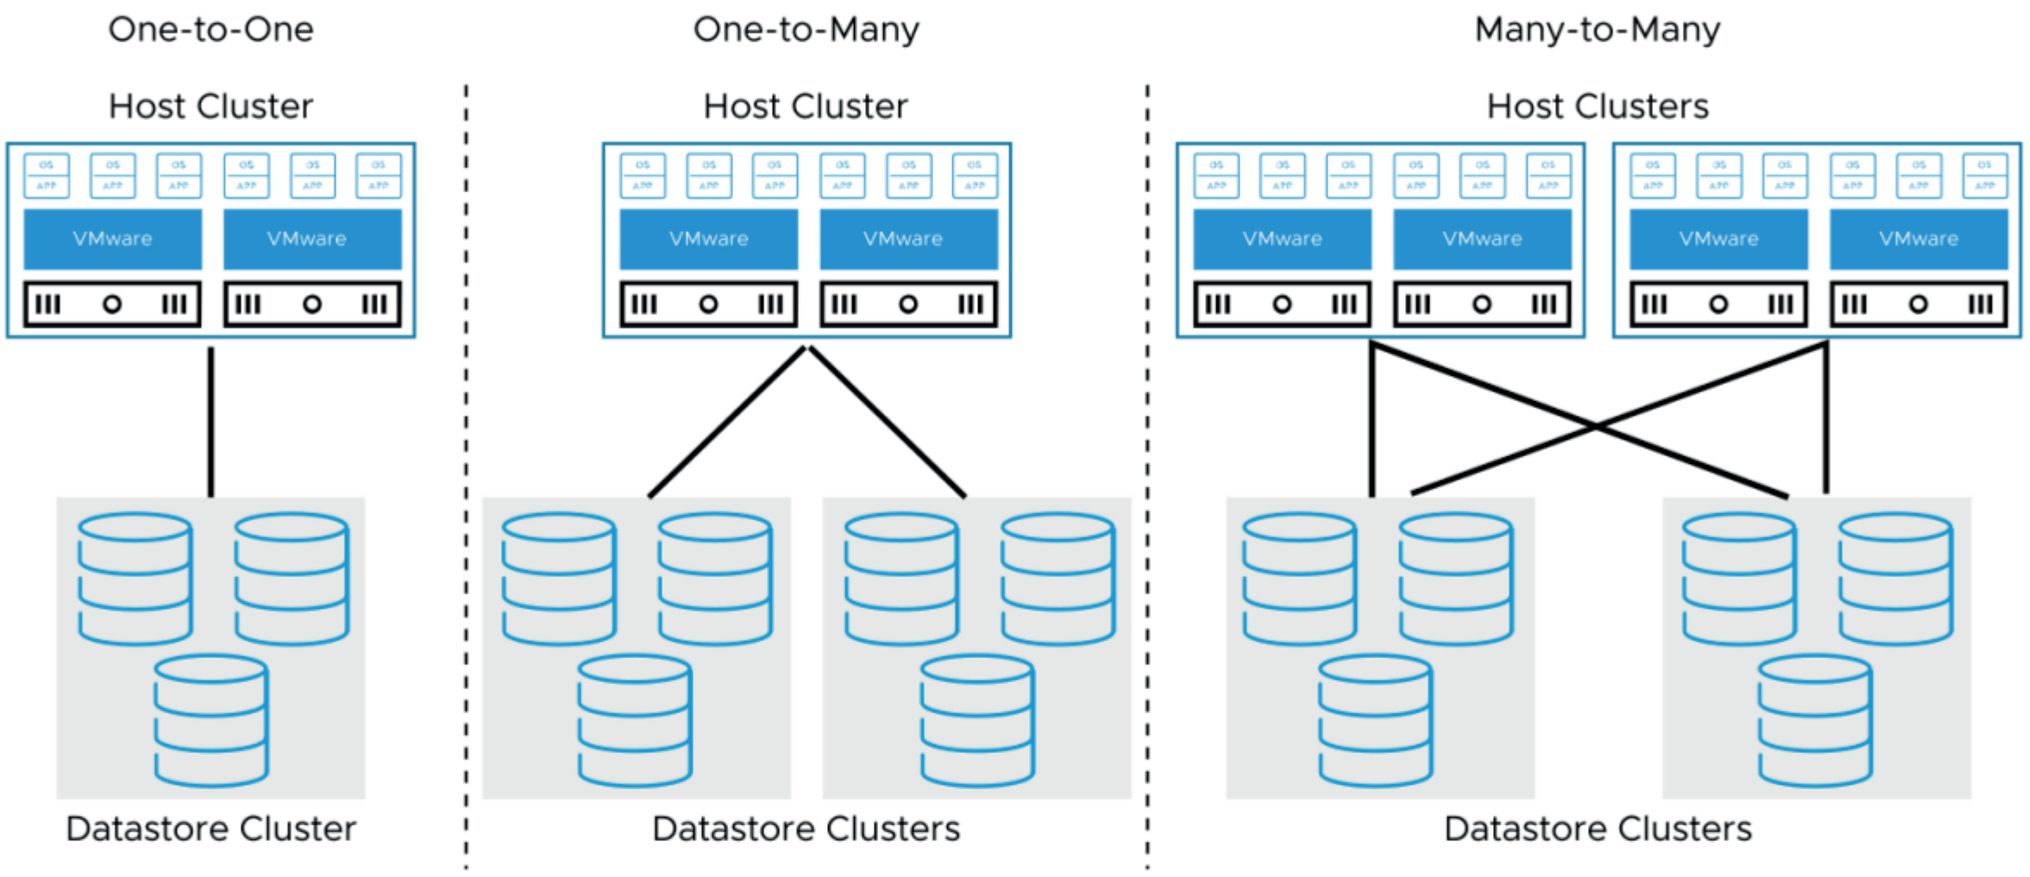

Host clusters and datastore clusters can coexist in the virtual infrastructure. A host cluster is a vSphere DRS or a vSphere HA cluster.

Load balancing by vSphere DRS and vSphere Storage DRS can occur at the same time. vSphere DRS balances VMs across hosts based on CPU usage, memory usage, and network usage (the latter if Network I/O Control version 3 is used). vSphere Storage DRS load balances VMs across storage, based on storage capacity and IOPS.

A standalone host (one that is not part of a host cluster) can also use a datastore cluster.

Technet24

5-161 vSphere Storage DRS Overview

With vSphere Storage DRS, you can manage the aggregated resources of a datastore cluster in the following ways:

- Initial placement of VMs based on storage capacity and, optionally, I/O latency
- Use of vSphere Storage vMotion to migrate VMs based on storage capacity and, optionally, I/O latency
- Use of affinity and anti-affinity rules to govern virtual disk location
- Use of datastore maintenance mode to clear a LUN of VM files

vSphere Storage DRS can be set up to work in either manual mode or fully automated mode.

vSphere Storage DRS manages the placement of VMs in a datastore cluster, based on the space usage of the datastores. It tries to keep usage as even as possible across datastores in the datastore cluster.

vSphere Storage DRS uses vSphere Storage vMotion to migrate VMs to maintain the balance across datastores.

Optionally, the user can configure vSphere Storage DRS to balance I/O latency across members of the datastore cluster as a way to help mitigate performance problems that are caused by I/O latency.

vSphere Storage DRS can be set up to work in either manual mode or fully automated mode:

- Manual mode presents migration and placement recommendations to the user, but nothing is executed until the user accepts the recommendation.
- Fully automated mode automatically handles initial placement and migrations based on runtime.

5-162 Initial Disk Placement

Initial placement occurs when vSphere Storage DRS selects a datastore in a datastore cluster on which to place a VM disk.

When VMs are created, cloned, or migrated, you select a datastore cluster rather than a single datastore. vSphere Storage DRS selects a member datastore based on capacity and, optionally, IOPS load.

By default, a VM's files are placed on the same datastore in the datastore cluster. This behavior can be changed by using vSphere Storage DRS anti-affinity rules.

When a VM is created, cloned, or migrated, and the user selects a datastore cluster, vSphere Storage DRS selects a member datastore in the datastore cluster based on storage use. vSphere Storage DRS tries to keep the member datastores evenly used.

Technet24

5-163 vSphere Storage DRS Affinity Rules

You can create vSphere Storage DRS anti-affinity rules to control which virtual disks should not be placed on the same datastore in a datastore cluster.

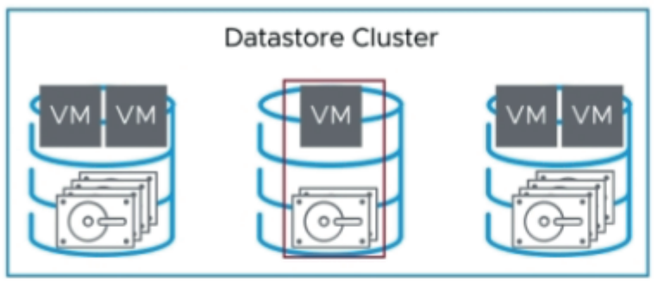

Intra-VM VMDK Affinity Rule (default):

- Keep virtual machine VMDKs together on the same datastore.
- Maximize VM availability when all disks are busy.

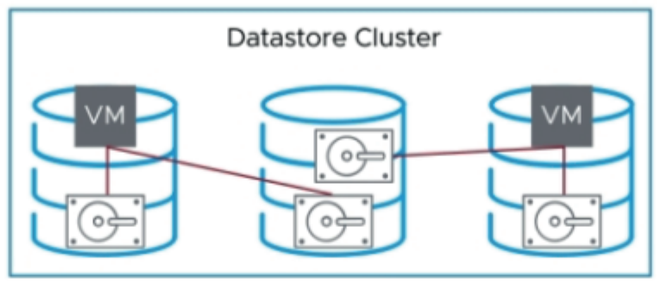

Intra-VM VMDK Anti-Affinity Rule:

- Keep virtual machine VMDKs on different datastore.
- Apply to all or a subset of VMDKs.

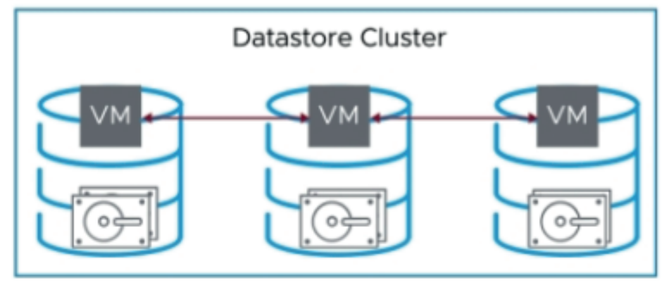

VM Anti-Affinity Rule:

- Keep VMs on different datastores.
- Maximize availability of a set of redundant VMs.

By default, all of a VM's disks can be on the same datastore.

A user might want the virtual disks on different datastores. For example, a user can place a system disk on one datastore and the data disks on another. In this case, the user can set up a VMDK anti-affinity rule, which keeps a VM’s virtual disks on separate datastores.

VM anti-affinity rules keep VMs on separate datastores. This rule is useful when redundant VMs must always be available.

5-164 Storage Migration Recommendations

vCenter displays migration recommendations on the vSphere Storage DRS Recommendations page for datastore clusters that have manual mode set.

Migration recommendations are executed under the following circumstances:

- When the I/O latency threshold is exceeded
- When the space utilization threshold is exceeded

If enabled, IOPS load history is checked every eight hours by default.

vSphere Storage DRS selects a datastore based on utilization and, optionally, I/O latency.

Load balancing is based on IOPS workload, which ensures that no datastore exceeds a particular VMkernel I/O latency level.

Storage migration recommendations work as follows:

- vSphere Storage DRS provides as many recommendations as necessary to balance the space and, optionally, the IOPS resources of the datastore cluster.
- Reasons for migration recommendations include balancing space usage in the datastore, reducing datastore I/O latency, and balancing datastore IOPS load.
- vSphere Storage DRS can also make mandatory recommendations based on the following conditions:
  - A datastore is out of space.
  - vSphere Storage DRS anti-affinity rules or affinity rules are being violated.
  - A datastore is entering maintenance mode.

To achieve optimal balance of VMs across datastores in the datastore cluster, vSphere Storage DRS considers moving powered-off VMs to other datastores.

Technet24

5-165 Datastore Correlation Detector

Datastore correlation refers to datastores that are created on the same physical set of spindles.

vSphere Storage DRS detects datastore correlation by performing the following operations:

- Measuring individual datastore performance
- Measuring combined datastore performance

If latency increases on multiple datastores when load is placed on one datastore, the datastores are considered to be correlated.

Anti-affinity rules can use correlation detection to ensure that the VMs or virtual disks are on different spindles.

The purpose of datastore correlation is to help vSphere Storage DRS in determining where to move a VM. For example, you gain little advantage by moving a VM from one datastore to another if both datastores are supported by the same set of physical spindles on the array.

The datastore correlation detector uses the I/O injector to determine whether a source and a destination datastore are using the same back-end spindles.

The datastore correlation detector works by monitoring the load on one datastore and monitoring the latency on another. If latency increases on other datastores when a load is placed on one datastore, the datastores are correlated.

The datastore correlation detector can also be used for anti-affinity rules, ensuring that VMs and virtual disks are not only kept on separate datastores but also kept on different spindles on the back end.

Correlation is determined by a long-running background process.

Datastore correlation is enabled by default.

5-166 Configuring vSphere Storage DRS Runtime Settings

You can configure vSphere Storage DRS thresholds to determine when migrations are performed or recommended.

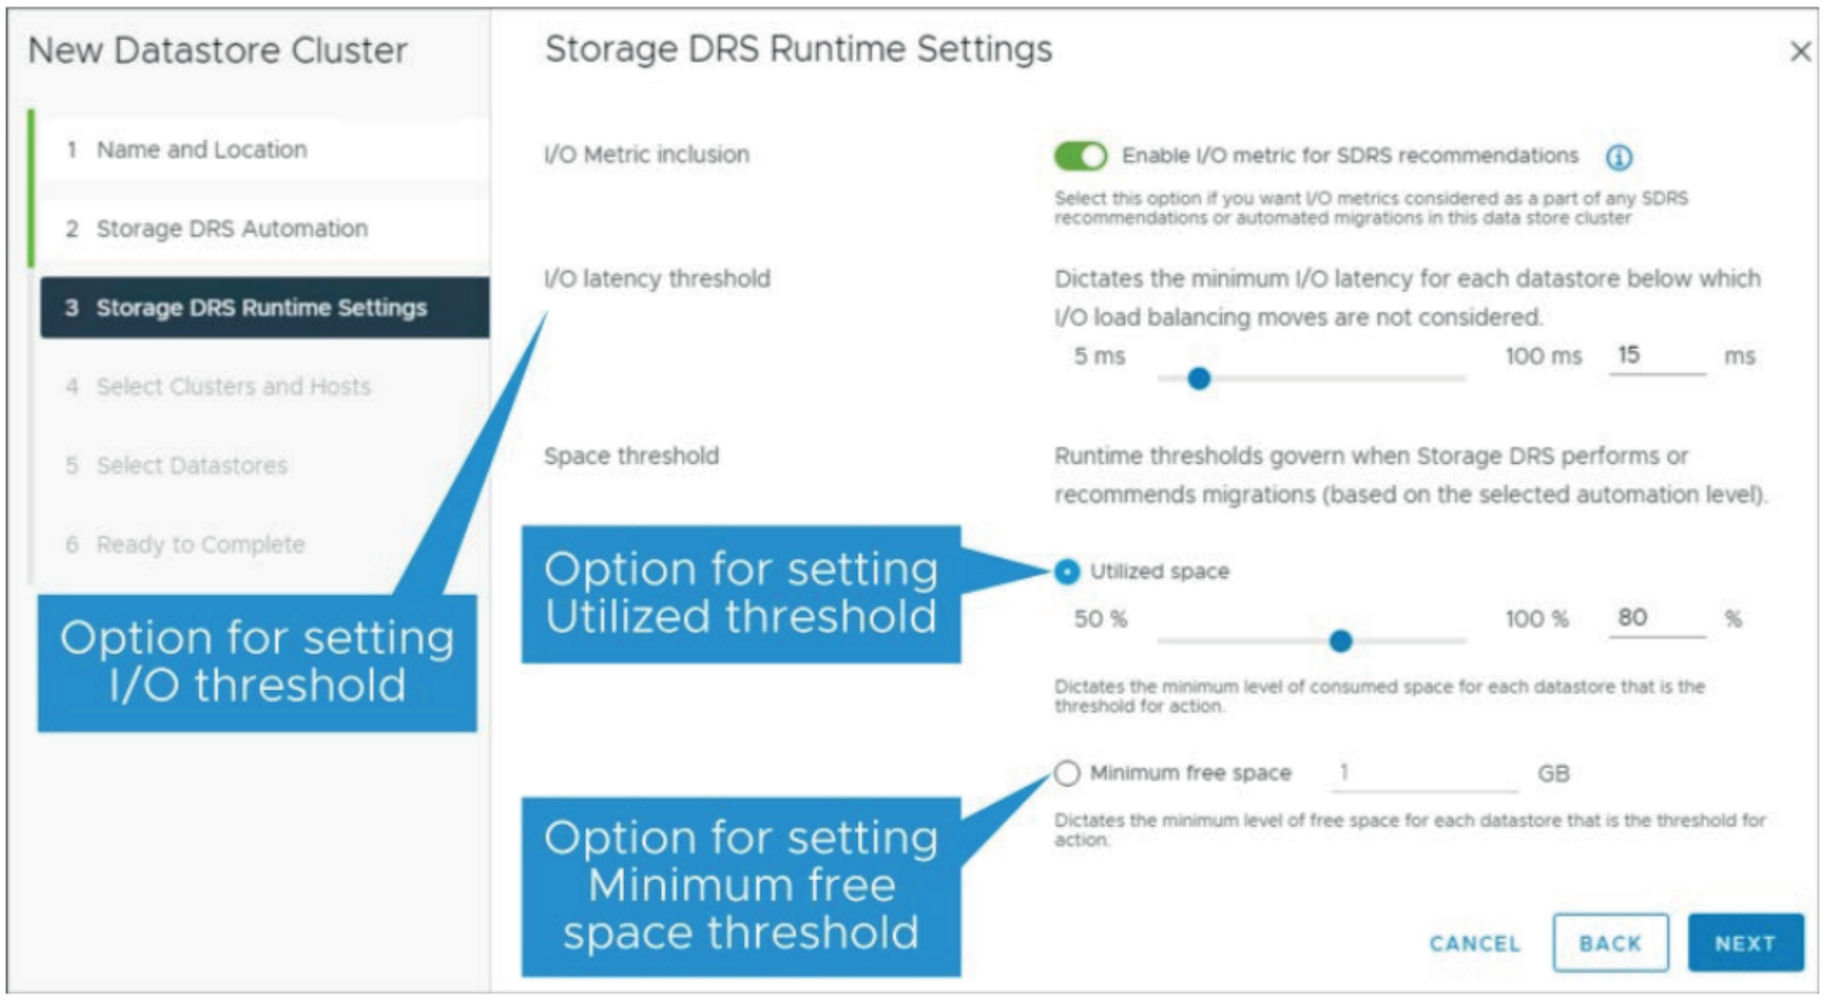

- When you create datastore cluster, you select or deselect the **Enable I/O metric for SDRS recommendations** check box to activate or deactivate IOPS metrics use in recommendations or automated migrations.
- When you enable I/O load balancing, Storage I/O Control is enabled for all the datastores in the datastore cluster if it is not already enabled for the cluster.
- When this option is deselected, you deactivate the following functions:
  - IOPS load balancing among datastores in the datastore cluster.
  - Initial placement for virtual disks based on the IOPS metric. Space is the only consideration when placement and balancing recommendations are made.
- vSphere Storage DRS runtime settings include the following options:
  - I/O latency threshold: Indicates the minimum I/O latency for each datastore below which I/O load-balancing moves are not considered. This setting is applicable only if the **Enable I/O metric for SDRS recommendations** check box is selected.
  - Space threshold: Determines the minimum levels of consumed space and free space for each datastore that is the threshold for action.

Technet24

5-167 vSphere Storage DRS Maintenance Mode

◊ With vSphere Storage DRS maintenance mode, you can take a datastore out of use to service it.

◊ vSphere Storage DRS maintenance mode evacuates VMs from a datastore placed in maintenance mode:

- Registered VMs (on or off) are moved.
- Templates, unregistered VMs, ISO images, and nonvirtual machine files are not moved.

Apply	Recommendation	Reason	Utilization % (source)	Utilization % (destination)	I/O Latency Before
☑	Migrate hard disk Hard disk 1, for "Photon-02" from "Gold" to "Silver"	Datastore entering maintenance mode	Before: 75.8 After: 64.7	Before: 29.8 After: 45.6	Source: 1.7 Destination: 1.6
☑	Migrate configuration files for "Photon-04" from "Gold" to "Silver"	Datastore entering maintenance mode	Before: 64.7 After: 53.6	Before: 45.6 After: 61.3	Source: 1.7 Destination: 1.6
	Migrate hard disk Hard disk 1, for "Photon-04" from "Gold" to "Silver"	Datastore entering maintenance mode	Before: 64.7 After: 53.6	Before: 45.6 After: 61.3	Source: 1.7 Destination: 1.6

With vSphere Storage DRS, you can place a datastore in maintenance mode. vSphere Storage DRS maintenance mode is available to datastores in a datastore cluster that is enabled for vSphere Storage DRS. Standalone datastores cannot be placed in maintenance mode.

The datastore does not enter maintenance mode until all files on the datastore are moved. You must manually move these files off the datastore so that the datastore can enter vSphere Storage DRS maintenance mode.

If the datastore cluster is placed in fully automated mode, VMs are automatically migrated to other datastores.

If the datastore cluster is placed in manual mode, migration recommendations appear in the vSphere Client. The virtual disks cannot be moved until the recommendations are accepted.

5-168 Backups and vSphere Storage DRS

Backing up VMs can add latency to a datastore.

Using the vSphere Client, you can schedule a task to deactivate vSphere Storage DRS behavior during the backup.

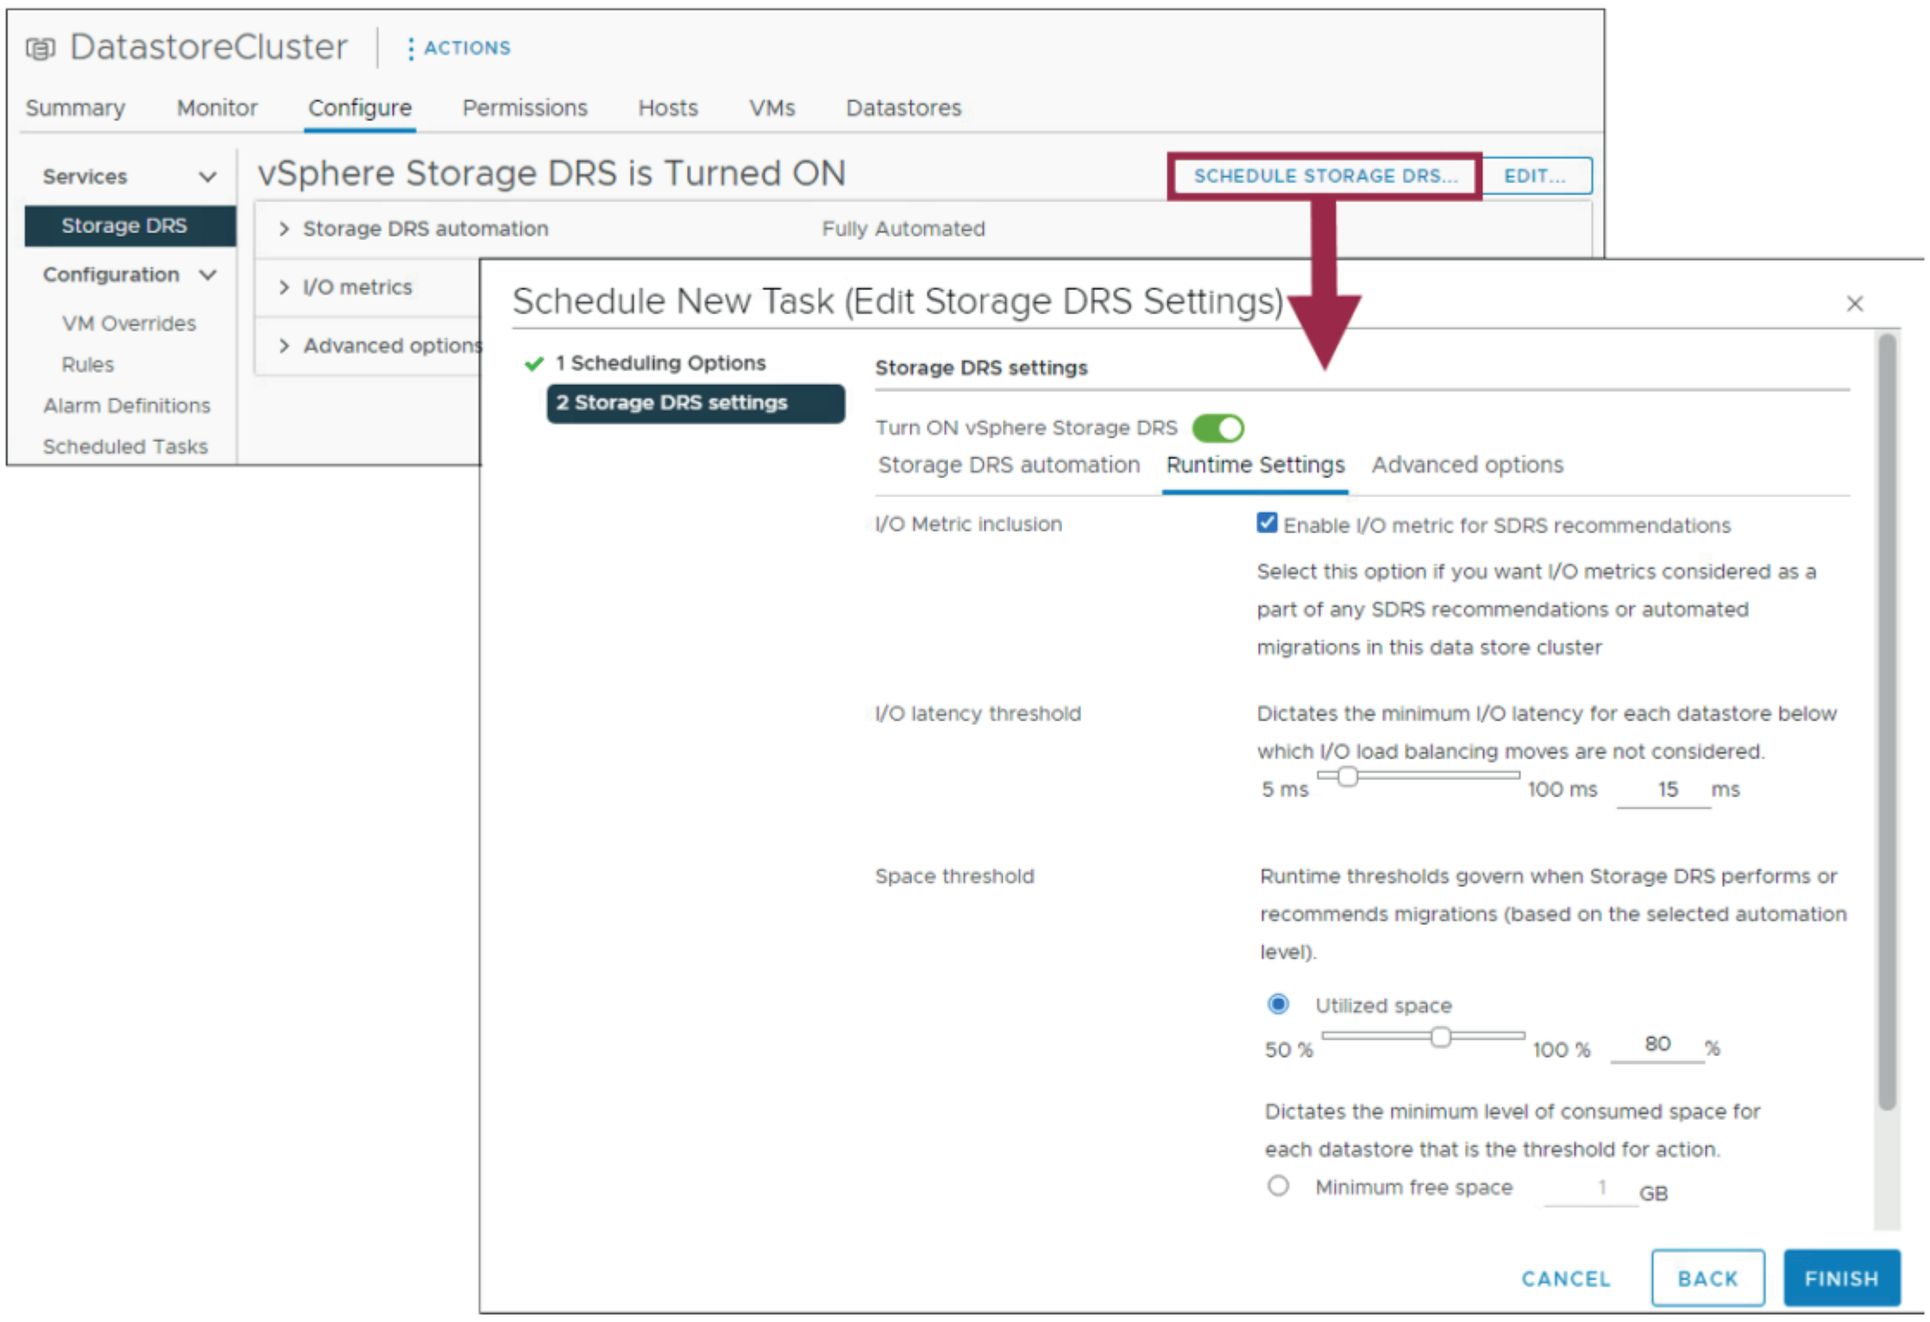

You can configure scheduled tasks to change vSphere Storage DRS behavior.

You can use scheduled tasks to change the vSphere Storage DRS configuration of the datastore cluster to match enterprise activity.

For example, if the datastore cluster is configured to perform migrations based on I/O latency, you might deactivate the use of I/O metrics by vSphere Storage DRS during the backup window.

You can reenable I/O metrics use after the backup window closes.

Technet24

5-169 vSphere Storage DRS and Storage I/O Control

For vSphere Storage DRS, Storage I/O Control gathers latency statistics for an ESXi host, sent to vCenter, and stored in the vCenter database.

With these statistics, vSphere Storage DRS can decide whether a VM should be migrated to another datastore.

With Storage I/O Control, you can balance I/O load in a datastore cluster that is enabled for vSphere Storage DRS.

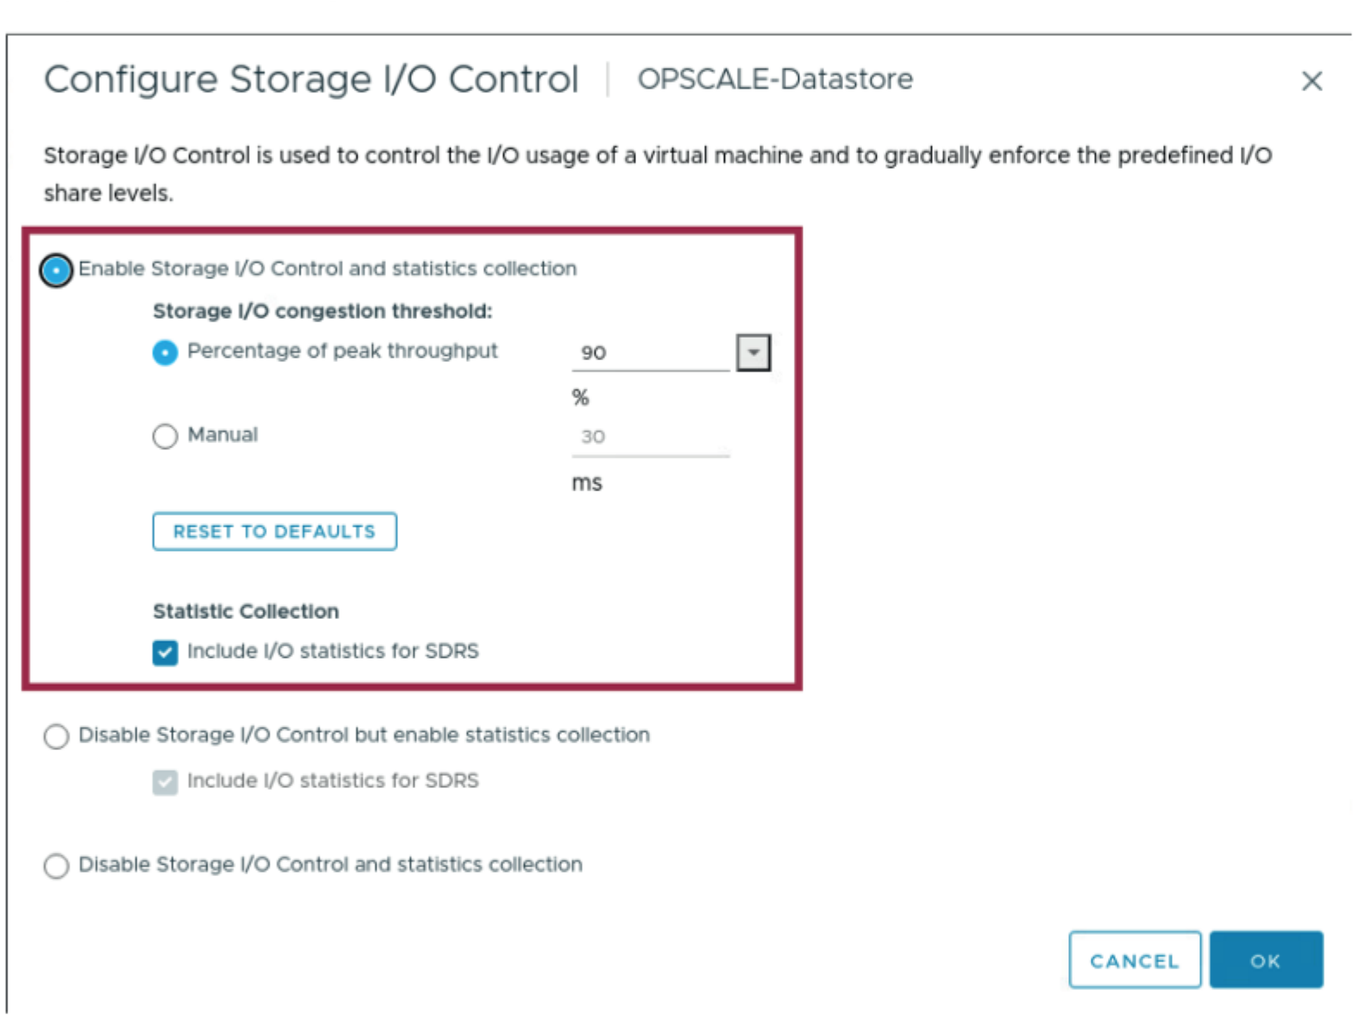

5-170 Activity: vSphere Storage DRS (1)

Consider the datastore cluster example. How do you configure vSphere Storage DRS to eliminate the single point of failure for the primary DNS and secondary DNS servers?

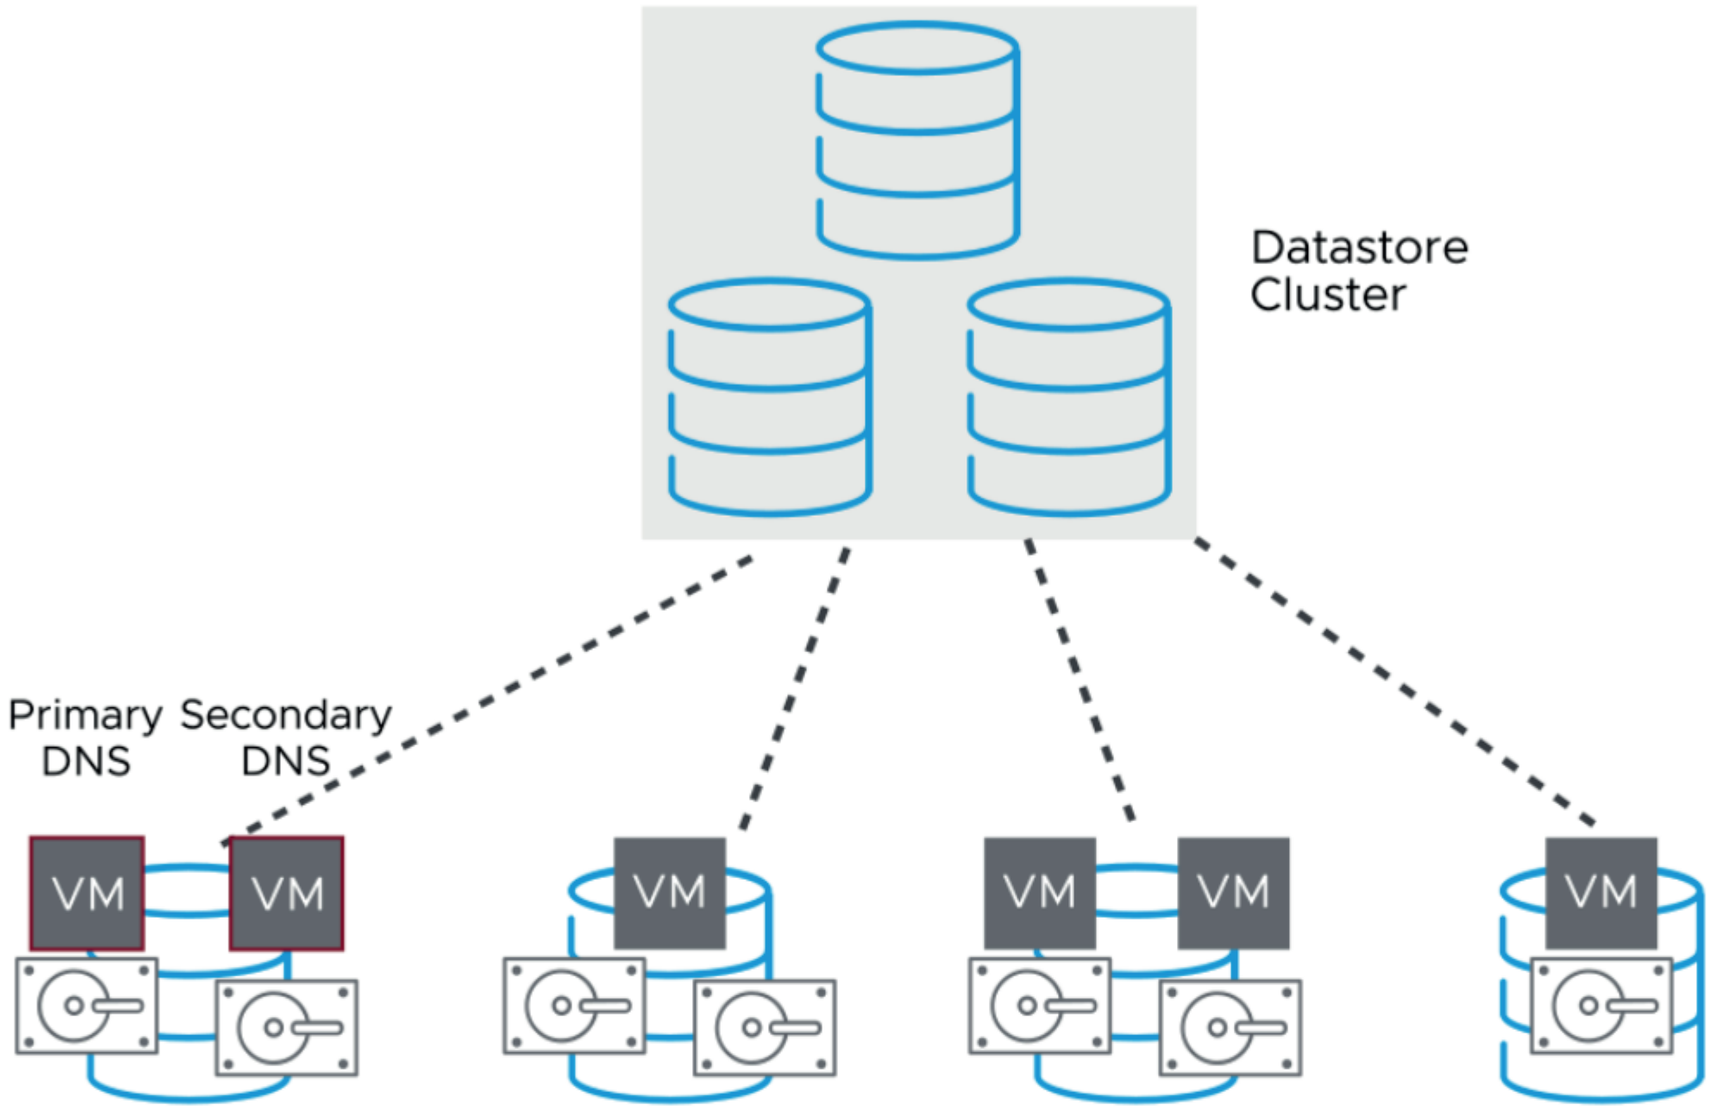

Technet24

5-171 Activity: vSphere Storage DRS (2)

You create a VM anti-affinity rule for the primary DNS and secondary DNS servers.

A VM anti-affinity rule ensures that the VMs listed in the rule are always placed on different datastores.

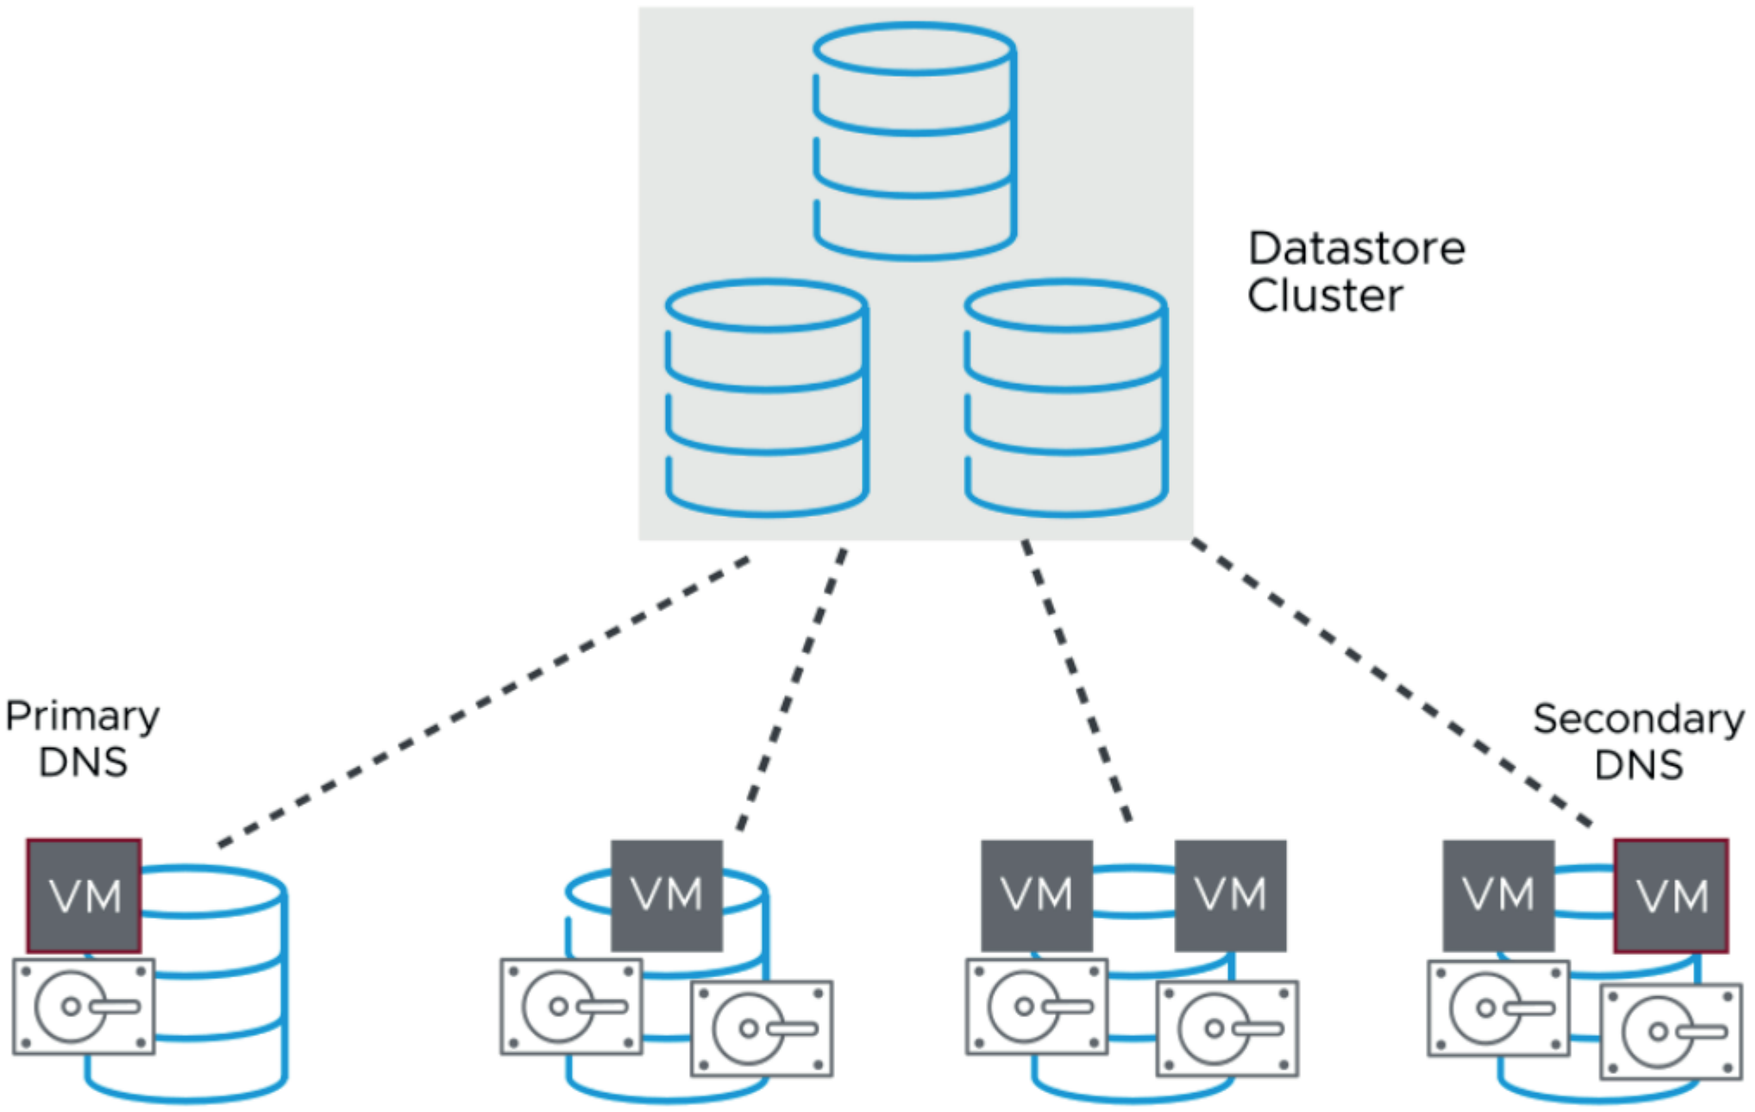

5-172 Review of Learner Objectives

- Create a datastore cluster
- Configure vSphere Storage DRS
- Explain how Storage I/O Control and vSphere Storage DRS complement each other

5-173 Key Points

- vSAN clusters direct-attached server disks to create shared storage designed for VMs.
- With a VMFS6 datastore, the ESXi host automatically sends the SCSI `unmap` command to the storage array, and the unmap rate is configurable.
- With vSphere API for Storage Awareness, storage vendors can provide information about the capabilities of their storage arrays to vCenter.
- With vSphere APIs for I/O Filtering, VMware and third-party vendors can create data services such as caching and replication.
- Storage policy based management abstracts storage and data services delivered by vSAN, vSphere Virtual Volumes, I/O filters, or traditional storage.
- Storage I/O Control provides cluster-wide storage I/O prioritization.
- A datastore cluster that is enabled for vSphere Storage DRS is a collection of datastores working together to balance storage capacity and I/O latency.

Questions?

Technet24

Module 6

vCenter and ESXi Operations

6-2 Importance

vCenter helps you centrally manage multiple ESXi hosts, virtual machines and their networking and storage. If you do not properly deploy, configure, and manage vCenter, your environment might experience reduced administrative efficiency or manageability.

6-3 Module Lessons

1. Backing Up and Restoring vCenter
2. vCenter High Availability
3. vSphere Configuration Profiles
4. VMware Certificate Management

Technet24

6-4 Lesson 1: Backing Up and Restoring vCenter

6-5 Learner Objectives

- Create a vCenter backup schedule
- Restore vCenter from backup
- Configure a backup of vCenter Appliance

6-6 About vCenter Backup and Restore

vCenter backup and restore operations protect data. These operations work in the following ways:

- Use the vCenter Management Interface
- Remove dependency on third-party backup solutions
- Support FTPS, HTTP, HTTPS, SFTP, FTP, NFS, and SMB
- Include the option to encrypt backups
- Restore directly from the vCenter GUI installer
- Restore a vCenter instance to a brand-new appliance

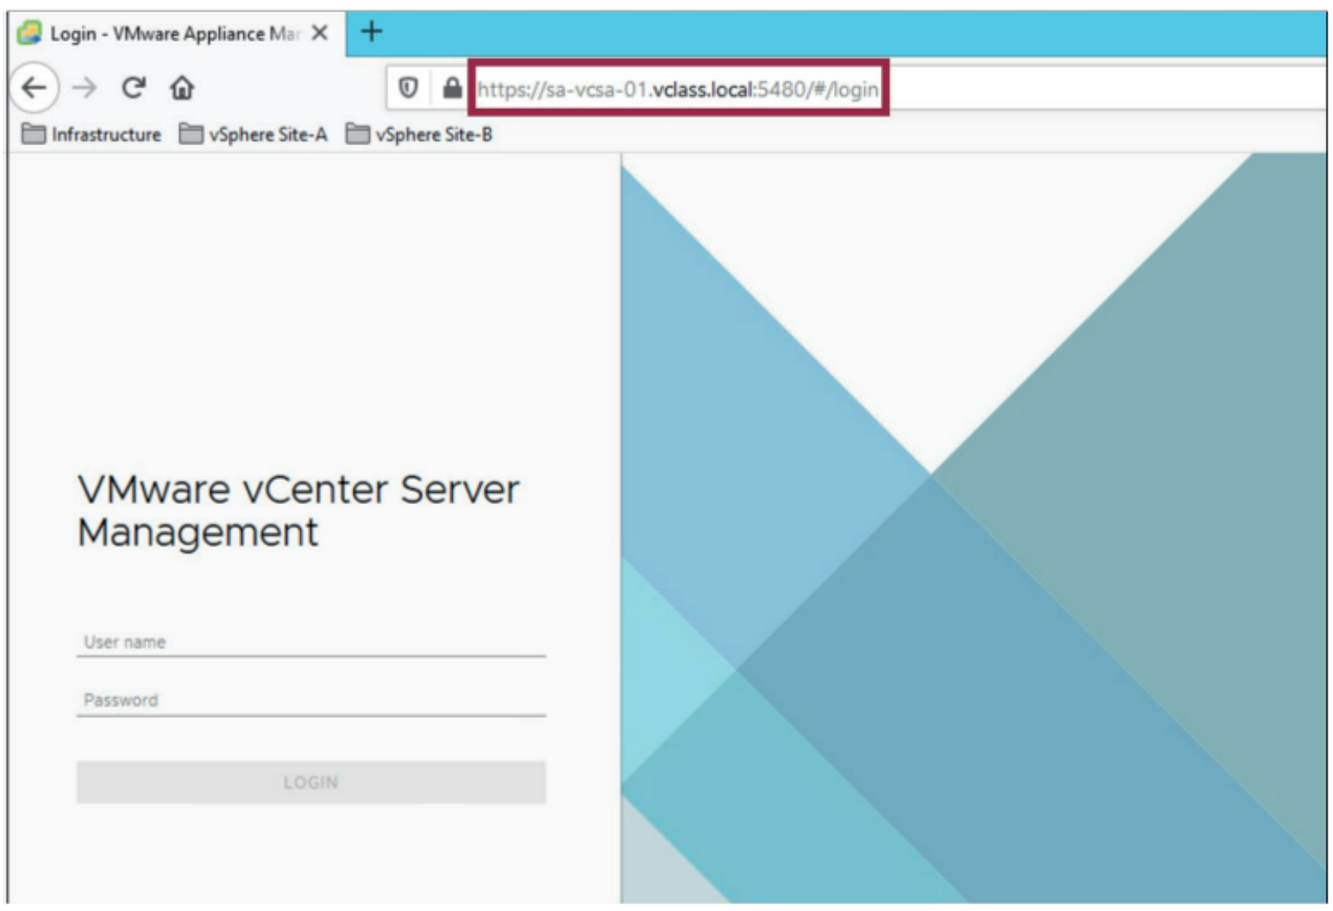

The vCenter Management Interface supports backing up key parts of the appliance. You can protect vCenter data and minimize the time required to restore data center operations.

The backup process collects vCenter key files and the vCenter database into a tar bundle and compresses the bundle to reduce the network load. To minimize the storage impact, the transmission is streamed without caching in the appliance. To reduce the total time required to complete the backup operation, the backup process handles the different components in parallel.

You can encrypt the compressed file before transmission to the backup storage location. When you choose encryption, you must supply a password that can be used to decrypt the file during restoration.

The backup operation always includes the vCenter database and system configuration files, so that a restore operation has all the data to recreate an operational appliance. Optionally, you can specify that a backup operation should include Statistics, Events, and Tasks from the current state of the data center. Current alarms are always included in a backup.

Technet24

6-7 Method for vCenter Backup and Restore

File-based backup and restore:

- Use the vCenter Management Interface to create a file-based backup.
- Restore the backup through the GUI installer of the appliance.
- Schedule the file-based backup and restore.

6-8 Scheduling Backups

You can schedule automatic file-based backups.

The backup scheduler supports:

- A retention policy to keep all backups or a defined number of backups
- Daily, weekly, or custom schedule

Failed backups trigger an alarm in the vSphere Client.

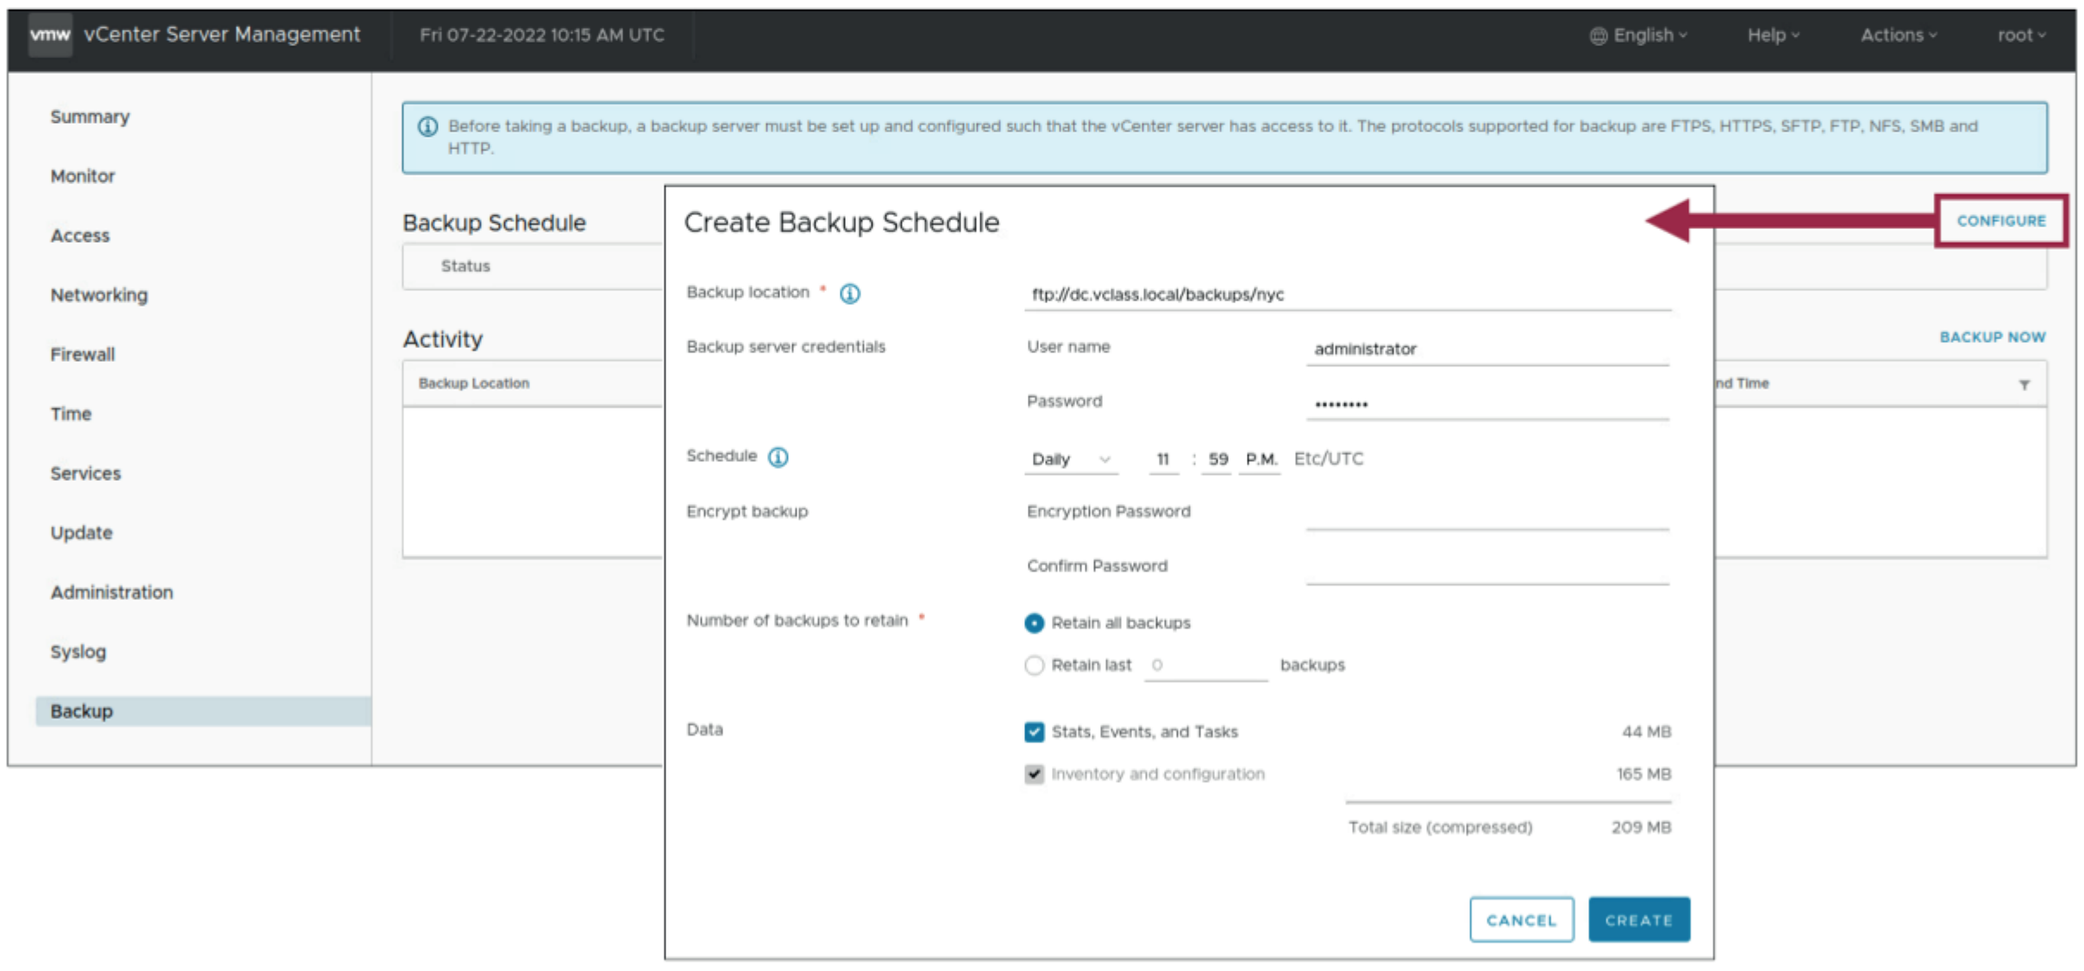

You can set up a file-based backup schedule to perform periodic backups:

- The schedule can be set up with information about the backup location, recurrence, and retention for the backups.
- You can set up only one schedule at a time.

The backup destination should be large enough for N+1 backups.

6-9 Viewing the Backup Schedule

You can view the existing defined backup schedule from the vCenter Management Interface.

The backup schedule can be edited, deactivated, or deleted.

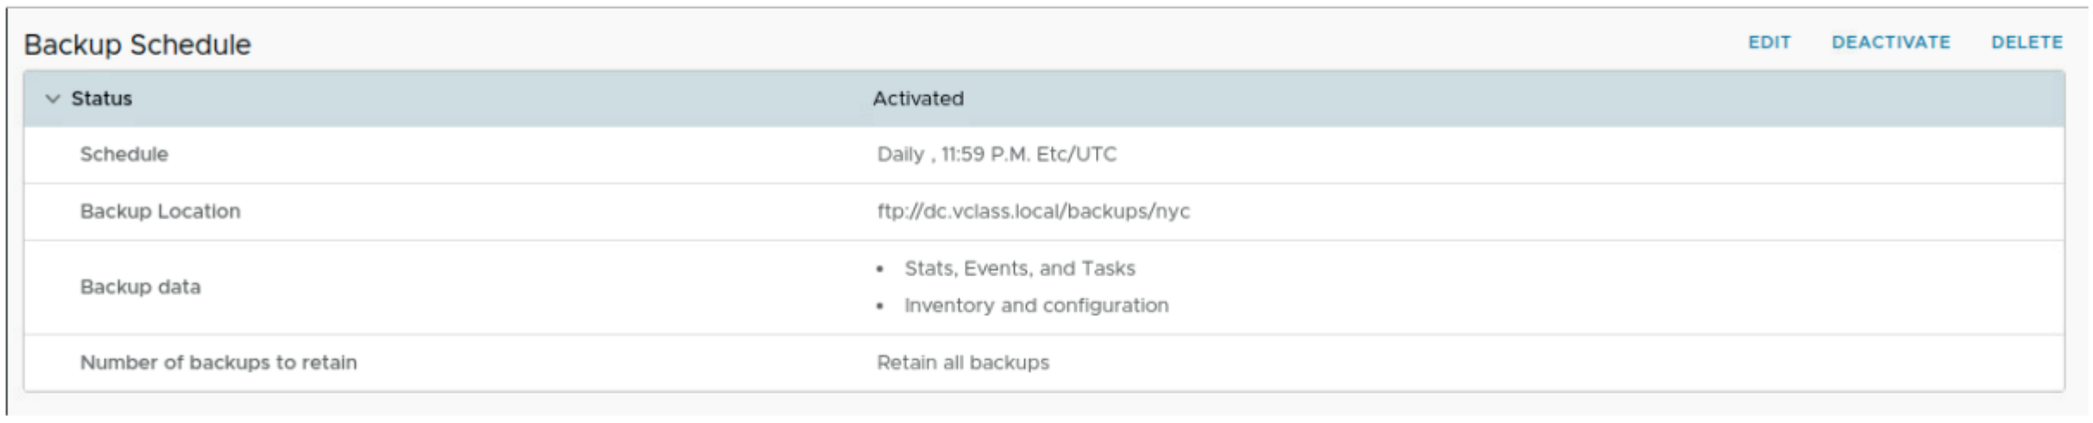

6-10 File-Based Backup of vCenter

You can perform a file-based backup manually.

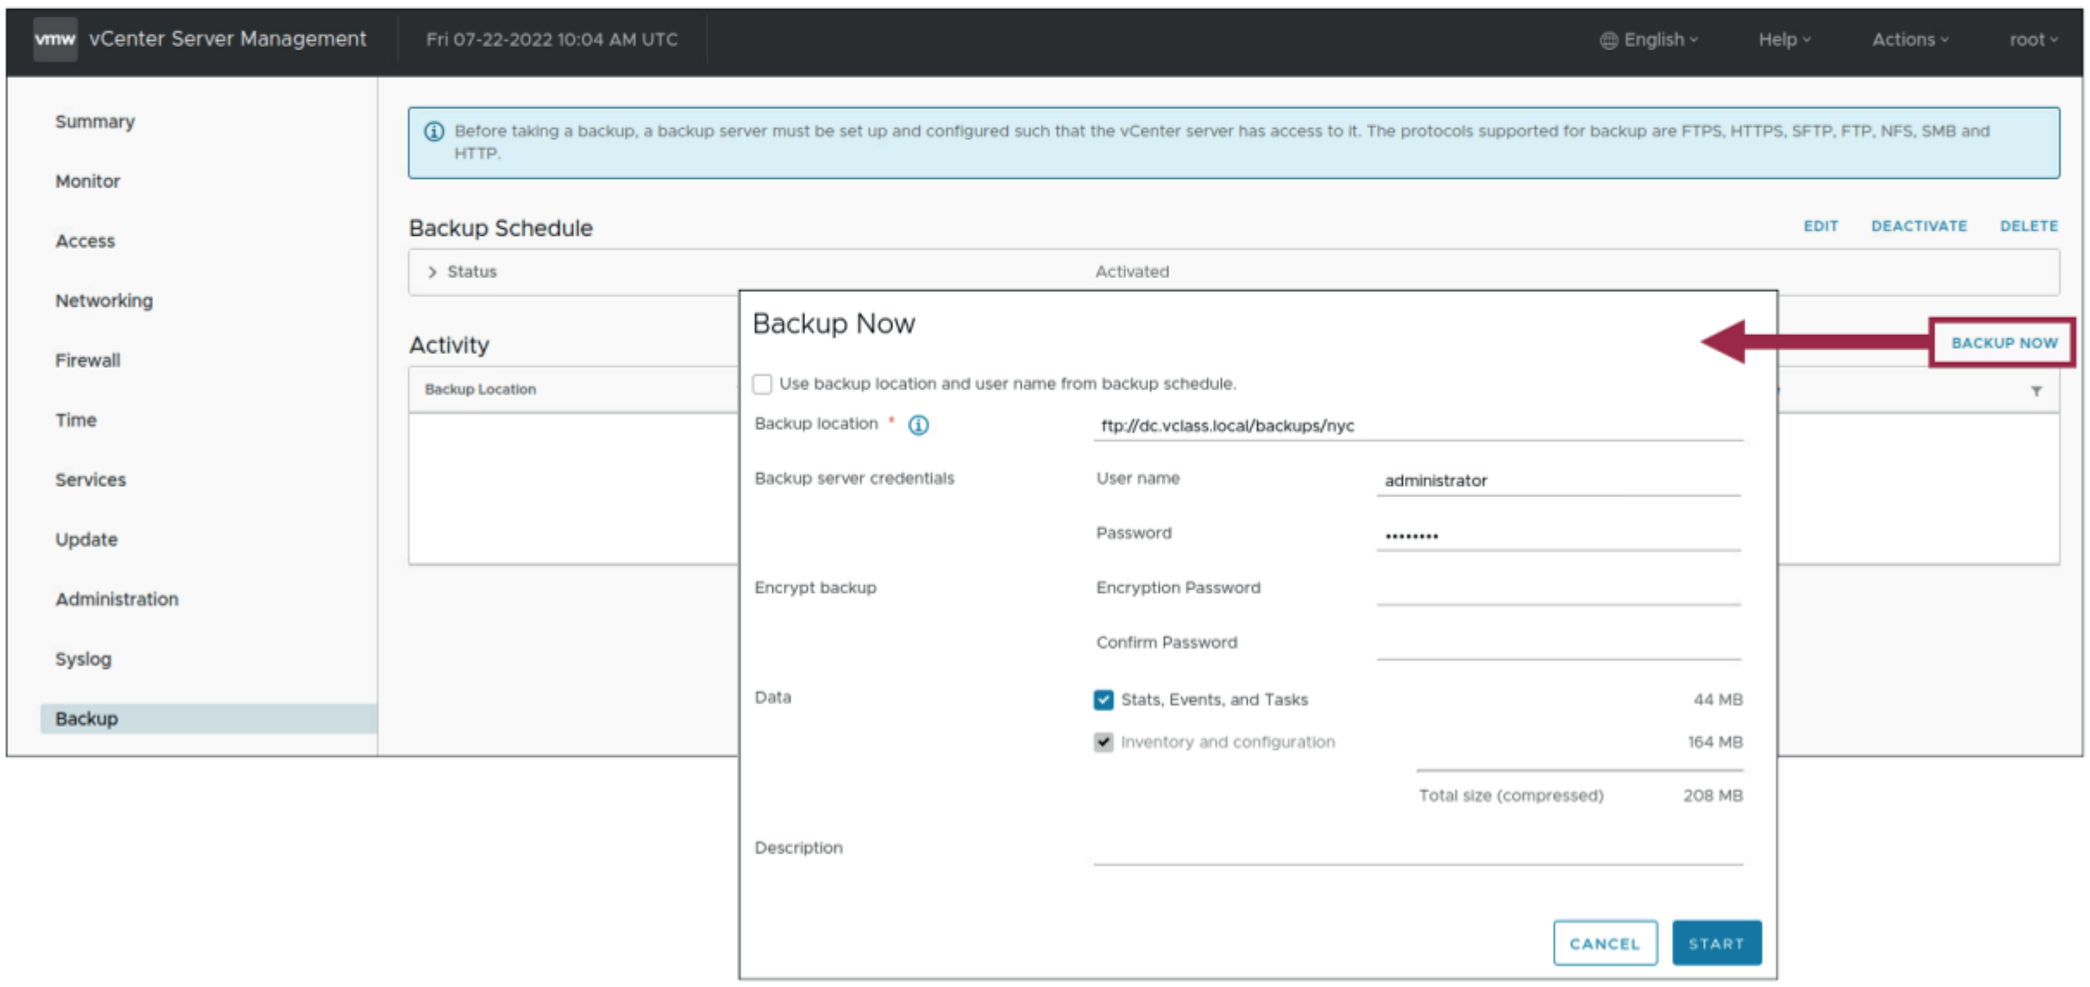

You use the vCenter Management Interface to perform a file-based backup of the vCenter core configuration, inventory, and historical data of your choice. The backed-up data is streamed over the selected protocol to a remote system. The backup is not stored on vCenter.

When specifying the backup location, use the following syntax: `protocol://<server-address>:<port-number>/folder/subfolder`.

Technet24

6-11 File-Based Restore of vCenter

You can use the vCenter GUI installer to restore vCenter to an ESXi host or a vCenter instance.

The restore procedure occurs in stages:

1 A new appliance is deployed.

2 The newly deployed vCenter is populated with the data stored in the file-based backup.

To perform a successful restore:

- The encryption password (if used) must be entered manually.
- The size of the appliance must be the same or larger as at the time of the backup.
- vCenter version numbers must match.

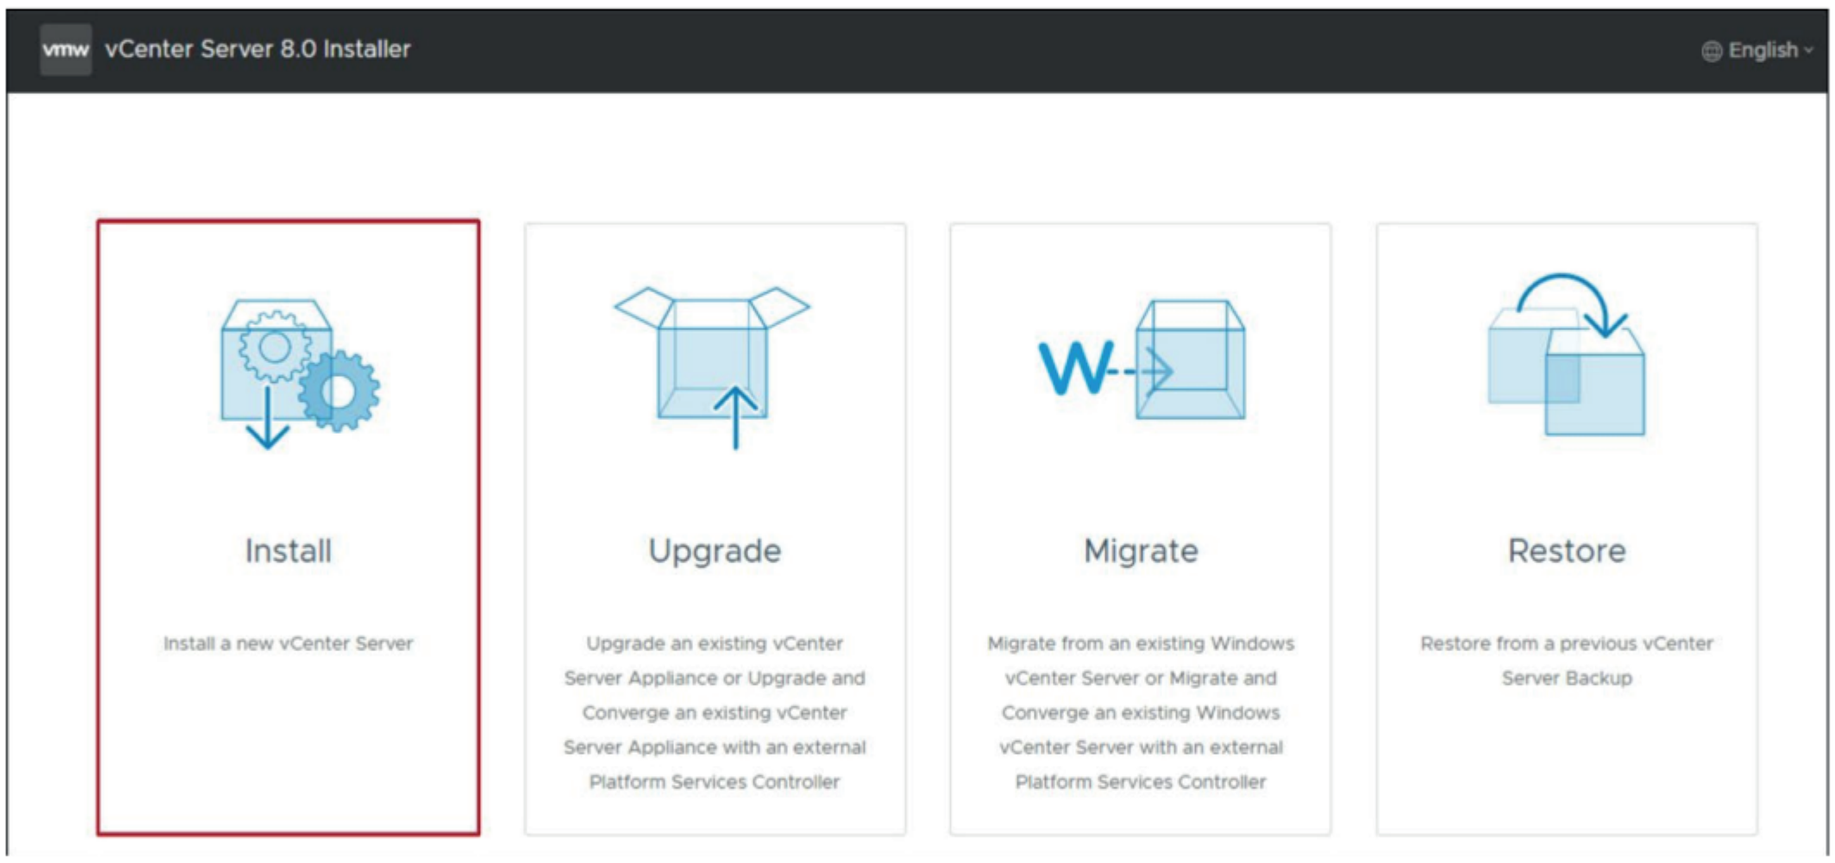

You can perform a file-based restore only for a vCenter instance that you previously backed up by using the vCenter Management Interface. You can perform the restore operation by using the GUI installer of vCenter. The process automates the process of of deploying a new vCenter instance and copying the data from the file-based backup to the new appliance.

You can also perform a restore operation by deploying a new vCenter instance and using the vCenter Management Interface to copy the data from the file-based backup to the new appliance.

The vCenter GUI installer does not support restore from a backup with the NFS or SMB protocol. To perform a restore from an NFS or SMB protocol, you use the vCenter Management Interface.

For more information, see "Restore vCenter Server from a File-Based Backup" at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-vcenter-installation/GUID-F02AF073-7CFD-45B2-ACC8-DE3B6ED28022.html .

6-12 Lab 13: Backing Up vCenter Appliance

Access the vCenter Server Appliance Management Interface and create a backup of vCenter Appliance.

1. Backup vCenter Appliance

6-13 Review of Learner Objectives

- Create a vCenter backup schedule
- Restore vCenter from backup
- Configure a backup of vCenter Appliance

Technet24

6-14 Lesson 2: vCenter High Availability

6-15 Learner Objectives

- Recognize the importance of vCenter High Availability
- Explain how vCenter High Availability works
- Identify vCenter High Availability requirements

6-16 Importance of Keeping vCenter Highly Available

High availability is an important characteristic for many VMware and third-party solutions that depend on vCenter as the primary management platform:

- vCenter is the foundation for the virtual desktop infrastructure.
- Backup and disaster recovery solutions rely on vCenter.

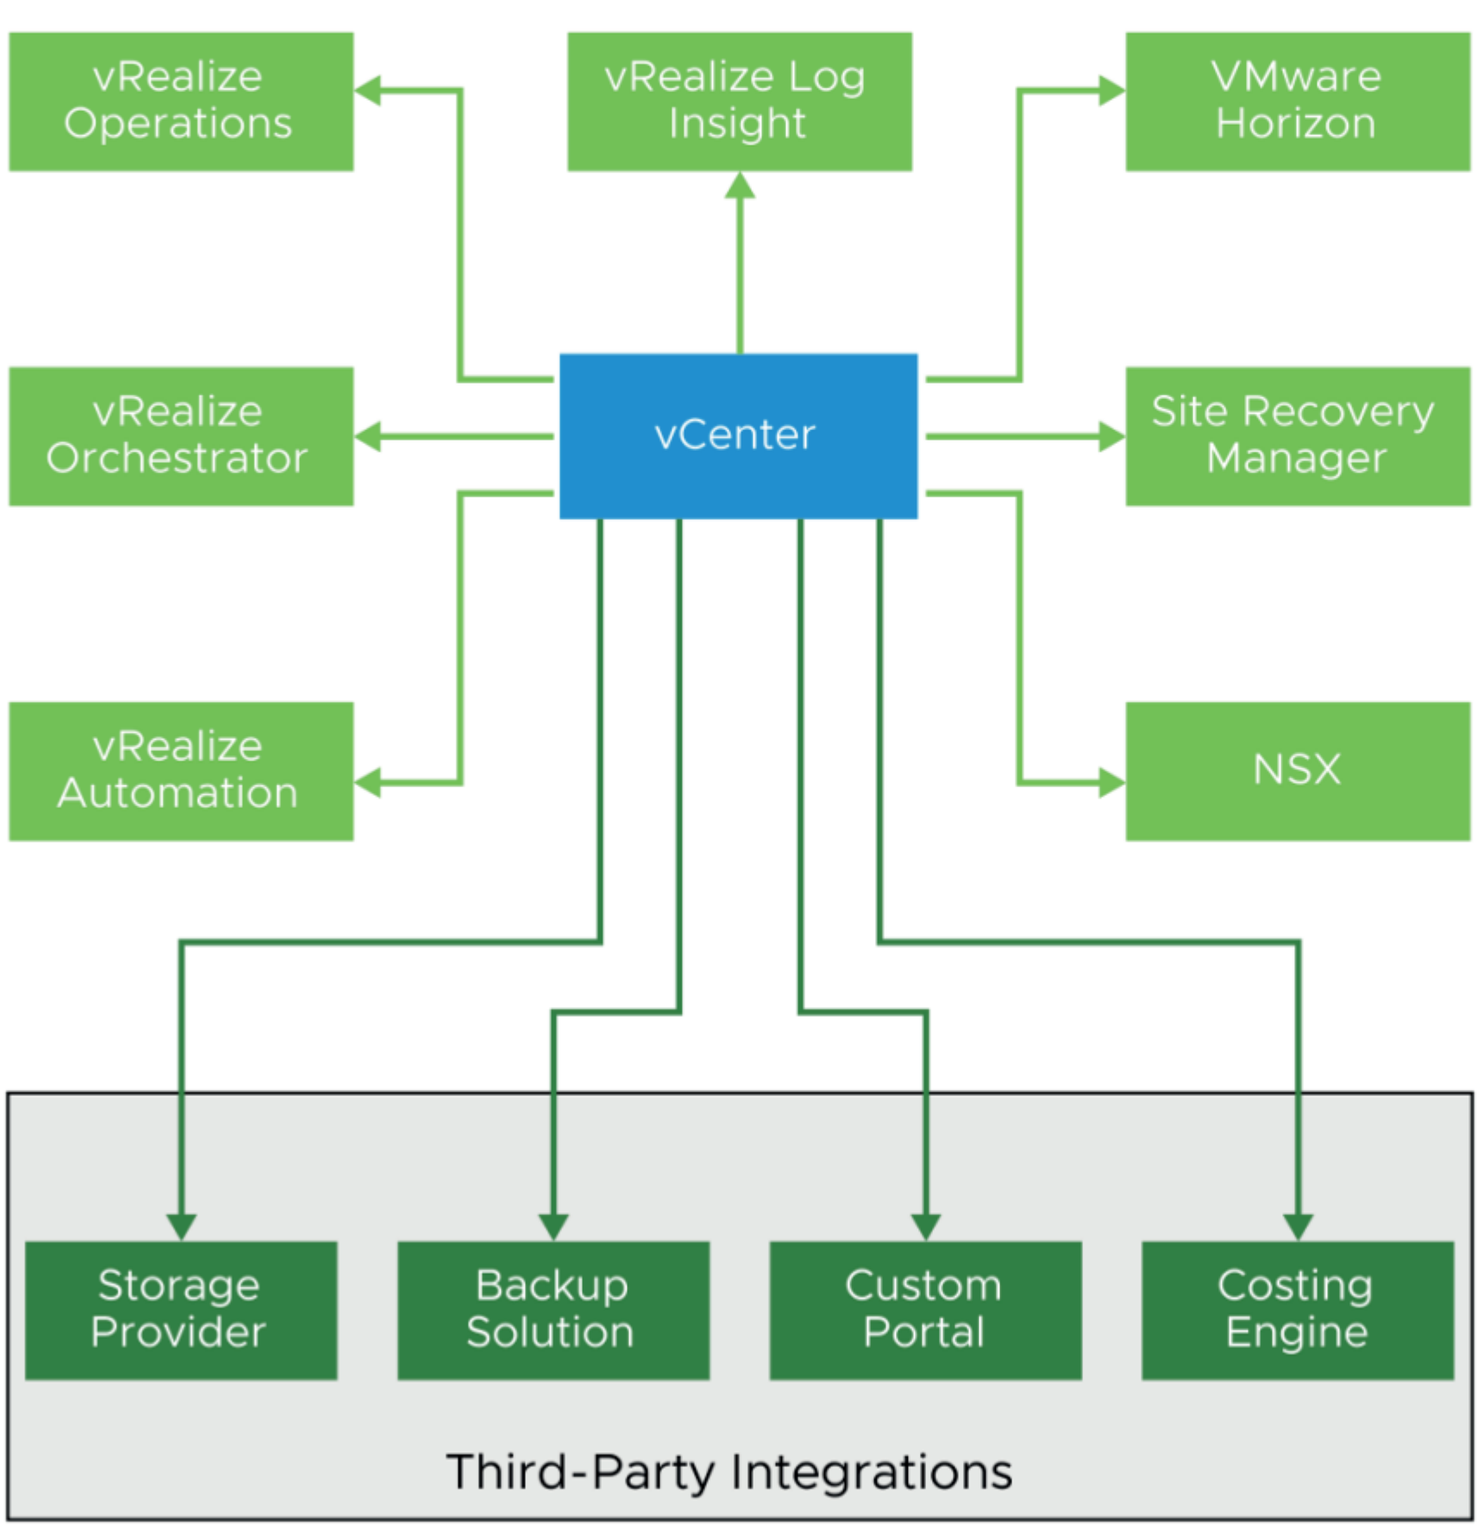

vSphere is a virtualization platform that forms the foundation for building and managing an organization's virtual, public, and private cloud infrastructures. vCenter sits at the heart of vSphere and provides services to manage various components of a virtual infrastructure, such as ESXi hosts, virtual machines, and storage and networking resources. As large virtual infrastructures are built using vSphere, vCenter becomes an important element in ensuring the business continuity of an organization. vCenter must protect itself from a set of hardware and software failures in an environment and must recover transparently from such failures.

Technet24

6-17 About vCenter High Availability

vCenter High Availability protects vCenter against both hardware and software failures.

vCenter High Availability forms a cluster of nodes:

- Active node: Runs the active vCenter instance
- Passive node: Automatically takes over the role of the Active node if a failure occurs
- Witness node: Provides a quorum to protect against a split-brain situation

vCenter High Availability is built in to vCenter and is included with the standard license.

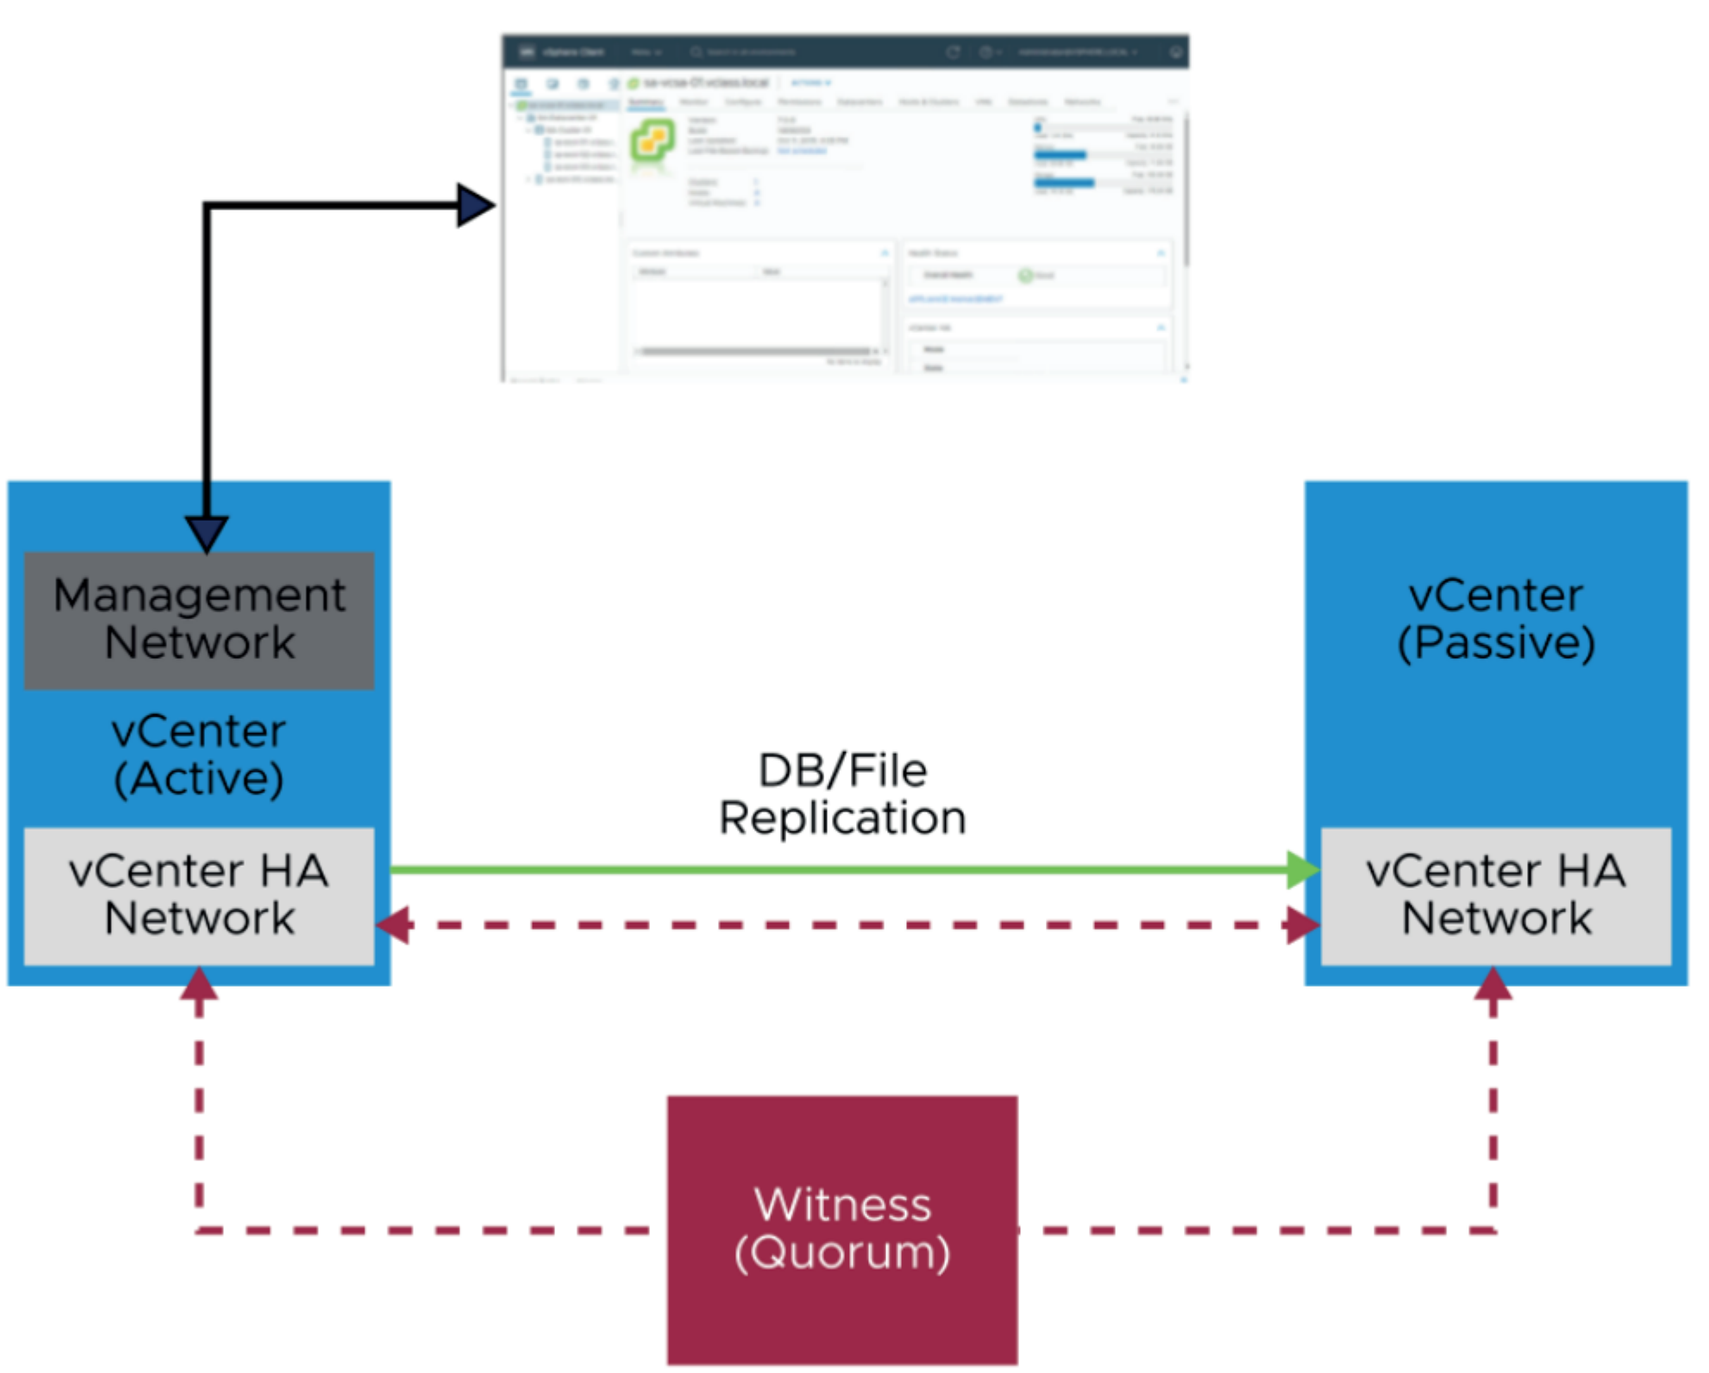

With vCenter High Availability, vCenter will recover automatically from failure with minimal downtime.

6-18 Scenario: Active Node Failure

If the active node fails, the passive node takes over the role of the active node. The cluster is considered to be running in a degraded state.

The animation demonstrates what happens if an active node fails. To play the animation, go to https://vmware.bravais.com/s/PlUBZn2zCO7HE5qN2fm4.

The active node runs the active instance of vCenter. The node uses an IP address on the Management network for the vSphere Client and ESXi hosts to connect to.

If the active node fails (because of a hardware, software, or network failure), the passive node takes over the role of the active node. The IP address to which the vSphere Client was connected is switched from the failed node to the new active node. The new active node starts serving client requests. Meanwhile, the user must log back in to the vSphere Client for continued access to vCenter.

Because only two nodes are up and running, the vCenter High Availability cluster is considered to be running in a degraded state and subsequent failover cannot occur. A subsequent failure in a degraded cluster means vCenter services are no longer available. A passive node is required to return the cluster to a healthy state.

Technet24

6-19 Scenario: Passive Node Failure

If the passive node fails, the active node continues to operate normally. However, the cluster is considered to be running in a degraded state.

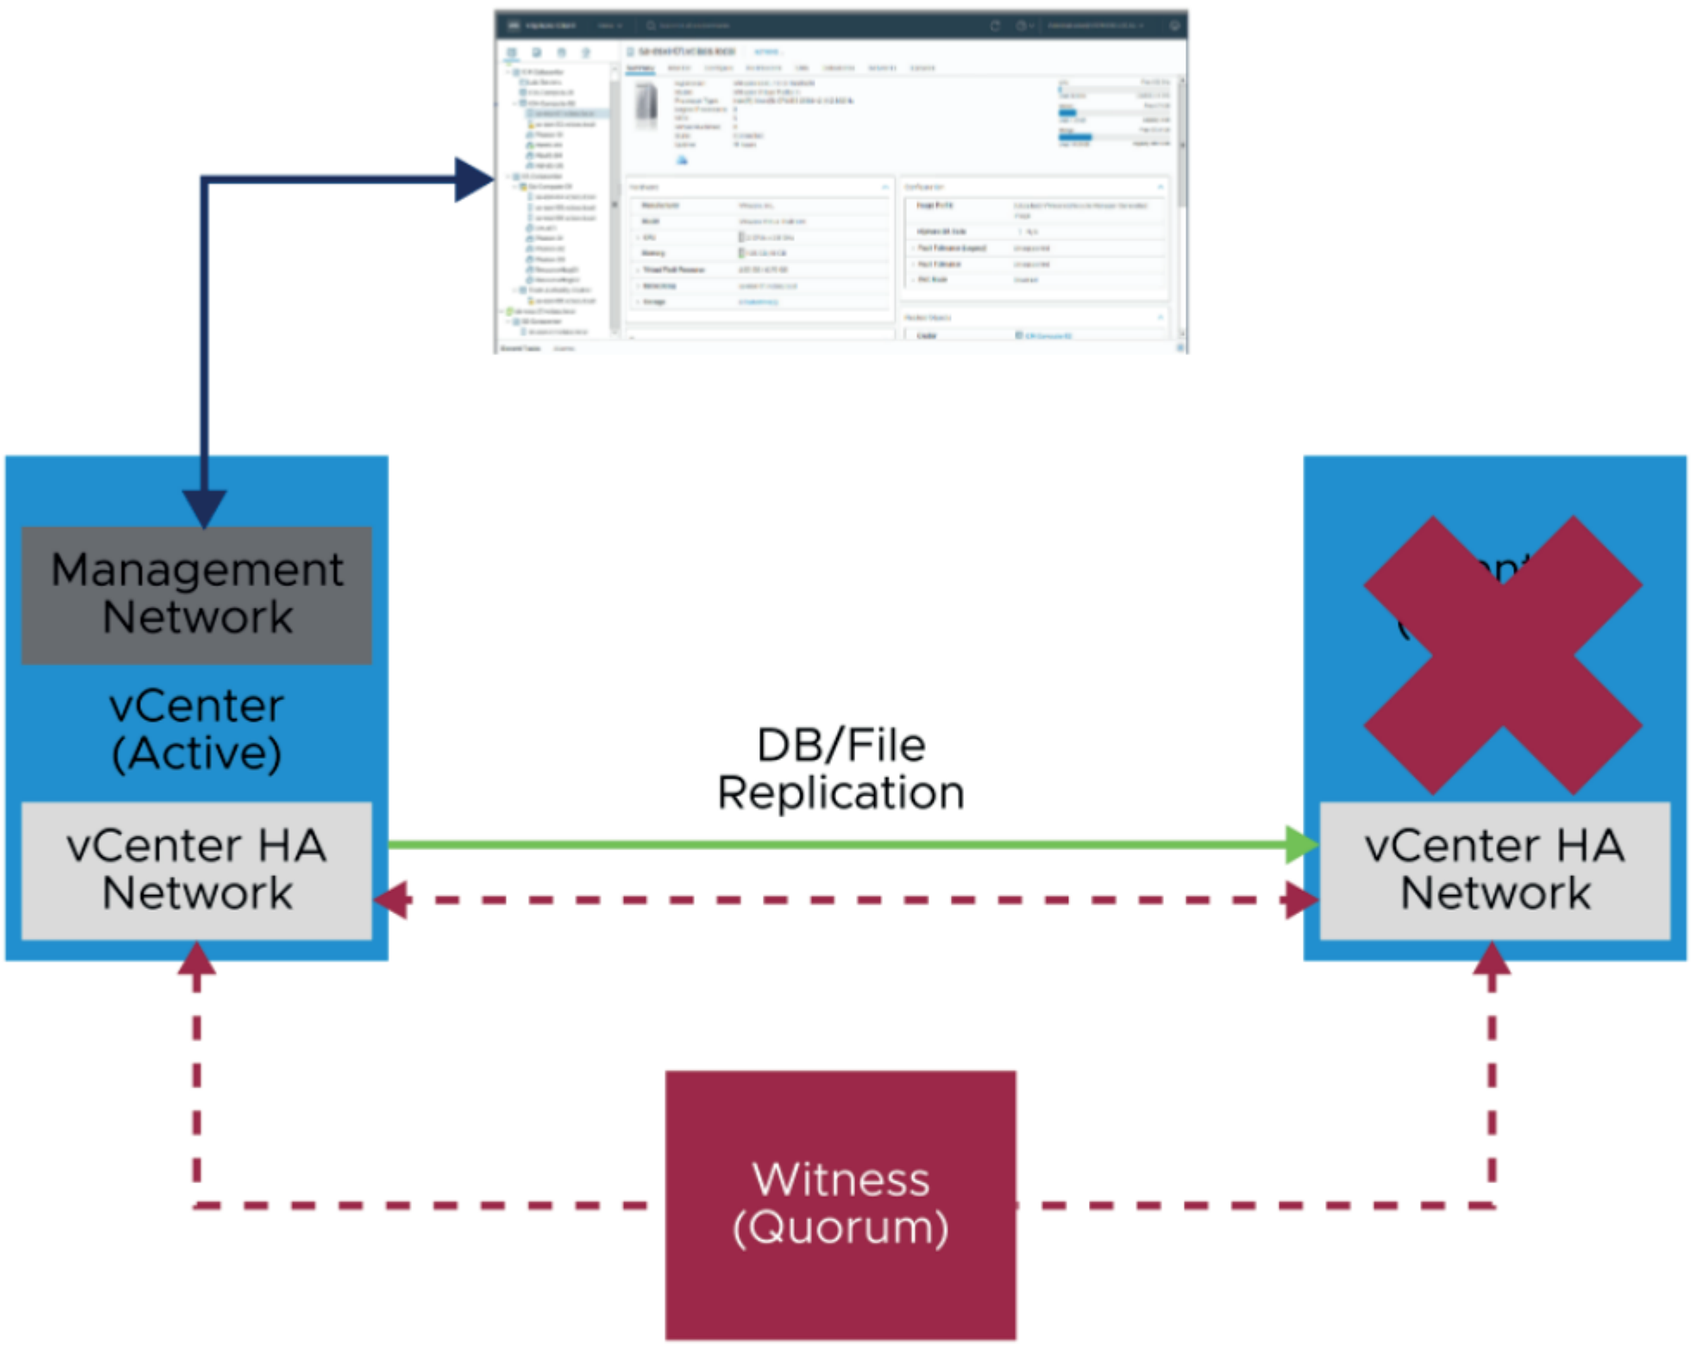

If the passive node fails, the active node continues to operate as normal. Because no disruption in service occurs, users can continue to access the active node using the vSphere Client.

Because the passive node is down, the active node is no longer protected. The cluster is considered to be running in a degraded state because only two nodes are up and running. A subsequent failure in a degraded cluster means vCenter services are no longer available. A passive node is required to return the cluster to a healthy state.

6-20 Scenario: Witness Node Failure

If the witness node fails, the active node continues to operate normally. However, the cluster is considered to be running in a degraded state.

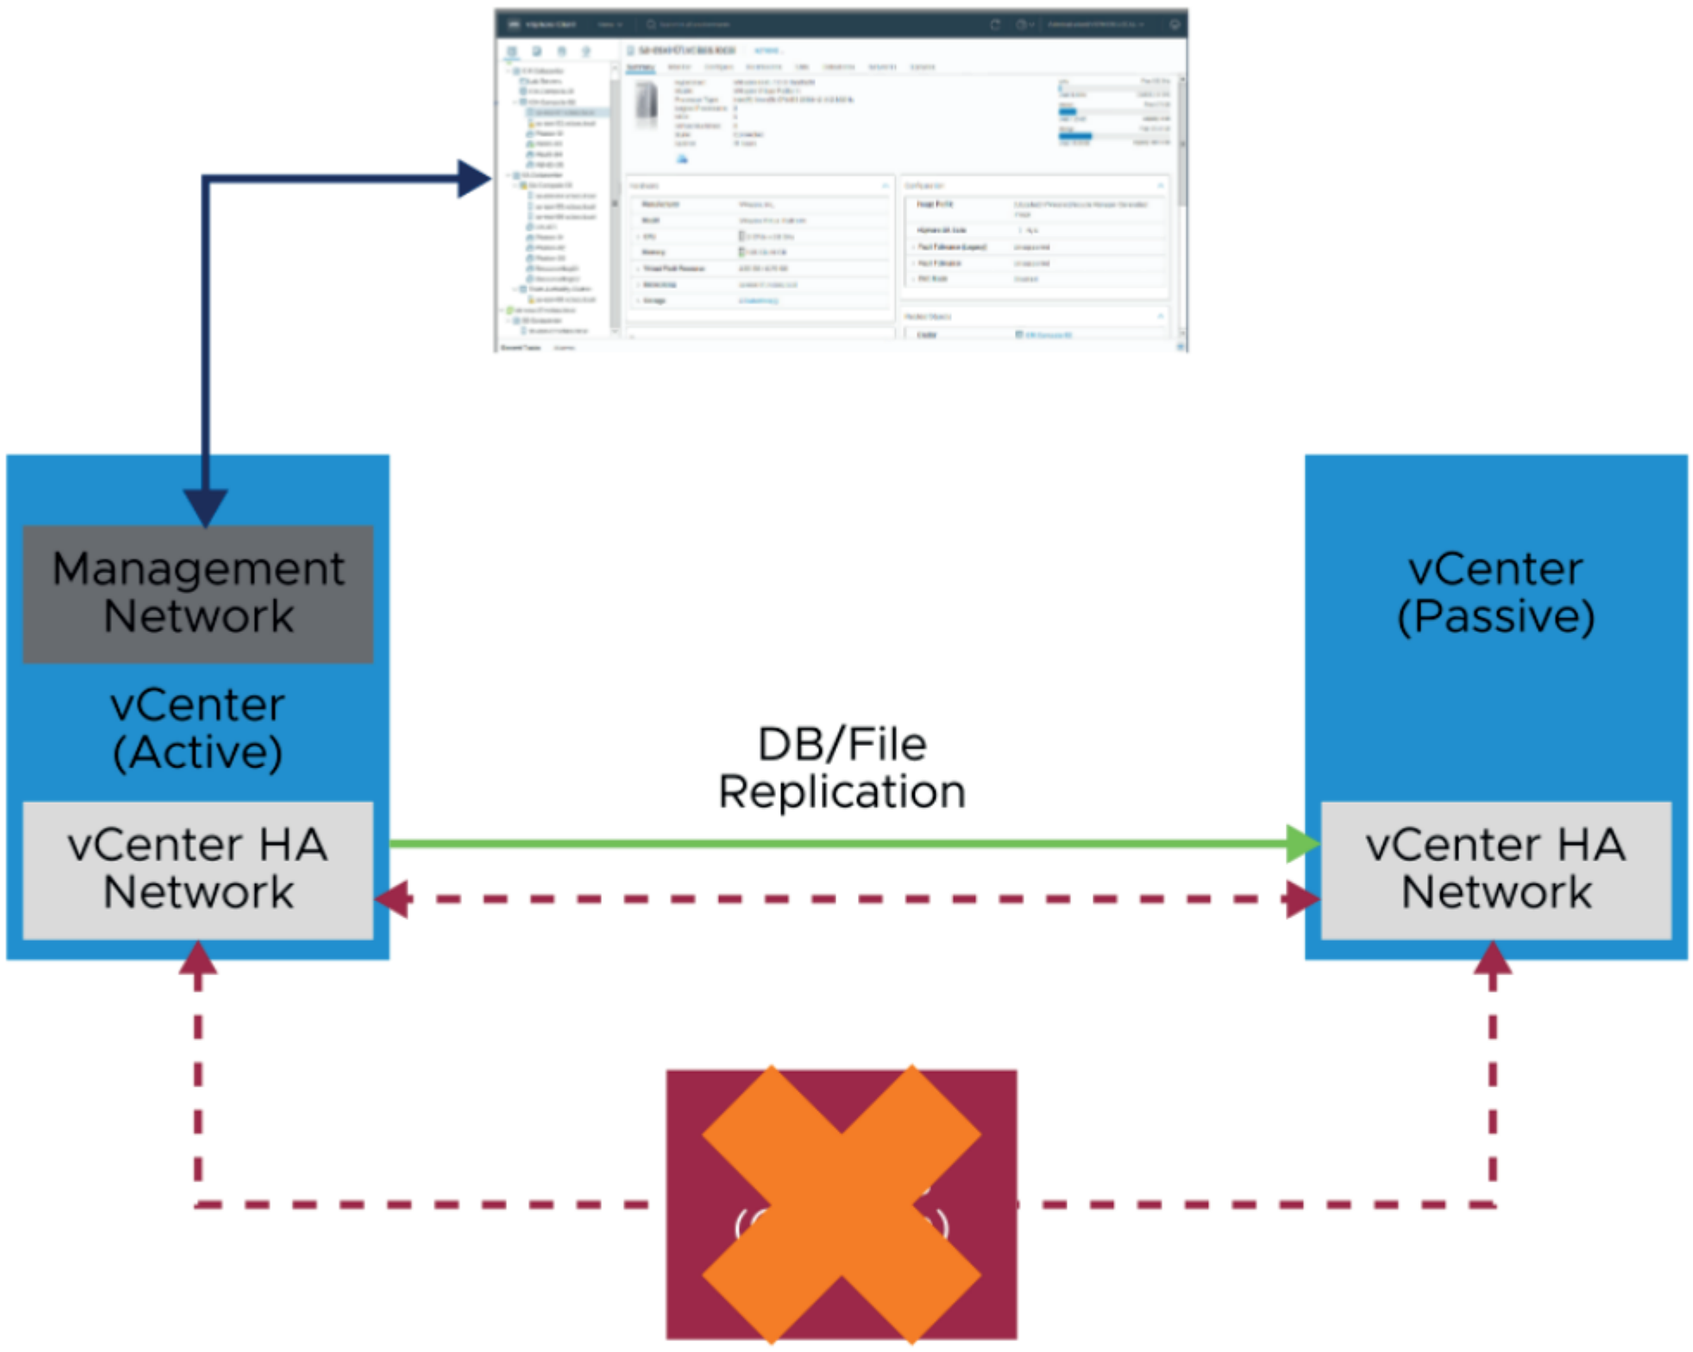

The witness node is used to maintain quorum.

If the witness node fails, the active node continues to operate without disruption in service. Because only two nodes are up and running, the cluster is considered to be running in a degraded state and failover cannot occur. A subsequent failure in a degraded cluster means vCenter services are no longer available. The witness node is required to return the cluster to a healthy state.

Technet24

6-21 Benefits of vCenter High Availability

vCenter High Availability provides many benefits:

- vCenter is made more resilient.
- Protection against hardware, host, and application failures is provided.
- Recovery occurs in minutes. End-to-end downtime is minimized.
- Active-passive architecture provides transparent failover:
  - Recovery point objective: No data loss.
  - Recovery time objective: Within minutes.
- Nodes can be geographically distant (less than 10 milliseconds latency).
- No shared storage is required.
- No third-party technologies are required.

6-22 vCenter High Availability Requirements

Component	Requirements
ESXi	• Minimum of three ESXi hosts is recommended.
vCenter	• Deployment size small or larger is required to meet the RTO. • Enough disk space to collect and store support bundles for all three nodes on the active node.
Network connectivity	• Network latency between the three nodes must be less than 10 milliseconds. • The vCenter HA network must be on a different subnet than the management network.
Licensing	• A single vCenter Standard license.

For more information about the vCenter High Availability requirements, see *vSphere Availability* at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-availability-guide.pdf.

6-23 Demonstration: Configuring vCenter High Availability

Your instructor will run a demonstration.

Demonstration - https://vmware.bravais.com/s/erHyJPGv9PAkR7gLsLjy

6-24 Review of Learner Objectives

- Recognize the importance of vCenter High Availability
- Explain how vCenter High Availability works
- Identify vCenter High Availability requirements

Technet24

6-25 Lesson 3: vSphere Configuration Profiles

6-26 Learner Objectives

- Recognize benefits of using vSphere Configuration Profiles
- Explain how the Host Profiles feature works
- Configure and manage vSphere Configuration Profiles

6-27 About vSphere Configuration Profiles (1)

Introduced in vSphere 8, administrators can use vSphere Configuration Profiles to test the functionality and provide feedback during the development process.

From vSphere 8.0 Update 1 and later, vSphere Configuration Profiles are fully supported.

With vSphere Configuration Profiles, administrators can:

- Manage the configuration of all hosts in a cluster collectively
- Set up a desired configuration that you can apply to all hosts in a single operation
- Use a reference host and make its configuration the desired configuration for an entire cluster

In vSphere 8, vSphere Configuration Profiles was introduced as a Technology Preview and not fully supported.

6-28 About vSphere Configuration Profiles (2)

Administrators can enable vSphere Configuration Profiles only on clusters that you manage with a single image.

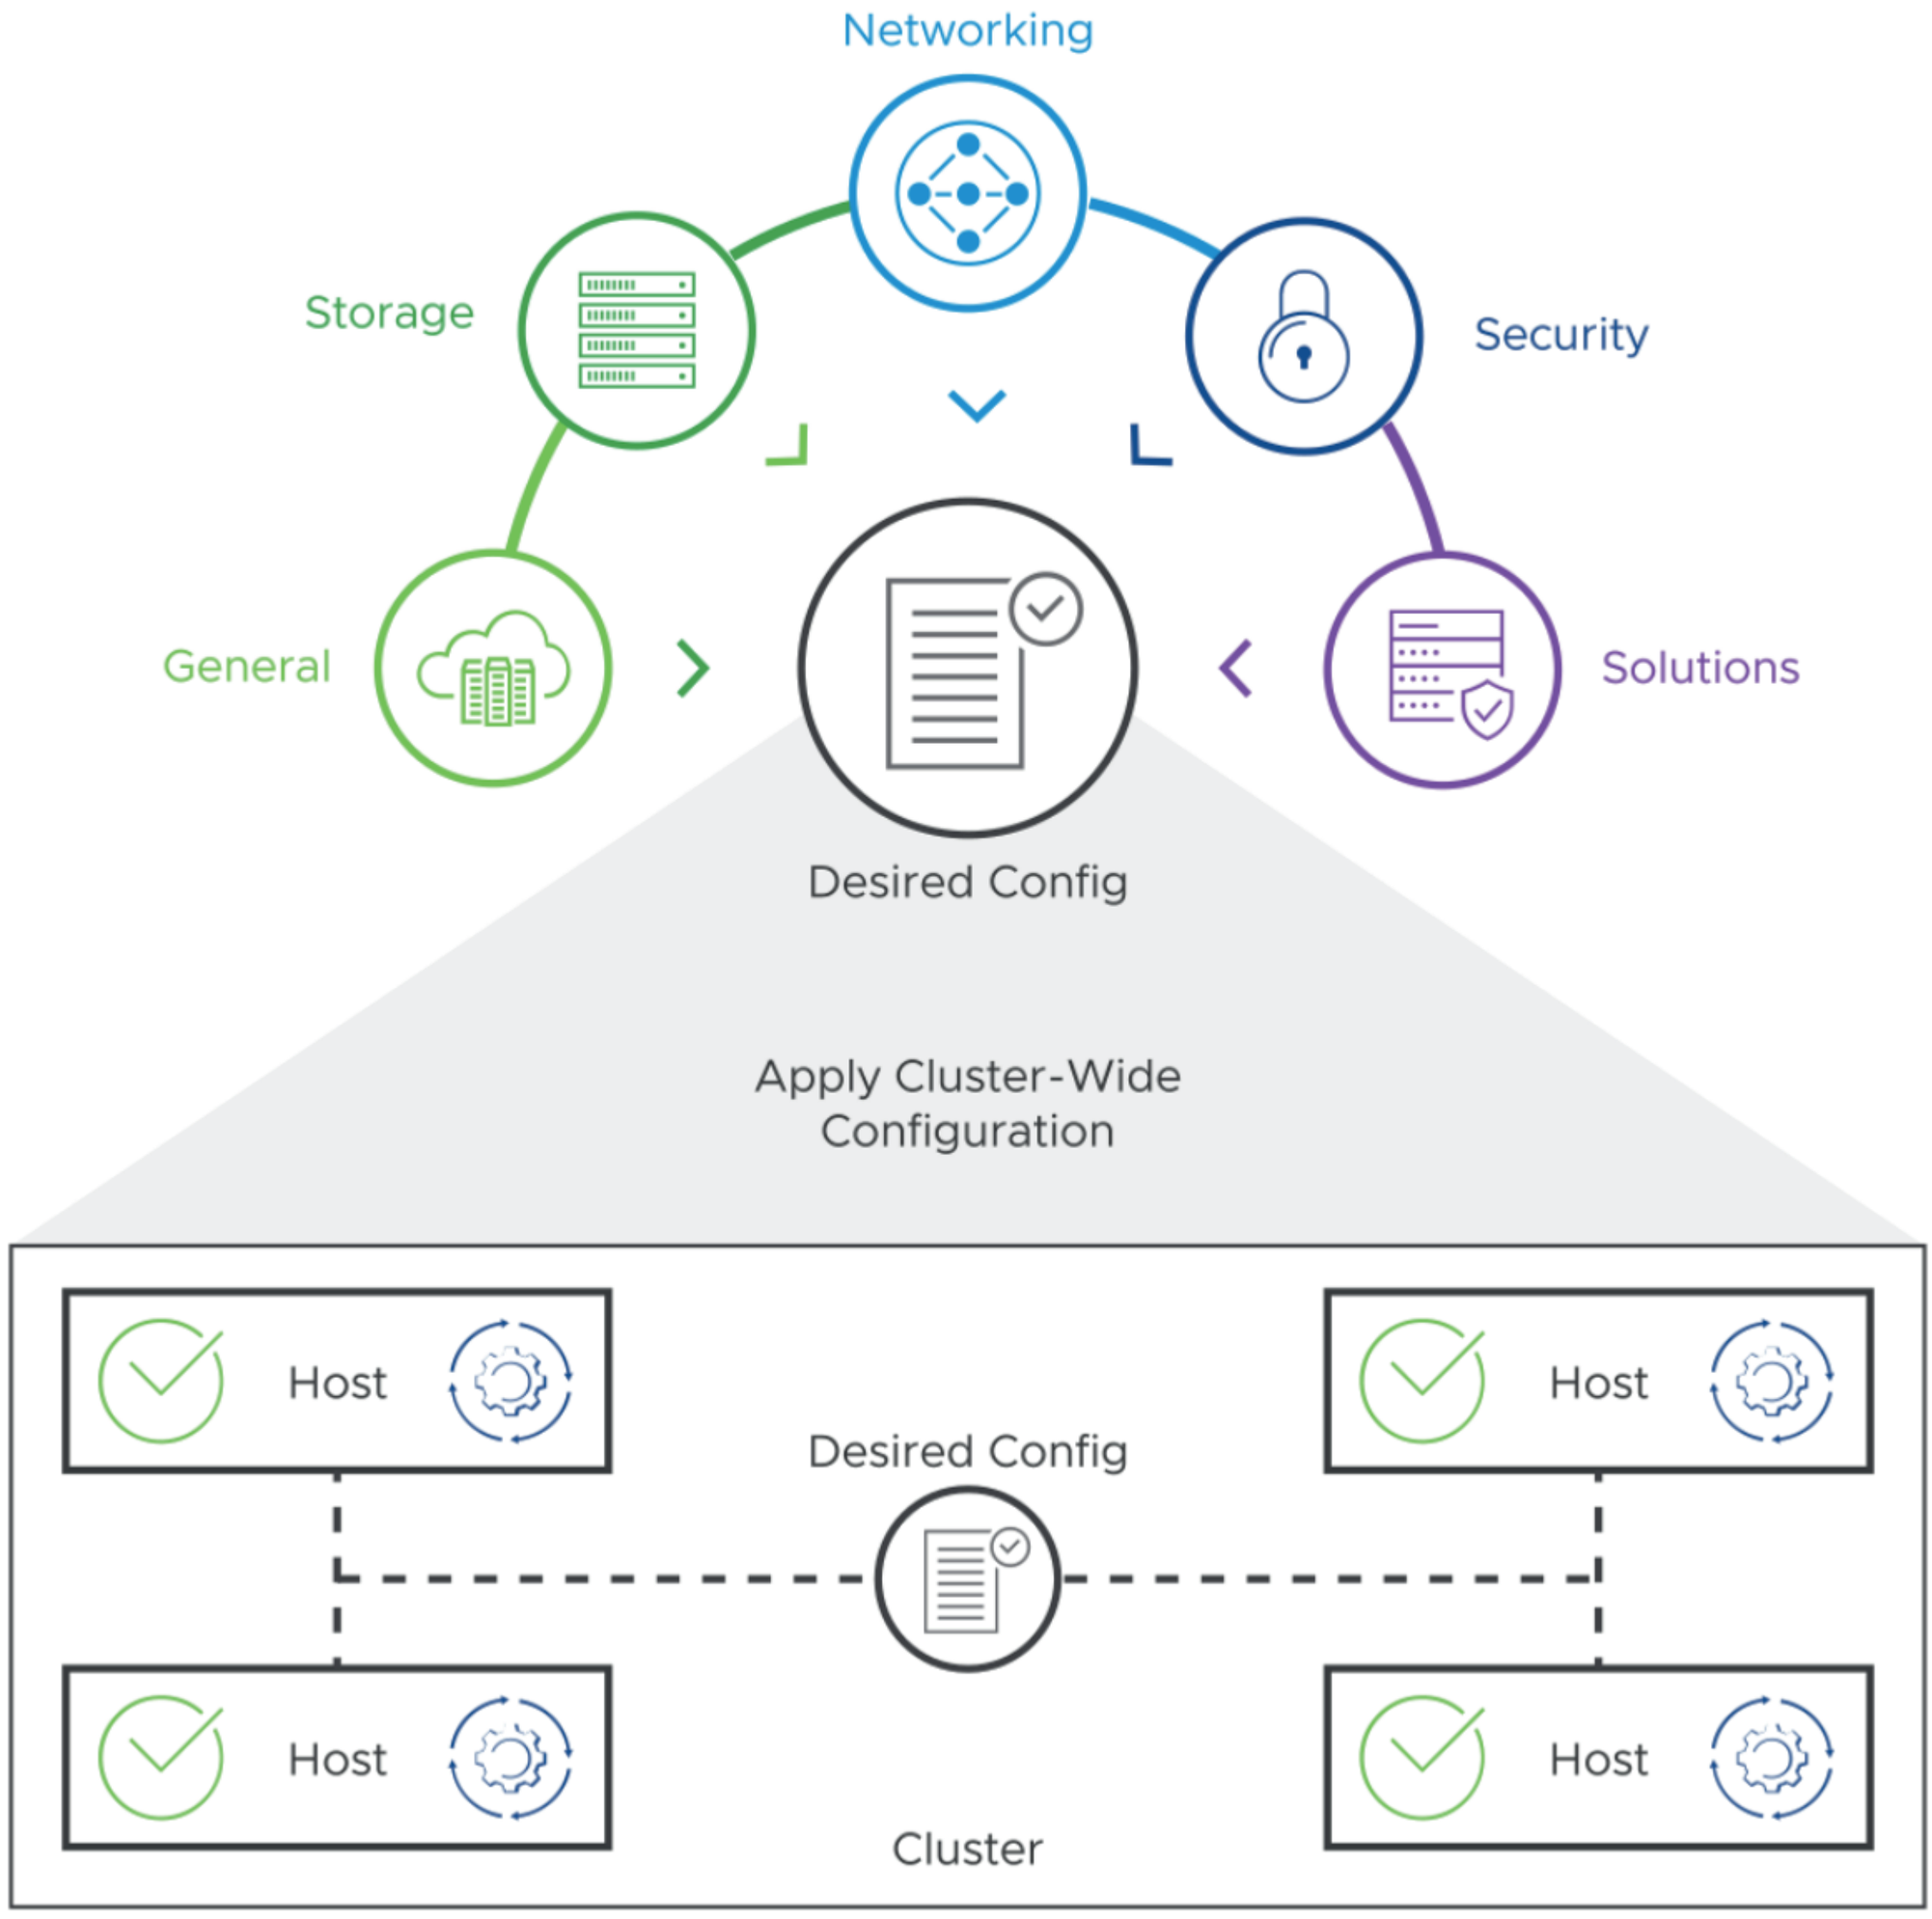

Technet24

6-29 Requirements for vSphere Configuration Profiles

To use vSphere Configuration Profiles, you must ensure that your environment meets the following requirements:

- vCenter Server 8.0 Update 1
- ESXi 8.0 Update 1
- Enterprise Plus license

6-30 Limitations to vSphere Configuration Profiles

vSphere Configuration Profiles come with the following limitations:

- Once activated on a cluster, vSphere Configuration Profiles cannot be deactivated.
- Does not support a cluster if NSX is also enabled
- Does not support a cluster with hosts with DPU installed.
- Some ESXi services startup and runtime states are not currently managed by vSphere Configuration Profiles.
    - For example, the SSH service startup and runtime settings are not included in a cluster configuration document.

For example, the SSH service startup and runtime settings are not included in a cluster configuration document.

6-31 Cluster-Level Configuration

In vSphere 8 Update 1 and later, managing the ESXi host configuration at the cluster level can be found in **Desired State > Configuration** under the cluster object's Configure tab.

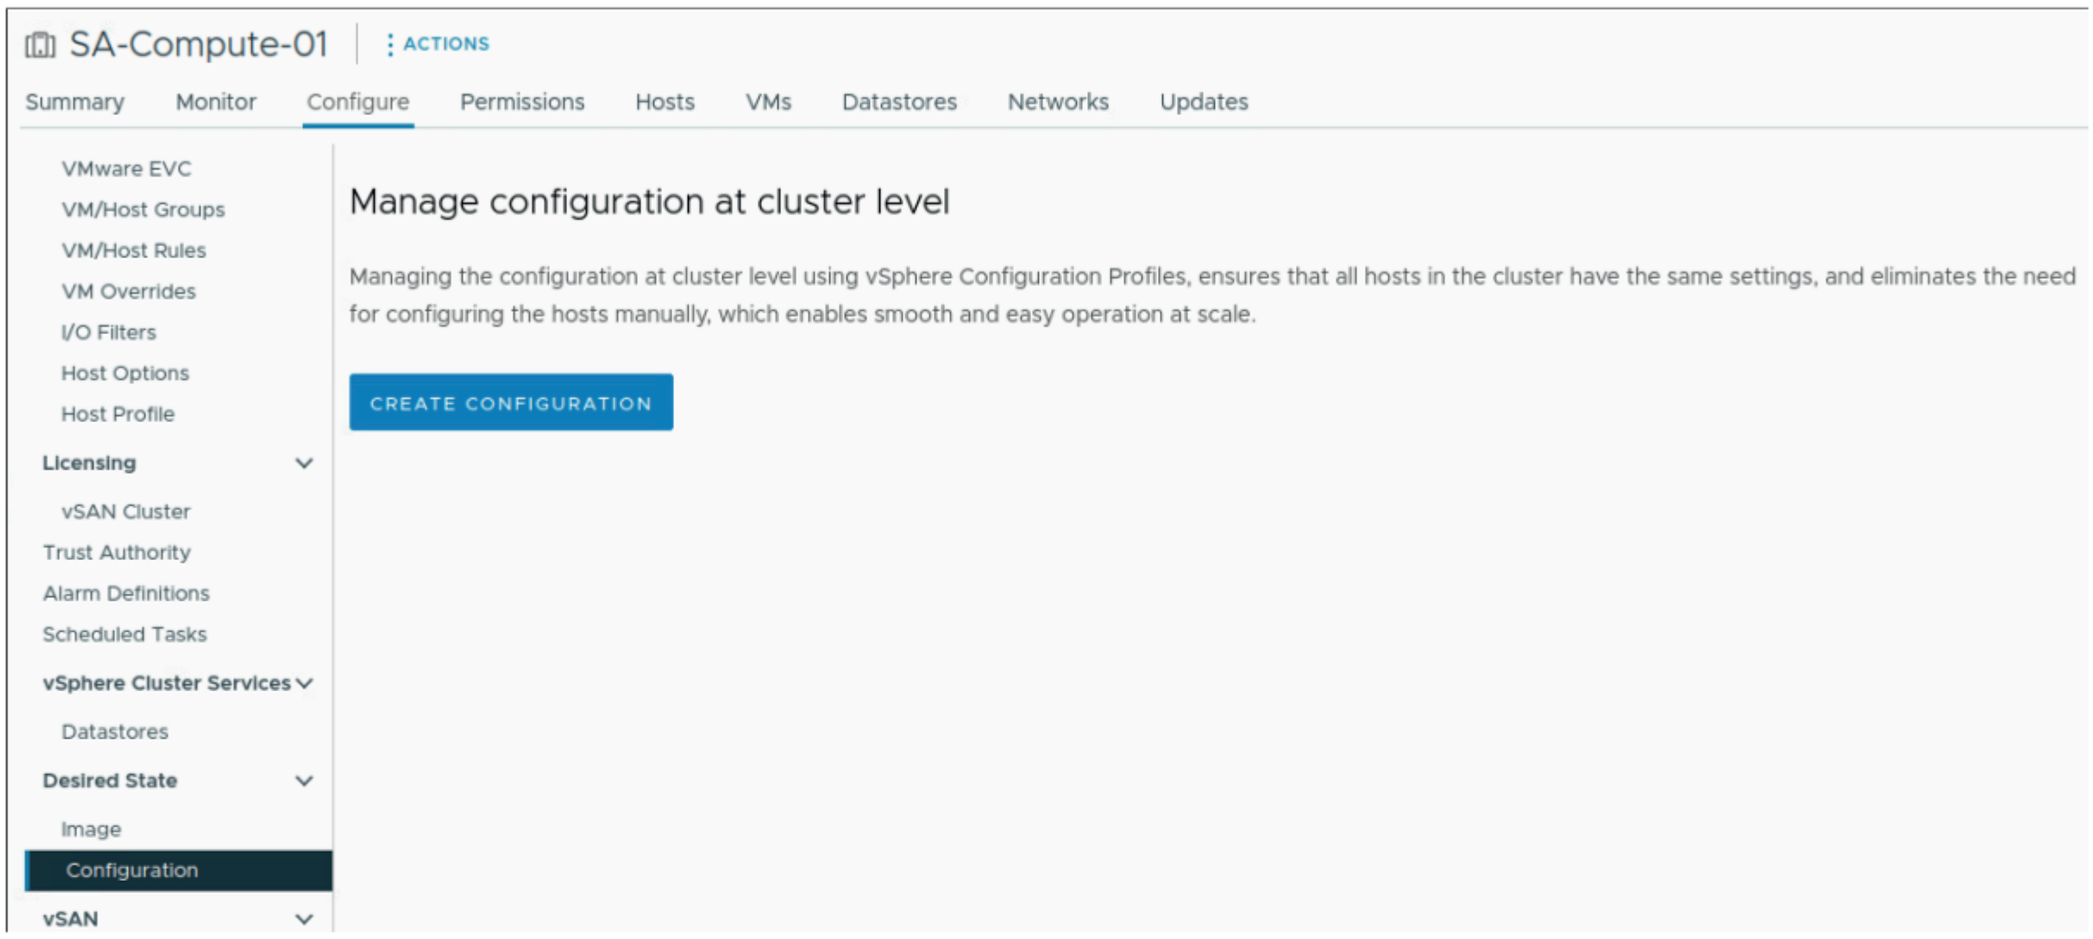

Technet24

6-32 Using vSphere Configuration Profiles

The configuration workflow starts with eligibility checks on the cluster and hosts.

- If the cluster is not eligible, the configuration process will not continue
- If the cluster is eligible, administrators can import a configuration from a host in the cluster or from a JSON file.

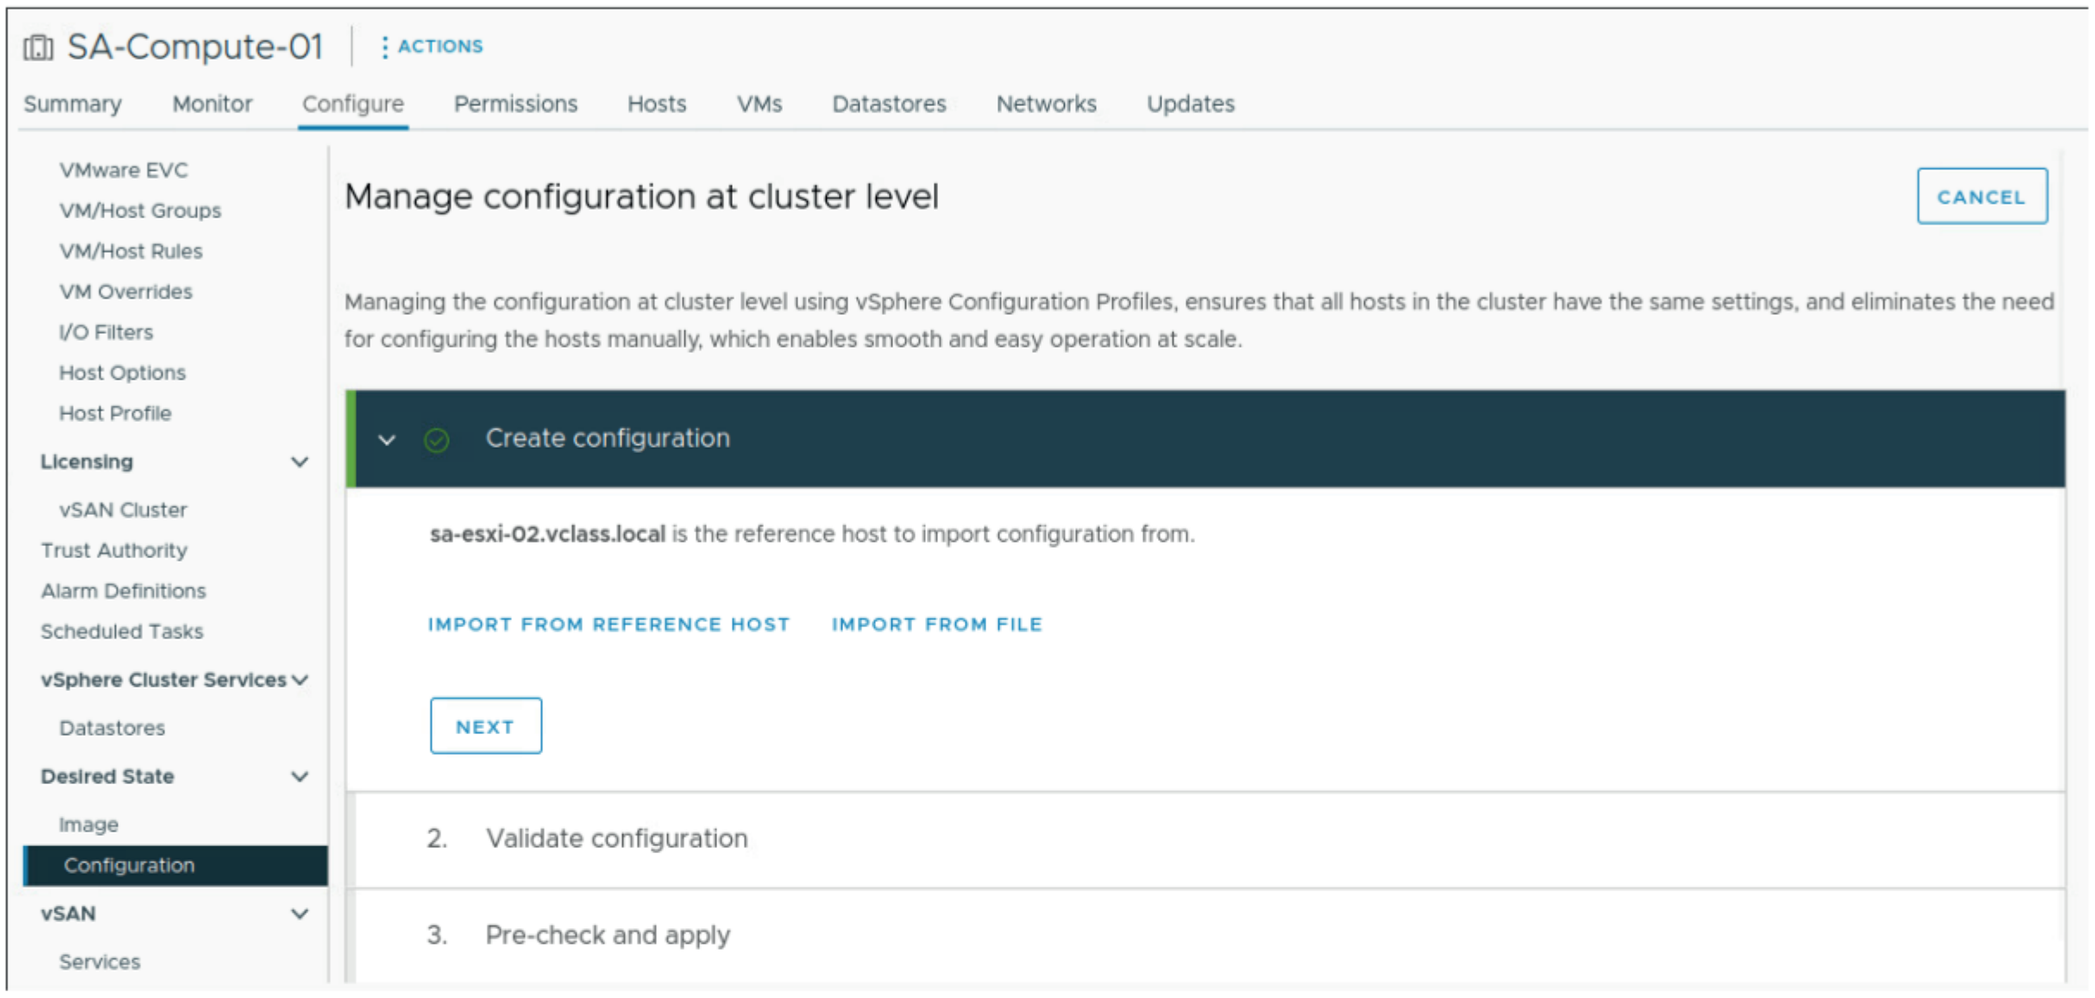

Reference - Using vSphere Configuration Profiles to Manage Host Configuration at a Cluster Level

6-33 View Host Settings in the Desired Configuration

You can view information about how each host setting is configured in the current desired configuration for the cluster.

For each host setting defined in the configuration document, you can see the values as defined for all hosts in the cluster and also any host overrides.

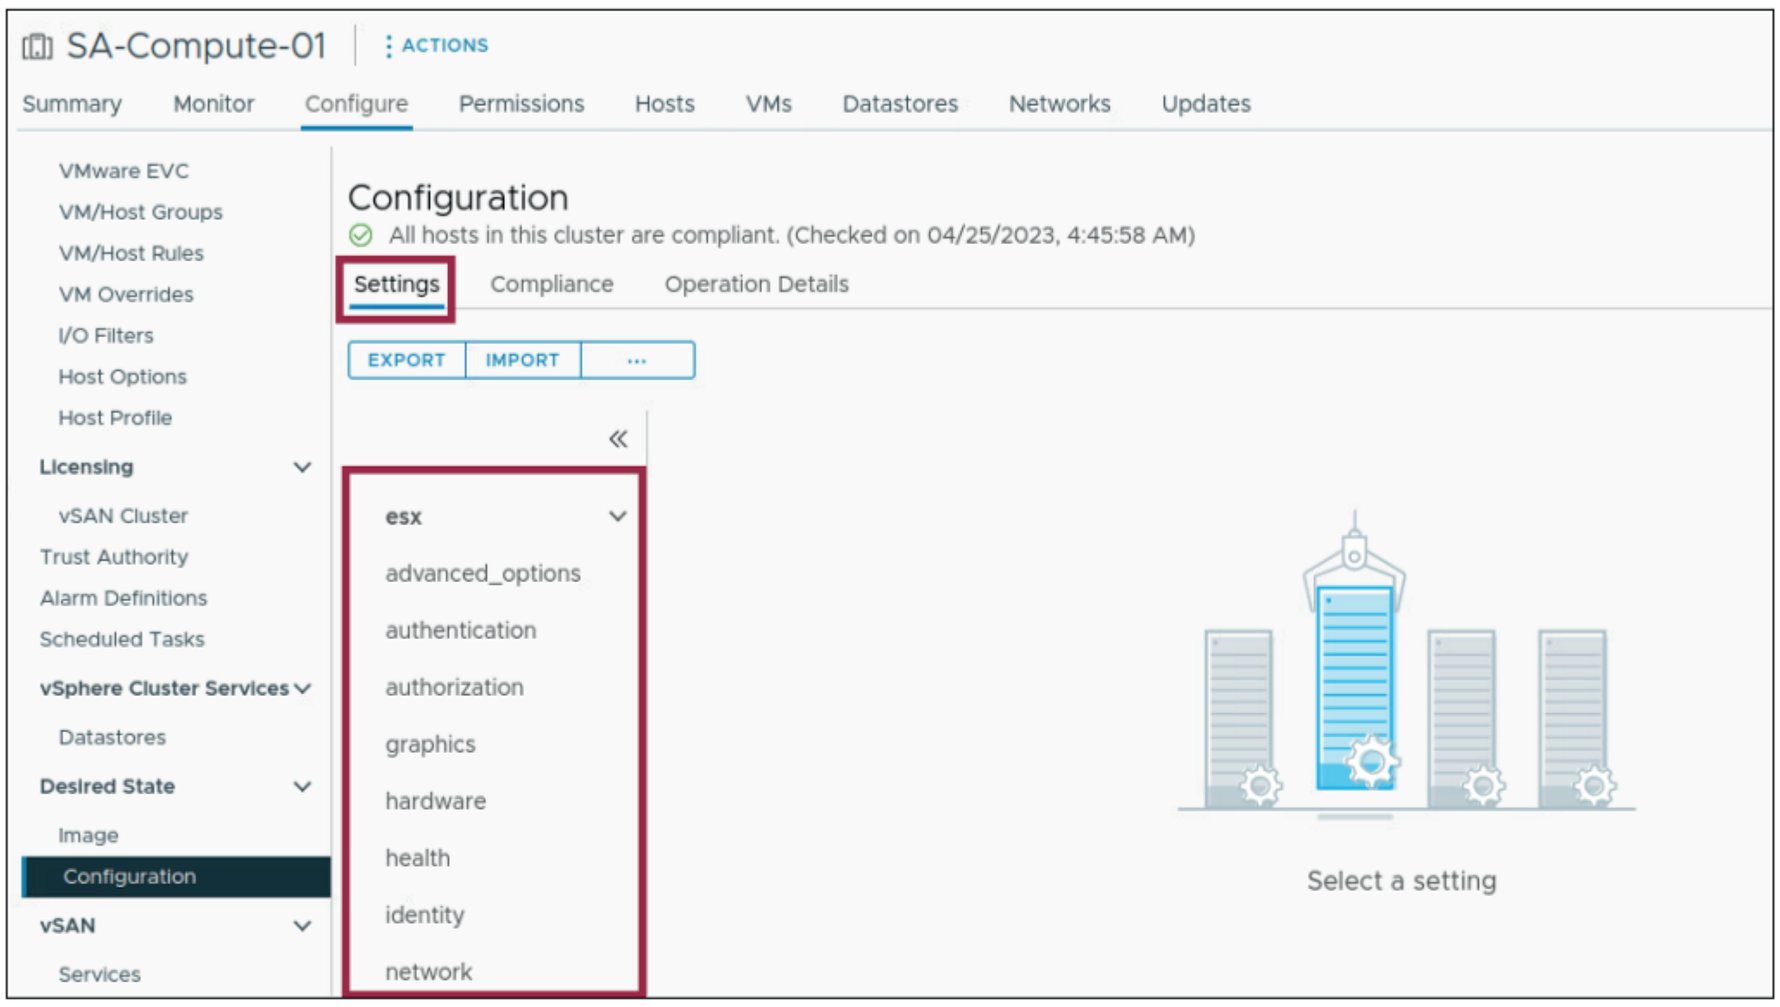

The common settings, the host-specific settings, and the host overrides are configured manually in the configuration document, which is an editable JSON file.

Technet24

6-34 Configuration Schema and the Configuration Document

Administrators can set up and manage the desired host configuration for an entire cluster by working with the configuration document for the respective cluster.

The configuration document is a JSON file that you can download to your local machine and edit with a JSON editor tool.

A valid configuration document contains a `profile` section and, optionally, a `host-specific` and `host-override` sections:

- The `profile` section contains a common configuration that is applicable to all hosts in the cluster.
- The `host-specific` section represents configuration that can only be specified per host.
- The `host-override` section represents configuration that is overridden for a specific host in the cluster.

6-35 Check Host Compliance Against the Desired Configuration for The Cluster

You can check the compliance of each host in a cluster against the desired configuration set for the entire cluster

The Check Cluster Configuration Compliance task runs automatically in the following events:

- When you make changes to the desired configuration of the cluster.
- When you add a host to the cluster.

6-36 Run a Remediation Pre-Check

You run a remediation pre-check to ensure that the desired configuration is valid and that the cluster is in good health.

During a remediation pre-check, vSphere Configuration Profiles runs various health checks on each host and the entire cluster. The remediation pre-check also computes the remediation impact on the hosts.

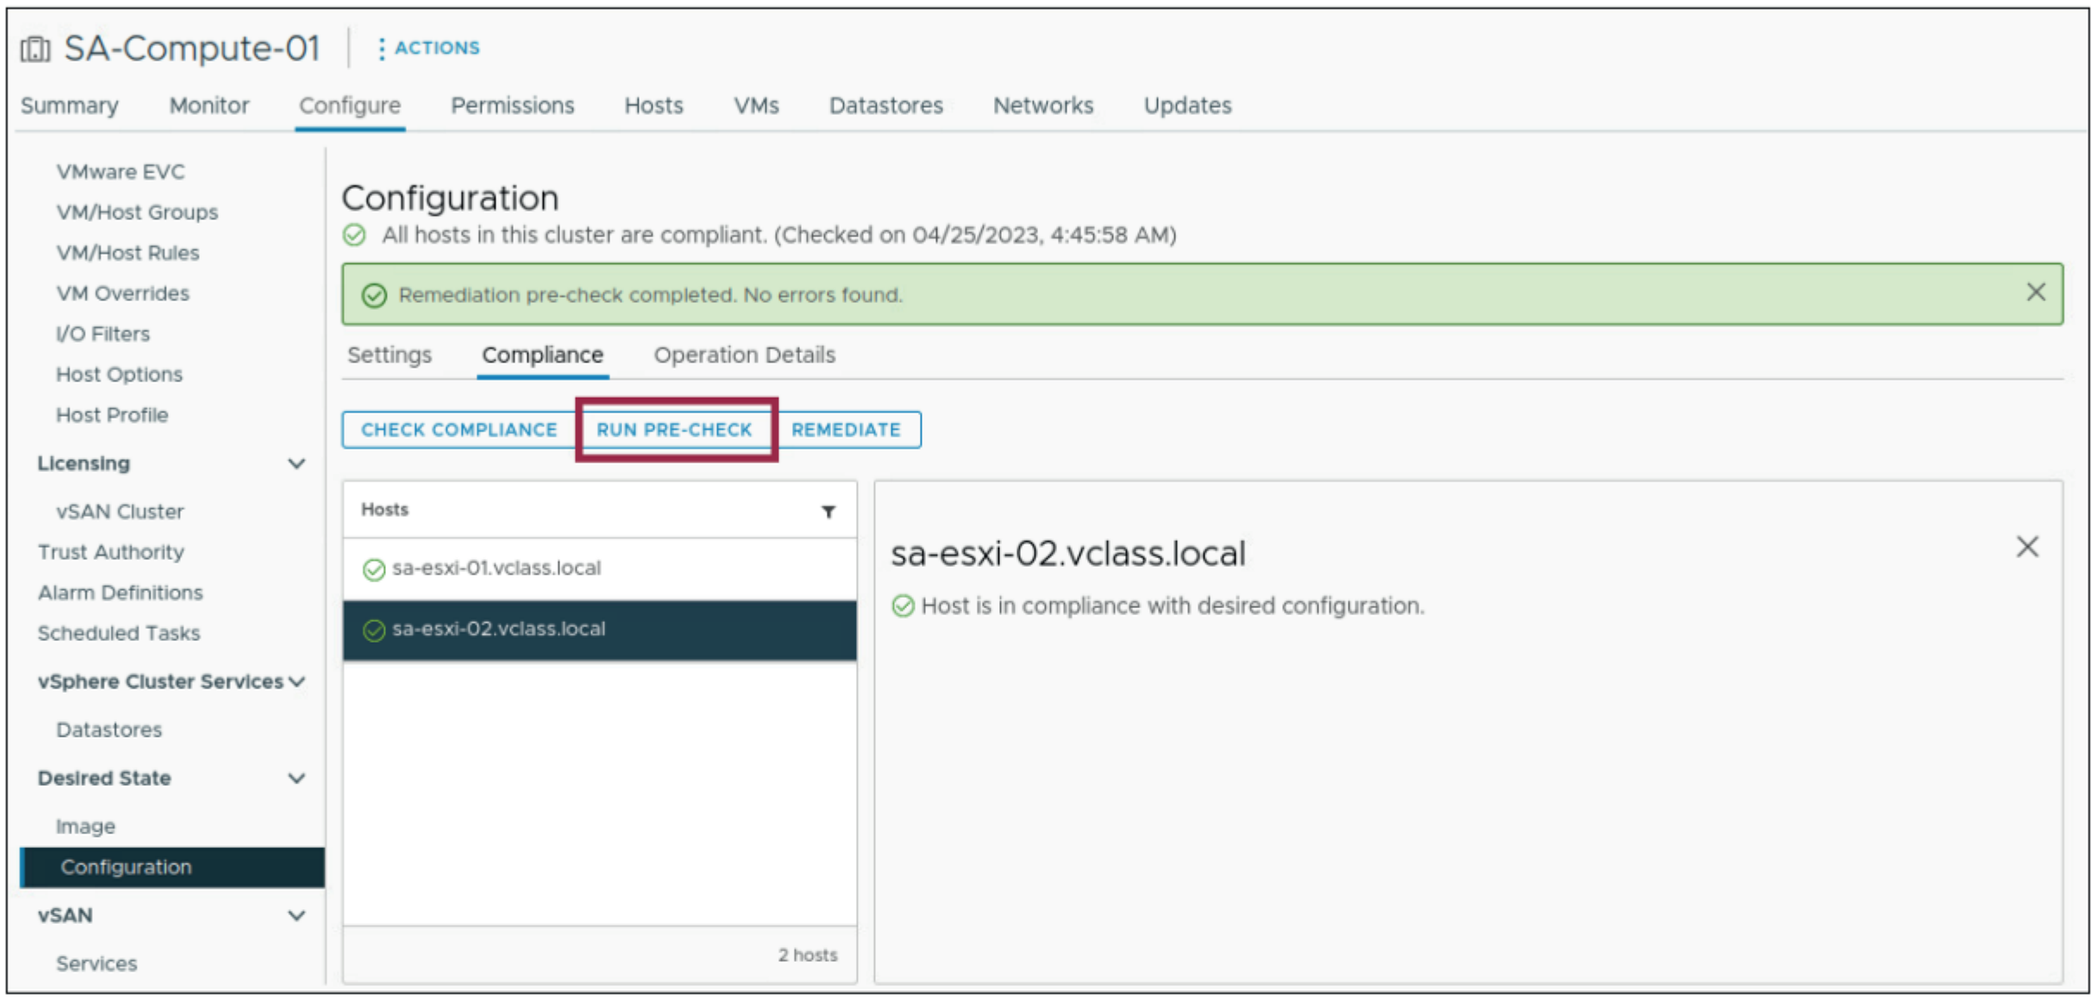

Technet24

6-37 Remediate a Cluster Against the Desired Configuration

You can remediate a cluster against the desired configuration to make all non-compliant hosts compliant with the configuration for the cluster.

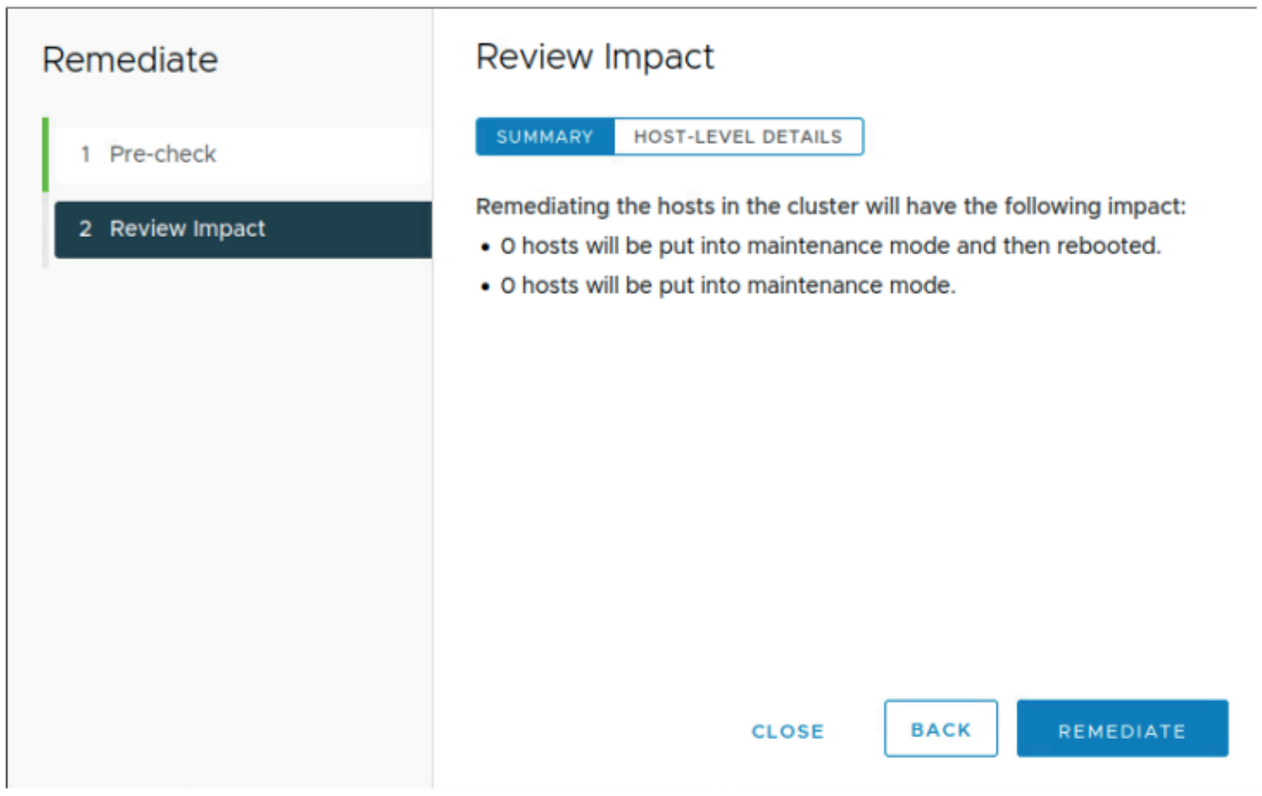

6-38 Consistent Configuration Across Multiple Sites

You can achieve consistent configuration across multiple vCenters in your environment.

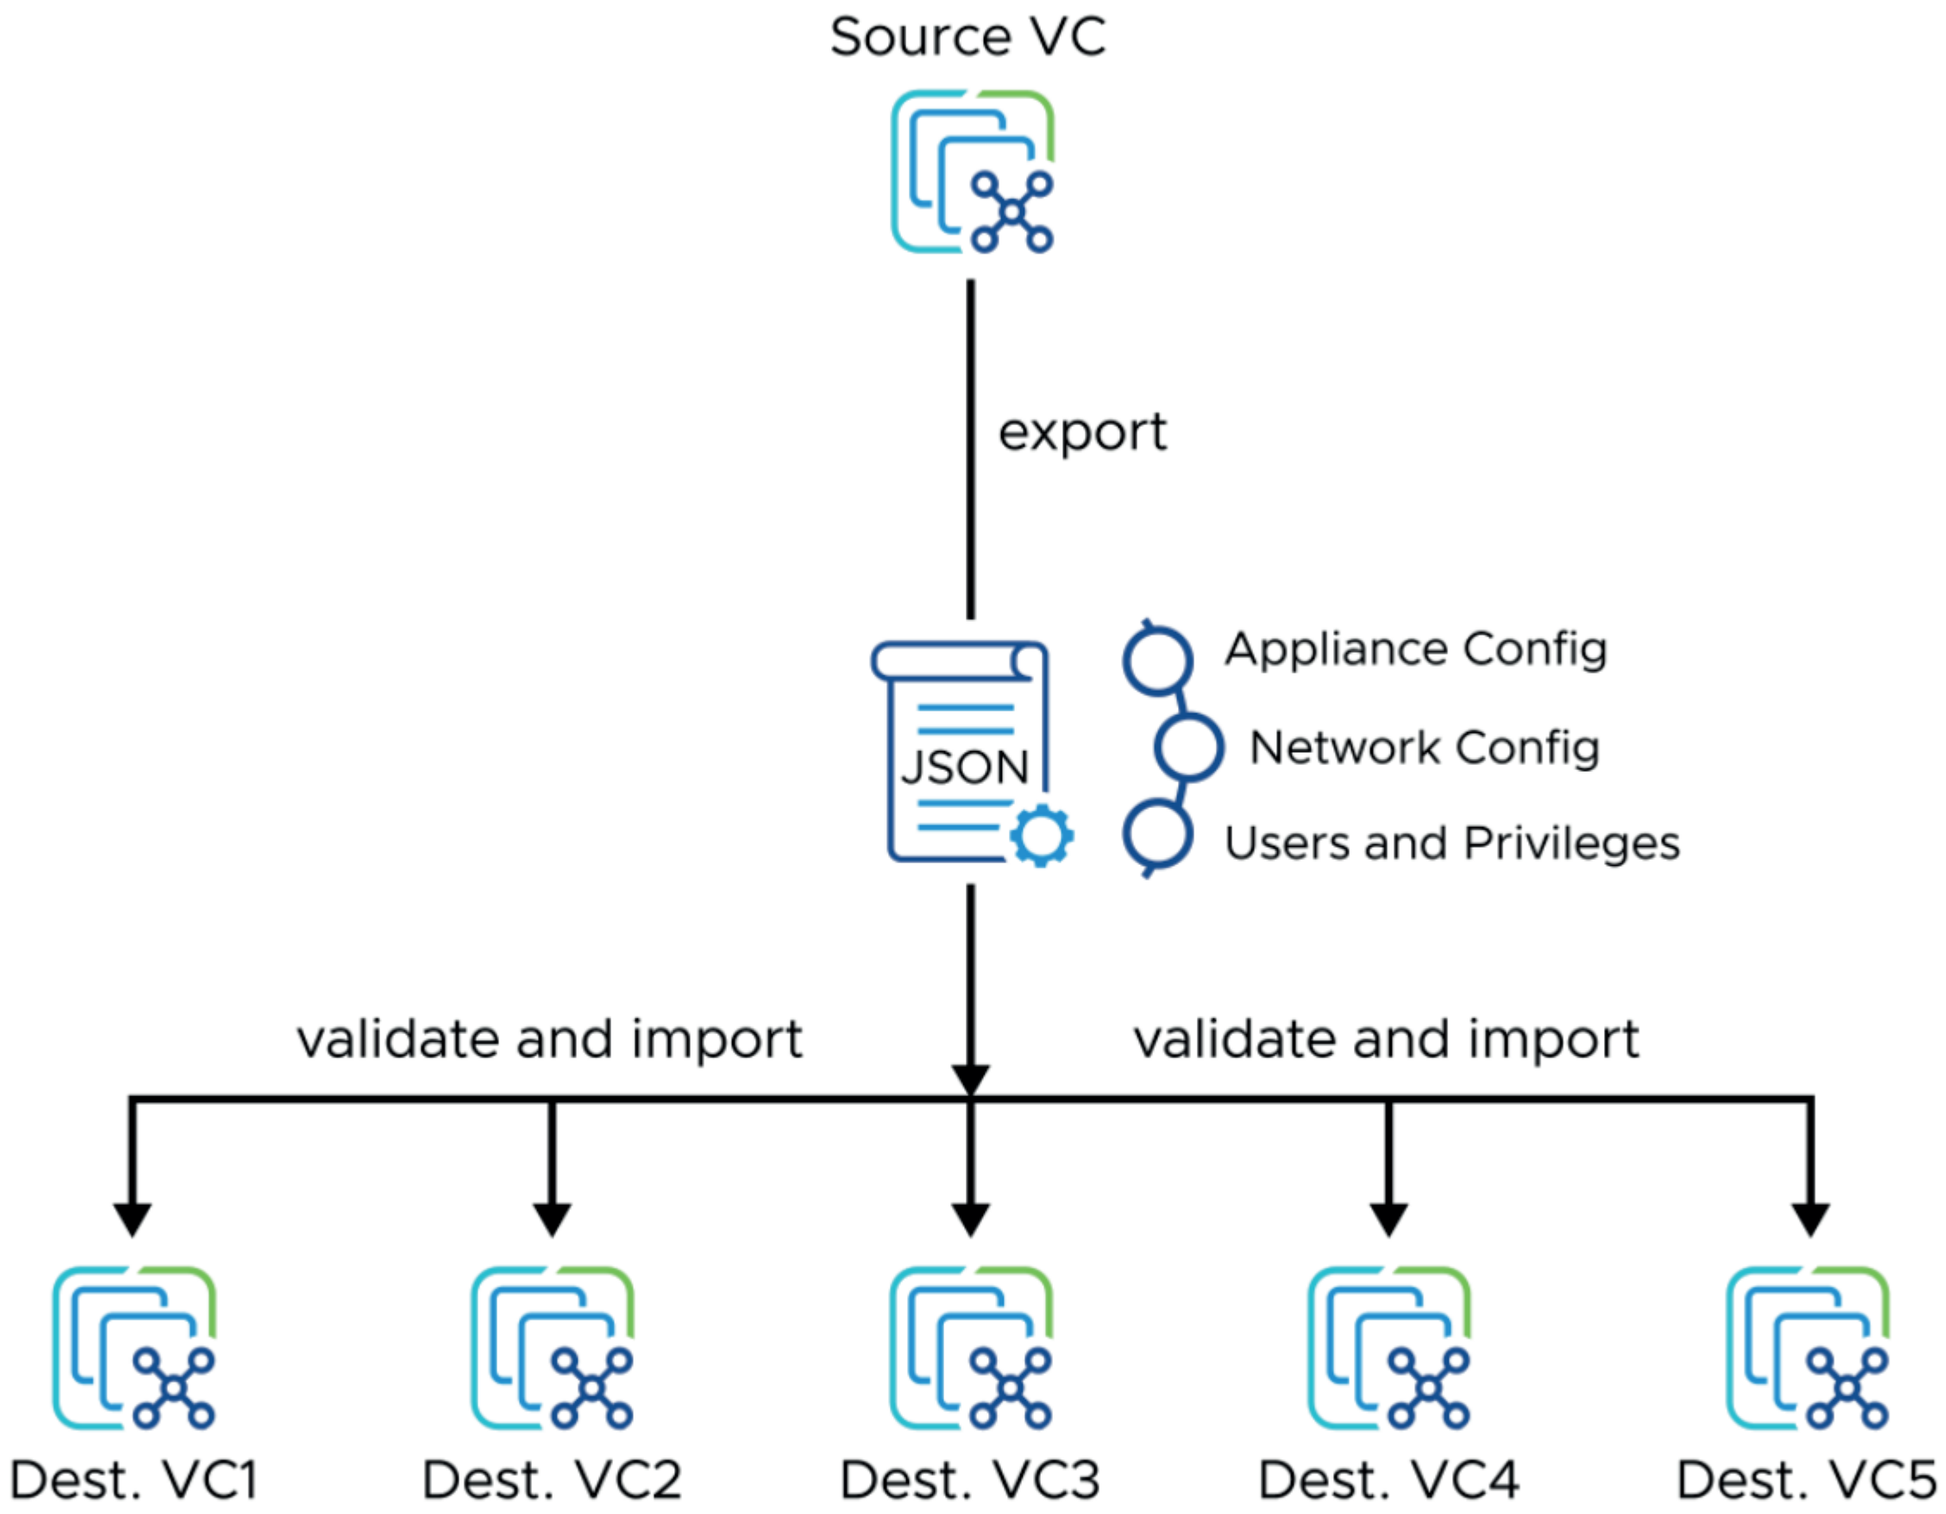

Technet24

6-39 vSphere Distributed Switch Support

vSphere Configuration Profiles supports a co-existence with vSphere Distributed Switch configurations.

This means that vSphere Distributed Switch and distributed port group information is still managed by vCenter and is compatible with using vSphere Configuration Profiles.

For example, if you create or modify distributed port groups, you do not need to add this information to the cluster configuration document.

The distributed port group is managed by the vSphere Distributed Switch independent of the cluster.

However, if you attach vmkernel interfaces to distributed port groups, then the vmkernel interface configuration needs to be managed by vSphere Configuration Profiles.

Reference - Announcing vSphere Configuration Profiles

6-40 Transitioning a Cluster with a vSphere Distributed Switch

An existing cluster configured to use Host Profiles can be converted to instead using vSphere Configuration Profiles.

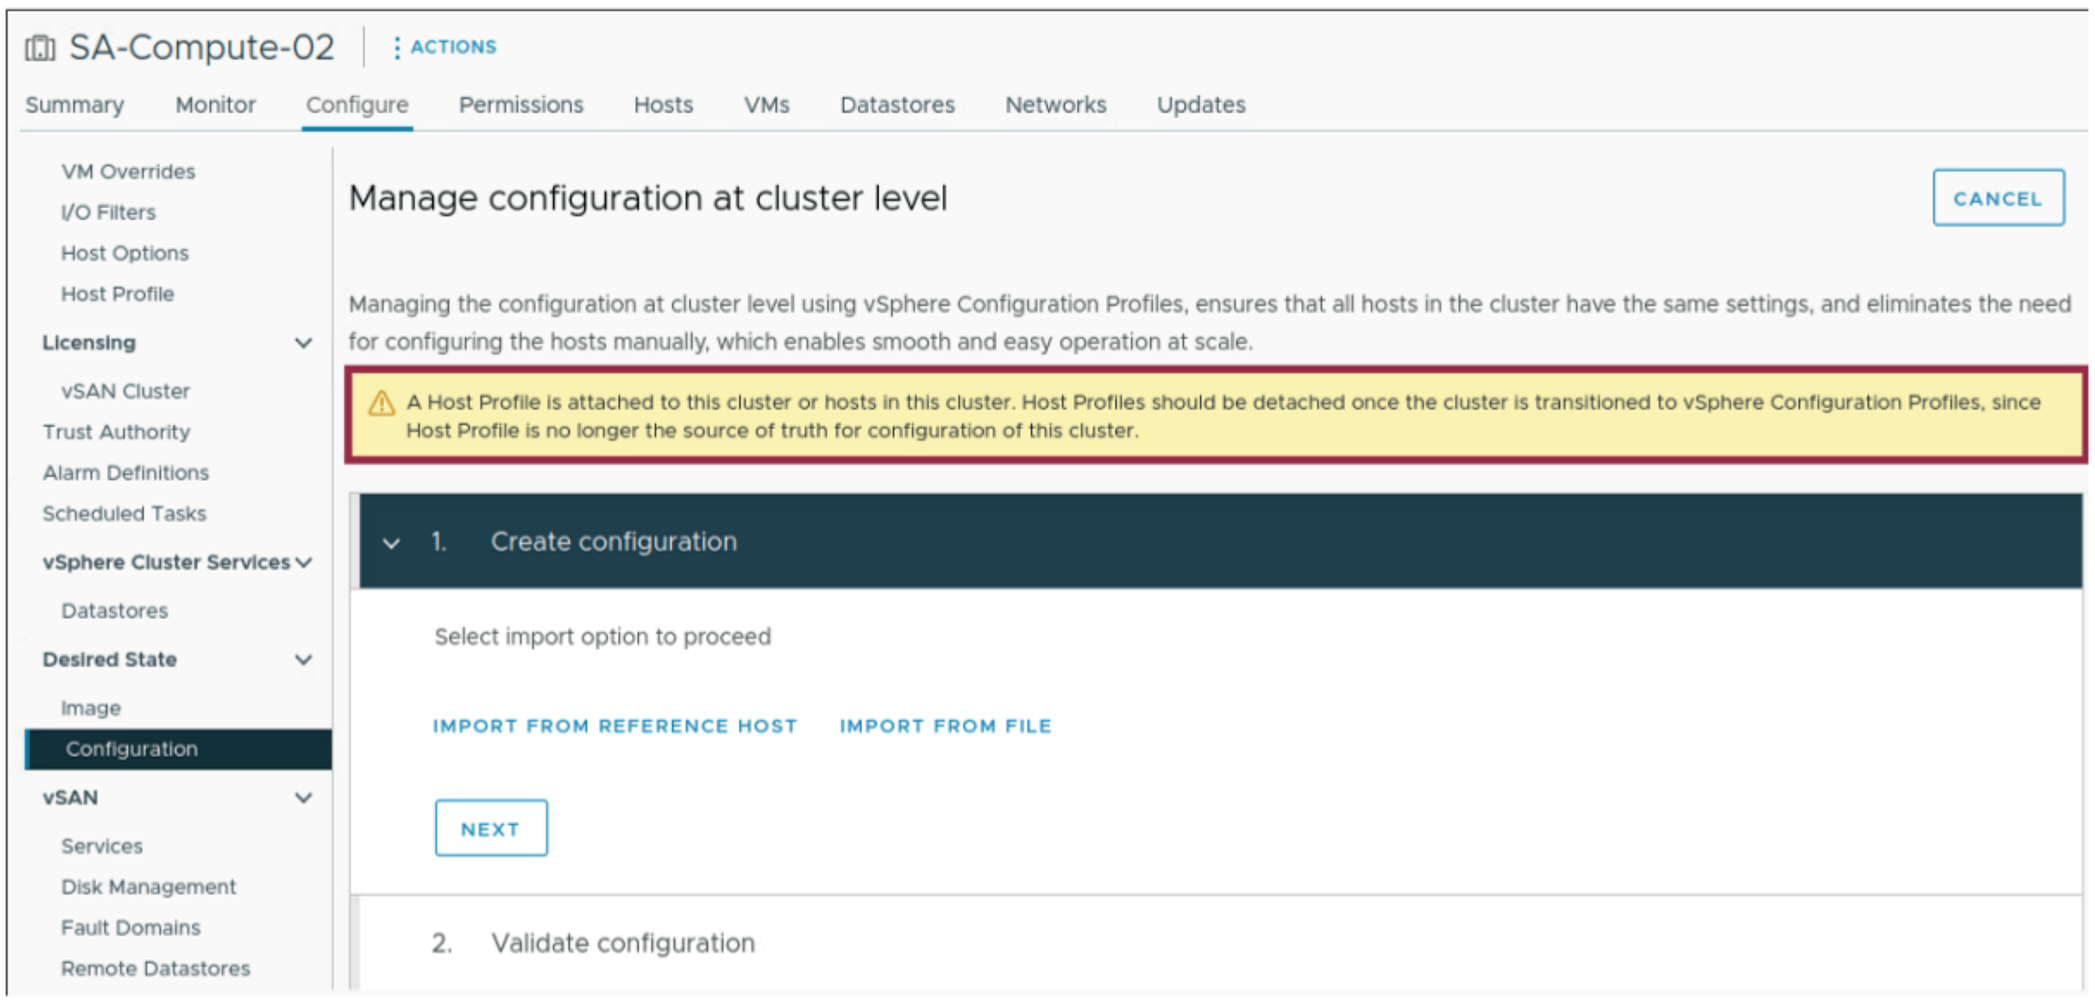

6-41 vSphere Configuration Profiles Service Health

vSphere Configuration Profiles is provided by the VMware vSphere Update Manager service.

Log messages are primarily written to the **vmware-vum-server-XX.log** file located at `/var/log/vmware/vmware-updatemgr/vum-server/` on the vCenter.

```
root@sa-vcsa-01 [ ~ ]# ls /var/log/vmware//vmware-updatemgr/vum-server/
2023_04_14.request.log            2023_04_22.request.log            vmware-vum-server-10.log.gz  vmware-vum-server-19.log.gz  vmware-vum-server-4.log.gz
2023_04_15.request.log            2023_04_23.request.log            vmware-vum-server-11.log.gz  vmware-vum-server-1.log.gz   vmware-vum-server-5.log.gz
2023_04_16.request.log            2023_04_24.request.log            vmware-vum-server-12.log.gz  vmware-vum-server-20.log.gz  vmware-vum-server-6.log.gz
2023_04_17.request.log            2023_04_24.request.log.160526036  vmware-vum-server-13.log.gz  vmware-vum-server-21.log.gz  vmware-vum-server-7.log.gz
2023_04_18.request.log            2023_04_25.request.log            vmware-vum-server-14.log.gz  vmware-vum-server-22.log.gz  vmware-vum-server-8.log.gz
2023_04_18.request.log.112143741  hostUpgrade                       vmware-vum-server-15.log.gz  vmware-vum-server-23.log.gz  vmware-vum-server-9.log.gz
2023_04_18.request.log.112721728  hsm-service.log                   vmware-vum-server-16.log.gz  vmware-vum-server-24.log     vmware-vum-server.log
2023_04_20.request.log            imageservice.log                  vmware-vum-server-17.log.gz  vmware-vum-server-2.log.gz   vum-log4j-filedownload.log
2023_04_21.request.log            imageservice.log.1.gz             vmware-vum-server-18.log.gz  vmware-vum-server-3.log.gz   vum-log4j-fileupload.log
root@sa-vcsa-01 [ ~ ]#
```


Technet24

6-42 About Host Profiles

Introduced back in vSphere 4.0, Host profiles are now the predecessor of vSphere Configuration Profiles.

Host Profiles are also used to automate and centrally managed mechanism for host configuration and configuration compliance.

Similarly to vSphere Configuration Profiles, Host profiles have many benefits:

- Simplified setup and change management for ESXi hosts
- Easy detection of noncompliance with standard configurations
- Automated remediation

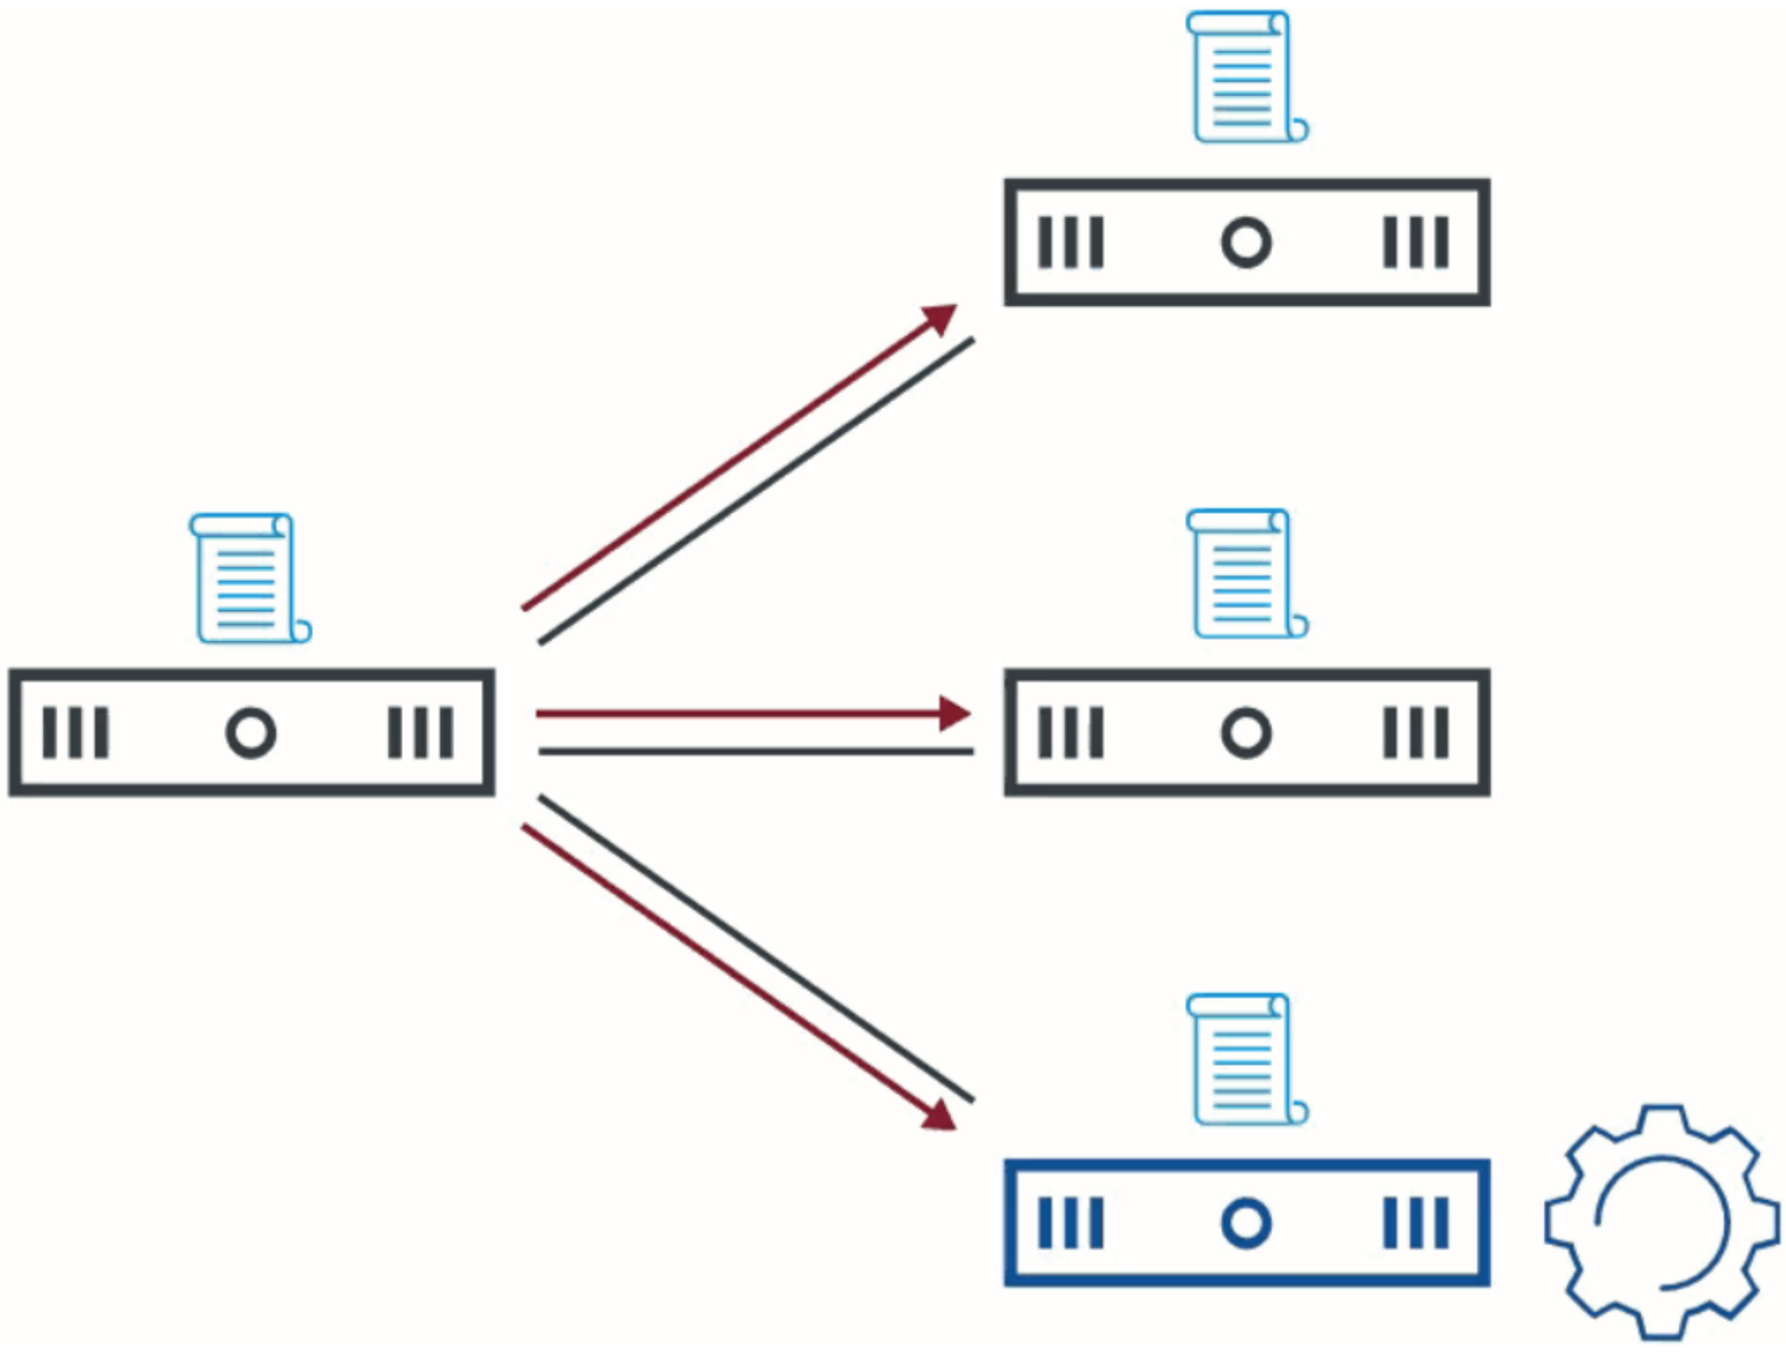

Over time, your vSphere environment expands and changes. New hosts are added to the data center and must be configured to coexist with other hosts in the cluster, for example:

- Storage: Datastores such as VMFS and NFS, iSCSI initiators and targets, and multipathing
- Networking: Virtual switches, port groups, physical NIC speed, security, NIC teaming policies, and so on
- Licensing: Edition selection, license key input, and so on

- Date and time configuration: Network Time Protocol (NTP) server
- Security settings: Firewall settings, Active Directory configuration, ESXi services, and so on
- DNS and routing: DNS server, IP default gateway, and so on

Traditionally you would modify the configuration of a single host:

- Manually, using the vSphere Client or the command prompt
- Automatically, using scripts

Using the Host Profiles feature, you can export configuration settings from a primary reference host and save them as a portable set of policies, called a host profile. You can use this profile to quickly configure other hosts in the data center. Configuring hosts with this method reduces the setup time of new hosts. Multiple steps are reduced to a single click. Host profiles also eliminate the need for specialized scripts to configure hosts.

The vCenter system uses host profiles as host configuration baselines so that you can monitor changes to the host configuration, detect discrepancies, and fix them.

As a licensed feature of vSphere, Host Profiles are available only with Enterprise Plus licensing. If you see errors, ensure that you have the appropriate vSphere licensing for your hosts.

Technet24

6-43 Implementing Host Profiles

You start with a reference host to create a host profile that can be applied to new hosts.

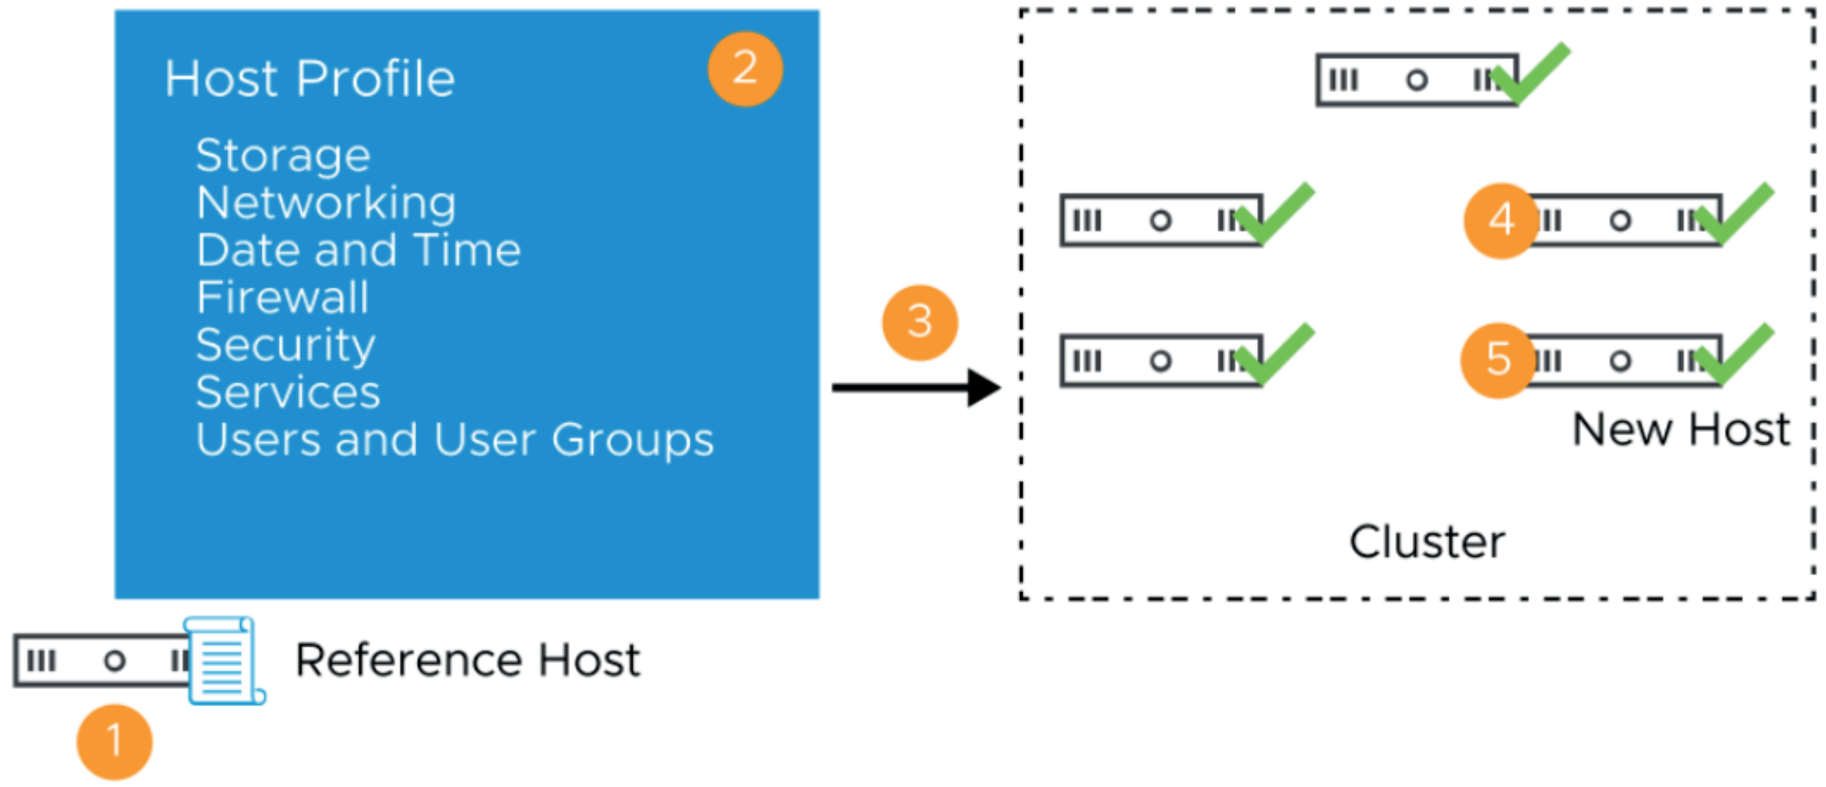

Follow this basic workflow to implement host profiles:

1. Set up and configure a host to be used as the reference host. A reference host is the host from which the profile is created.
2. Use the Create Profile wizard to create a profile from the designated reference host. vCenter retrieves the configuration settings of the reference host, including NFS mounts, and vSphere vMotion ports. You can edit the profile to further customize the policies in the profile.
3. Attach the host profile to a host or a cluster of hosts. You can attach a profile to standalone hosts, or to one or more clusters of hosts.
4. Check host compliance to the reference host's profile. If all hosts are compliant with the host profile, the hosts are correctly configured.
5. Check new hosts for compliance against the host's profile. You can easily apply the host profile to other hosts or clusters of hosts that are not compliant.

You can also import and export a host profile to a profile file to a host profile that is in the VMware profile format (.vpf). This feature is useful for sharing host profiles across multiple vCenter instances.

6-44 Monitoring for Compliance

When a new host is added to a cluster, the host is checked for compliance against the host profile that was applied. A cluster can also be manually checked for compliance against a host profile.

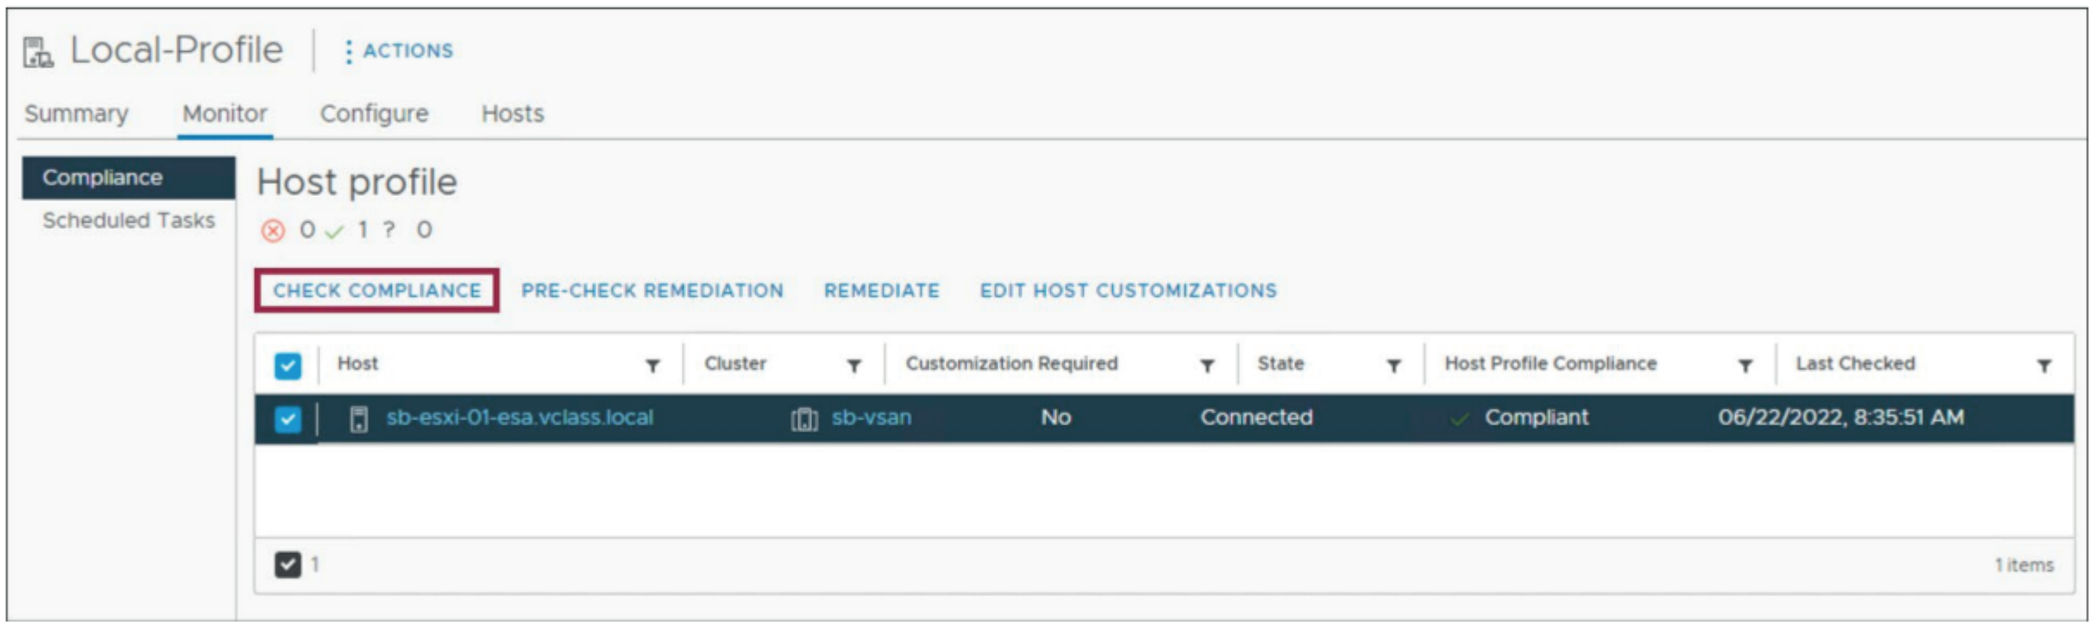

After the host profile is created and associated with a set of hosts or clusters, you can verify the compliance status in various places in the vSphere Client:

- Host profile's main view: Displays the compliance status of hosts.
- Host **Summary** tab: Displays the compliance status of the selected host.
- Host profile's **Monitor** tab: After you check compliance, this view displays the compliance status of the selected hosts. This view also displays any inconsistencies between the host values and the host profile values (shown on the slide).

You can also check compliance for individual components or settings. These built-in cluster compliance checks are useful for vSphere DRS and vSphere DPM.

Technet24

6-45 Standardization Across Multiple vCenter Instances

Organizations that span multiple vCenter instances can export the host profile from one vCenter instance and import the host profile for use in other vCenter instances.

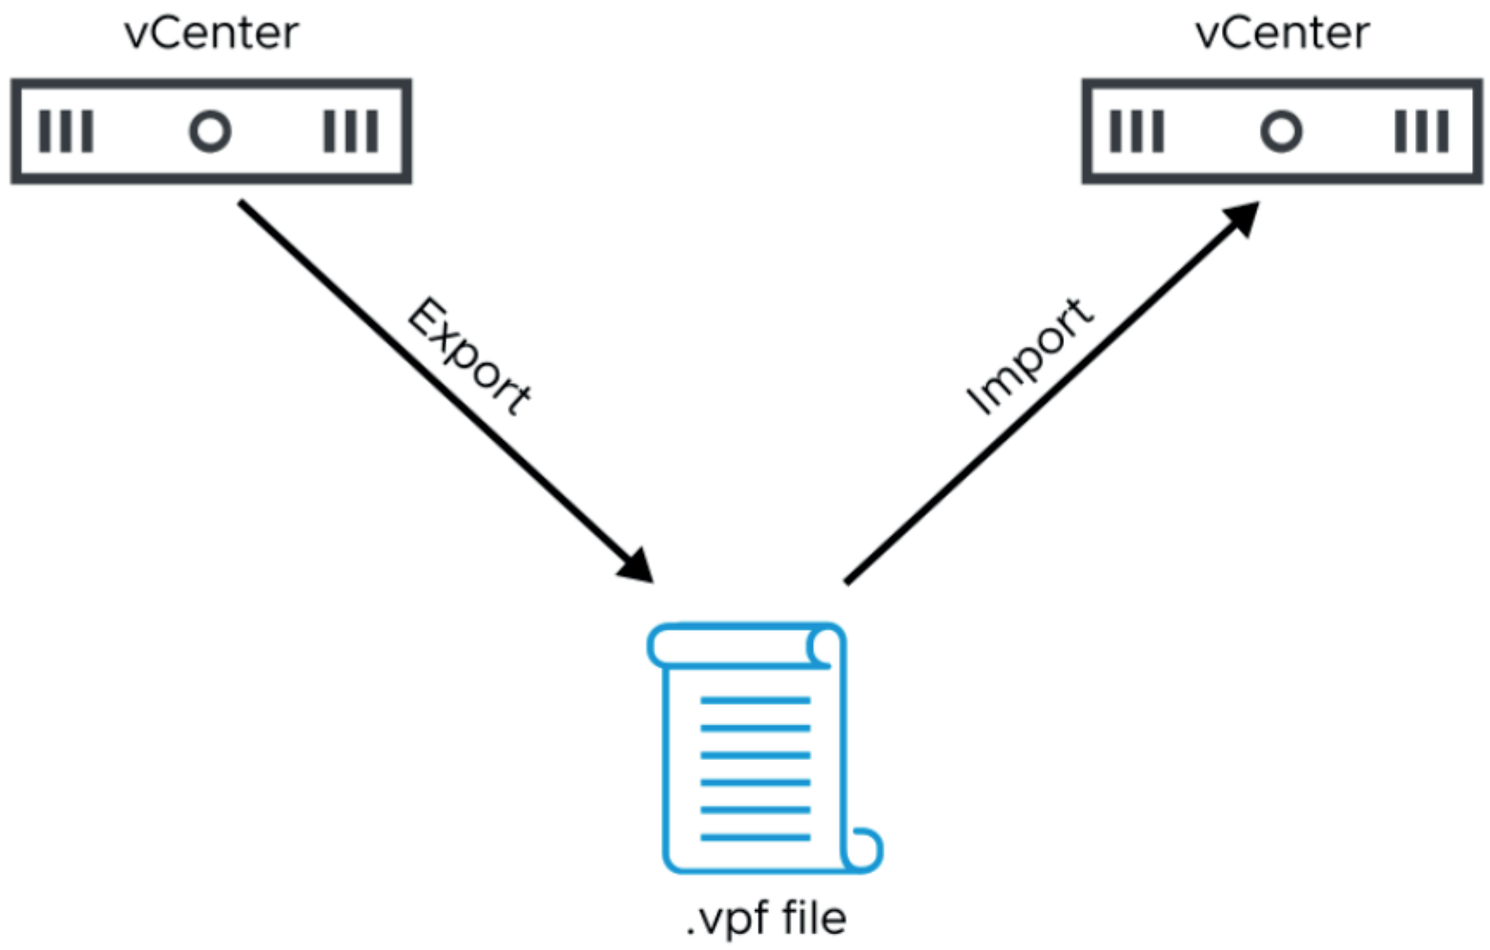

Some organizations might have large environments that span many vCenter instances. If you want to standardize a single host configuration across multiple vCenter environments, you can export the host profile from one vCenter instance so that other vCenter instances can import the profile for use in their environments. Host profiles are not replicated, shared, or kept synchronized across vCenter instances connected by Enhanced Linked Mode.

Files are exported in the VMware profile format (.vpf).

NOTE

When a host profile is exported, administrator and user profile passwords are not exported. This restriction is a security measure that stops passwords from being exported in plain text. After the profile is imported, you are prompted to reenter values for the password, and the password is applied to a host.

6-46 Lab 14: Using vSphere Configuration Profiles

Transition a cluster from using baselines to using images and use vSphere configuration profiles to manage all the hosts in the cluster:

1. Configure a Cluster with a Single Image
2. Configure a Cluster with vSphere Configuration Profiles
3. Remediate the Hosts in a Cluster
4. View the Configuration Document

6-47 Review of Learner Objectives

- Recognize benefits of using vSphere Configuration Profiles
- Explain how the Host Profiles feature works
- Configure and manage vSphere Configuration Profiles

Technet24

6-48 Lesson 4: VMware Certificate Management

6-49 Learner Objectives

- Use VMware Certificate Authority and VMware Endpoint Certificate Store to configure vSphere security certificate management
- Describe solution users and their certificate requirements
- Describe common VMware CA modes
- Replace VMware CA certificates with enterprise or third-party commercial certificates
- Use the vSphere Client to manage vSphere certificates

6-50 Certificate Use Case

Your browser does not recognize or trust the VMware CA root certificate.

To eliminate certificate warning messages from appearing every time you connect to vCenter, you must add the VMware CA root certificate to your browser's list of trusted certificates.

Refer to vendor documentation for your browser.

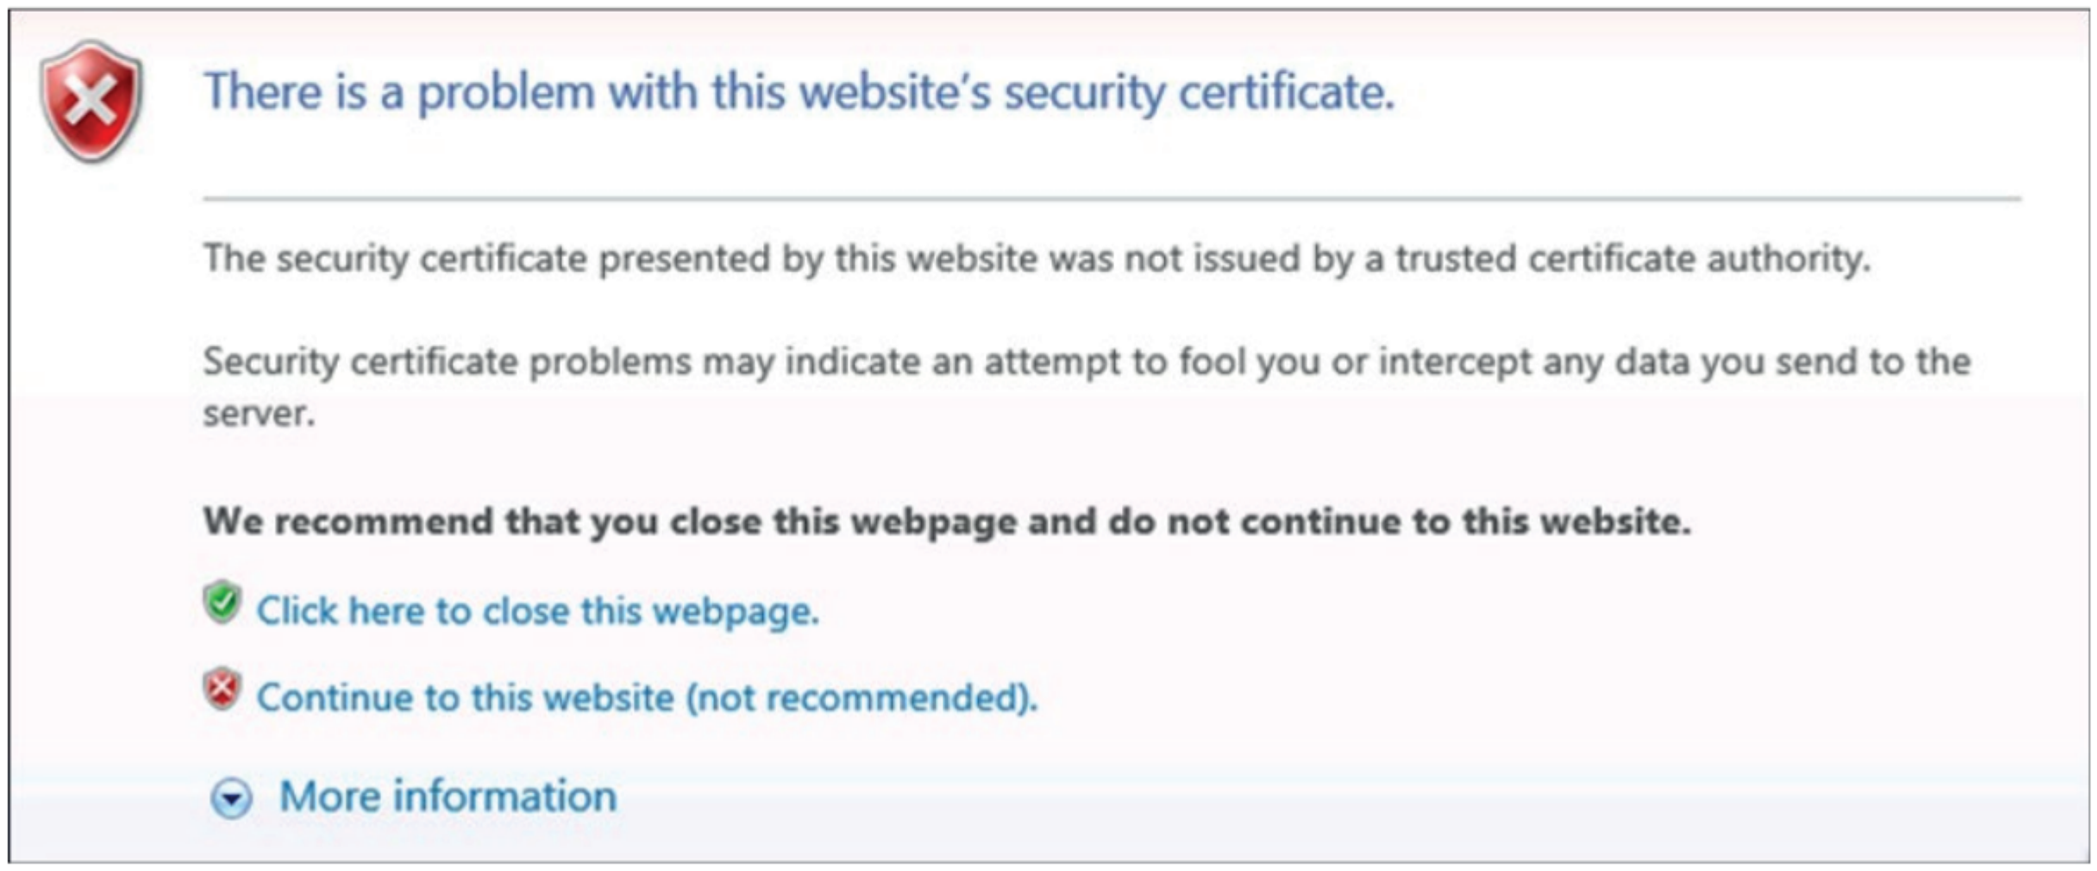

6-51 About Certificate Authorities

Certificate authorities (CAs) issue public-key certificates:

- Commercial CAs charge to issue certificates.
- Community CAs, such as CAcert.org, provide free certificates.

Some organizations deploy their own CA and issue self-signed certificates.

To trust a certificate signed by a CA, clients such as web browsers must have the CA's root certificate installed:

- The root certificates of large commercial CAs are preinstalled in modern browsers.

Public key cryptographic systems use pairs of keys. The data is encrypted using a private key, which is known only to the owner. The data can be decrypted only by using a public key, which might be widely disseminated.

CAs are important in public key infrastructure systems. When a Secure Sockets Layer (SSL) or Transport Layer Security (TLS) client connects to a server, the server sends its public key to the client to authenticate the server. However, the server cannot send its public key in plain text because such text is easily corruptible by attackers who can inject themselves into the communication.

Instead, the server sends an X.509 certificate to the client. The X.509 certificate contains the server's name (the subject), its public key, and other information. X.509 certificates are signed by a trusted CA, which means that they are encrypted with the CA's private key.

The client trusts the CA because the client already has the CA's public key, which is either preinstalled or manually installed by an administrator. For example, browsers such as Safari, Firefox, Internet Explorer, and Chrome have public keys of the most common CAs preinstalled.

SSL or TLS use is not limited to websites and browsers, although these uses are common and well known. Any client or server software solution that requires secure authentication of the parties can use SSL or TLS in the manner described here.

Technet24

6-52 About VMware CA

VMware CA issues and validates certificates.

The main purpose of VMware CA is to issue certificates so that vSphere components can establish trust relationships.

The main consumers of certificates from VMware CA are services that run as part of vCenter.

Services request certificates to establish their own identity and to establish the identity of other vSphere entities.

By default, VMware CA acts as a root CA. It issues and validates certificates for vSphere components such as ESXi hosts and solutions. VMware CA can handle all certificate management for you. VMware CA provisions vCenter components and ESXi hosts with certificates that use VMware CA as the root certificate authority.

6-53 vSphere Certificate Types

vSphere uses several types of certificates:

- VMCA Root certificate
- Machine certificates:
  - vCenter
  - ESXi
- Solution user certificates

vSphere components use SSL to communicate securely with one another and with ESXi hosts. SSL communications ensure data confidentiality and integrity. Data is protected and cannot be modified in transit without detection.

vCenter services, such as the vSphere Client, also use certificates for their initial authentication to vCenter Single Sign-On. vCenter Single Sign-On provisions each component with a SAML token that the component uses for authentication from that point on.

VMware products use standard X.509 version 3 (X.509v3) certificates to encrypt session information that is sent over SSL between components.

For more information about vSphere certificate requirements, see *Where vSphere Uses Certificates* at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-authentication/GUID-3AF7757E-A30E-4EEC-8A41-28DA72102520.html.

Technet24

6-54 Chain of Trust (1)

If you trust the CA, you implicitly trust all the certificates issued by that CA.

- To trust a certificate, you must trust the chain of trust. One of the following conditions must be true:
  - You must say that you explicitly trust the certificate itself.
  - You must say that you explicitly trust the CA that issued it.
- In a self-signed certificate, the issuer and the user are the same system:
  - To use self-signed certificates, every user (client) system must install and explicitly trust every self-signed certificate that is used in the network.
  - Every time a new service is brought online, all clients must individually install and trust every self-signed certificate in the network.
- An in-house or commercial CA eliminates the requirement that each client system installs every self-signed certificate as long as:
  - All client systems trust the CA.
  - All certificates come from that trusted CA.

6-55 Chain of Trust (2)

This configuration is known as a two-node chain of trust.

The two nodes are:

- The client system
- The CA

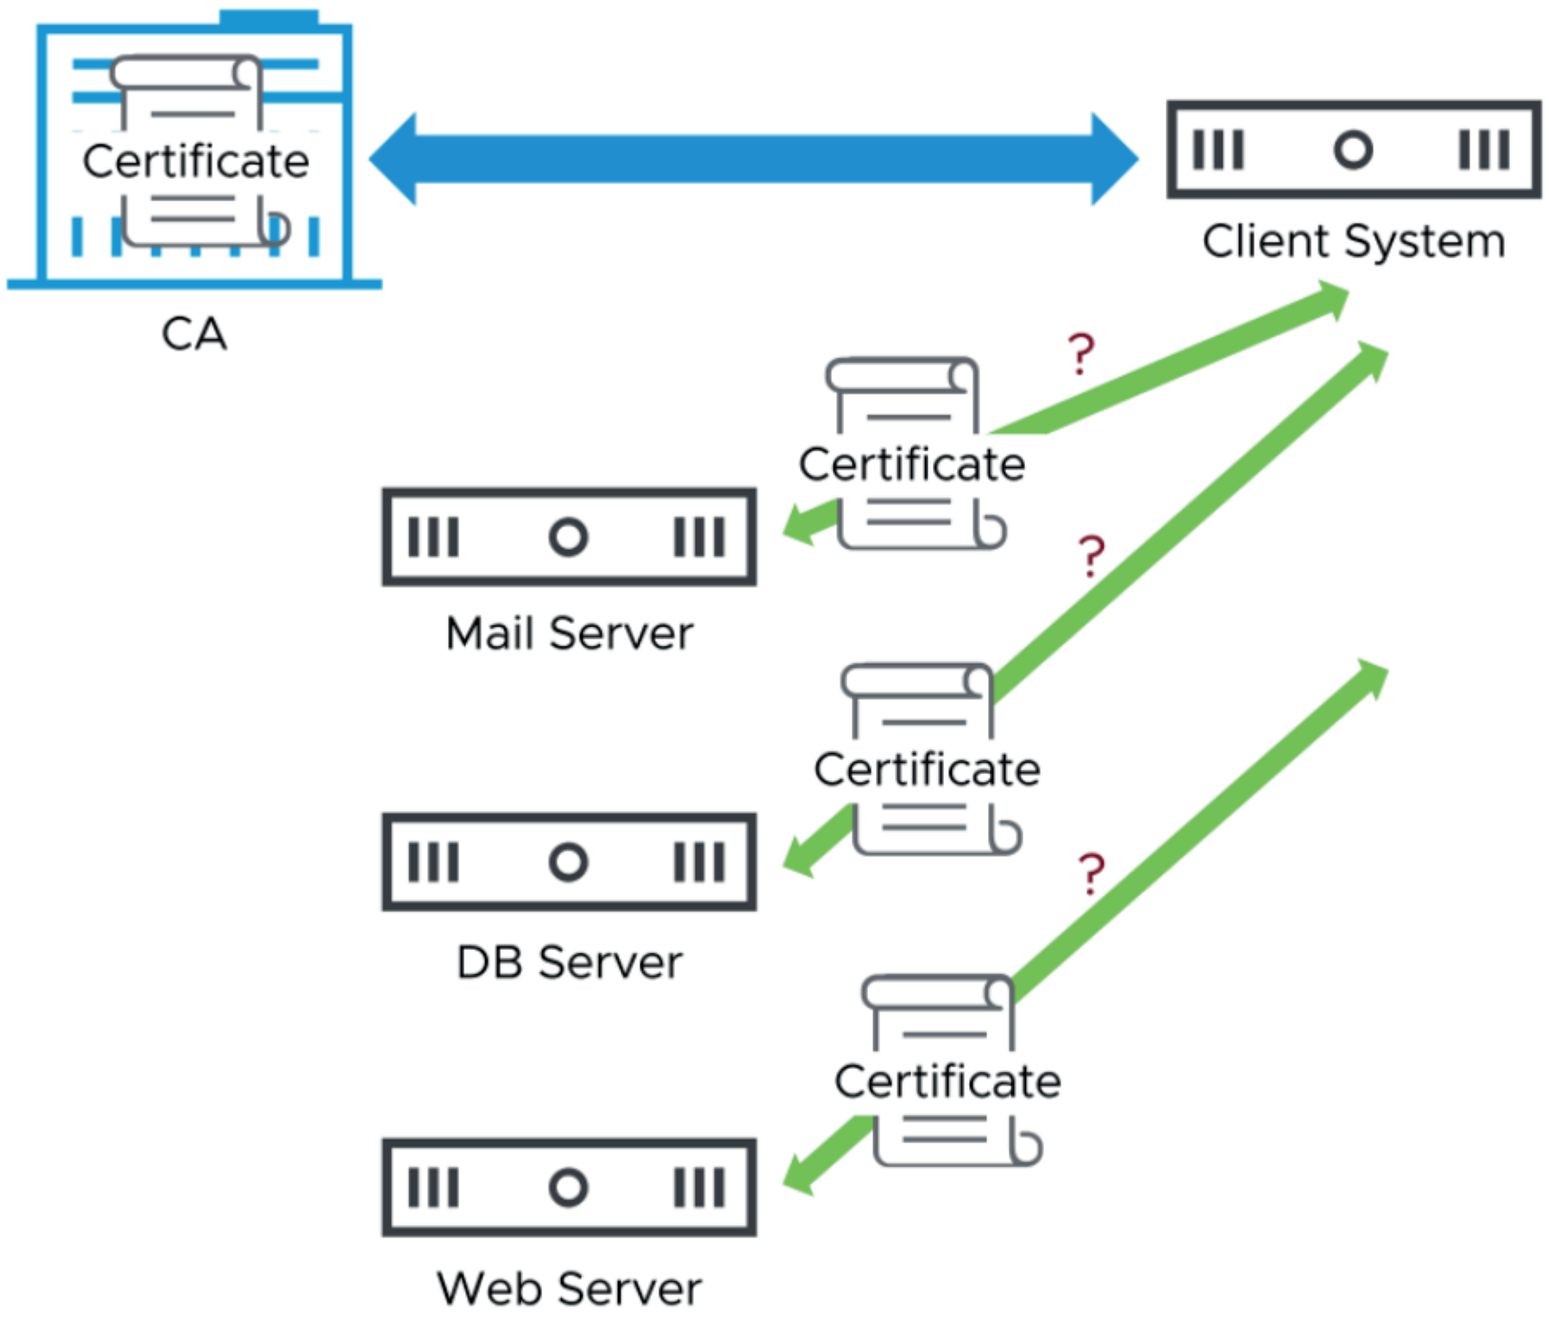

The client trusts the server because the Certificate Authority (CA) issued the server its certificate and the client trusts the CA.

Technet24

6-56 Chain of Trust (3)

This configuration is known as a three-node chain of trust.

The three nodes are:

- The client system
- The ICA
- The CA

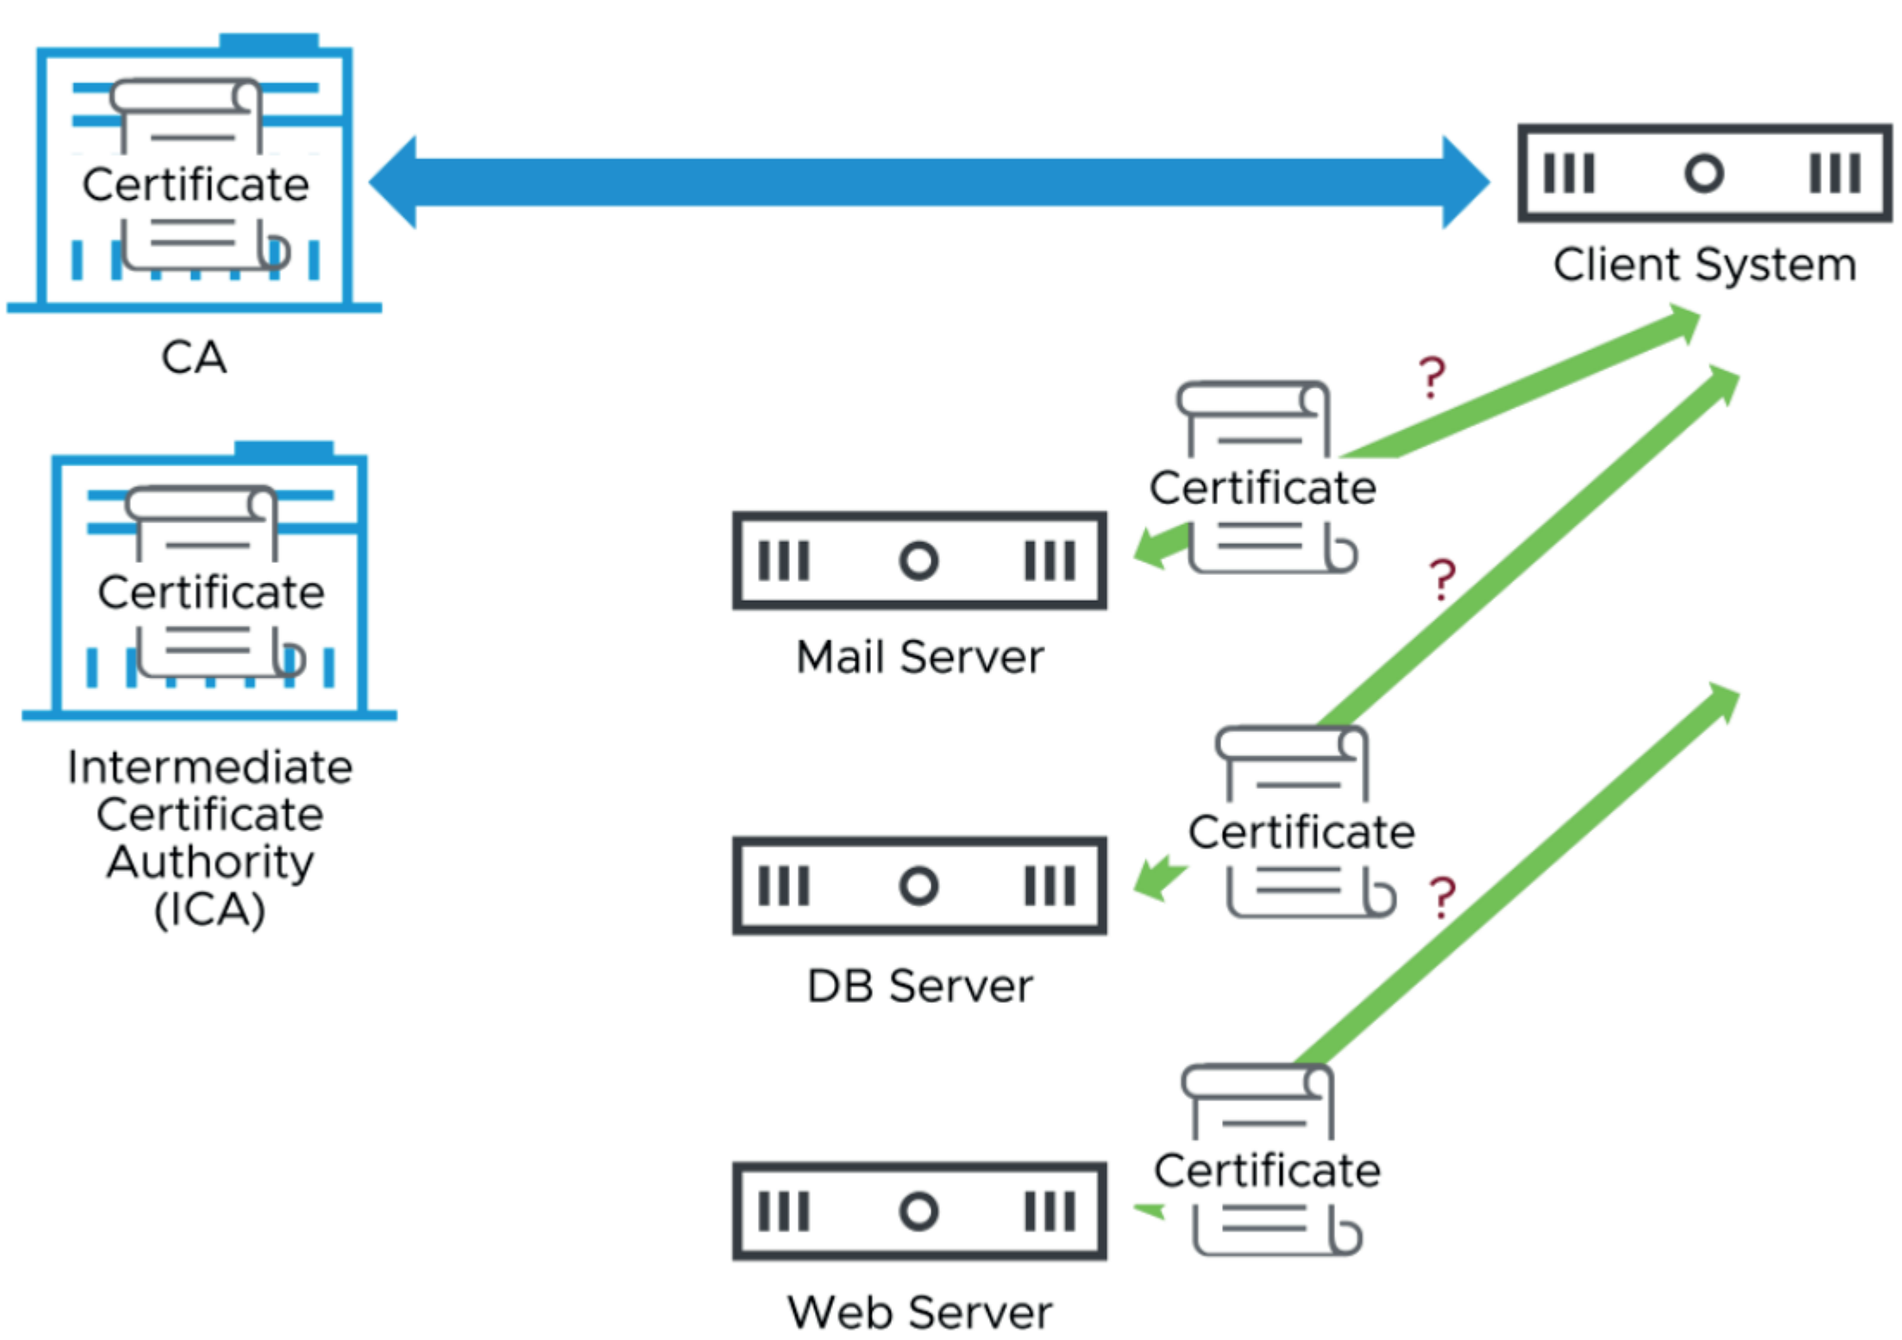

The client trusts the server because the Intermediate Certificate Authority (ICA) issued the server its certificate, the CA issued the ICA its certificate, and the client trusts the CA.

6-57 VMware CA Modes

VMware CA can be configured to run in several different modes:

- Root CA (default): VMware CA acts as a root CA. VMware CA uses its own self-signed root certificate to create and issue certificates to vCenter, ESXi, and other components that require them.
- Subordinate CA: VMware CA is an ICA to an existing enterprise or external third-party CA for all vSphere certificates.
- Hybrid: VMware CA is subordinate to an existing enterprise or external third-party CA for only the VMCA root certificate and vCenter machine certificate.
- Custom: All certificates issued by VMware CA are replaced with custom certificates from a corporate or commercial CA. This option bypasses VMware CA.

By default, VMware CA is configured to act as a root CA that can use its self-signed root certificate to issue and sign other certificates. However, you can also configure VMware CA as a subordinate CA. In subordinate CA mode, VMware CA is an intermediate CA in the chain of trust that continues up to an enterprise CA or an external third-party CA.

If you do not want to use VMware CA, you can bypass VMware CA and install custom certificates on all your systems that were issued by an external or enterprise CA .

Use subordinate mode or bypass mode if you want to use trusted certificates. All your custom certificates (except for host certificates) must be installed in VECS.

Technet24

6-58 VMware CA: Root CA Mode

Root CA mode for VMware CA is the default mode:

- This mode uses a self-signed root certificate.
- The root certificate is a CA certificate that can issue other certificates.
- All vSphere certificates are created from this root CA certificate.

Normally, a CA has a self-signed certificate at the root, especially if the certificate is the first one created in a new domain.

VMware CA uses a self-signed root certificate. It issues certificates to vCenter, ESXi, and so on, and manages these certificates. These certificates have a chain of trust that stops at the VMware CA root certificate. VMware CA is not a general-purpose CA, and its use is limited to VMware components.

6-59 Replacing the VMware CA Certificate in Root CA Mode (1)

To replace the VMware CA certificate in root CA mode, follow these high-level steps.

You are not required to replace the certificates that were issued and signed by the previous VMware CA root CA certificate. When you replace the VMware CA root CA certificate, it does not automatically invalidate the previous certificate. It also does not remove the old certificate, so the old certificate remains valid in your vSphere environment.

However, when you want to replace the root CA certificate, you typically also want to replace all your existing certificates (vCenter, ESXi hosts etc) , so that they are signed by the current VMware CA root CA certificate. You can replace all these certificates at once by using the vSphere Certificate Manager tool.

6-60 Replacing the VMware CA Certificate in Root CA Mode (2)

Use the vSphere Client to access the Certificate Management page, where you can find the VMware CA.

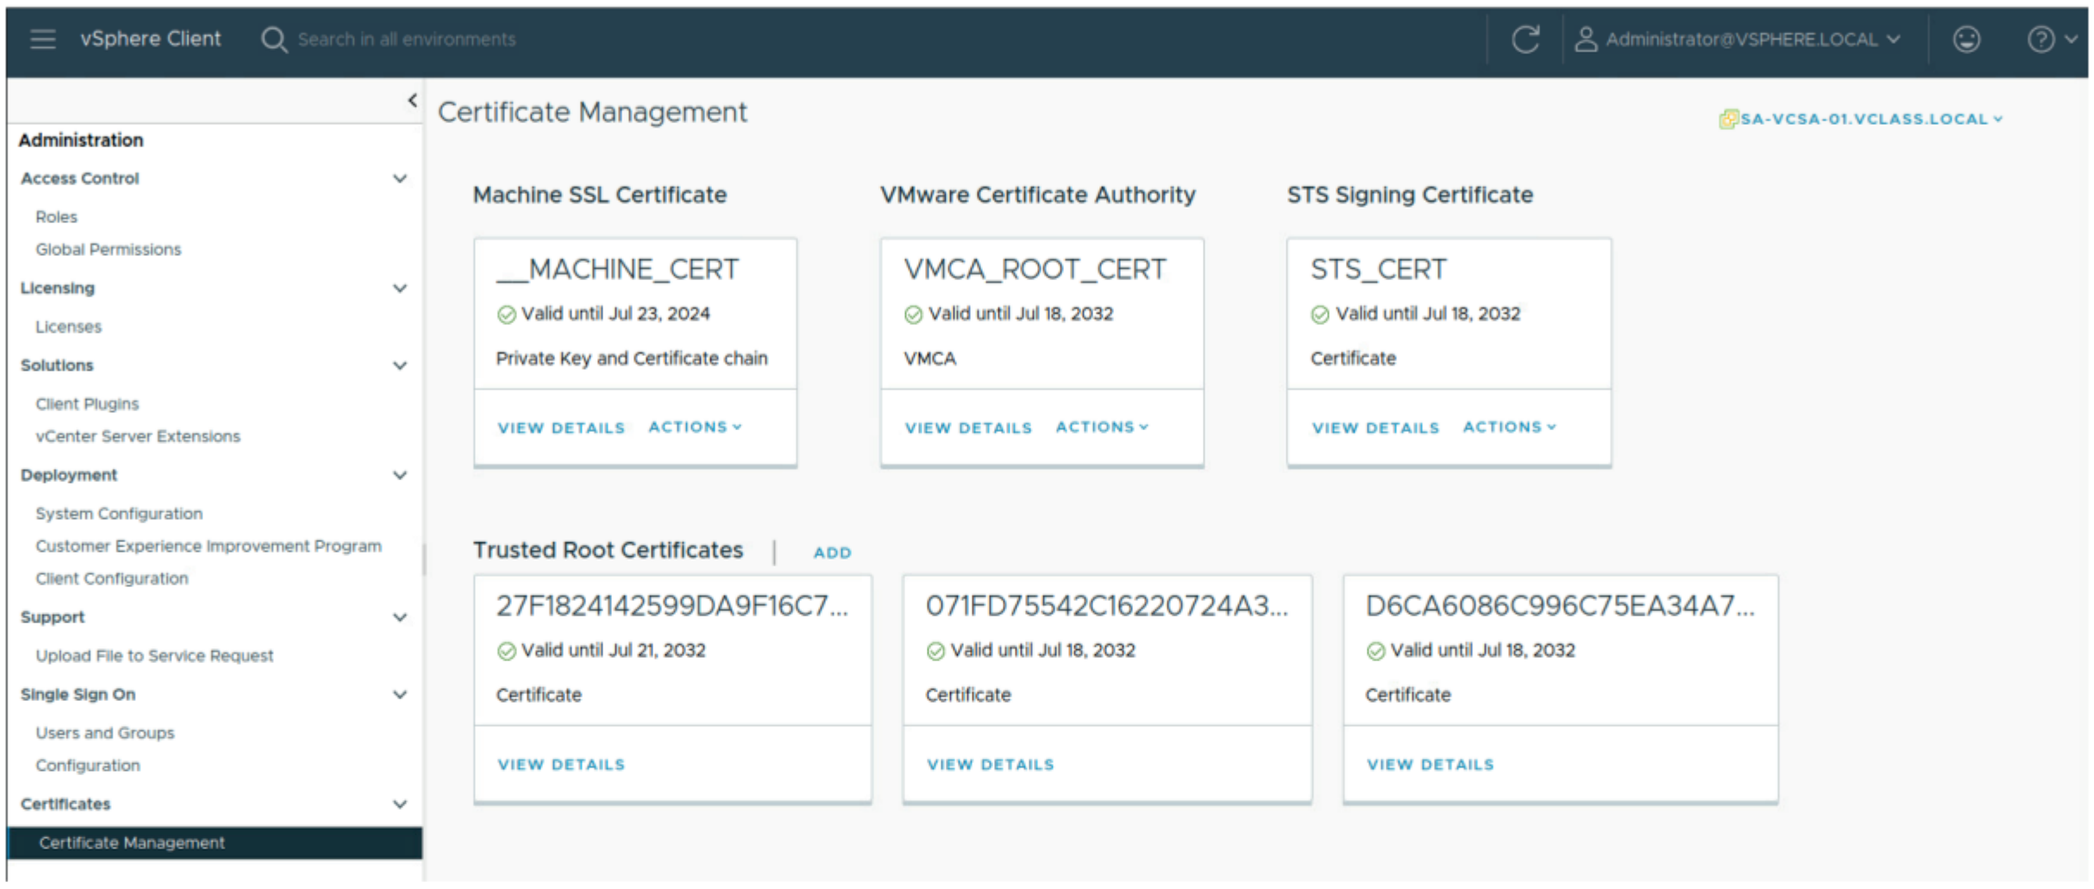

Technet24

6-61 VMware CA: Subordinate CA Mode

When VMware CA is in subordinate CA mode:

- You replace the default self-signed CA certificate that was created during installation with a trusted certificate from a corporate or commercial CA .
- You typically replace all issued default certificates after implementation.

This mode requires a certificate signing request (CSR) that is issued from VMware CA. The CSR is used in an enterprise or commercial CA to generate a new issuing certificate.

An enterprise CA signs the CSR that VMware CA generates. You then configure VMware CA to use this certificate and keys.

In subordinate CA mode, you replace the VMware CA root certificate with a third-party CA-signed certificate that includes VMware CA in the certificate chain. The replacement certificate can be signed by your enterprise CA or by a third-party, external CA.

All certificates that VMware CA will now include the full chain. You can replace existing certificates with newly generated certificates, although this step is not required. The original certificate is still valid if it has not expired or been revoked, and you can continue to use the certificates that were signed with the original certificate.

The subordinate CA mode combines the security of a third-party, CA-signed certificate with the convenience of the automated certificate management in VMware CA.

Issue a certificate means the acts of a certification authority in creating a certificate and notifying the subscriber listed in the certificate of the contents of the certificate.

6-62 Replacing the VMware CA Certificate with a Subordinate CA Certificate

To replace the VMware CA root CA certificate with an enterprise or third-party CA certificate, follow these high-level steps.

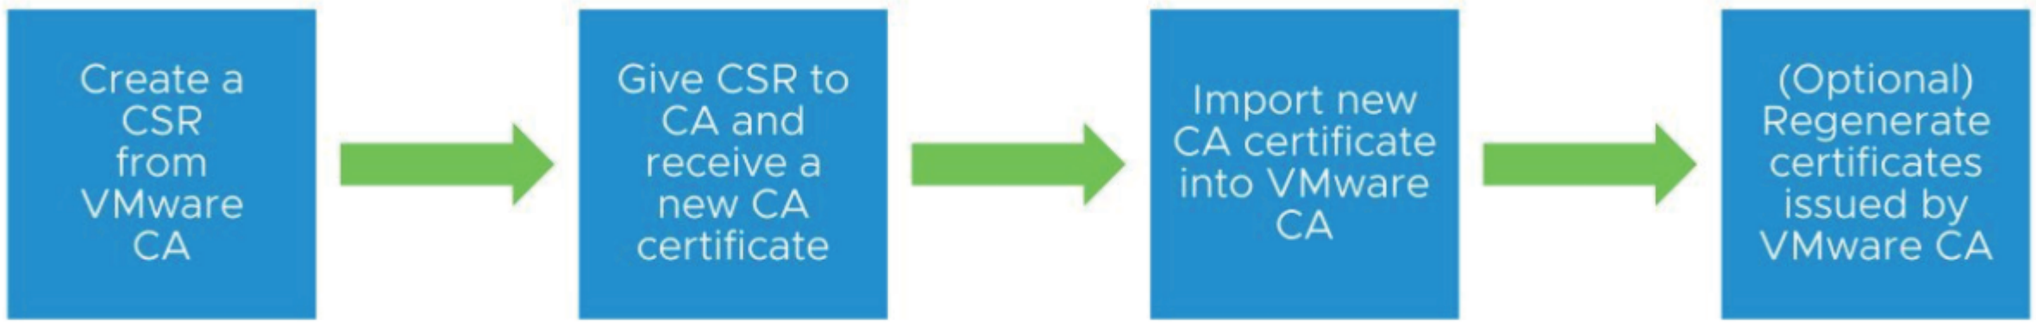

To configure your VMware CA as a subordinate CA, you must generate a certificate signing request (CSR) that you can present to your enterprise or third-party CA so that the CA can create an appropriate certificate for you. You must import an enterprise or third-party CA certificate into VMware CA.

6-63 VMware CA: Hybrid Mode

In Hybrid mode, the VMCA root certificate is replaced by a trusted certificate, the vCenter is then issued with a trusted certificate by the VMCA, but the ESXi hosts continue to use their default, untrusted certificates:

- Company policy may not allow intermediate CAs.
- Replace the vCenter Machine SSL Certificate with a custom certificate that was signed by the enterprise CA and saved in VECS.
- All other vSphere traffic is secured by the default self-signed certificates issued by VMware CA.

This mode requires a certificate signing request (CSR) that is issued by VMware CA.

An enterprise CA uses the CSR to sign a new certificate that will be used to replace the vCenter machine certificate.

For more information about replacing certificates, see *Replacing vSphere Certificates* at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-authentication/GUID-4469A6D3-048A-471C-9CB4-518A15EA2AC0.html.

Technet24

6-64 Hybrid Mode: Replacing a vCenter Machine Server Certificate with a Custom Certificate

Replacing the vCenter machine certificate with a custom certificate signed by an enterprise CA follows these high-level steps:

6-65 Certificate Replacement Options

For vCenter solutions, you must replace certificates in the following situations:

- In root CA mode, when you want to replace the VMware CA certificate with a new certificate
- In subordinate CA mode, when you replace the VMware CA certificate with a CA certificate that is issued by your enterprise or a third-party CA
- In custom mode, when you replace all your certificates with enterprise or third-party certificates

With the introduction of VMware CA in vSphere 6.0, managing your vSphere certificates is much easier. The situations described on the slide are the three most common ones.

6-66 Certificate Manager (1)

You can use the vSphere Client or the vSphere Certificate Manager to perform certificate tasks.

Option	Usage
1	Replace machine SSL certificate with custom CA certificate.
2	Replace VMware CA root certificate with custom CA certificate and replace all certificates.
3	Replace machine SSL certificate with VMware CA certificate.
4	Regenerate a new VMware CA root certificate and replace all certificates.
5	Replace solution user certificates with custom certificates.
6	Replace solution user certificates with VMware CA certificates.
7	Revert last-performed operation by republishing old certificates.
8	Reset all certificates.

With vSphere 7 and later, you can use the vSphere Client to replace the certificates in your vSphere environment and to reset all vSphere certificates. You can also use the vSphere Certificate Manager.

For more information about certificate management, see *vSphere Certificate Management Overview* at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-authentication/GUID-3D0DE463-D0EC-442E-B524-64759D063E25.html.

Technet24

6-67 Certificate Manager (2)

To start the vSphere Certificate Manager, you log in to the vCenter system and run the command at the command prompt.

```
root@sa-vcsa-01 [ ~ ]# /usr/lib/vmware-vmca/bin/certificate-manager
```

```
*** Welcome to the vSphere 7.0 Certificate Manager  ***

                -- Select Operation --

1. Replace Machine SSL certificate with Custom Certificate

2. Replace VMCA Root certificate with Custom Signing
   Certificate and replace all Certificates

3. Replace Machine SSL certificate with VMCA Certificate

4. Regenerate a new VMCA Root Certificate and
   replace all certificates

5. Replace Solution user certificates with
   Custom Certificate
   NOTE: Solution user certs will be deprecated in a future
   release of vCenter. Refer to release notes for more details.

6. Replace Solution user certificates with VMCA certificates

7. Revert last performed operation by re-publishing old
   certificates

8. Reset all Certificates

Note : Use Ctrl-D to exit.
Option[1 to 8]:
```

6-68 Replacing All VMware CA Certificates with Custom Certificates

You can replace all certificates that are issued by VMware CA with custom certificates if your site's policies require it.

When you replace VMware CA with custom certificates:

- VMware CA is not in your certificate chain.
- You are responsible for replacing all certificates yourself and for tracking certificate expirations.

To use custom certificates, send CSRs to your certificate provider and request the following items:

- A root certificate.
- A machine SSL certificate for each machine. The CSR for machine certificates must include the host name and the IP address for the machine.
- Solution user certificates should not include the IP address, host name, or email address. Each certificate should have the name of the solution user.

Replacing VMware CA issued certificates with custom enterprise or third-party certificates is not recommended. However, your own site security policies might dictate that you use only a specific CA for all your site's systems.

When you replace the VMware CA certificates with custom certificates, you increase your management tasks, which can increase the risk of errors that lead to security vulnerabilities. You are responsible for all certificate management tasks. Because you do not use VMware CA, you cannot rely on VMware CA to manage your certificates for you.

Technet24

6-69 ESXi Certificate Expiration

You can view information about certificate expiration for certificates that are signed by VMCA or a third-party CA in the vSphere Client.

A yellow alarm is raised if the certificate is in the Expiring Shortly state (less than eight months). A red alarm is raised if the certificate is in the Expiration Imminent state (less than two months)

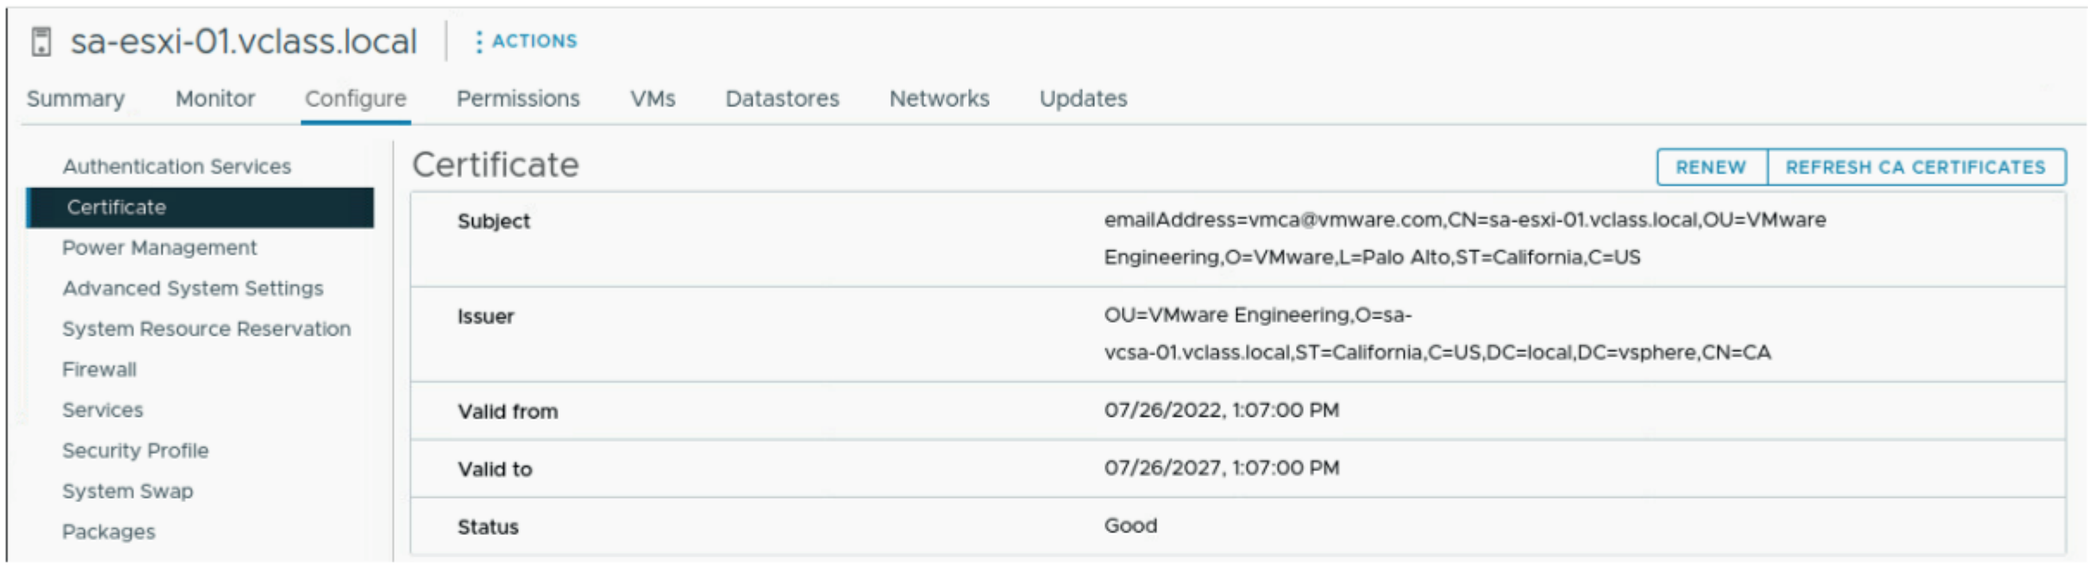

6-70 ESXi Certificate Replacement Options

Several options for ESXi certificate replacement are available:

- Replace with certificates issued by VMware CA:
  - VMware CA provisions the host certificates.
  - Host certificates include the full chain.
- Use enterprise or third-party certificates:
  - You can use your own custom certificates.
  - You must have an enterprise CA.

You can renew your certificates when they are about to expire, or if you want to provision the host with a new certificate for other reasons.

If you do not renew the certificate before it expires, disconnecting the host and reconnecting it causes vCenter to renew the certificate. The act of re-adding the host to vCenter reestablishes trust, and enables vCenter to unconditionally issue the renewed certificate.

6-71 Renewing ESXi Certificates

If VMCA assigned certificates to your ESXi hosts, you can renew or refresh those certificates from the vSphere Client.

- Renewing certificates:
  - Typically, you renew certificates when they are about to expire, but you can provision a host with a new certificate for other reasons.
  - Renewing retrieves a new signed certificate for the host from VMCA.
  - If the certificate is already expired, you must disconnect the host and reconnect it.
  - By default, vCenter renews the certificates of a host with the Expired, Expiring immediately, or Expiring status each time the host is added to the inventory or reconnected.

6-72 Refreshing ESXi Certificates

Refreshing CA certificates pushes all certificates in the TRUSTED_ROOTS store in the vCenter VECS store to the host. The TRUSTED_ROOTS store holds trusted root certificates of all VMware CA instances, as well as third-party trusted root certificates.

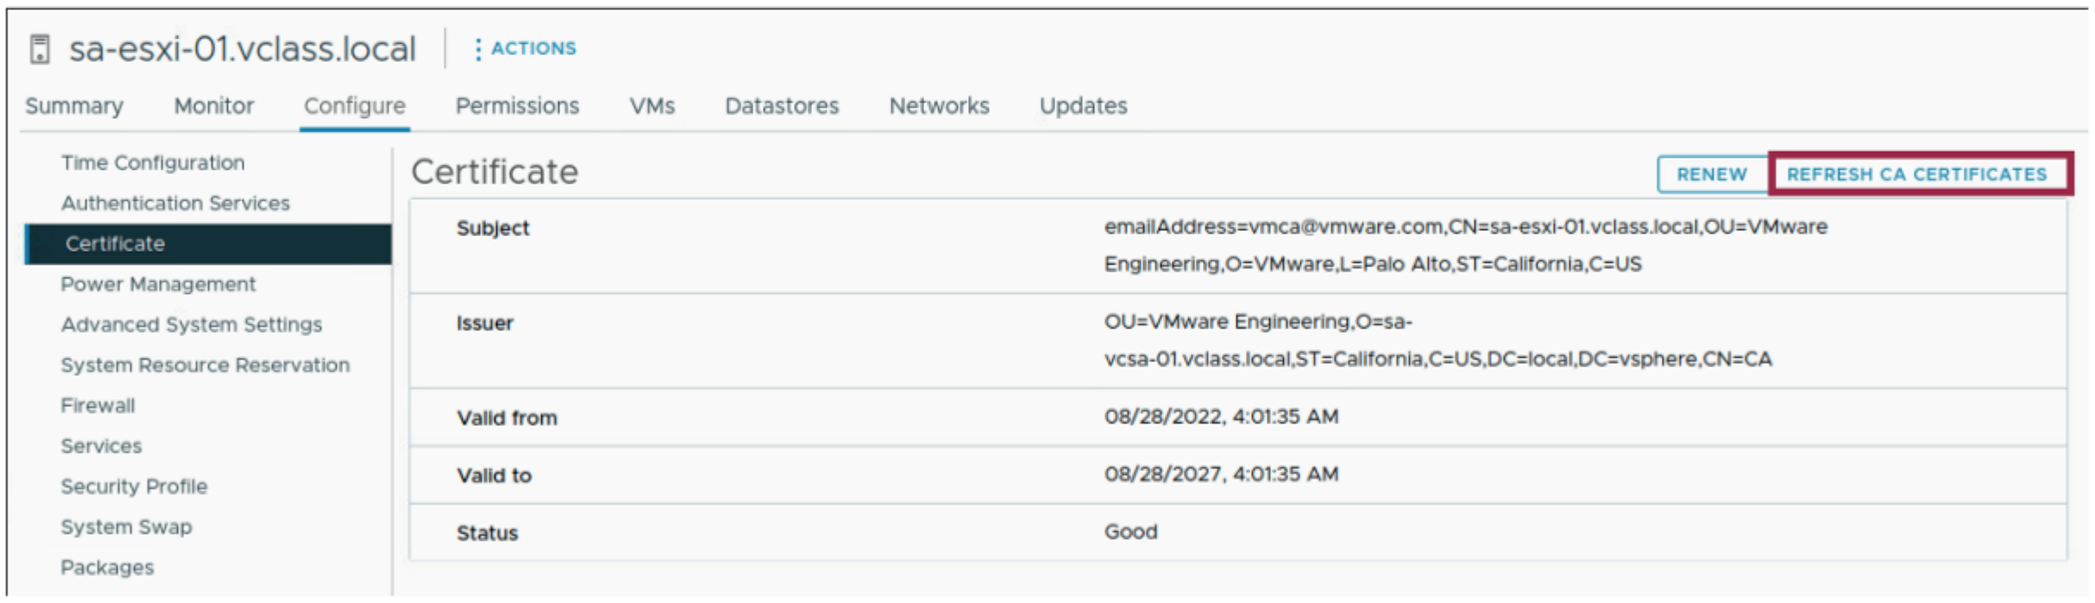

To renew or refresh CA certificates from the vSphere Client:

1. Browse to the host in the vSphere Client inventory.
2. Click the **Configure**.
3. Under System, click **Certificate**.
   You can view detailed information about the selected host's certificate.
4. Choose to **Renew** or **Refresh**.
5. Click Yes to confirm.

Technet24

6-73 VMware CA Availability

In the unlikely event that VMware CA is not available, some operations do not work:

- Adding new ESXi hosts
- Refreshing or renewing certificates on any ESXi host

Sometimes, certificate refreshing on certificate expiry might fail.

If VMware CA is unavailable for short periods of time, operations are not disrupted.

VMware CA is embedded with vCenter. When VMware CA is not available, you lose the ability to manage your certificates. However, issued certificates are still valid. Your ongoing vSphere services or functions should not be disrupted.

6-74 Lab 15: Working with Certificates

Generate and replace a vCenter certificate using the vSphere Client:

1. Examine the Machine SSL Certificate
2. Generate a Certificate Signing Request
3. Replace a Machine SSL Certificate with a Pregenerated CA Certificate

6-75 Review of Learner Objectives

- Use VMware Certificate Authority and VMware Endpoint Certificate Store to configure vSphere security certificate management
- Describe solution users and their certificate requirements
- Describe common VMware CA modes
- Replace VMware CA certificates with enterprise or third-party commercial certificates
- Use the vSphere Client to manage vSphere certificates

6-76 Key Points

- You can use the vCenter Management Interface to monitor appliance resource use and perform a file-based backup of the appliance.
- You can use vSphere Configuration Profiles to test the functionality and provide feedback during the features development process.
- vSphere certificate management simplifies the process of certificate replacement.
- VMware CA provisions each ESXi host and each vCenter service with certificates that are signed by VMware CA.
- vCenter High Availability is built in to vCenter and protects the appliance from both hardware and software failures.

Questions?

Technet24

300

Module 7

vSphere Monitoring

7-2 Importance

Although the VMkernel works proactively to avoid resource contention, maximizing performance requires both analysis and ongoing monitoring of your vSphere environment.

7-3 Module Lessons

1. Resource Monitoring Tools
2. Monitoring Resource Use
3. Monitoring vCenter and Performance
4. Using Alarms
5. Proactive Management with VMware Skyline

Technet24

7-4 Lesson 1: Resource Monitoring Tools

7-5 Learner Objectives

- Describe the performance-tuning methodology
- Identify resource-monitoring tools
- Use vCenter performance charts to view performance

7-6 Performance-Tuning Methodology

You can tune the performance of your vSphere environment:

- Assess performance:
  - Use appropriate monitoring tools.
  - Record a numerical benchmark before changes.
- Identify the limiting resource.
- Make more resources available:
  - Allocate more resources.
  - Reduce competition.
  - Log your changes.
- Benchmark again.

Do not make casual changes to production systems.

The best practice for performance tuning is to take a logical step-by-step approach:

- For a complete view of the performance situation of a VM, use monitoring tools in the guest operating system and in vCenter.
- Identify the resource that the VM relies on the most. This resource is most likely to affect the VM's performance if the VM is constrained by it.
- Give a VM more resources or decrease the resources of other VMs.
- After making more of the limiting resource available to the VM, take another benchmark and record changes.

Be cautious when making changes to production systems because a change might negatively affect the performance of the VMs.

7-7 Resource-Monitoring Tools

Many resource-monitoring and performance-monitoring tools are available for use with vSphere.

Inside the Guest OS
Perfmon DLL **Task Manager** **The `top` command (Linux)**

Outside the Guest OS
vCenter performance charts **vRealize Operations** **vSphere/ESXi system logs** **`resxtop` and `esxtop`**

Tools in the guest operating system are available from sources external to VMware and are used in various VMware applications. Many tools used outside of the guest OS are made available by VMware for use with vSphere and other applications.

A partial list of these resource-monitoring tools is shown on the slide.

Technet24

7-8 Guest Operating System Monitoring Tools

To monitor performance in the guest operating system, use tools that you are familiar with, such as Windows Task Manager.

Windows Task Manager

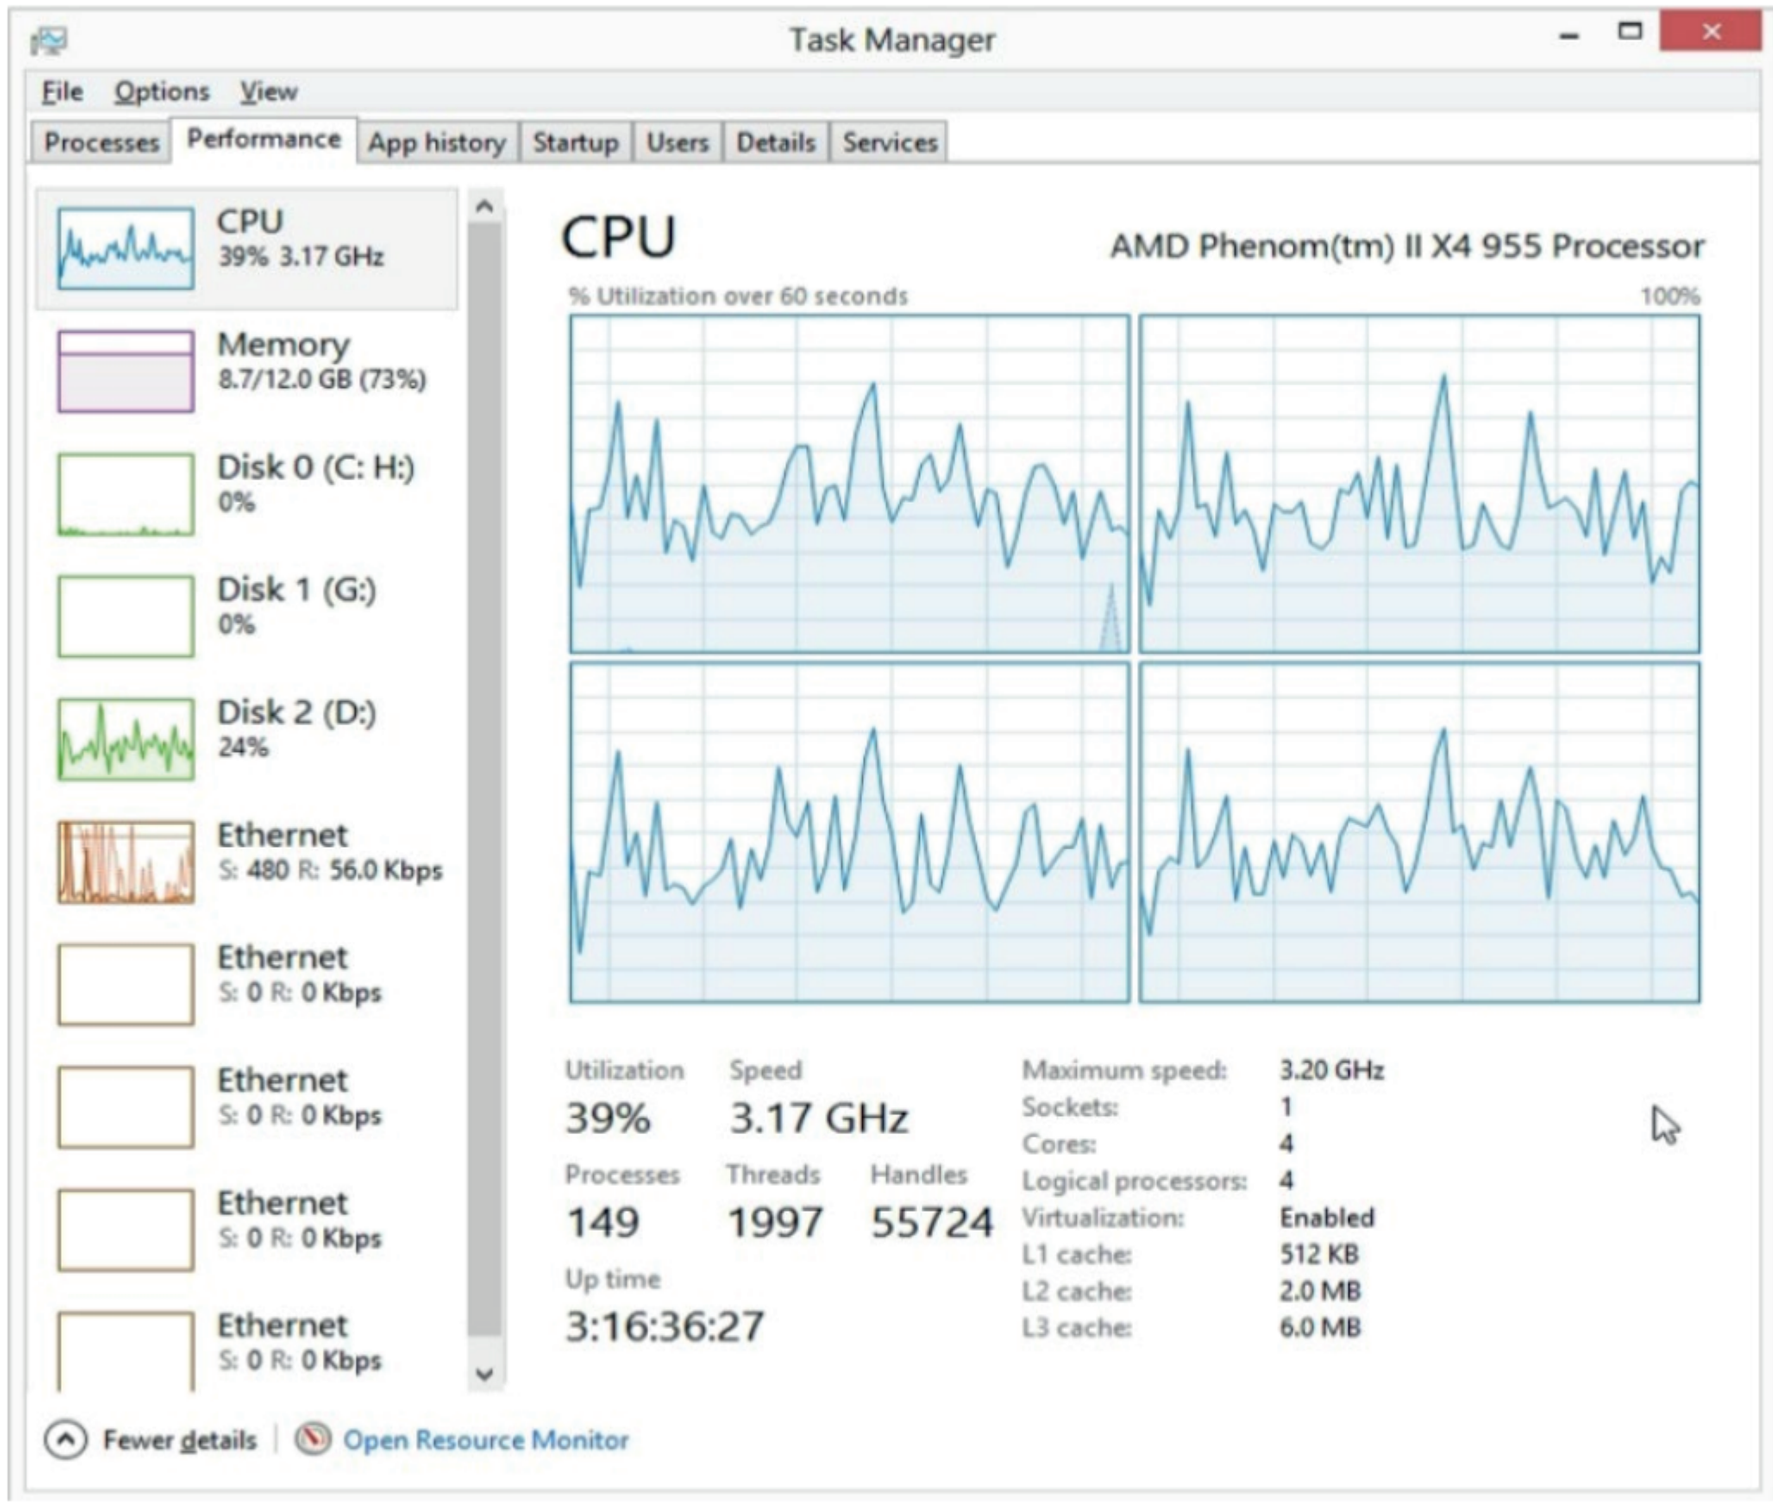

Windows Task Manager helps you measure CPU and memory use in the guest operating system.

The measurements that you take with tools in the guest operating system reflect resource usage of the guest operating system, not necessarily of the VM itself.

7-9 Using Perfmon to Monitor VM Resources

The Perfmon DLL in VMware Tools provides VM processor and memory objects for accessing host statistics in a VM.

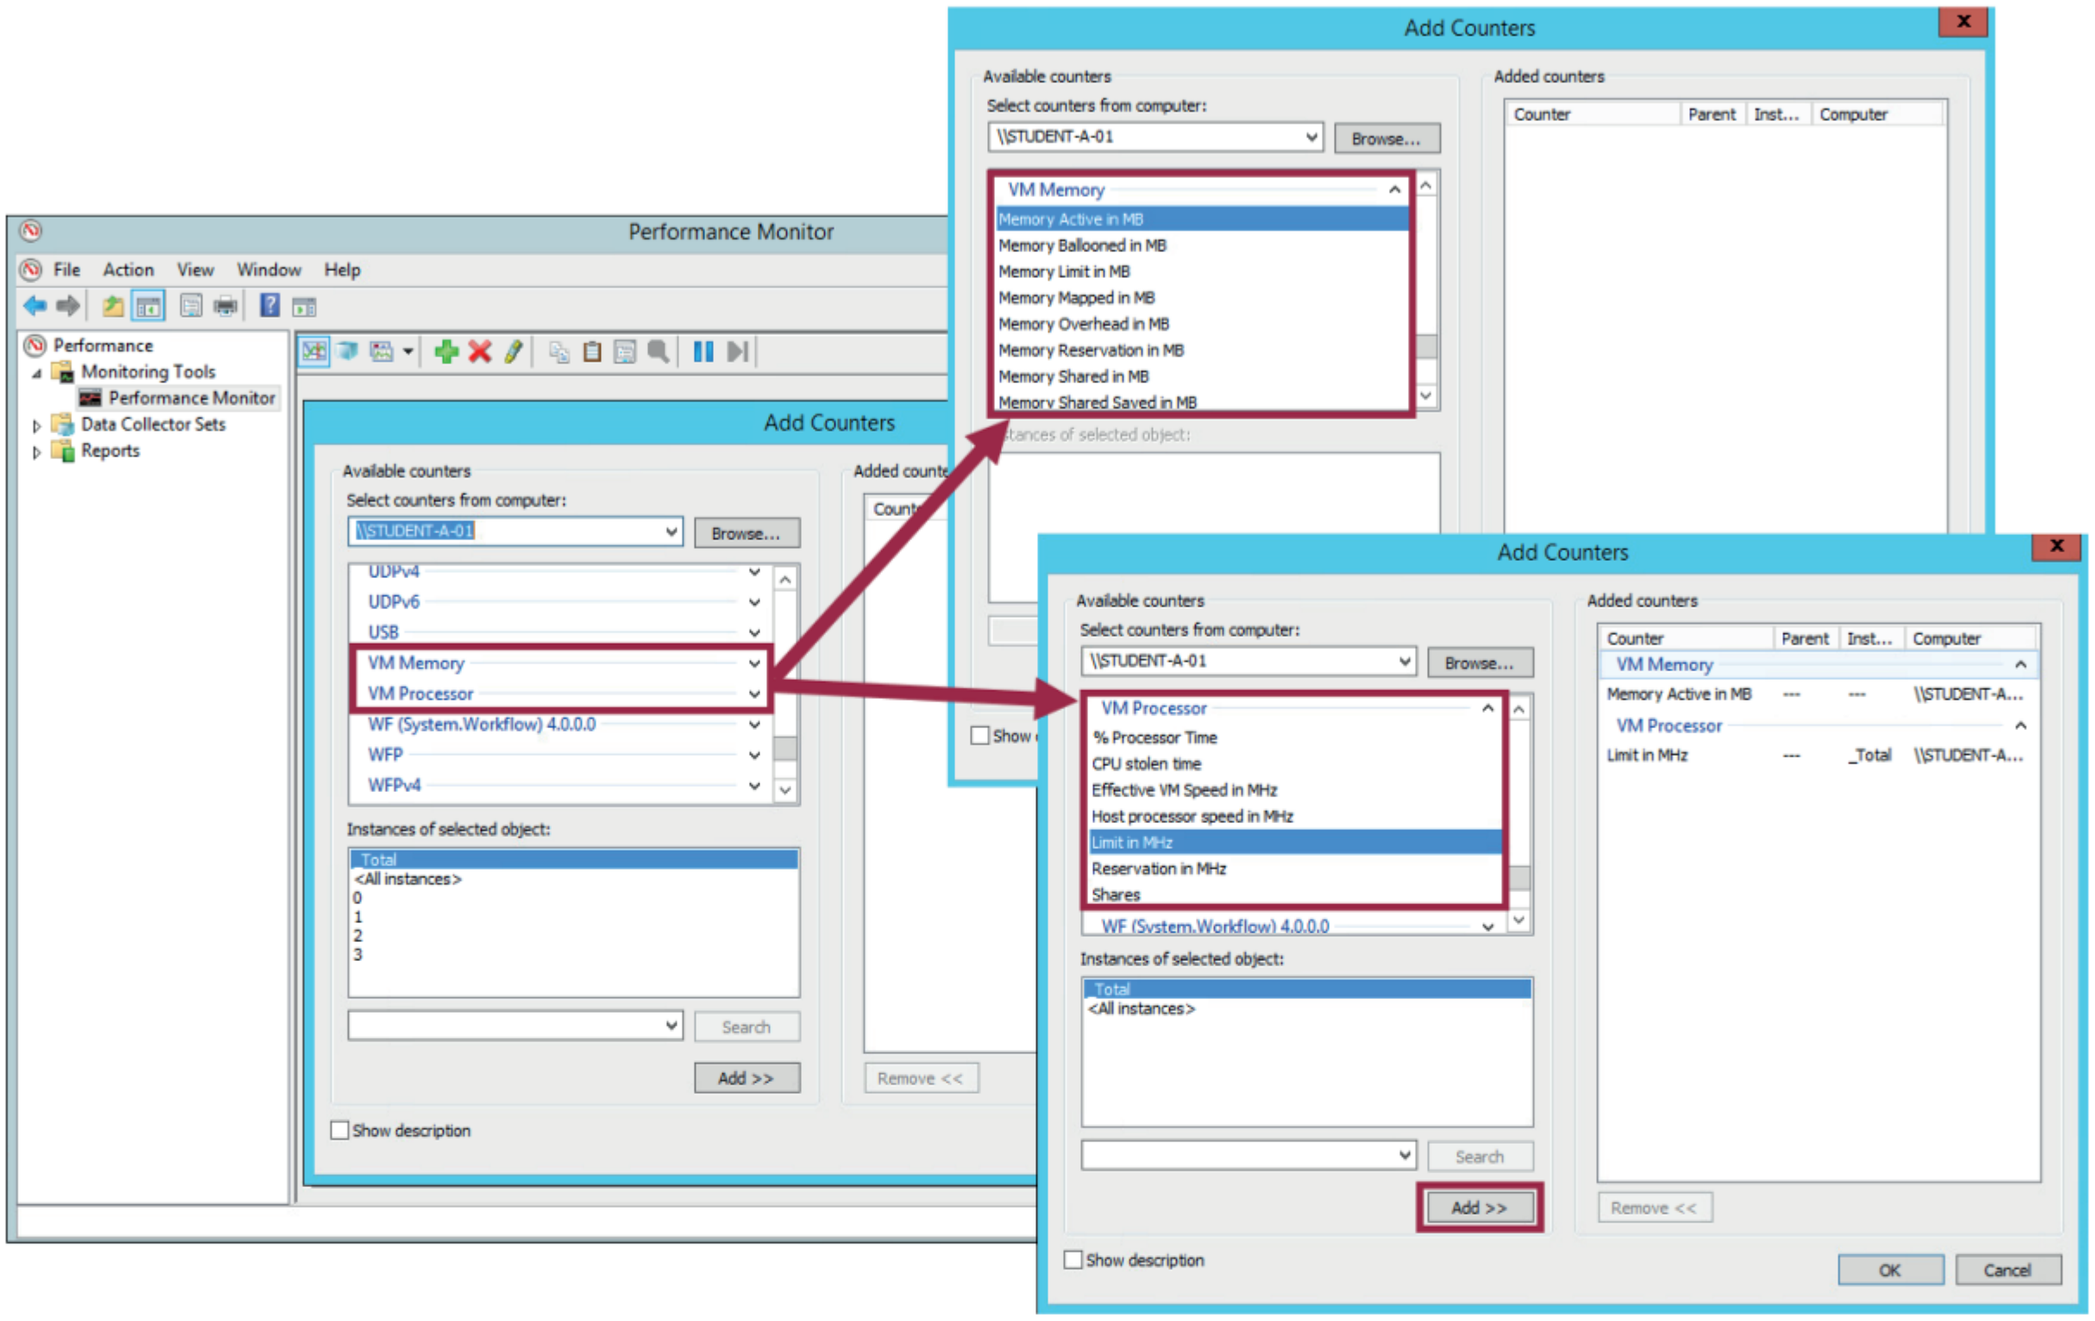

VMware Tools includes a library of functions called the Perfmon DLL. With Perfmon, you can access key host statistics in a guest VM. Using the Perfmon performance objects (VM Processor and VM Memory), you can view actual CPU and memory usage and observed CPU and memory usage of the guest operating system.

For example, you can use the VM Processor object to view the % Processor Time counter, which monitors the VM's current virtual processor load. Likewise, you can use the Processor object and view the % Processor Time counter (not shown), which monitors the total use of the processor by all running processes.

Technet24

7-10 Using esxtop to Monitor VM Resources

The `esxtop` utility is the primary real-time performance monitoring tool for vSphere:

- Can be run from the host's local vSphere ESXi Shell as `esxtop`
- Can be run remotely using `resxtop`
- Works like the `top` performance utility in Linux operating systems

In this example, you enter lowercase **c** and uppercase **V** to view CPU metrics for VMs.

```
 2:48:37pm up 7 days 13:36, 493 worlds, 3 VMs, 3 vCPUs; CPU load average: 0.03, 0.03, 0.02
PCPU USED(%): 1.8 3.9 AVG: 2.8
PCPU UTIL(%): 2.2 4.2 AVG: 3.2
```

ID	GID	NAME	NWLD	%USED	%RUN	%SYS	%WAIT	%VMWAIT	%RDY	%IDLE	%OVRLP	%CSTP
52127	52127	VM2-1	8	0.15	0.14	0.00	800.00	0.72	0.03	99.67	0.00	0.00
43316	43316	VM1-2	8	0.11	0.10	0.01	800.00	0.74	0.07	99.67	0.00	0.00
55009	55009	VM2-2	8	0.11	0.10	0.01	800.00	0.73	0.12	99.63	0.00	0.00

You can run the `esxtop` utility by using vSphere ESXi Shell to communicate with the management interface of the ESXi host. You must have root user privileges.

is a command-line utility or tool that runs on linux and provides a detailed look at how ESXi uses resources in real time

7-11 Monitoring Inventory Objects with Performance Charts

The vSphere statistics subsystem collects data on the resource usage of inventory objects, which include:

- Clusters
- Hosts
- Datastores
- Networks
- Virtual machines

Data on a wide range of metrics is collected at frequent intervals, processed, and archived in the vCenter database. You can access statistical information through command-line monitoring utilities or by viewing performance charts in the vSphere Client.

7-12 Working with Overview Performance Charts

The overview performance charts display the most common metrics for an object in the inventory.

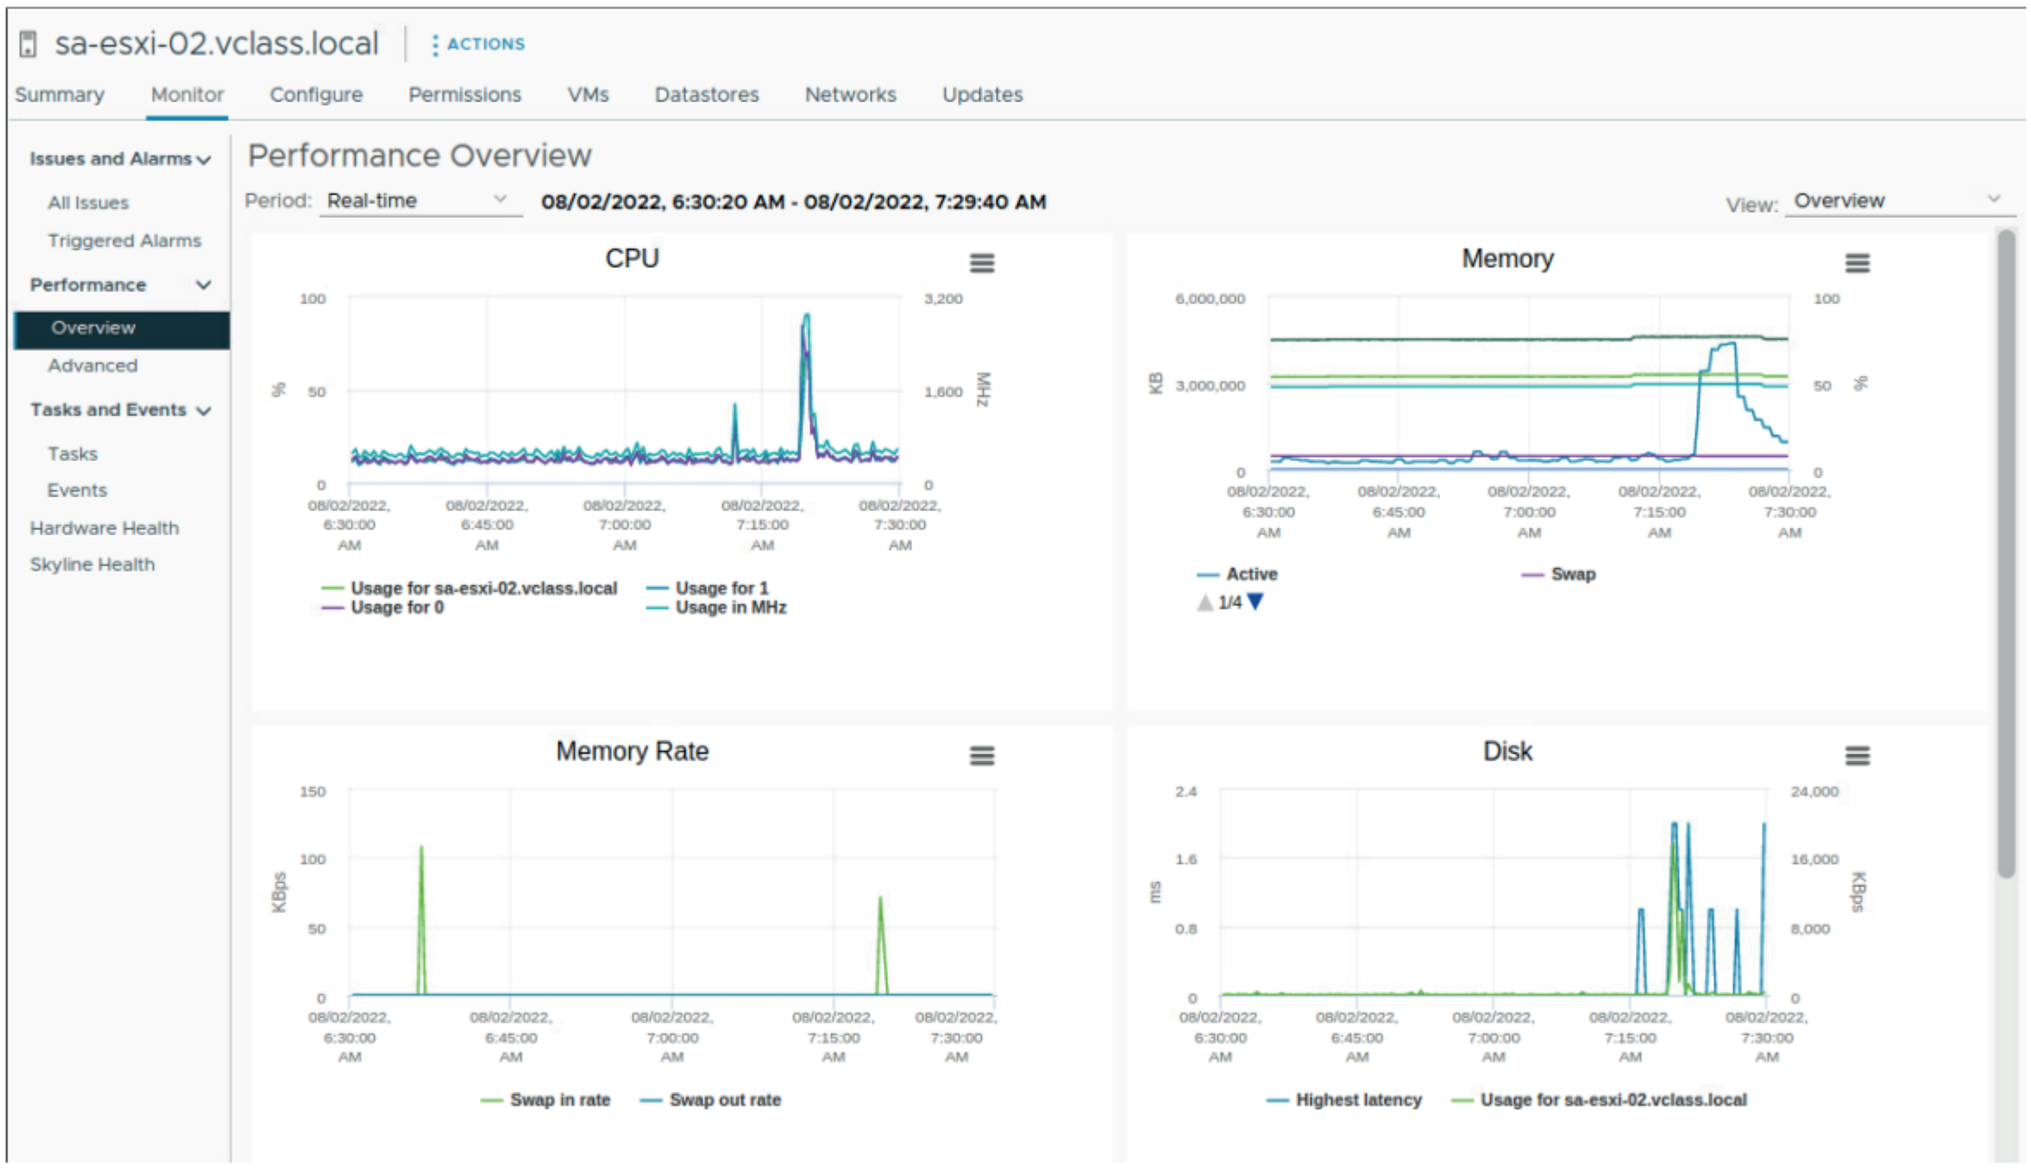

You can access overview and advanced performance charts in the vSphere Client.

Overview performance charts show the performance statistics that VMware considers most useful for monitoring performance and diagnosing problems.

Depending on the object that you select in the inventory, the performance charts provide a quick visual representation of how your host or VM is performing.

Technet24

7-13 Working with Advanced Performance Charts

Advanced charts support data counters that are not necessarily available in other performance charts.

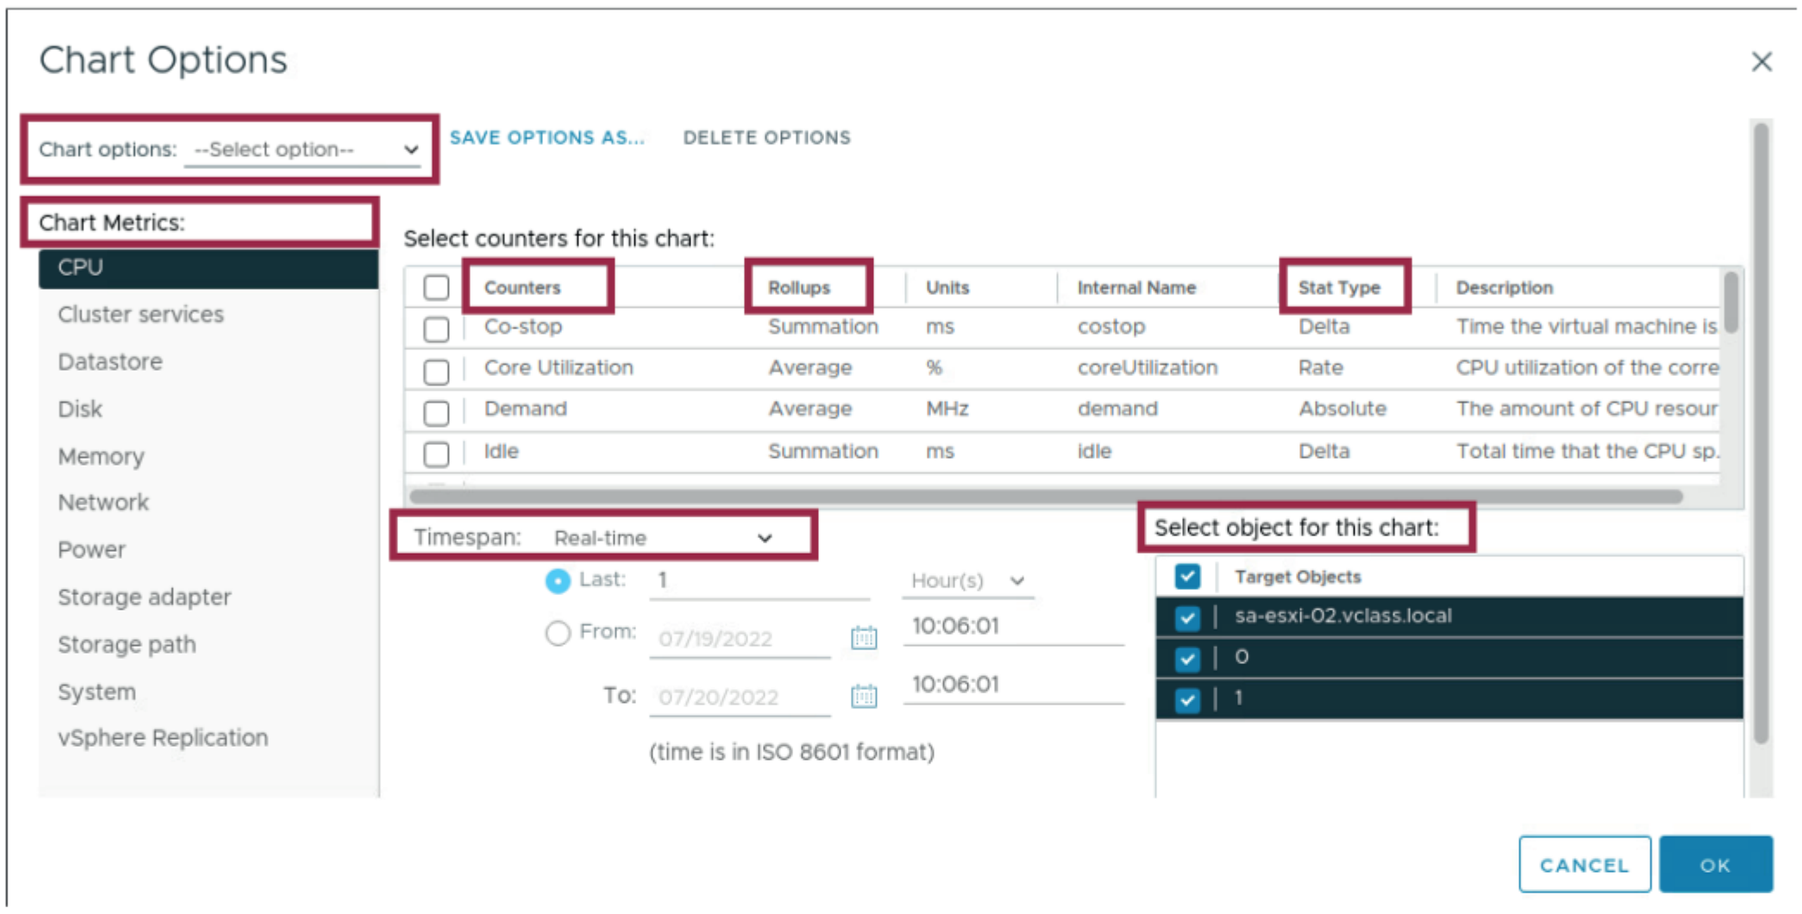

In the vSphere Client, you can customize the appearance of advanced performance charts.

Advanced charts have the following features:

- More information than overview charts: Point to a data point in a chart to display details about that specific data point.
- Customizable charts: Change chart settings. Save custom settings to create your own charts.
- Save data to an image file or a spreadsheet.

To customize advanced performance charts, select **Advanced** under Performance. Click the **Chart Options** link in the Advanced Performance pane.

Advanced performance charts are available on an object's **Performance** > **Advanced** page and are highly configurable in Chart Options.

7-14 Chart Options: Real-Time and Historical

vCenter stores statistics at different specificities.

Time Interval	Data Frequency	Number of Samples	Stored In
Real-time (past hour)	20 seconds	180	vCenter Memory
Past day	5 minutes	288	vCenter Database
Past week	30 minutes	336	vCenter Database
Past month	2 hours	360	vCenter Database
Past year	1 day	365	vCenter Database

Real-time information is information that is generated for the past hour at 20-second intervals. Historical information is generated for the past day, week, month, or year, at varying specificities.

By default, vCenter has four archiving intervals: day, week, month, and year. Each interval specifies a length of time that statistics are archived in the vCenter database.

You can configure which intervals are used and for what period of time. You can also configure the number of data counters that are used during a collection interval by setting the collection level.

Together, the collection interval and the collection level determine how much statistical data is collected and stored in your vCenter database.

For example, using the table, past-day statistics show one data point every 5 minutes, for a total of 288 samples. Past-year statistics show 1 data point per day, or 365 samples.

Real-time statistics are not stored in the database. They are stored in a flat file on ESXi hosts and in memory on vCenter instances. ESXi hosts collect real-time statistics only for the host or the VMs that are available on the host. Real-time statistics are collected directly on an ESXi host every 20 seconds.

If you query for real-time statistics, vCenter queries each host directly for the data. vCenter does not process the data at this point. vCenter only passes the data to the vSphere Client.

On ESXi hosts, the statistics are kept for 30 minutes, after which 90 data points are collected. The data points are aggregated, processed, and returned to vCenter. vCenter then archives the data in the database as a data point for the day collection interval.

To ensure that performance is not impaired when collecting and writing the data to the database, cyclical queries are used to collect data counter statistics. The queries occur for a specified collection interval. At the end of each interval, the data calculation occurs.

Technet24

7-15 Chart Types: Bar and Pie

Depending on the metric type and object, performance metrics are displayed in different types of charts, such as bar charts and pie charts.

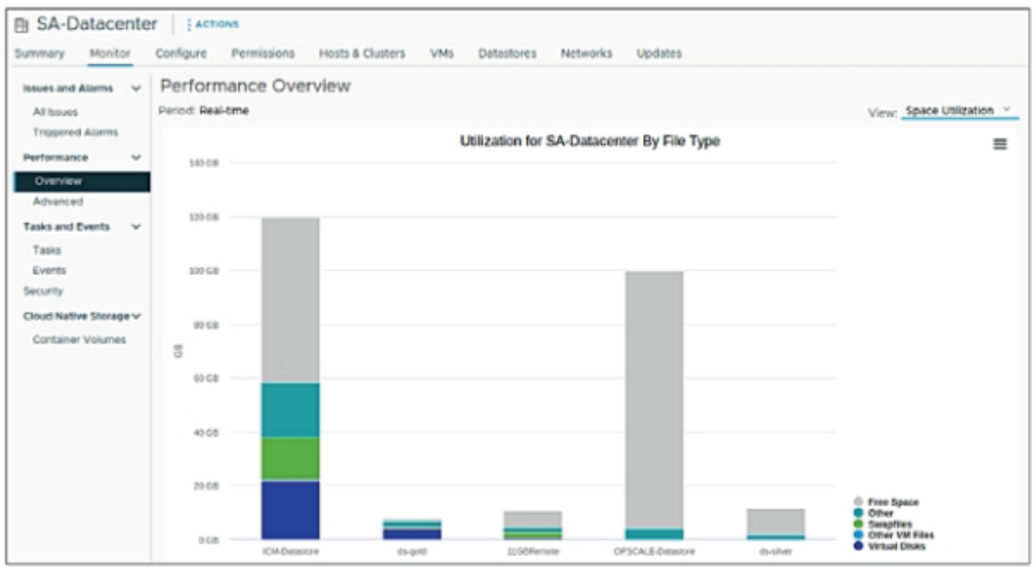

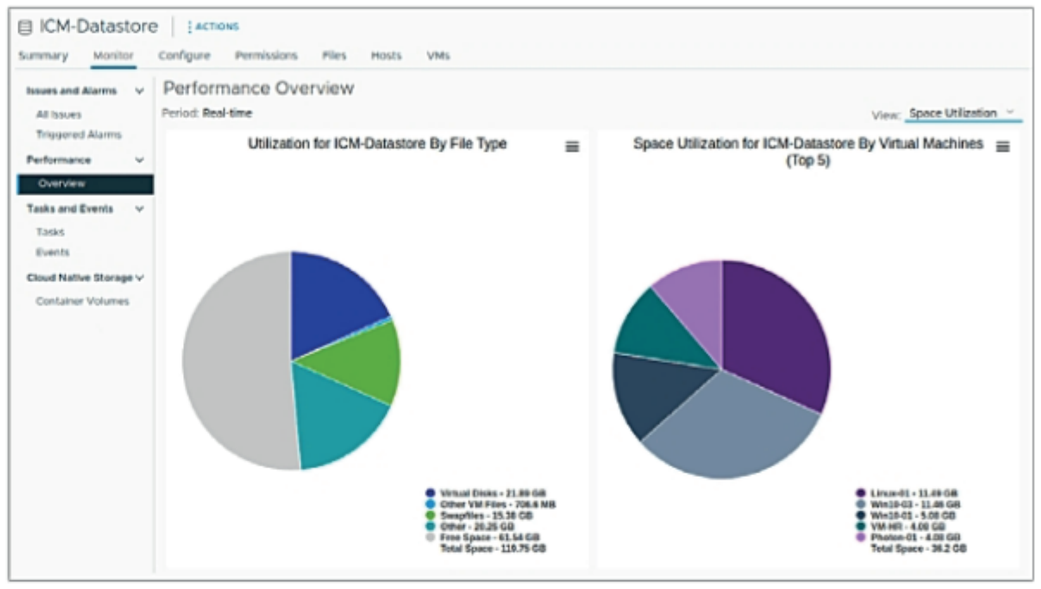

Bar charts display storage metrics for datastores in a selected data center. Each datastore is represented as a bar in the chart. Each bar displays metrics based on the file type: virtual disks, other VM files, snapshots, swap files, and other files.

Pie charts display storage metrics for a single object, based on the file types or VMs. For example, a pie chart for a datastore can display the amount of storage space occupied by the VMs that take up the largest space.

7-16 Chart Types: Line

A line chart can display multiple metrics for a single inventory object, for example, metrics for each CPU on an ESXi host.

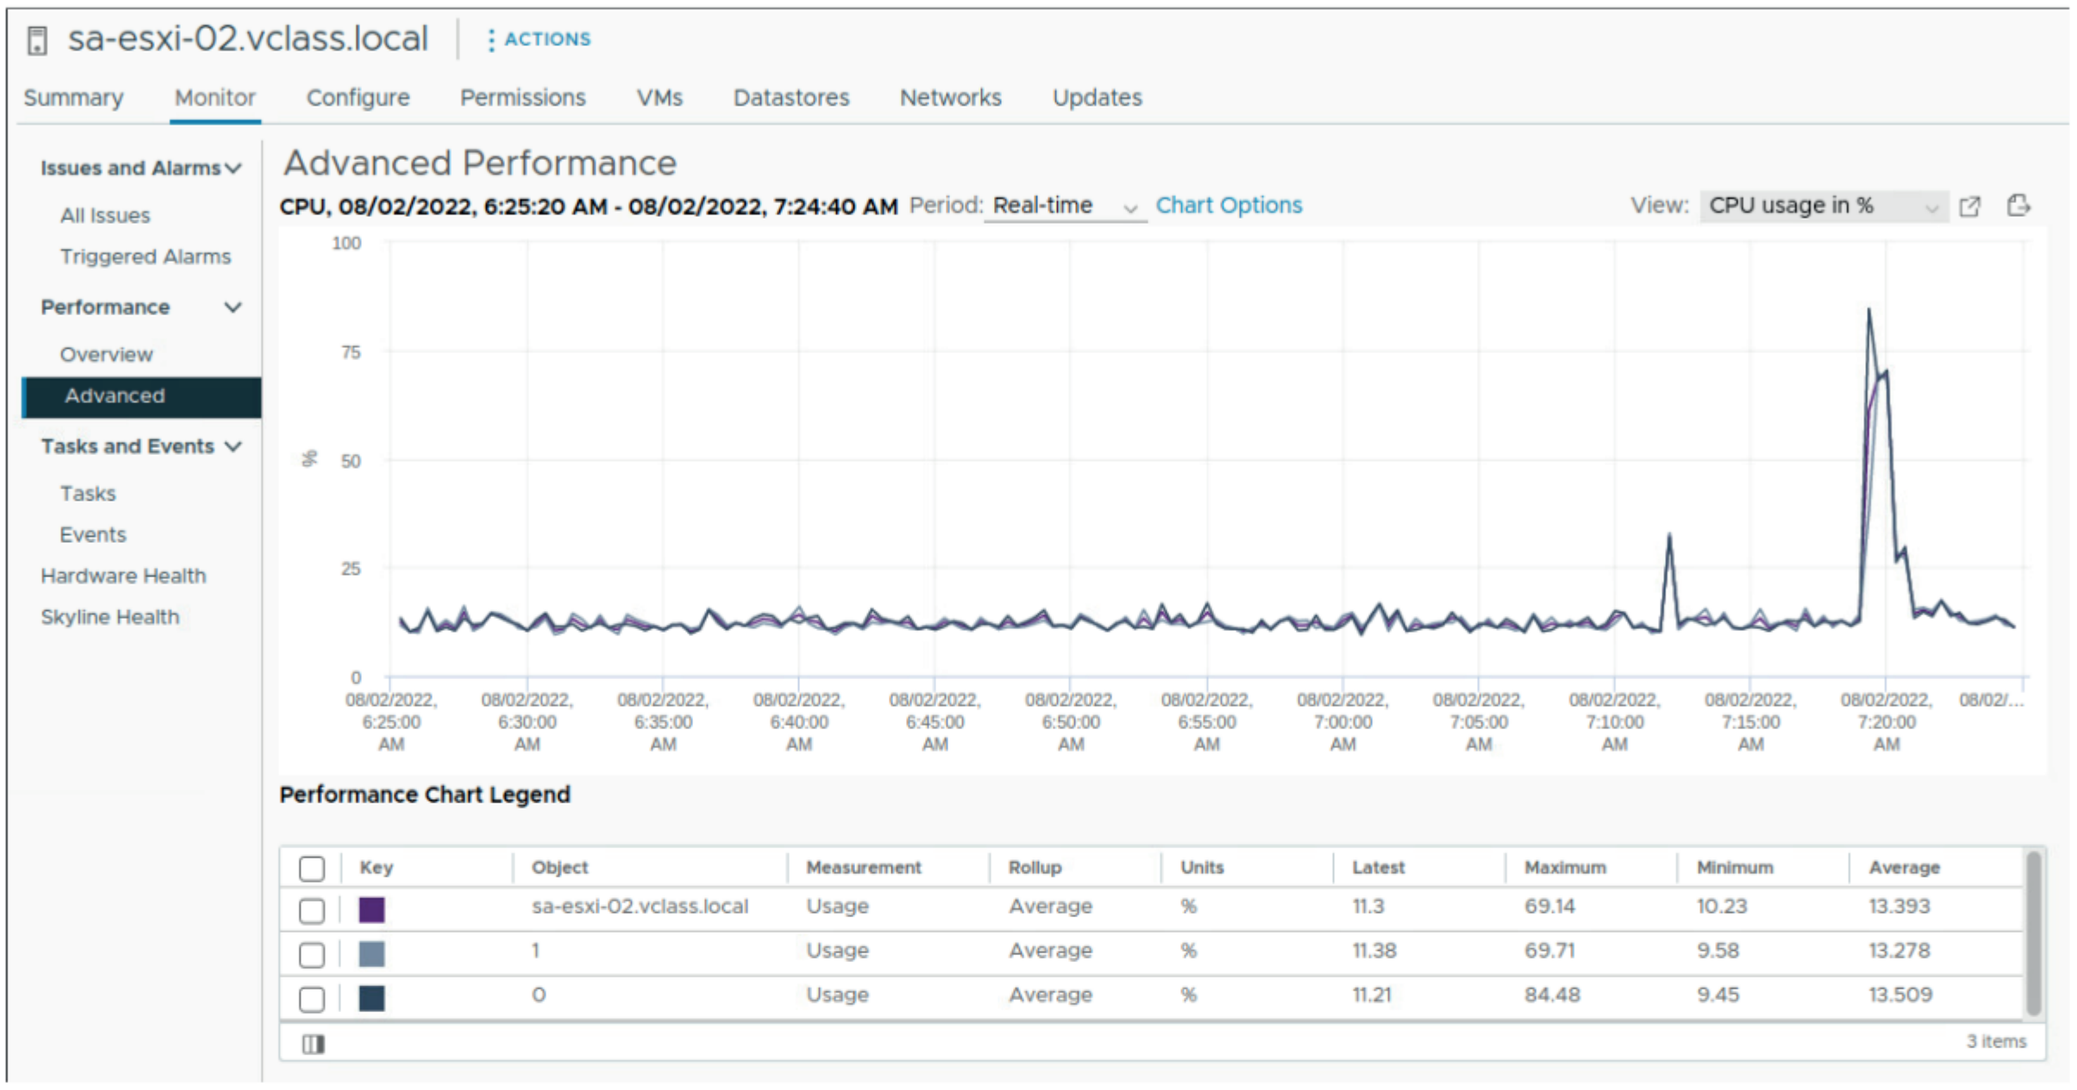

In a line chart, the data for each performance counter is plotted on a separate line in the chart. For example, a CPU chart for a host can contain a line for each of the host's CPUs. Each line plots the CPU's usage over time.

Technet24

7-17 Chart Types: Stacked

Stacked charts are useful for comparing resource allocation and usage across hosts, VMs or multiple objects.

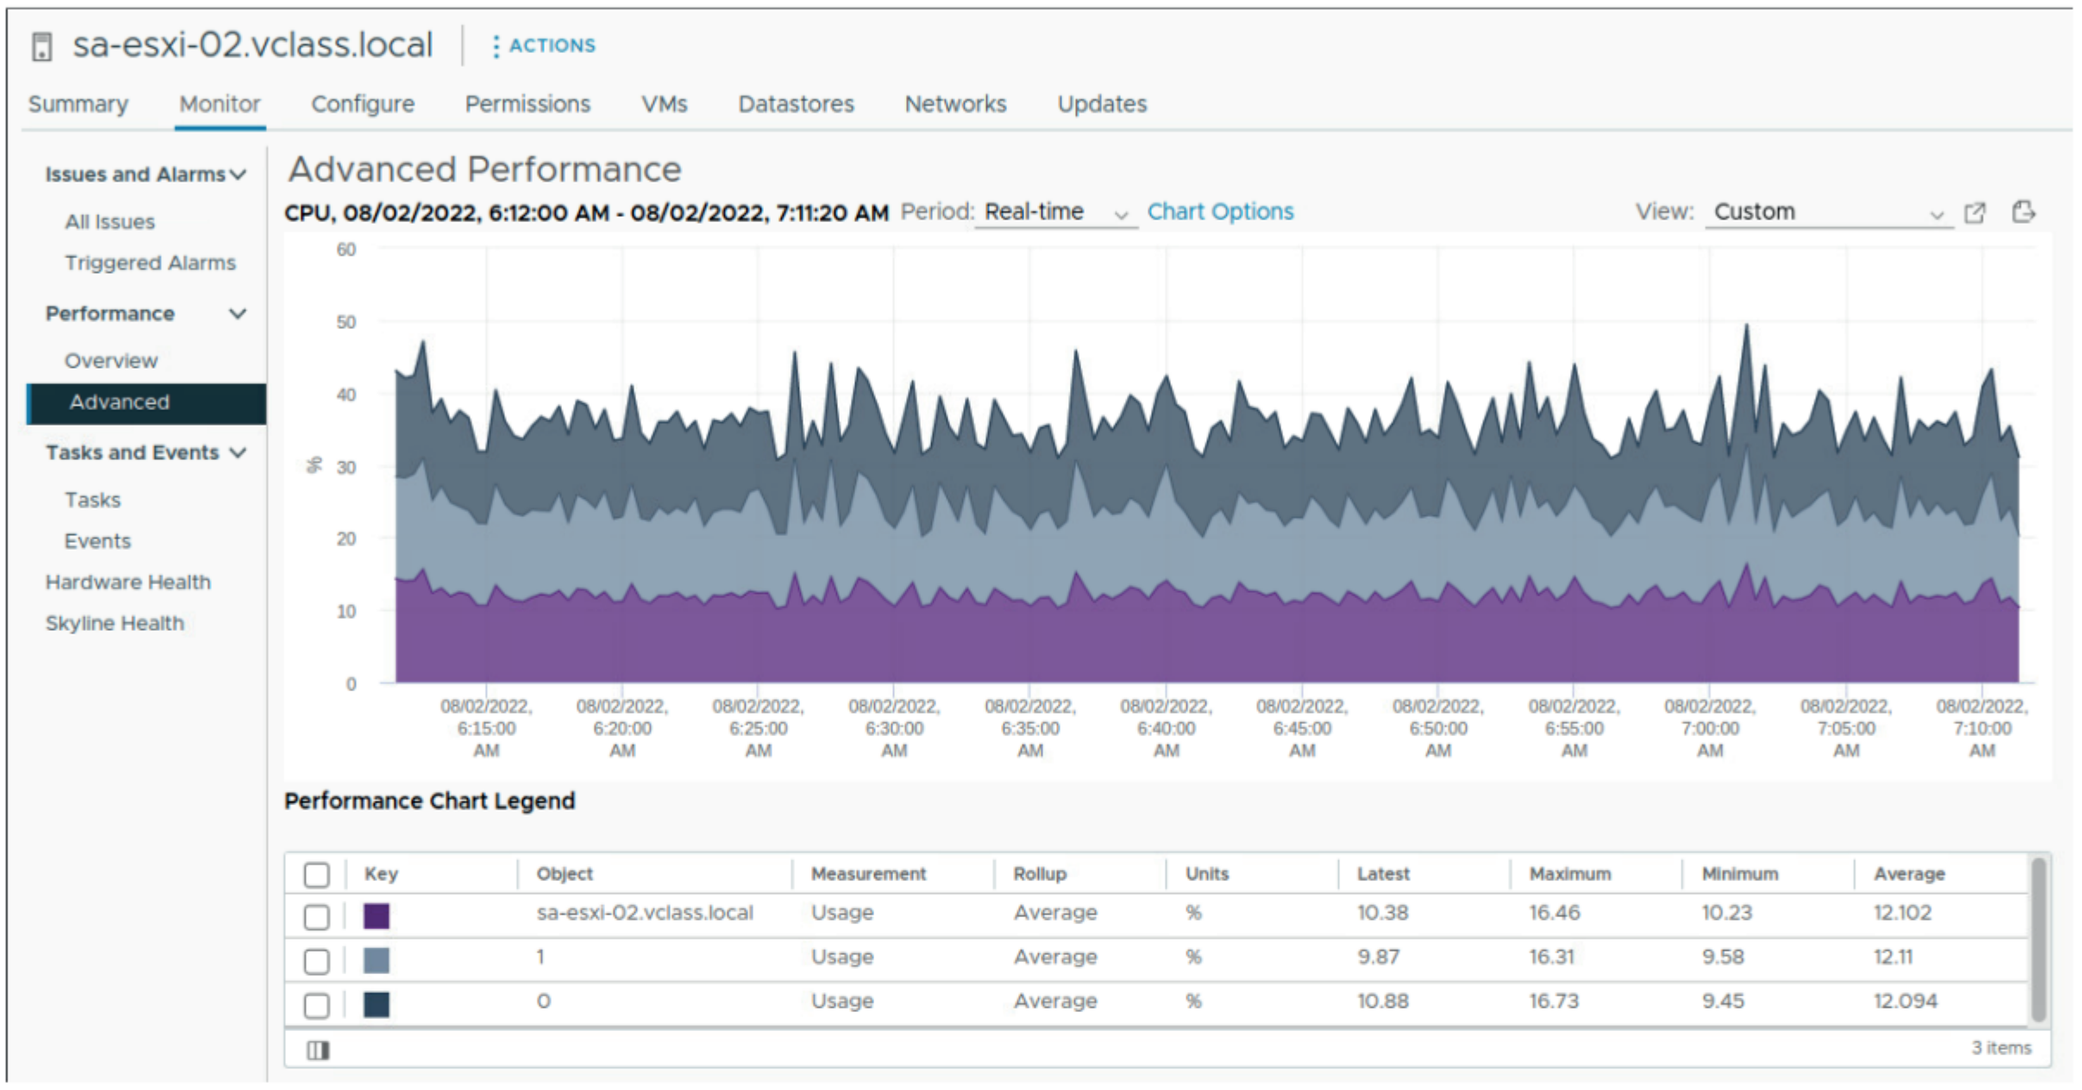

Stacked charts can only view a single metric, unlike line charts that can view multiple metrics.

Line charts can typically show more information, but stacked charts are typically easier to read

7-18 Chart Types: Stacked Per VM

Per-VM stacked graphs are available only for hosts.

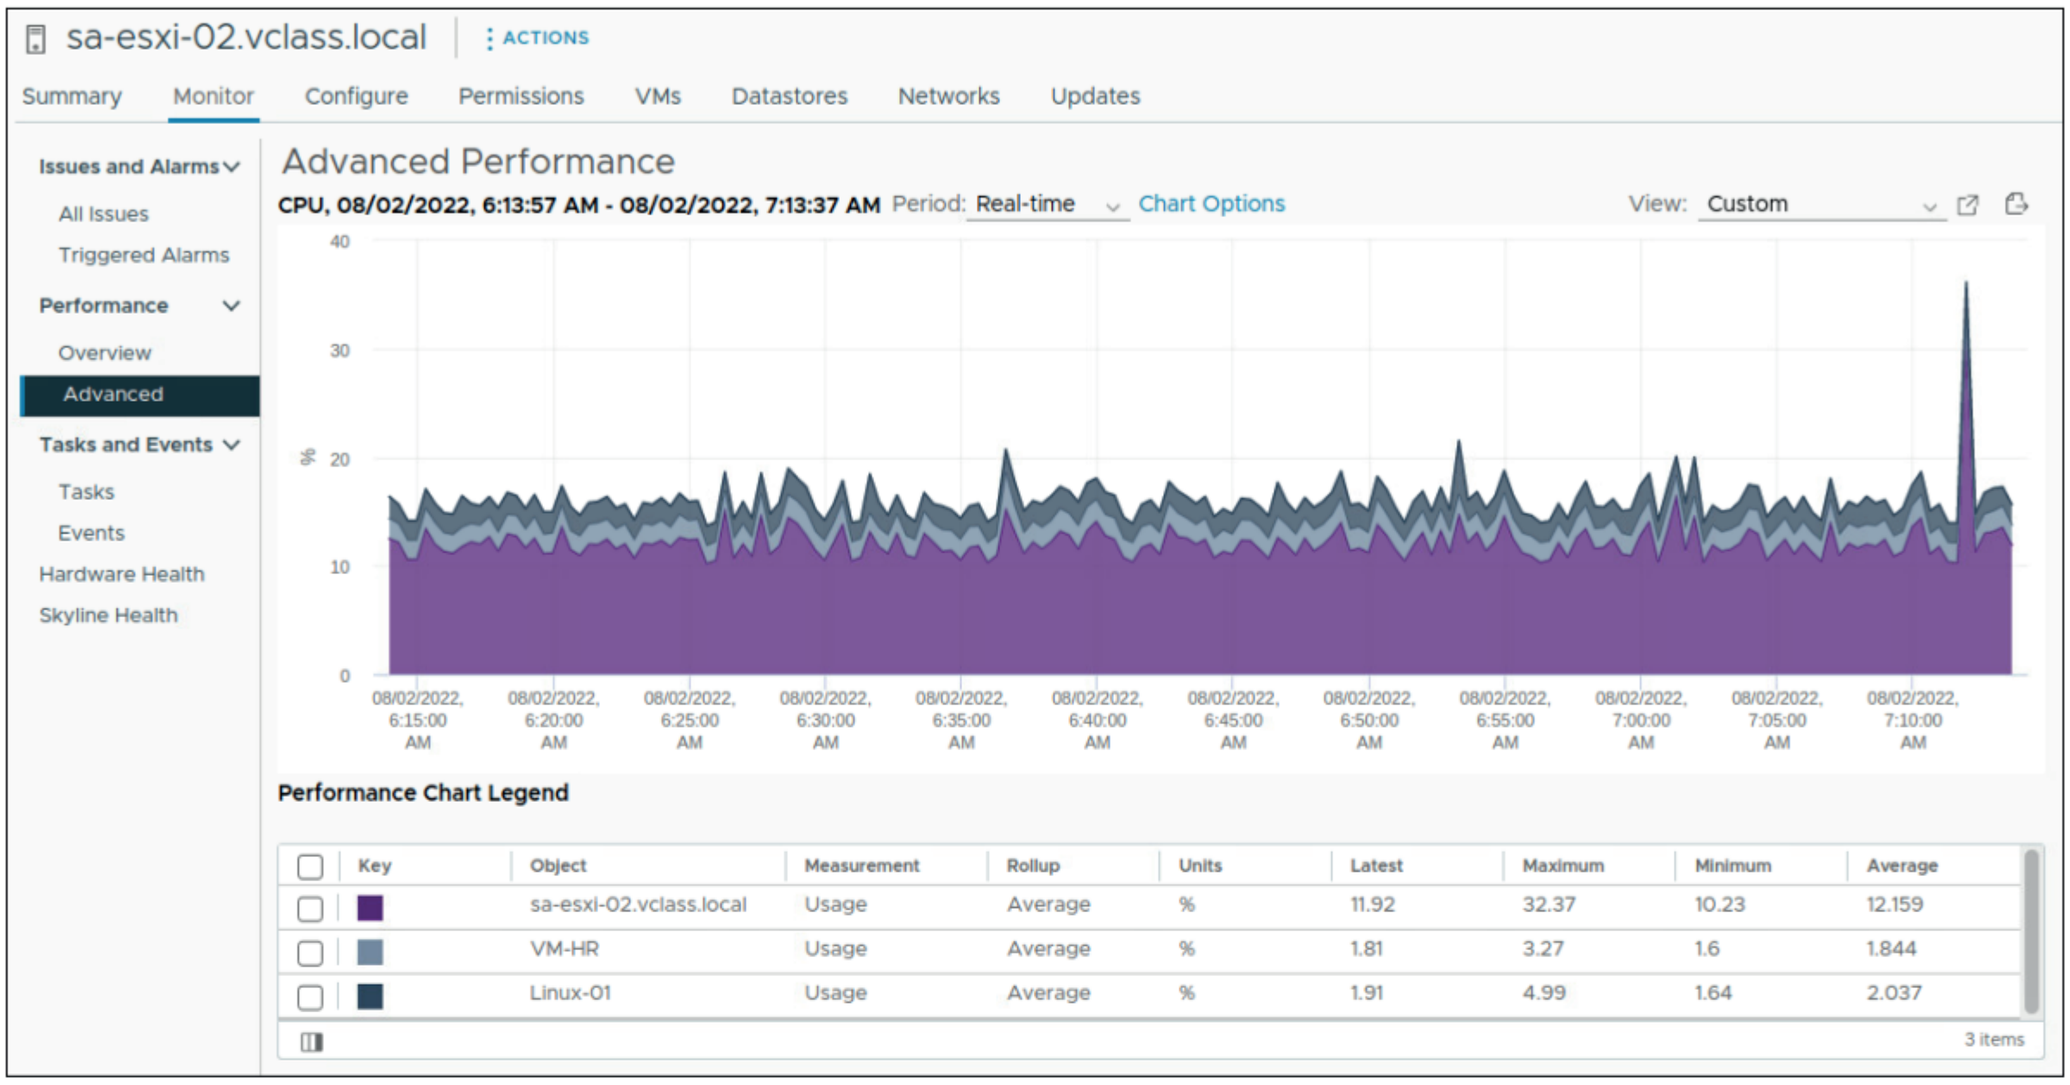

Stacked charts display metrics for the child objects that have the highest statistical values. All other objects are aggregated, and the sum value is displayed with the term Other. For example, a host's stacked CPU usage chart displays CPU usage metrics for the five VMs on the host that are consuming the most CPU resources. The Other amount contains the total CPU usage of the remaining VMs. The metrics for the host itself are displayed in separate line charts. By default, the 10 child objects with the highest data counter values appear.

Technet24

7-19 Saving Charts

You click the **Save Chart** icon above the graph to save performance chart information.

You can save information in PNG, JPEG, SVG, and CSV formats.

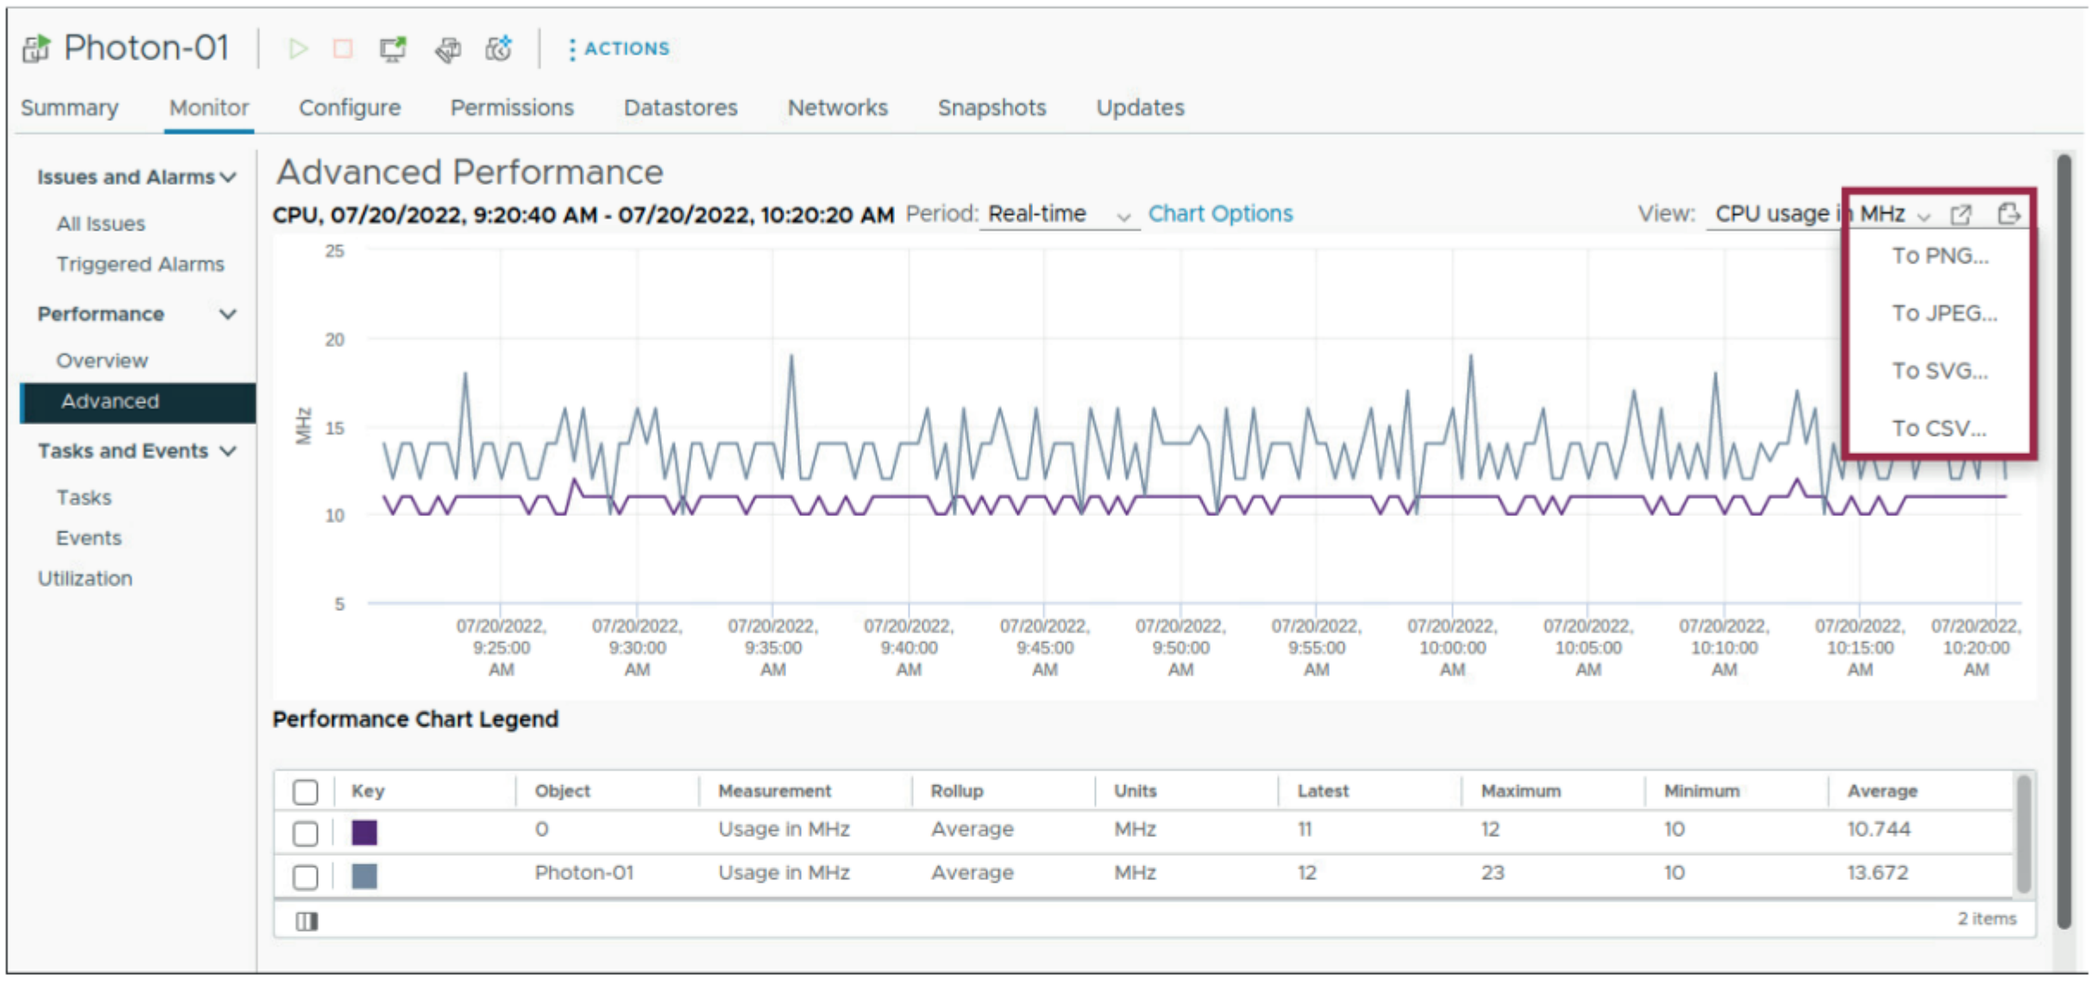

In the vSphere Client, you can save data from the advanced performance charts to a file in various graphics formats or in CSV format. When you save a chart, you select the file type and save the chart to the location of your choice.

7-20 About Objects and Counters

Performance charts graphically display CPU, memory, disk, network, and storage metrics for devices and entities managed by vCenter.

Objects are instances or aggregations of devices:

- Examples:
  - Average vCPU for the VM
  - vmhba1
  - Aggregation over all NICs

Counters identify which statistics to collect:

- Examples:
  - CPU: Used time, ready time, usage (%)
  - NIC: Network packets received
  - Memory: Memory swapped

In vCenter, you can determine how much or how little information about a specific device type is displayed. You can control the amount of information a chart displays by selecting one or more objects and counters.

An object refers to an instance for which a statistic is collected. For example, you might collect statistics for an individual CPU, all CPUs, a host, or a specific network device.

A counter represents the actual statistic that you are collecting. An example is the amount of CPU used or the number of network packets per second for a given device.

Technet24

7-21 About Statistics Types

The statistics type is the unit of measurement that is used during the statistics interval.

Statistics Type	Description	Example
Rate	Value over the current interval	CPU use (MHz)
Delta	Change from previous interval	CPU ready time
Absolute	Absolute value, independent of interval	Memory active

The statistics type refers to the measurement that is used during the statistics interval and is related to the unit of measurement.

The statistics type is one of the following:

- Rate: Value over the current statistics interval
- Delta: Change from the previous statistics interval
- Absolute: Absolute value (independent of the statistics interval)

For example, CPU usage is a rate, CPU ready time is a delta, and memory active is an absolute value.

7-22 About Rollup

Rollup is the conversion function between statistics intervals:

- Real-time (20s) statistics are rolled up to create one data point every 5 minutes.
- The result is 12 data points every hour and 288 data points every day.
- After 30 minutes, the six data points collected are aggregated and rolled up as a data point for the 1-Week time range.
- 1-Day statistics are rolled up to create one data point every 30 minutes.
- The result is 48 data points every day and 336 data points every week.
- Every 2 hours, the 12 data points collected are aggregated and rolled up as a data point for the 1-Month time range.

7-23 Rollup Types

The rollup type determines the type of statistical values returned for the counter:

Rollup Type	Conversion Function	Sample Statistic
Average	Average of data points	CPU use (average)
Summation	Sum of data points	CPU ready time (milliseconds)
Latest	Last data point	Uptime (days)
Minimum	Current or average minimum data point	Available in all counters when vCenter statistics collection is set to level 4
Maximum	Current or average maximum data point	

Data is displayed at different specificities according to the historical interval. Past-hour statistics are shown at a 20-second specificity, and past-day statistics are shown at a 5-minute specificity. The averaging that is done to convert from one time interval to another is called rollup.

The minimum and maximum values are collected and displayed only in collection level 4. Minimum and maximum rollup types are used to capture peaks in data during the interval. For real-time data, the value is the current minimum or current maximum. For historical data, the value is the average minimum or average maximum.

For example, the following information for the CPU usage chart shows that the average is collected at collection level 1 and that the minimum and maximum values are collected at collection level 4:

- Counter: Usage
- Unit: Percentage (%)
- Rollup Type: Average (Minimum/Maximum)
- Collection Level: 1 (4)

Statistics levels include summation and latest:

- Summation: The collected data is summed. The measurement displayed in the performance chart represents the sum of data collected during the interval.
- Latest: The data that is collected during the interval is a set value. The value displayed in the performance chart represents the current value.

For example, if you look at the CPU Used counter in a CPU performance chart, the rollup type is summation. So, for a given 5-minute interval, the sum of all the 20-second samples in that interval is represented.

Technet24

7-24 Review of Learner Objectives

- Describe the performance-tuning methodology
- Identify resource-monitoring tools
- Use vCenter performance charts to view performance

7-25 Lesson 2: Monitoring Resource Use

7-26 Learner Objectives

- Monitor the key factors that can affect a virtual machine's performance
- Use performance charts to view and improve performance

7-27 Interpreting Data from Tools

vCenter monitoring tools and guest OS monitoring tools provide different points of view.

CPU Usage Chart for VM

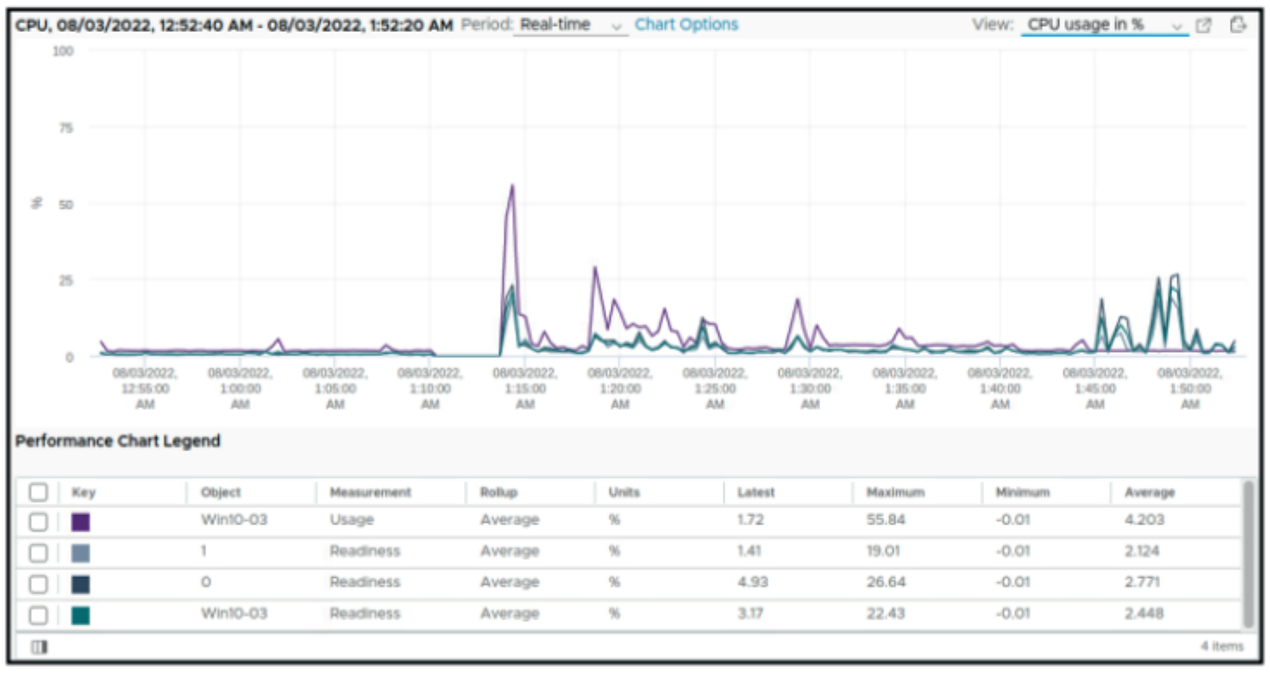

Task Manager in Guest Operating System

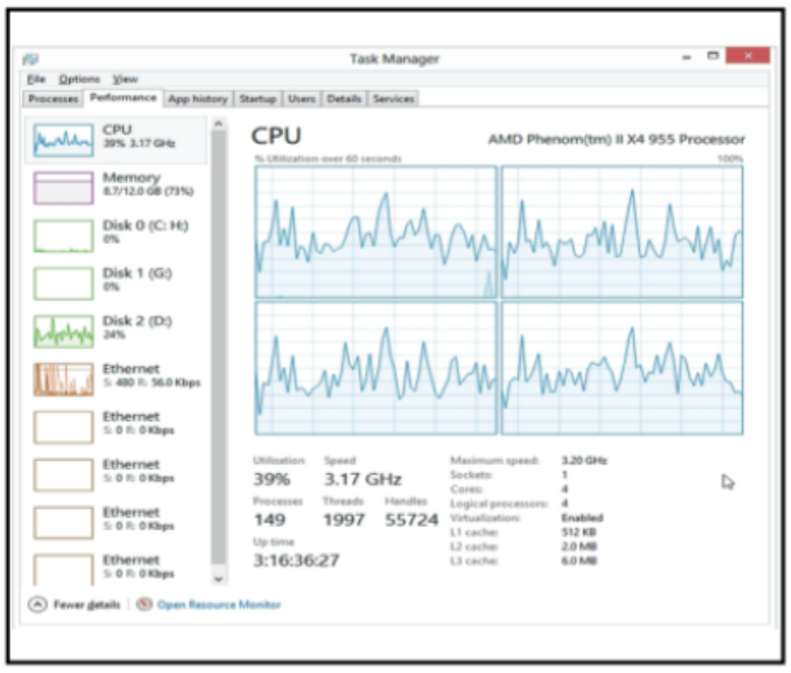

The key to interpreting performance data is to observe the range of data from the perspective of the guest operating system, the VM, and the host.

The CPU usage statistics in Task Manager, for example, do not give you the complete picture. View CPU usage for the VM and the host on which the VM is located.

Use the performance charts in the vSphere Client to view this data.

Technet24

7-28 CPU-Constrained VMs (1)

If CPU use is continuously high, the VM is constrained by the CPU. However, the host might have enough CPU for other VMs to run.

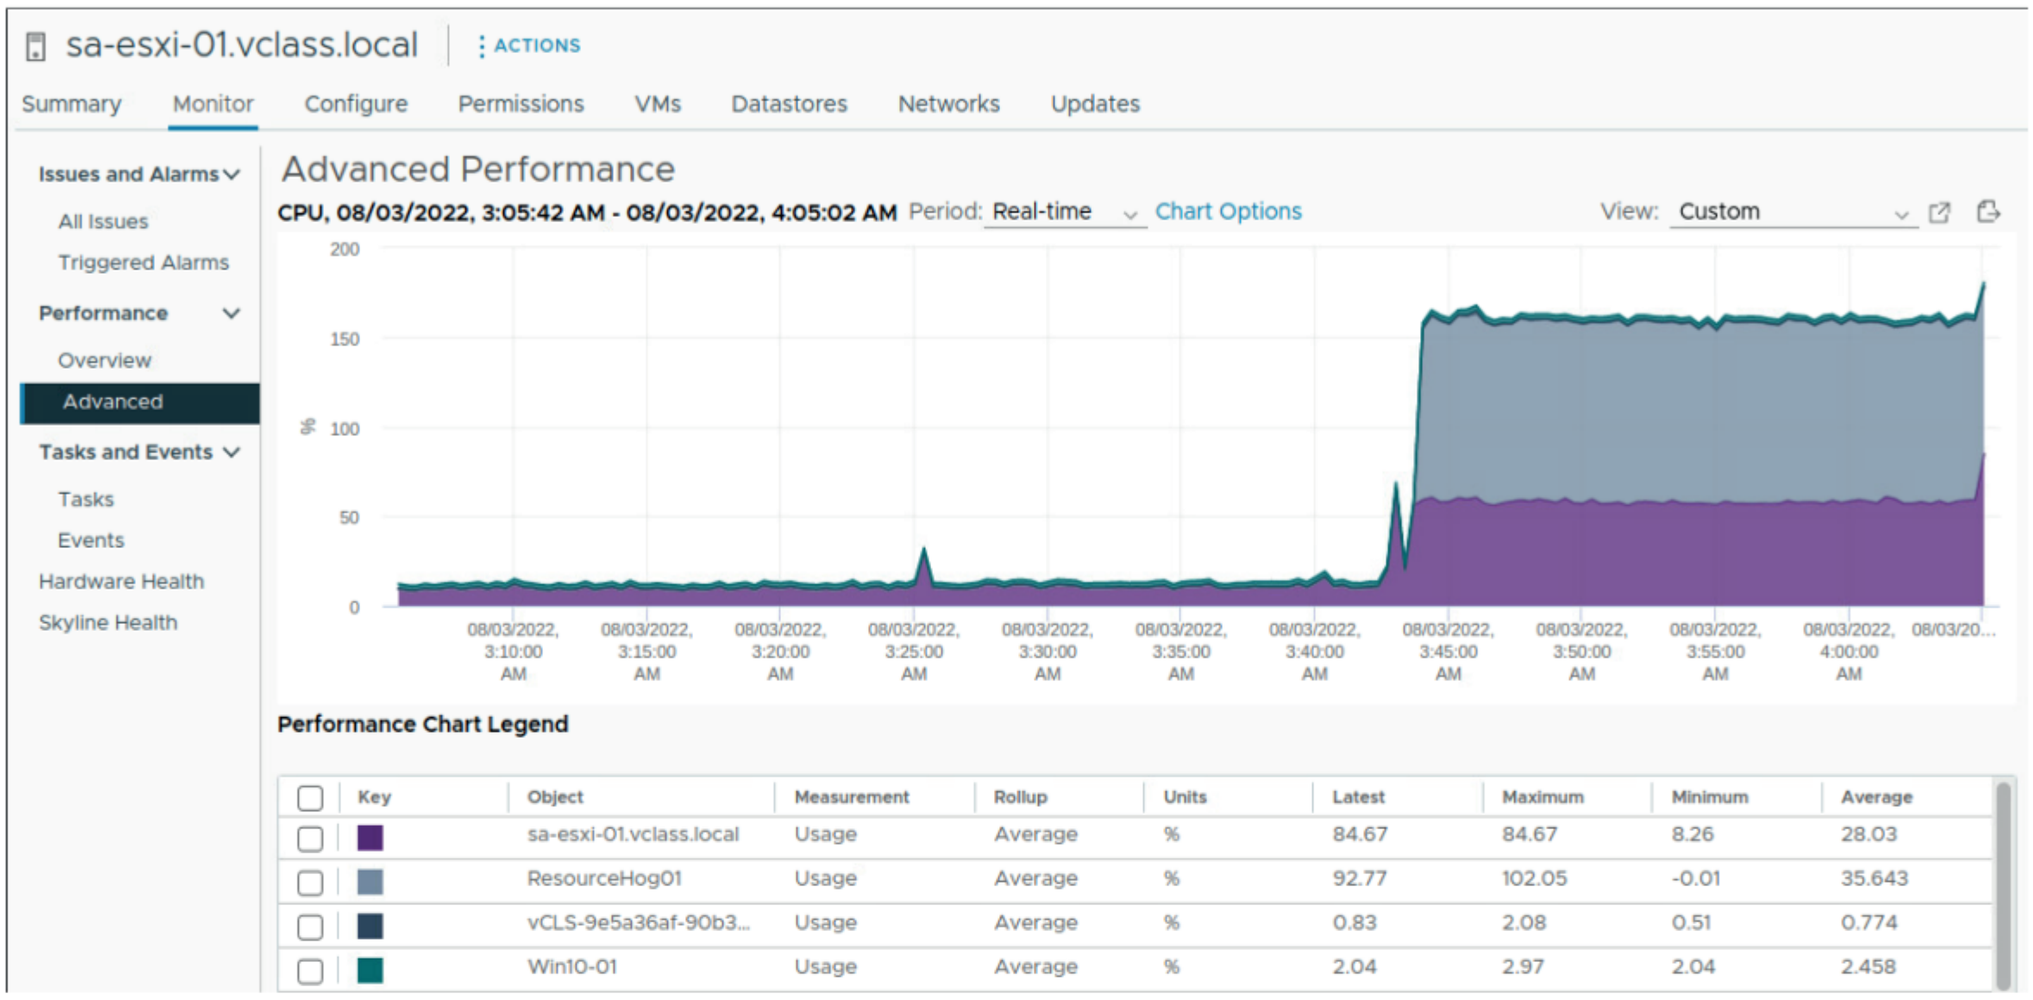

Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
	sa-esxi-01.vclass.local	Usage	Average	%	84.67	84.67	8.26	28.03
	ResourceHog01	Usage	Average	%	92.77	102.05	-0.01	35.643
	vCLS-9e5a36af-90b3...	Usage	Average	%	0.83	2.08	0.51	0.774
	Win10-01	Usage	Average	%	2.04	2.97	2.04	2.458

If CPU use is high, check the VM's CPU usage statistics. Use either the overview charts or the advanced charts to view CPU usage. The slide displays an advanced chart tracking a VM's CPU usage.

If a VM's CPU use remains high over a period of time, the VM is constrained by CPU. Other VMs on the host might have enough CPU resources to satisfy their needs.

If more than one VM is constrained by CPU, the key indicator is CPU ready time. Ready time refers to the interval when a VM is ready to execute instructions but cannot because it cannot get scheduled onto a CPU. Several factors affect the amount of ready time:

- Overall CPU use: You are more likely to see ready time when use is high because the CPU is more likely to be busy when another VM becomes ready to run.
- Number of resource consumers (in this case, guest operating systems): When a host is running a larger number of VMs, the scheduler is more likely to queue a VM behind VMs that are already running or queued.

A good ready time value varies from workload to workload. To find a good ready time value for your workload, collect ready time data over time for each VM. When you have this ready time data for each VM, estimate how much of the observed response time is ready time. If the shortfalls in meeting response-time targets for the applications appear largely because of the ready time, take steps to address the excessive ready time.

7-29 CPU-Constrained VMs (2)

Multiple VMs are constrained by the CPU if the following conditions are present:

- High virtual machine vCPU Usage
- Relatively high virtual machine vCPU readiness values for the VMs

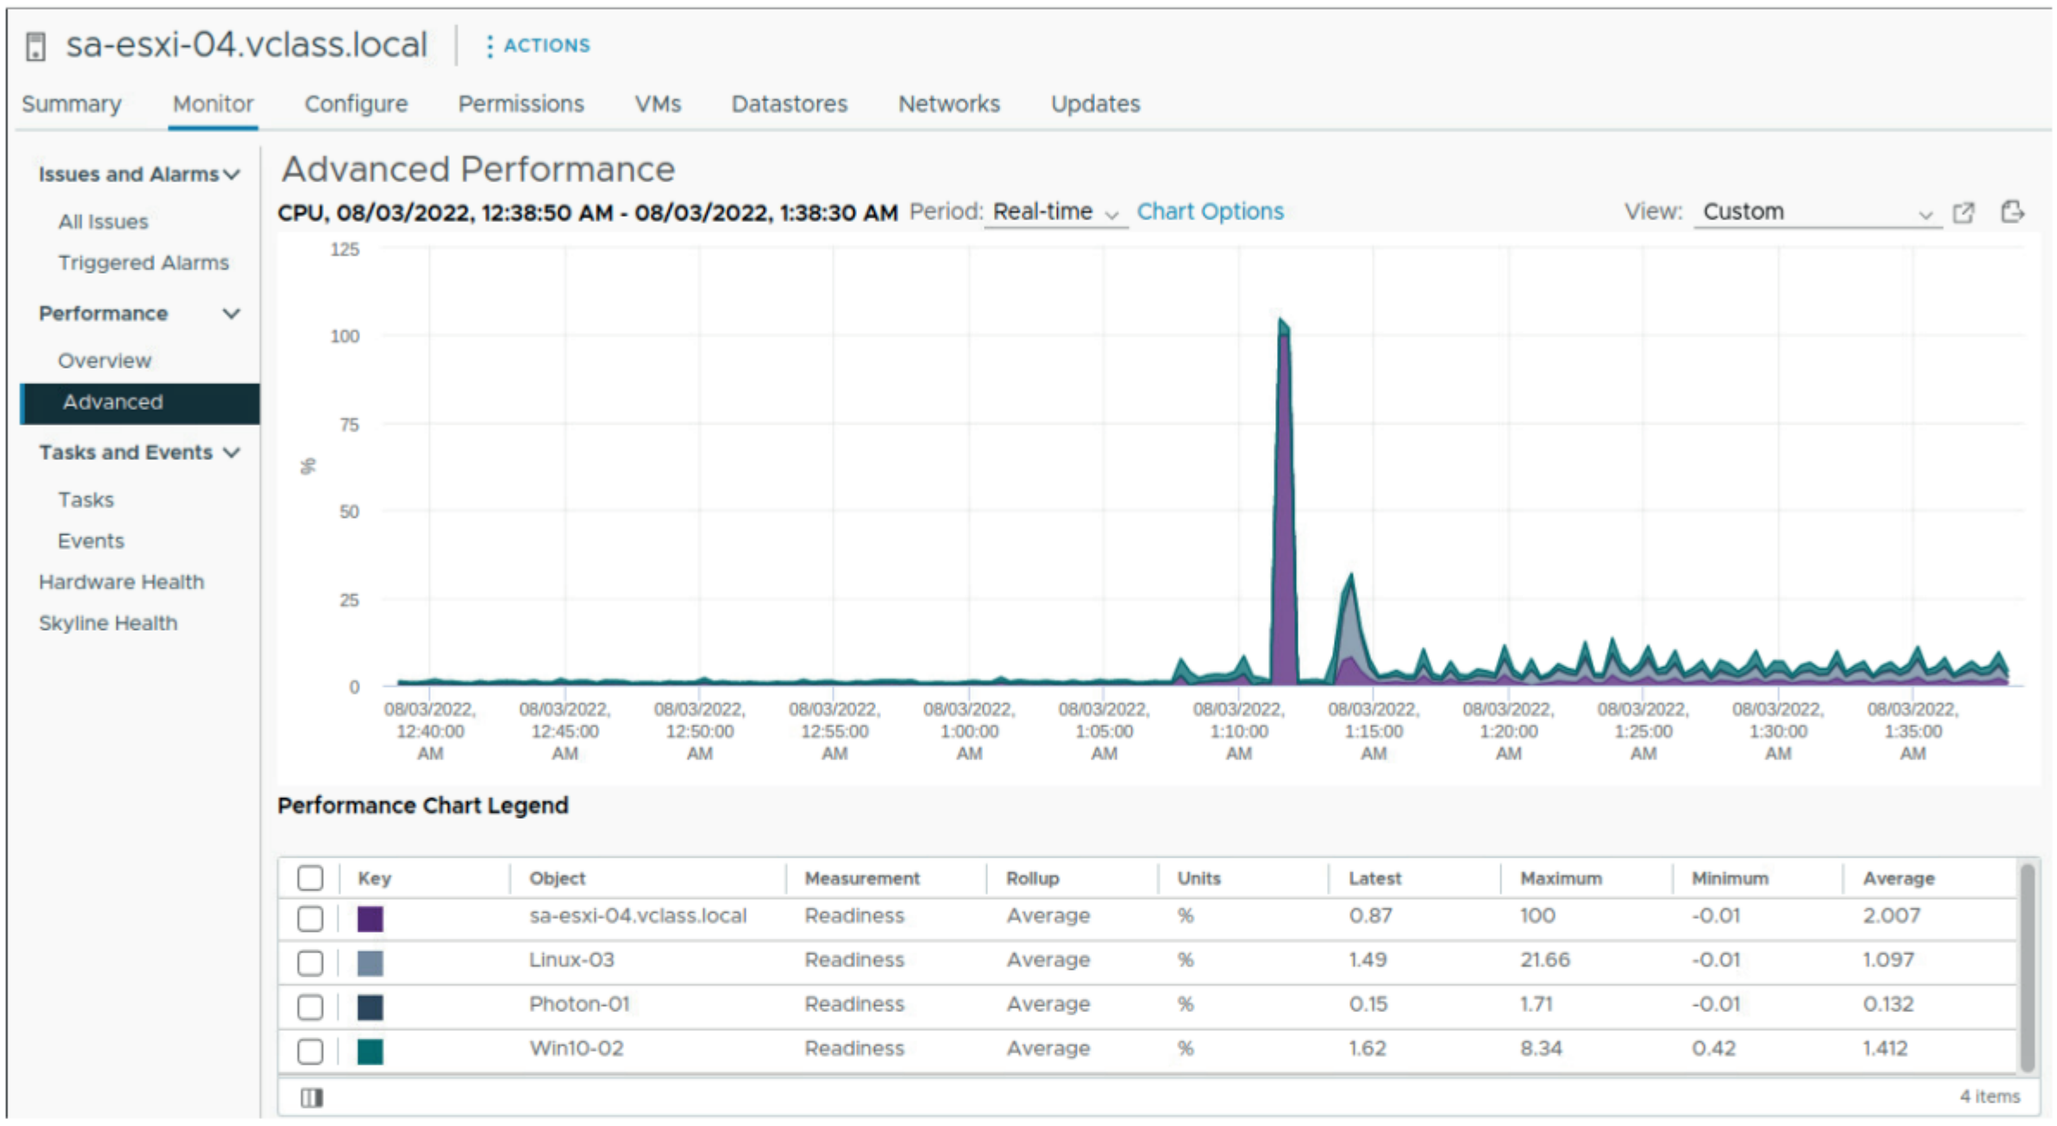

To determine whether a VM is being constrained by CPU resources, view CPU usage in the guest operating system using, for example, Task Manager or using performance charts.

If more than one VM is constrained by CPU, the key indicator is CPU readiness. CPU readiness is the percent of time that the VM cannot run because it is contending for access to the physical CPUs.

You are more likely to see readiness values when use is high because the CPU is more likely to be busy when another VM becomes ready to run.

A good readiness value varies from workload to workload.

Technet24

7-30 Memory-Constrained VMs (1)

Compare a VM's memory consumed and granted values to determine whether the VM is memory-constrained.

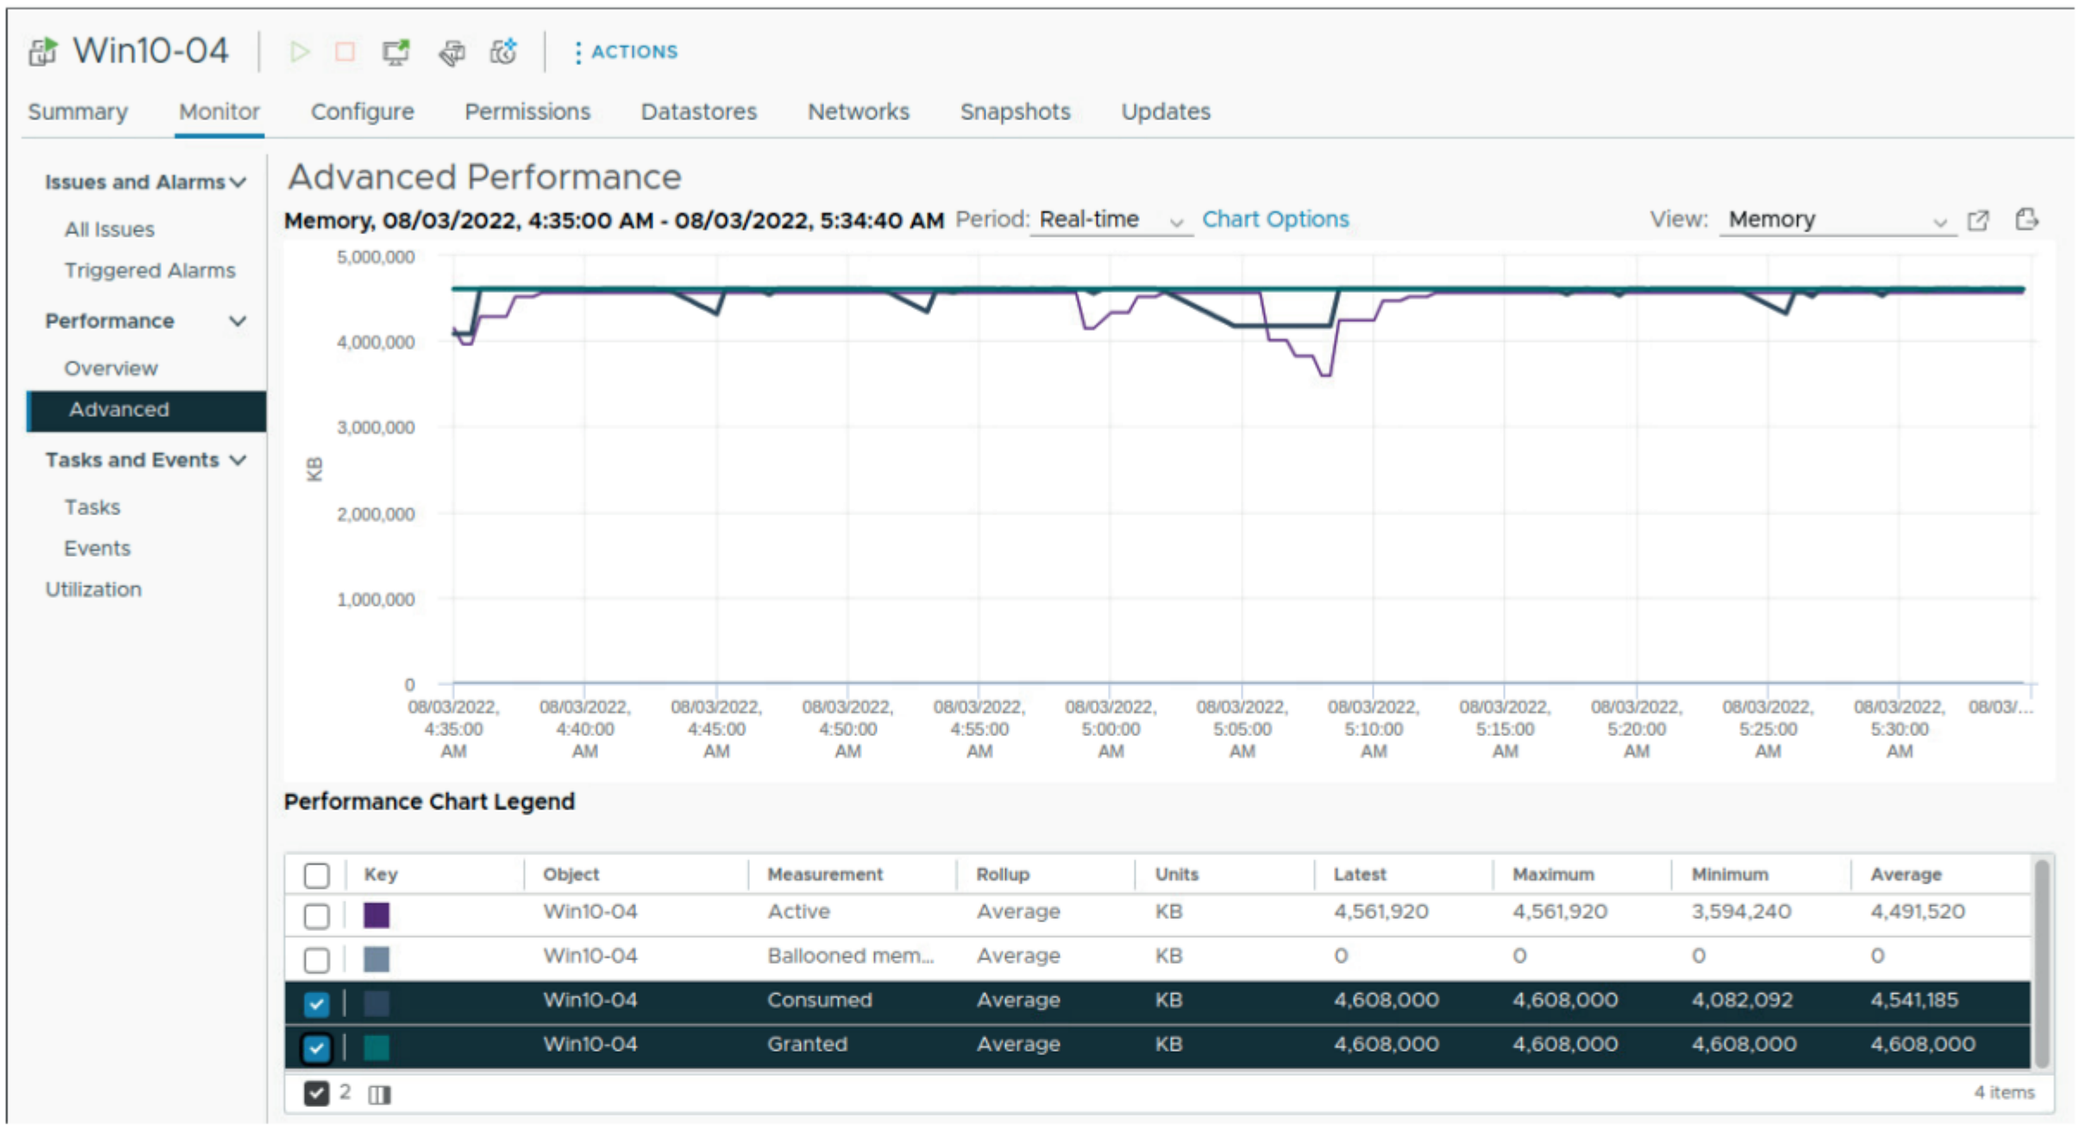

Consumed Memory is the amount of memory used on the ESXi in total. This includes memory used by the VMKernel, other vSphere Services and Management agents, plus the total consumed memory for all running virtual machines.

Granted memory is the total memory available to the VM.

consumed memory = memory granted - memory saved from page sharing

7-31 Memory-Constrained VMs (2)

If a VM actively uses its entire memory allocation, the VM might be memory-constrained, and you should consider increasing the VM's memory size.

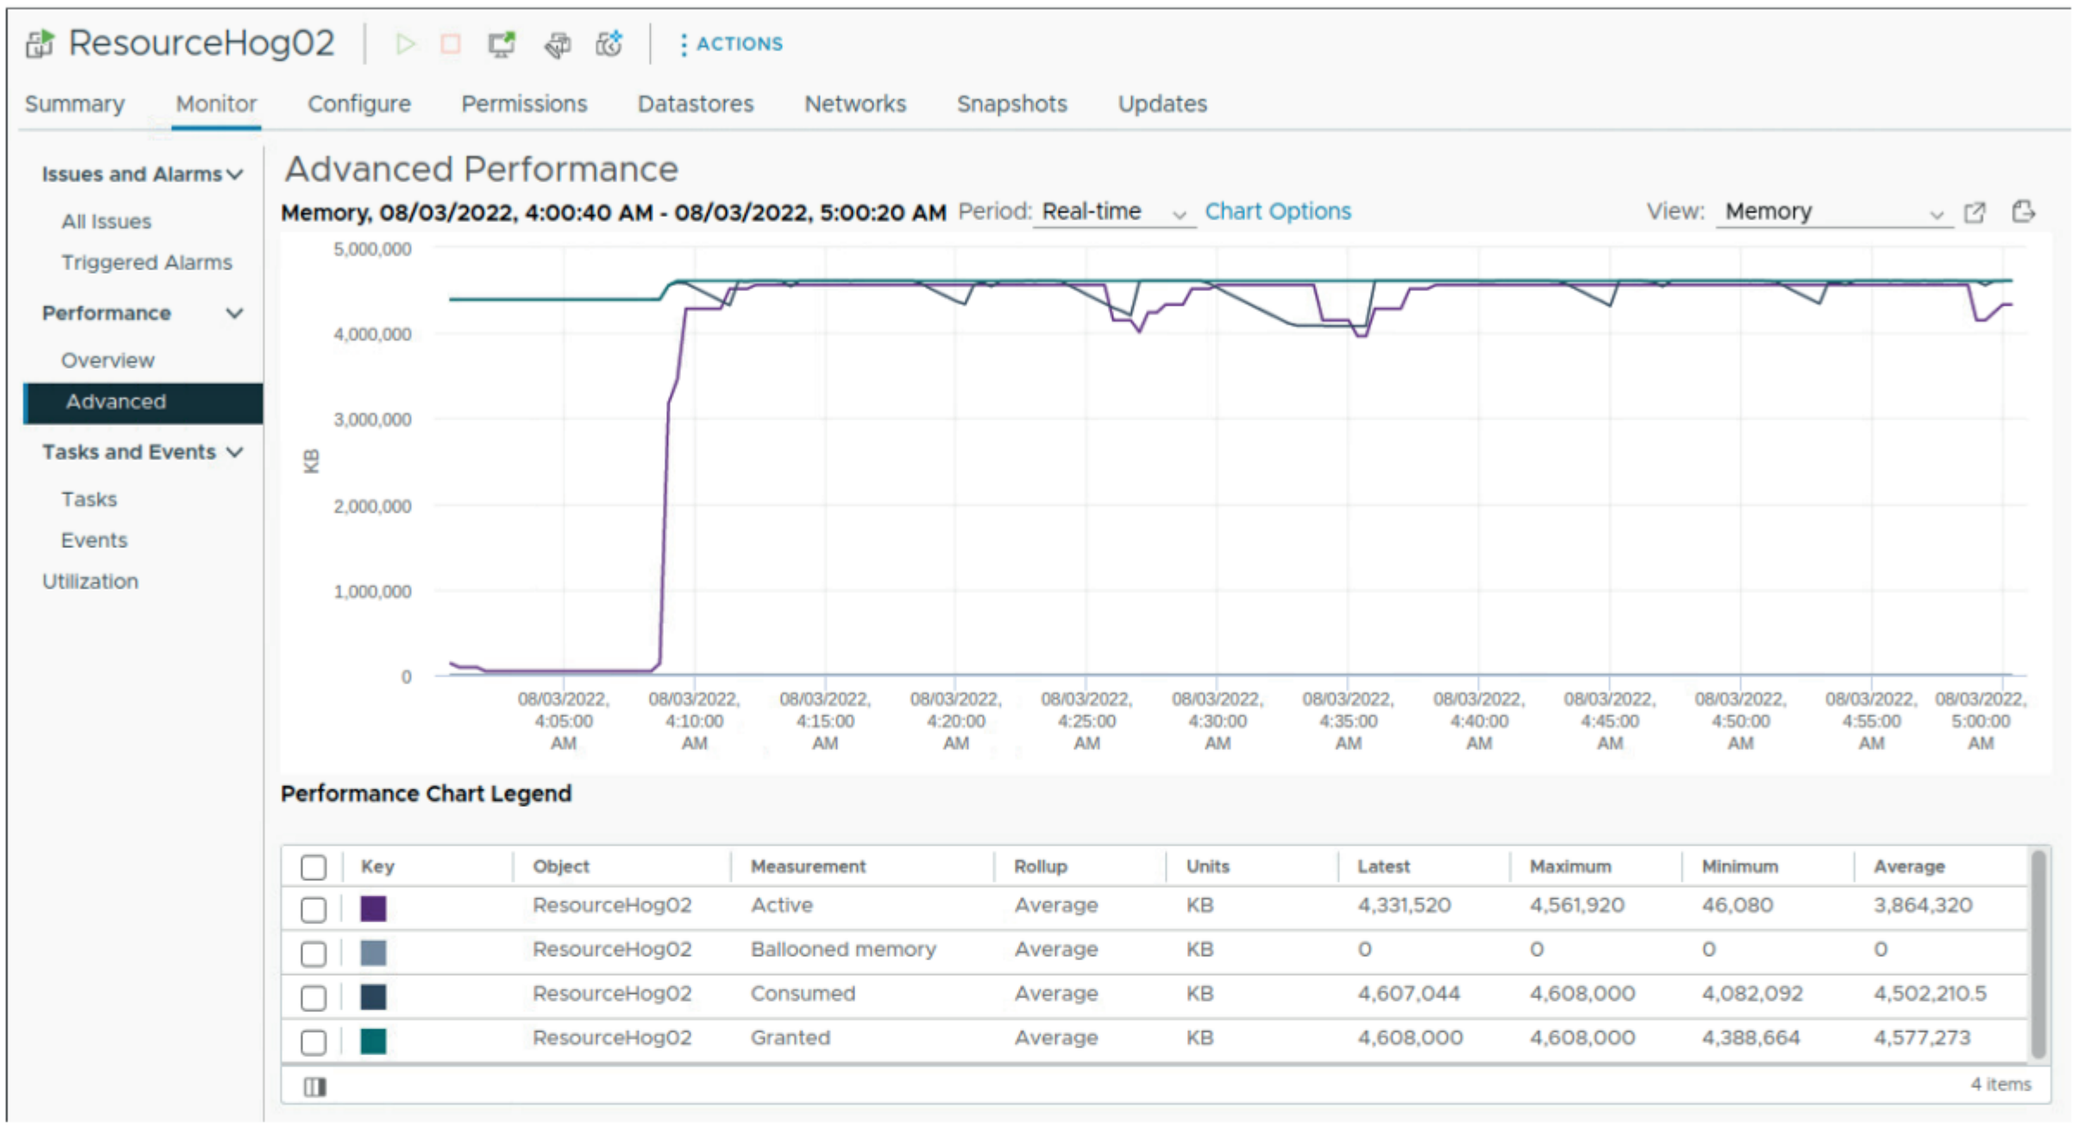

Technet24

7-32 Memory-Constrained Hosts

Any evidence of ballooning or swapping is a sign that your host might be memory-constrained.

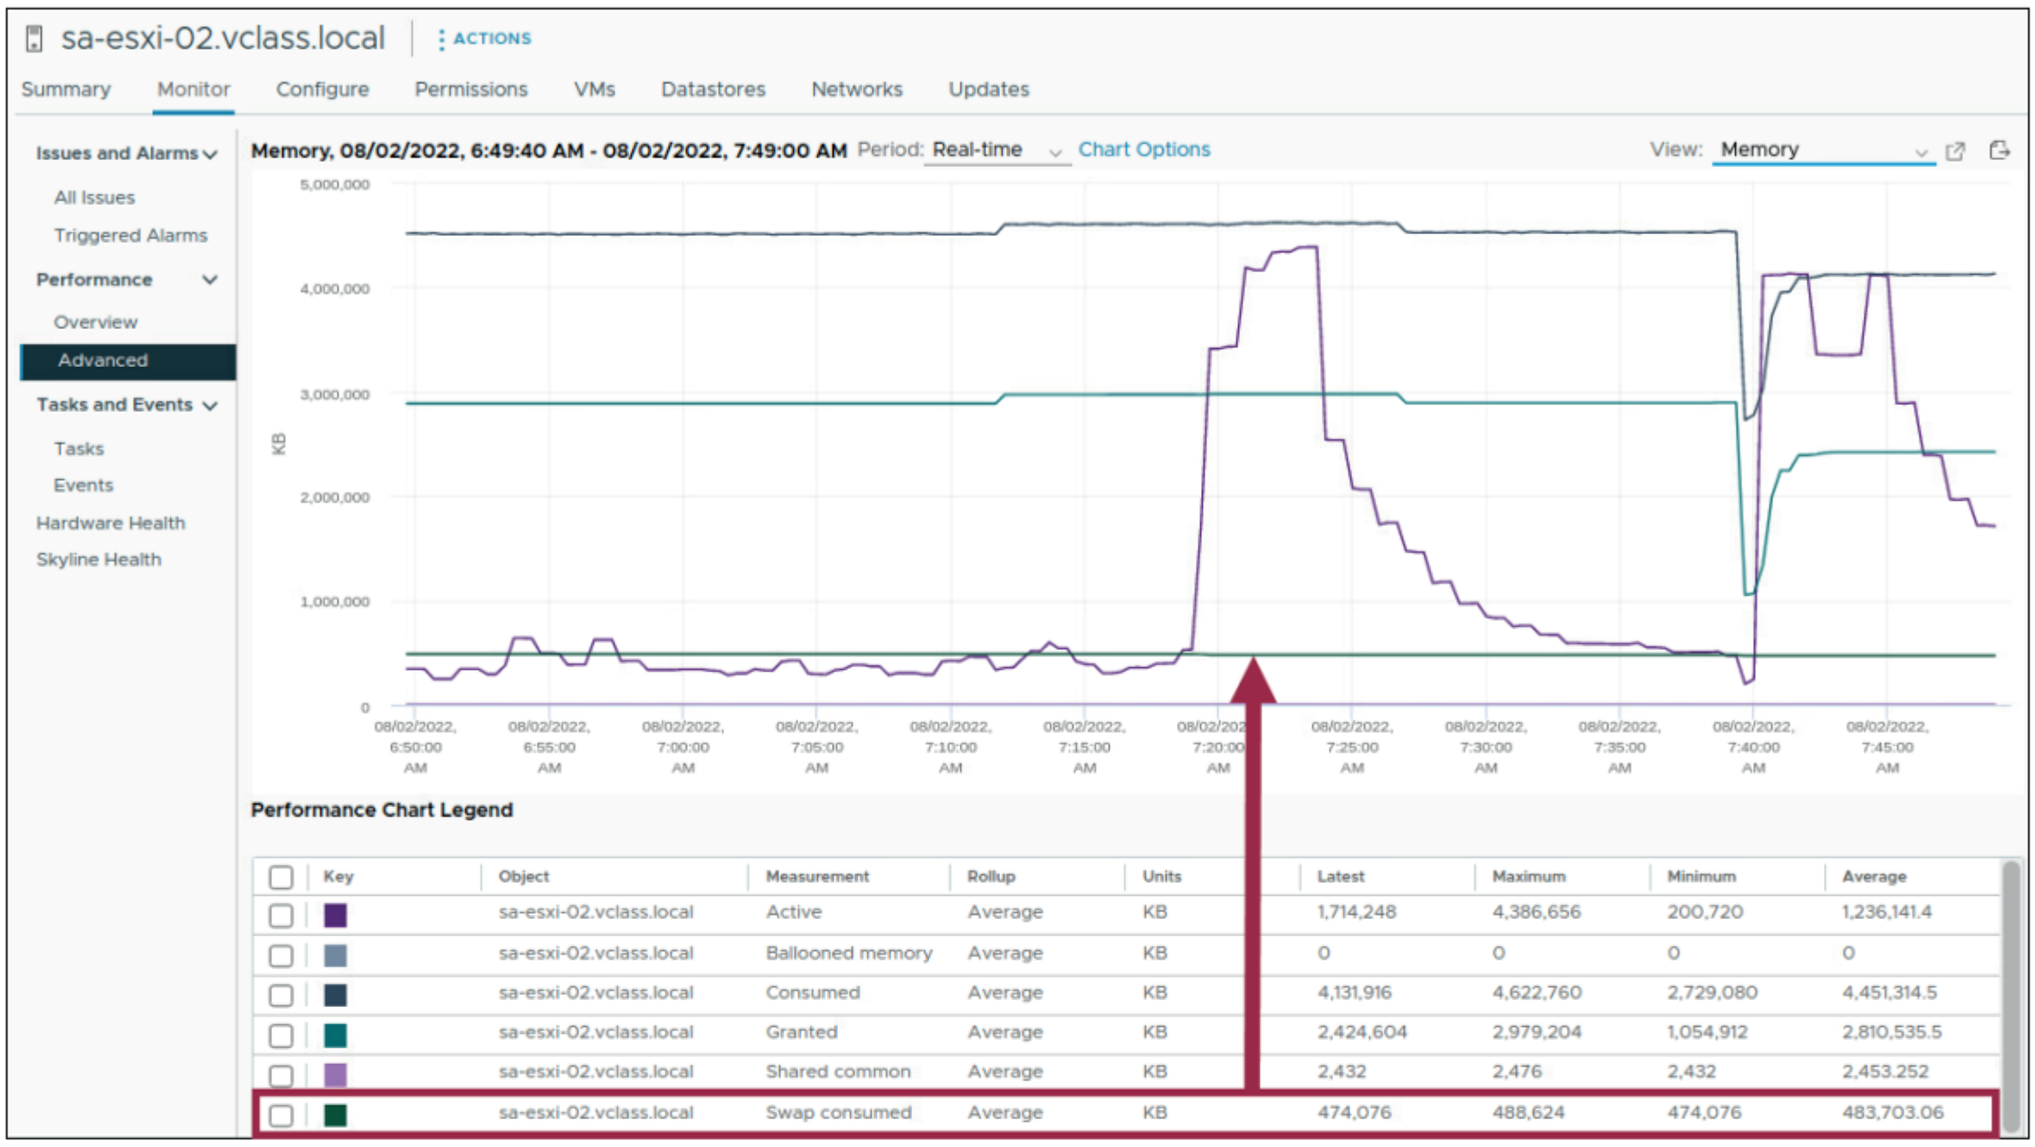

Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
	sa-esxi-02.vclass.local	Active	Average	KB	1,714,248	4,386,656	200,720	1,236,141.4
	sa-esxi-02.vclass.local	Ballooned memory	Average	KB	0	0	0	0
	sa-esxi-02.vclass.local	Consumed	Average	KB	4,131,916	4,622,760	2,729,080	4,451,314.5
	sa-esxi-02.vclass.local	Granted	Average	KB	2,424,604	2,979,204	1,054,912	2,810,535.5
	sa-esxi-02.vclass.local	Shared common	Average	KB	2,432	2,476	2,432	2,453.252
	sa-esxi-02.vclass.local	Swap consumed	Average	KB	474,076	488,624	474,076	483,703.06

You might see VMs with high ballooning activity and VMs being swapped in and out by the VMkernel. This serious situation indicates that the host memory is overcommitted to the point of contention and must be increased.

No VMs are having their memory ballooned at this point, only swapped.

7-33 Disk-Constrained VMs

Disk-intensive applications can saturate the storage or the path.

If you suspect that a VM is constrained by disk access, take these actions:

- Measure the throughput and latency between the VM and storage.
- Use the advanced performance charts to monitor throughput and latency:
  - Read rate and write rate
  - Read latency and write latency

Disk performance problems are commonly caused by saturating the underlying physical storage hardware. You can use the vCenter advanced performance charts to measure storage performance at different levels. These charts provide insight about a VM performance. You can monitor everything from the VM's datastore to a specific storage path.

If you select a host object, you can view throughput and latency for a datastore, a storage adapter, or a storage path. The storage adapter charts are available only for Fibre Channel storage. The storage path charts are available for Fibre Channel and iSCSI storage, not for NFS.

If you select a VM object, you can view throughput and latency for the VM’s datastore or specific virtual disk.

To monitor throughput, view the Read rate and Write rate counters. To monitor latency, view the Read latency and Write latency counters.

Technet24

7-34 Monitoring Disk Latency

To determine disk performance problems, monitor two disk latency data counters:

- Kernel command latency:
  - This counter is the average time that is spent in the VMkernel per SCSI command.
  - High numbers (greater than 2 milliseconds or 3 milliseconds) represent either an overworked array or an overworked host.
- Physical device command latency:
  - This counter is the average time that the physical device takes to complete a SCSI, NVMe and NVMe-of command.
  - High numbers represent a slow or overworked array, for example:
    - For spinning disks (HDDs), greater than 15 milliseconds or 20 milliseconds
    - For SSDs, greater than 3 milliseconds or 4 milliseconds

To determine whether your vSphere environment is experiencing disk problems, monitor the disk latency data counters. Use the advanced performance charts to view these statistics. In particular, monitor the following counters:

- Kernel command latency: This data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value should be 0 through 1 millisecond. If the value is greater than 4 milliseconds, the VMs on the ESXi host are trying to send more throughput to the storage system than the configuration supports.
- Physical device command latency: This data counter measures the average amount of time, in milliseconds, for the physical device to complete a SCSI or NVMe command.

Local vs SAN-based storage:

- Local storage will have shorter paths and more bandwidth and therefore lower latency.
- SANs will cover larger distances and could therefore result in larger latency.

7-35 Network-Constrained VMs

Network-intensive applications often bottleneck on path segments outside the ESXi host:

- Example: WAN links between server and client

If you suspect that a VM is constrained by the network, take these actions:

- Verify that VMware Tools is installed and that VMXNET3 is the virtual network adapter.
- Measure the effective bandwidth between the VM and its peer system.
- Check for dropped receive packets and dropped transmit packets.
- Check where packets are dropped and in which direction.

Like disk performance problems, network performance problems are commonly caused by saturating a network link between client and server. Use a tool such as Iometer, or a large file transfer, to measure the effective bandwidth.

Network performance depends on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use the advanced performance charts to examine the droppedTx and droppedRx network counter values of a VM.

In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU that is required to process the data. In some instances, large packets can result in high network latency. When network packets are small, more packets are transferred, but the network speed is slower because more CPU is required to process the data.

Technet24

7-36 Lab 16: Monitoring Virtual Machine Performance

Use the system monitoring tools to review the CPU workload:

1. Create a CPU Workload
2. Use Performance Charts to Monitor CPU Use

7-37 Review of Learner Objectives

- Monitor the key factors that can affect a virtual machine's performance
- Use performance charts to view and improve performance

7-38 Lesson 3: Monitoring vCenter and Performance

7-39 Learner Objectives

- Review vCenter components and services
- Describe the factors that influence vCenter performance
- Use vCenter tools to monitor vCenter resource use

7-40 Monitoring vCenter Services

You can use the vCenter Management Interface to monitor the health and state of the vCenter services. You can restart, start, or stop services from this interface.

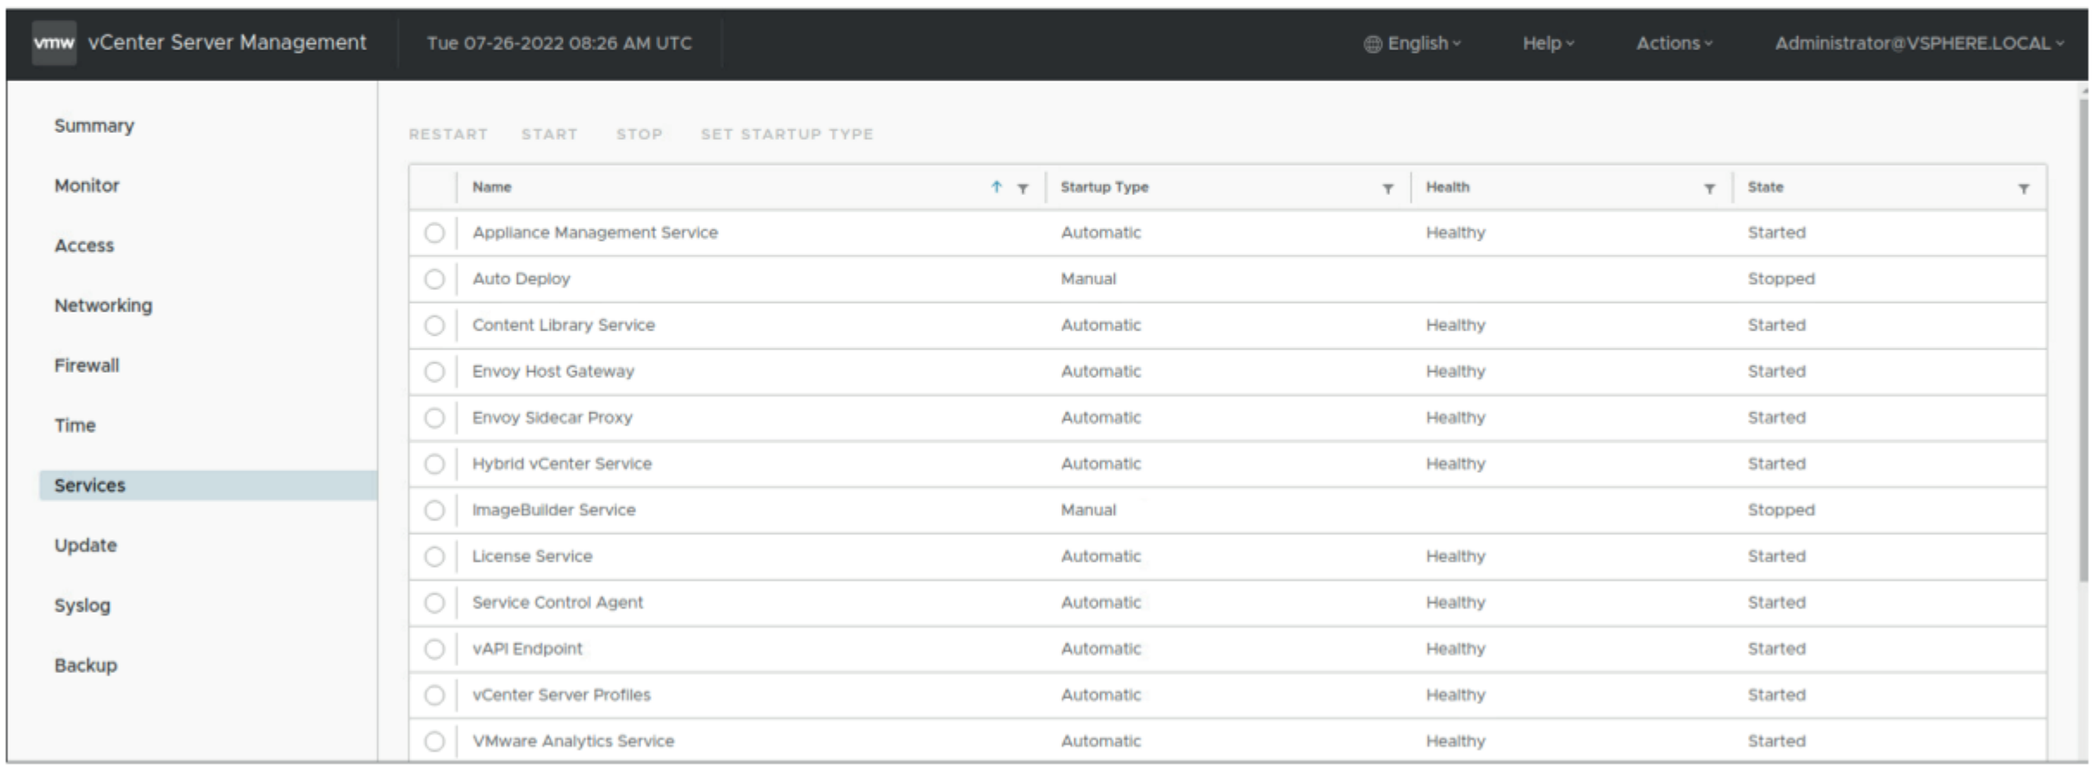

Technet24

7-41 Monthly Patch Updates for vCenter

VMware provides monthly security patches for vCenter:

- Critical vulnerability patches are delivered on a monthly release cycle.
- Important and low vulnerabilities are delivered with the next available vCenter patch or update.

You can configure the vCenter to perform automatic checks for available patches in the configured repository URL at a regular interval.

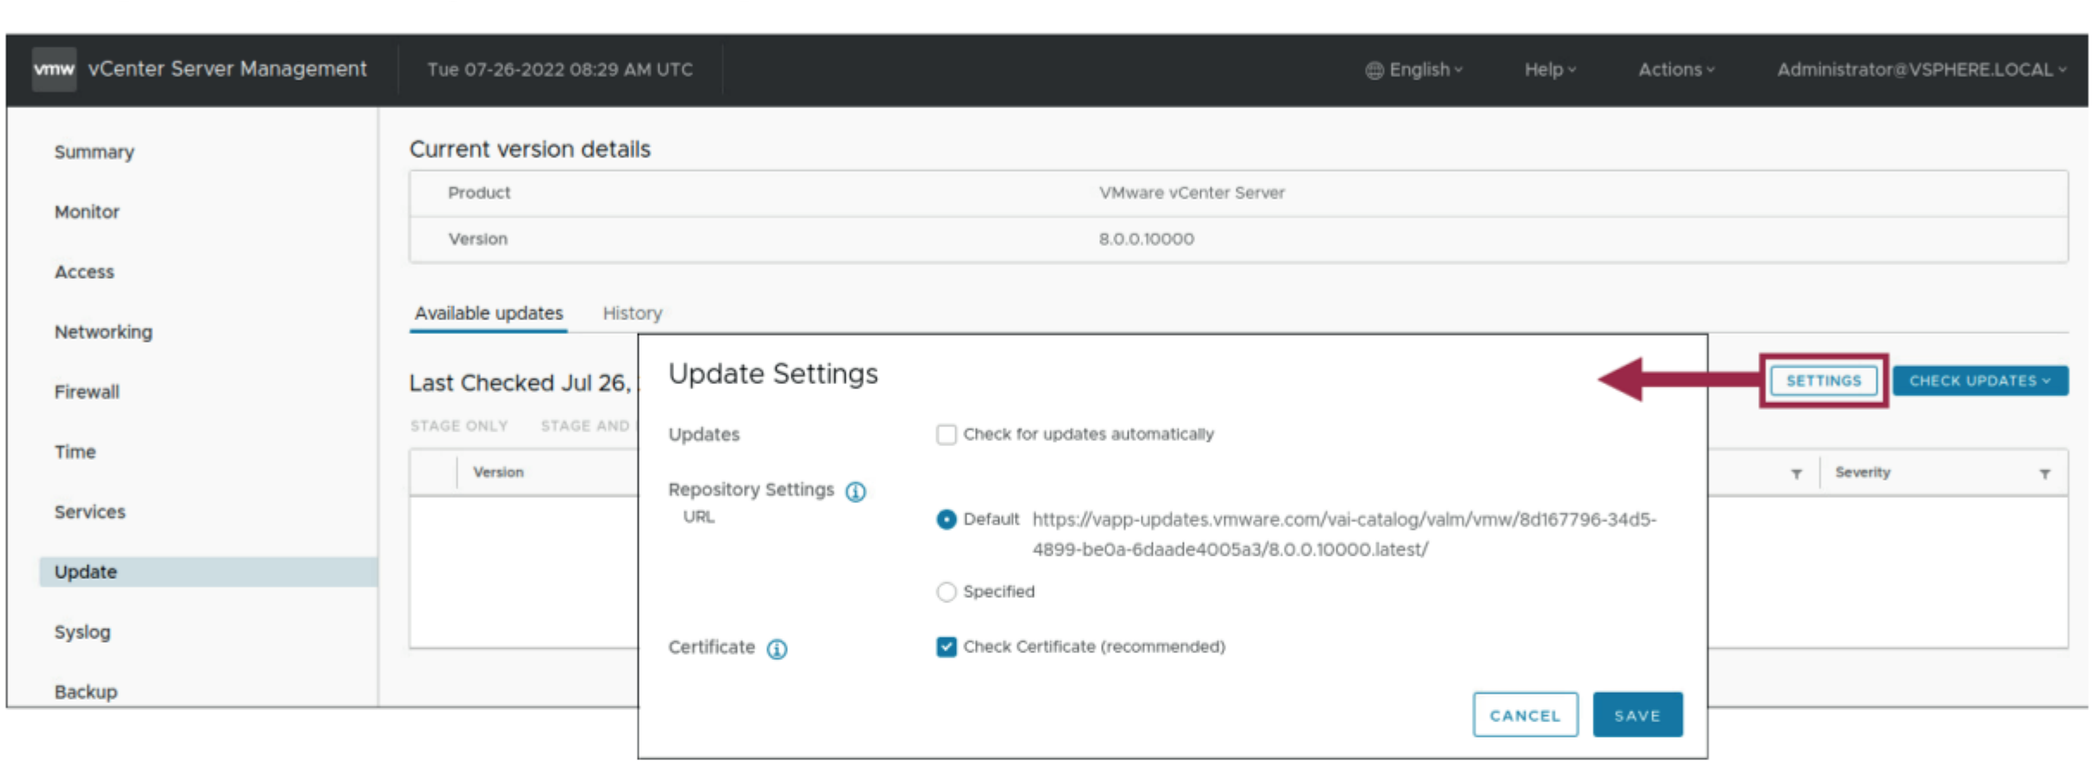

If a vCenter patch or update occurs in the same time period as the monthly security patch, the monthly security patch is rolled into the vCenter patch or update.

Checking for vCenter updates can be automated by enabling the "Check for updated automatically". You will need to initiate this process manually.

7-42 vCenter Performance Considerations

Consider the following factors when monitoring and maintaining acceptable levels of performance for vCenter and its services:

- CPU and memory performance
- Network performance
- Database performance
- User interface performance

7-43 vCenter CPU Performance

For optimal CPU performance, vCenter should not exceed 70 percent CPU usage.

If vCenter CPU usage consistently exceeds 70 percent, perform the following steps:

- Identify which processes are using the most CPU:
  - If the `vpxd` process is consuming most of the CPU, consider the following solutions:
    - Add more CPUs to vCenter VM.
    - Add another vCenter instance and balance the load across vCenter instances.
  - If a service is consuming most of the CPU, you might need to increase the heap memory size for that service.
- Examine plug-ins, extensions, or the custom API code that communicates with vCenter:
  - You might need to increase resources to accommodate these extensions for them to operate optimally.

Consistently high CPU usage can cause poor performance. Verify which processes, such as `vpxd` and `vsphere-ui`, use the most CPU. CPU usage that is consistently higher than 70 percent can cause a vCenter process to perform more queries than necessary.

The following actions might require increasing the amount of resources used by the vCenter instance beyond the guidance provided in the VMware documentation:

- Adding extensions to vCenter
- Adding code to the vSphere Client
- Adding another extension point from where to get data
- Running more services might require more CPUs and memory.

  For example, with NSX, running the vSphere ESX Agent Manager (EAM) service is known to increase demand on resources.

Enabling CPU and Memory Hot Add settings on the vCenter VM, will allow you to add CPU and/or Memory without downtime.

Technet24

7-44 vCenter Memory Performance

Monitor both CPU and memory usage in vCenter to ensure that all services have enough resources:

- If vCenter memory usage consistently exceeds 70 percent, perform the steps:
  - Verify if vCenter is swapping.
  - Verify which services have high CPU usage.
    High CPU usage might indicate excessive swapping because CPU is needed to perform the swapping operation.
- More memory might need to be added if high usage is consistent.
- If memory is overcommitted in your cluster, consider setting reservations for vCenter.

vCenter are critical components of your software-defined data center (SDDC). For maximum performance and uptime, ensure that vCenter is provisioned with enough resources to keep its systems running optimally.

7-45 Monitoring CPU and Memory Use with the VAMI

Use the virtual appliance management interface (VAMI) to monitor CPU and memory usage.

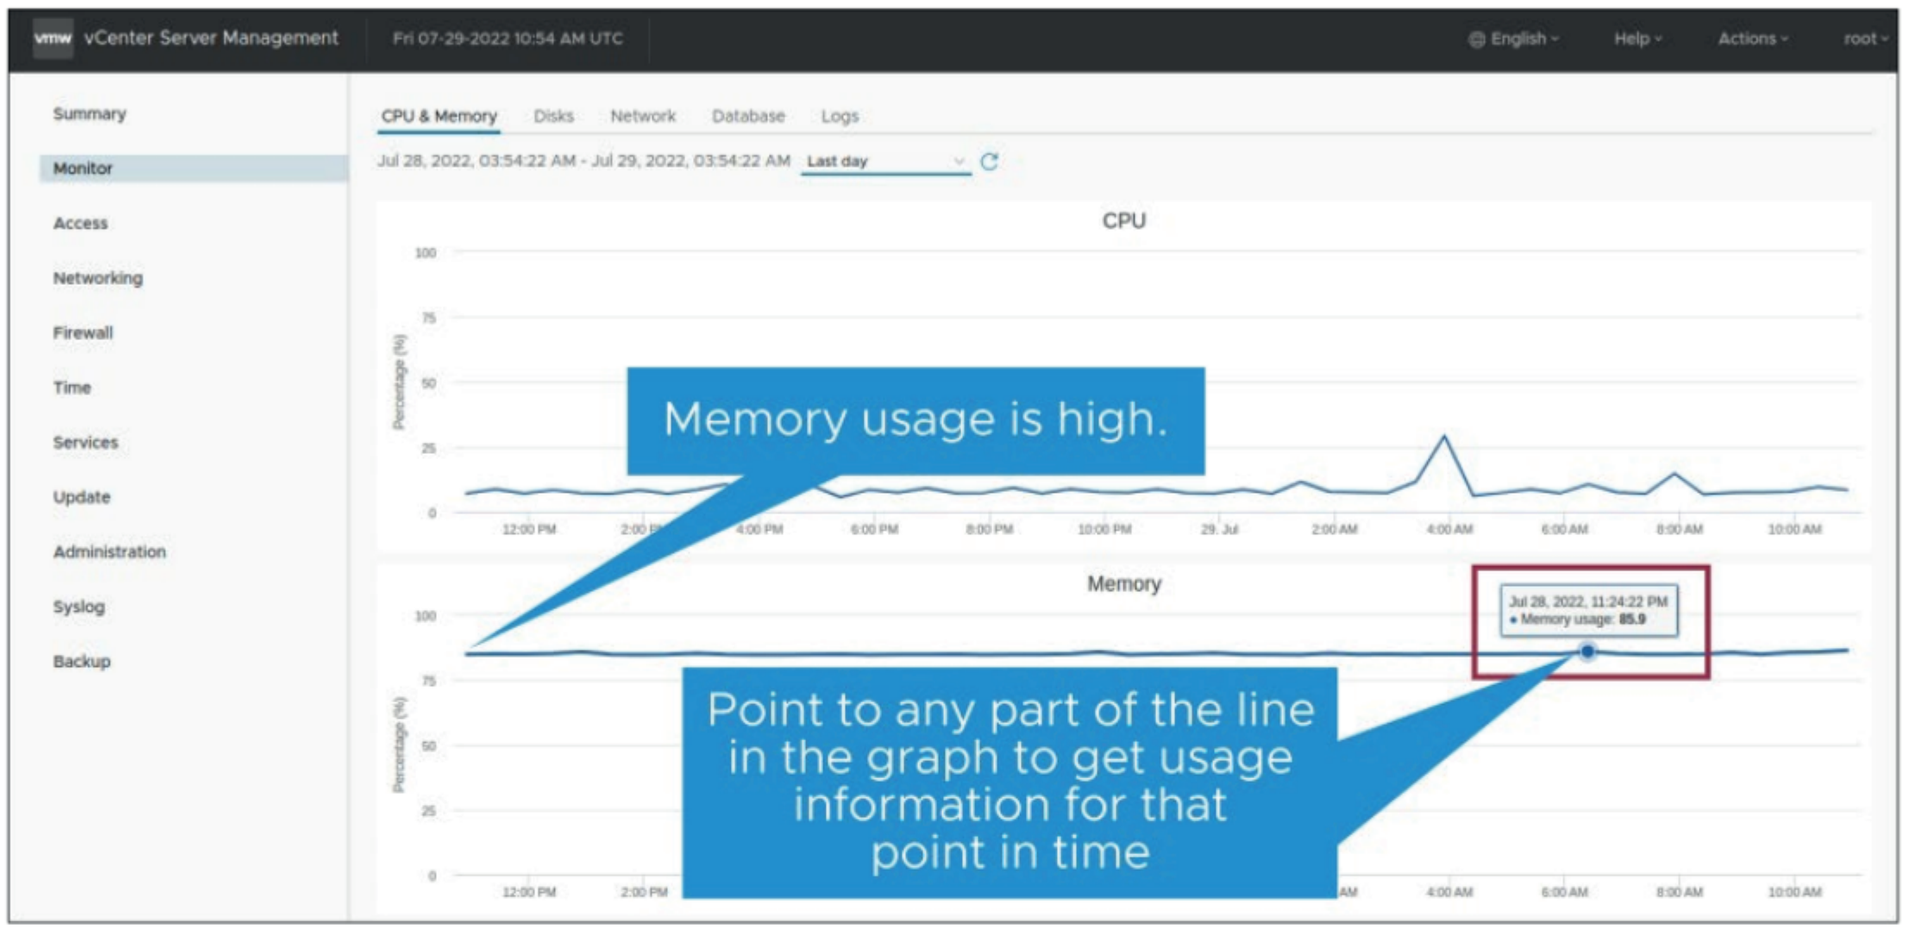

To access the virtual appliance management interface (VAMI), open a web browser and navigate to `https://<vCenter_Server_Name>:5480`. Click **Monitor** in the left pane. The **CPU & Memory** tab displays by default.

Technet24

7-46 Using vimtop to Monitor vCenter Appliance Resources

`vimtop` is a text-based tool that provides a simple way to monitor the resource usage of vCenter Appliance.

The `vimtop` tool as the following characteristics:

- Similar to the `esxtop` command
- Shows overall resource information about vCenter Server Appliance
- Shows individual services and their resource usage

To start `vimtop` you must first start the Bash shell from the vCenter Appliance command line.

```
Command> shell
Shell access is granted to root
root@sa-vcsa-01 [ ~ ]# vimtop
```

SA-VCSA-01

```
VIMTopVIMTOP VERSION: uptime:  0 days 01:01:45, load average:  0.23,  0.20,  0.34
Tasks:    36 all,    1 run,   35 sleep,    0 block,    0 stop,    0 zombie
CPUs:  29.2%us   6.8%sy   0.0%ni 359.3%id   0.0%wa   0.0%hi   0.0%si   0.0%st  8380MHz
 CPU      %USR     %SYS      %NI      %IDL      %IOW      %IRQ     %SIRQ    %STEAL       MHZ
 3       14.61     1.95      0.00     83.74     0.00      0.00      0.00      0.00   2095.08
 2        6.82     0.97      0.00     89.59     0.00      0.00      0.00      0.00   2095.08
 1        4.87     2.92      0.00     90.56     0.00      0.00      0.00      0.00   2095.08
Memory:   14347908KB total   11400408KB used     134144KB free     91812KB buffs
Swap:     26206204KB total      90368KB used   26115836KB free   2721544KB cached
Vmalloc:  33554431MB total          0MB used   33554431MB free         0MB chunk
(Enter),(H)elp,(P)ause,(S)et period,(W)rite config,P(r)ocesses,Dis(k)s,Netw(o)rk,(Q)uit
 PID      NAME                          THREADS     %CPU      MHZ     %MEM       MEM       VIRT
                                        num -1        %      MHz        %        MB         MB
 3413     Java Process                      104    18.50      387     5.05    706.89    3136.86
 3328     Trust Manager                      73     3.90       81     2.75    385.81    2643.81
 2171     Directory Service                  36     2.92       61     0.44     61.26   22751.96
 41777    VIM Top                             8     1.95       40     0.37     51.86     524.02
 2104     Authentication Framework           57     1.95       40     0.20     27.95    2571.49
 3377     vSphere Profile-Driven St         236     0.97       20     3.81    533.41    3181.97
 8941     vSphere H5 Web Client             254     0.97       20     6.51    911.50    3531.73
 1825     Observability                      56     0.97       20     0.37     51.84    2393.35
 9159     Certificate Manager                43     0.97       20     2.75    385.14    2713.29
 4467     Licensing Manager                  57     0.00        0     2.01    282.29    2608.66
 6304     Update Manager                     30     0.00        0     5.55    777.13    2333.68
 2992     Java Process                       83     0.00        0     2.06    289.02    2535.98
 7199     Update Manager Java subpr          29     0.00        0     2.43    340.17    6790.14
 4737     vCenter Server                    193     0.00        0     2.95    413.86    2453.39
 9055     Java Process                      323     0.00        0     2.15    301.45    2606.58
 4670     Service Control Agent              49     0.00        0     1.53    214.89    2516.45
 5056     vCenter Services                  387     0.00        0     4.59    643.66    3651.15
 8552     Observability vApi                 32     0.00        0     1.58    221.59    2509.92
 8820     Hybrid vCenter Service             46     0.00        0     2.05    287.82    2536.38
 1748     VMware Virtual Machine To           4     0.00        0     0.06      7.99     241.92
 2578     ESXi Agent Manager                 96     0.00        0     2.07    290.25    2587.73
 2870     Appliance Management Serv           4     0.00        0     1.26    176.32     463.35
 3283     vAPI Endpoint                      66     0.00        0     3.95    553.83    2832.83
```

To access `vimtop` , enter **`vimtop`** at the Bash shell command prompt.

To see the `vimtop` help menu, enter **`h`**.

To quit `vimtop`, enter **`q`**.

7-47 vimtop Commands and Columns

By default, `vimtop` displays the Processes window.

You can add additional statistics to the display by entering **c** and selecting the desired counters.

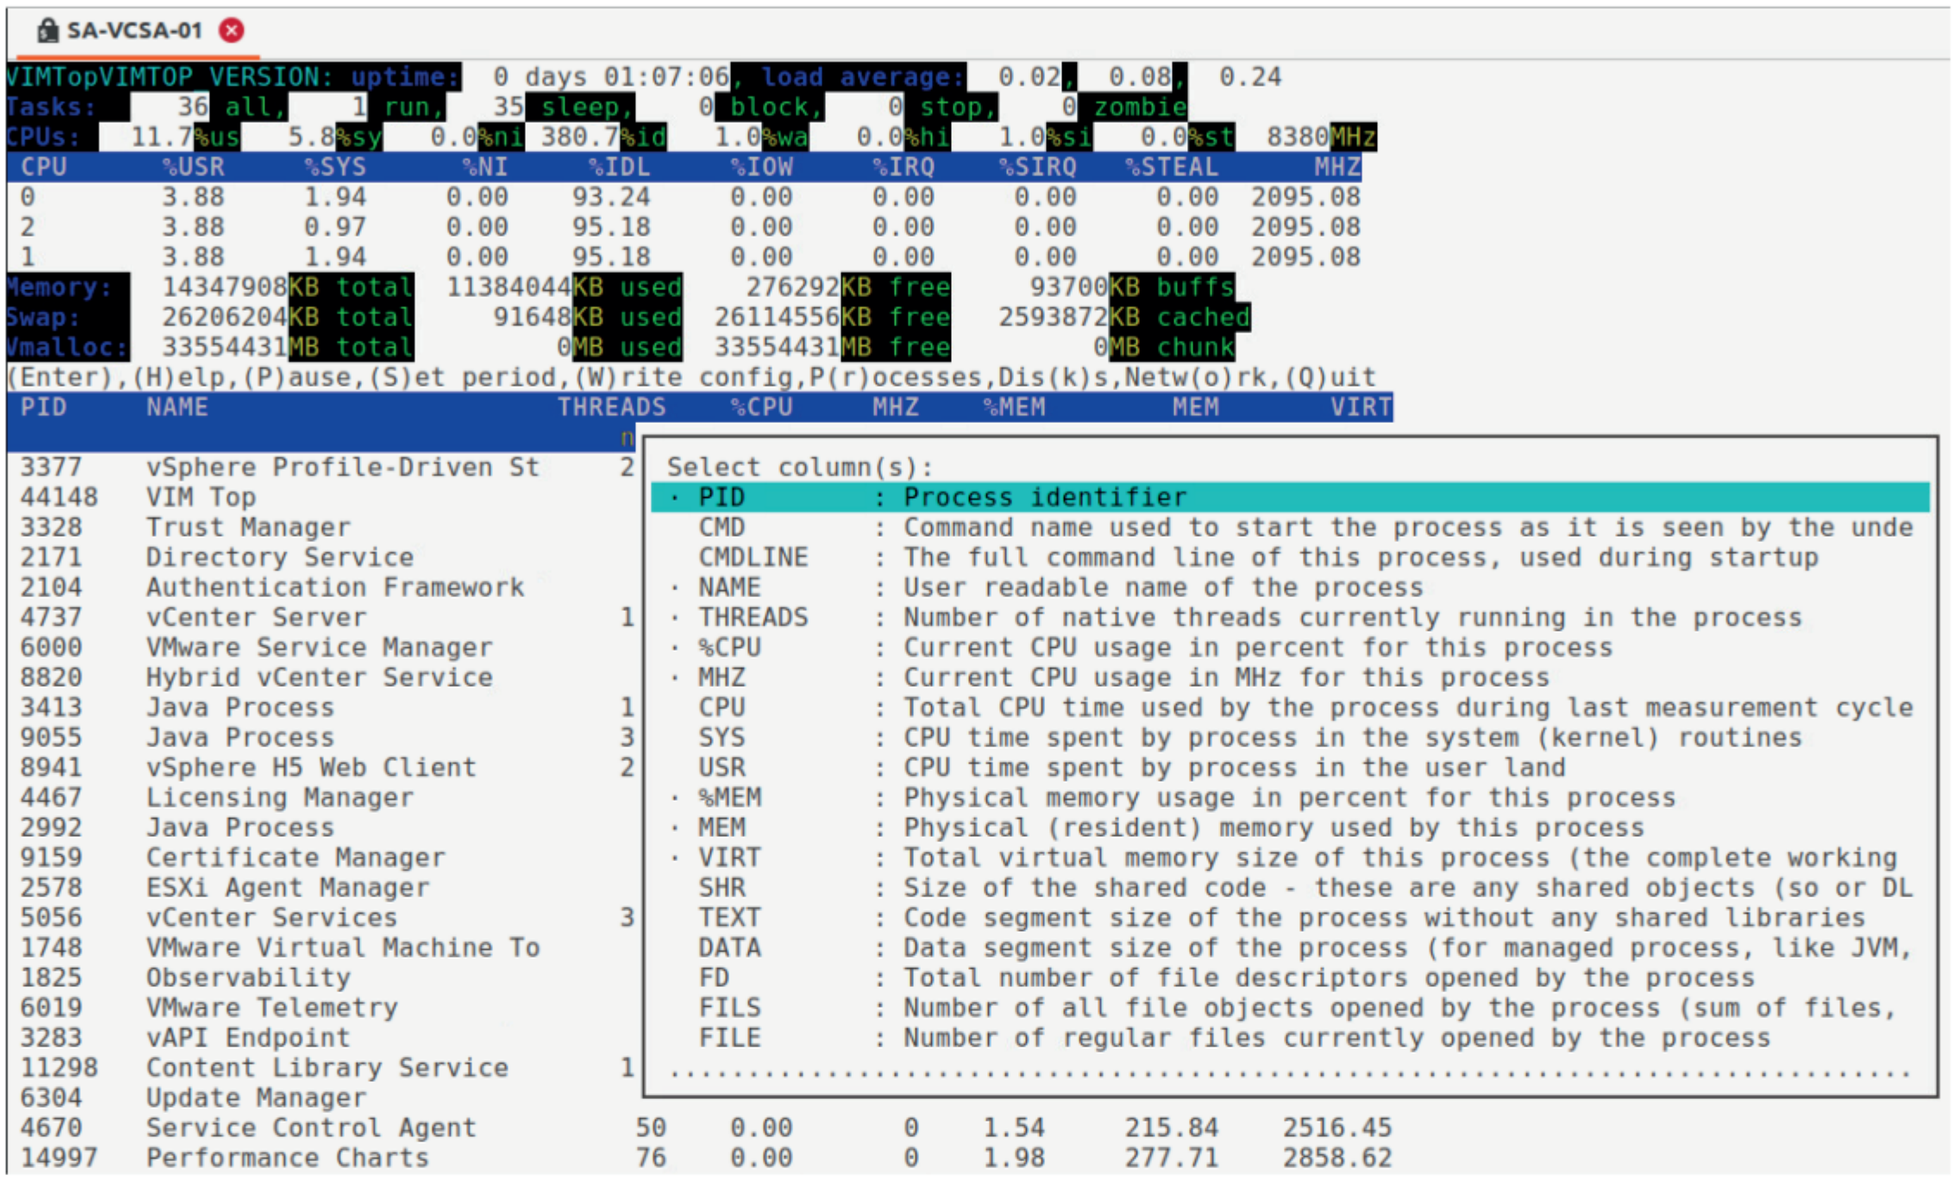

To add a metric (or column) to the `vimtop` display, perform the following steps:

1. Enter **c** to display the column selection pane.
2. Use the up and down arrow keys to highlight a metric in the column selection panel.
3. Press the space bar to select the metric.
   A tilde (**~**) appears to the left of the metric name.
4. Press the Esc key to exit the column selection panel.

Technet24

7-48 Monitoring CPU Use with vimtop

The Processes window shows overall CPU metrics and individual process CPU metrics.

To return to this window, enter **r**.

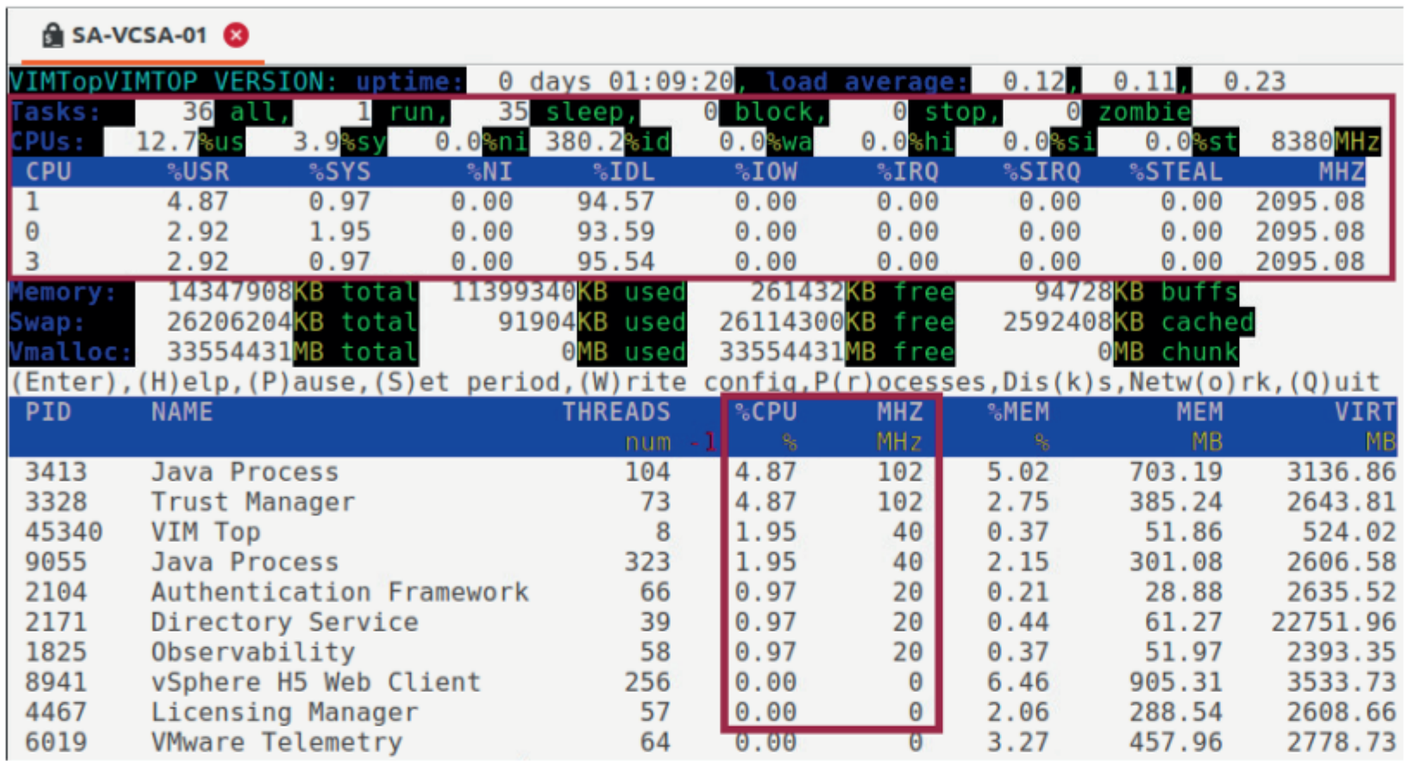

You can monitor aggregate CPU use from the overview pane or examine CPU use of individual process in the task pane.

Some columns that display include:

- Name: User-friendly name of the process (service).
- %CPU: Current CPU use (in percentage) for this process. If the value is 100 percent, the process is using an entire core.
- MHZ: Current CPU use (in megahertz) for this process.

7-49 Activity: CPU and Memory Use (1)

As a general guideline, you should keep CPU and memory use by a vCenter instance below what percentage?

◊ ❒ 60%

◊ ❒ 70%

◊ ❒ 80%

◊ ❒ 90%

7-50 Activity: CPU and Memory Use (2)

As a general guideline, you should keep CPU and memory use by a vCenter instance below what percentage?

◊ ❒ 60%

◊ ✓ 70%

◊ ❒ 80%

◊ ❒ 90%

7-51 Monitoring Memory Use with vimtop

The Processes window also displays overall memory and swap metrics. This window displays memory use for individual vCenter Server services.

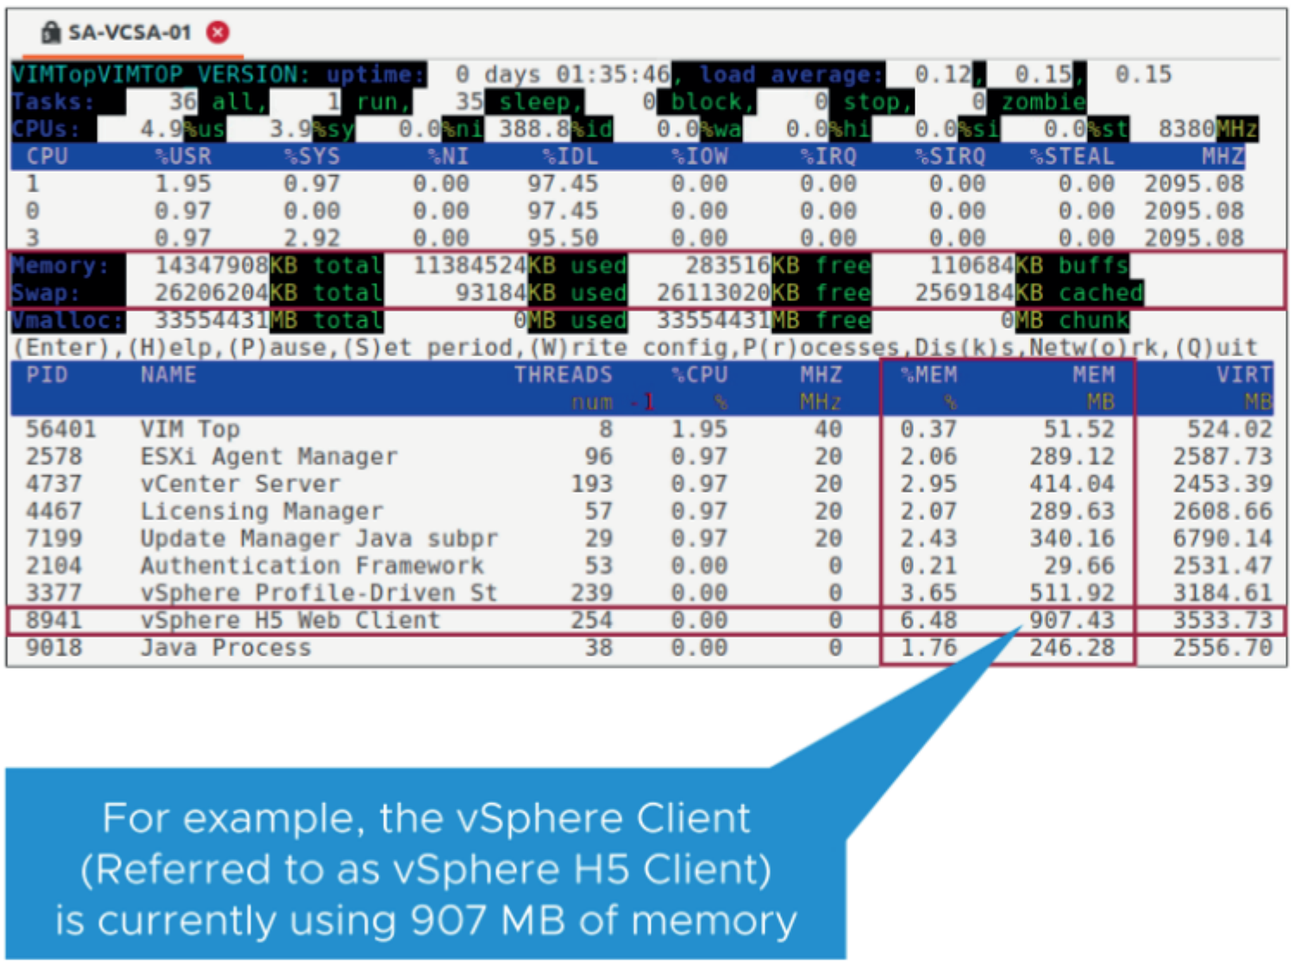

For example, the vSphere Client (Referred to as vSphere H5 Client) is currently using 907 MB of memory

Some columns that display include:

- %MEM: Physical memory use (in percentage) for this process.
- MEM: Physical or resident memory (in MB) used by this process.
- VIRT: Total virtual memory size (in MB). This value represents the complete working set, which includes resident and swapped memory.

Technet24

7-52 Viewing Service Memory Allocation with cloudvm-ram-size

The `cloudvm-ram-size` command shows the amount of memory allocated to the vCenter Server services.

The `cloudvm-ram-size -l` command displays the memory allocated (in MB) for each service.

For example, the vSphere Client service (referred to as vSphere-ui) is allocated 928 MB of memory

```
root@sa-vcsa-01 [ ~ ]# cloudvm-ram-size -l
OS                      =     2299
vmdird                  =       24
vmware-netdumper        =       20
vmware-rbd-watchdog     =      100
vmware-imagebuilder     =       50
applmgmt                =      208
vmware-envoy            =       81
vmware-envoy-hgw        =       99
vmware-rhttpproxy       =       25
vmafdd                  =       24
vmware-vsan-health      =      100
vsphere-ui              =      928
vmware-vapi-endpoint    =      402
vmware-cis-license      =      196
vmware-eam              =      175
lookupsvc               =      119
vmware-vpxd-svcs        =     1063
vmware-perfcharts       =      418
vmware-sca              =      128
vmware-sps              =      614
vmware-stsd             =      589
vmware-content-library  =      491
vmcad                   =      124
vmware-vpostgres        =      864
vmware-postgres-archiver =        5
vmware-vdtc             =       35
vmware-vpxd             =     1858
vmware-vsm              =      165
vmonapi                 =       15
vmware-vmon             =        5
vmware-updatemgr        =      362
vmware-statsmonitor     =       10
vcha                    =       46
vmware-hvc              =      120
vmware-trustmanagement  =      131
vmware-certificatemanagement =      274
vmware-certificateauthority =      100
vmware-pod              =      200
observability           =       80
observability-vapi      =      115
vmware-infraprofile     =      206
```

7-53 Viewing Service Heap Size

vCenter Server consists of several Java services, such as the vSphere Client service (vsphere-ui). A Java service uses a heap (pre-reserved area of memory) to store data objects.

Proper heap size is important for Java services to perform optimally.

For example, the vSphere Client might need more heap size in these scenarios:

- Multiple vCenter Server instances in enhanced linked mode
- Many extensions and plug-ins
- Large inventories

The following command displays the heap size for a specific service:

- `cloudvm-ram-size -J service_name`

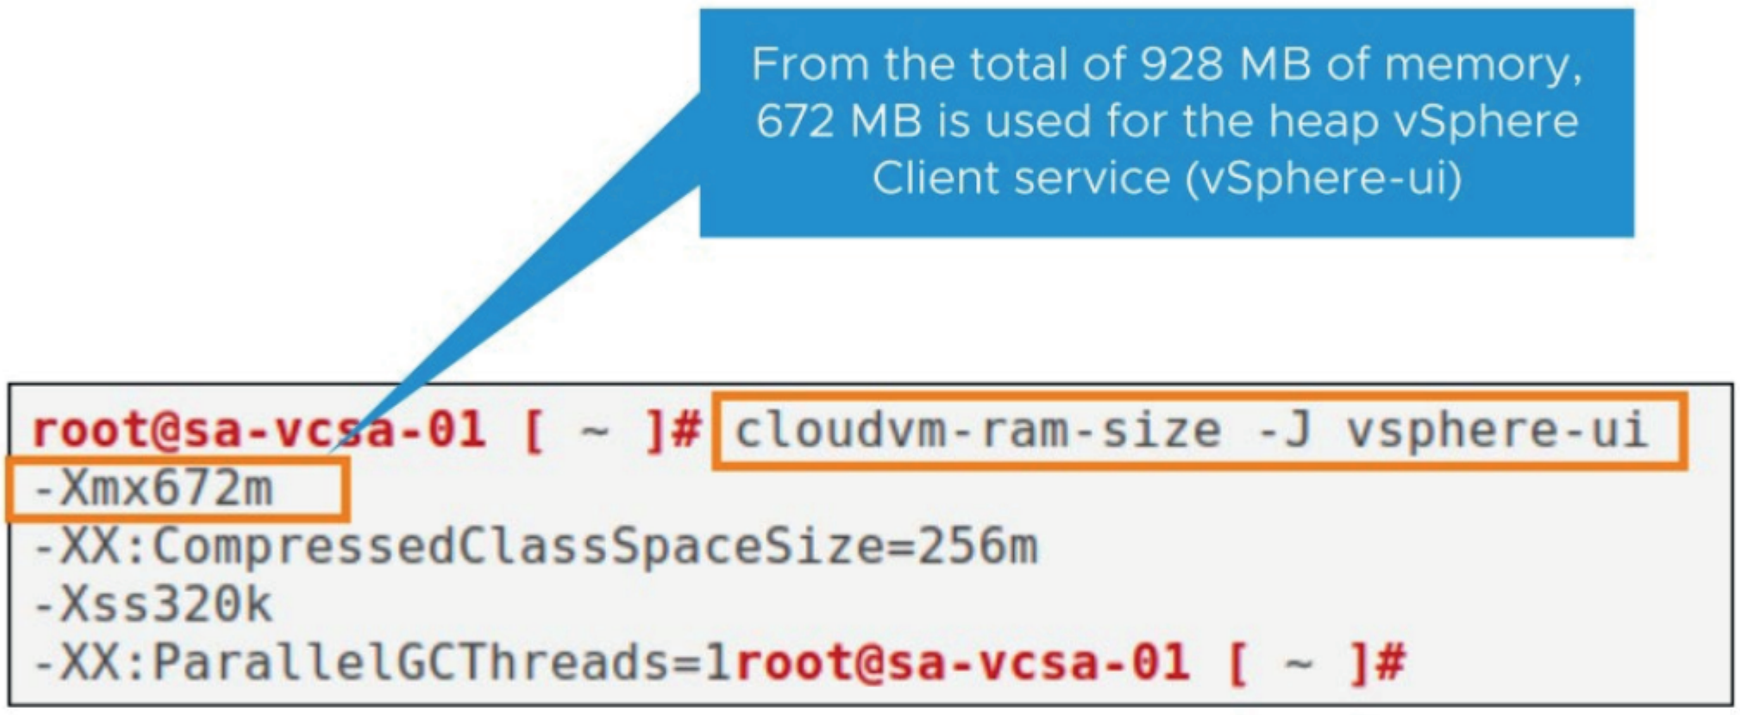

From the total of 928 MB of memory, 672 MB is used for the heap vSphere Client service (vSphere-ui)

Technet24

7-54 Changing Service Heap Size

You can increase the heap size of a vCenter Server service to resolve performance issues.

To change the heap size, use one of the following options:

- Resize the vCenter Server instance's memory and reboot.
  Heaps are automatically resized for all services.
- Resize the individual service and restart the service.
  No reboot is required.

  To resize an individual service, use the following command:

  ```
  cloudvm-ram-size -C <heap_size> <service_name>
  ```

  For example, the following command changes the heap size of the vSphere Web Client service to 700 MB:

  ```
  cloudvm-ram-size -C 700 vsphere-ui
  ```

If you do not want to take the downtime involved with rebooting vCenter Server after adding memory, increase the heap size of each individual service and restart that service. The heap is automatically resized.

7-55 vCenter Network Performance

Many vCenter operations involve communication between vCenter and the ESXi hosts, but when multiple vCenters are configured in an Enhanced Linked Mode Single Sign-on domain, lookup operations require communications from vCenter to vCenter:

- vCenter-to-vCenter network bandwidth is typically more critical than vCenter-to-host bandwidth.
- In most cases, operation latency is bound by host latency.

7-56 Monitoring Network Activity with the VAMI

Use the VAMI to monitor network activity, which includes byte rate, packets dropped, errors detected, and packet rate

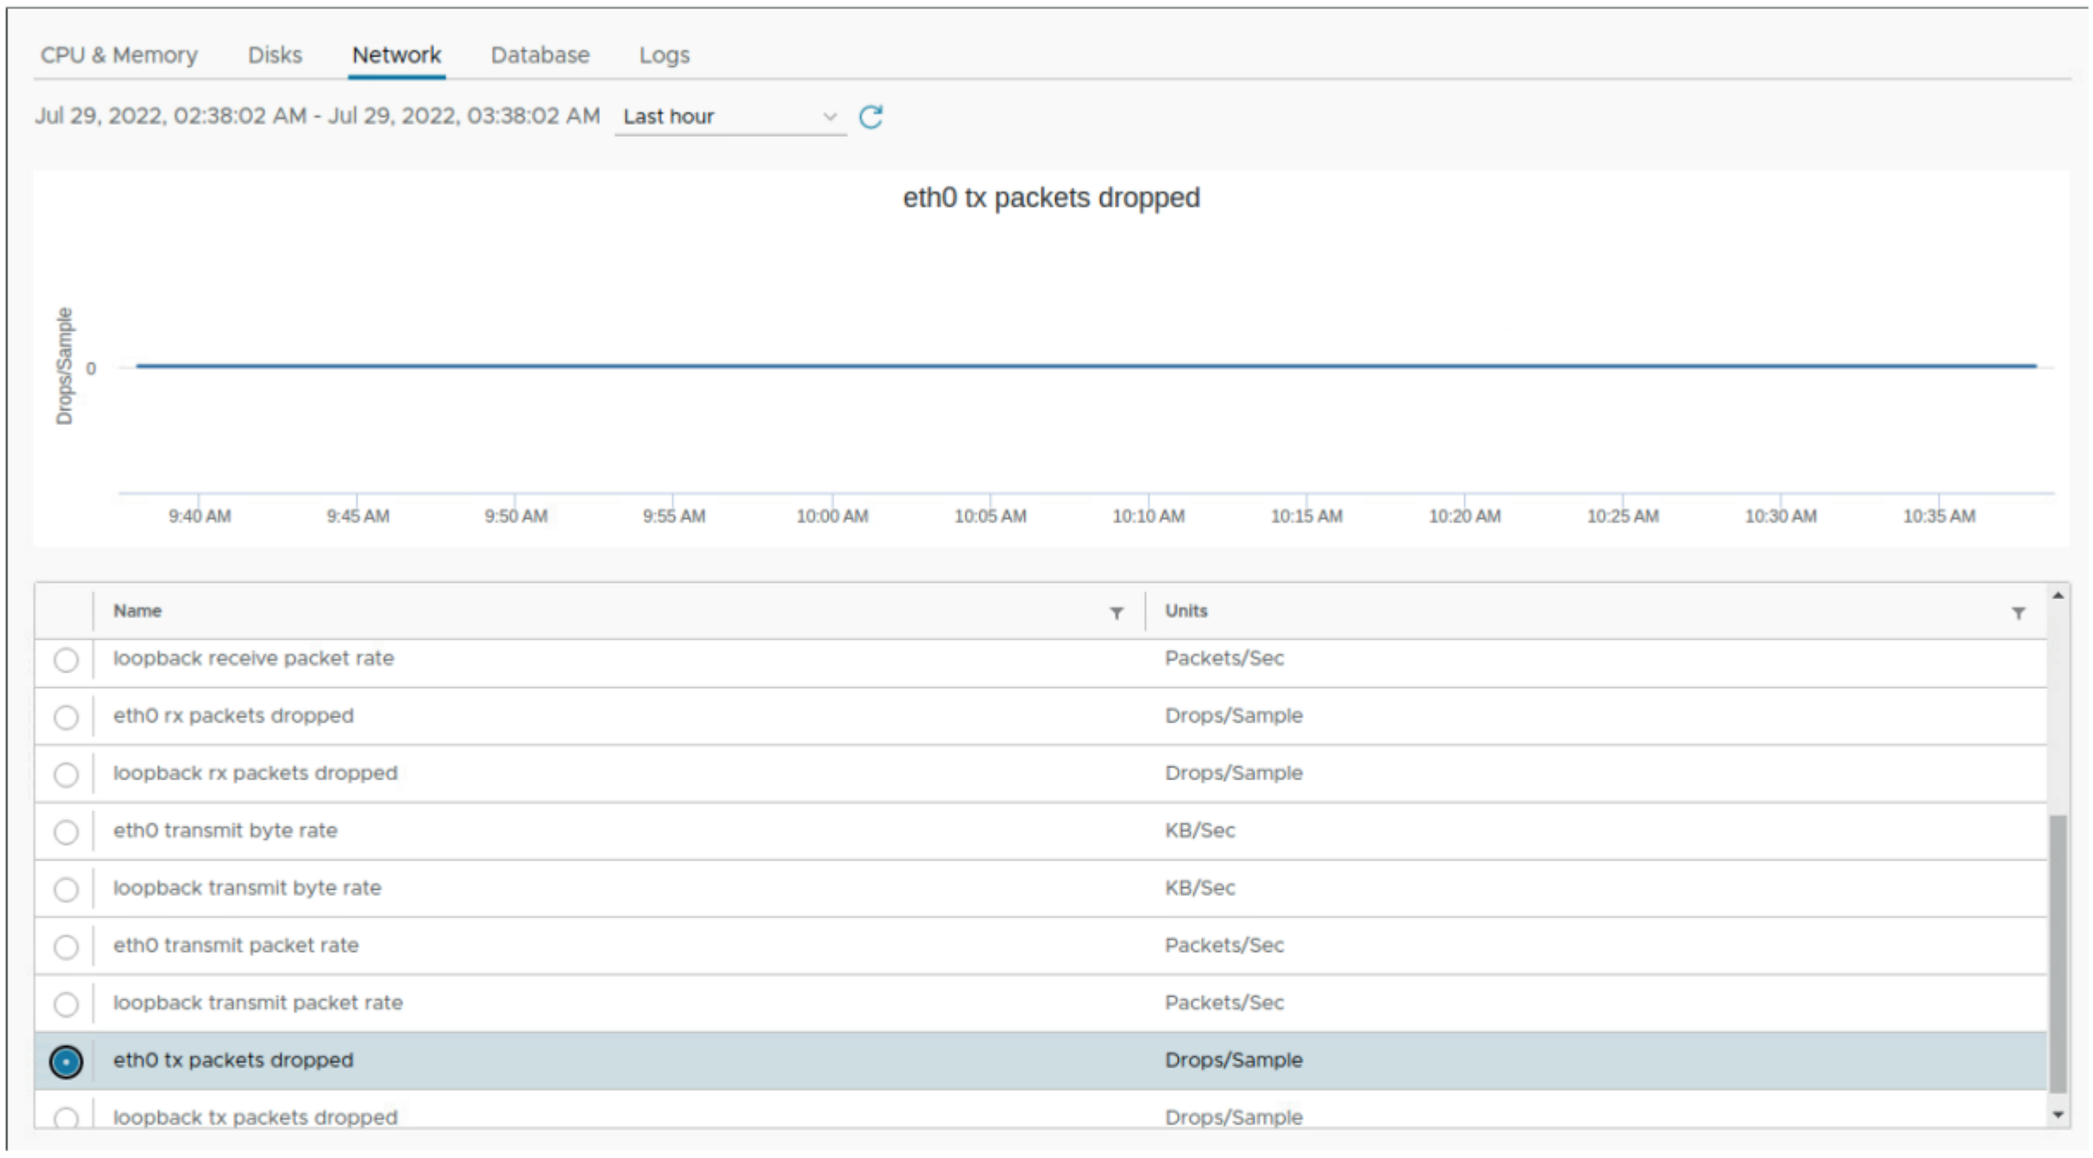

Technet24

7-57 Monitoring Network Activity with vimtop

In the `vimtop` window, enter **o** to display network metrics.

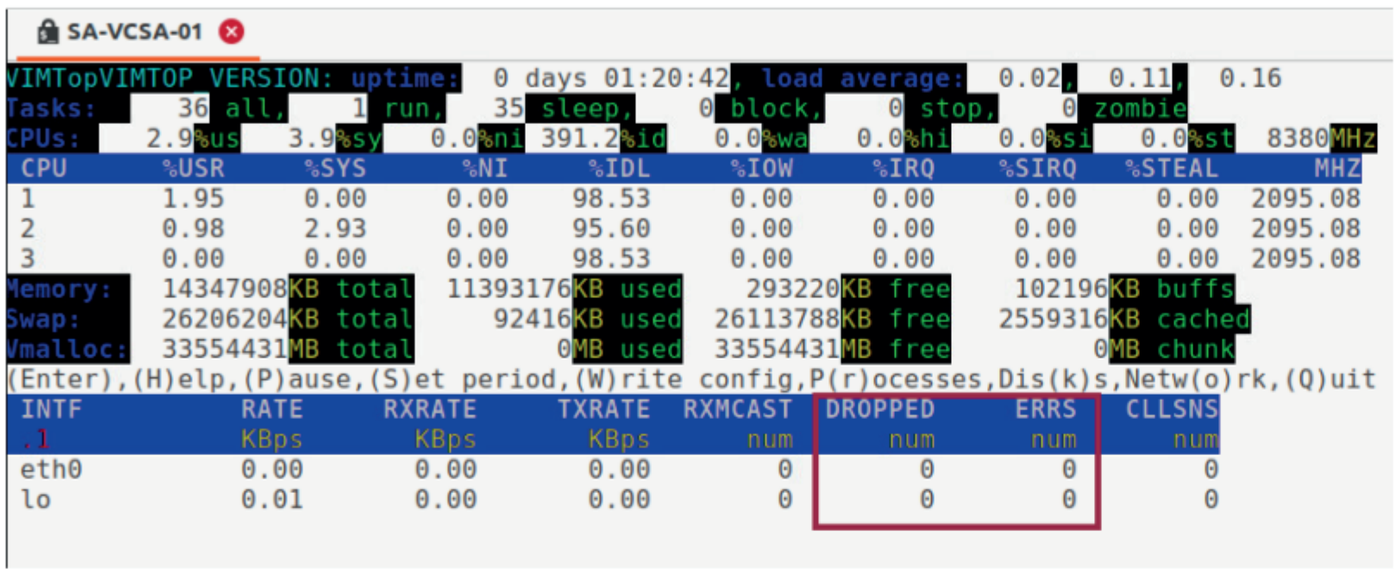

You can monitor the following columns for network activity:

- DROPPED: Number of dropped packets through this network interface because the system ran out of buffers during the last measurement cycle.
- ERRS: Total number of transmit faults and receive faults on this interface.

7-58 vCenter Database Performance (1)

Several factors should be considered for database performance:

- Under normal vCenter usage, the database resources (CPU, memory, disk, and network) should not be saturated.
- Transient high CPU usage (for example, during a rollup) is not a problem.
- Consistently high CPU usage is not normal.

The following symptoms occur with high CPU usage:

- Rollups or Top-N calculations of a large inventory are slow.
- Historical queries of a large inventory are slow.

CPU usage at 100 percent for long periods is not normal. When you see consistently high CPU usage, verify that nothing basic is malfunctioning.

If you have a large inventory, you will have many statistics. If you do not have a good enough disk subsystem, sufficient memory, or sufficient CPU, then rollups or Top_N calculations might be slow.

A large inventory does not only mean a high number of hosts and VMs, it might mean many datastores, networks, or devices per VM. Some queries require many joins or full scans on tables, which is not ideal but does occur. In such cases, the database statistics might need to be recomputed or reindexed to make the query engine operate more efficiently.

7-59 vCenter Database Performance (2)

If you see slow historical queries and consistently high database CPU usage:

- Verify the vPostgres database:
  - Recompute database statistics on database tables.
  - Purge database tables, lower the retention policy, or lower the statistics level.
  - Check the vpxd log, database log, database partition sizes, and SEAT tables.

If you are using a third-party tool (which uses the VMware vSphere® Web Services SDK) to obtain historical data:

- Verify queries.
- Check events, tasks, and statistics tables.

If you experience slow rollups or Top-N calculations:

- Add CPUs.
- Verify memory utilization.
- Verify rollup status.
- Verify alarms.

You can verify the `vpxd` log for diagnosis. You might see the `[VdbStatement] Execution elapsed time: ...` and `SQL execution took too long` statements in the log. You should expect to see a few of these events (often for Top_N, statistics, and events), each of which typically lasts between three and four seconds. If these events are extremely frequent and each lasts 10 or more seconds, this behavior often indicates problems in your database. Verify the resources (I/O, CPU, and memory) and that the database is correctly provisioned.

Technet24

7-60 Effects of Changing Statistics Level on Database Traffic (1)

vCenter uses a collection level, or statistics level, to determine the amount of data gathered and stored in the database.

When you increase the statistics level, the storage and system requirements will change.

The configuration of VMs and hosts, as well as the number of datastores, directly affects the number of statistics that are generated.

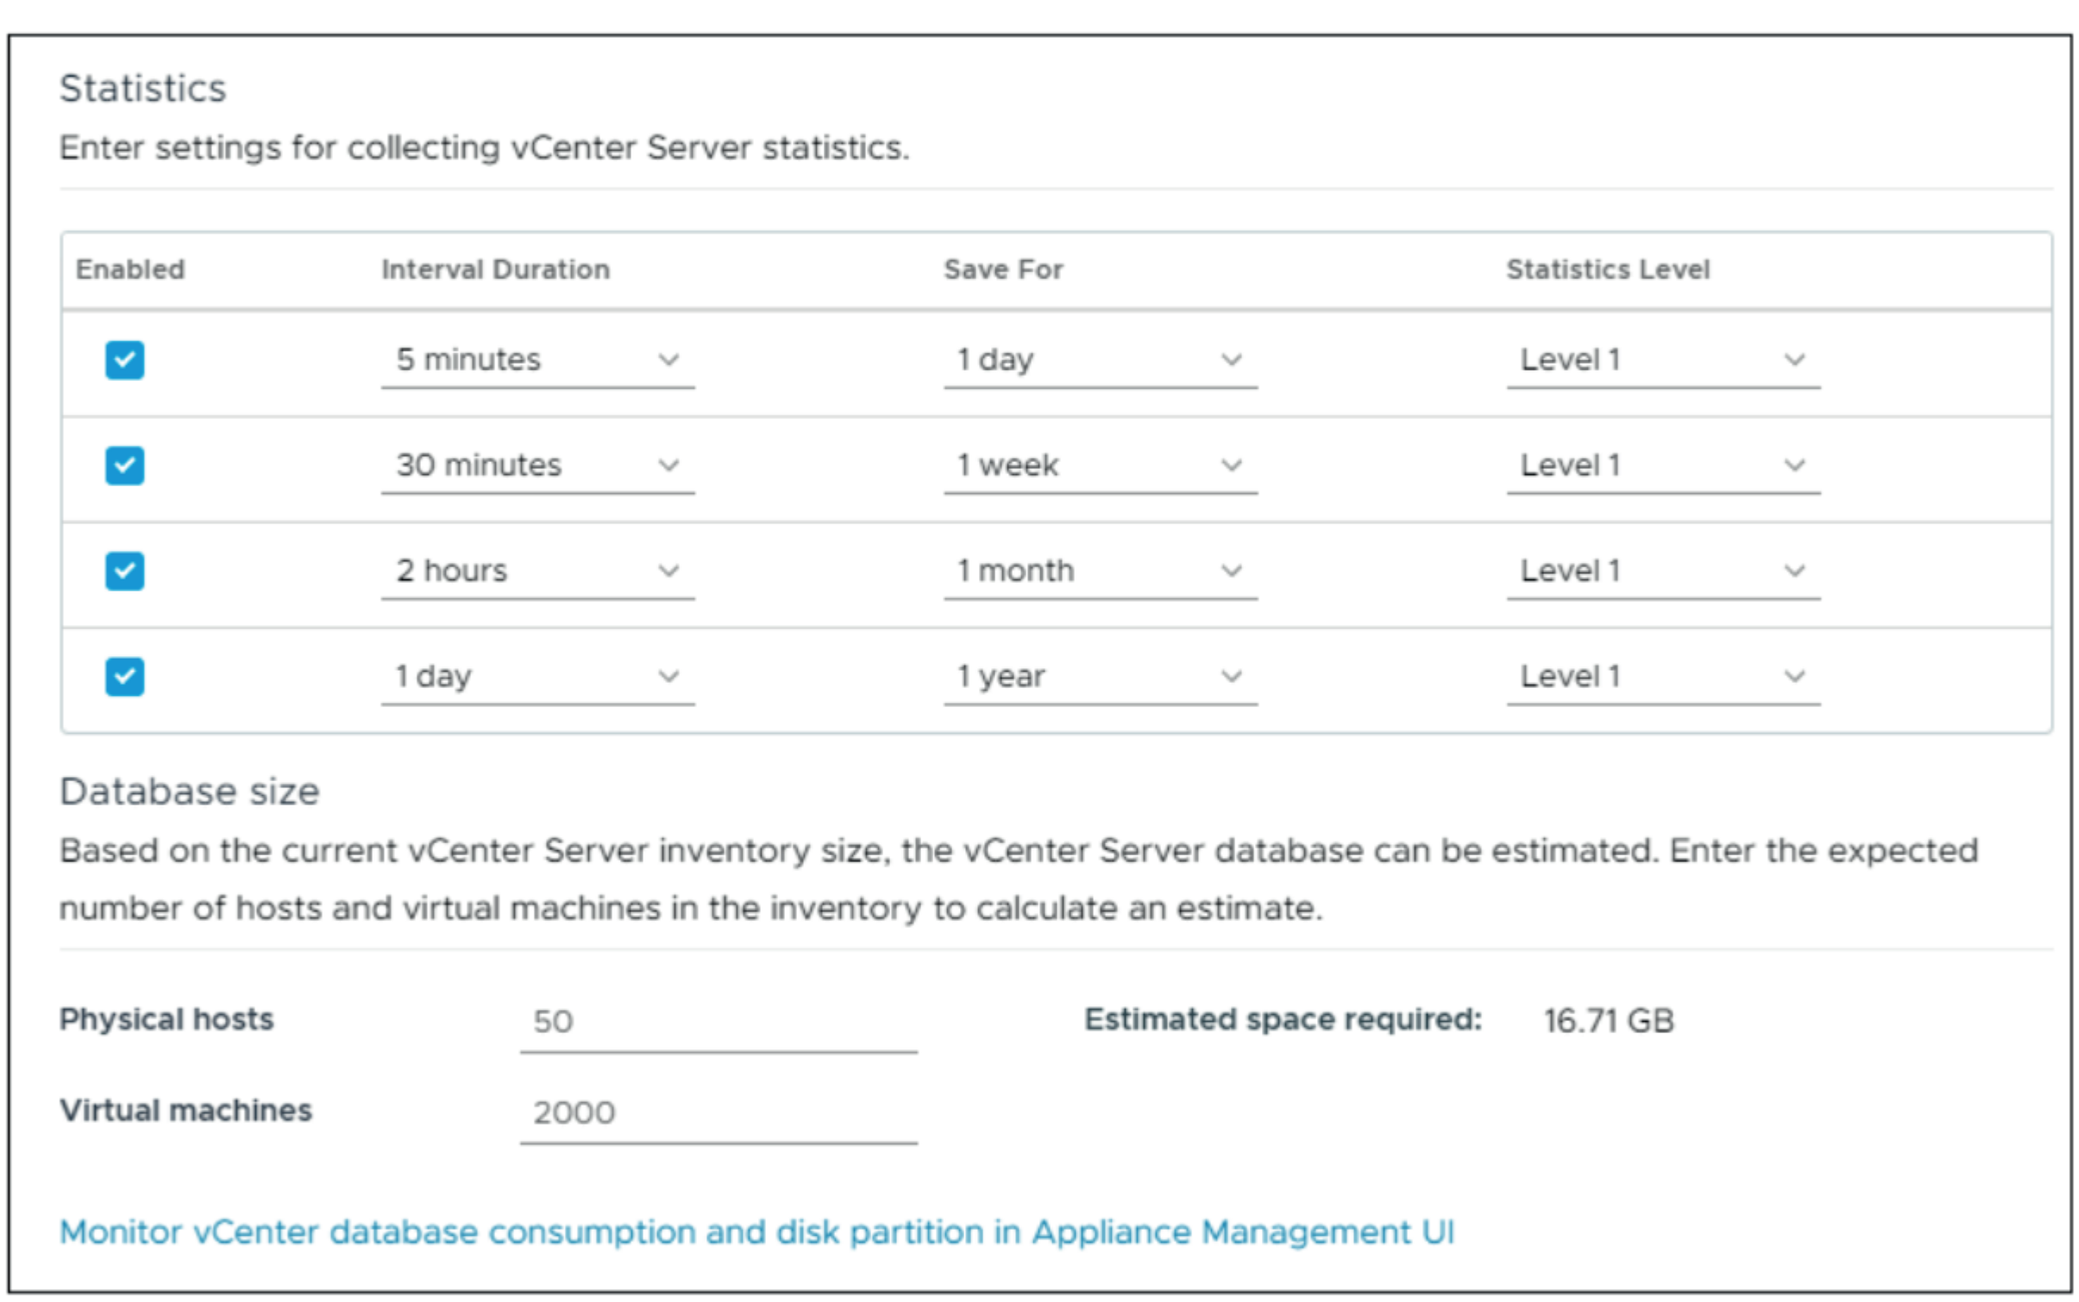

7-61 Effects of Changing Statistics Level on Database Traffic (2)

When changing from level 1 to level 2, the number of storage and network counters increases significantly.

Hardware	VM 1	VM 2	ESXi Host
CPUs	2	2	48 (logical)
Virtual disks	11	1	13
Datastores	1	1	9
NICs	1	1	3

Statistics Level	VM 1	VM 2	ESXi Host
1	67	34	223
2	231	148	858
3	263	184	1,779
4	348	196	1,967

The tables show what happens when different statistics levels are used.

The configuration of two VMs and a host is shown in the first table. The configuration of your VMs and hosts is critical to how many statistics are generated. In the example, the VM configurations are the same except for the number of virtual disks.

The second table shows what happens when you increase the statistics level. Of the four statistics levels, 1 is the least detailed and 4 is the most detailed. The specific numbers are not as important as the significant increase in statistics from level 1 to level 2. Level 1 reports on aggregated statistics, such as aggregated CPU, aggregated disk latencies, and so on. Level 2 reports more detailed statistics, such as per-NIC and per-datastore statistics.

In the second table, you can see that VM1 has far more statistics than VM2, which is because VM1 has many more disks than VM2. Several datastore and network counters (such as read/write requests per second and packets sent/received) are introduced at level 2. As a result, the larger the number of disks and networks in your configuration, the more statistics that are generated.

In your environment, ensure that you size the infrastructure and database optimally to accommodate the load of a level.

Technet24

7-62 vCenter Disk Use (1)

The vCenter database is the main disk bandwidth consumer and is write-intensive.

The highlighted partitions should be stored on high-speed storage.

Disk (VMDK)	Minimum Size	Mount Point	Purpose
VMDK1	12 GB	/ (10 GB), /boot (123 MB), SWAP (1 GB)	Kernel images and boot loader configurations
VMDK2	1.8 GB	/tmp	Temporary files
VMDK3	25 GB	SWAP	Used when system is out of memory to swap to disk
VMDK4	25 GB	/storage/core	Core dumps from vpxd process
VMDK5	10 GB	/storage/log	vCenter logs
VMDK6	10 GB	**/storage/db**	VMware Postgres database storage location
VMDK7	15 GB	**/storage/dblog**	VMware Postgres database logging location
VMDK8	10 GB	**/storage/seat**	Statistics, events, alarms, and tasks for VMware Postgres

The substantial source of disk I/O tends to be the database. The main disk bandwidth consumers are the following partitions:

- `/storage/db` is the location of the vCenter database.
- `/storage/dblog`
- `/storage/seat`

Ensure that these partitions are on high-speed storage and have sufficient space.

7-63 vCenter Disk Use (2)

Disk (VMDK)	Minimum Size	Mount Point	Purpose
VMDK9	1 GB	/storage/netdump	vSphere Netdump collector repository for ESXi dumps
VMDK10	10 GB	/storage/autodeploy	vSphere Auto Deploy repository
VMDK11	10 GB	/storage/imagebuilder	vSphere ESXi Image Builder repository
VMDK12	100 GB	/storage/updatemgr	vSphere Lifecycle Manager repository
VMDK13	50 GB	/storage/archive	Repository used for archiving files
VMDK14	10 GB	/storage/vtsdb	VMware vTSDB service repository that stores the stats
VMDK15	5 GB	/storage/vtsdblog	VMware vTSDB service repository that stores the logs of the service
VMDK16	100 GB	/storage/lifecycle	Workload Control Plane service stage directory that stores the binaries for install, update, and upgrade
VMDK17	24.5 GB	/storage/lvm_snapshot	Allow's vCenter to be more efficiently upgraded

Technet24

7-64 Monitoring Disk Usage with the VAMI

Use the VAMI to monitor disk use by partition.

Point to a partition to view information on used space, available space, and total space.

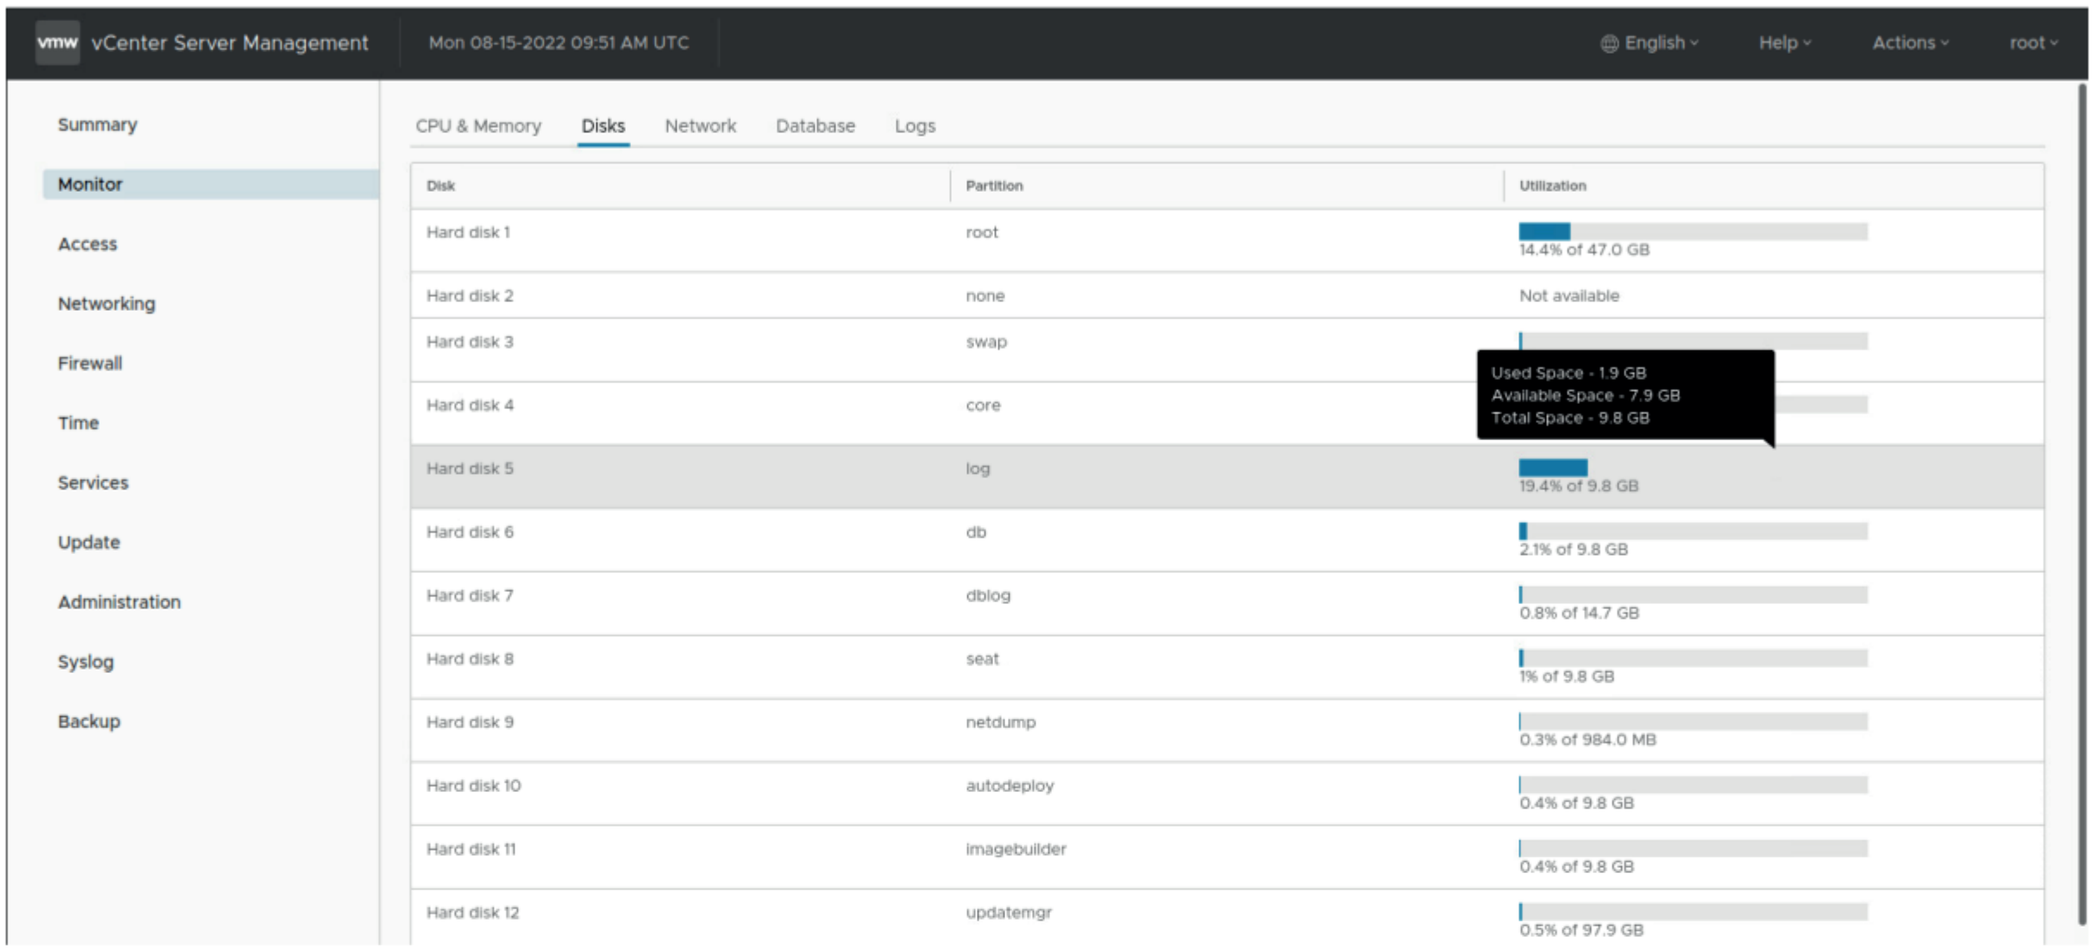

7-65 Monitoring Disk Usage with the df Command

You can also monitor disk use from the vCenter Server Appliance command line by using the **df -h** command.

```
root@sa-vcsa-01 [ ~ ]# df -h
Filesystem                                  Size  Used Avail Use% Mounted on
devtmpfs                                    6.9G     0  6.9G   0% /dev
tmpfs                                       6.9G  812K  6.9G   1% /dev/shm
tmpfs                                       6.9G  1.3M  6.9G   1% /run
tmpfs                                       6.9G     0  6.9G   0% /sys/fs/cgroup
/dev/mapper/vg_root_0-lv_root_0              47G  7.0G   38G  16% /
/dev/sda3                                   488M   30M  423M   7% /boot
/dev/sda2                                    10M  2.0M  8.1M  20% /boot/efi
/dev/mapper/vg_lvm_snapshot-lv_lvm_snapshot  25G   45M   24G   1% /storage/lvm_snapshot
tmpfs                                       6.9G  5.6M  6.9G   1% /tmp
/dev/mapper/core_vg-core                    9.8G   37M  9.3G   1% /storage/core
/dev/mapper/netdump_vg-netdump              985M  2.5M  915M   1% /storage/netdump
/dev/mapper/log_vg-log                      9.8G  1.3G  8.0G  14% /storage/log
/dev/mapper/seat_vg-seat                    4.9G  113M  4.5G   3% /storage/seat
/dev/mapper/db_vg-db                        9.8G  186M  9.1G   2% /storage/db
/dev/mapper/dblog_vg-dblog                   15G  121M   14G   1% /storage/dblog
/dev/mapper/autodeploy_vg-autodeploy        985M  2.5M  915M   1% /storage/autodeploy
/dev/mapper/imagebuilder_vg-imagebuilder    4.9G   20M  4.6G   1% /storage/imagebuilder
/dev/mapper/updatemgr_vg-updatemgr          9.8G  471M  8.8G   5% /storage/updatemgr
/dev/mapper/vtsdb_vg-vtsdb                  9.8G   69M  9.2G   1% /storage/vtsdb
/dev/mapper/lifecycle_vg-lifecycle          9.8G  3.7G  5.7G  39% /storage/lifecycle
/dev/mapper/vtsdblog_vg-vtsdblog            4.9G   52M  4.6G   2% /storage/vtsdblog
/dev/mapper/archive_vg-archive              9.8G  521M  8.8G   6% /storage/archive
root@sa-vcsa-01 [ ~ ]#
```

The **df** command displays file system use information, such as total size, used space, and available space. The **-h** option shows sizes in powers of 1024, for example, in megabytes, gigabytes, and so on.

Technet24

7-66 Monitoring Disk Activity with vimtop

In the `vimtop` window, enter **k** to display disk metrics.

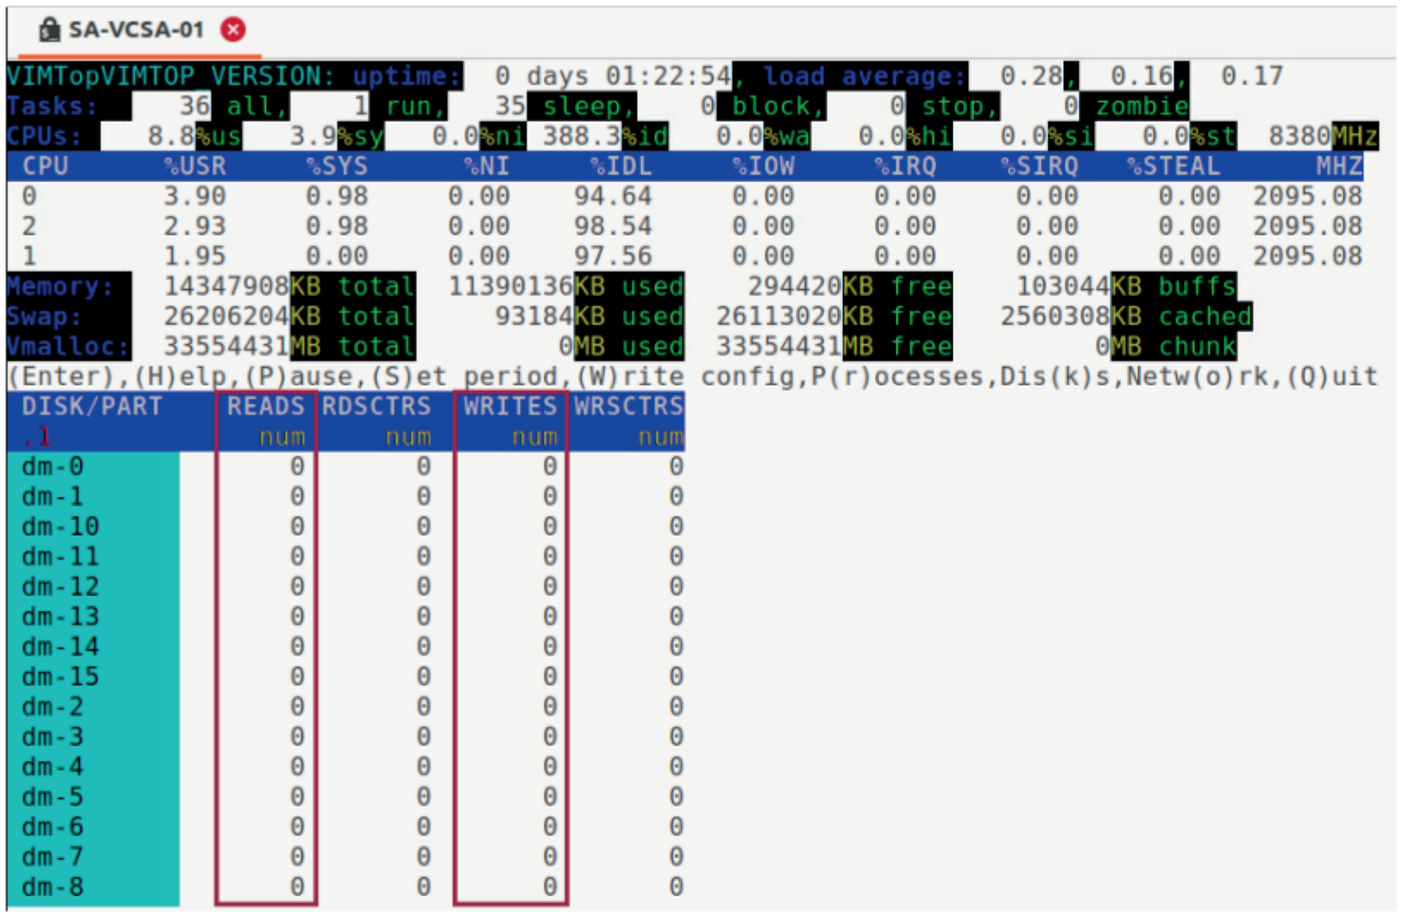

The following columns show:

- READS: Number of reads issued to this partition and completed successfully during the last measurement interval.
- WRITES: Number of writes issued to this partition and completed successfully during the last measurement interval.

7-67 Monitoring Disk usage with the Appliance Shell

Using Appliance Management CLI, the command:

`com.vmware.appliance.version1.resources.storage.stats.list`

has to be executed from the `Command>` location. The command does not work if you have entered inside the shell prompt.

```
Command> com.vmware.appliance.version1.resources.storage.stats.list
Stats:
  1:
      Name: /
      Usedkbytes: 7273468
      Totalkbytes: 49275540
      Usedinodes: 86612
      Totalinodes: 3145728
  2:
      Name: /boot
      Usedkbytes: 30692
      Totalkbytes: 499656
      Usedinodes: 314
      Totalinodes: 32768
  3:
      Name: /boot/efi
      Usedkbytes: 1978
      Totalkbytes: 10202
      Usedinodes: 0
      Totalinodes: 0
  4:
      Name: /storage/lvm_snapshot
      Usedkbytes: 45080
      Totalkbytes: 25663716
      Usedinodes: 11
      Totalinodes: 1638400
  5:
      Name: /storage/core
      Usedkbytes: 36904
      Totalkbytes: 10247444
      Usedinodes: 17
      Totalinodes: 655360
```

See VMware knowledge base article 2145603 at http://kb.vmware.com/kb/2145603.

Technet24

7-68 Increase Disk Space for vCenter

- From the ESXi Host Client, power off the vCenter VM.
- Right-click the VM.
- Click Edit Settings.
- Select Virtual Disk.
- Increase the size of the disk. Ensuring there is no snapshot attached to VM.
- Power on VM.
- Use the BASH Shell or Appliance Shell to expand the logical partition(s).
  - From the BASH Shell:

    ```
    /usr/lib/applmgmt/support/scripts/autogrow.sh
    ```
  - From the Appliance Shell:

    ```
    com.vmware.appliance.system.storage.resize
    ```

7-69 Monitoring Database Activity with the VAMI

Use the VAMI to monitor space usage trends for SEAT (stats, events, alarms, and tasks) activity, database log activity, and inventory data (core) activity.

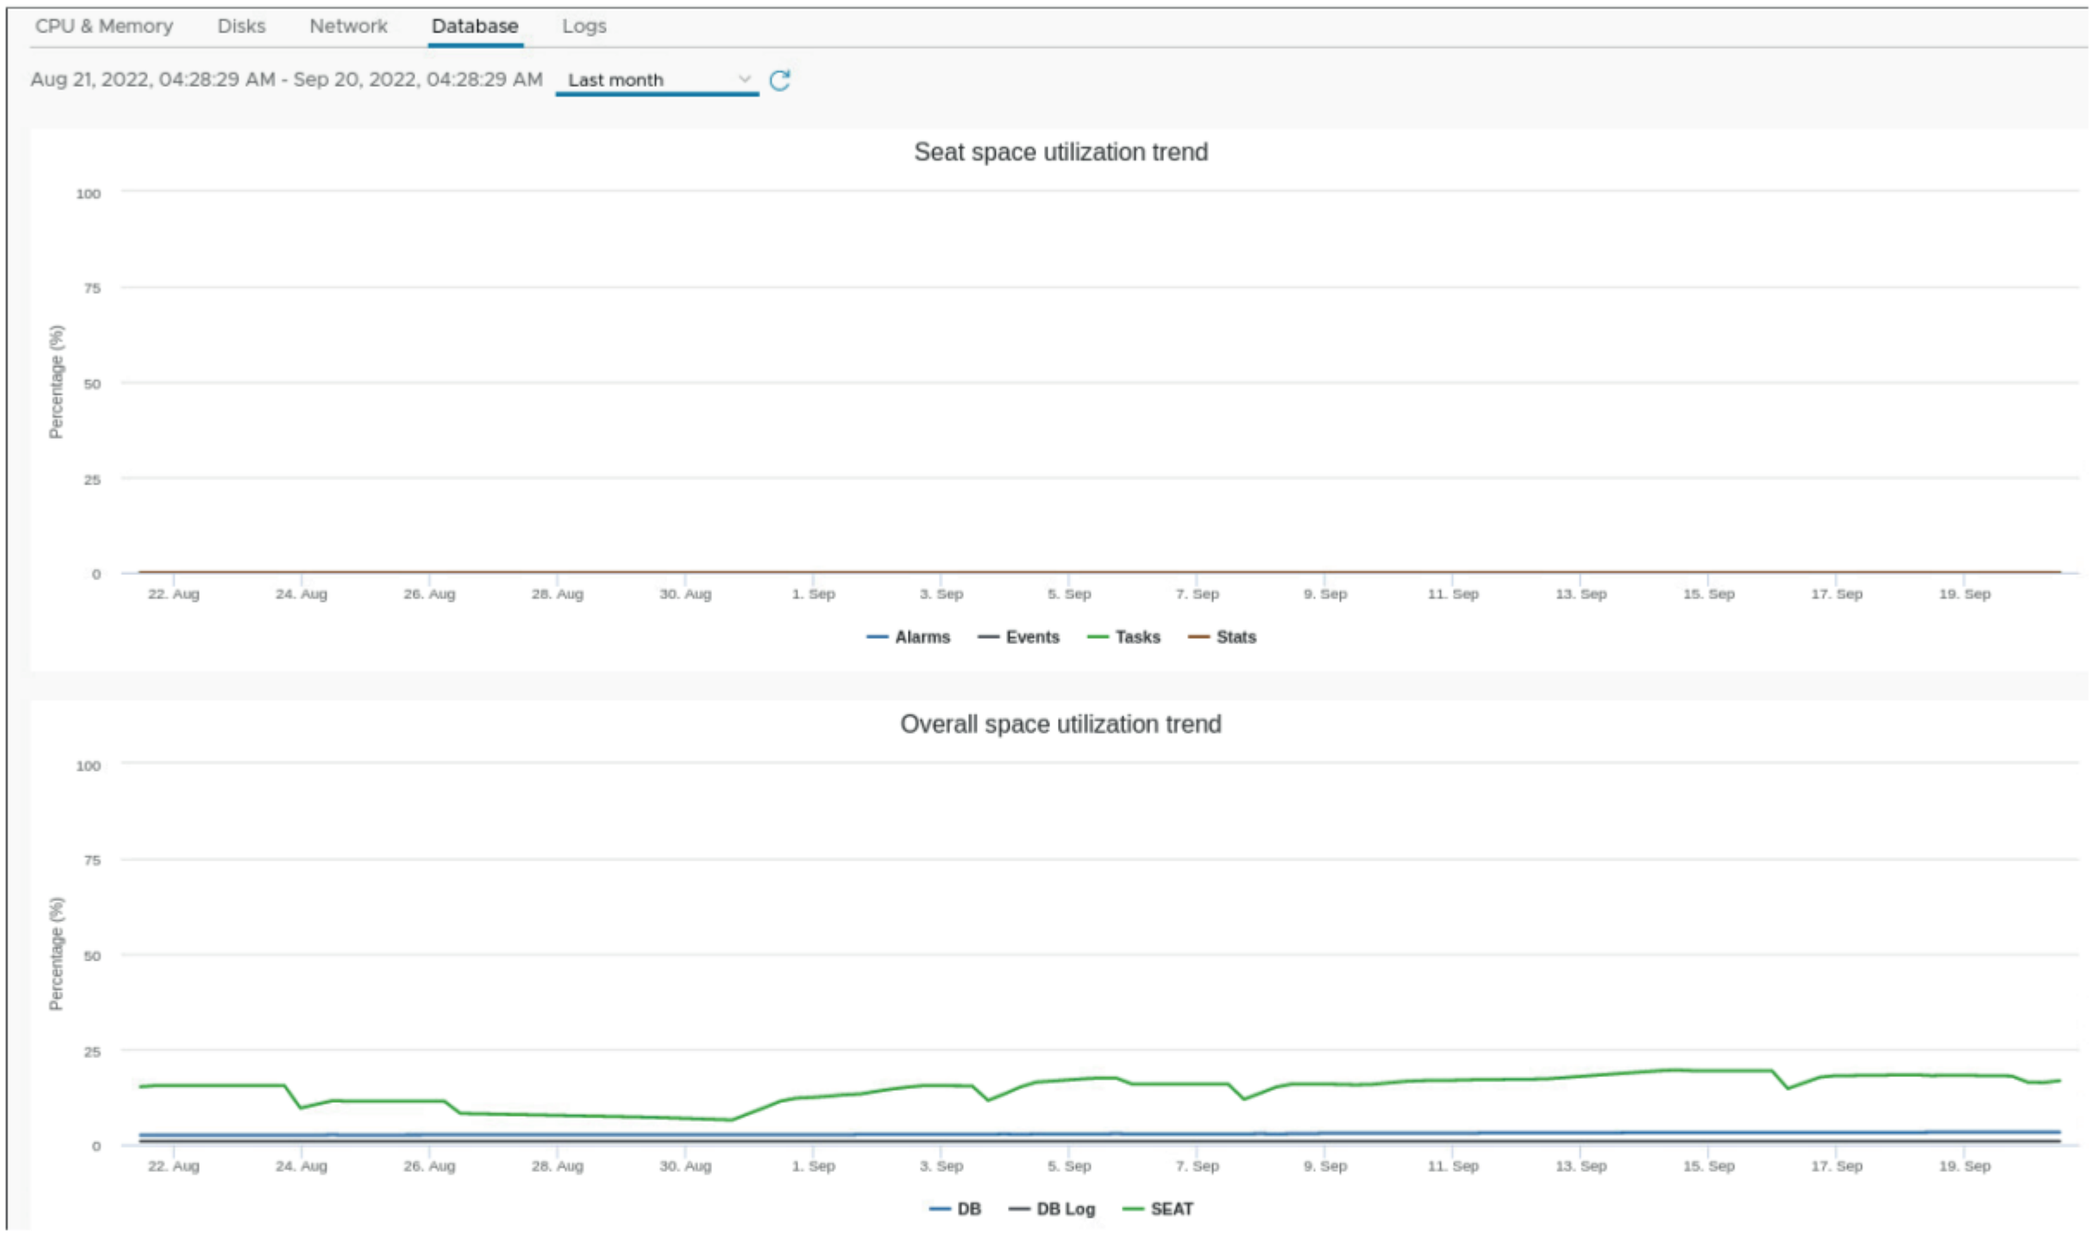

7-70 Track Your Progress Towards Sustainability Goals

Introduced in vSphere 8.0, the vSphere Client allows administrators to track and measure your environmental energy usage, carbon emissions, and other sustainability metrics in real-time.

Metrics are now available to capture power consumption of a host, including power associated with virtual machine workloads.

Key features include:

- Improved Sustainability
- Cost Savings
- Enhanced Reporting

Technet24

7-71 NEW: Enhanced vSphere Green Metrics

Introduced in vSphere 8.0 Update 1, administrators can get the power metrics for each of the VMs and gain access to an aggregated view of their power consumption data.

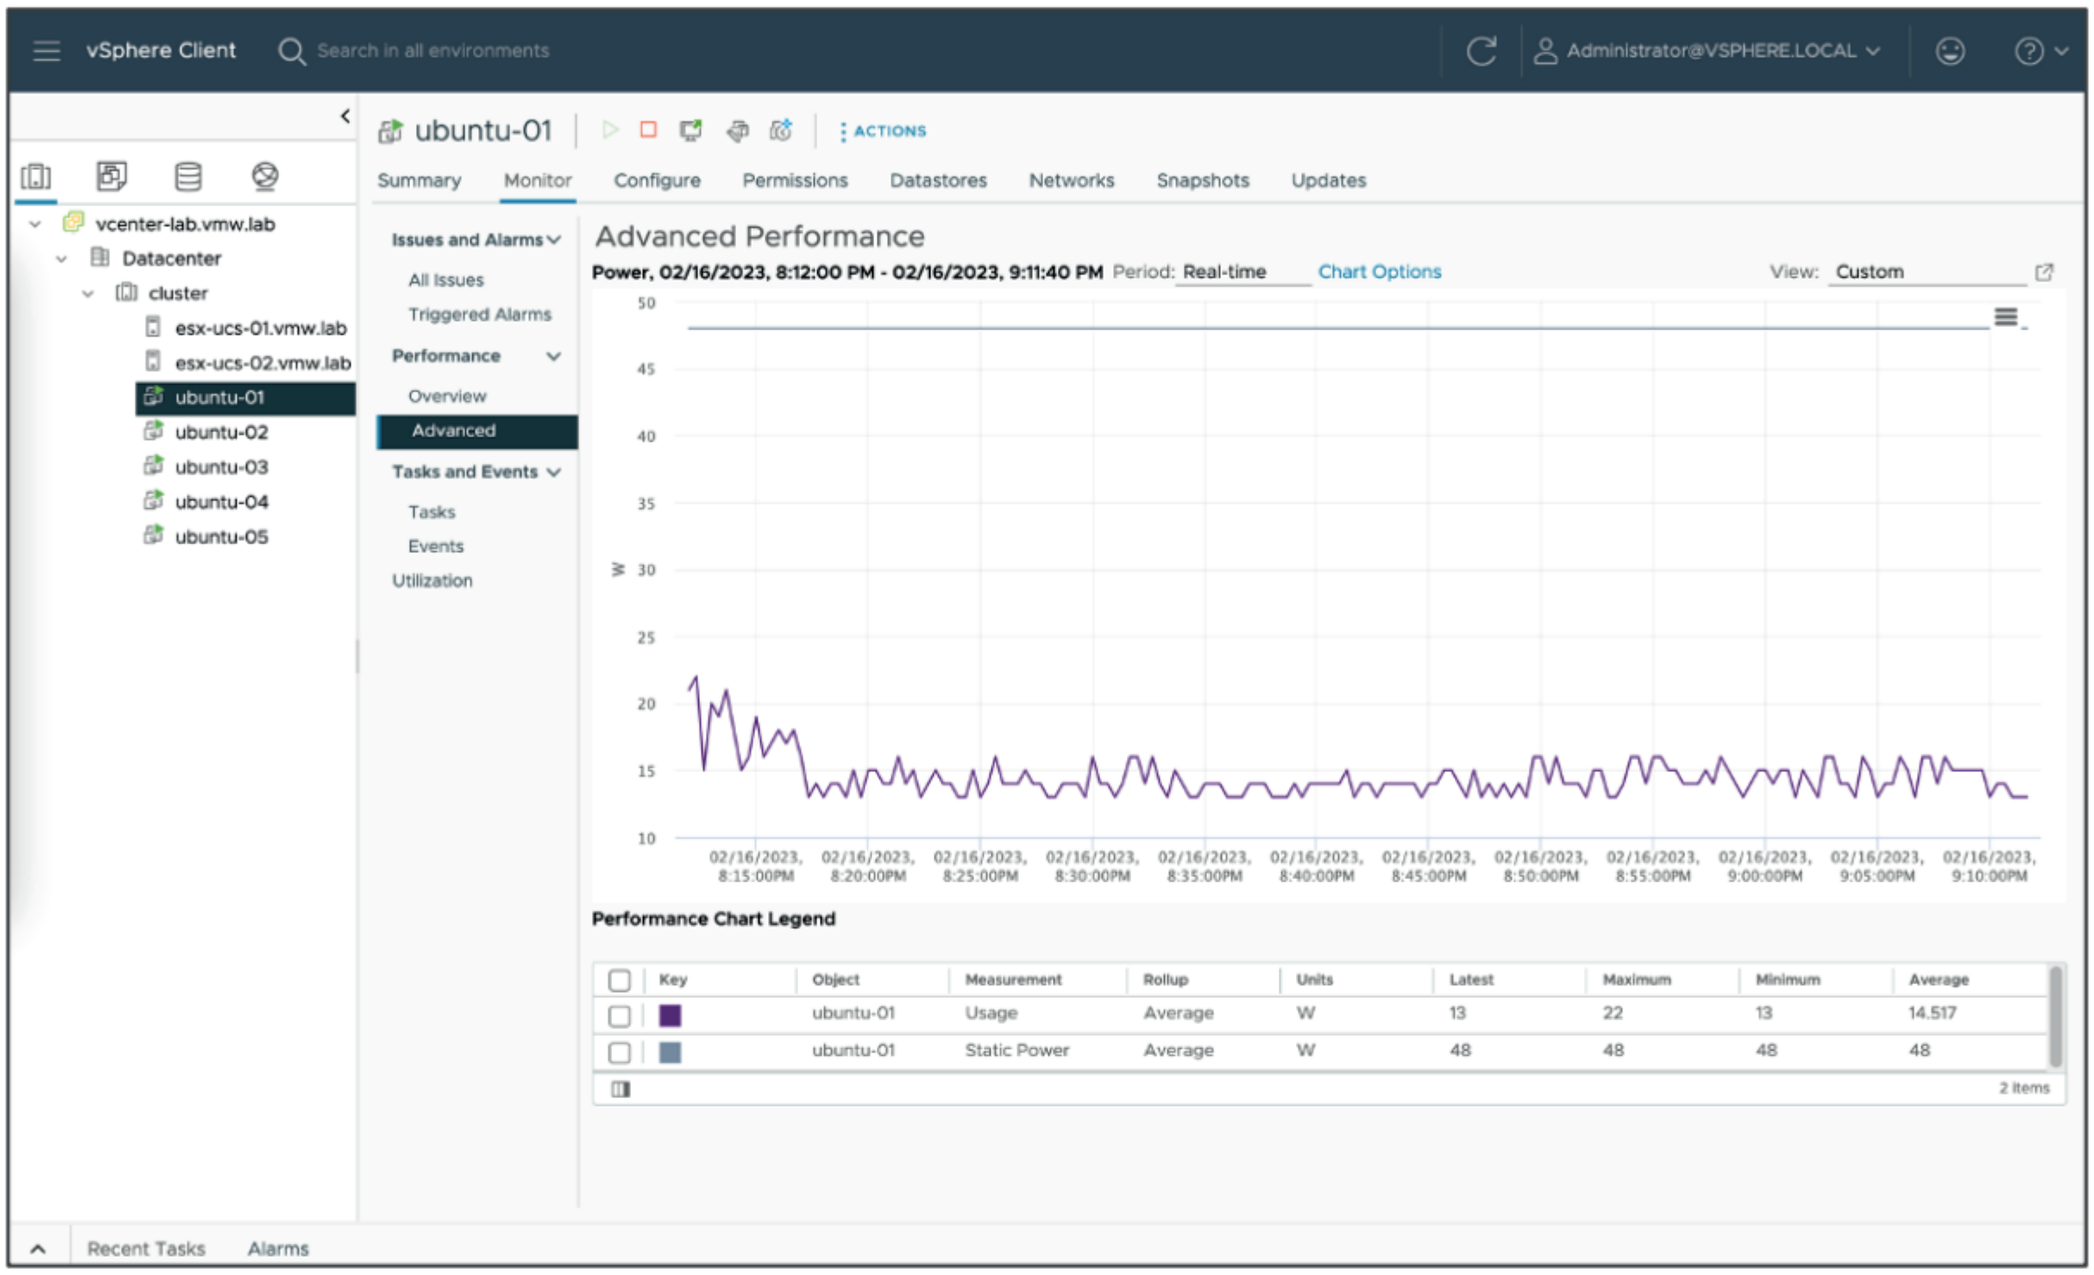

7-72 Activity: Network Performance (1)

Which of the following factors can affect network performance in your environment?

- ◊ ❒ Network bandwidth and latency
- ◊ ❒ Number of CPU cores assigned to a virtual machine
- ◊ ❒ Virtual machine uptime
- ◊ ❒ Amount of RAM assigned to a virtual machine

7-73 Activity: Network Performance (2)

Which of the following factors can affect network performance in your environment?

◊ ✓ Network bandwidth and latency

◊ ❒ Number of CPU cores assigned to a virtual machine

◊ ❒ Virtual machine uptime

◊ ❒ Amount of RAM assigned to a virtual machine

7-74 Activity: vCenter Tools (1)

Which of the following tools can you use to monitor the free space in the vCenter database? Select all that apply.

◊ ❒ df

◊ ❒ vCenter Advanced Performance Charts

◊ ❒ vSphere Client

◊ ❒ VAMI

Technet24

7-75 Activity: vCenter Tools (2)

Which of the following tools can you use to monitor the free space in the vCenter database? Select all that apply.

◊ ✓ df

◊ ❒ vCenter Advanced Performance Charts

◊ ❒ vSphere Client

◊ ✓ VAMI

7-76 Review of Learner Objectives

- Review vCenter components and services
- Describe the factors that influence vCenter performance
- Use vCenter tools to monitor vCenter resource use

7-77 Lesson 4: Using Alarms

7-78 Learner Objectives

- Use predefined alarms in vCenter
- View and acknowledge alarms
- Create custom alarms

7-79 About Alarms

An alarm is a notification that is sent in response to an event or condition that occurs with an object in the inventory.

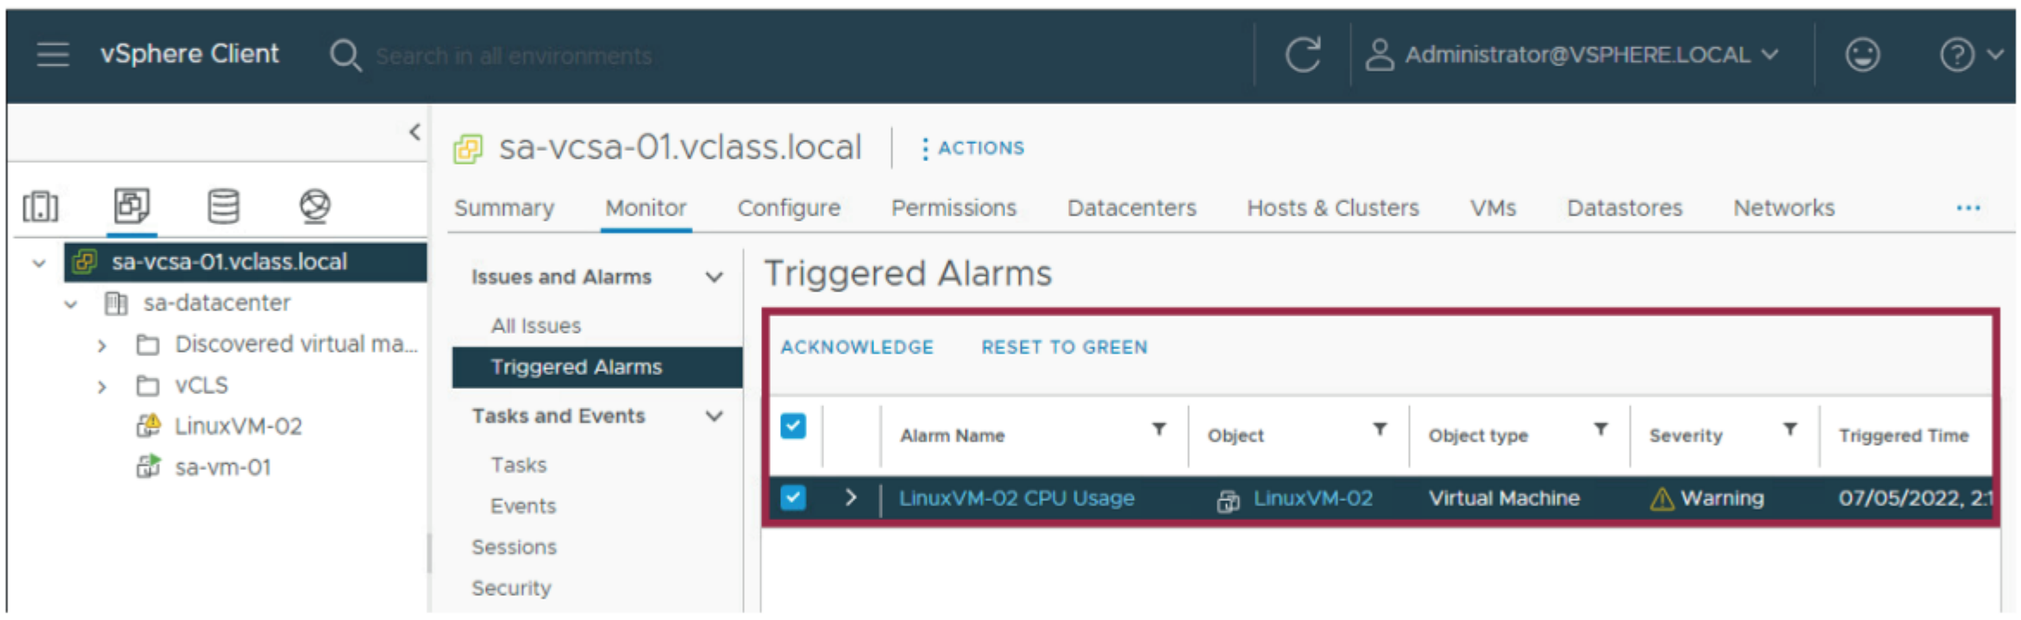

You can acknowledge an alarm to let other users know that you take ownership of the issue. For example, a VM has an alarm set to monitor CPU use. The alarm is configured to send an email to an administrator when the alarm is triggered. The VM CPU use spikes, triggering the alarm, which sends an email to the administrator. The administrator acknowledges the triggered alarm to let other administrators know the problem is being addressed

After you acknowledge an alarm, the alarm actions are discontinued, but the alarm does not get cleared or reset when acknowledged. You reset the alarm manually in the vSphere Client to return the alarm to a normal state, or the alarm will return to normal state due to its reset rules.

Technet24

7-80 Predefined Alarms

You can access many predefined alarms for various inventory objects, such as hosts, virtual machines, datastores, networks, and so on.

You can edit predefined alarms, or you can make a copy of an existing alarm and modify the settings as needed.

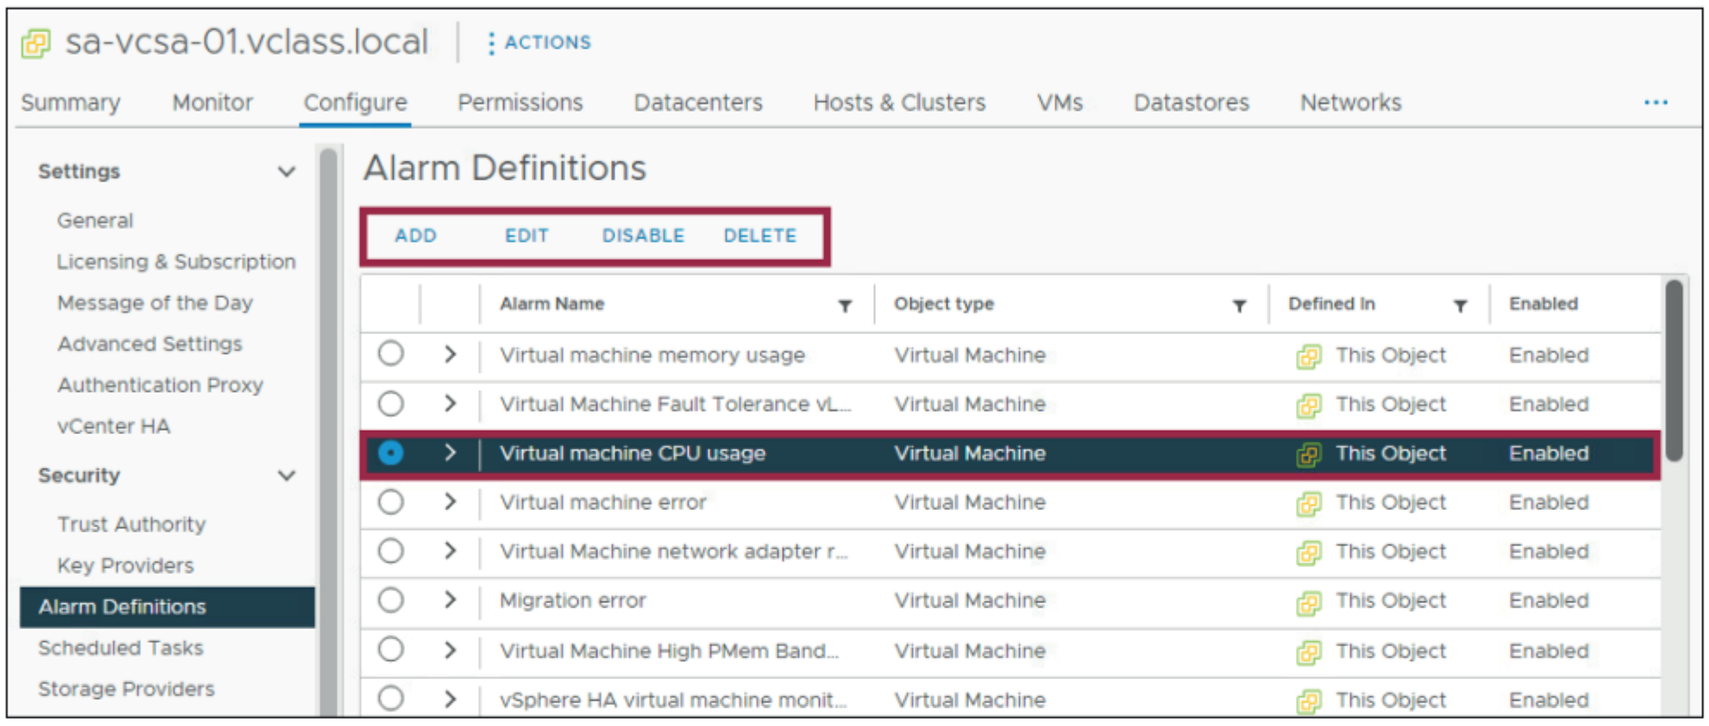

To make a copy of an alarm, choose the object in the inventory. Next, you will need to navigate to **Security** > **Alarm Definitions.** Finally, choose the alarm and click **ADD** and continue through the wizard.

7-81 Creating a Custom Alarm

In addition to using predefined alarms, you can create custom alarms in the vSphere Client.

The selected object in the inventory determines the scope of the alarm.

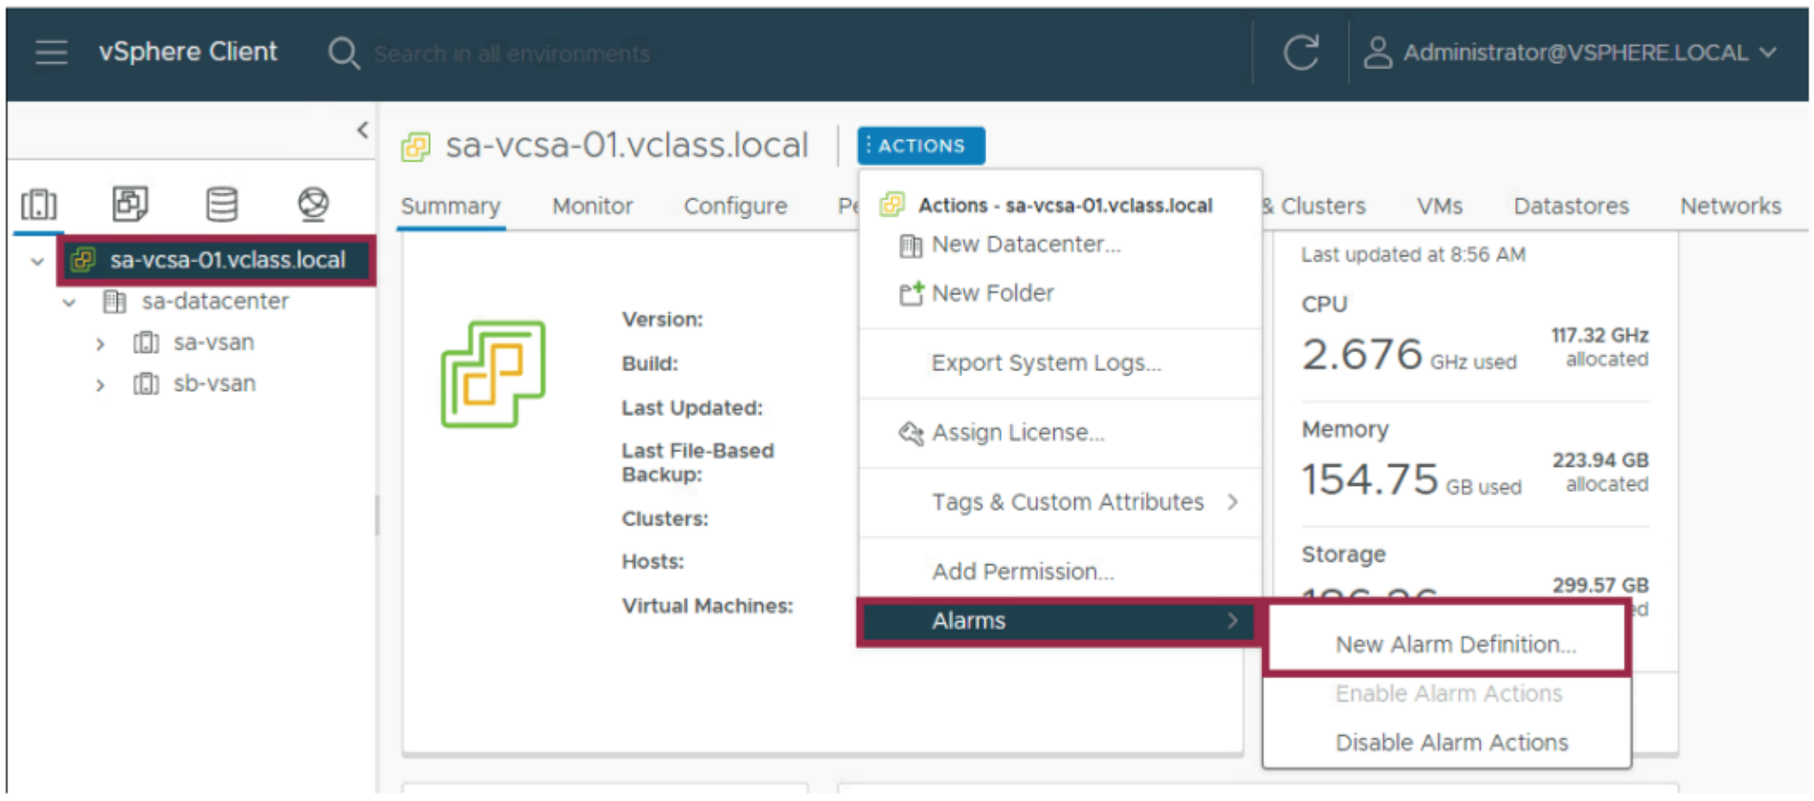

If the predefined alarms do not address the event, state, or condition that you want to monitor, define custom alarm definitions instead of modifying predefined alarms.

Technet24

7-82 Defining the Alarm Target Type

On the Name and Targets page, you name the alarm, give it a description, and select the type of inventory object that this alarm monitors.

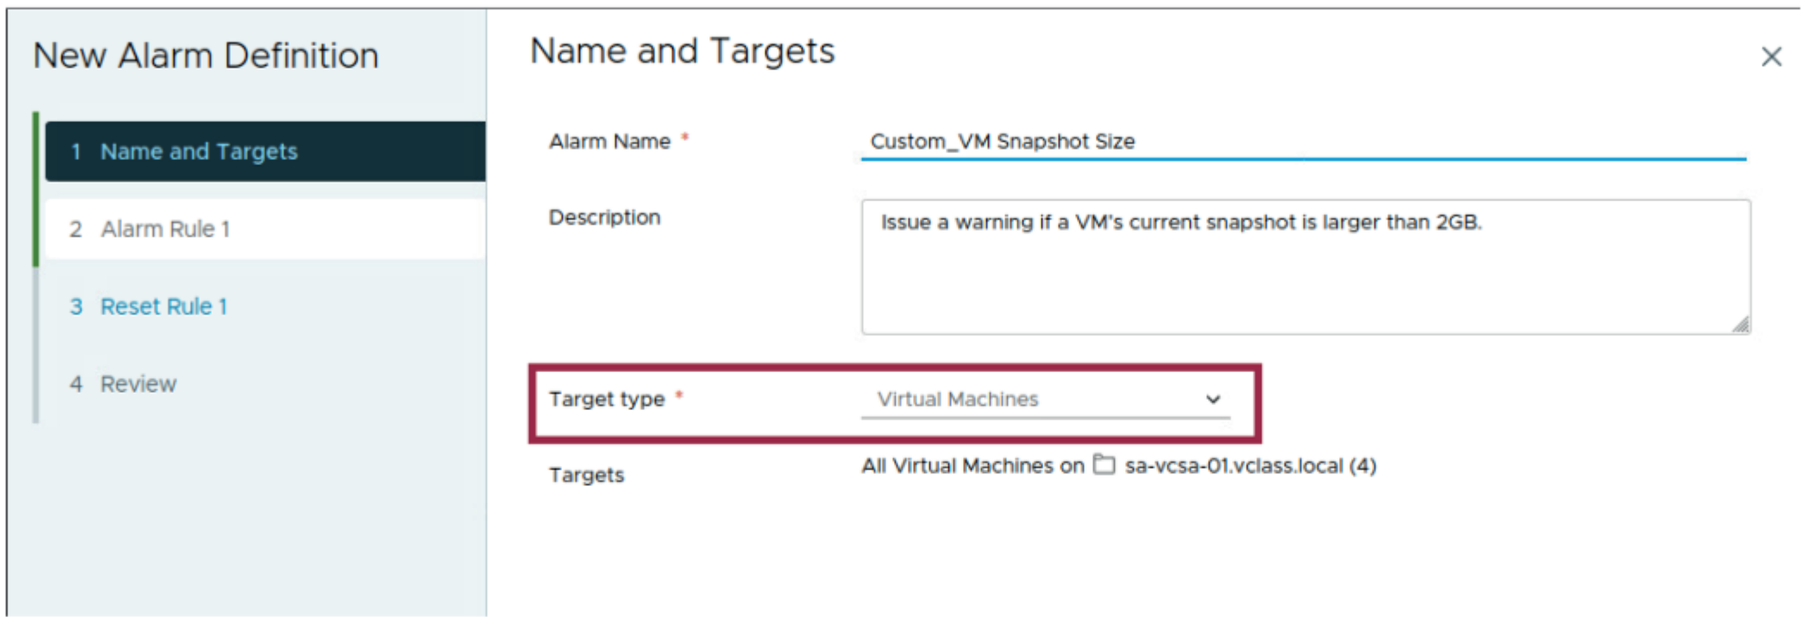

Depending upon the object where the alarm is being defined, you can create custom alarms for the following target types:

- Virtual machines
- Hosts, clusters, and data centers
- Datastores and datastore clusters
- Distributed switches and distributed port groups
- vCenter

7-83 Defining the Alarm Rule: Condition-based Trigger

An alarm rule must contain at least one trigger.

A trigger can monitor the current condition or state of an object, for example:

- A VM's current snapshot is more than 2 GB.
- A host is using 90 percent of its total memory.
- A datastore is disconnected from all hosts.

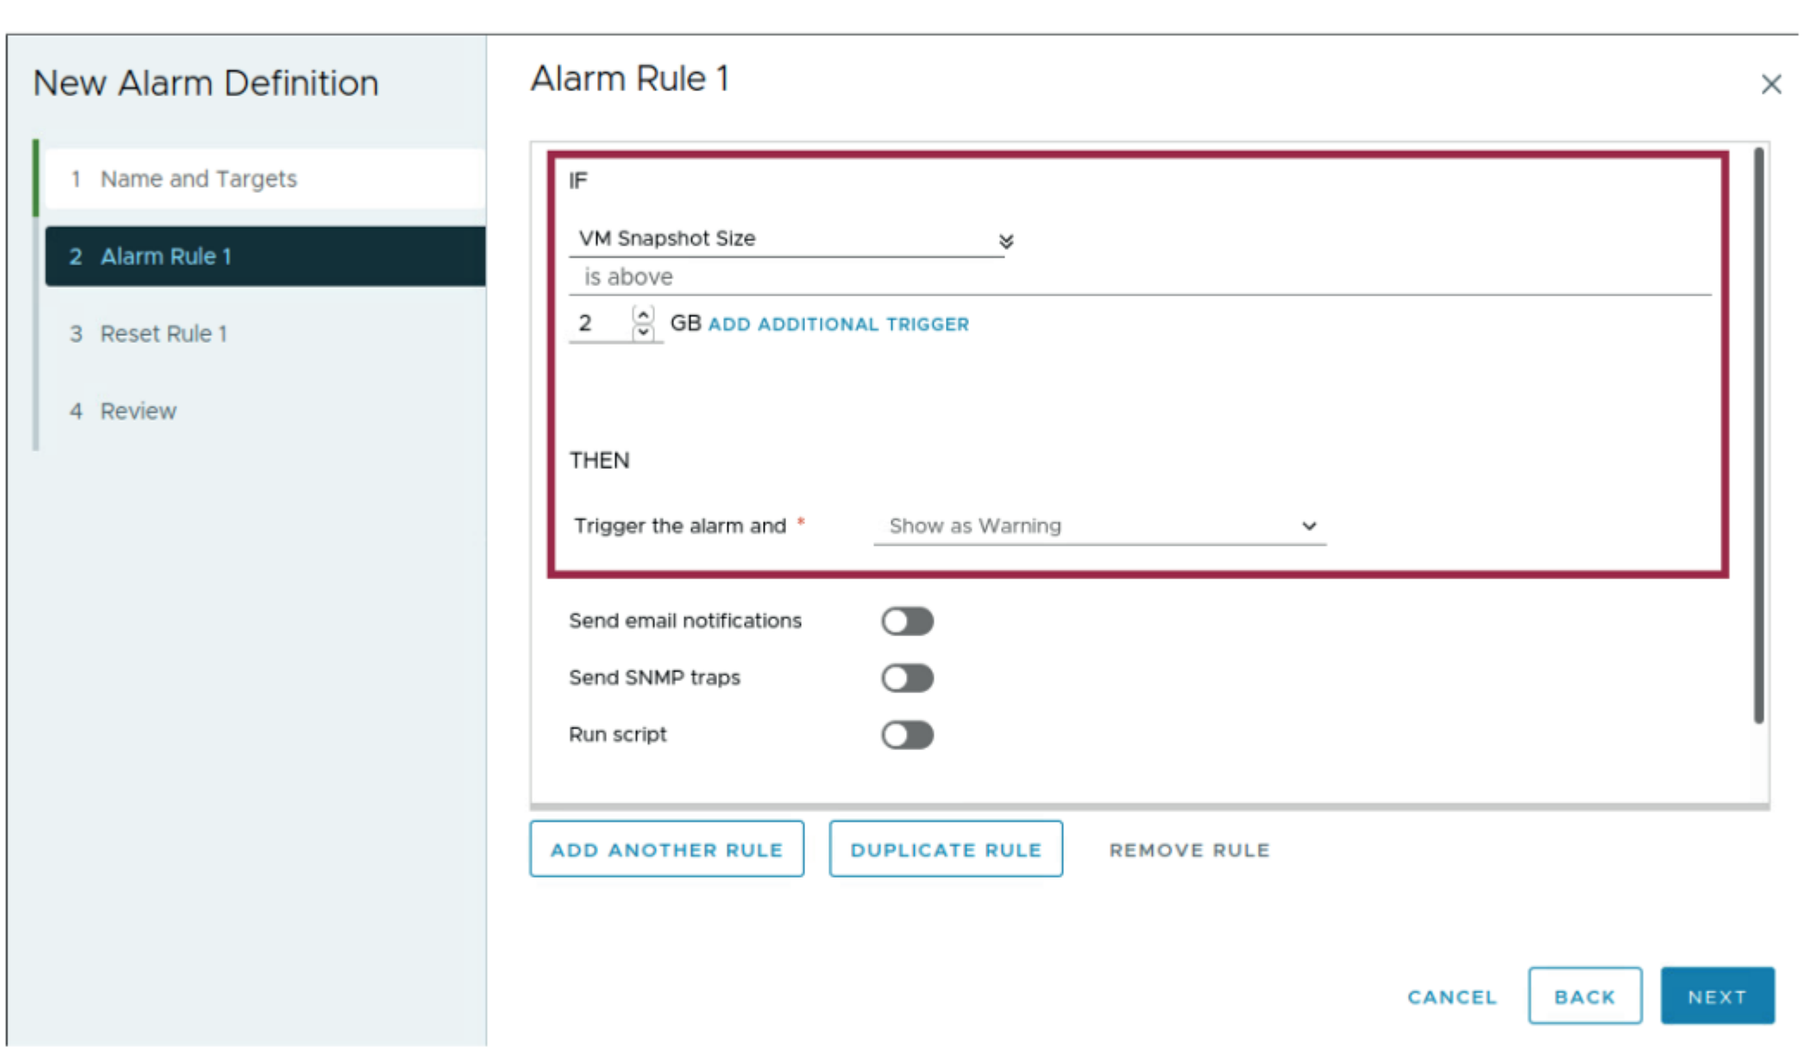

You configure the alarm trigger to show as a warning or critical event when the specified criteria are met:

- You can monitor the current condition or state of virtual machines, hosts, and datastores. Conditions or states include power states, connection states, and performance metrics such as CPU and disk use.
- You can monitor events that occur in response to operations occurring with a managed object in the inventory or vCenter itself. For example, an event is recorded each time a VM (which is a managed object) is cloned, created, deleted, deployed, and migrated.

Technet24

7-84 Defining the Alarm Rule: Event-based Trigger

A trigger can monitor events that occur in response to operations occurring on a managed object, for example:

- The health of a host's hardware changes.
- A license expires in the data center.
- A host leaves the distributed switch.

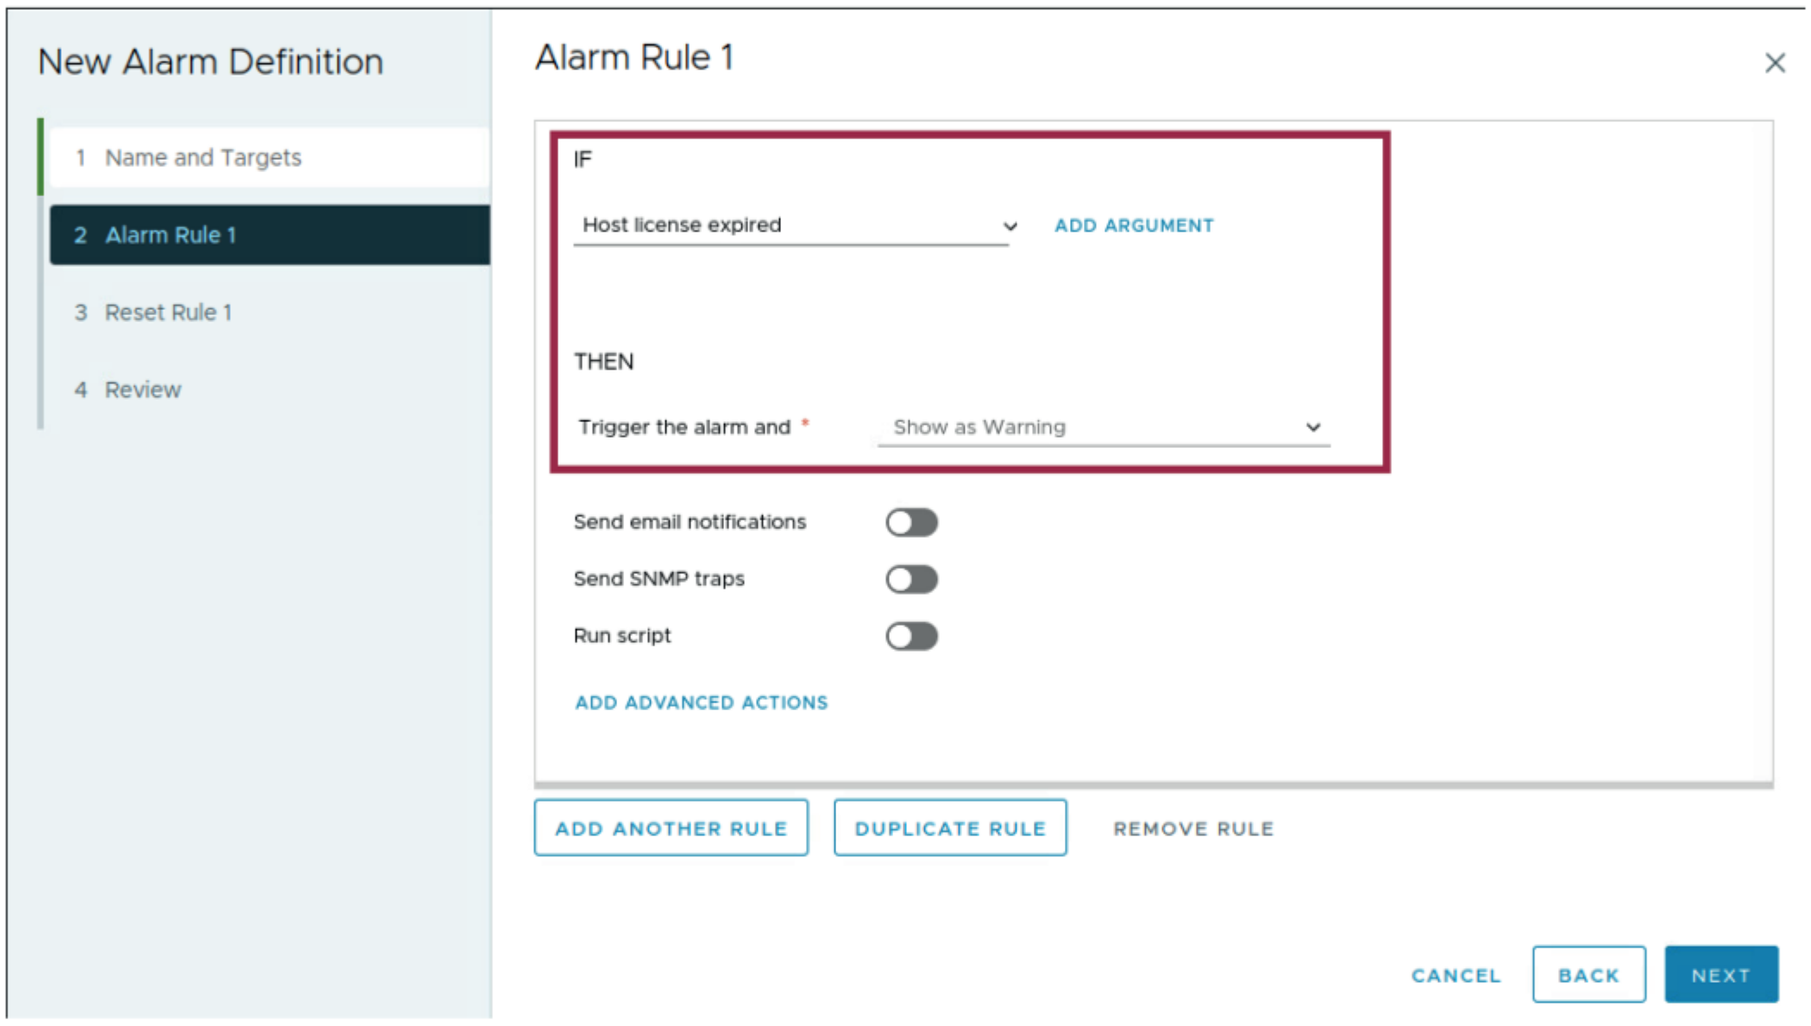

7-85 Defining the Alarm Rule: Setting the Notification

You configure the notification method to use when the alarm is triggered. The methods are sending an email, sending an SNMP trap, or running a script.

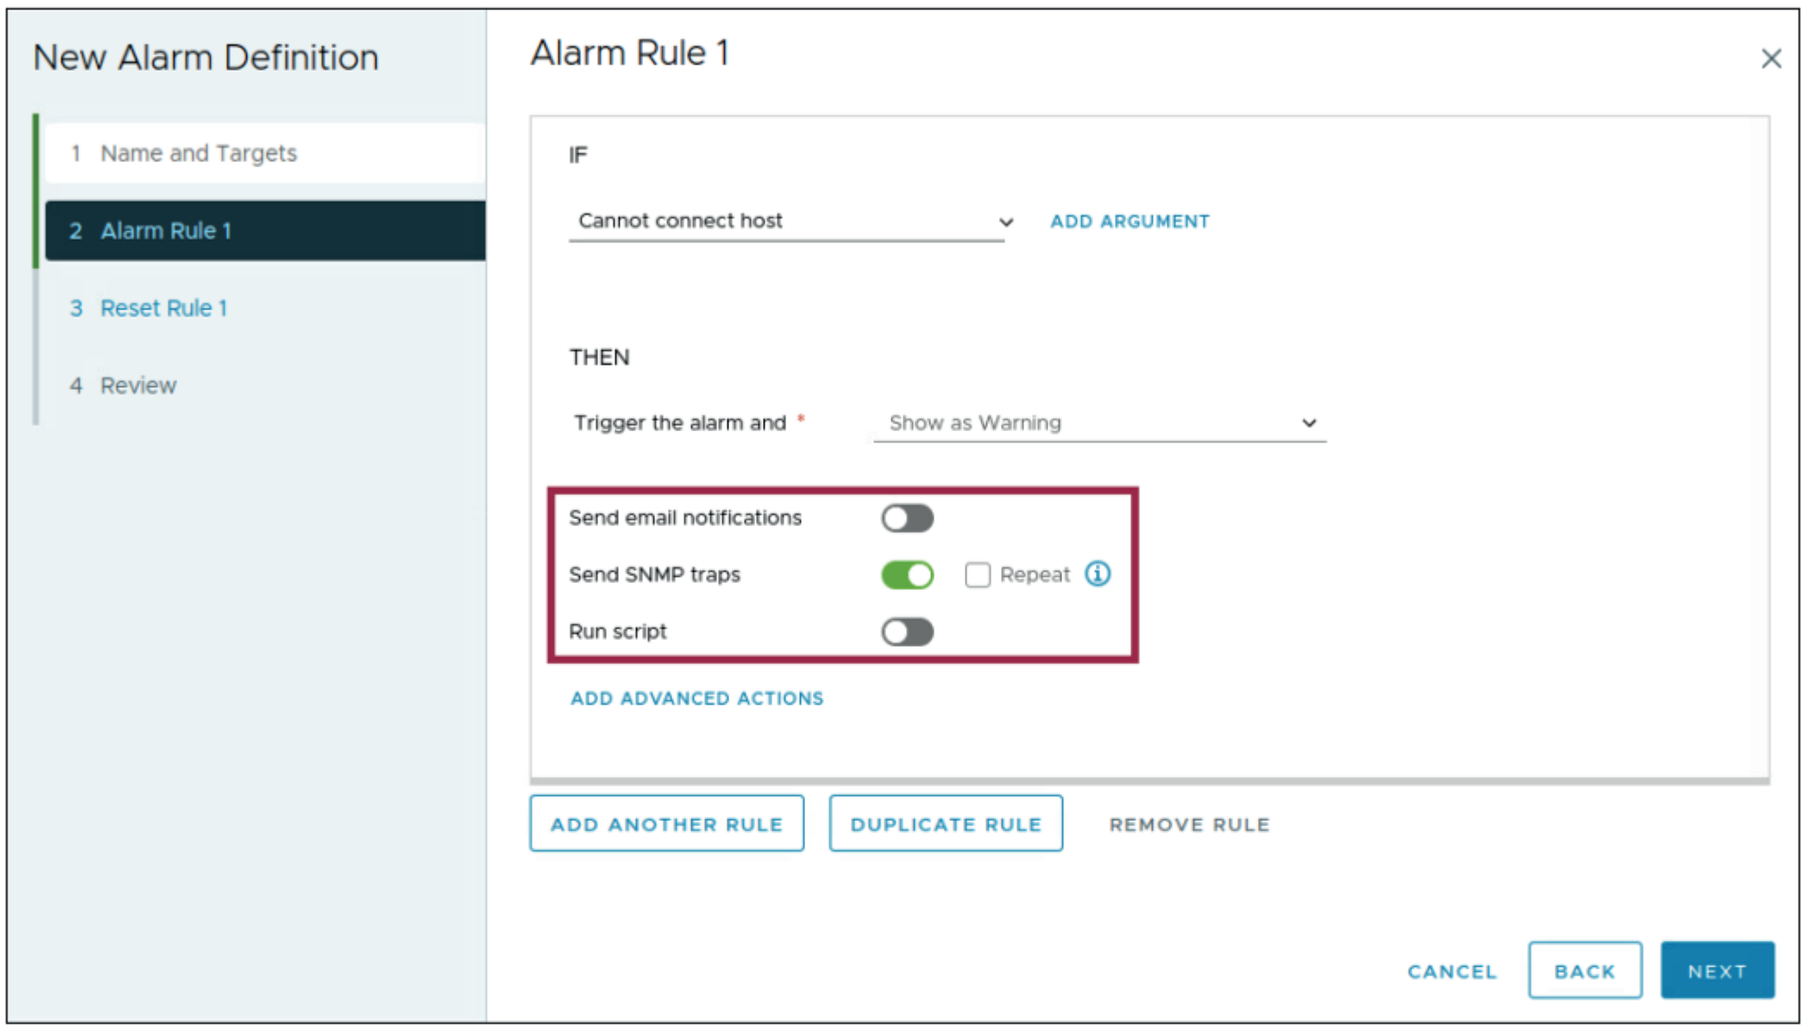

Technet24

7-86 Defining the Alarm Rule: Add Advanced Options

You can define the advanced actions for virtual machine and hosts. There are different sets of advanced actions based on target types of virtual machines and hosts.

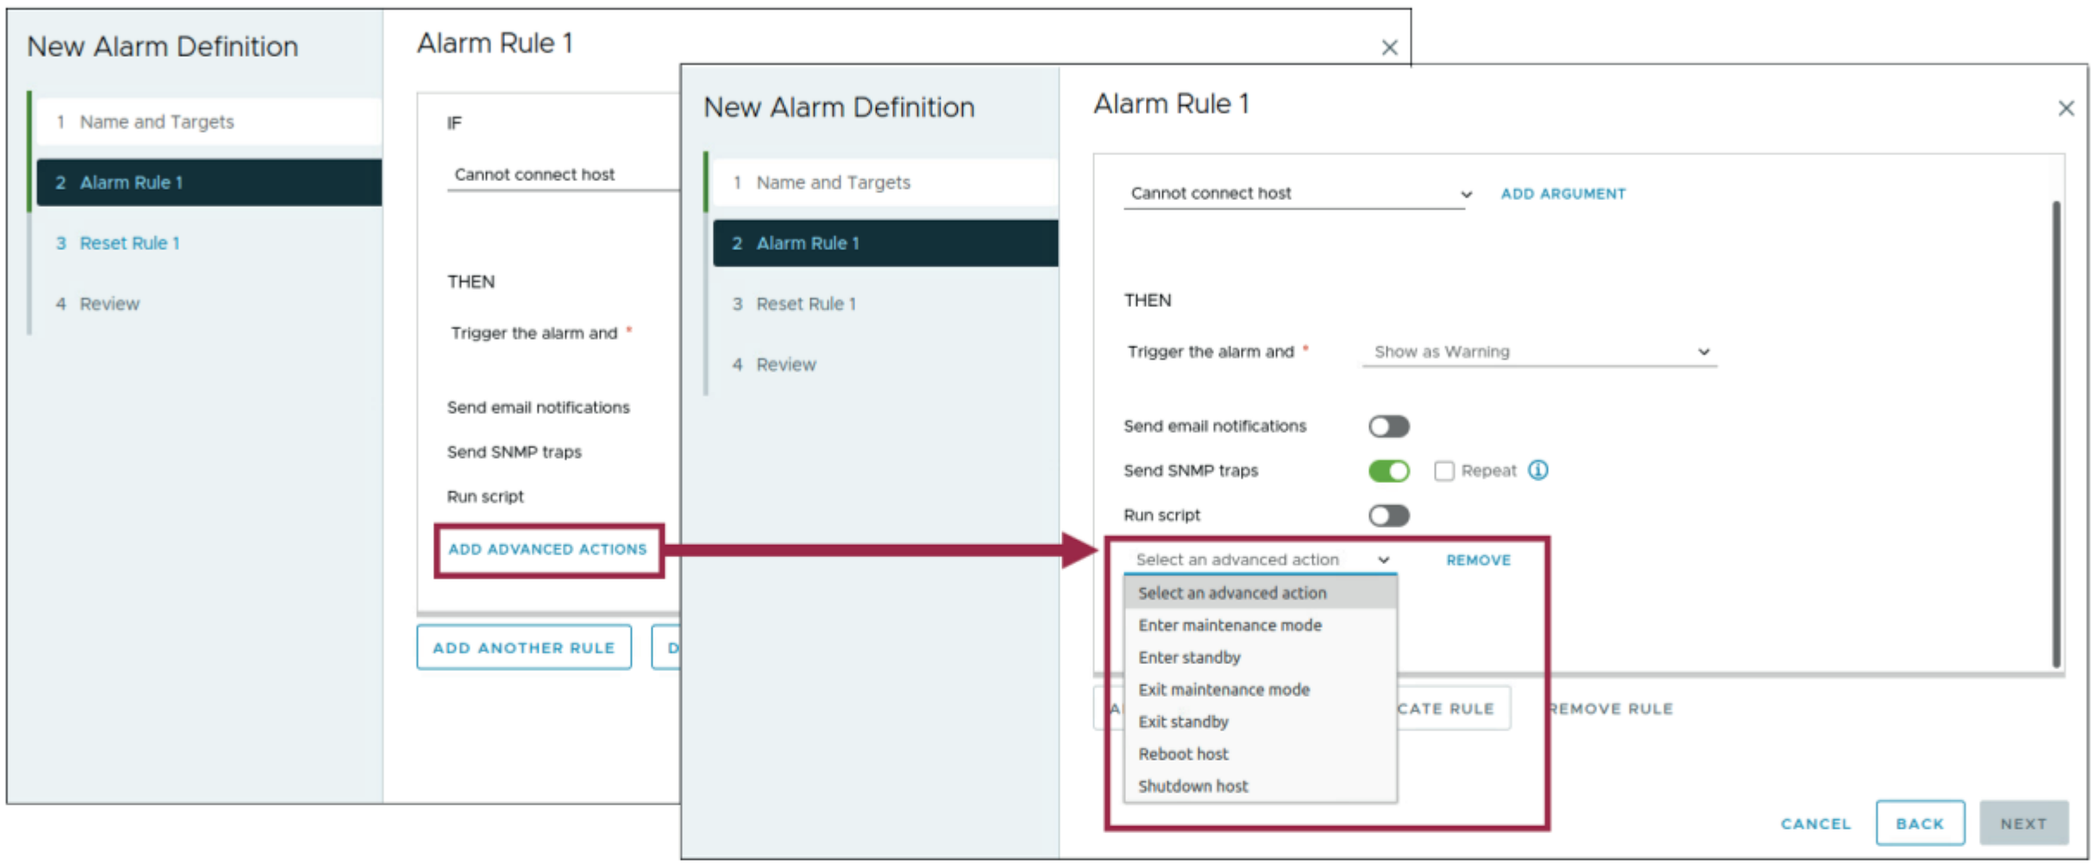

7-87 Defining the Alarm Reset Rules

You can select and configure the events, states, or conditions to reset the alarm to normal.

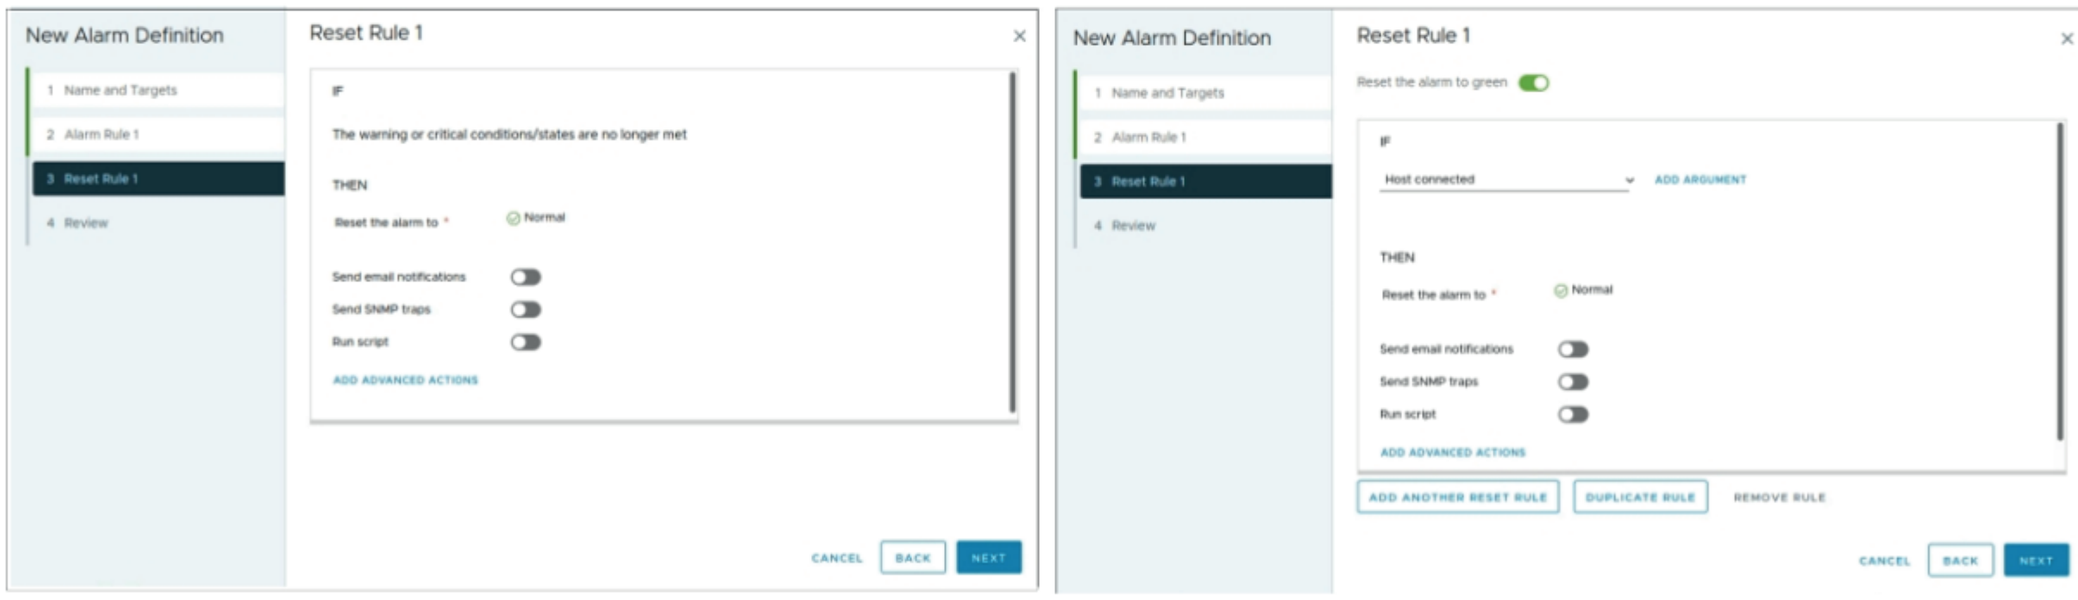

Sometimes, as in this example on your left, you can access only one option to reset the alarm.

On the right, you have an event based alarm, which is more configurable.

7-88 Activating the Alarm

On the Review page, the new alarm definition is activated by default.

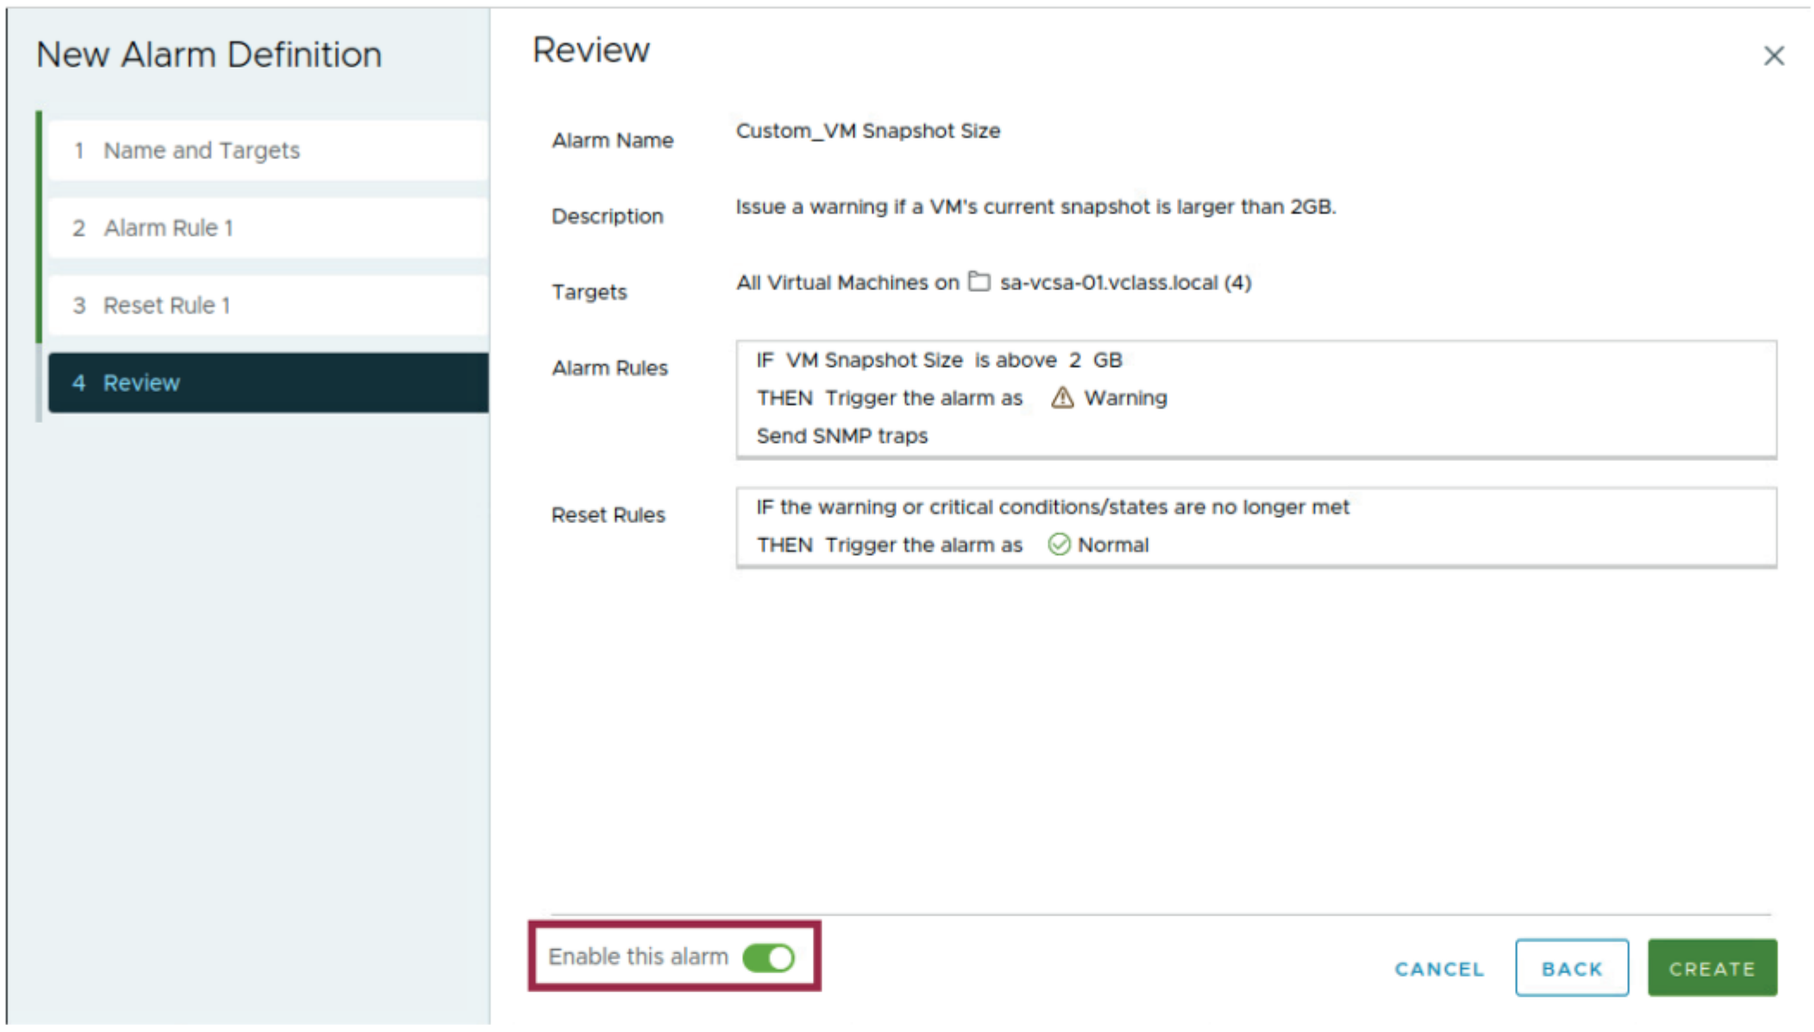

7-89 Triggered Alarms

When triggered, an alarm appears in the vSphere Client.

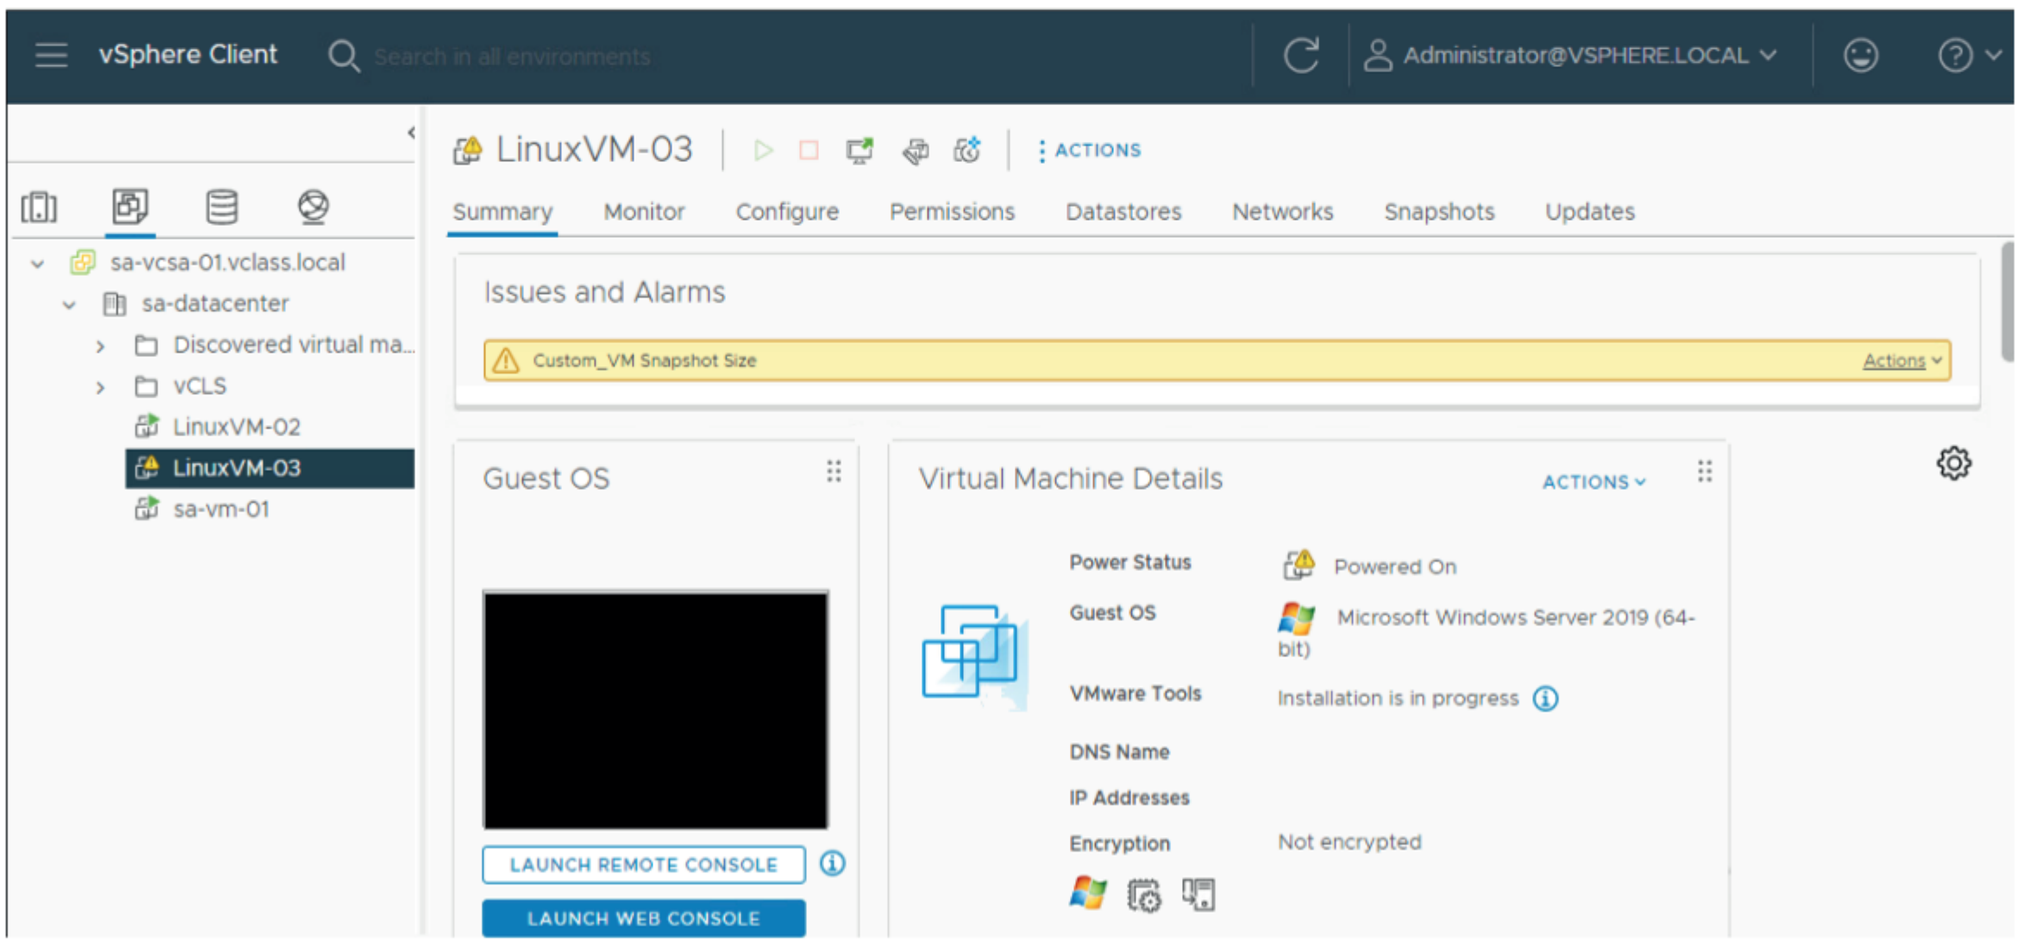

Technet24

7-90 Configuring vCenter Notifications

If you use email or SNMP traps as the notification method, you must configure vCenter to support these notification methods.

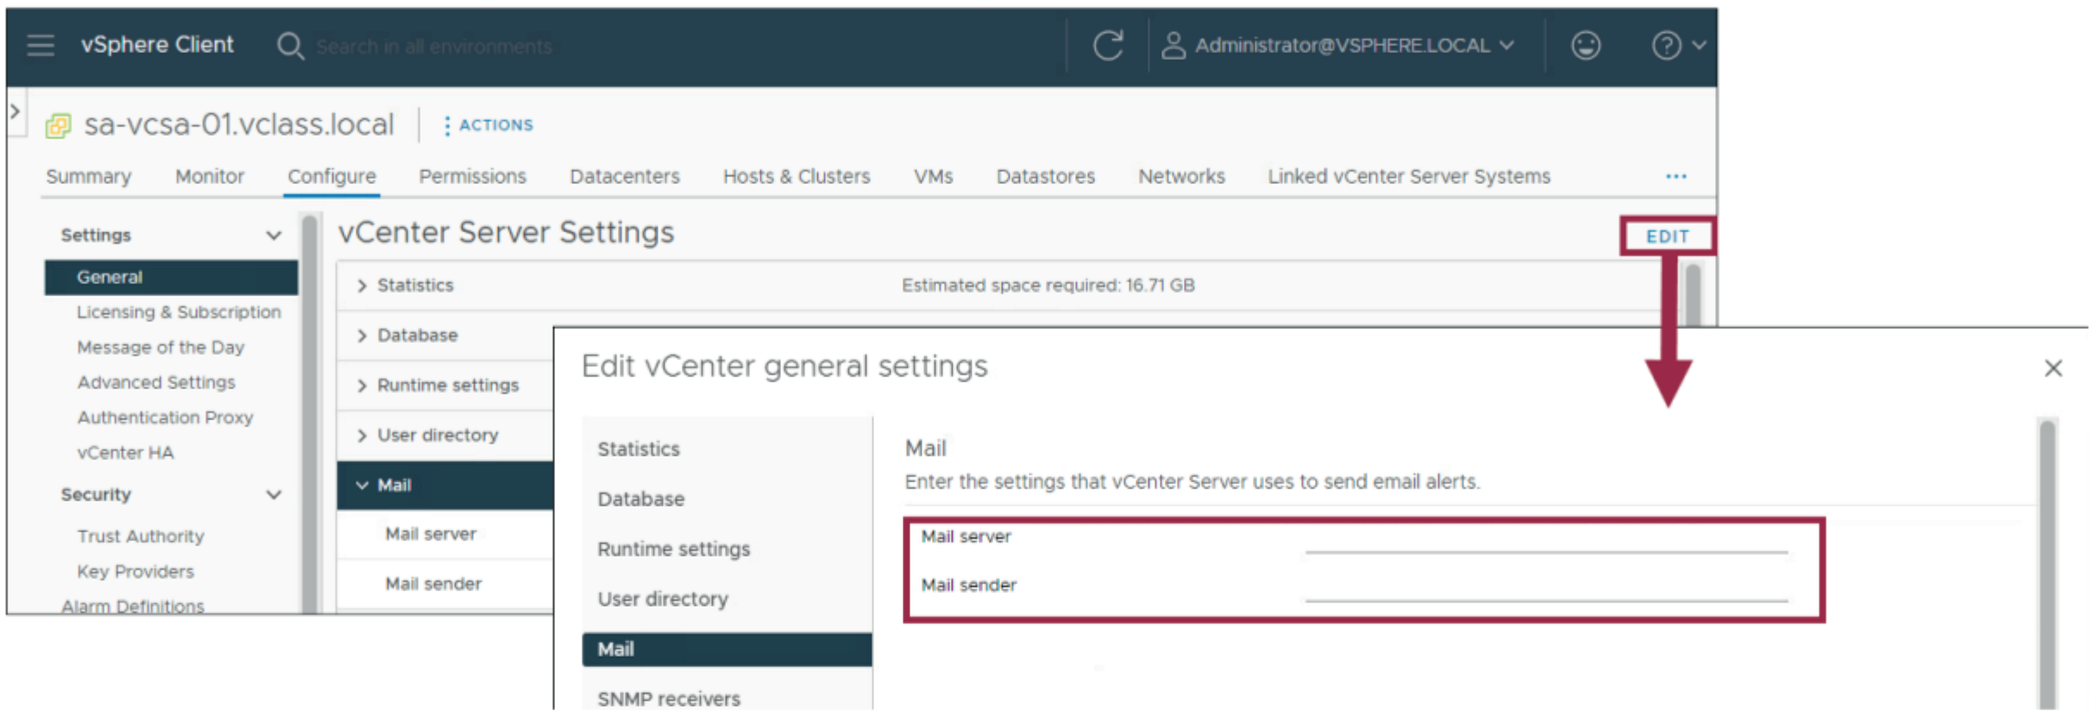

To configure email, specify the mail server FQDN or IP address and the email address of the sender account.

You can configure up to four receivers of SNMP traps. They must be configured in numerical order. Each SNMP trap requires a corresponding host name, port, and community.

7-91 Lab 17: Using Alarms

Create alarms to monitor virtual machine events and conditions:

1. Create a Virtual Machine Alarm to Monitor a Condition
2. Trigger the Virtual Machine Alarm
3. Create a Virtual Machine Alarm to Monitor an Event
4. Trigger the Virtual Machine Alarm
5. Deactivate Virtual Machine Alarms
6. Knowledge Check

7-92 Review of Learner Objectives

- Use predefined alarms in vCenter
- View and acknowledge alarms
- Create custom alarms

Technet24

7-93 Lesson 5: Proactive Management with VMware Skyline

7-94 Learner Objectives

- Describe the benefits and capabilities of VMware Skyline
- Explain how VMware Skyline works
- Identify the types of health information provided by Skyline Health
- Recognize uses for Skyline Advisor Pro

7-95 About VMware Skyline

VMware Skyline is a support technology that provides predictive analysis and proactive recommendations to help avoid problems. VMware Skyline provides the following benefits:

- Issue avoidance:
  - Proactively identifies potential issues based on environment-specific configuration, details, and usage.
  - Resolves issues before they occur, improving environment reliability and stability.
- Shortens time to resolution:
  - Environment-specific, data-driven analytics accelerate problem resolution.
- Personalized recommendations:
  - Resolution is specific to your environment.
- No additional cost:
  - You receive additional value with your current support subscription (Basic, Production, or Premier support).

VMware Skyline shortens the time it takes to resolve a problem so that you can get back to business quickly. The VMware Technical Support engineers can use VMware Skyline to view your environment's configuration and the specific, data-driven analytics to help speed up problem resolution.

7-96 VMware Skyline Family

The VMware Skyline family includes Skyline Health and Skyline Advisor Pro.

Skyline Health
All VMware Customers
Key capabilities: • vSphere and vSAN findings • Available in the vSphere Client • Supports vSphere 6.7 and later

Skyline Advisor Pro
Production and Premier Support Customers
Key capabilities: • Supports vSphere, vSAN, NSX for vSphere, vRealize Operations Manager, and VMware Horizon • Supports vSphere 5.5 and later • Identification of VMware Validated Design, VxRail, and VMware Cloud Foundation deployments • Automates log transfers with Log Assist • Uses cloud-based ID and access
Premier Support Customers
Key capabilities: • Advanced findings and reporting • Tailored remediation plans

With Basic Support, you can access Skyline findings and recommendations for vSphere and vSAN by using Skyline Health in the vSphere Client (version 6.7 and later).

With Production or Premier Support, you can use Skyline Advisor Pro and the full functionality of Skyline (including Log Assist).

Technet24

With Premier Support, you receive additional Skyline features that are not available with Production Support, for example:

- An advanced set of findings and recommendations
- Scheduled and custom operational summary reports that provide an overview of the proactive findings and recommendations
- All additional benefits of Premier Support, including a designated support team, direct access to senior-level technical support engineers, assistance with multivendor troubleshooting and onsite support services, such as Mission Critical Support (MCS), Healthcare Critical Support (HCS), and Carrier Grade Support (CGS)

Skyline supports vSphere, NSX for vSphere, vSAN, VMware Horizon, and vRealize Operations Manager. A Skyline management pack for vRealize Operations Manager is also available. If you install this management pack, you can see Skyline proactive findings and recommendations within the vRealize Operations Manager client.

The identification and tagging of VxRail and VMware Validated Design deployments help you and VMware Technical Support to better understand and support multiproduct solutions.

Skyline identifies all ESXi 5.5 objects within a vCenter instance and provides additional information in VMware knowledge base article 51491 at https://kb.vmware.com/s/article/51491. This article details the end of general support for vSphere 5.5.

For versions of vSphere, vSAN, NSX for vSphere, VMware Horizon, and vRealize Operations Manager that are supported by Skyline, see the Skyline Collector Release Notes at https://docs.vmware.com/en/VMware-Skyline/index.html.

7-97 Skyline Privacy and Security Policy

VMware is committed to keeping your data private and secure:

- The Skyline Collector appliance is separate from your environment:
  - Your production infrastructure is not exposed to the Internet.
- Collected data is forwarded to VMware using an encrypted channel:
  - The data is only temporarily stored in the Skyline Collector appliance.
- Data is warehoused in a secure back-end system in the United States and is managed by VMware:
  - Access to your data is restricted to permitted users only.
- Information is limited to data such as host names:
  - Data that could identify an individual is not collected.

Skyline collects data in a private and secure manner. Your production infrastructure is not exposed to the Internet. All collected data is tagged with legally identifiable customer information. No customer data is permanently stored in the Skyline Collector appliance or shared with third parties.

7-98 Skyline Health for vSphere (1)

Skyline Health provides findings for your vSphere environment using both hourly in-house healthchecks and online healthchecks

- Findings are based on telemetry data collected from your vSphere environment.
- Data is analyzed and then provided in the vSphere Client.
- Discovered findings can be related to stability and incorrect configurations in vSphere.

In vSphere 6.7 and later environments, you can use vSphere health and vSAN health services to monitor your vSphere 6.7 environments for configuration consistency, best practice adherence, vulnerability discovery, and proactive analytics. vSphere health and vSAN health services are part of VMware Skyline Health.

Technet24

7-99 Skyline Health for vSphere (2)

You access Skyline Health with the vSphere Client.

For Skyline Health for vCenter to function correctly:

- You must join the Customer Experience Improvement Program (CEIP).
- Your vCenter instance must have Internet connectivity.

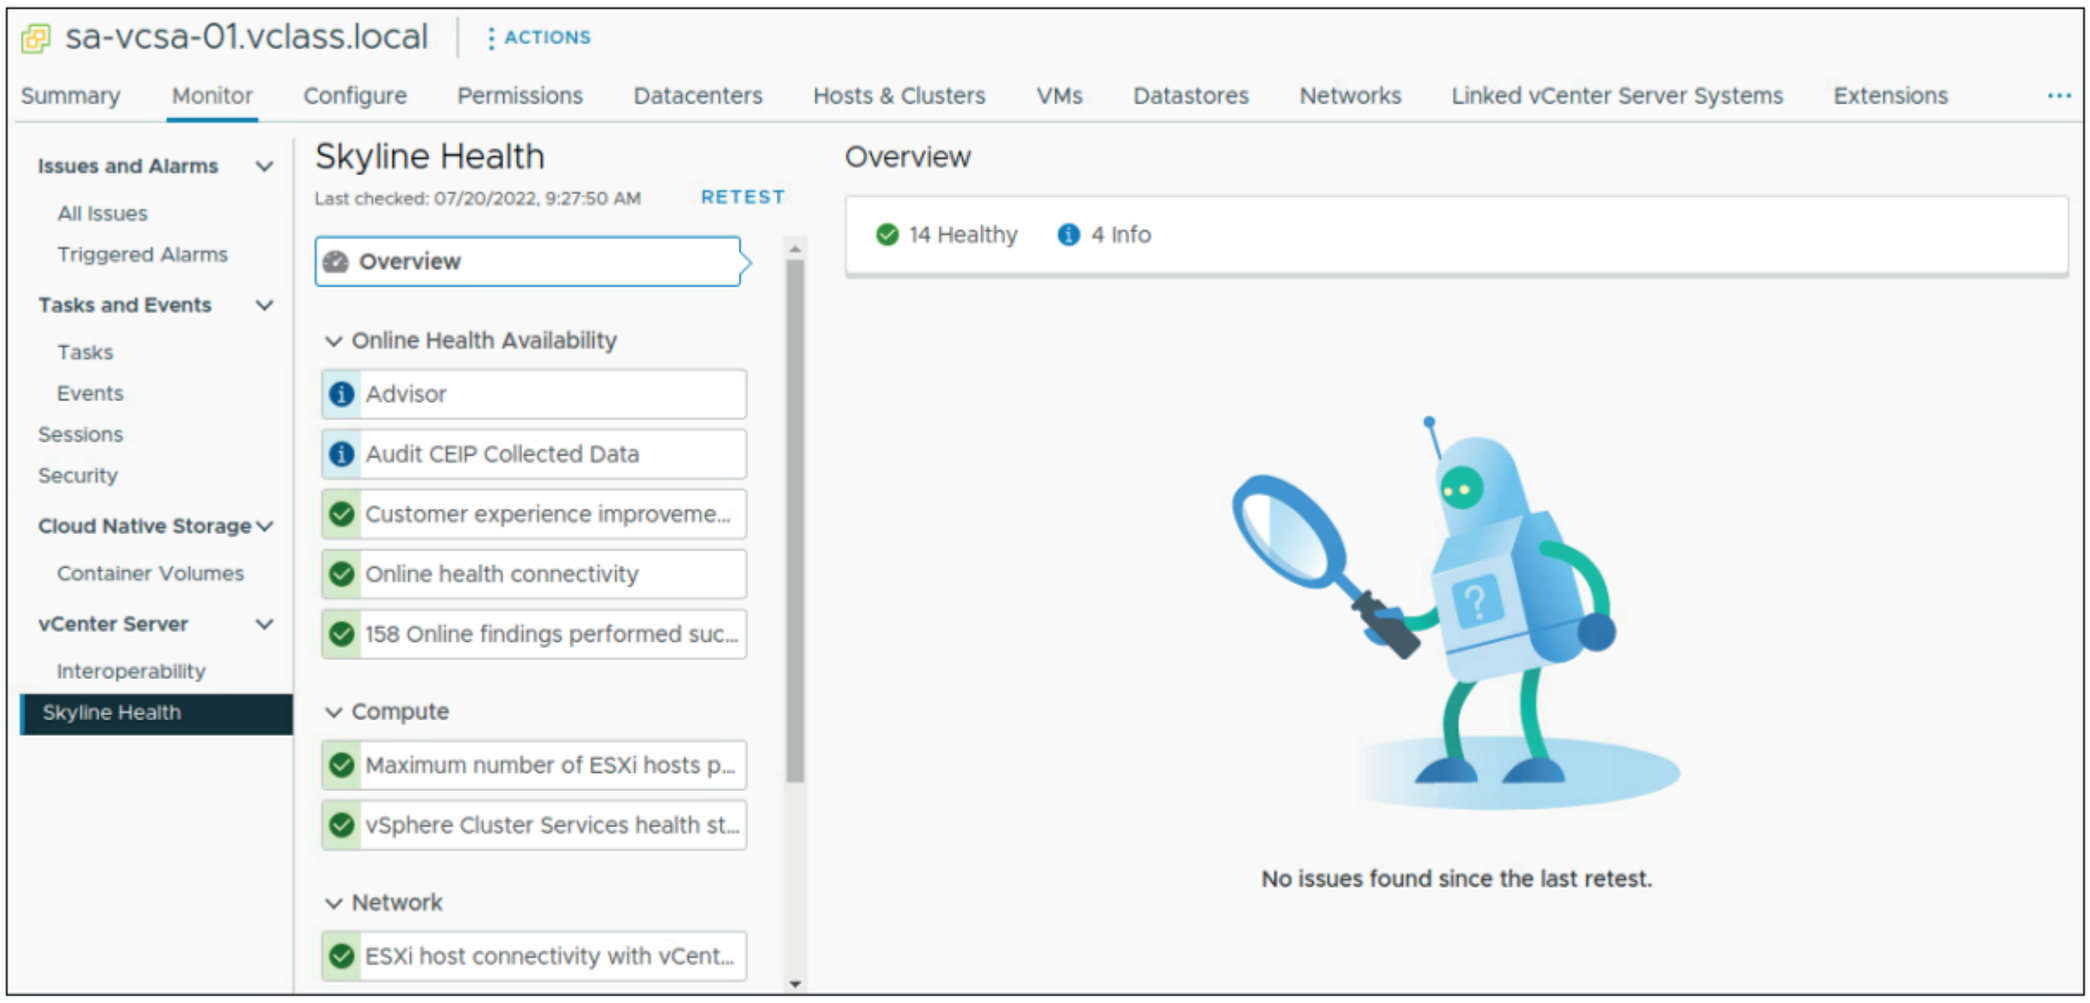

To view vSphere health information, click the **Monitor** tab for a vCenter instance and select **Skyline Health**.

To view vSAN health information, click the **Monitor** tab for a vSAN cluster or vSAN enabled ESXi host and select **Skyline Health**.

7-100 Skyline Health for vSAN

For vSAN online healthchecks, Skyline Health findings have the following characteristics:

- Findings are based on data from thousands of vSAN deployments.
- Findings draw on VMware knowledge base articles and best practices.
- Discovered findings can be related to misconfigurations and the general health and well-being of hosts in a vSAN cluster.

vSAN Health includes in-product health checks that do not require you to join CEIP. Nor do they require internet connectivity.

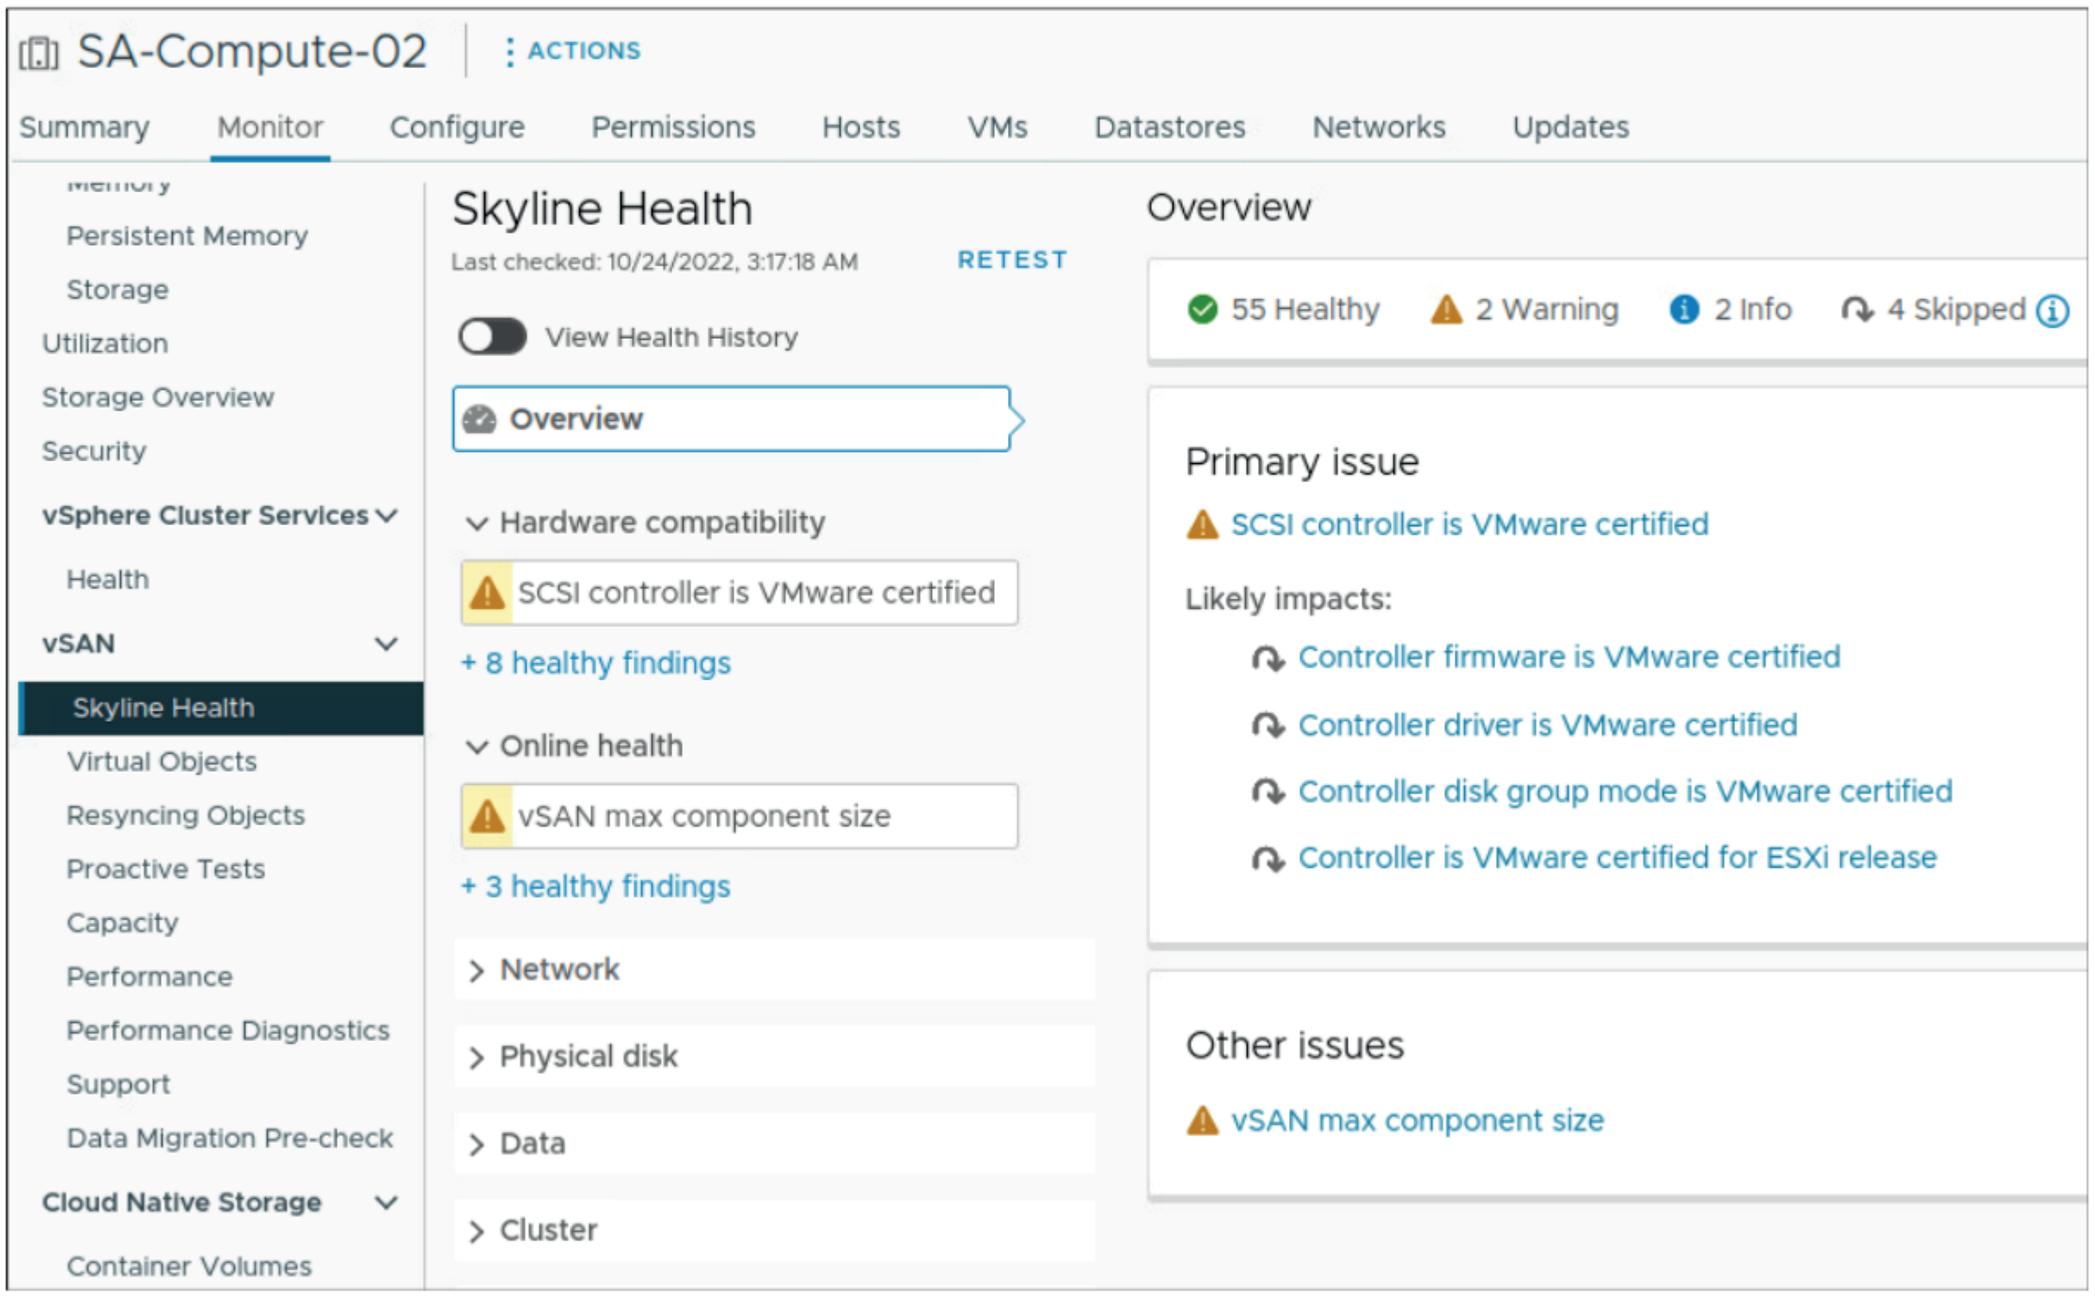

Technet24

7-101 Skyline Components: Collector and Advisor

Skyline Advisor Pro and Skyline Collector work together to provide proactive support:

- Skyline Collector is a standalone virtual appliance that collects diagnostic data about your vSphere, NSX for vSphere, vSAN, VMware Horizon, and vRealize Operations Manager environments.
- Skyline Advisor Pro is a web-based portal where you can view proactive findings and recommendations on demand.

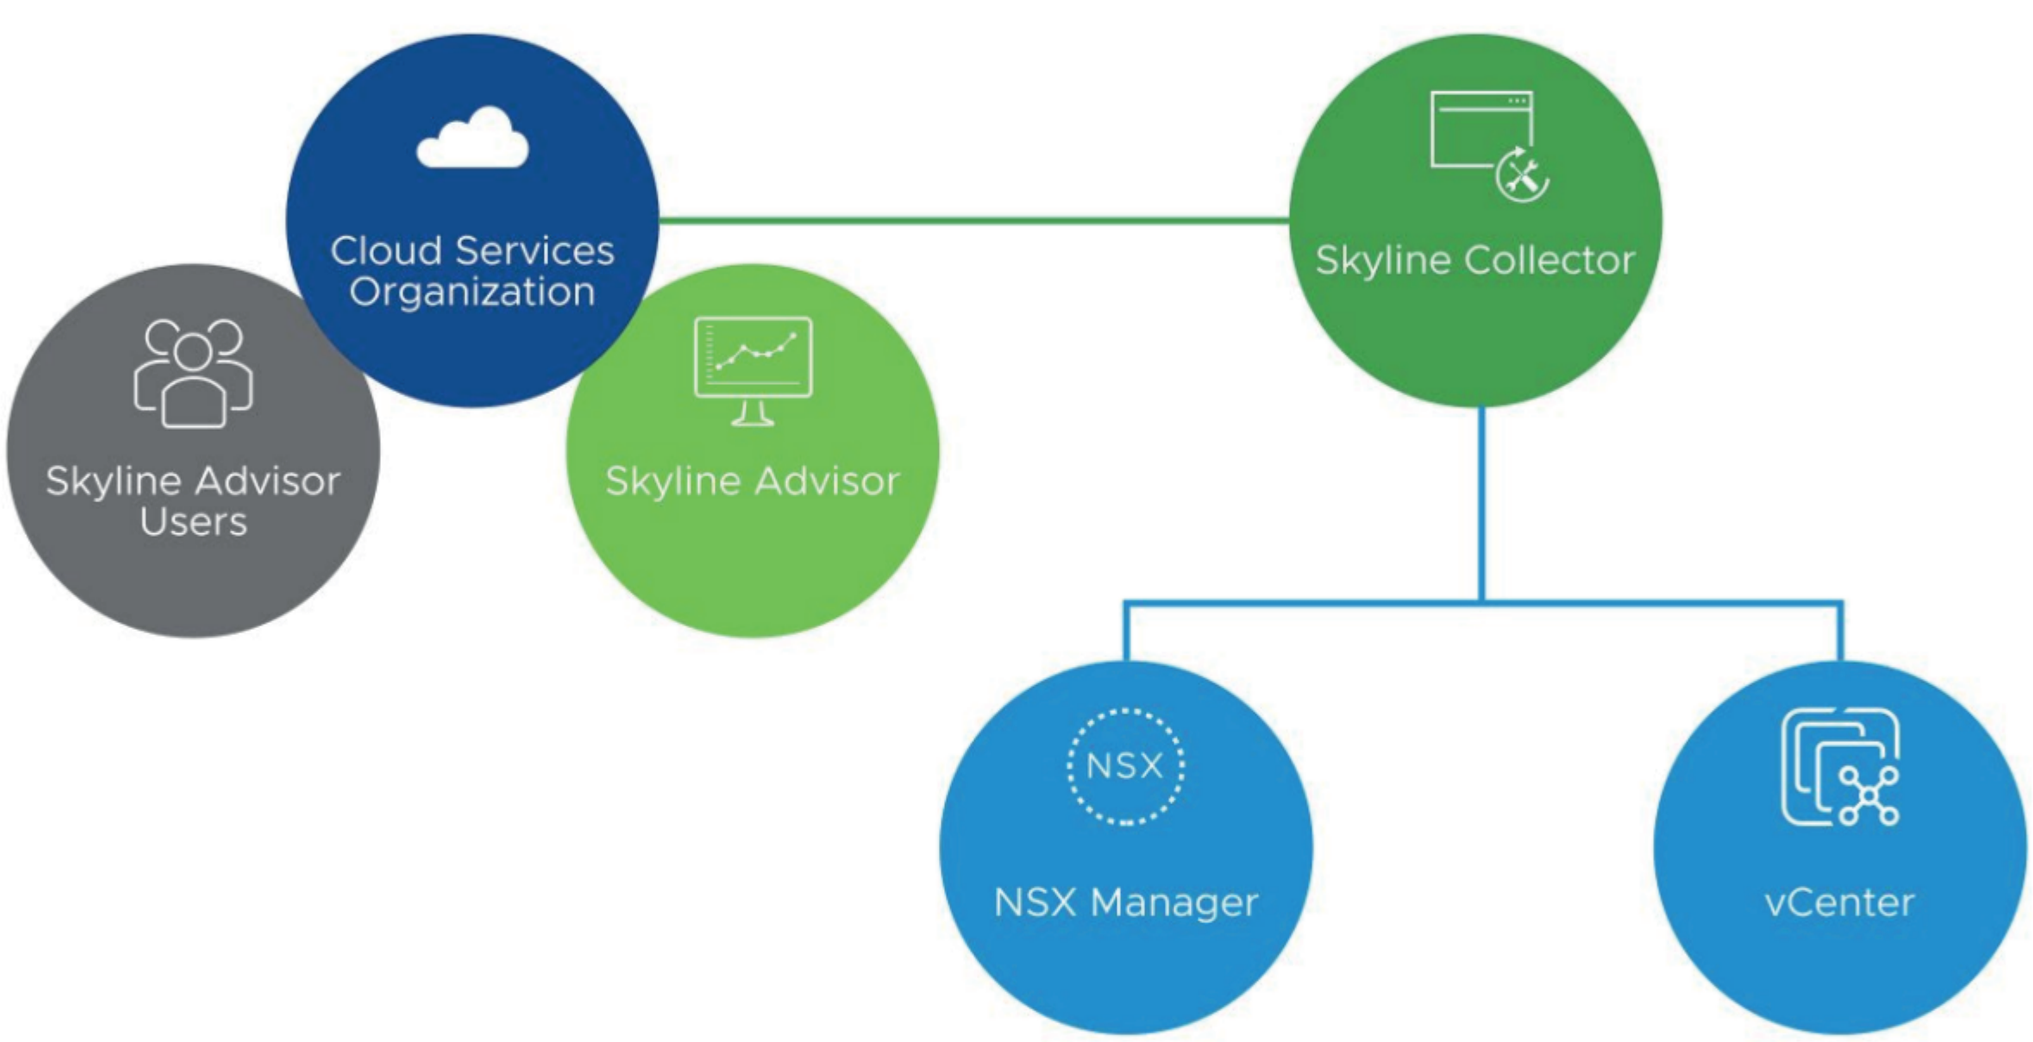

You can access Skyline Advisor Pro directly using this link: https://skyline.vmware.com/advisor.

You can also access Skyline Advisor Pro from the list of services in the console of your Cloud Services organization.

7-102 VMware Skyline Networking Design

This sample networking design has the following features:

- Skyline Collector instance has routed access to the management network.
- An upstream router provides external connectivity.
- An HTTP proxy server can be used for outbound connectivity.

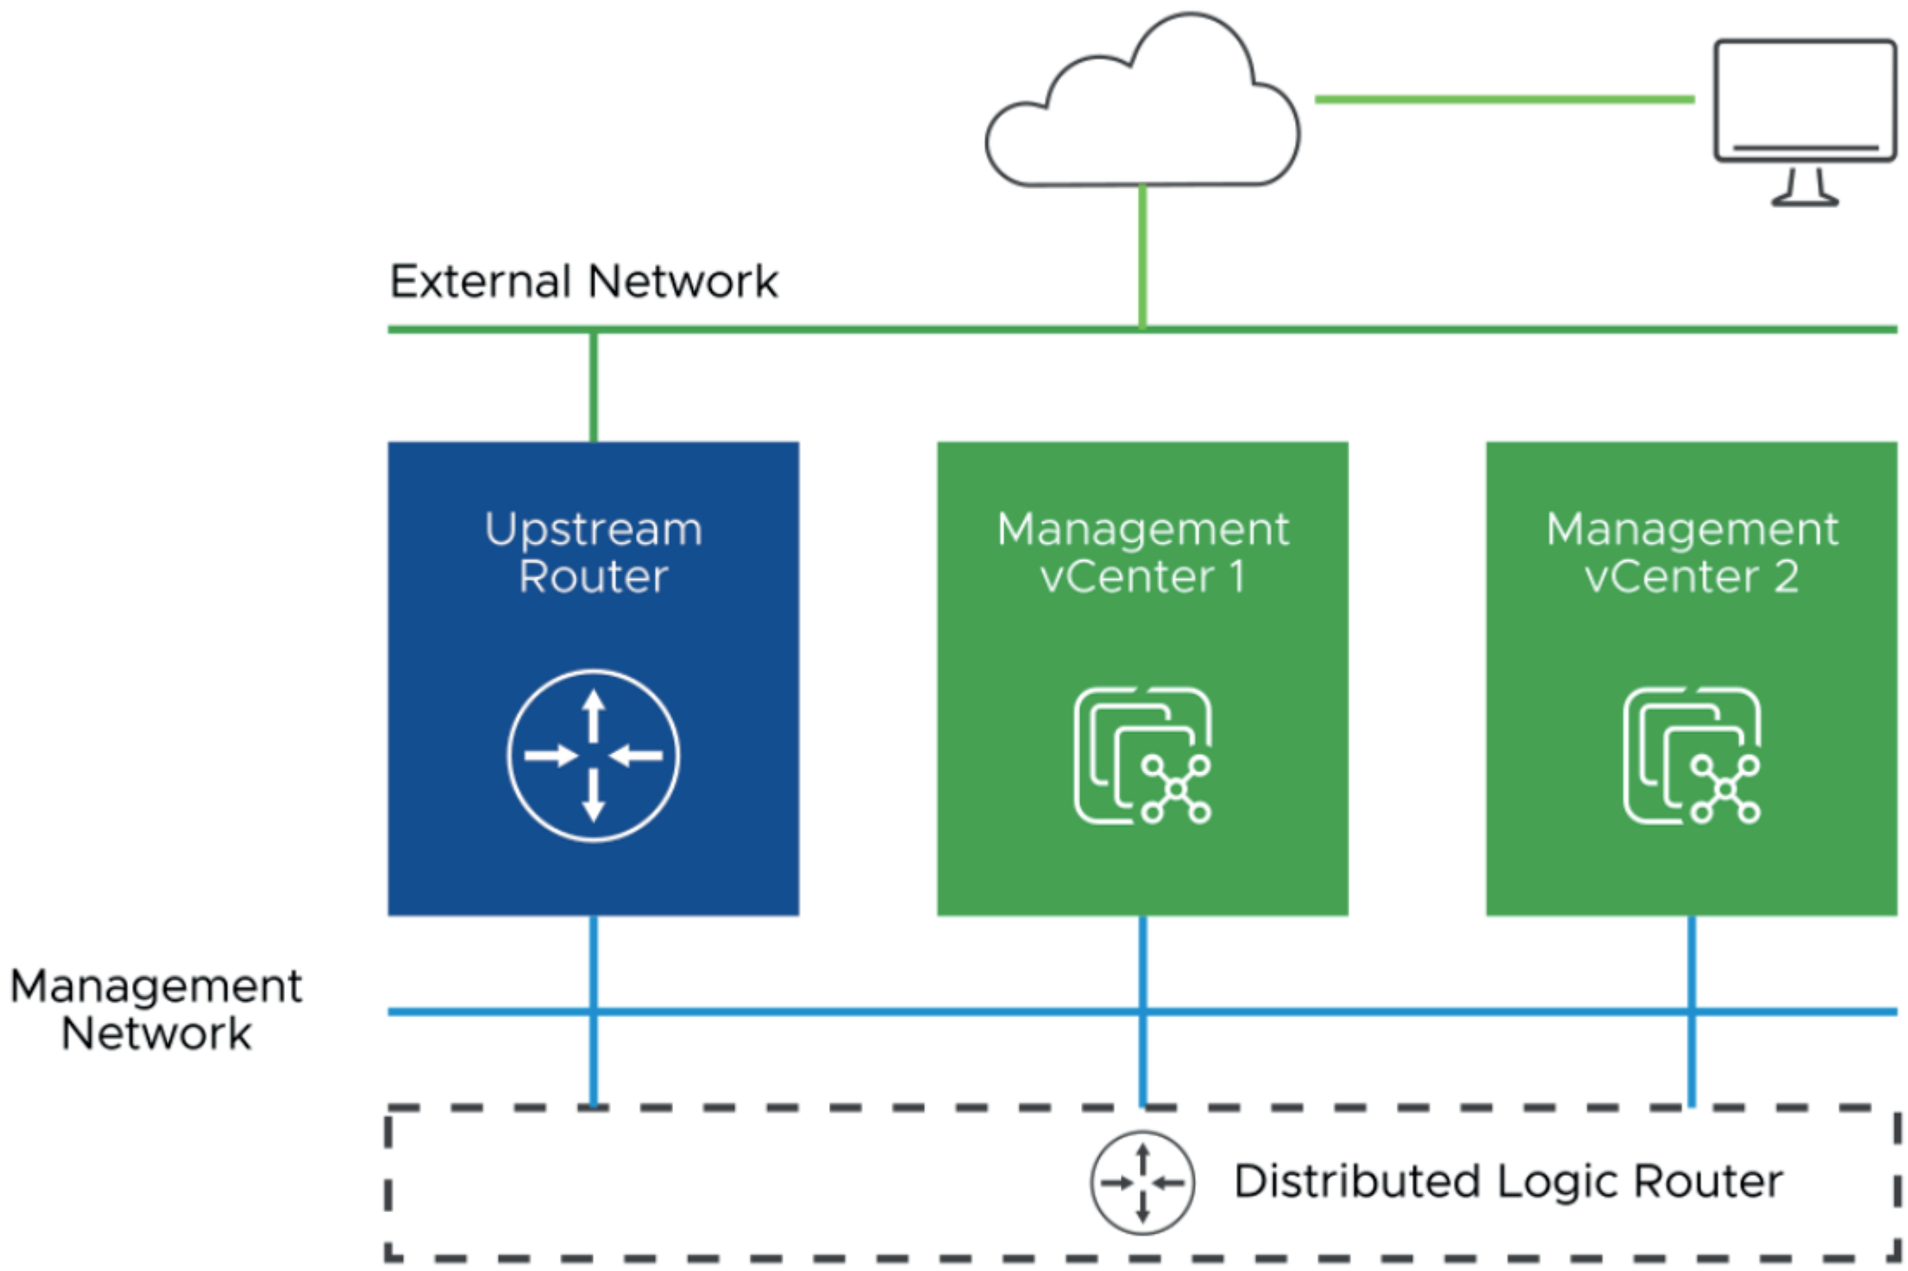

Technet24

7-103 How VMware Skyline Works

1. Install and configure:
   - You install the Skyline Collector appliance into your environment.
   - You configure Skyline Collector with one or more endpoints (for vCenter, NSX Manager, Horizon View, vRealize Operations).
2. Collect:
   - Skyline Collector starts collecting product usage data and capturing patterns, events, trends, configuration, design-compliance, and cross-product information.
3. Analyze:
   - VMware Skyline analyzes the environment details and product usage and compares this data against VMware best practices and the VMware knowledge base.
4. Act:
   - VMware Skyline provides proactive findings and recommendations through the Skyline Advisor Pro, a web-based portal.

Skyline Collector is a lightweight appliance that takes 15 to 30 minutes to install. The installation involves adding Skyline Collector to a new or existing Cloud Services organization, associating your support entitlement with the organization, downloading the Skyline Collector appliance from My VMware, and using the vSphere Client to deploy the appliance from an OVF template.

For complete details about how to install and configure Skyline Collector, see the *VMware Skyline Collector User Guide* at https://docs.vmware.com/en/VMware-Skyline-Collector/3.2/user-guide.pdf.

Skyline Collector gathers and aggregates product usage information such as configuration, feature, and performance data, while listening for changes and events within the customer's environment. This information is encrypted and then sent to VMware for analysis.

7-104 Skyline Advisor Pro

Skyline Advisor Pro is a self-service web application with which you can view proactive findings and recommendations on-demand, within a web browser.

To activate Skyline Advisor Pro, upgrade to Skyline Collector 3.0.

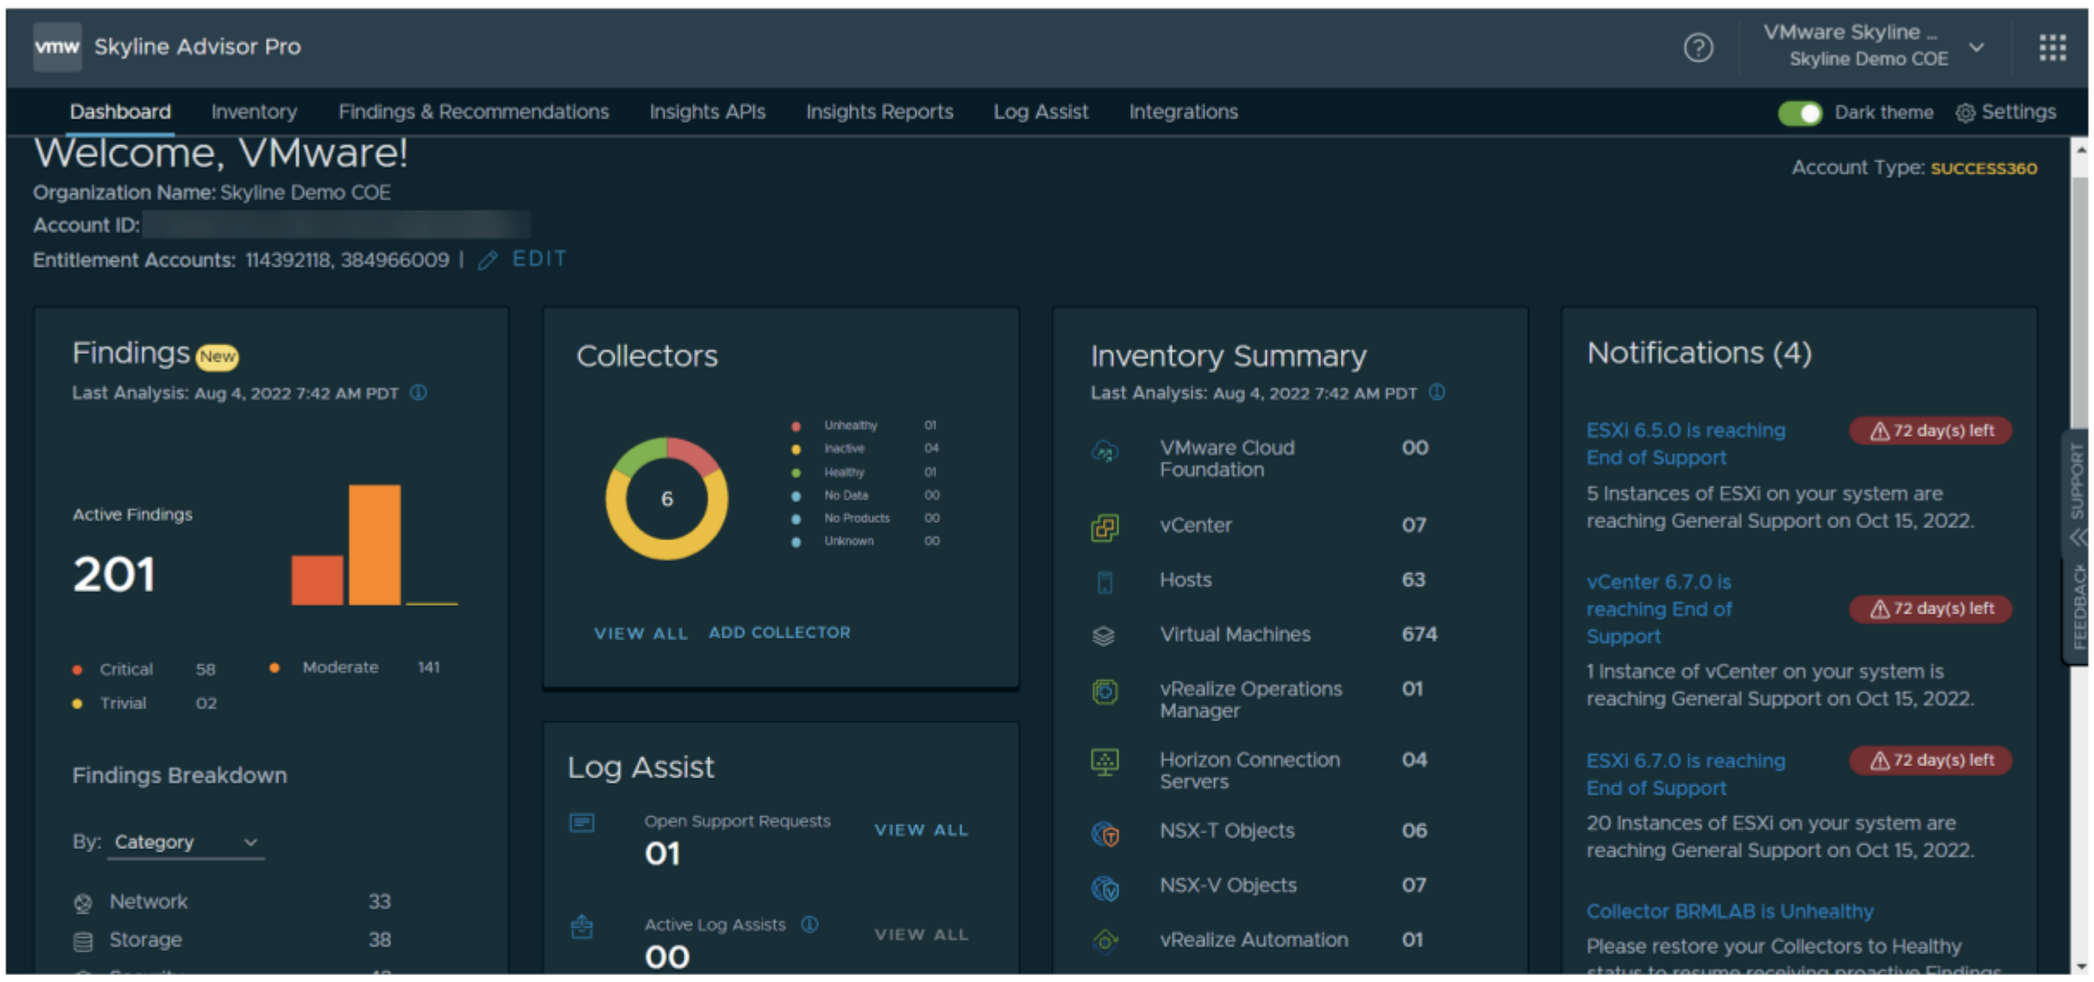

The following information is available within Skyline Advisor Pro.

- Account Details, including total number of vCenters, ESXi hosts and virtual machines being analyzed for potential issues.
- Deployed Skyline Collectors.
- Inventory details for VMware vSphere, NSX-V, Horizon 7, and vRealize Operations Manager.
- Download scheduled and custom Operational Summary Reports (OSR) (Premier Support customers only).
- Streamlined support log bundle upload capabilities with Skyline Log Assist.

Technet24

7-105 Skyline Advisor Pro: Proactive Findings

Proactive findings are at the heart of VMware Skyline.

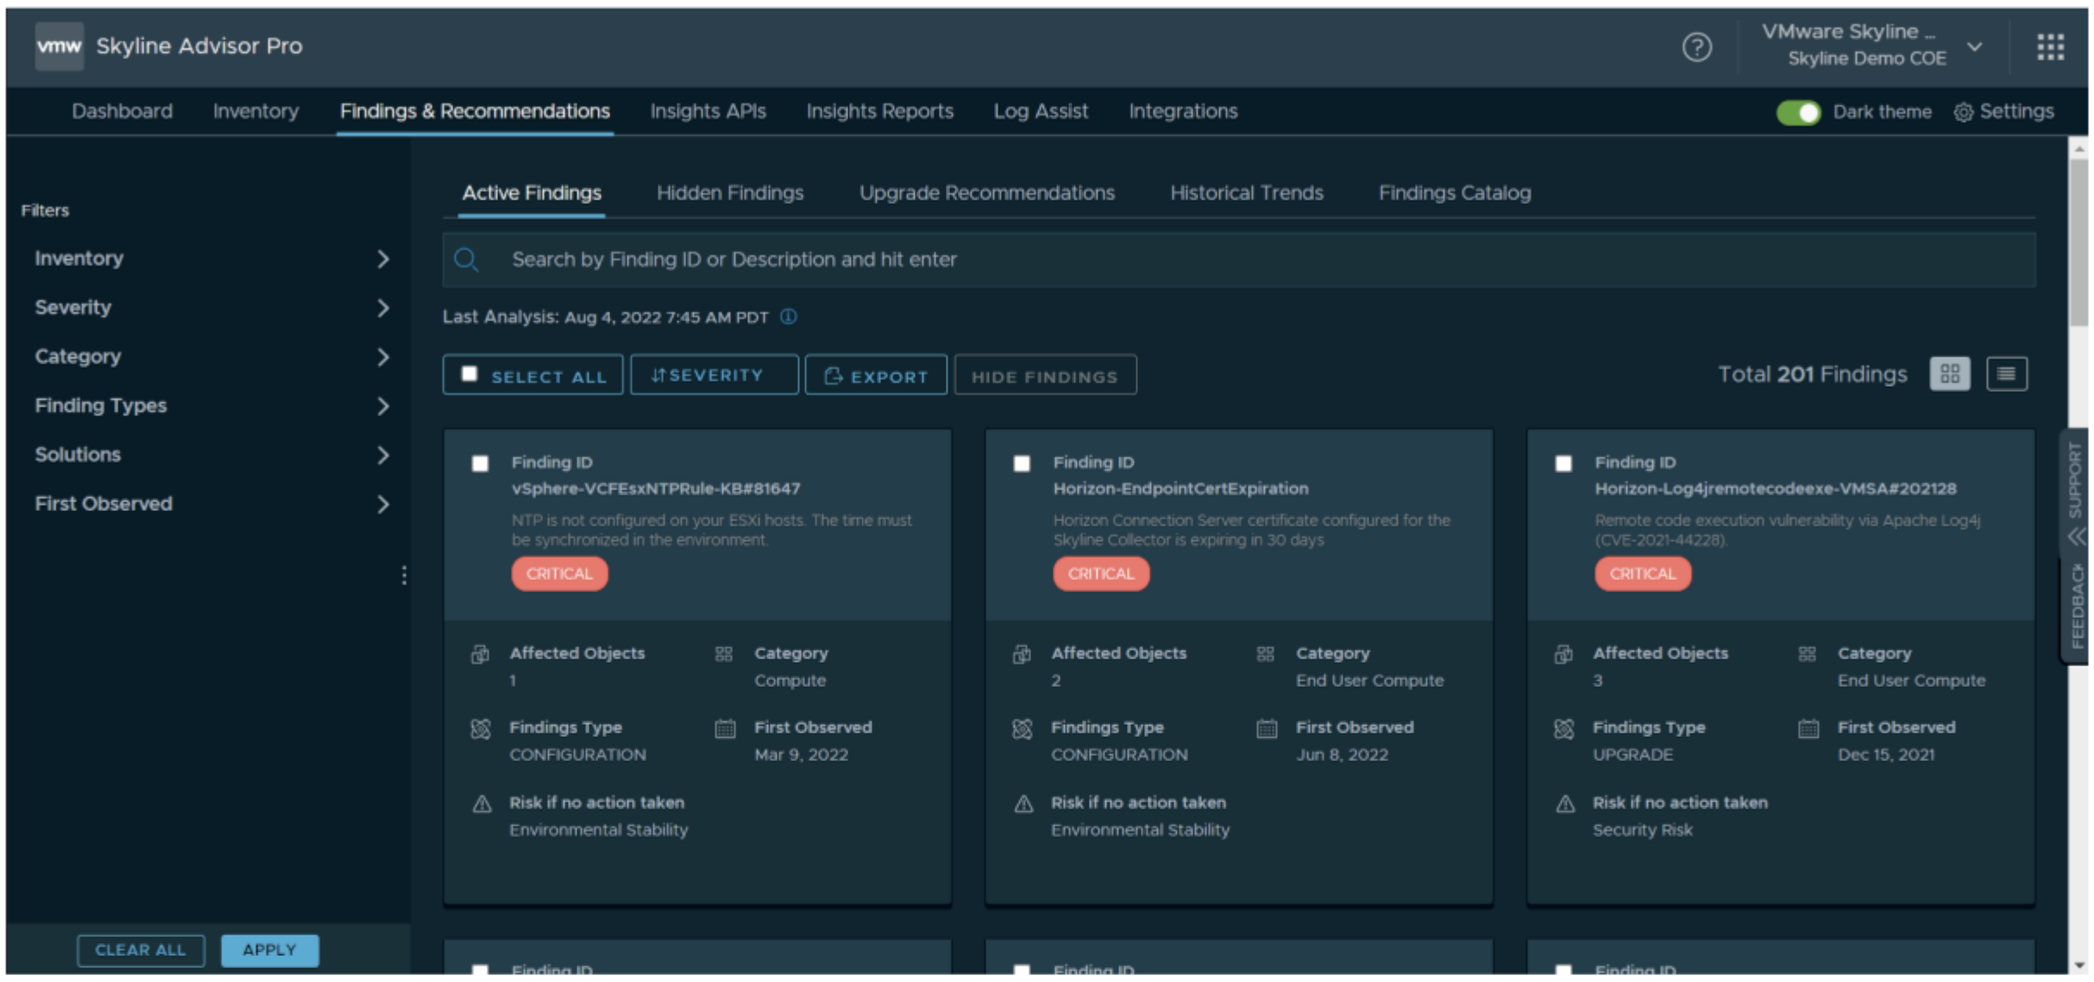

On the **Findings & Recommendations** tab, you can find all potential issues in your environment. These issues should be remediated to keep your environment secure, stable, and productive. Findings are ranked by severity, and you can filter by several different options.

7-106 Log Assist

When using VMware Skyline Log Assist, you do not need to upload your log files manually.

Support log bundles can be automatically uploaded to VMware Technical Support.

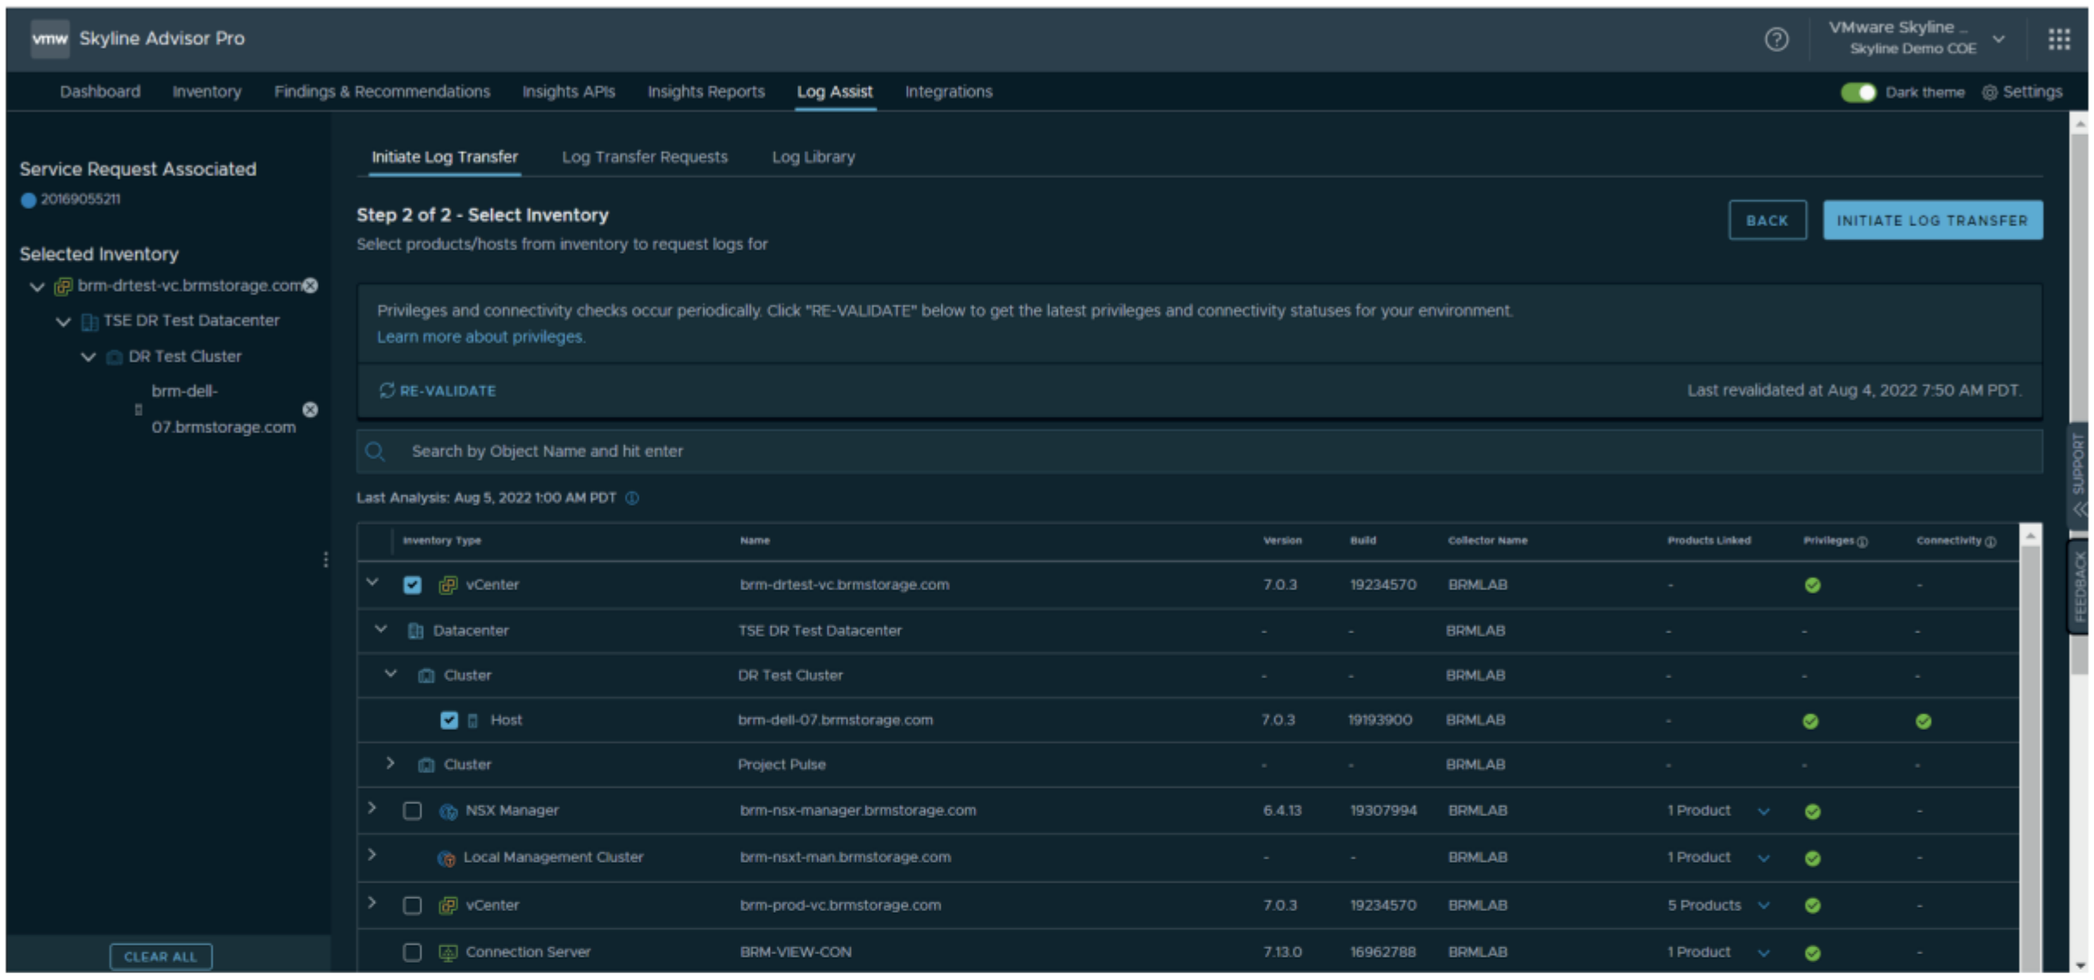

Log Assist helps to streamline the time to resolution from when you open a support request.

VMware Technical Support engineers can request the log file bundles associated with your support request. When you approve their request, the log files are automatically uploaded to the Technical Support team, and they can begin troubleshooting right away.

You can also use Log Assist to initiate the automatic upload yourself. Log Assist eliminates the time-consuming process of manually uploading log files using SSH File Transfer Protocol (SFTP) or VMware Customer Connect .

7-107 Skyline Features Comparison by Support Level

Features	Basic Support	Production Support	Premier Support
Skyline Health	✓	✓	✓
Basic findings and recommendations	✓	✓	✓
Moderated Skyline community	✓	✓	✓
Skyline Advisor Pro		✓	✓
Log Assist		✓	✓
Proactive support for all supported products		✓	✓
Export of findings, recommendations, and affected objects to a comma-separated file		✓	✓
Advanced findings and recommendations			✓
Custom operational summary reports			✓
Guided remediation from VMware Global Support Services			✓

7-108 Integrated Skyline Health Diagnostics Plug-in

Introduced in vSphere 8 Update 1, the vSphere Client can be used to easily deploy the Skyline Health Diagnostics appliance and register it with the vCenter.

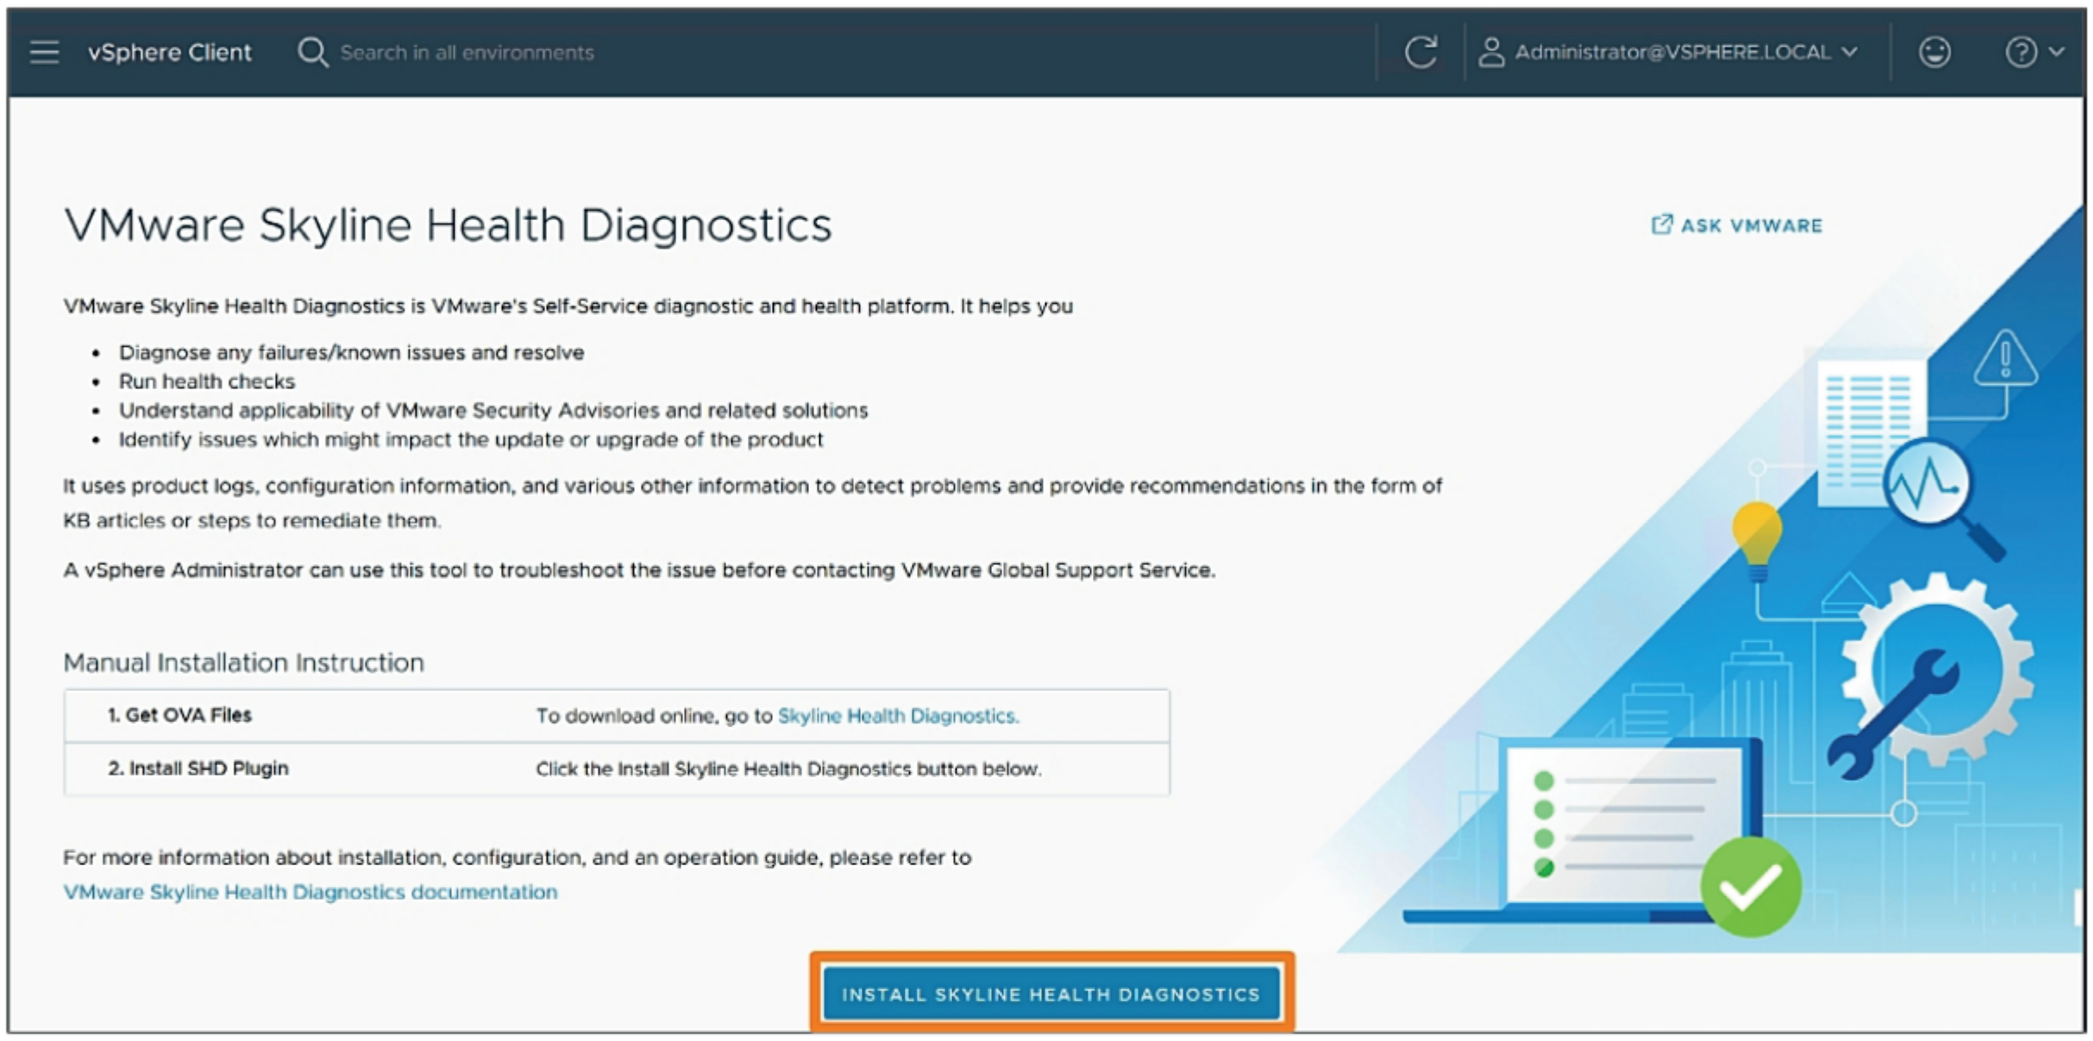

7-109 Benefits of Using VMware Skyline Health Diagnostics

Skyline Health Diagnostics provides many benefits:

- Based on symptoms, the tool provides a Knowledge Base article or remediation steps to resolve the issue.
- Self-service improves the time to get the recommendations to resolve the issue.
- Quick recommendation/remediation helps to get the infrastructure back from failure and ensures that the business runs with a less disruption.

Technet24

7-110 VMware Skyline Health Diagnostics Architecture

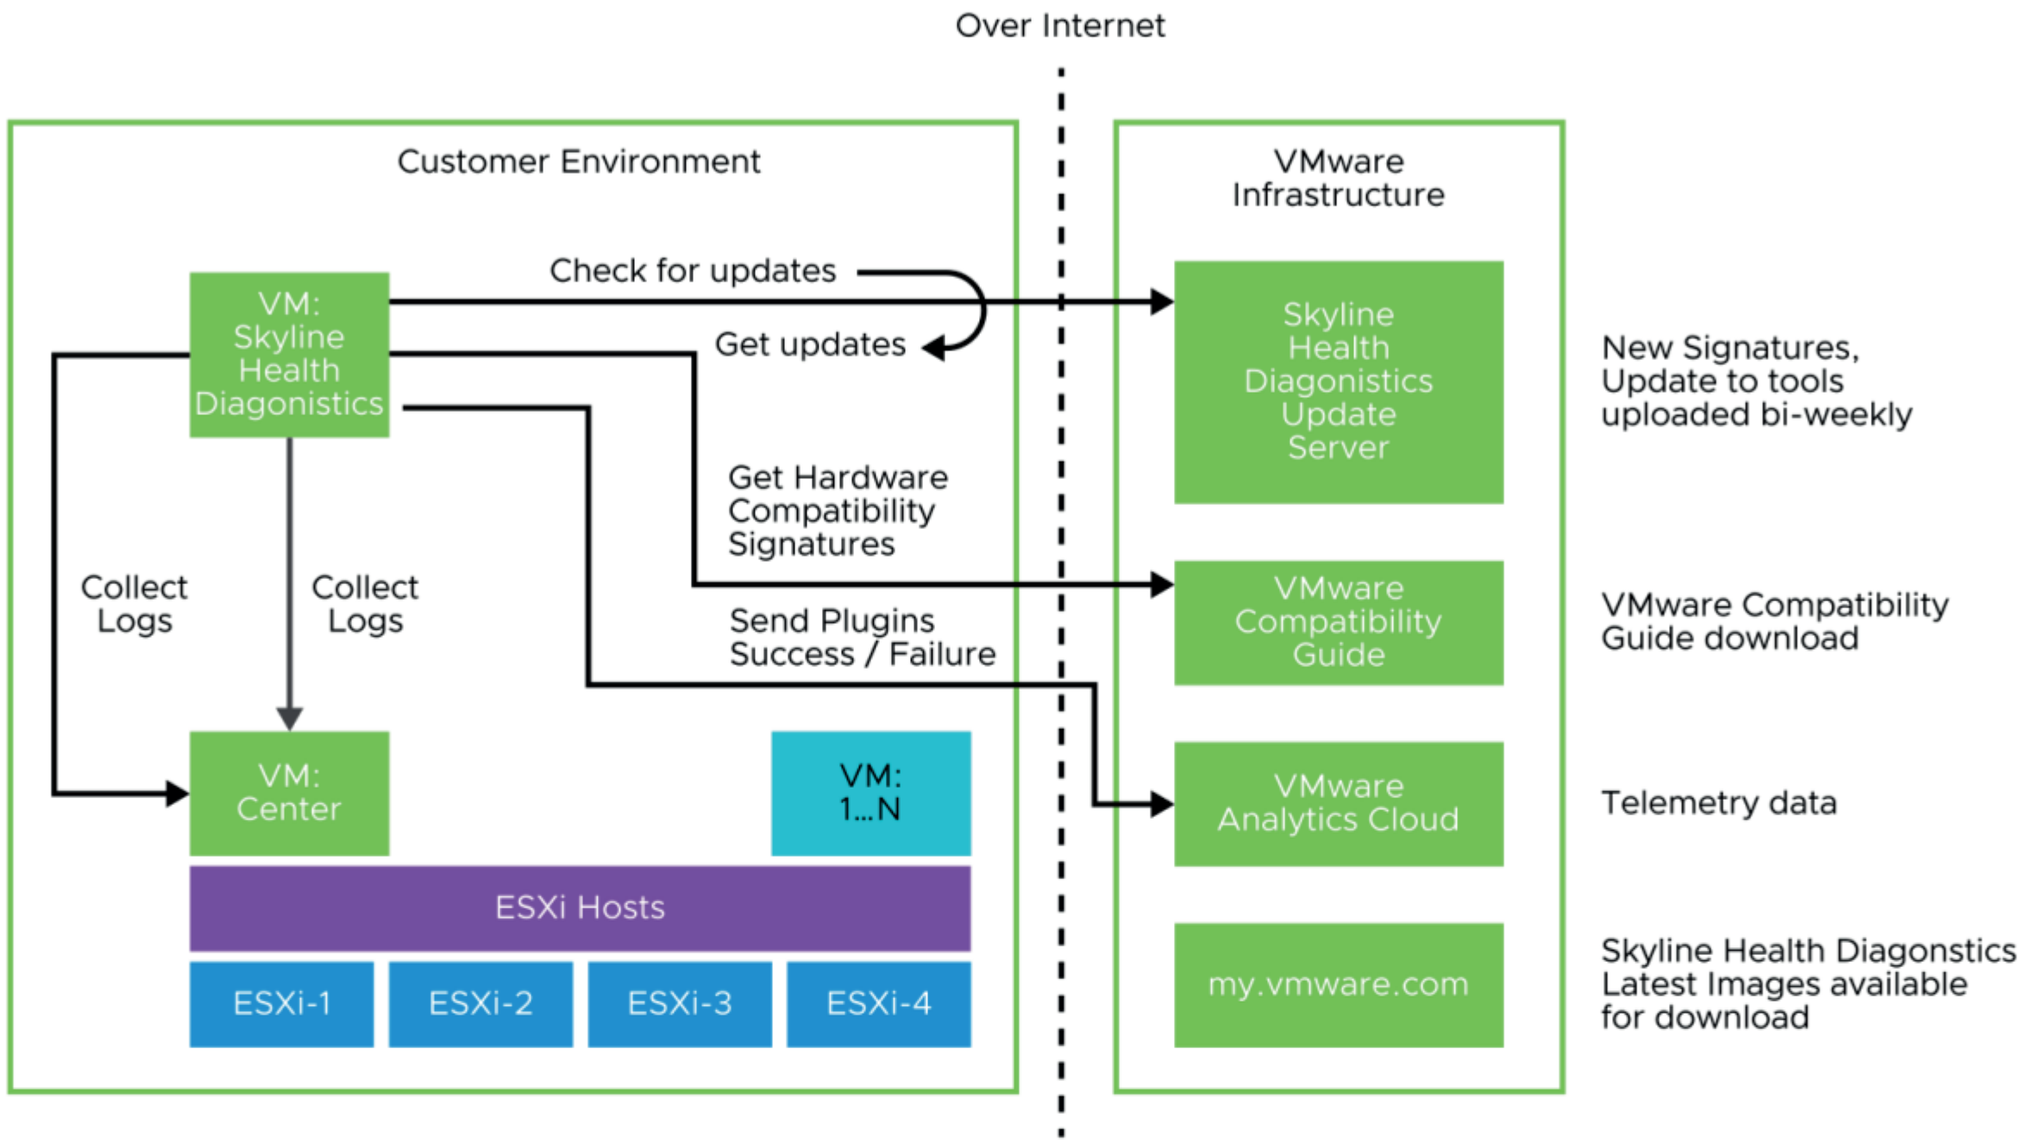

The VMware Skyline Health Diagnostics get deployed as a virtual appliance.

When you experience an issue with VMware Products, you can collect the log bundle and run the analysis using the Diagnostics Platform.

Diagnostics Platform has many signatures to detect the issue. When a signature matches the information in log files, you are notified.

Reference VMware Docs, VMware Skyline Health Diagnostics for vSphere Documentation - https://docs.vmware.com/en/VMware-Skyline-Health-Diagnostics/index.html

7-111 Review of Learner Objectives

- Describe the benefits and capabilities of VMware Skyline
- Explain how VMware Skyline works
- Identify the types of health information provided by Skyline Health
- Recognize uses for Skyline Advisor Pro

7-112 Key Points

- The key to interpreting performance data is to observe the range of data from the perspective of the guest operating system, the virtual machine, and the host.
- You use alarms to monitor the vCenter inventory objects and send notifications when selected events or conditions occur.
- VMware Skyline is a proactive support technology that provides predictive analysis and recommendations to help you avoid problems.

Questions?

Technet24

Module 8

vSphere Security and Access Control

8-2 Importance

By default, vSphere components are secured by certificates, authorization, access control, and compliance to security standards. You can further harden your vSphere environment against security threats by controlling settings for vCenter systems, ESXi hosts, VMs, and the vSphere network.

8-3 Module Lessons

1. Securing vSphere
2. Configuring ESXi Host Access and Authentication
3. vSphere with Single Sign-On and LDAP Authentication
4. vCenter Identity Federation

Technet24

8-4 Lesson 1: Securing vSphere

8-5 Learner Objectives

- Explain the importance of the vSphere Security Configuration Guide
- Identify recommendations for vCenter security
- Recognize strategies for securing the vSphere management network
- Identify recommendations for ESXi host security
- Recognize guidelines for general VM and guest operating system protection

8-6 About Securing vSphere

Securing vSphere involves aspects of security for the vCenter system, ESXi hosts, VMs, and the vSphere network. The VMware website provides several security resources.

VMware security advisories provide information about security vulnerabilities that are reported in VMware products. You can sign up to receive new and updated advisories by email.

Topic	Resource
vSphere security	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-security-guide.pdf
VMware security policy	https://www.vmware.com/security/
Corporate security response	https://www.vmware.com/support/policies/security_response.html
VMware security advisories	https://www.vmware.com/security/advisories.html

8-7 About the vSphere Security Configuration Guide

The vSphere Security Configuration Guide helps you to configure your vSphere environment according to security best practices.

Over time, vSphere has become more secure by default.

The vSphere Security Configuration Guide is the baseline for security hardening of VMware vSphere itself, and the core of VMware security best practices. It has long served as guidance for vSphere Administrators looking to protect their infrastructure.

For the latest version of the vSphere Security Configuration Guide, see VMware Security Hardening Guides at https://www.vmware.com/security/hardening-guides.html.

Perform a daily audit against the vSphere Security Configuration Guide to identify security vulnerabilities in your vSphere environment.

Technet24

8-8 Controlling Access to vCenter Systems

Strictly control access to vCenter components to increase security for the system.

vCenter access control includes the following best practices:

- Use named accounts and grant the Administrator role only to accounts that are required to have it.
- Monitor Privileges of vCenter Administrator Users
- Restrict datastore browser access to users who really need that privilege.
- Instruct users of the vSphere Client to never ignore certificate verification warnings.
- Custom roles should be configured to ensure that users only have the privileges they need.
- Set the vCenter password policy.
- Limit vCenter network connectivity by putting the vCenter system only on the management network.

The vCenter root user does not have administration rights. Instead, you assign the vCenter Administrator role to a select number of named vCenter administrator accounts.

Assign the Browse Datastore privilege only to users or groups who really need the privilege. Users with the privilege can view, upload, or download files on datastores associated with the vSphere deployment through the vSphere Client.

Verify vSphere Client certificates. Without certificate verification, the user might be subject to man-in-the-middle attacks.

By default, vCenter automatically changes the vpxuser password every 30 days. You can change this frequency in the vSphere Client.

Ensure that vSphere management traffic is on a restricted network.

For the complete list of vCenter best practices, see *vSphere Security* at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-security-guide.pdf.

8-9 vCenter Access Control

To secure access to vCenter Photon OS, apply these best practices:

- Do not allow users to log directly into the vCenter instance.
- Allow access only to users that have legitimate tasks to perform on vCenter.
- Control access to SSH and the command line.

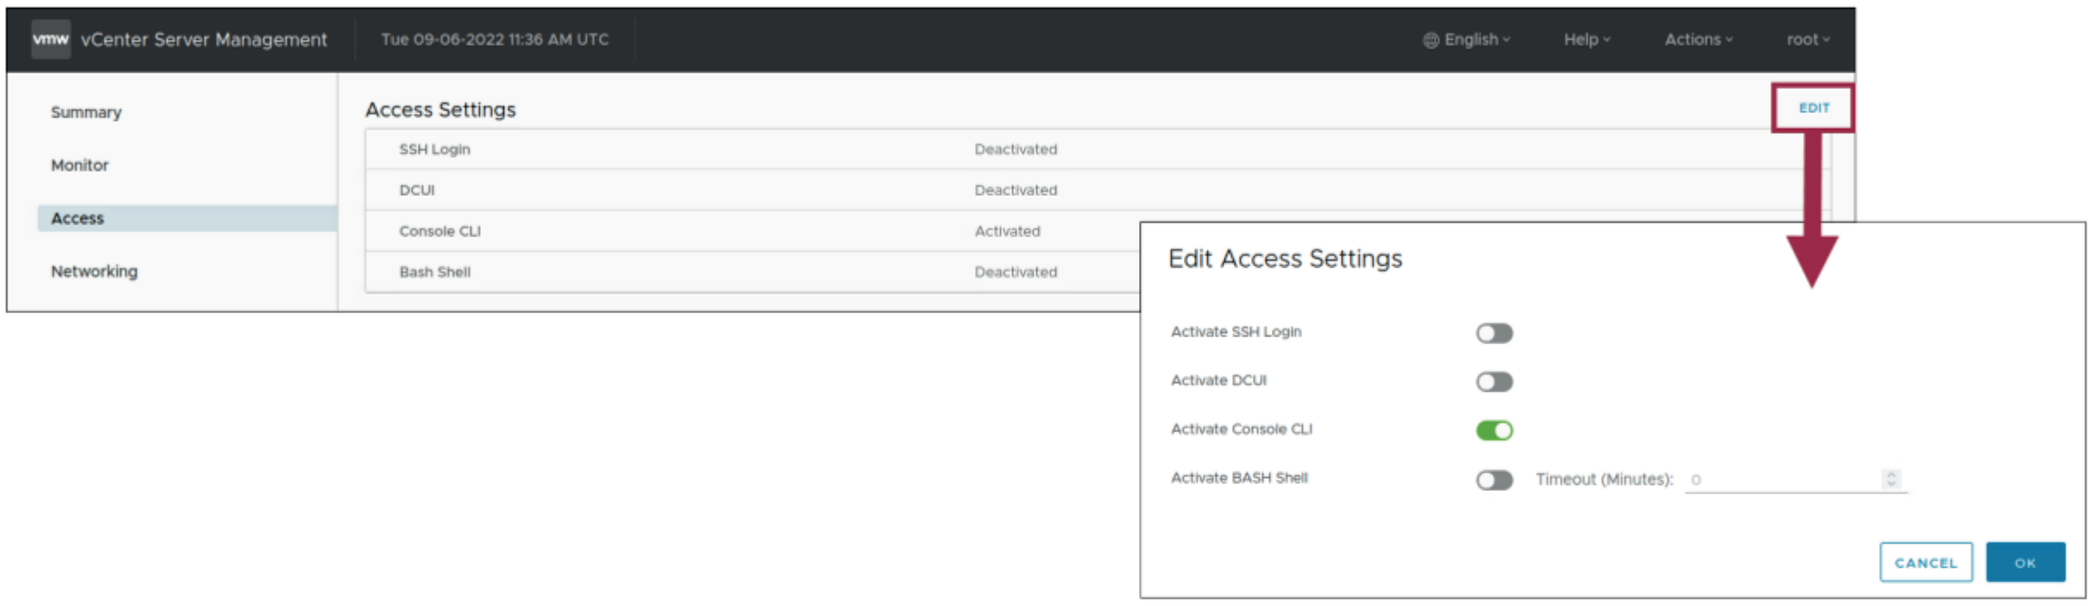

Users who can log in to the vCenter Photon OS with root privileges can cause harm by, either intentionally or unintentionally, by altering settings and modifying processes. Those users also have potential access to vCenter credentials, such as the SSL certificate.

Consider auditing login events regularly to ensure that only authorized users are accessing vCenter.

Use the VAMI to control access to SSH, the DCUI, the Appliance Shell, and the Bash shell on a vCenter Appliance instance.

Technet24

8-10 Securing vCenter

Ensure that vCenter networking is secure:

- Configure Network Time Protocol (NTP) to ensure that all systems use the same relative time source.
- Restrict access to components that are required to communicate with vCenter Appliance.

Your vCenter instance should be installed with the latest updates and security patches:

- VMware regularly releases patches on a regular basis for vCenter.
- Download patch ISO images from https://customerconnect.vmware.com/patch.

Follow security best practices for vCenter Appliance instances:

- Ensure that all systems use the same relative time source. Synchronized systems are essential for certificate validation. Network Time Protocol (NTP) also makes it easier to track an intruder in log files.
- Restrict access to components that are required to communicate with the vCenter Appliance instance. Blocking access from unnecessary systems reduces the potential for attacks on the operating system.

8-11 Securing vSphere Management Networks

Securing vSphere management networks resembles securing physical networks, but with some special characteristics:

- Securing the physical switches that the ESXi hosts are connected to prevents attackers from gaining access to the host and its VMs.
- You can secure standard switch ports and vSphere distributed switches with security policies - Reject promiscuous mode, Reject MAC address changes and Reject forged transmits.
- You can use physically isolated networks. If physical isolation is not possible, use VLANs.

Create and install the vCenter instance on a management network whenever possible. Avoid putting the vCenter system on production or storage networks, or on a network with access to the Internet.

vCenter systems require network connectivity with the following systems:

- ESXi hosts
- Other vCenter systems in the same vCenter Single Sign-On domain
- Infrastructure services such as DNS, AD, and NTP
- Other systems that run components that are essential to the functionality of the vCenter system
- The machines from which the vSphere Client is being accessed

For the best protection of your hosts, ensure that physical switch ports are configured with the Spanning Tree Protocol deactivated and that the nonnegotiate option is configured for trunk links between external physical switches and virtual switches in Virtual Switch Tagging mode.

Technet24

8-12 Securing ESXi Hosts

To secure ESXi hosts, VMware constrains vSphere parameters and settings.

Consider these guidelines for maintaining ESXi host security:

- Manage access to the ESXi host:
  - Avoid using the root user account.
  - Use AD to manage users instead of maintaining local user accounts.
  - Always use the vSphere Client to administer ESXi hosts managed by vCenter.
  - Only use VMware Host Client to administer ESXi hosts if vCenter is unavailable.
- Manage the use of security services and functions:
  - Keep the ESXi host firewall enabled.
  - Leave the SSH and vSphere ESXi Shell services deactivated.
  - Enable lockdown mode to further increase the security of ESXi hosts.

Create at least one named user account and assign it full administrative privileges. This user account can be either a local account or a domain account that belongs to the ESX Admins domain group. Use this account instead of the root account. Set a complex password for the root account and limit its use.

You can configure ESXi hosts to use AD to manage users. Using AD alleviates maintaining local user accounts on each ESXi host.

ESXi includes a firewall that is enabled by default. You can manage firewall settings by using the Security Profiles UI in the vSphere Client. Many essential security settings for ESXi hosts can be customized using security profiles. Using security profiles is especially useful for single host management.

You can use SSH to restrict, control, and secure access to an ESXi host. SSH access and vSphere ESXi Shell access are deactivated on ESXi hosts by default.

Lockdown mode is deactivated by default. When you use lockdown mode, management of ESXi hosts can only be performed through vCenter.

You manage certificates from the vSphere Client. ESXi hosts can use VMware CA, a custom certificate authority, or thumbprint mode.

For the complete list of ESXi hypervisor and ESXi host security best practices, see *vSphere Security* at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-security-guide.pdf.

8-13 VM Protection Overview

Employ the same security measures for a VM as you do for an equivalent physical server:

- Keep all security measures current and patched.
- Install antivirus and antimalware software.
- Remove unnecessary virtual hardware in the VM and deactivate unnecessary functions in the Guest OS.
- Ensure that the copy and paste operations between guest operating systems and the remote console are deactivated.
- Restrict users from running commands on a VM.
- Use templates to deploy VMs.
- Keep the operating system patches up to date.

Using templates to deploy VMs removes the risk of misconfiguration. All VMs are created with a known baseline level of security.

8-14 Review of Learner Objectives

- Explain the importance of the vSphere Security Configuration Guide
- Identify recommendations for vCenter security
- Recognize strategies for securing the vSphere management network
- Identify recommendations for ESXi host security
- Recognize guidelines for general VM and guest operating system protection

Technet24

8-15 Lesson 2: Configuring ESXi Host Access and Authentication

8-16 Learner Objectives

- Configure the ESXi firewall
- Activate and deactivate lockdown mode on an ESXi host
- Configure user logins to authenticate with directory services

8-17 About ESXi Host Security Properties

The ESXi hypervisor architecture has many built-in security features, such as CPU isolation, memory isolation, firewall, and device isolation.

You can open and close ports on the firewall, activate and deactivate services such as SSH and configure lockdown mode.

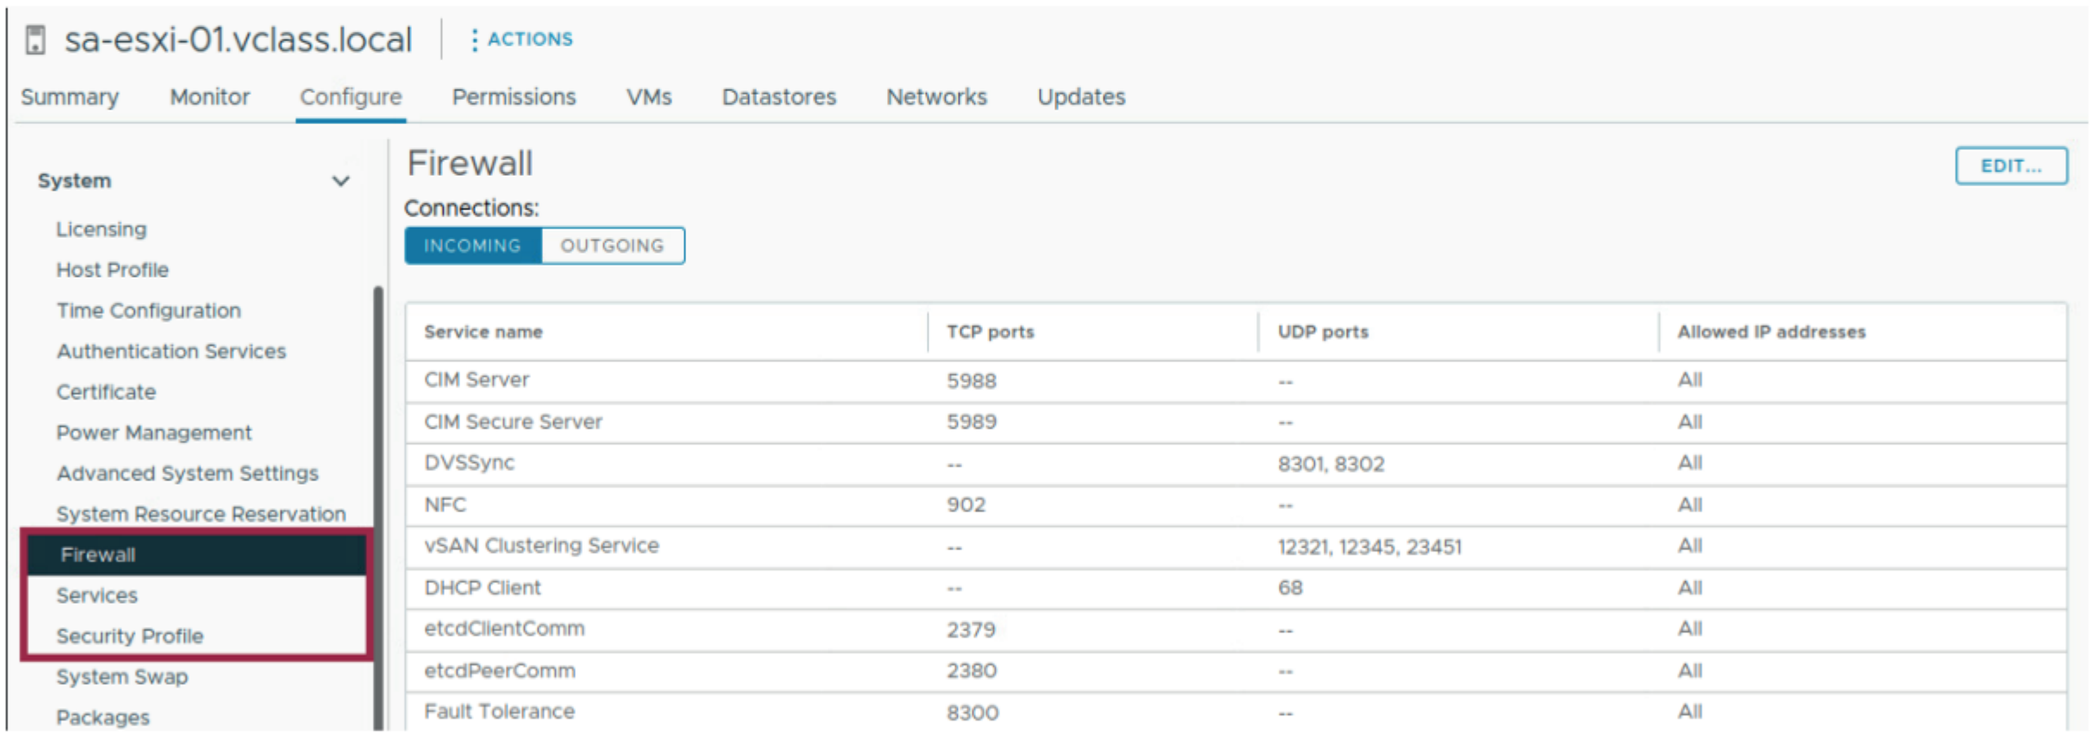

An ESXi host is protected with a firewall. You can open ports for incoming and outgoing traffic, as needed, but should restrict access to services and ports. Using the ESXi lockdown mode and limiting access to vSphere ESXi Shell can further contribute to a more secure environment.

8-18 Configuring the ESXi Firewall

To minimize the risk of an attack through the management interface, ESXi includes a service-oriented, stateless firewall. You can define the IP addresses or range of IP addresses.

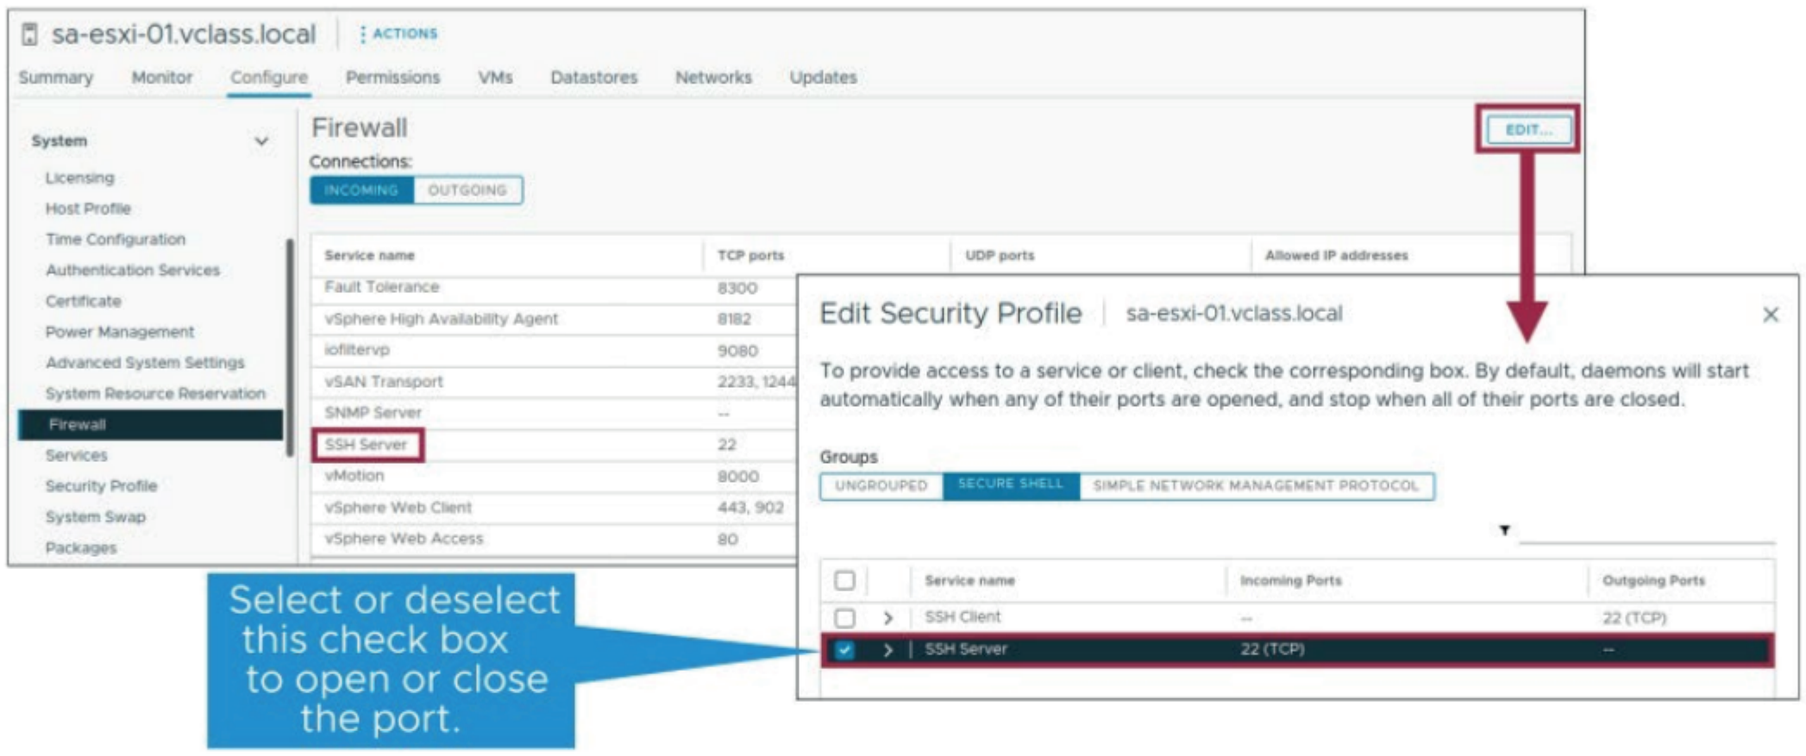

Firewalls control access to devices in their perimeter by closing all communication pathways, except for the pathways that the administrator explicitly or implicitly designates as authorized. The pathways, or ports, that administrators open in the firewall allow traffic between devices on different sides of the firewall.

With the ESXi firewall engine, rule sets define port rules for each service. For remote hosts, defining the IP addresses or range of IP addresses allows access to each service.

You can configure the ESXi firewall by using the vSphere Client or vSphere CLI. You configure firewall management at the command line by using the **`esxcli network firewall`** namespace.

Technet24

8-19 Configuring ESXi Services

You configure services to start, stop, or restart on an ESXi host. You also control whether the services start and stop with the host or must be started manually.

In this example, vSphere ESXi Shell and SSH services are stopped. These services are stopped by default.

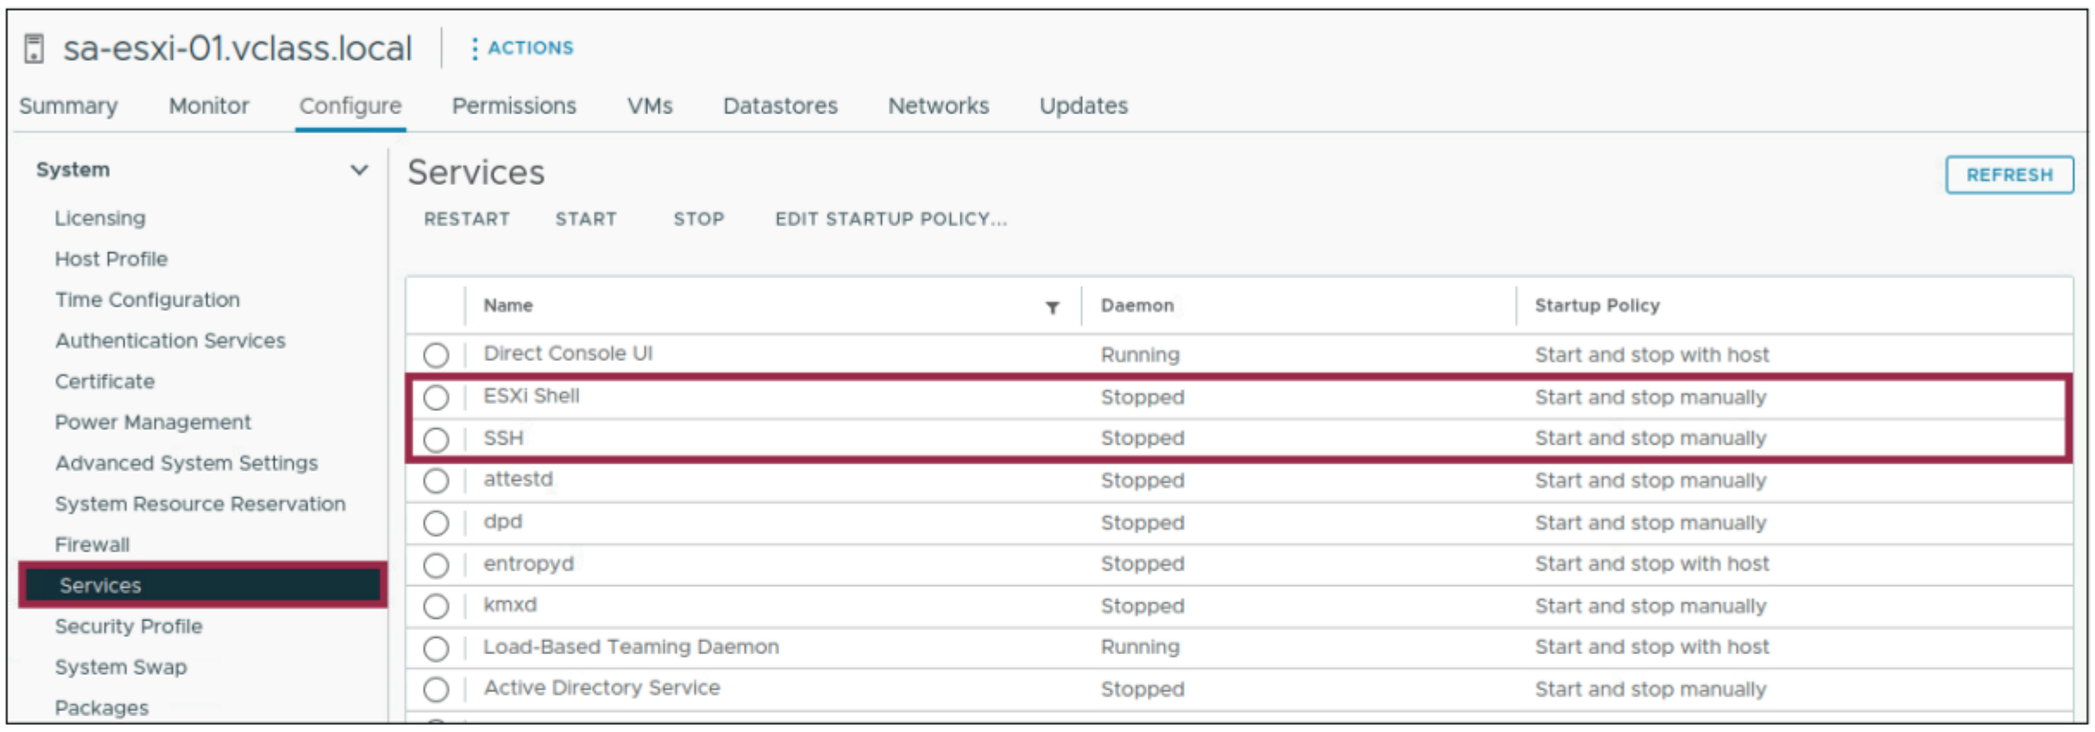

8-20 About Lockdown Mode (1)

To increase the security of your ESXi hosts, you can put your hosts in lockdown mode.

In lockdown mode, direct access to the ESXi host is either restricted or denied.

The available modes are:

- Normal
- Strict

8-21 About Lockdown Mode (2)

By default, lockdown mode is deactivated.

You can enable lockdown mode via the vSphere Client, either when the ESXi host is added to the inventory or later by editing the host's Security Profile.

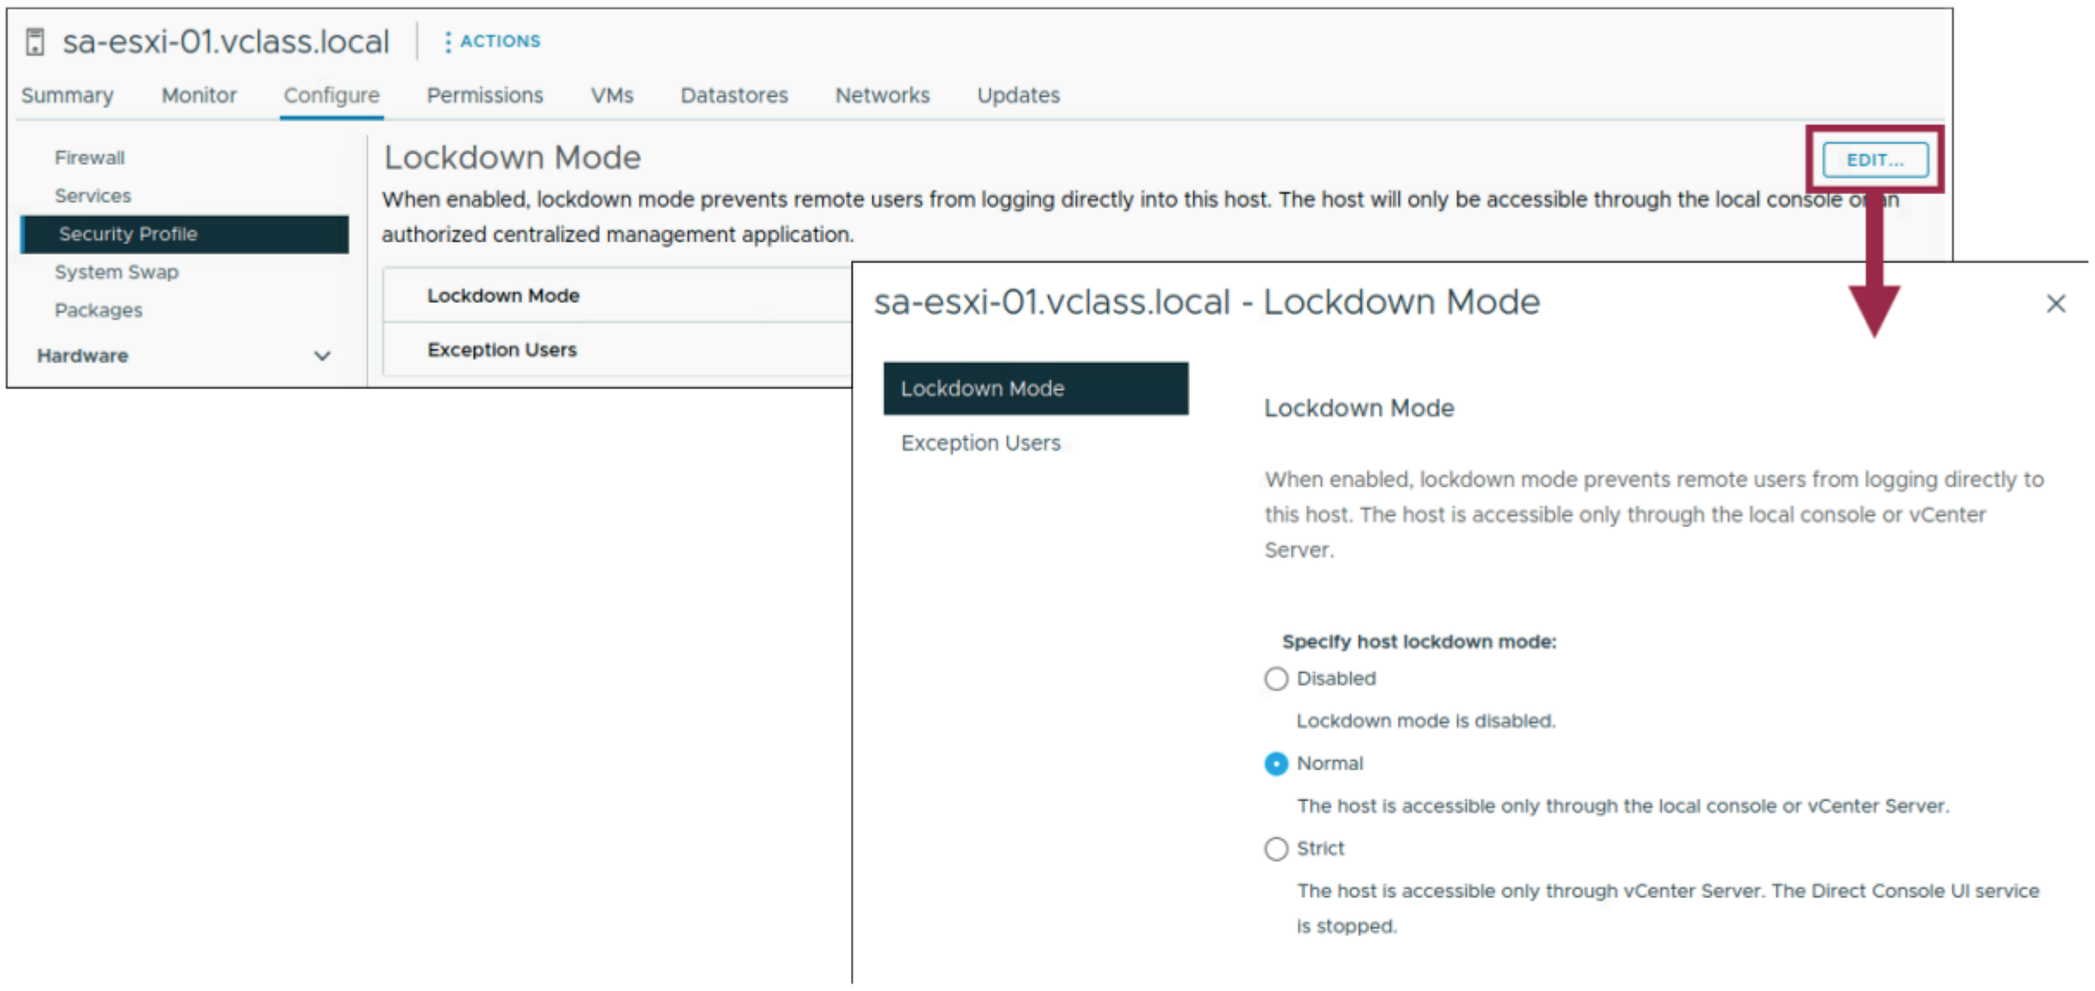

Technet24

8-22 About Lockdown Mode (3)

When you enable lock mode (normal or strict), users cannot perform operations directly against the host. All management must be performed through vCenter via the vSphere Client.

You can enable lockdown mode:

- When adding the host to the vCenter
- By using the Direct Console User Interface (DCUI) - Disable or Enable Normal mode only
- By using a host's security profile in the vSphere Client

In lockdown mode, all host operations must be performed through vCenter. The host cannot be accessed directly through the VMware Host Client or the ESXi Shell (local or remote).

8-23 Normal Lockdown Mode Behavior

When normal lockdown mode is enabled on an ESXi host:

- The host can be accessed only by vCenter, for example, by using the vSphere Client.
- Access to the DCUI is blocked for all users except:
  - Users on the Exception Users list, but they must have administrator privileges
  - Users defined in the DCUI.Access advanced system setting:
    - By default, the root user is defined in this setting.
    - Non-administrative users can also be defined in this setting.
    - Non-administrative users can only activate or deactivate lockdown mode.
- If vSphere ESXi Shell and SSH services are enabled, users with administrator privileges who are on the Exception Users list can log in to the host directly using SSH or the local vSphere ESXi Shell.

In normal lockdown mode, users listed in the DCUI.Access advanced system setting can log in to the Direct Console User Interface (DCUI). By being on this list, users have emergency access to the DCUI if the connection to vCenter is lost. These users do not require administrative privileges on the host. Nonadministrative DCUI.Access users can only deactivate or activate lockdown mode.

8-24 Strict Lockdown Mode Behavior

When strict lockdown mode is enabled on an ESXi host:

- The host can be accessed only by vCenter, for example, by using the vSphere Client.
- The DCUI service is deactivated and cannot be started. Therefore, no users can log in to the DCUI.
- If vSphere ESXi Shell and SSH services are enabled, users with administrator privileges who are on the Exception Users list can log in to the host directly using SSH or the local vSphere ESXi Shell.

For strict or normal lockdown mode to be an effective security measure, ensure that vSphere ESXi Shell and SSH services are deactivated.

If vCenter is unavailable, the vSphere Client is also unavailable. Therefore, hosts that are in strict lockdown mode are not manageable.

Technet24

8-25 Integrating ESXi with Active Directory

You can configure an ESXi host to join an Active Directory (AD) domain so that you can manage users and groups from a single user interface.

When the ESXi host is added to AD, the 'ESX Admins' domain group is assigned full administrative access to the host, if this group exists.

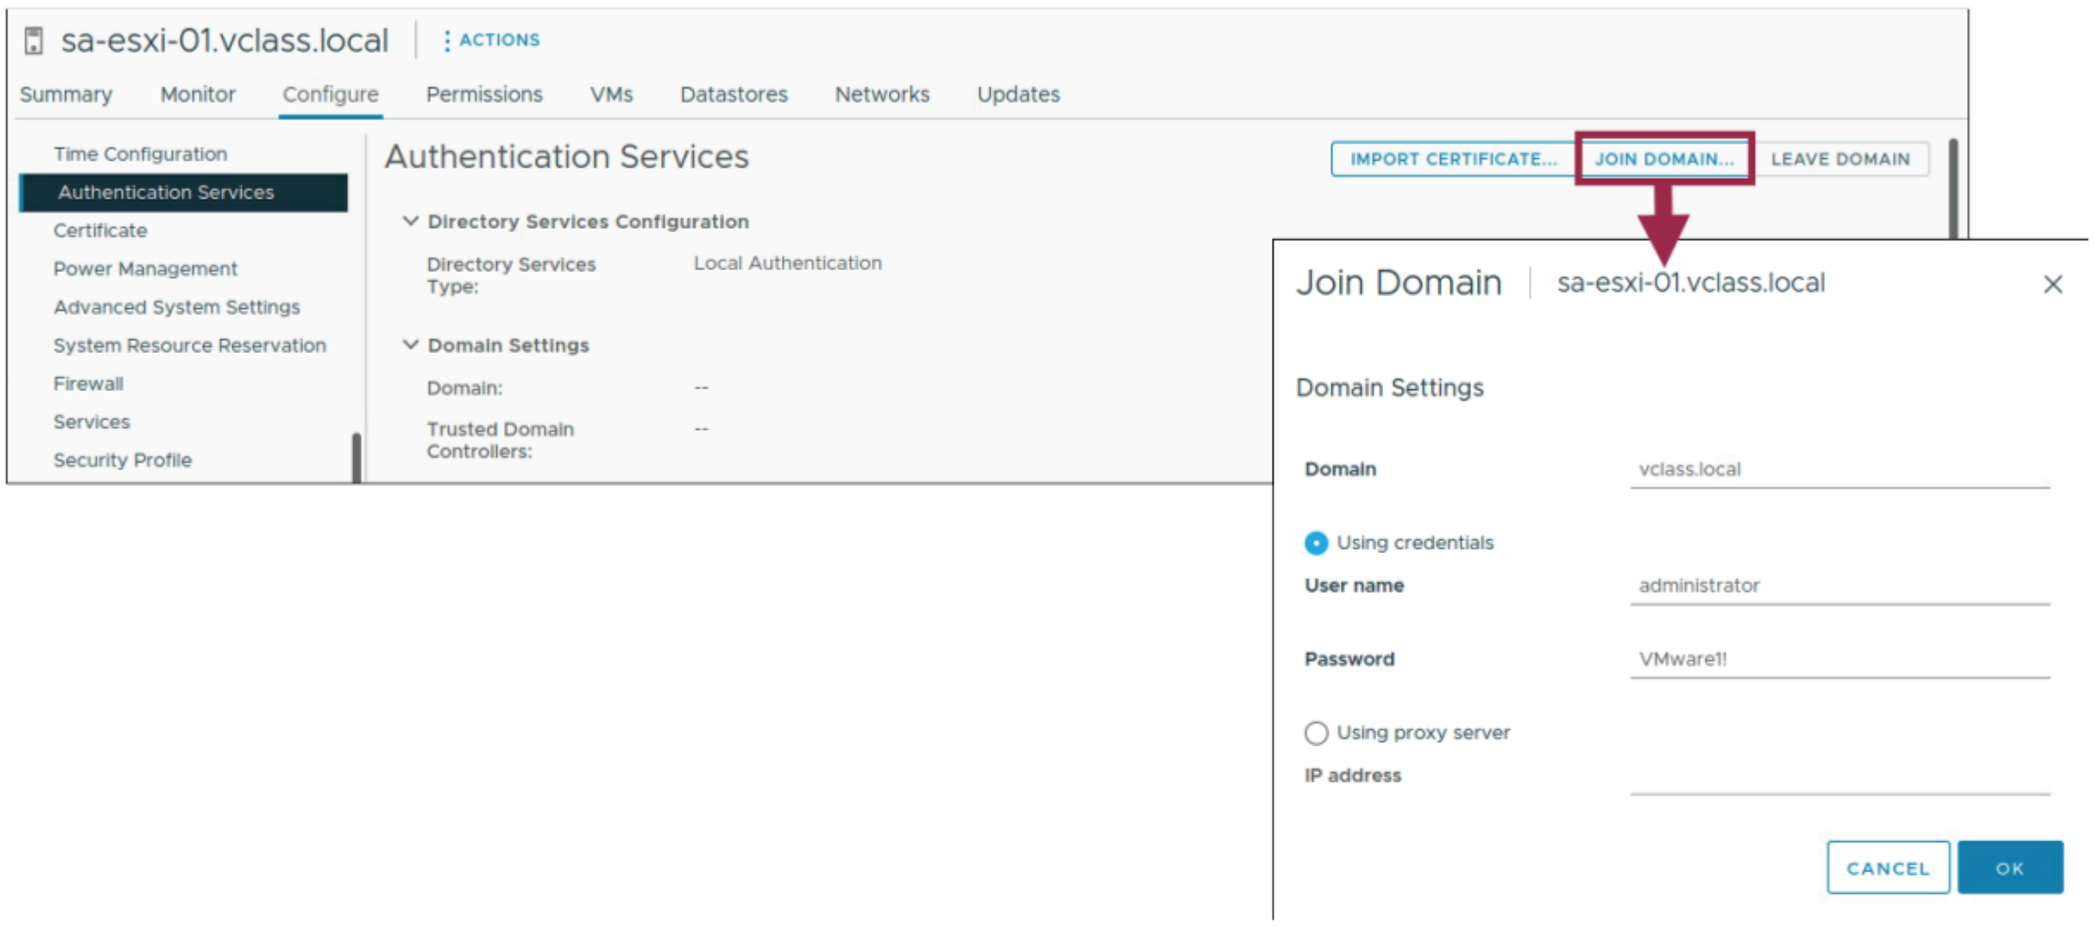

Although day-to-day vSphere management operations are done while logged in to vCenter using the vSphere Client, the user sometimes requires direct access to the ESXi host, for example, when accessing local log files.

Whenever asked to provide credentials (for example, when using VMware Host Client to log in directly to the ESXi host), you can enter the user name and password of a user in the domain to which the host is joined. The advantage of this user account management model is that you can use AD to manage your ESXi host user accounts. This model is easier and more secure than trying to manage accounts independently per host.

The only user that is defined on the ESXi host is root. The root password is not mapped to an AD account. The initial root password is typically set during ESXi installation, but it can be changed later through VMware Host Client, vSphere ESXi Shell, or the DCUI.

You can create other domain groups if you only want to grant administrative rights to a subset of hosts. You give that domain group administrator rights on the specific hosts and add the domain users to the group.

8-26 VMware vSphere Authentication Proxy Service

You can automate the process of adding ESXi hosts to an AD domain by using the vSphere Authentication Proxy service. This service provides support for joining auto-deployed ESXi hosts to an AD domain by using an account with delegated privileges.

If you are provisioning hosts with Auto Deploy, you can set up a reference host that points to Authentication Proxy.

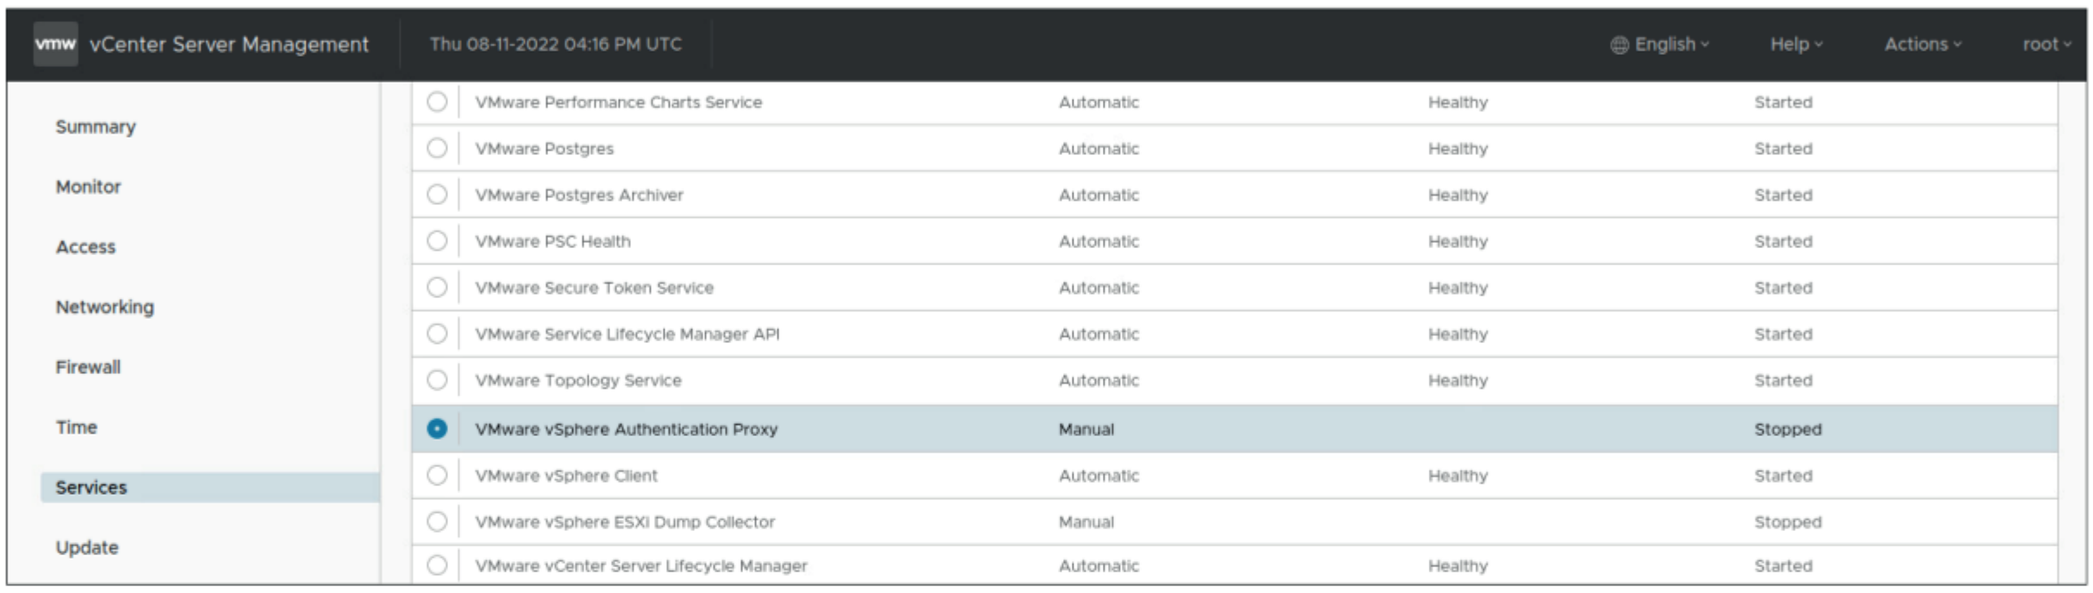

You can use vSphere Authentication Proxy to join the host to an AD domain, instead of adding the host directly to the AD domain. A chain of trust exists between vSphere Authentication Proxy and the host, which allows vSphere Authentication Proxy to join the host to the AD domain.

vSphere Authentication Proxy is especially useful when used with vSphere Auto Deploy. You can set up a reference host that points to vSphere Authentication Proxy and set up a rule that applies the reference host profile to any ESXi host provisioned with vSphere Auto Deploy. Even if you use vSphere Authentication Proxy in an environment that uses certificates, provisioned by VMware CA or third-party certificates, the process works seamlessly if you follow the instructions for using custom certificates with vSphere Auto Deploy.The vSphere Authentication Proxy service is available on each vCenter system. By default, the service is not running. If you want to use vSphere Authentication Proxy in your environment, you can start the service from the virtual appliance management interface (VAMI) or from the command line.

Technet24

8-27 Lab 18: Configuring Lockdown Mode

Configure and test lockdown mode:

1. Start the SSH Service
2. Enable and Test Lockdown Mode
3. Disable Lockdown Mode
4. Knowledge Check

8-28 Review of Learner Objectives

- Configure the ESXi firewall
- Activate and deactivate lockdown mode on an ESXi host
- Configure user logins to authenticate with directory services

8-29 Lesson 3: vSphere with Single Sign-On and LDAP Authentication

8-30 Learner Objectives

- Identify the services used for vCenter Single Sign-On
- Recognize LDAP server identity source settings
- Recognize methods for two-factor authentication on vCenter

8-31 vCenter Single Sign-On

vCenter Single Sign-On is an authentication broker and security token exchange infrastructure.

vCenter SSO issues a token when a user authenticates.

The user can use the token to authenticate to vCenter Server services.

8-32 vCenter Single Sign-On Services

To effectively manage vCenter Single Sign-On, you need to understand the underlying services and how they work.

- VMware Directory Service (VMDir)
- VMware Certificate Authority (VMCA)
- VMware Authentication Framework Daemon (VMAFD)
- Lookup Service (lookupsvc)
- Identity Management Service (vIDM)

```
root@sb-vcsa-01 [ ~ ]# service-control --status --all
Running:
 applmgmt lookupsvc lwsmd observability observability-vapi pschealth vlcm vmafdd vmcad vmdird vmonapi vmware-analytics vmware-
certificateauthority vmware-certificatemanagement vmware-cis-license vmware-content-library vmware-eam vmware-envoy vmware-hvc
 vmware-infraprofile vmware-perfcharts vmware-pod vmware-postgres-archiver vmware-rhttpproxy vmware-sca vmware-sps vmware-stat
smonitor vmware-stsd vmware-topologysvc vmware-trustmanagement vmware-updatemgr vmware-vapi-endpoint vmware-vdtc vmware-vmon v
mware-vpostgres vmware-vpxd vmware-vpxd-svcs vmware-vsan-health vmware-vsm vsphere-ui vstats vtsdb wcp
Stopped:
 vmcam vmware-imagebuilder vmware-netdumper vmware-rbd-watchdog vmware-vcha
```


Technet24

8-33 VMware Directory Service (VMDir)

VMware Directory Service (VMDir) is an internal (local) LDAP repository in vCenter that contains user identities, groups, and configuration data.

The VMware Directory Service is also used for certificate management.

The log files for VMDir (`vmdird`) are located on vCenter at `/var/log/vmware/vmdir` and `/var/log/vmware/vmdird`

8-34 VMware Certificate Authority (VMCA)

VMCA includes the VMware Endpoint Certificate Store (VECS) and several other authentication services.

The VMCA also provisions vCenter components and ESXi hosts with certificates that use VMCA as the root certificate authority.

VMCA issues certificates only to clients that can authenticate to vCenter Single Sign-On in the same domain.

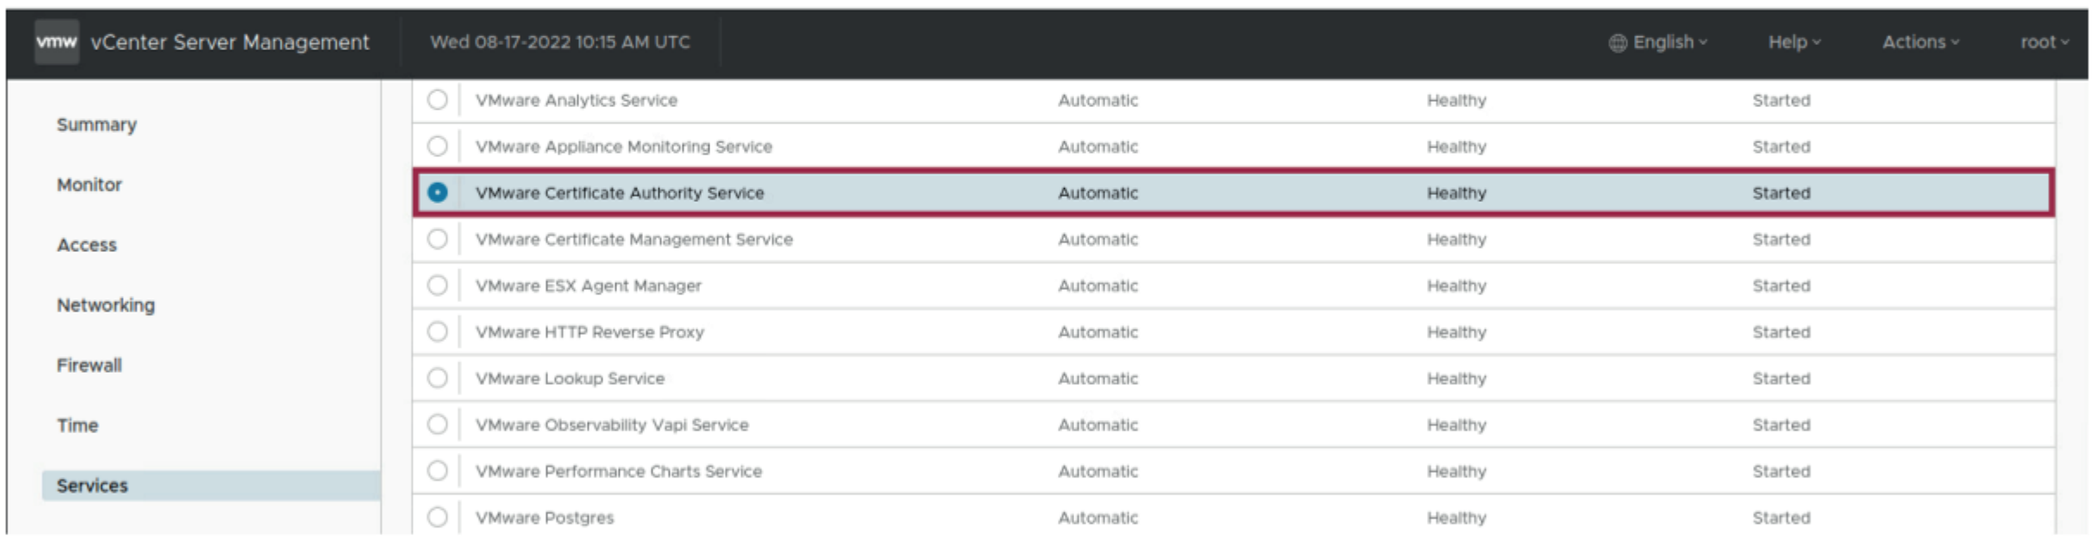

8-35 Lookup Service (lookupsvc)

The Lookup Service provides a central location for services to publish their functionalities as service endpoints.

vCenter Services and external solutions can query the Lookup Service to list the services and endpoints for specific functionality.

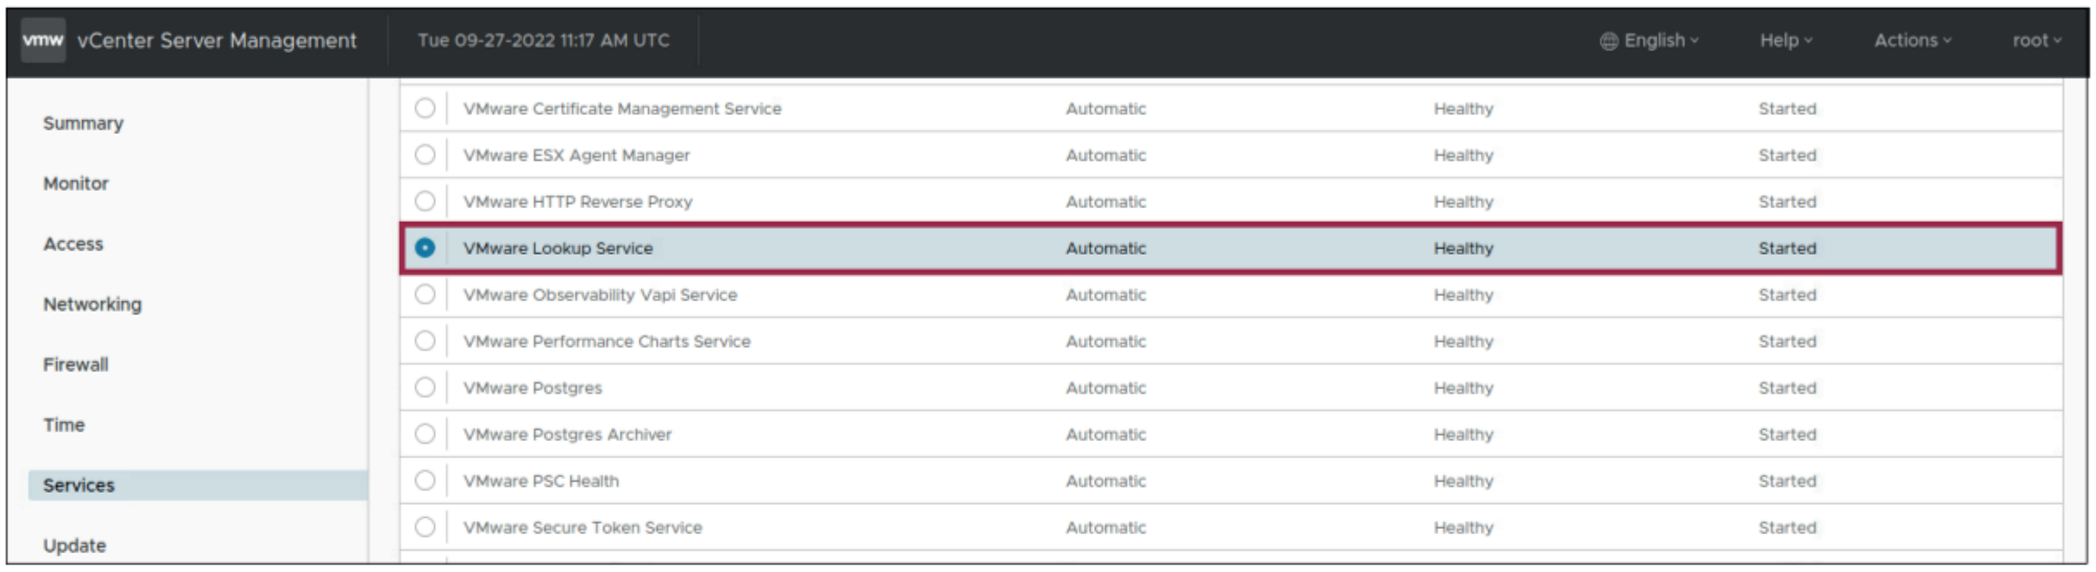

8-36 Identity Management Service (vIDM)

The Identity Management Service handles identity sources such as AD and OpenLDAP-based directory services and STS authentication requests.

The vCenter Server Security Token Service (STS) is a Web service that issues, validates, and renews security tokens.

The log files for the Identity Management Service (`vmware-identity-sts`) are located on the vCenter at `/var/log/vmware/sso`

Technet24

8-37 VMware Authentication Framework Daemon (VMAFD)

VMware Endpoint Certificate Store (VECS) runs as part of the VMware Authentication Framework Daemon.

VECS serves as a local (client-side) repository for certificates, private keys, and other certificate information that can be stored in a keystore.

You can also explicitly manage certificates and keys in VECS using vecs-cli commands.

```
root@sa-vcsa-01 [ ~ ]# /usr/lib/vmware-vmafd/bin/vecs-cli store list
MACHINE_SSL_CERT
TRUSTED_ROOTS
TRUSTED_ROOT_CRLS
machine
vsphere-webclient
vpxd
vpxd-extension
hvc
data-encipherment
SMS
wcp
APPLMGMT_PASSWORD
BACKUP_STORE
root@sa-vcsa-01 [ ~ ]#
```

For more information on vecs-cli commands, see *vecs-cli Command Reference* - https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-authentication/GUID-80C30209-028D-4F92-9551-CDB6B25A8403.html

8-38 About Identity Sources

Identity sources are used to authenticate users who are accessing the vSphere environment, including the vCenter and ESXi hosts.

Adding an identity source to vSphere allows it to search for user accounts and other objects within the directory to grant permissions and access to resources.

Identity sources provide centralized management of user accounts and authentication, improving security and reducing administrative overhead.

8-39 LDAP Server Identity Source Settings

You can choose from the following identity sources on your vCenter.

Identity Source	Description
System Domain	Contains the vCenter Single Sign-On internal users.
Local OS (Default)	All local operating system users.
Active Directory over LDAP	vCenter Single Sign-On supports multiple Active Directory over LDAP identity sources.
OpenLDAP	vCenter Single Sign-On supports multiple OpenLDAP identity sources.

Technet24

8-40 Identity Sources for vCenter with vCenter Single Sign-On

You can use identity sources to allow Single Sign-On to authenticate against external domains.

A domain is a repository for users and groups that the vCenter Single Sign-On server can use for user authentication.

In vSphere 7.0 Update 2 and later, you can enable Federal Information Processing Standards (FIPS) on vCenter. AD over LDAP and IWA are not supported when FIPS is enabled. Use external identity provider federation when in FIPS mode.

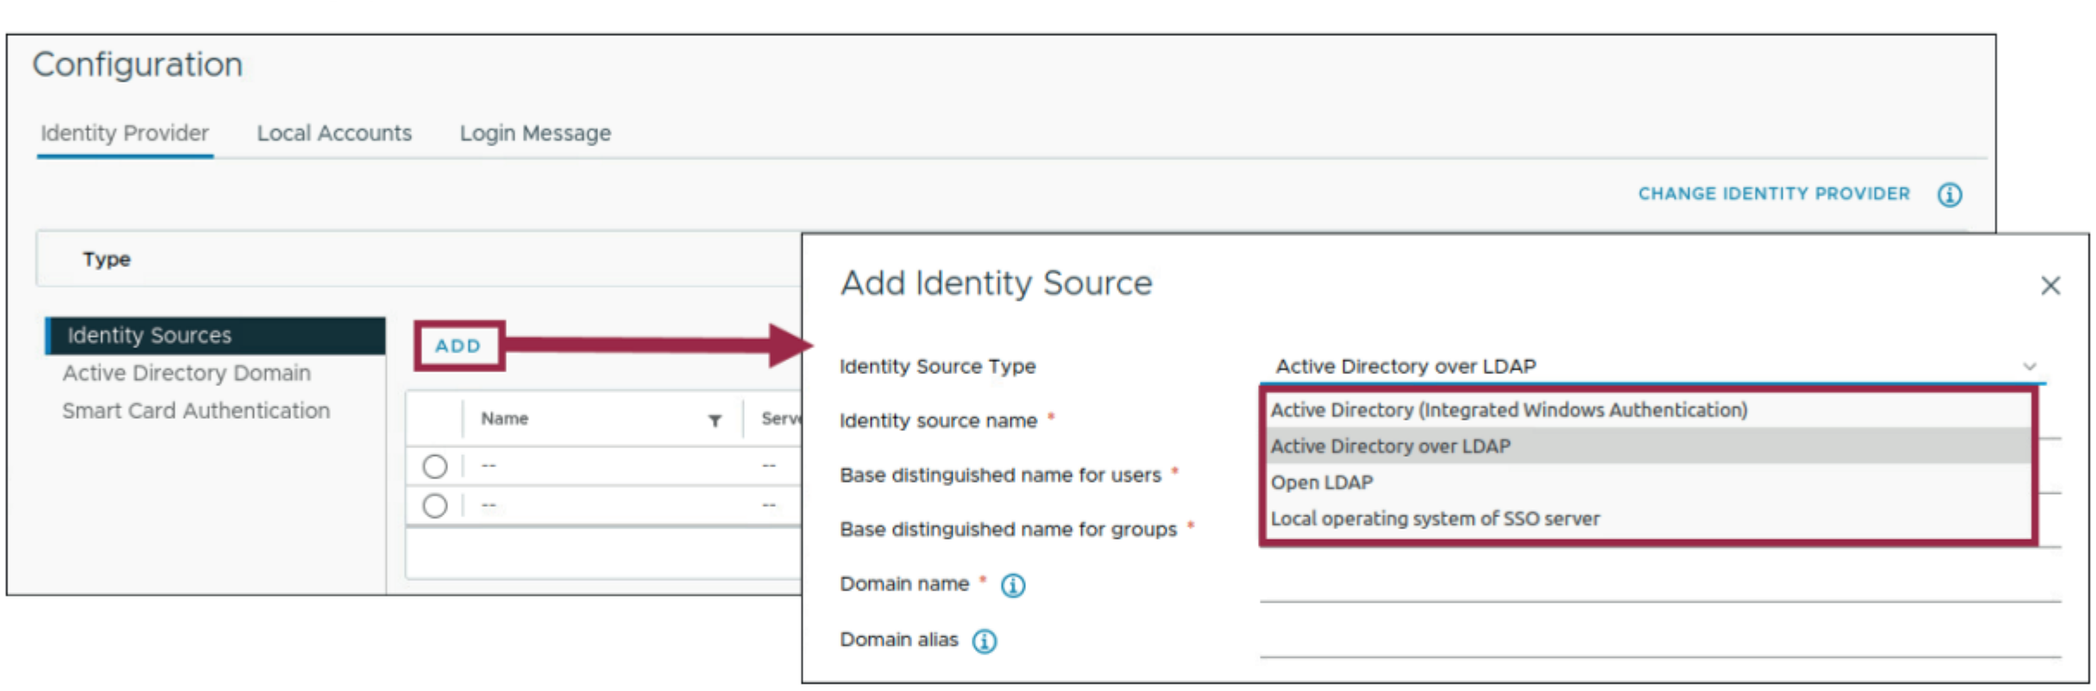

The user and group data is stored in Active Directory, OpenLDAP, or locally to the operating system of the machine where vCenter Single Sign-On is installed. Every instance of vCenter Single Sign-On has its own identity source, (default: vsphere.local). The name of this domain can be configured only during installation This identity source is internal to vCenter Single Sign-On.

See *Configuring vCenter Single Sign-On Identity Sources* at https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-authentication/GUID-75D4E587-3F9B-4B50-96DA-D6DB6D1781D7.html.

8-41 vCenter Two-Factor Authentication

vCenter allows you to authenticate as a user in an identity source that is known to vCenter Single Sign-On, or by using Windows session authentication.

You can also authenticate by using a smart card (UPN-based Common Access Card or CAC), or by using an RSA SecurID token.

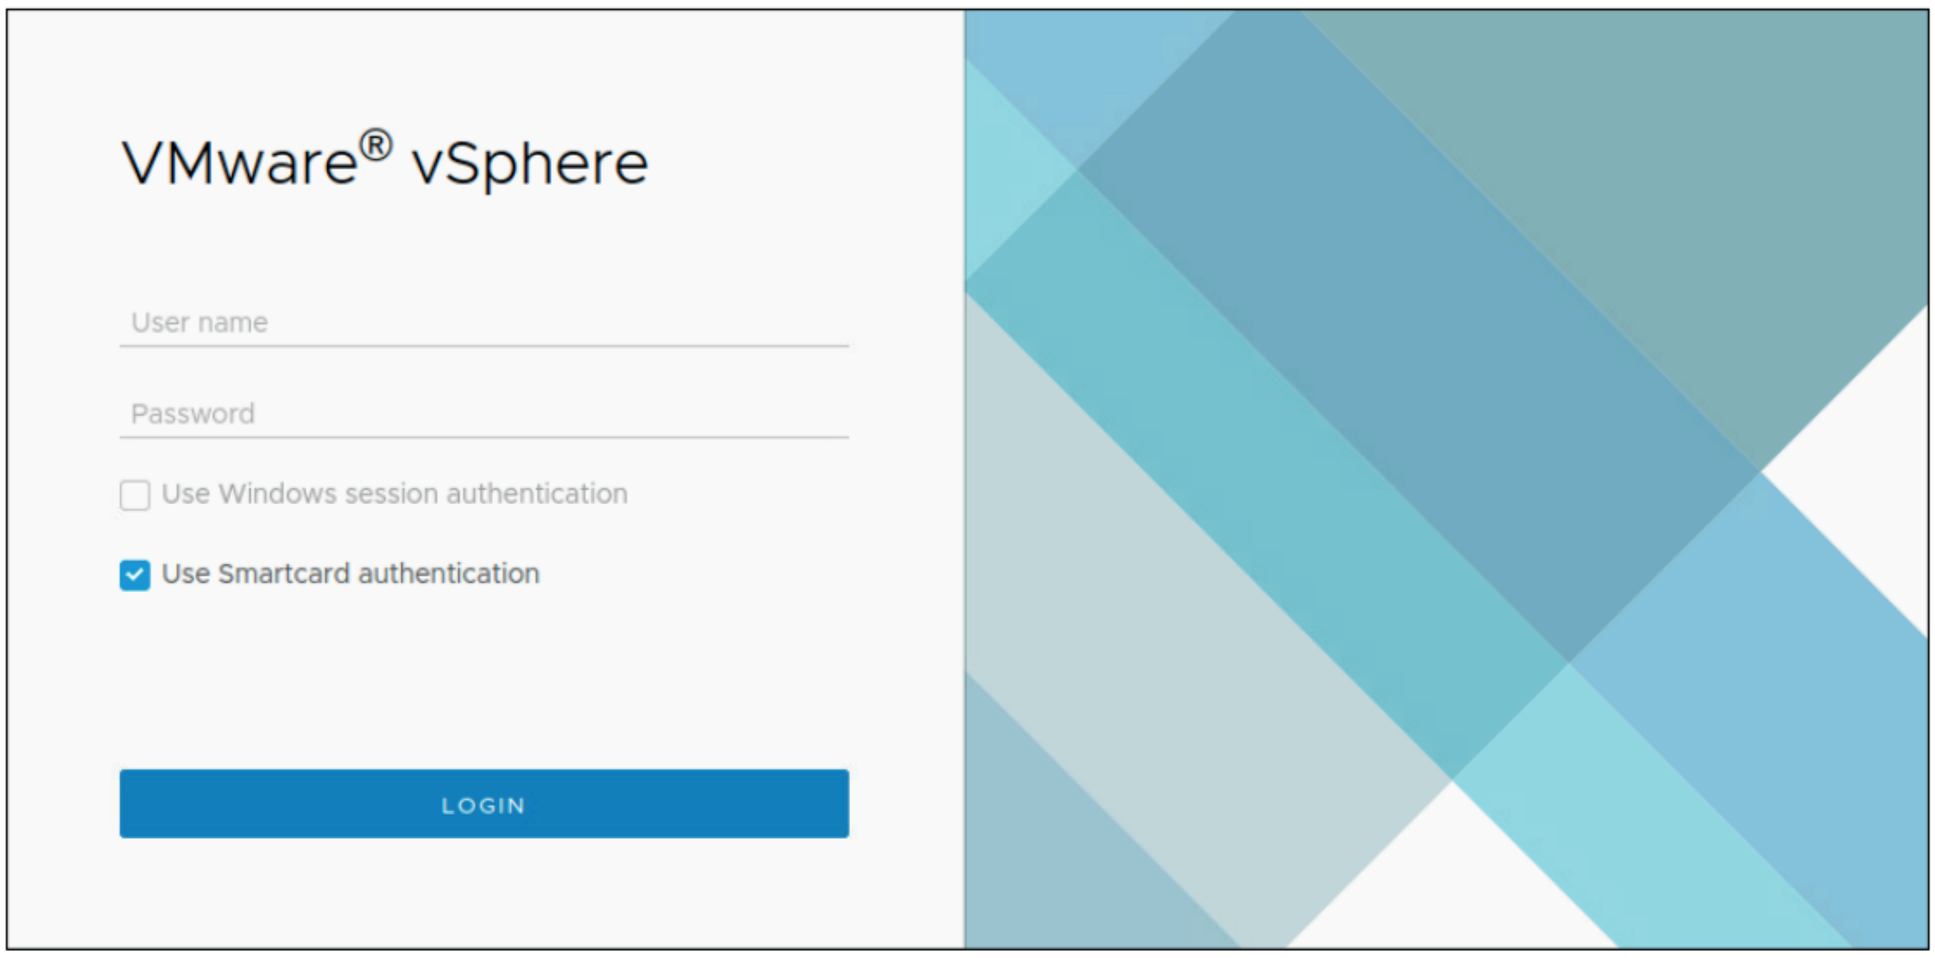

8-42 Smart Card Authentication

A smart card is a small plastic card with an embedded integrated circuit chip

Many government agencies and large enterprises use smart cards such as Common Access Card (CAC) to increase the security of their systems and to comply with security regulations.

You may use a smart card in environments where each machine includes a smart card reader. Smart card hardware drivers that manage the smart card are typically preinstalled.

To enable smart card authentication on your vCenter you must:

- Set up the reverse proxy.
- Enable and configure smart card authentication.

Technet24

8-43 Logging in with Smart Card Authentication

When you log in to your vCenter you are prompted to authenticate with a smart card and PIN combination, as follows.

- When you insert the smart card into the smart card reader, the browser reads the certificates on the card.
- The browser prompts the user to select a certificate, then prompts the user for the PIN for that certificate.
- vCenter Single Sign-On checks whether the certificate on the smart card is known.
- If the certificate is known to vCenter Single Sign-On, and is not a revoked certificate, the user is authenticated and can perform tasks for which that the user has permissions.

8-44 RSA SecurID Authentication

For RSA SecurID authentication, your environment must include a correctly configured RSA Authentication Manager.

If your vCenter is configured to point to the RSA server, and if RSA SecurID Authentication is enabled, users can log in with their username and token.

Prerequisites:

- RSA Authentication Manager version 8.0 or later is required.
- Verify that the identity source that RSA Manager uses has been added to vCenter Single Sign-On.
- Verify that the RSA Authentication Manager system can resolve the vCenter host name, and that the vCenter system can resolve the RSA Authentication Manager host name.

8-45 Review of Learner Objectives

- Identify the services used for vCenter Single Sign-On
- Recognize LDAP server identity source settings
- Recognize methods for two-factor authentication on vCenter

8-46 Lesson 4: vCenter Identity Federation

8-47 Learner Objectives

- Describe identity federation
- Recognize use cases for identity federation
- Describe the architecture of identity federation
- Configure identity federation

8-48 About vCenter Identity Provider Federation for ADFS

vCenter 7.0 introduced support for identity provider (IDP) federation for Active Directory Federation Services (ADFS).

When vCenter is configured to use an external IDP, vCenter SSO is bypassed and the user credentials are passed from the browser directly to the external IDP.

In this federation authentication setup, vCenter does not perform the following functions:

- Handling the user's credentials
- Creating the initial identity token
- Managing or updating user and group information in the identity source

8-49 vCenter Identity Provider Federation Use Cases

You can use an external IDP to replace vCenter Single Sign-On as the Identity Provider (IDP) and enable single sign-on (SSO) with federated infrastructure and applications.

You can use the vCenter IDP federation in the following scenarios:

- When using SSO with the existing ADFS-integrated infrastructure and applications
- For improving data center security because vCenter never handles the user’s credentials
- When using the authentication mechanisms, such as multi-factor authentication, that are supported by the external IDP

Technet24

8-50 Prerequisites for vCenter Identity Provider Federation

vCenter supports ADFS as an external IDP.

ADFS requirements:

- ADFS for Windows Server 2016 or greater must be deployed.
- ADFS must be connected to Active Directory.
- An application group for vCenter must be created in ADFS as part of the configuration process.
- The ADFS server certificate should be signed by a trusted certificate authority.

vCenter supports only one external IDP (one ADFS source) - the local OS and vsphere.local identity source.

You cannot use multiple external IDPs.

The vCenter IDP federation uses OpenID Connect (OIDC) for user login to vCenter.

8-51 vCenter Identity Provider Federation Components

The vCenter 7.0 IDP federation configuration includes the following vSphere and third-party components:

- vCenter
- An external IDP service that is configured on vCenter
- An ADFS server and associated Microsoft AD domain
- An ADFS application group
- AD groups and users that have permissions on vCenter

8-52 User Login with the vCenter Built-In Identity Provider

vCenter, vCenter Single Sign-On, and AD interact during the login workflow

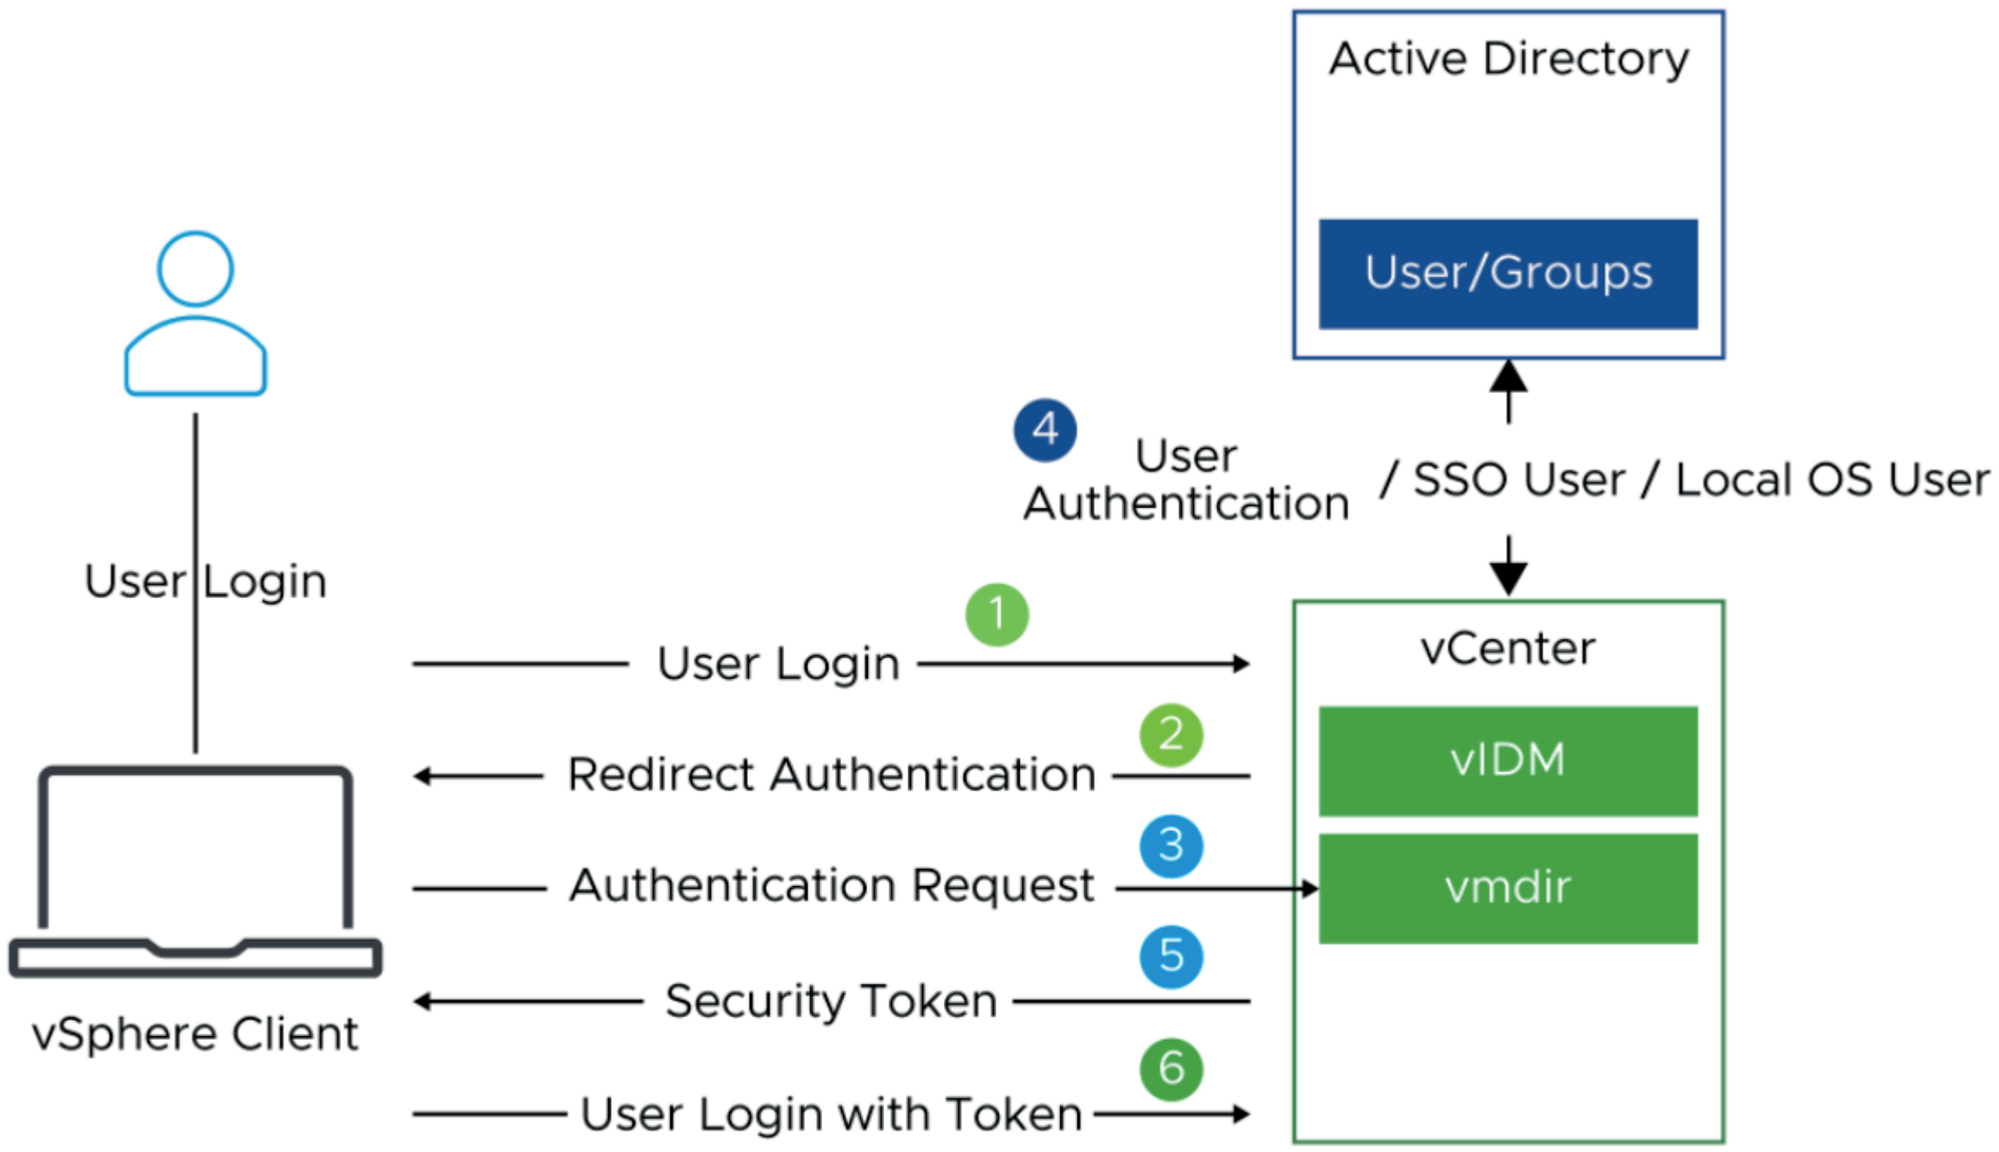

Before vCenter IDP Federation for ADFS, the interaction between vCenter, vCenter Single Sign-On, and AD during the login workflow is as follows:

1. The user logs in to vCenter on the usual vCenter landing page.
2. vCenter redirects the authentication request to the vCenter Single Sign-On service
3. If needed, the vCenter Single Sign-On service prompts the user to log in with Active Directory credentials
4. The vCenter Single Sign-On service authenticates the user with Active Directory
5. The vCenter Single Sign-On service signs all tokens with a signing certificate, and stores the token signing certificate on disk. The certificate for the service itself is also stored on disk.
6. vCenter uses the token to log in the user.

Technet24

8-53 vSphere Identity Sources with Multifactor Authentication (MFA)

Identity sources with MFA provide an additional layer of security for authenticating users accessing the vSphere environment.

To enable MFA with vSphere identity sources, administrators can use third-party solutions that integrate with vSphere, such as:

- RSA SecurID
- Microsoft ADFS
- Okta

These solutions typically require additional configuration to integrate with vSphere, but once configured, they can provide an additional layer of security for the entire vSphere environment.

8-54 Okta Federated Identity Management for vCenter

vSphere 8 Update 1 expands vCenter support for 3rd party identity providers to Okta, in addition to Active Directory Federation Services (ADFS).

- Admins using Okta can log in once for both vCenter and NSX Manager.
- This capability improves both efficiency and security for the environment.

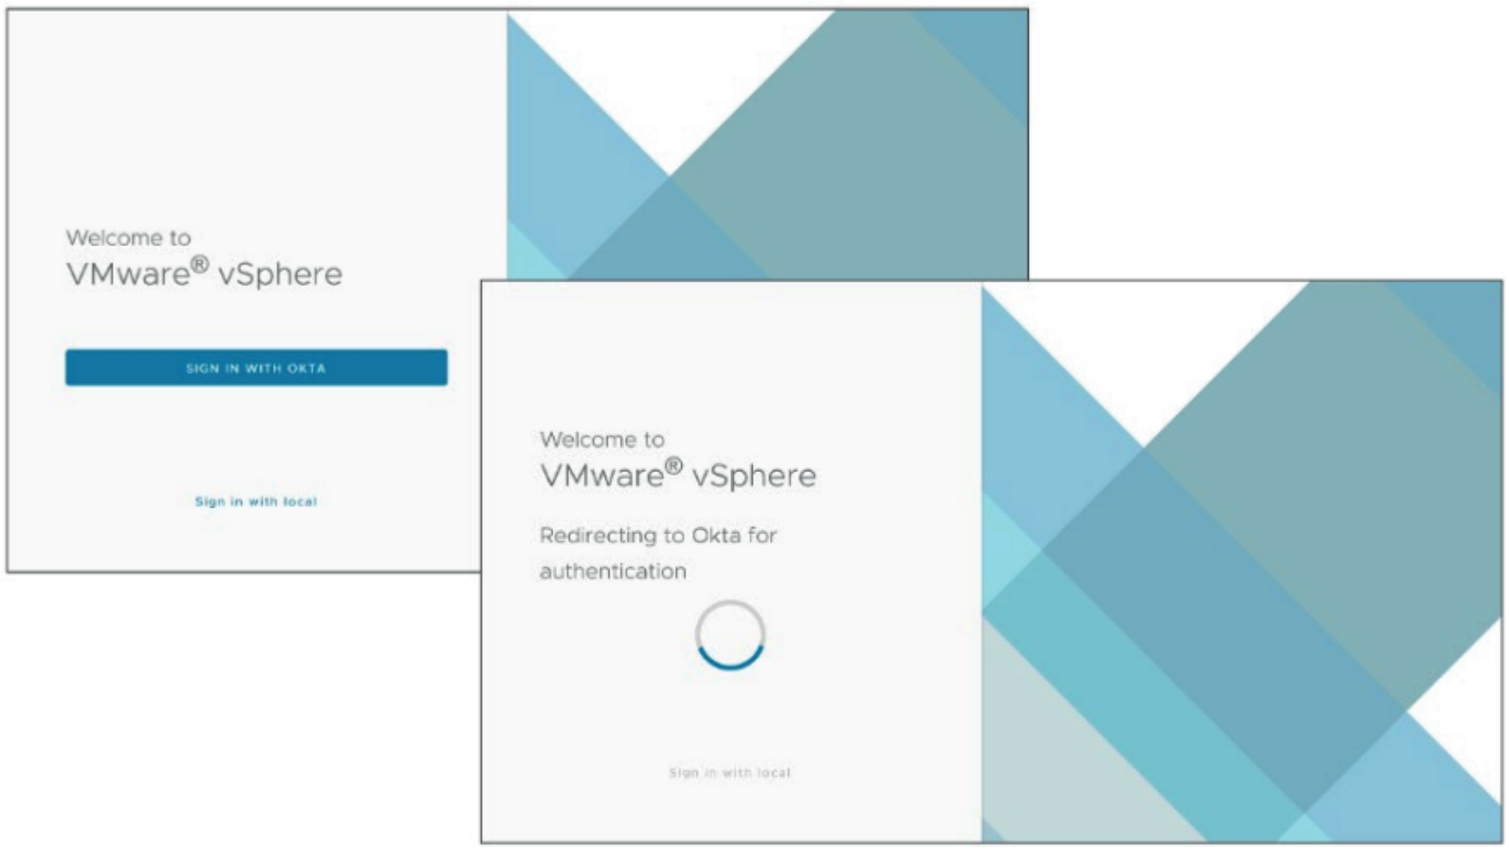

Technet24

8-55 Okta Prerequisites

Okta requirements:

- Have a dedicated domain space, for example, https://your-company.okta.com
- You must create the following Okta applications to perform OIDC logins and manage user and group permissions:
  - An Okta native application with OpenID Connect as the sign-on method
  - A System for Cross-domain Identity Management (SCIM) 2.0 application with an OAuth 2.0 Bearer Token

Okta connectivity requirements:

- vCenter must be able to connect to the Okta discovery endpoint, and the authorization, token, JWKS, and any other endpoints advertised in the discovery endpoint metadata
- Okta must also be able to connect with vCenter to send user and group data for the SCIM provisioning

Reference VMware Doc, Configure vCenter Server Identity Provider Federation for Okta - https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-authentication/GUID-88933505-9299-49FB-9C30-56E43683099B.html

8-56 vCenter Prerequisites

vCenter Server requirements:

- vSphere 8.0 Update 1
- On the vCenter where you want to create the Okta identity source, verify that the VMware Identity Services are activated

vSphere privileges requirements:

- You must have the privilege to create, update, or delete a vCenter Identity Provider that is required for federated authentication

8-57 Configure vCenter Identity Provider Federation for Okta (1)

In Okta, you create an OpenID Connect application in Okta and assign groups and users to the OpenID Connect application.

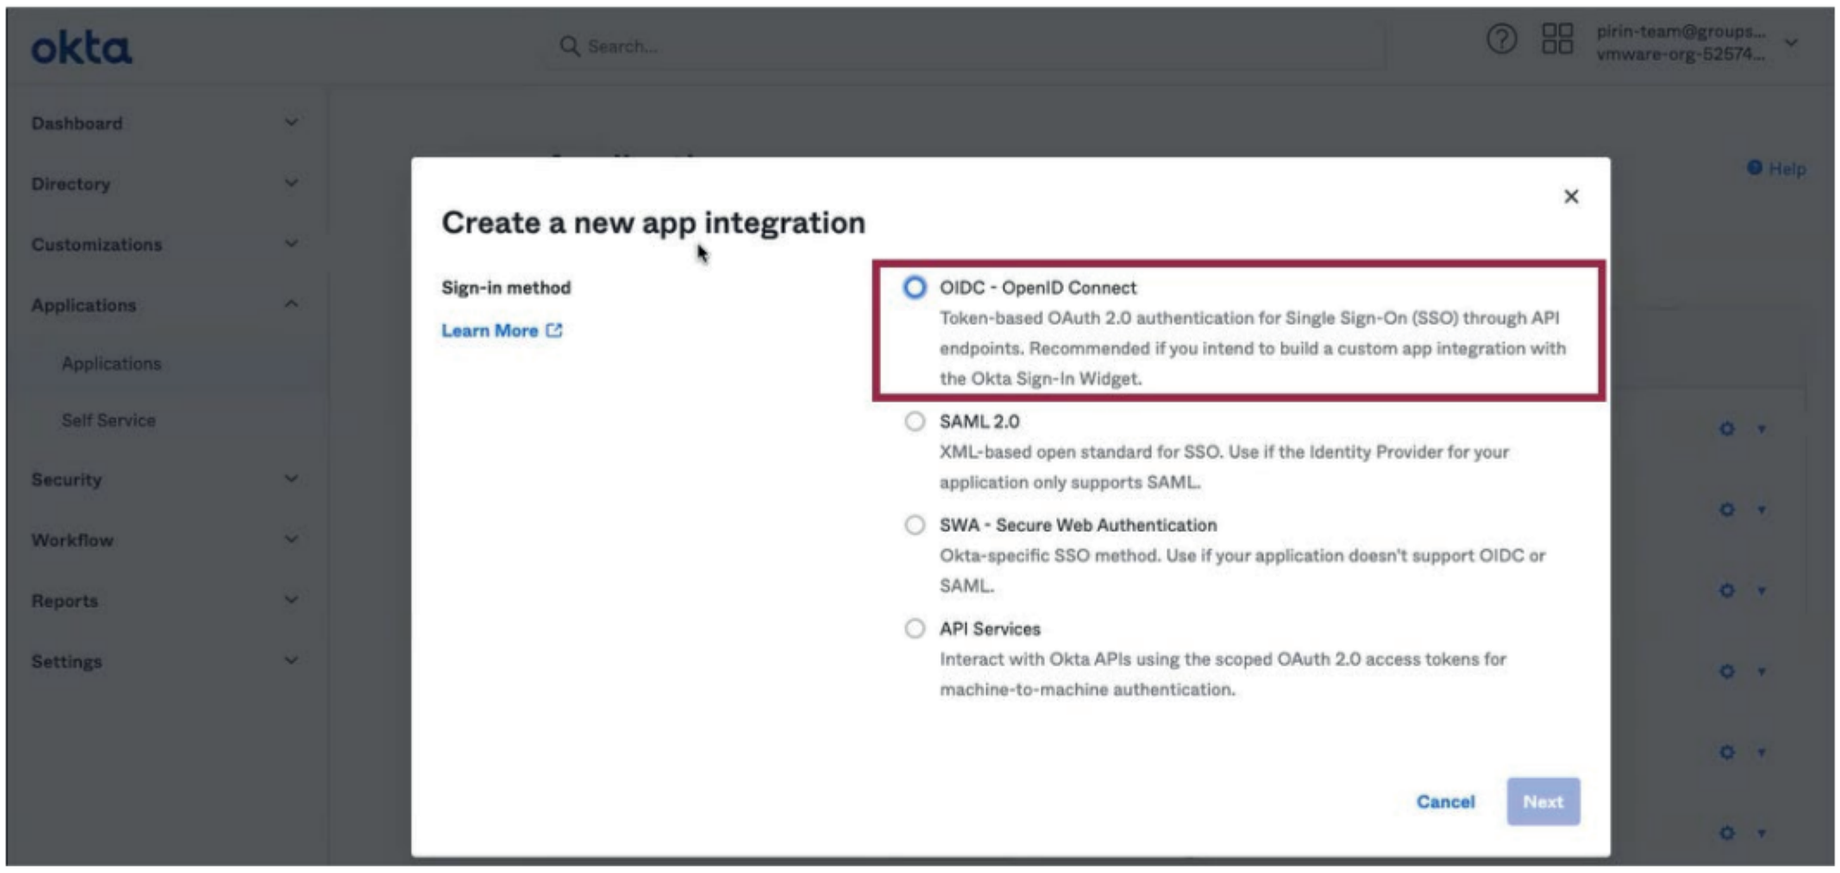

8-58 Configure vCenter Identity Provider Federation for Okta (2)

In the vSphere Client, you create and configure the identity provider.

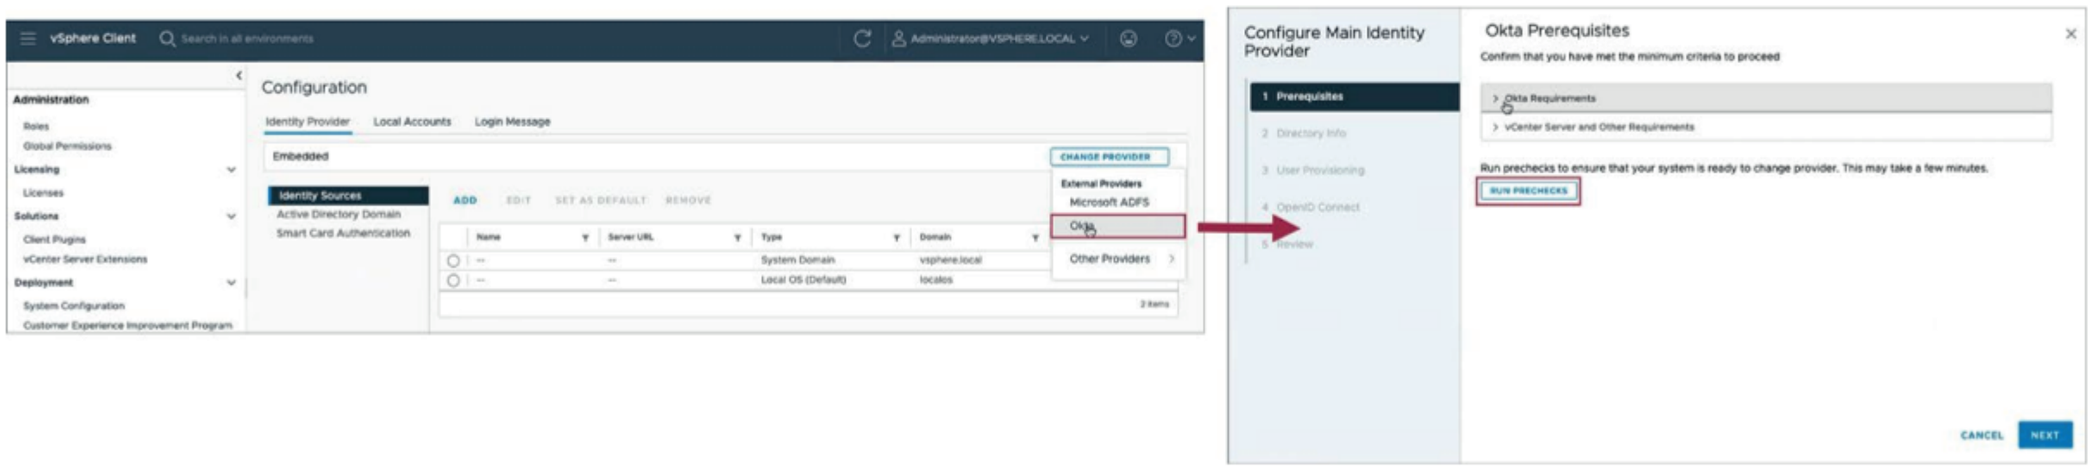

Technet24

8-59 Configure vCenter Identity Provider Federation for Okta (3)

After Okta identity provider configuration on vCenter, you now update the Okta OpenID Connect application with the Redirect URI that you copy from the **Okta Identity Provider** Configuration page.

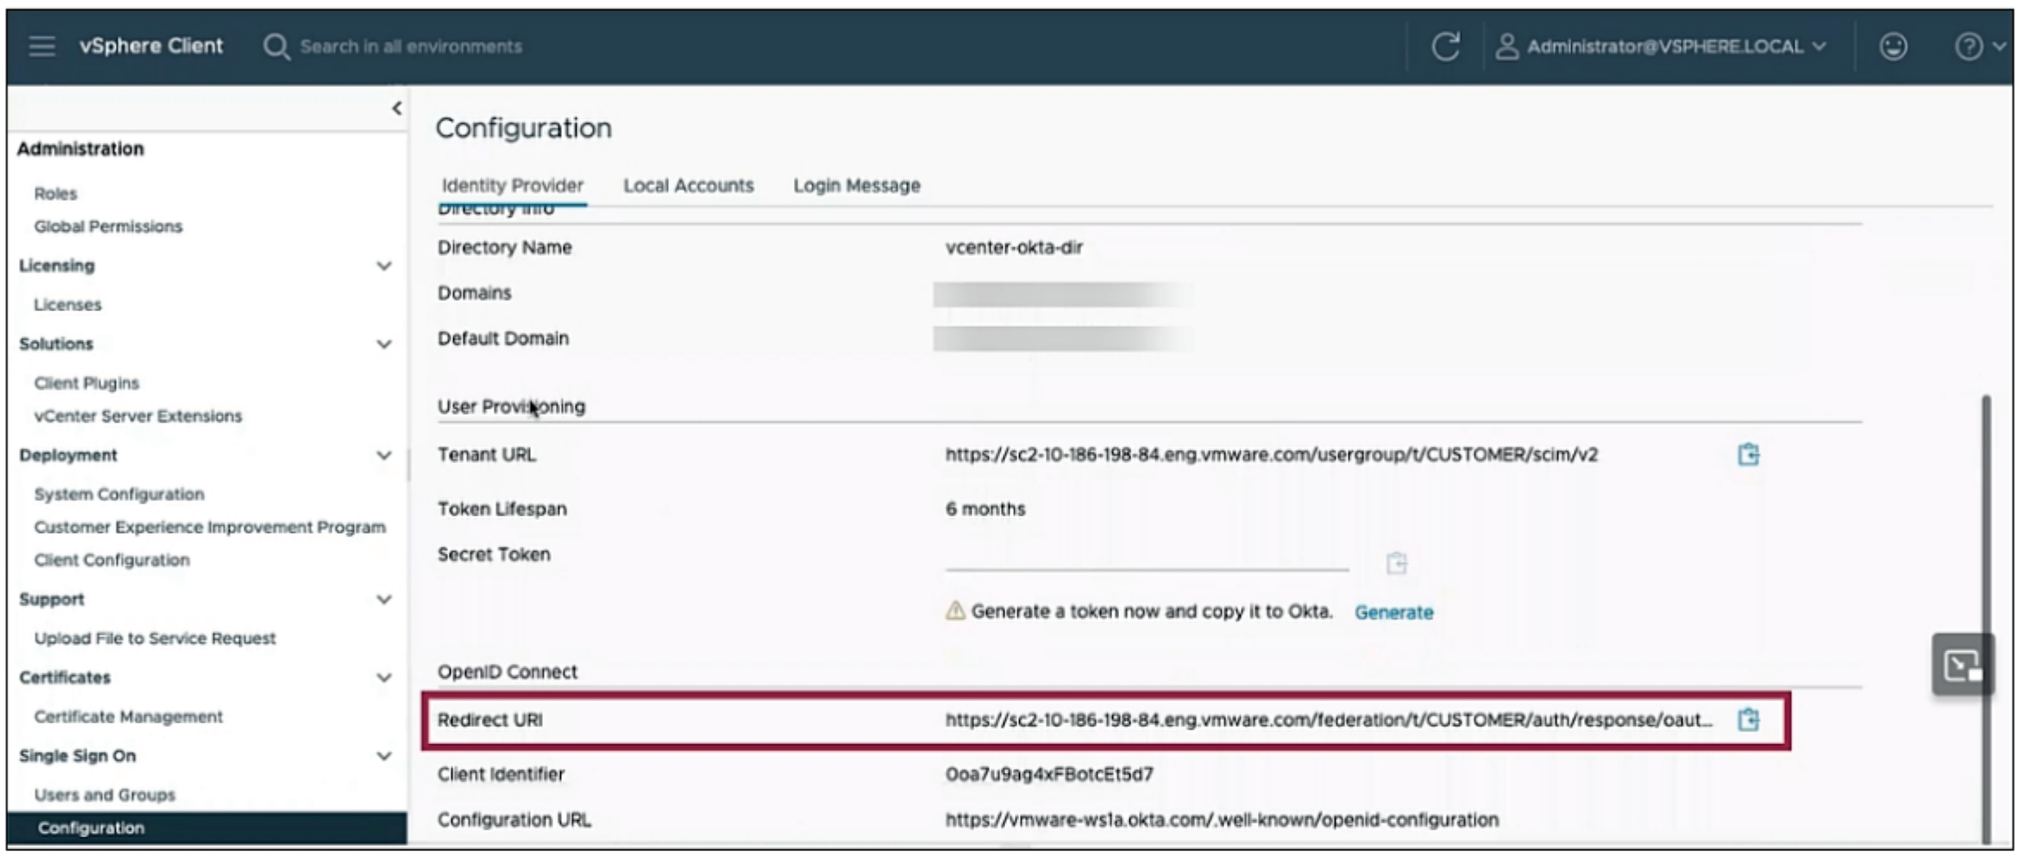

8-60 Configure vCenter Identity Provider Federation for Okta (4)

In Okta, from the Okta Application catalog, you create the SCIM 2.0 Application and Push Users and Groups to vCenter

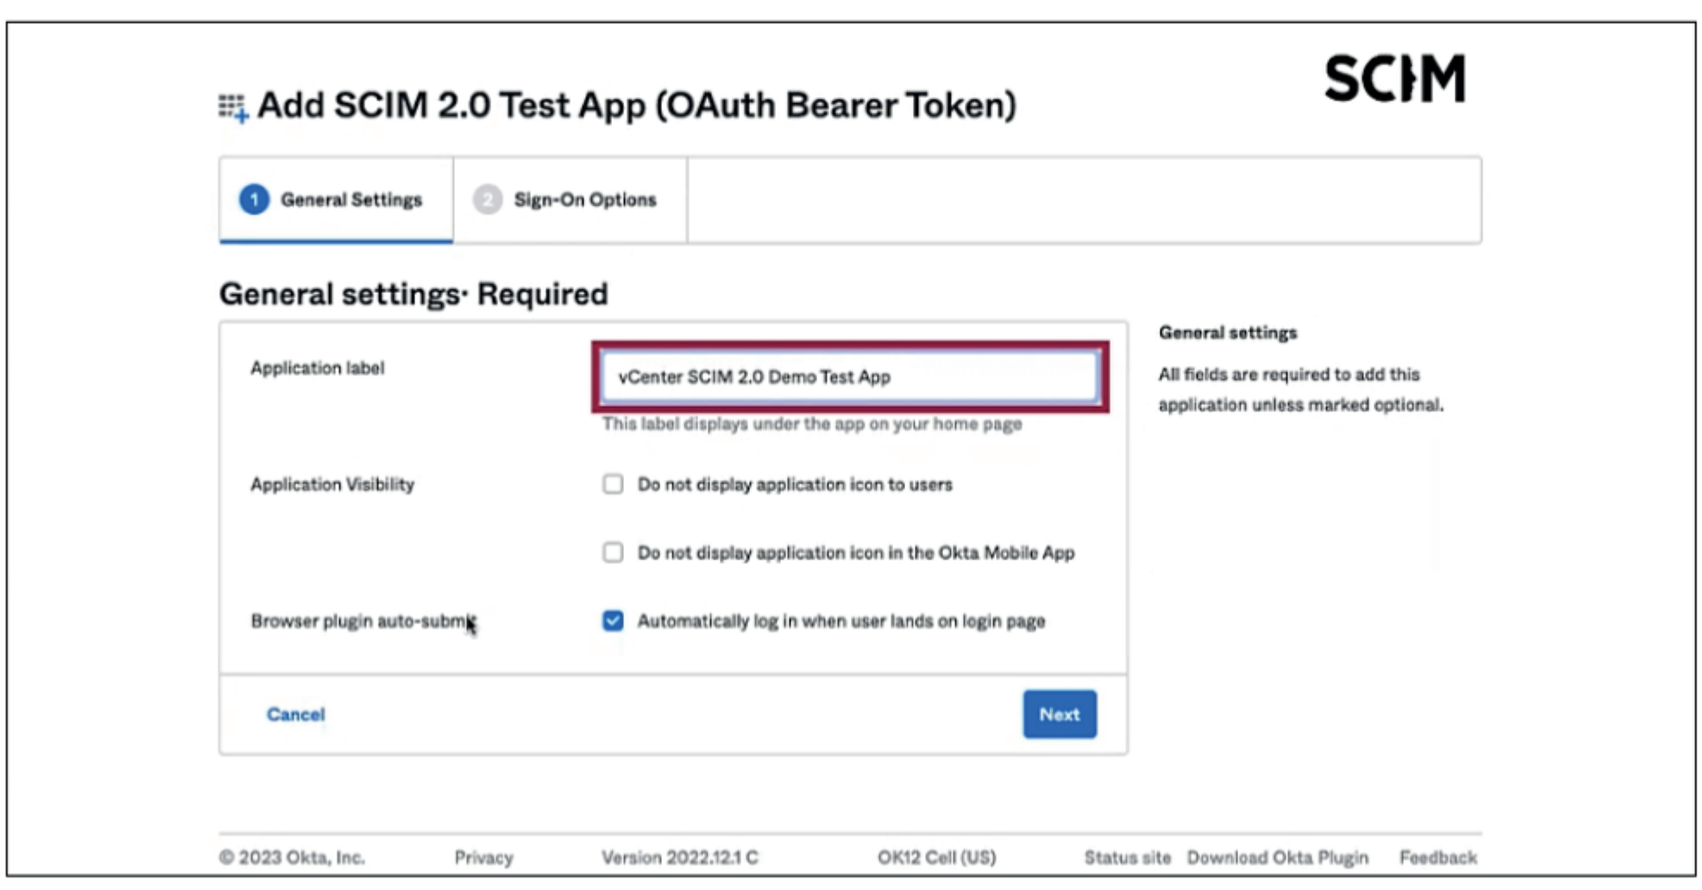

8-61 Configure vCenter Identity Provider Federation for Okta (5)

In the vSphere Client, you configure the group membership for Okta Authorization and verify the login.

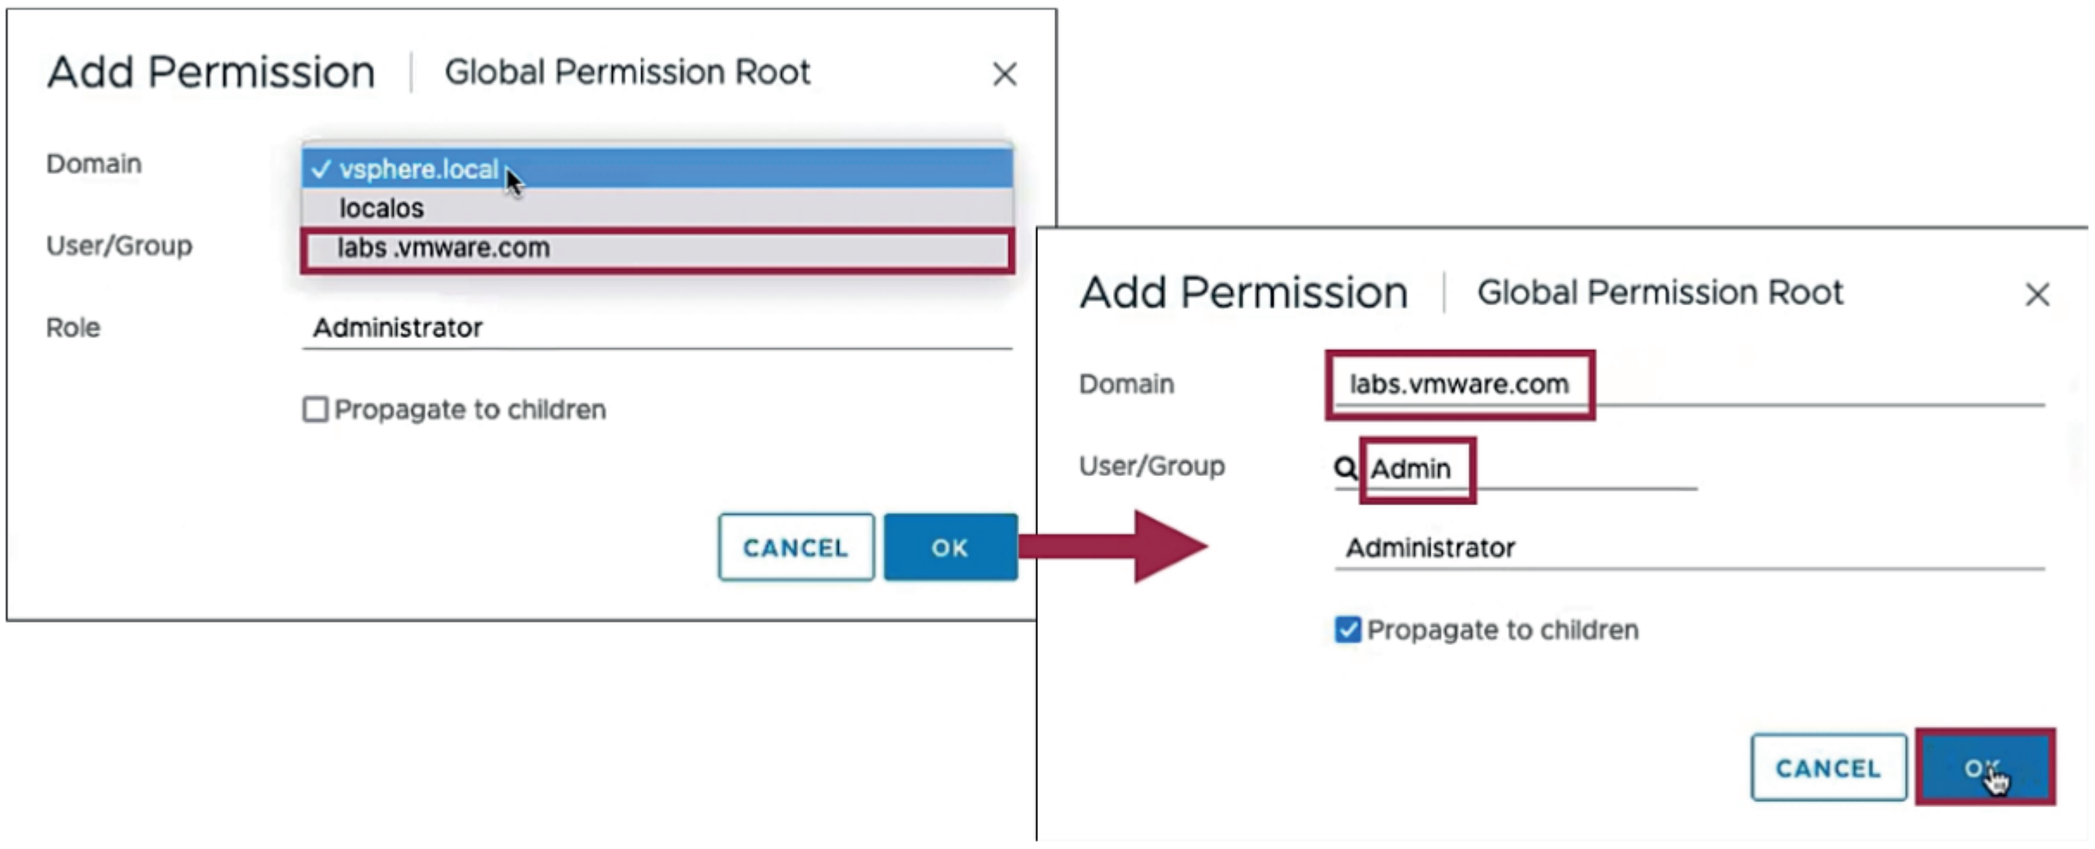

Technet24

8-62 Login Workflow of the vCenter Identity Provider Federation for ADFS

vCenter, ADFS, and AD interact during the login workflow for vCenter IDP federation for ADFS.

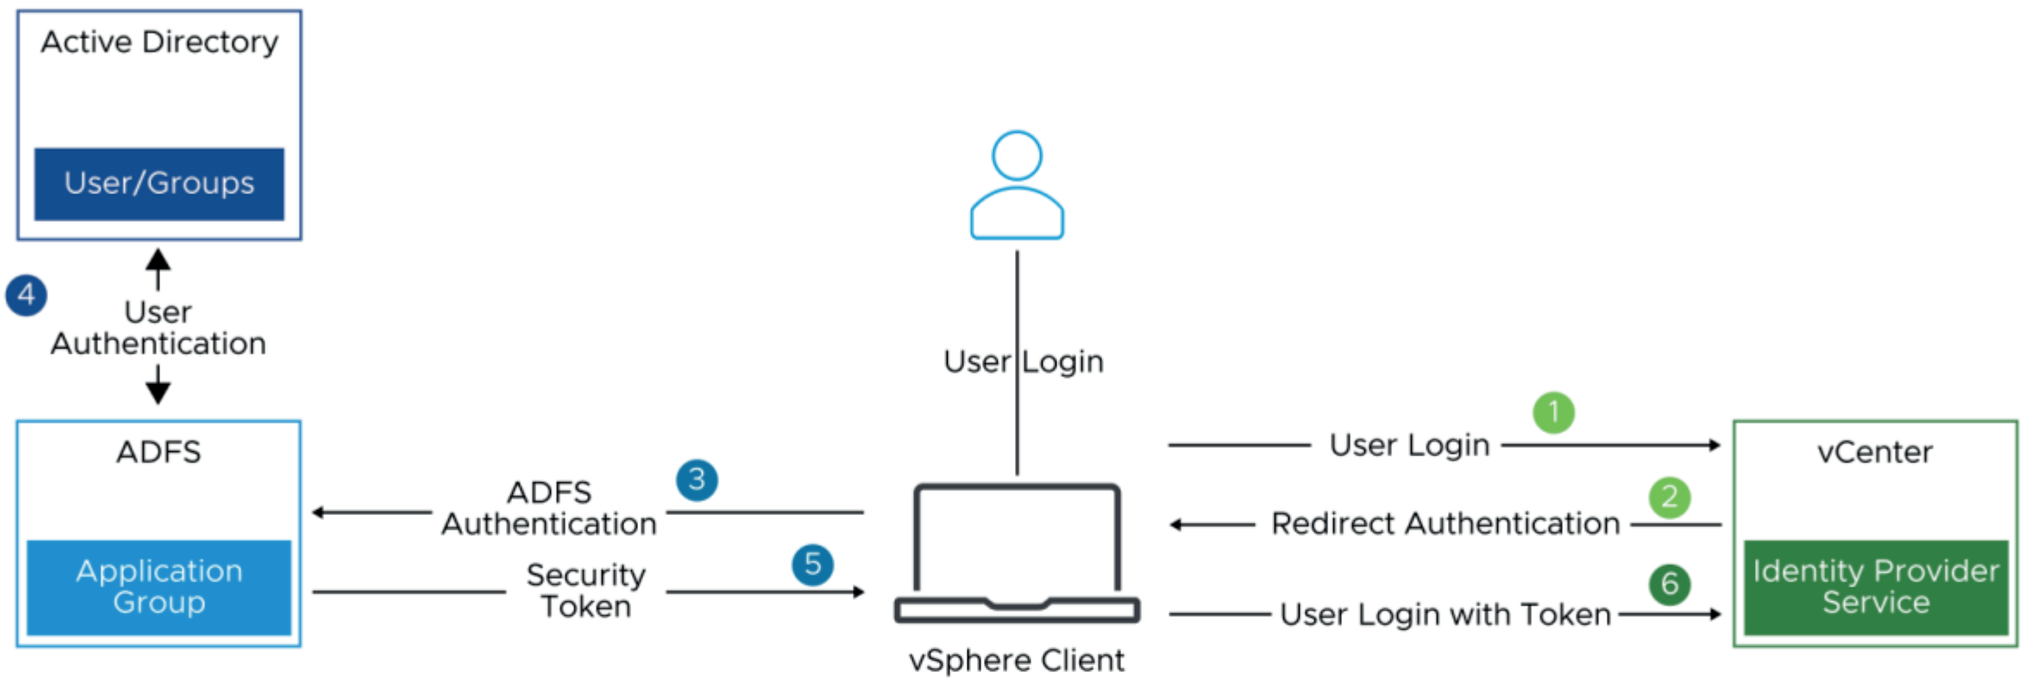

vCenter, ADFS, and AD interact during the login workflow as follows:

1. The user connects to vCenter on the usual vCenter landing page.
2. vCenter redirects the authentication request to ADFS based upon the domain provided.
3. If needed, ADFS prompts the user to log in with Active Directory credentials.
4. ADFS authenticates the user with Active Directory.
5. ADFS issues a security token with group information from Active Directory.
6. vCenter uses the token to log in the user.

If the external IDP is unreachable, the login process falls back to the vCenter landing page, showing an appropriate information message. Administrative users can still log in using their local accounts (local OS or vsphere.local identity sources).

8-63 Configuring vCenter Identity Provider Federation in vSphere (1)

In the vSphere Client, you obtain the redirect URIs from vCenter.

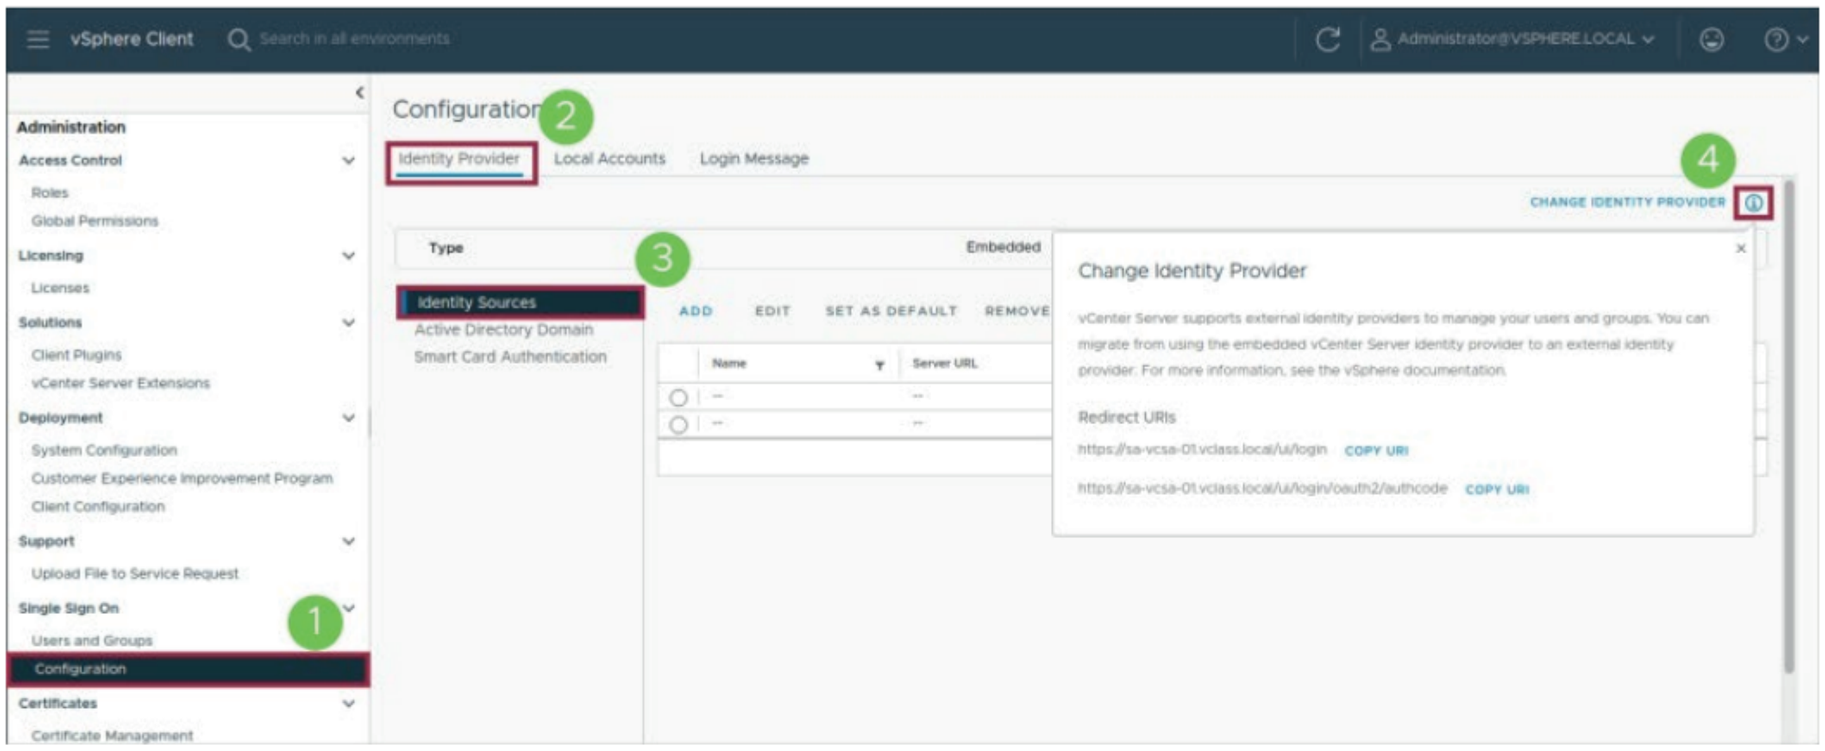

Technet24

8-64 Configuring vCenter Identity Provider Federation in vSphere (2)

The vCenter IDP federation is not enabled by default. After vCenter installation, you have the option to manually configure it.

The high-level steps to configure IDP federation are as follows:

1. Obtaining the redirect URIs from vCenter
2. Creating an ADFS application group and configuring it for vCenter in the Windows configuration
3. Creating an IDP on vCenter in the vSphere Client
4. Configuring group memberships from the Users and Groups tab or privileges through global or object permissions in the vSphere Client for ADFS authorization
5. Verifying login of an AD user to vCenter

To configure group membership for vCenter for AD FS Authorization.

- Click the **Groups** tab from Administration > Single Sign On > Users and Groups.
- Click the Administrators group and click Add Members.
- Select the domain from the drop-down menu.
- In the text box below the drop-down menu, enter the first few characters of AD FS group that you want to add then wait for the drop-down selection to appear.
- Select the AD FS group and add it to the Administrators group.
- Click Save.

For more information about configuring *vCenter IDP Federation for vCente*r, see https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-authentication/GUID-C5E998B2-1148-46DC-990E-A5DB71F93351.html.

8-65 Configuring vCenter Identity Provider Federation in vSphere (3)

In the vSphere Client, you can change the vCenter identity provider.

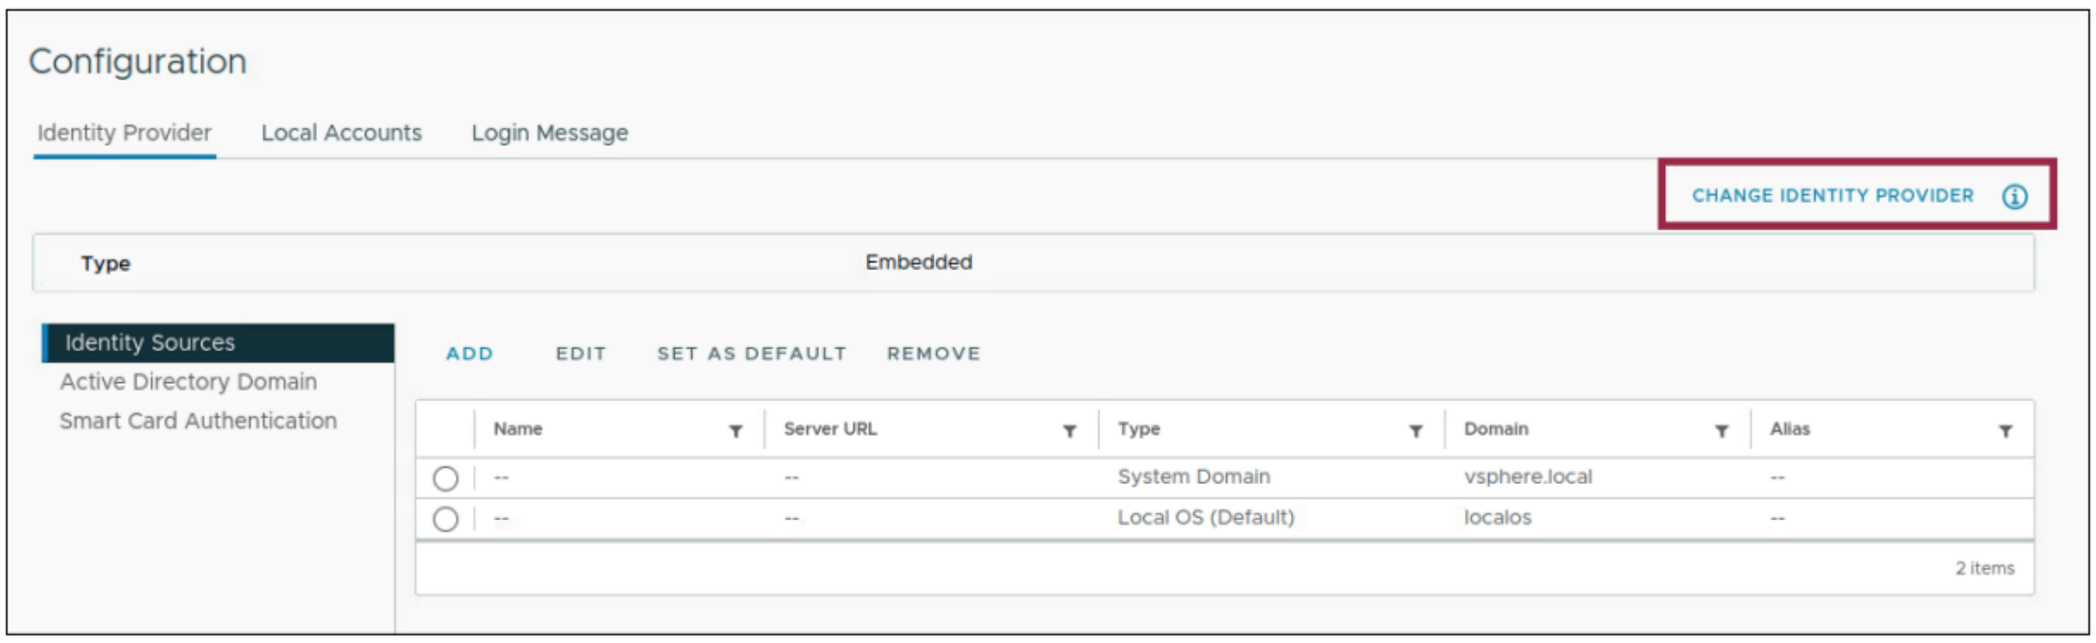

8-66 Configuring vCenter Identity Provider Federation in vSphere (4)

In the vSphere Client, you provide application group information.

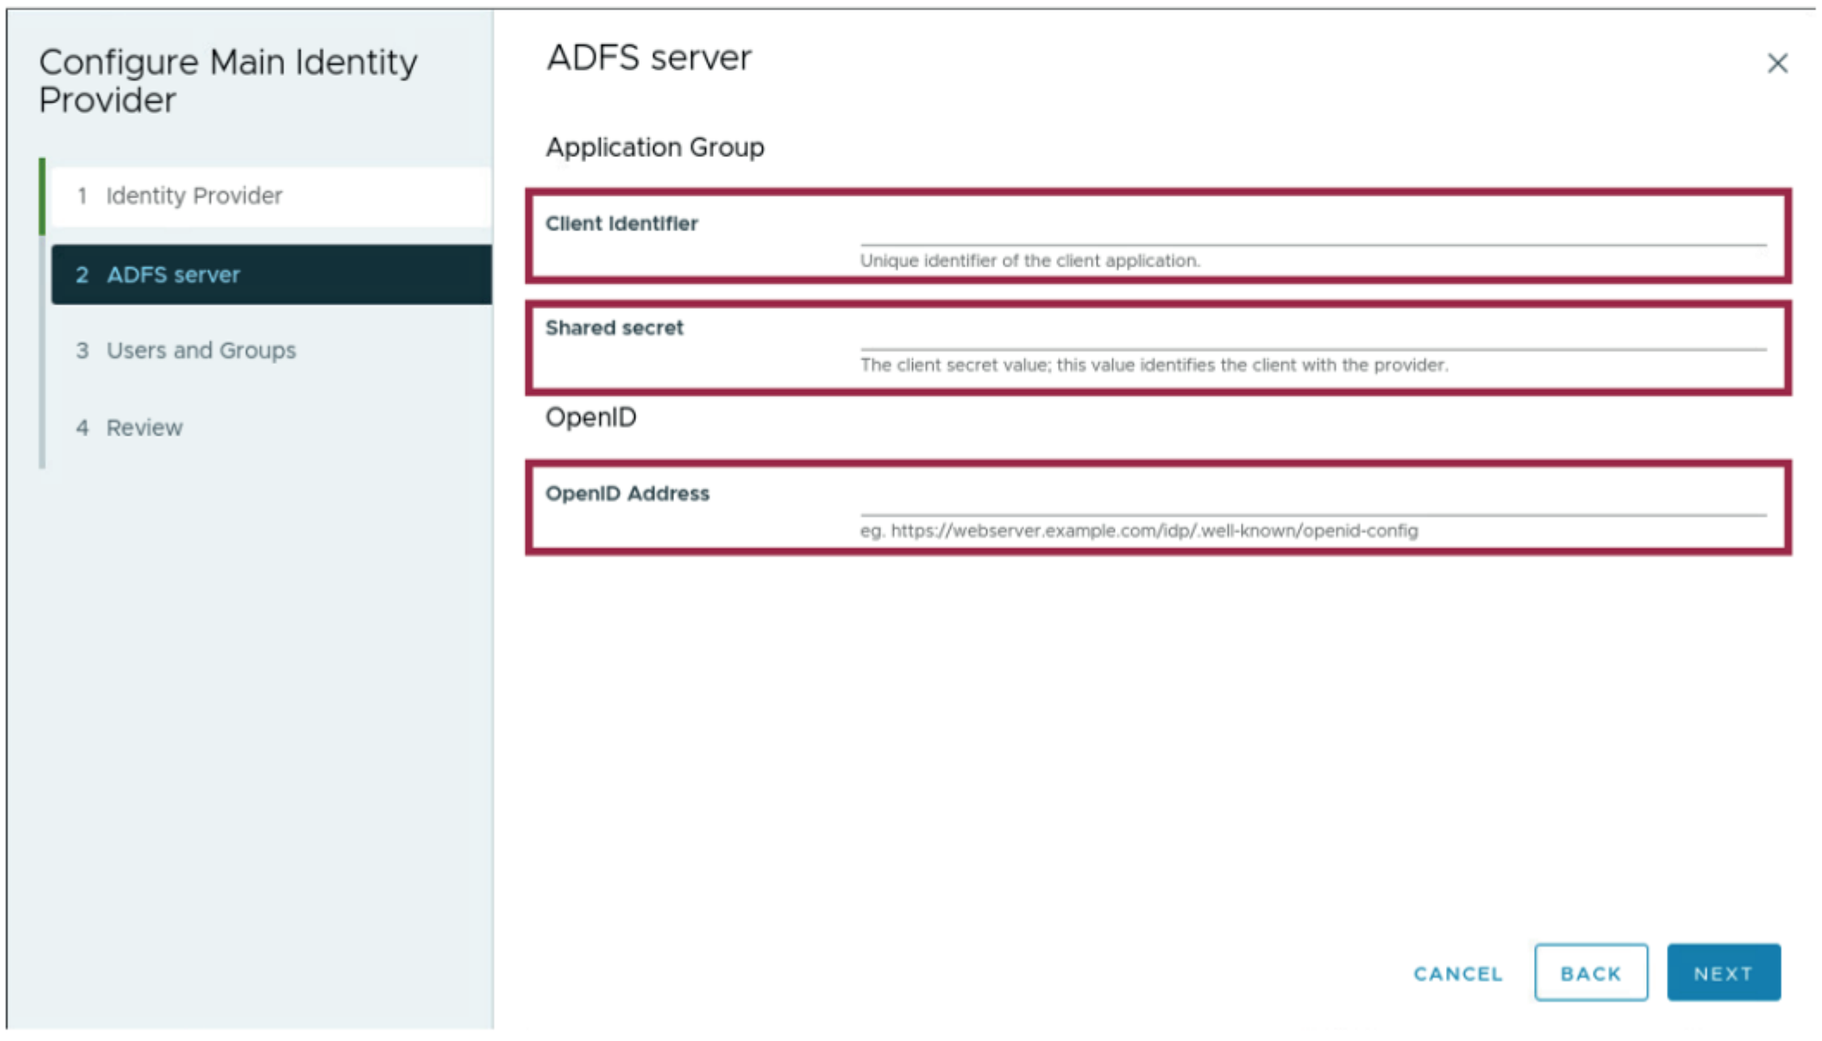

Technet24

8-67 Enabling ADFS on an Existing Enhanced Linked Mode Configuration

To enable ADFS on an existing Enhanced Linked Mode configuration:

1. Deploy a minimum of two vCenter nodes in an Enhanced Linked Mode configuration.
2. Configure ADFS on one of the linked vCenter nodes. The ADFS configuration is replicated to the other vCenter node.
3. Add all redirect URIs for both vCenter nodes to the configured OAuth Application Group in ADFS.

8-68 Lab 19: (Simulation) Configuring Identity Federation to Use Microsoft ADFS

Configuring Identity Federation to use Microsoft ADFS:

1. Configure vCenter Identity Provider Federation
2. Log In to vCenter Using an AD Account

8-69 Review of Learner Objectives

- Describe identity federation
- Recognize use cases for identity federation
- Describe the architecture of identity federation
- Configure identity federation

8-70 Key Points

- When defining a role, use the smallest number of privileges possible for better security and added control.
- Hardening a vSphere environment to be more secure involves configuring the vCenter system, ESXi hosts, VMs, and the vSphere network.
- Identity sources with MFA provide an additional layer of security for authenticating users accessing the vSphere environment.

Questions?

Technet24

Module 9

vSphere Trusted Environments and VM Encryption

9-2 Importance

With vSphere Trusted environments and Virtual Machine Encryption, you can encrypt your sensitive ESXi hosts and workloads in an even more secure way.

9-3 Module Lessons

1. Booting ESXi Securely
2. Securing Virtual Machines
3. VM Encryption
4. Key Providers
5. vSphere Trust Authority

Technet24

9-4 Lesson 1: Booting ESXi Securely

9-5 Learner Objectives

- Use UEFI Secure Boot for ESXi hosts
- Plan for secure boot and TPM 2.0 support for an ESXi host

9-6 UEFI Secure Boot for ESXi Hosts

UEFI Secure Boot is part of the Unified Extensible Firmware Interface (UEFI) firmware standard.

With UEFI Secure Boot enabled, a machine refuses to load any UEFI driver unless the operating system bootloader is cryptographically signed.

ESXi supports UEFI Secure Boot if it is enabled in the hardware.

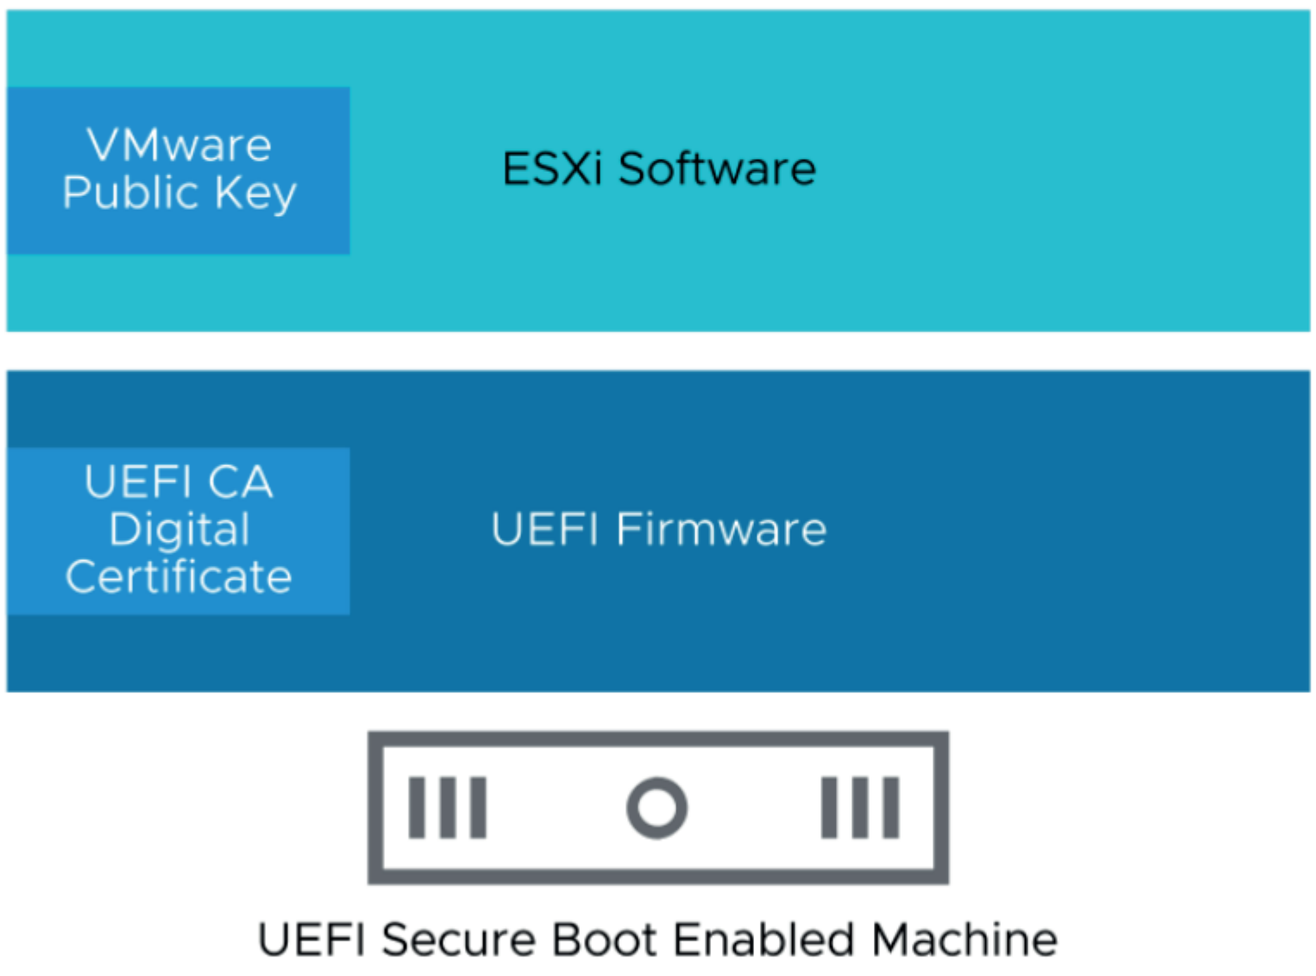

UEFI Secure Boot Enabled Machine

To ensure that boot code is executed without modification, Unified Extensible Firmware Interface (UEFI) uses a digital signature provided by a trusted code creator. The digital signature is embedded in every executable code section. Using public-private key pairs, the code creator signs the code with a private key, which can be checked against the public key in a preinstalled signature before the code is executed. If the executable code is marked as modified or invalid, the code is not executed in the boot path. The system might take an alternate boot path, or the user might be notified to take remedial actions.

9-7 UEFI Secure Boot Sequence for ESXi Hosts

ESXi supports UEFI Secure Boot at each level of the boot stack.

The boot sequence proceeds as follows:

1. The ESXi bootloader contains a VMware public key. The bootloader uses this key to verify the signature of the kernel and a small subset of the system that includes a secure boot VIB verifier.
2. The VIB verifier verifies every VIB package that is installed on the system. At this point, the entire system boots with the root of trust in certificates that are part of the UEFI firmware.

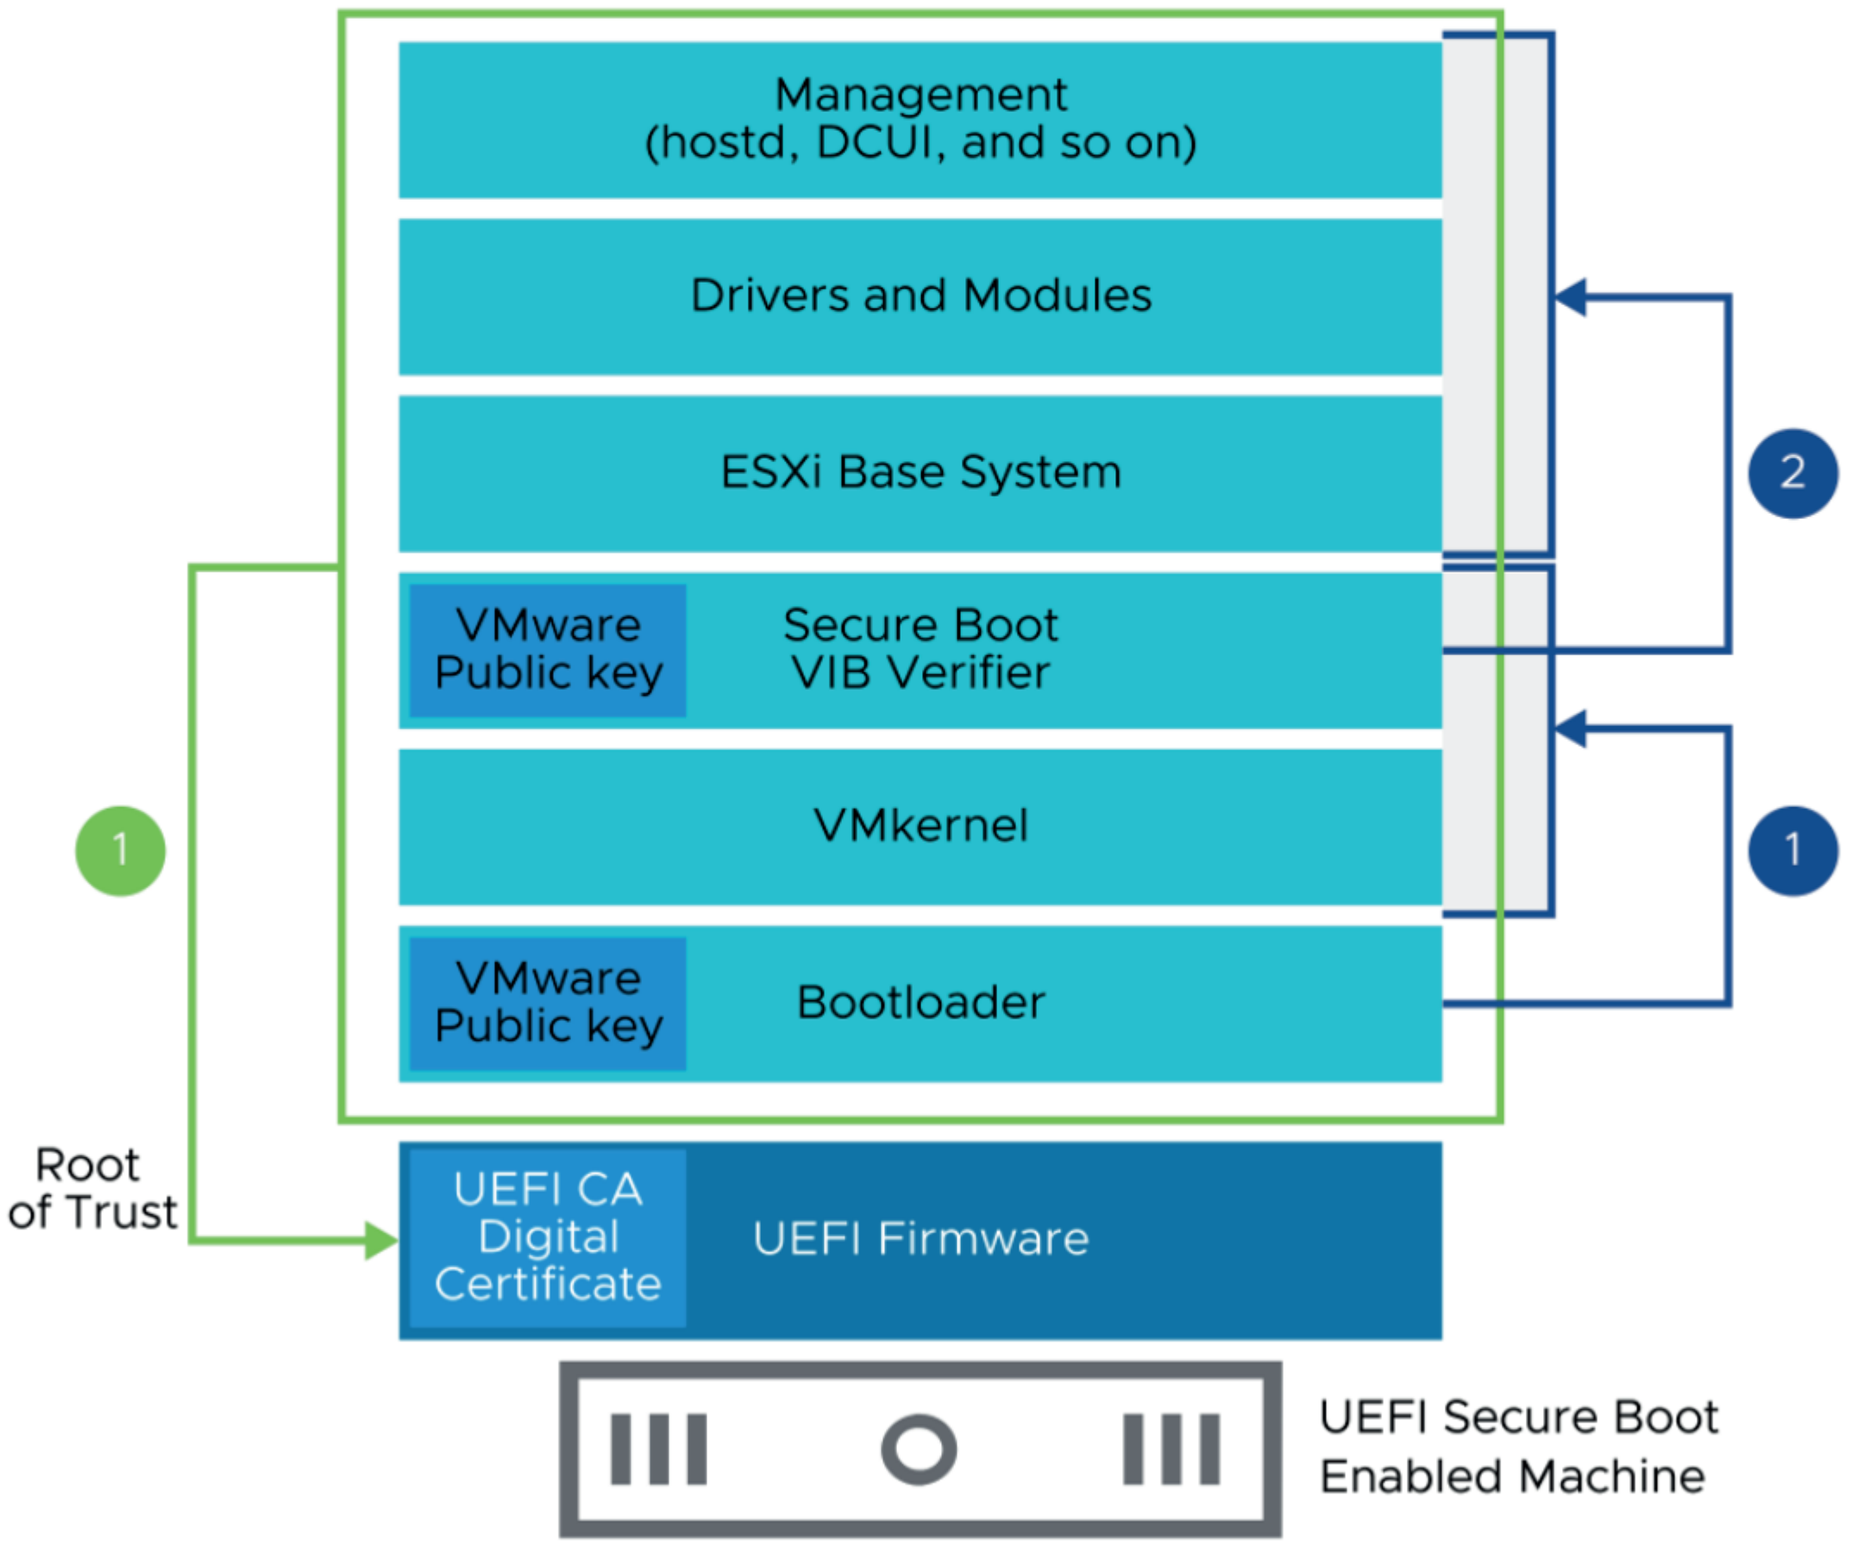

Technet24

9-8 Troubleshooting UEFI Secure Boot

After you install ESXi on a host machine in UEFI mode, the machine might fail to boot.

If a package (VIB or driver) has been tampered with, you will encounter a purple screen.

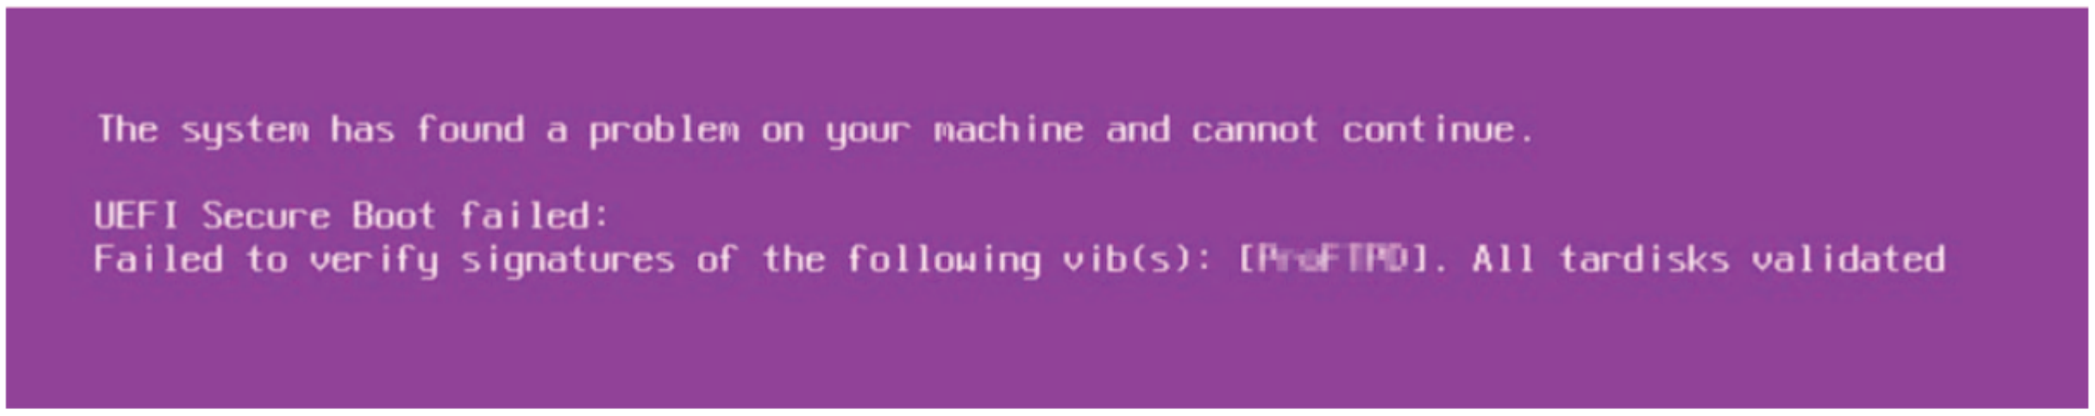

To resolve issues with secure boot, follow these steps.

- Reboot the host with secure boot disabled.
- Run the secure boot verification script.
  (For more information about how to *Run the Secure Boot Validation Script on an Upgraded ESXi Host* see https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-upgrade/GUID-DD075A38-8900-459F-BD9D-69DC87CCE11B.html)
- Examine the information in the /var/log/esxupdate.log file.

9-9 UEFI Secure Boot Validation Script

Running the secure boot validation script allows you to validate whether an ESXi host could be booted with secure boot or not.

Prior to running the script, you need to verify:

- The hardware supports UEFI secure boot
- All VIBs are signed with an acceptance level of at least PartnerSupported

If the ESXi host has any unsigned VIB packages, the validation script will fail.

```
[root@sa-esxi-04:~] /usr/lib/vmware/secureboot/bin/secureBoot.py
All vib signatures verified. All tardisks validated
[root@sa-esxi-04:~]
```

If you include VIBs at the CommunitySupported level, you cannot use secure boot.

9-10 TPM 2.0 on ESXi Hosts

vSphere 6.7 introduced support for TPM 2.0.

ESXi can use Trusted Platform Module (TPM) chips to enhance host security.

TPM protects users from software-based attacks that attempt to steal sensitive information by corrupting system and BIOS code, or by modifying the platform's configuration.

TPM is an industry-wide standard for secure cryptoprocessors. The Trusted Computing Group (TCG) is responsible for TPM technical specifications.

The dedicated microprocessor secures hardware by integrating cryptographic keys into devices.

You are required to have UEFI Secure Boot enabled to enable TPM 2.0.

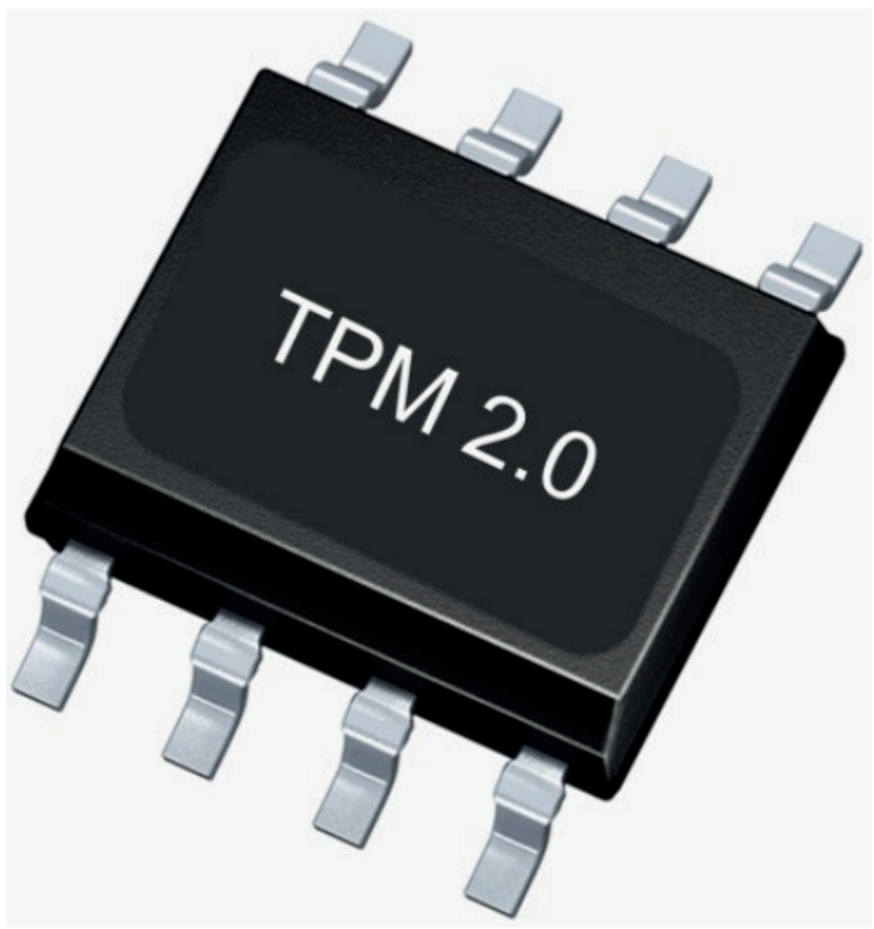

When you install or upgrade to vSphere 7.0 Update 2 or later, and an ESXi host has a TPM, the TPM seals the sensitive information by using a TPM policy based on PCR values for UEFI Secure Boot. This value is loaded during subsequent reboots if the policy is satisfied as true.

Technet24

9-11 ESXi Quick Boot Support on Servers with TPM 2.0 Processor

Hosts with TPM 2.0 enabled go through a secure boot and attestation process, which verifies the configuration of the host, and is a powerful way to detect and protect from malware and other misconfigurations.

Introduced in vSphere 8 Update 1, Quick Boot is now compatible.

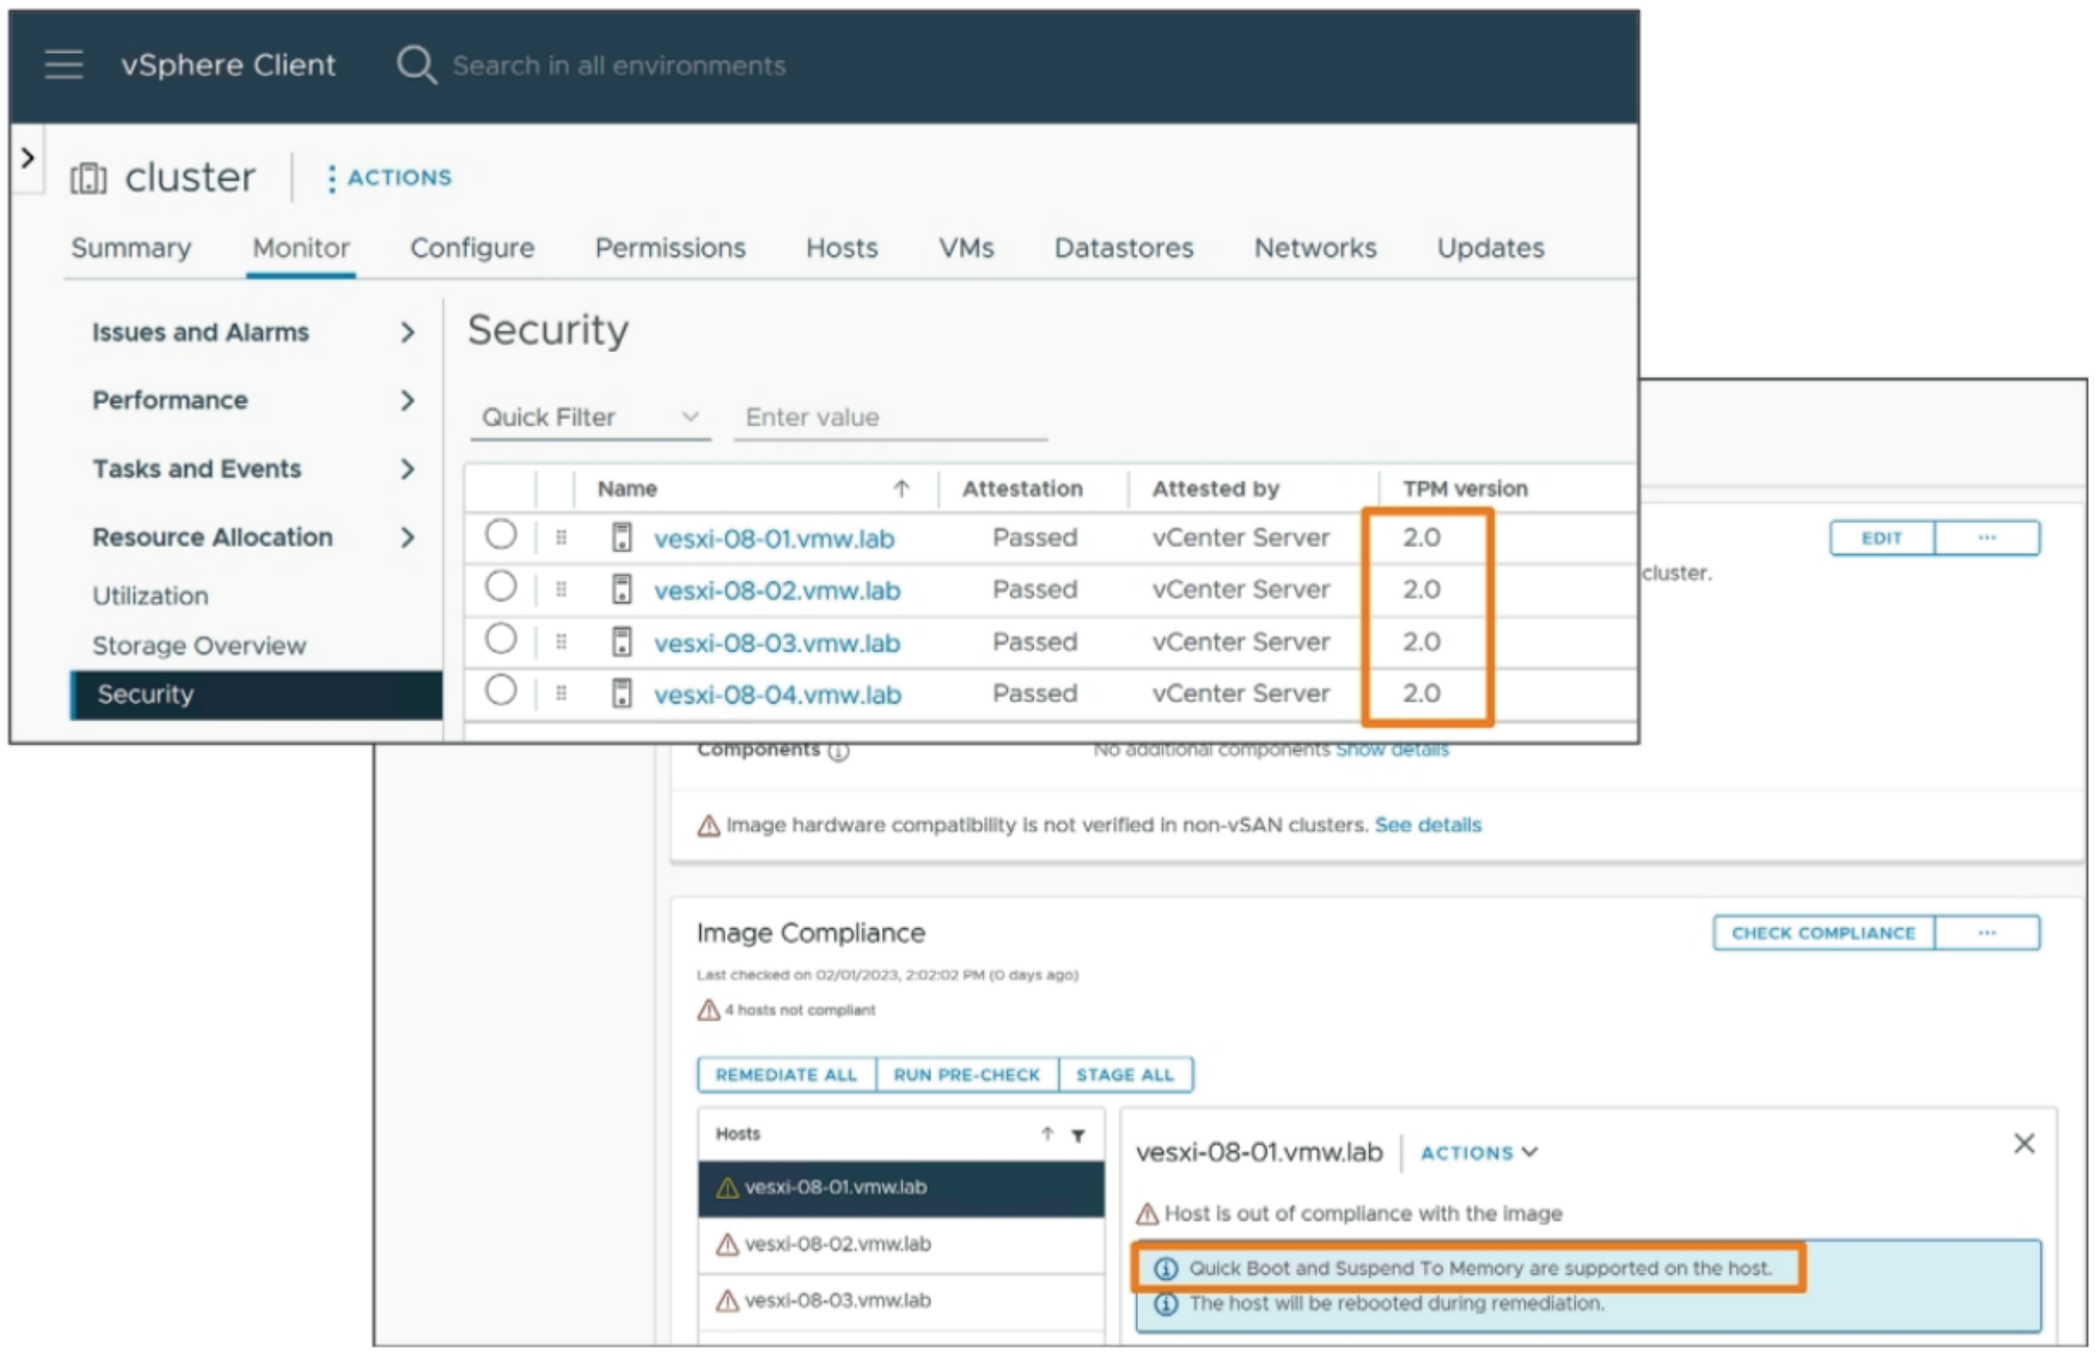

Quick Boot is used in lifecycle management activities like patching, upgrades, etc., and saves considerable time.

9-12 About Remote Attestation

TPM hardware attests to an ESXi host's identity. This process is called remote attestation.

Remote attestation authenticates and attests to the state of the host's software at a given point in time.

TPM hardware records and securely stores measurements of the hypervisor image:

- The measurements are stored in Platform Configuration Registers (PCRs).
- The recorded state of these measurements is called a quote.
- These measurements can be used to detect changes for anything that can be loaded into memory.

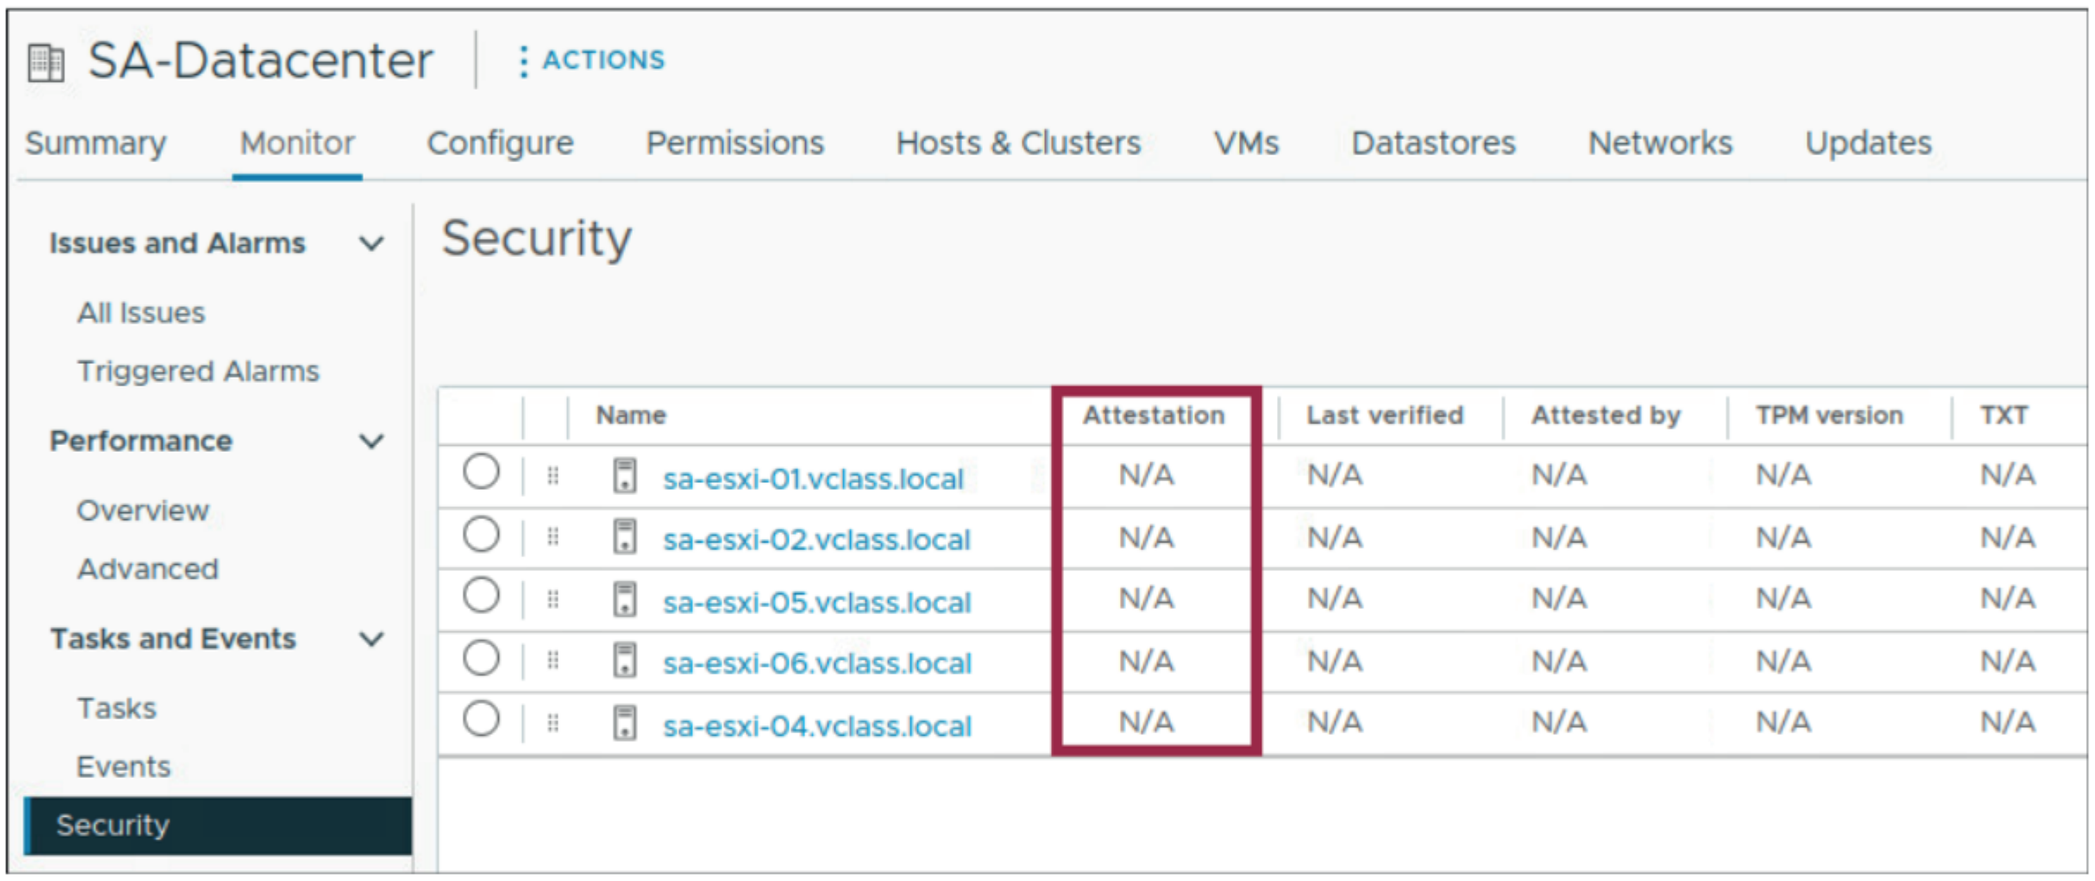

Technet24

9-13 Activity: ESXi Secure Boot (1)

Which of the following statements accurately describes the behavior of an ESXi host when UEFI Secure Boot is enabled and an invalid VIB signature is detected during the boot sequence?

◊ ❒ A purple screen is displayed, indicating that UEFI Secure Boot failed due to an invalid VIB.

◊ ❒ The host continues to boot normally without any indication of the failed UEFI Secure Boot.

◊ ❒ UEFI Secure Boot cannot be enabled on an ESXi host, so this scenario is not applicable

◊ ❒ The host automatically disables UEFI Secure Boot and proceeds with the boot process.

9-14 Activity: ESXi Secure Boot (2)

Which of the following statements accurately describes the behavior of an ESXi host when UEFI Secure Boot is enabled and an invalid VIB signature is detected during the boot sequence?

◊ ✓ A purple screen is displayed, indicating that UEFI Secure Boot failed due to an invalid VIB.

◊ ❒ The host continues to boot normally without any indication of the failed UEFI Secure Boot.

◊ ❒ UEFI Secure Boot cannot be enabled on an ESXi host, so this scenario is not applicable

◊ ❒ The host automatically disables UEFI Secure Boot and proceeds with the boot process.

9-15 Review of Learner Objectives

- Use UEFI Secure Boot for ESXi hosts
- Plan for secure boot and TPM 2.0 support for an ESXi host

9-16 Lesson 2: Securing VMs

9-17 Learner Objectives

- Identify the role of vTPM and VBS for securing VMs
- With Intel Software Guard Extensions (SGX), software programs can create private memory regions called enclaves
- Secure Encrypted Virtualization with Encrypted State (SEV-ES) is a hardware feature that is enabled in AMD CPUs
- Enable encrypted vSphere vMotion

9-18 VM Protection with UEFI Secure Boot

You can enable UEFI Secure Boot in a VM if the following prerequisites are met:

- The VM uses EFI firmware.
- The guest operating system supports UEFI Secure Boot.
- Virtual hardware version 13 or later.

To effectively support UEFI Secure Boot in a VM, see the guest OS vendor documentation on configuring and using EFI firmware.

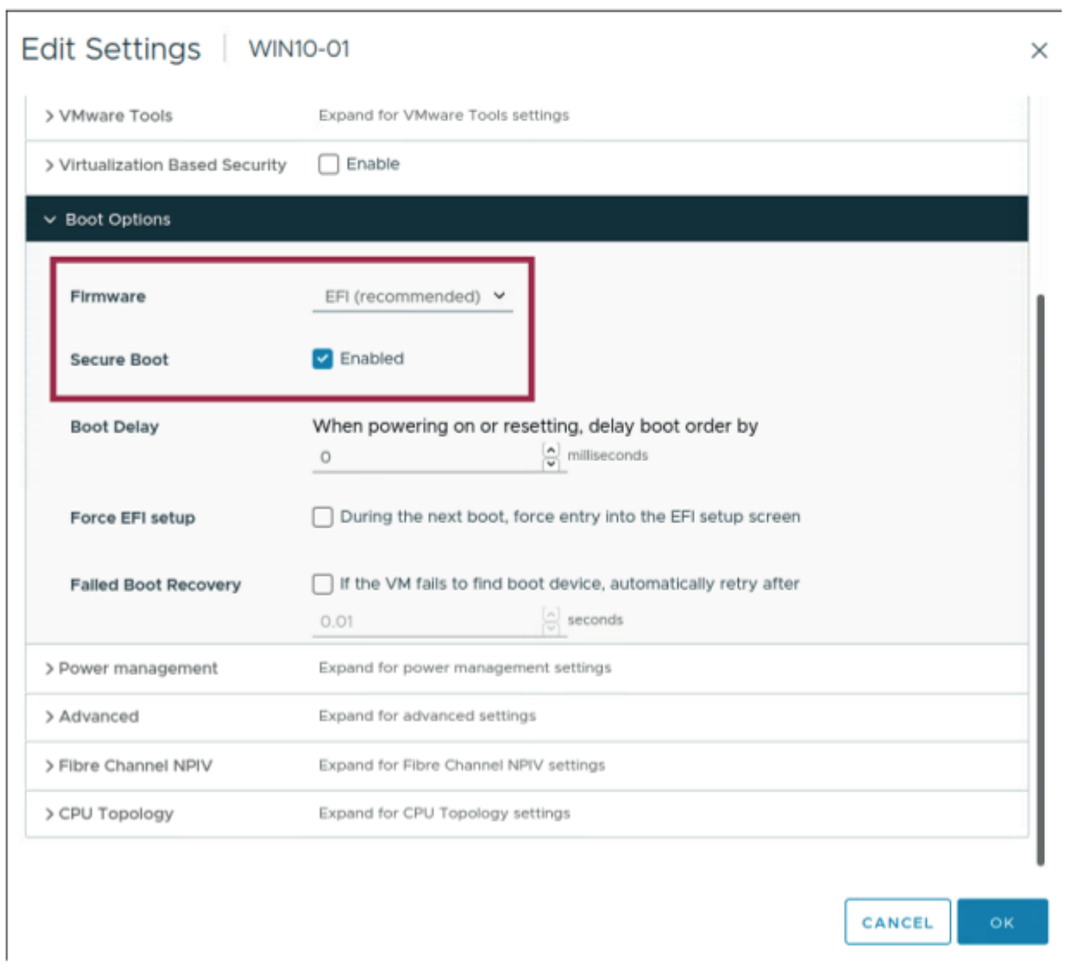

Technet24

For certain VM hardware versions and operating systems, you enable UEFI Secure Boot exactly as you enable it for a physical machine.

In an operating system that supports UEFI Secure Boot, each piece of boot software is signed, including the bootloader, the operating system kernel, and operating system drivers. The VM's default configuration includes several code-signing certificates:

- A Microsoft certificate that is used only for booting Windows
- A Microsoft certificate that is used for third-party code and signed by Microsoft, such as Linux bootloaders
- A VMware certificate that is used only for booting ESXi inside a VM

The VM's default configuration includes one certificate for authenticating requests to modify the UEFI Secure Boot configuration, including the UEFI Secure Boot revocation list from inside the VM, which is a Microsoft key encryption key (KEK) certificate. In most cases, replacing the existing certificates is not necessary.

For a list of operating systems that support UEFI Secure Boot, see the *VMware Compatibility Guide* at https://www.vmware.com/resources/compatibility.

9-19 Virtual TPM on ESXi Hosts

The VMware virtual TPM is compatible with TPM 2.0 and creates a TPM-enabled virtual chip for use by the virtual machine and the guest OS it hosts.

Virtual TPM supports clone and rekey operations with vSphere Native Key Provider, Standard Key Provider and vSphere Trust Authority.

vSphere 8 introduced compatibility with OVF format for VMs with vTPM devices.

9-20 Using vTPM in a VM (1)

vSphere 6.7 introduced support for the Virtual Trusted Platform Module (vTPM) device. With this device, you can add a TPM 2.0 virtual cryptoprocessor to a VM.

Virtual TPM supports clone and rekey operations with vSphere Native Key Provider, Standard Key Provider and vSphere Trust Authority.

A vTPM device is a software emulation of the TPM functionality:

- The guest operating system can create and store private keys in such a way that they are never exposed to this guest operating system.
- The guest operating system can use the private key for encryption or signing.

With a vTPM device, a third party can remotely attest to (validate) the identity of the firmware and the guest operating system.

vTPM has the following use cases:

- An operating system can verify that the firmware loaded was not compromised since the last run.
- An application can verify that the operating system did not load any malicious components.

Compromising the guest operating system usually compromises its secrets, but enabling a vTPM in the guest reduces this risk.

You can add a vTPM device to either a new VM or an existing VM.

Introduced in vSphere 8, is the compatibility with OVF format for VMs with vTPM devices.

Technet24

9-21 Using vTPM in a VM (2)

- vTPM uses VM encryption technology to secure virtual TPM data.
- When you configure vTPM, VM encryption automatically encrypts the VM's home directory (which contains *.nvram, *.vmx, *.vmsn, snapshots, core files, and so on) but not the disks.
- You can choose to add encryption explicitly for the VM and its disks using the VM Encryption feature.
- You can back up a VM that is enabled with a vTPM:
  - The backup must include all VM data, including the `.nvram` file.
  - If your backup does not include the `nvram` file, you cannot restore a VM with a vTPM.
  - Because the VM home files are encrypted, you must ensure that the encryption keys are available at the time of a restore.
- You can remove a vTPM from a VM. However, removing vTPM causes all encrypted information on a VM to become unrecoverable.
- Before removing a vTPM from a VM, deactivate any applications in the guest operating system that use vTPM.

vTPM can be enabled on a VM by editing the VM's settings with the vSphere Client.

vTPM can be removed from a VM but the VM must be powered off first. Before you remove vTPM, deactivate any applications that use vTPM. If you do not deactivate these applications, the VM might fail to boot when powered on.

9-22 Prerequisites for Adding vTPM

To use vTPM, your vSphere environment must meet specific requirements:

- Component requirements:
  - ESXi 6.7 of later
  - vCenter Server 6.7 and later for Windows virtual machines, vCenter Server 7.0 Update 2 and later for Linux virtual machines.
  - Native or external key management server that is configured in vCenter to encrypt a VM
- VM requirements:
  - EFI firmware
  - VM hardware version 14 or later
  - Windows 10 (64-bit) or Windows Server 2016 (64-bit) or later guest operating system

vTPM does not require a physical TPM 2.0 chip to be present on the ESXi host. However, if you want to perform host attestation, an external entity, such as a TPM 2.0 physical chip, is required.

Technet24

9-23 vSphere Fault Tolerance with vTPM

Introduced in vSphere 8 Update 1, Fault Tolerance for a VM employing a vTPM module is now supported

This capability helps administrators achieve continuous availability and security for mission-critical VMs.

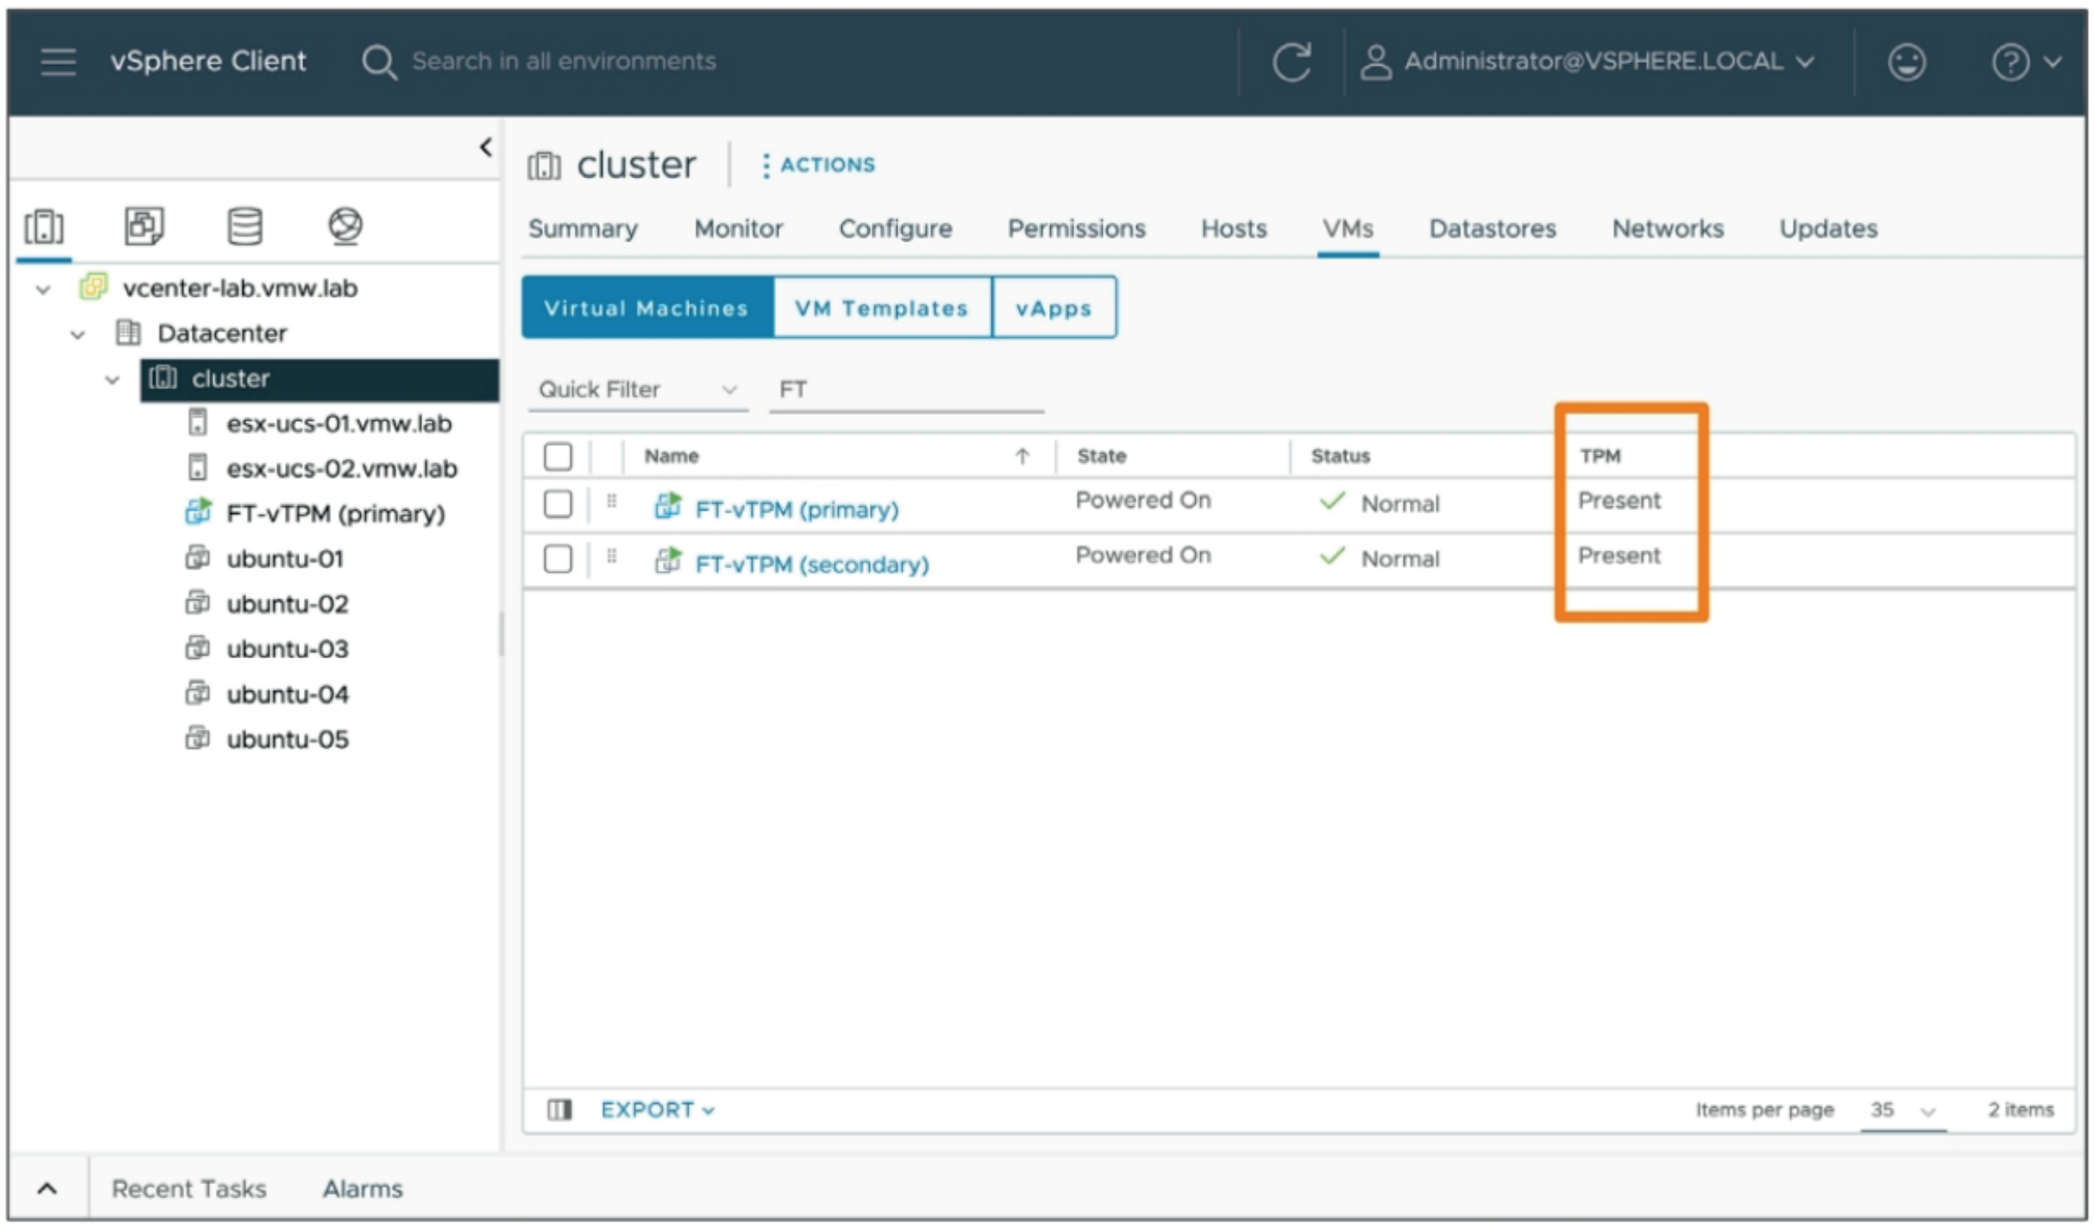

Fault Tolerance provides continuous availability for a VM by maintaining an identical VM that can quickly failover in the event of a failure.

9-24 Virtualization-Based Security for Windows Guest Operating Systems

Windows 10 and Windows Server 2016 (64-bit versions only) contain a security technology called virtualization-based security (VBS).

VBS uses Microsoft's Hyper-V virtualization technology to isolate core Windows operating system services in a virtualized environment.

This virtualized environment provides the following security benefits:

- An extra level of protection to the key services in Windows
- An isolated execution environment for user activity
- Enhanced host operating system protection from attacks

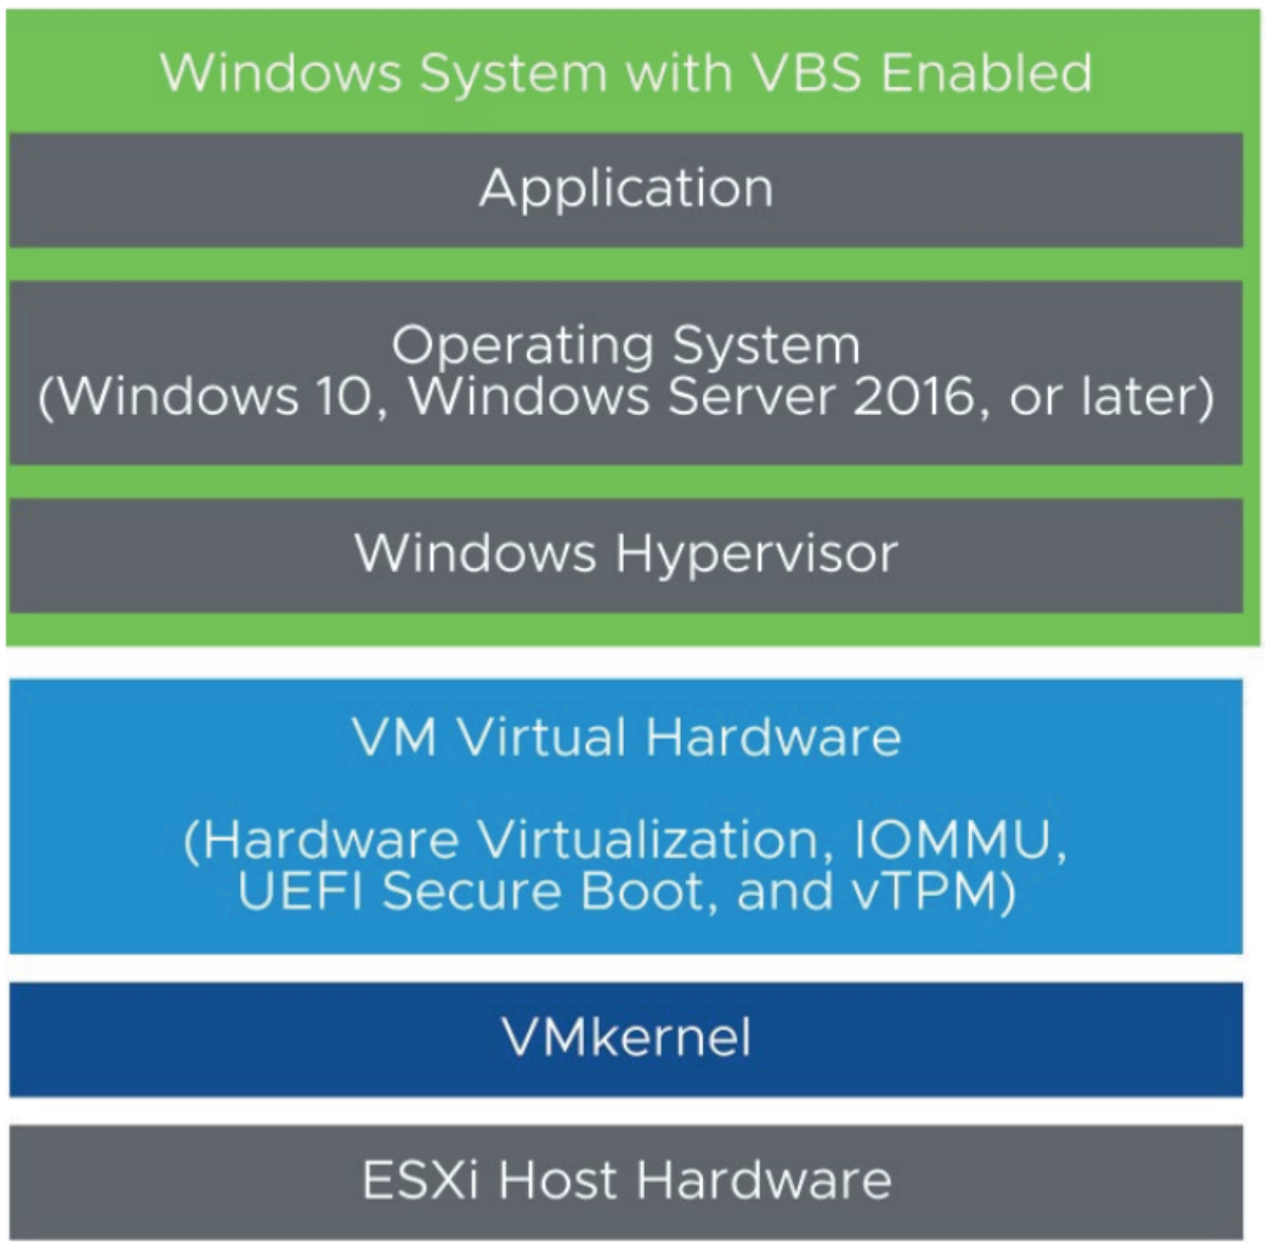

Microsoft virtualization-based security (VBS), a feature of Windows 10 and Windows Server 2016 and later operating systems, uses hardware and software virtualization to enhance system security by creating an isolated, hypervisor-restricted, specialized subsystem.

Technet24

With VBS, you can use the following Windows security features to harden your system and isolate key system and user secrets so that they are not compromised:

- Credential Guard: Aims to isolate and harden key system and user secrets against compromise
- Device Guard: Provides a set of features that work together to prevent and eliminate malware from running on a Windows system
- Configurable Code Integrity: Ensures that only trusted code runs from the boot loader onwards

For more information about virtualization-based security, see the Microsoft documentation.

9-25 VBS Requirements and Limitations

To use VBS in a VM, you must meet prerequisites:

- The VM must be of virtual hardware version 14 or later.
- The guest operating system must be either Windows 10 or Windows Server 2016, or later (64 bit).
- The VM must be on an ESXi host version 6.7 or later for Intel-based hosts and 7.0U2 or later for AMD-based hosts.

VMs that have VBS enabled use UEFI and secure boot by default.

The following features are not supported with VBS:

- vSphere Fault Tolerance
- PCI passthrough
- Hot add of CPU or memory

Using AMD CPUs for VBS requires vSphere 7.0 Update 2 or later.

9-26 Enabling VBS

To enable VBS when you create VM, select Windows 10 or Windows Server 2016 or later and select the **Enable Windows Virtualization Based Security** check box.

Enabling VBS exposes hardware-assisted virtualization, the input-output memory management unit (IOMMU), EFI, and UEFI Secure Boot to the guest operating system.

After you enable VBS for a VM, you must enable VBS within the guest operating system.

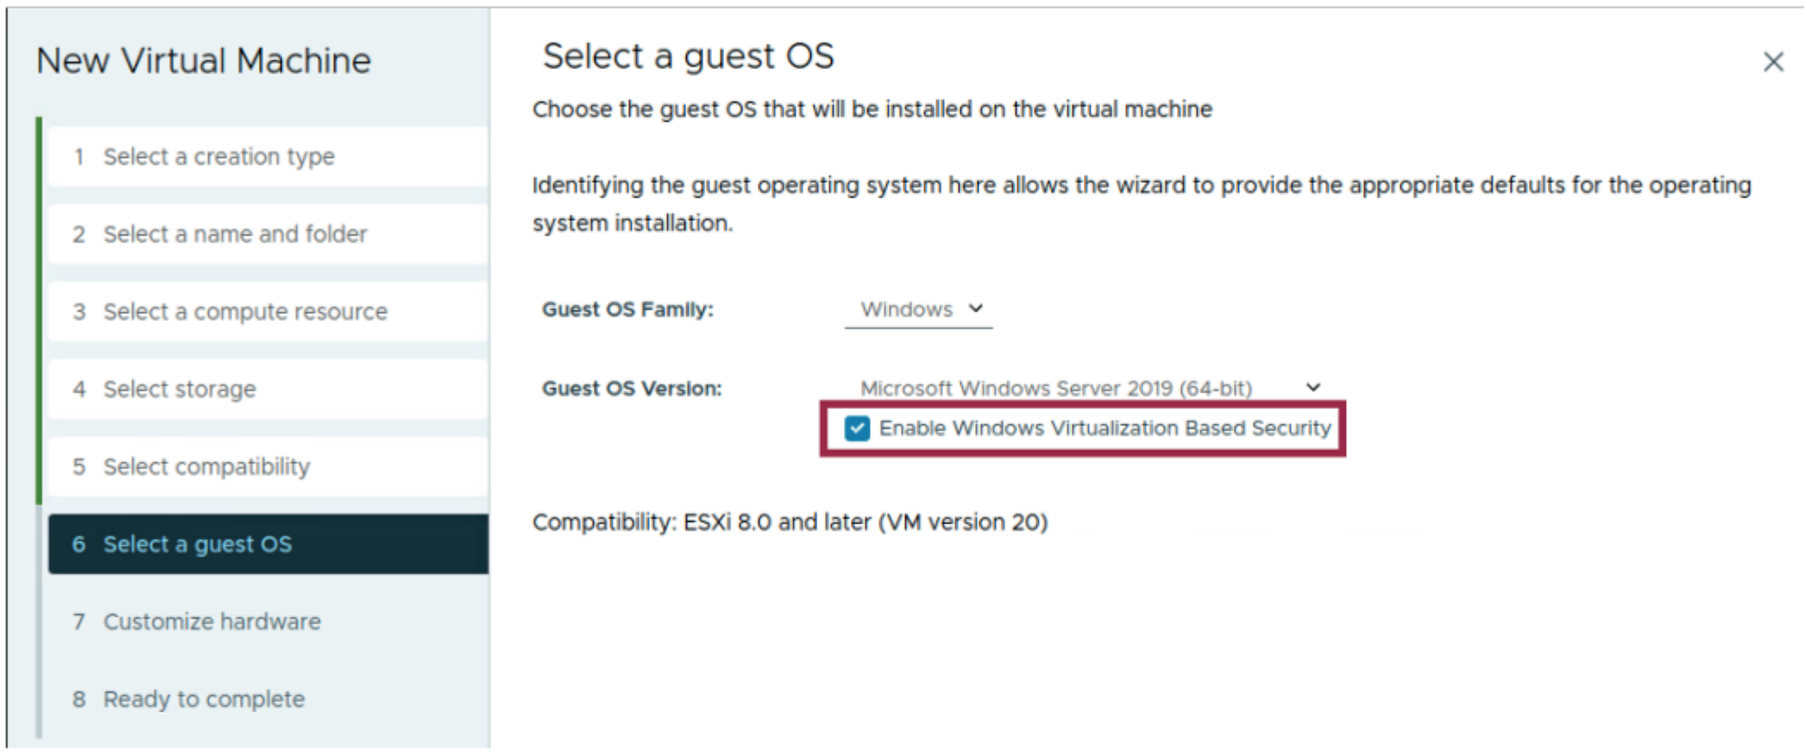

Technet24

9-27 About SGX

Intel Software Guard Extensions (SGX) technology meets the needs of the trusted computing industry. Trusted computing refers to various technologies and proposals for resolving computer security problems:

- With Intel SGX, software programs can create private memory regions called enclaves.
- Data in enclaves can be accessed only by the intended program.
- The enclave region is isolated from other programs, operating systems, hypervisors, and so on.
- Enclaves are used by the software program to secure data.
- Virtual SGX (vSGX) exposes Intel's SGX to VMs running in a vSphere environment.

The problems are resolved through hardware enhancements and associated software modifications. Several major hardware manufacturers and software vendors, collectively known as the Trusted Computing Group (TCG), cooperate in this venture and create specific plans.

For more information about the Trusted Computing Group, see https://trustedcomputinggroup.org.

Data inside an application can be compromised by rogue software or a bug in the operating system or in the hypervisor.

Programmers can use Intel SGX to create private memory regions called enclaves. Data in enclaves is protected from unauthorized access and modification by rogue software running at higher privilege levels. Intel SGX provides this protection without disrupting the ability of legitimate software to schedule and manage the use of platform resources.

vSGX is implemented as part of the vSphere core virtualization stack. With vSGX, applications can run on VMs to create their own enclaves.

9-28 About AMD SEV-ES

Secure Encrypted Virtualization with Encrypted State (SEV-ES) is a hardware feature that is enabled in AMD CPUs.

SEV-ES encrypts the VM guest memory and register state, protecting it against access from the hypervisor or other VMs running on the hypervisor.

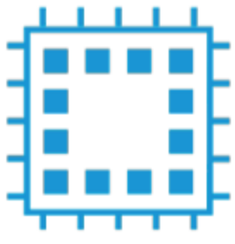

9-29 About Encrypted vSphere vMotion

Encrypted vSphere vMotion secures the confidentiality, integrity, and authenticity of data that is transferred during a vMotion task.

Encrypted vSphere vMotion supports all variants of vSphere vMotion for unencrypted VMs, including migration across different vCenter instances and versions.

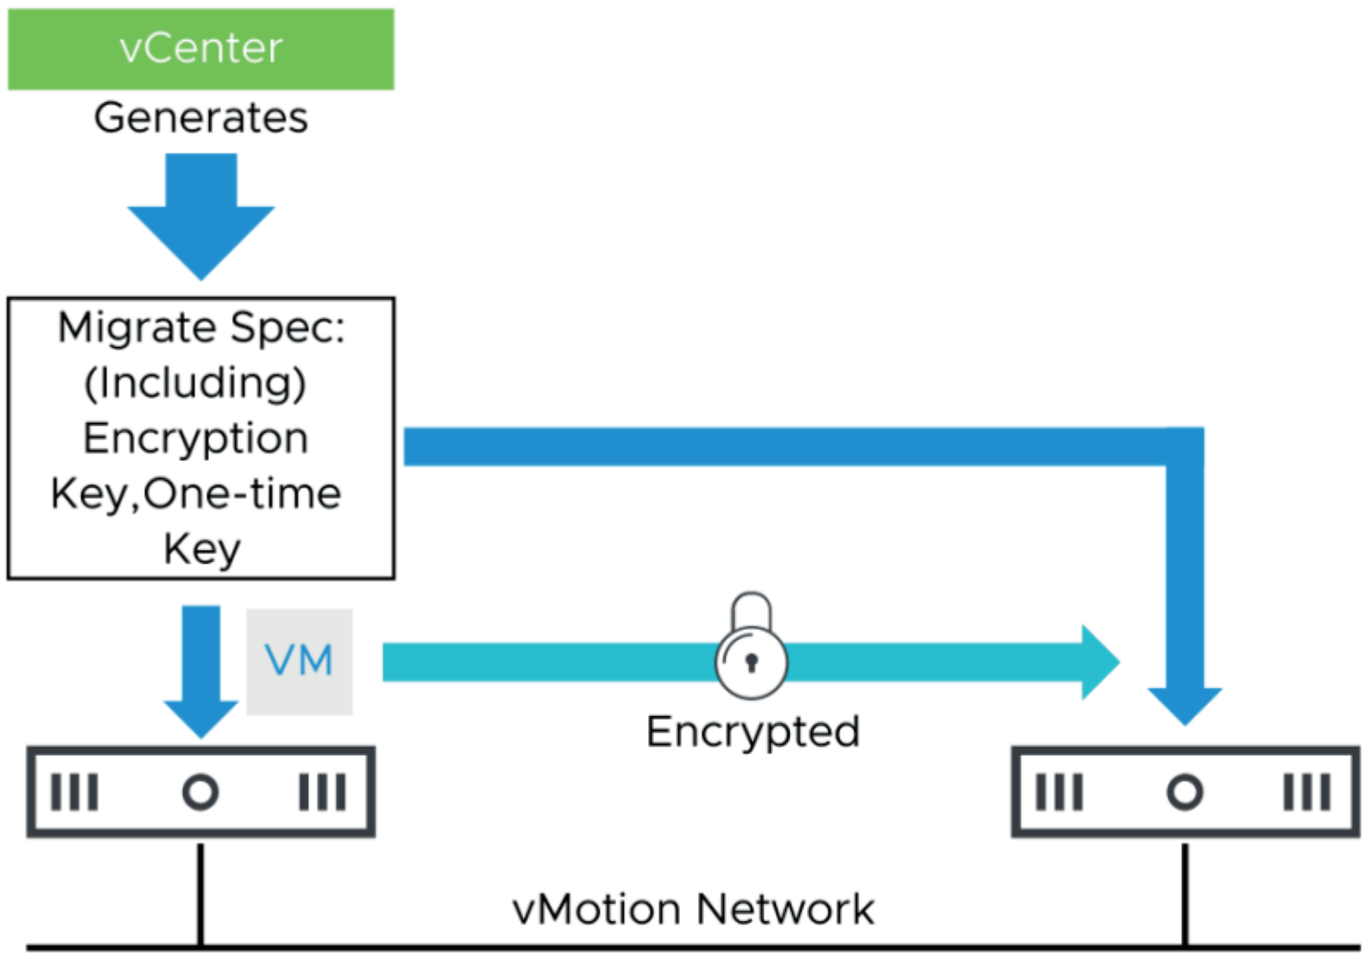

Encrypted vSphere vMotion secures communication on the vMotion network so that sensitive data cannot be discovered during memory migration.

vCenter generates a one-use (nonce) encryption key and includes this key in a communication (or migration specification) to both the source and the destination ESXi hosts. This key is used to encrypt and decrypt the data sent over the vMotion network.

In a vSphere vMotion migration across vCenter instances, the local vCenter instance generates the encryption key and includes the key in a communication over an TLS-secured channel to the remote vCenter instance. The two vCenter instances pass the required keys to their local ESXi hosts, and these hosts communicate directly using these keys.

For encrypted disks, the data is transmitted encrypted in all cases.

For unencrypted disks, the following applies: If disk data is transferred within a host, that is without changing the host, you change only the datastore, the transfer is unencrypted. If disk data is transferred between hosts and encrypted vMotion is used, the transfer is encrypted. If encrypted vMotion is not used the transfer is unencrypted

For VMs that are encrypted, migration with vSphere vMotion always uses encrypted vSphere vMotion. You cannot turn off encrypted vSphere vMotion for encrypted VMs.

Technet24

9-30 Configuring Encrypted vSphere vMotion

vSphere vMotion always uses encryption when migrating encrypted VMs.

For VMs that are not encrypted, you can edit the VM's settings to use one of the following states:

- Disabled: Do not use encrypted vSphere vMotion.
- Opportunistic (default): Use encrypted vSphere vMotion if the source and destination hosts support it.
- Required: Use only encrypted vSphere vMotion.

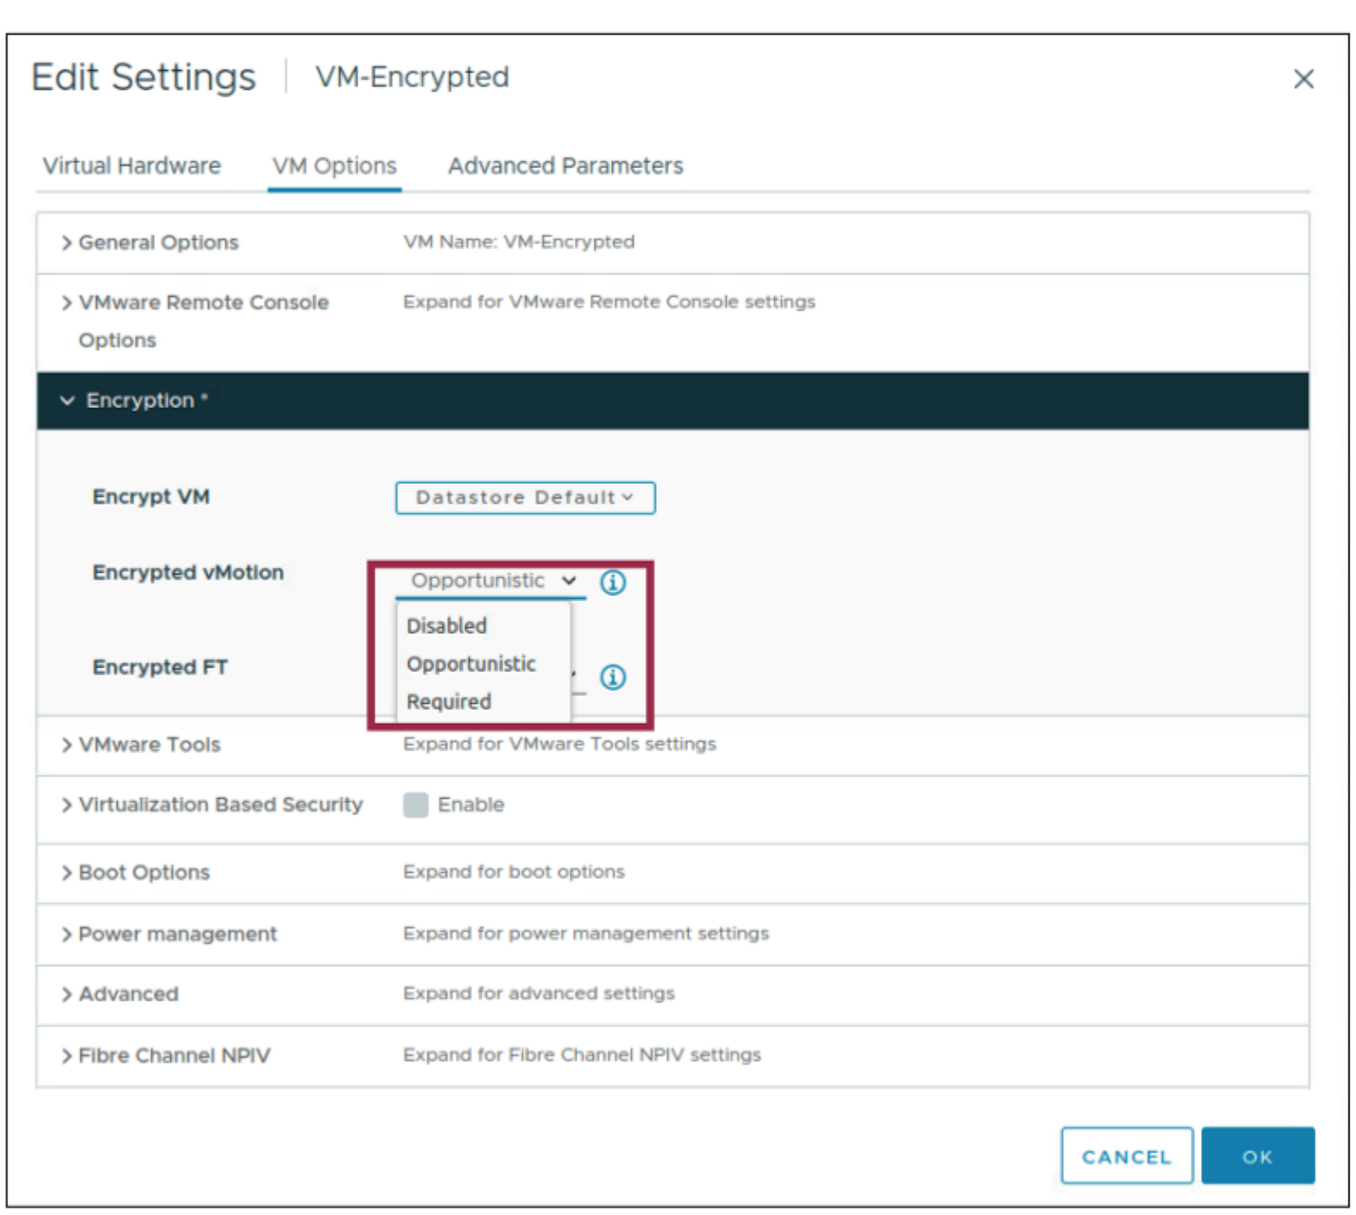

9-31 Review of Learner Objectives

- Identify the role of vTPM and VBS for securing VMs
- With Intel Software Guard Extensions (SGX), software programs can create private memory regions called enclaves
- Secure Encrypted Virtualization with Encrypted State (SEV-ES) is a hardware feature that is enabled in AMD CPUs
- Enable encrypted vSphere vMotion

9-32 Lesson 3: VM Encryption

9-33 Learner Objectives

- Describe the components of a VM encryption architecture
- Manage encrypted VMs
- Use encryption with instant clones
- Migrate Encrypted VMs
- List VM encryption events and alarms

Technet24

9-34 About VM Encryption

VM encryption provides several functions.

Function	Description
Encryption	• Protection of VM disks and metadata files, such as `.nvram` and `.vswp` • Multilayer key protection

Orchestration	• Simplified deployment using storage policies
	• Agnostic:
	  • Storage-agnostic
	  • Guest operating system-agnostic

Key control	• Key management provided by key servers and vSphere Native Key Provider • Use of Key Management Interoperability Protocol (KMIP) standard • Non-persistence of keys for added security
Access control	• Role for administrators without cryptography permissions • Cryptographic tasks authorized only to administrators with appropriate permissions

VM encryption protects not only the data on disk but also the swap file and any guest-specific information that might be contained in the `.vmx` or `.nvram` files. VM encryption also provides a multilayer key protection mechanism that makes breaking the security of the VM almost impossible.

VM encryption makes it simple to configure encryption by using the orchestration capabilities of vCenter storage policies. By using storage policies, you can apply encryption to any type of VM, regardless of the guest operating system or the storage on which it resides.

By using external, enterprise-grade key servers, exposure to risk is limited because vSphere does not need to manage keys or keep key data on disk anywhere in the vSphere environment. Using the industry-standard Key Management Interoperability Protocol (KMIP), vSphere can be integrated with various third-party key management server (KMS) vendors and products.

In vSphere 7.0 Update 2 and later, you can use the built-in vSphere Native Key Provider to enable encryption technologies, such as virtual TPMs (vTPM).

You can limit access to cryptographic operations by enforcing roles and permissions for users. Only users with the required permissions can perform cryptographic tasks such as encrypting or decrypting a VM.

9-35 Overview of VM Encryption

By enabling encryption on VMs, you can protect your confidential data.

VM encryption was introduced in vSphere 6.5.

- VM disks (VMDKs) are encrypted and can only be accessed with a digital key.
- vSphere access control restricts access to the digital key for decryption.
- In-flight data is also encrypted.

  For example, network communication between the encrypted virtual machine and other networked systems, including other virtual machines, physical machines, or external devices.

9-36 VM Encryption Interoperability (1)

vSphere 6.7 and later:

- Cloning is supported, conditionally.
  - — Full clones are supported.
  - — Linked clones are supported and the clone inherits the parent encryption state including keys.
- An encrypted VM can resume from a suspended state.
- An encrypted VM can revert to a memory snapshot.
- Encrypted VMs with a memory snapshot and suspended state can be migrated between ESXi hosts.
- Encrypted VMs must be powered off during most VM encryption operations, with the following exceptions:
  - — A powered-on encrypted VM can be cloned.
  - — Shallow recrypt can be performed with a powered-on encrypted VM.
- Not all backup solutions that use vSphere Storage APIs - Data Protection for virtual disk backup are supported.

Technet24

vSphere VM encryption has limitations in the devices and features that it can interoperate with in vSphere 6.5 and later releases:

- Full clones are supported. The clone inherits the parent encryption state, including keys. You can re-encrypt a full clone to use new keys or decrypt the full clone.
- Linked clones are supported, and the clone inherits the parent encryption state including keys. You cannot decrypt the linked clone or re-encrypt a linked clone with different keys.
- Starting with vSphere 6.7, you can resume from a suspended state of an encrypted VM or revert to a memory snapshot of an encrypted VM. You can migrate an encrypted VM with memory snapshot and suspended state between ESXi hosts.
- You cannot perform the following tasks on an encrypted VM:
  - For most VM encryption operations, the VM must be powered off. You can clone an encrypted VM, and you can perform a shallow recrypt while the VM is powered on.
  - You cannot encrypt a VM that has existing snapshots. Consolidate all existing snapshots before you perform the encryption.
- You can use vSphere VM encryption with IPv6 mode or in mixed mode. You can configure the KMS with IPv6 addresses. Both vCenter and the KMS can be configured with only IPv6 addresses.

For more information about VM encryption interoperability, see https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-security/GUID-C0AF1F3A-67B4-41A6-A933-7E52A3603D9D.html.

9-37 VM Encryption Interoperability (2)

Certain features do not work with vSphere VM encryption:

- Content library.
- vSphere ESXi Dump Collector.
- Using vSphere VM encryption with other VMware products such as Workstation.
- Sending output from an encrypted VM to a serial port or parallel port.
- Certain types of VMDK configurations are not supported.
  - A named virtual disk unassociated with a VM, also called First Class Disk.
  - RDM (Raw Device Mapping).
  - Multi-writer or shared disks (MSCS, WSFC, or Oracle RAC).

You cannot send output from an encrypted VM to a serial port or parallel port. Even if the configuration appears to succeed, output is sent to a file.

If a virtual disk is encrypted, and you select **Multi-writer** on the VM's Edit Settings page, the **OK** button is deactivated.

For more information about VM encryption interoperability, see https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-security/GUID-C0AF1F3A-67B4-41A6-A933-7E52A3603D9D.html.

Technet24

9-38 VM Encryption Enhancements

The following enhancements to VM encryption were introduced in vSphere 7:

- vSphere Fault Tolerance.
- Cross-vCenter migrations of encrypted VMs.
- Fault Tolerance for encrypted VMs without vTPM.
- Cluster-Wide Encryption.
- New APIs improve the management of the VM encryption environment.
- New events and alarms can be used to monitor the VM encryption environment.

The enhancements in vSphere 7 improve the general functionality of VM encryption. They are based on competitive analysis and feedback from VMware customers and partners.

To mitigate the risk of an administrator accessing data without authorization, VMs can be encrypted, and access to the encryption and decryption functionality can be provided to a subset of administrators. You can provide this type of encryption by using the inbuilt cryptographic privileges in vSphere.

9-39 Advantages of VM Encryption

VM encryption has several advantages over similar offerings on the market:

- It does not require in-guest agents.
- Its methodology is uniform across all guest operating systems.
- It protects all VM data, including swap files, not only virtual disk files.
- It enables detailed key control: VMs and disks can use different keys.
- It offers easy orchestration through VM storage policies.
- It is compatible with KMIP, an industry-standard language for the management of security keys. KMIP allows the interoperable exchange of data between different key management servers and clients.

9-40 Business Use Case: Securing VMs

A common requirement in a company's security policy is to protect VM data from administrative users:

- Problem:
  - A large company has several vSphere and storage administrators.
  - The company must protect its confidential data.
  - The company must reduce the risk of someone easily downloading a VMDK file or the entire VM to a removable storage device and leaving the company with this data.
- Solution:
  - With VM encryption, the company can secure confidential data on a VMDK file so that this data is unreadable without the digital key that was used to encrypt the disk.
  - The key is not readable in any file. It is secured by an extra layer of encryption.
  - The company grants only a limited number of people access to the key.

The digital key used to encrypt a VM is secured in an extra layer of encryption, much like taking the key needed to open a safe and placing it in another safe.

You allow regular vSphere administrators to continue with their daily activities without change, but you grant cryptographic privileges to a subset of these administrators.

Access to VM files does not mean that storage administrators or vSphere administrators can walk away with the VM (VMDK) file and read its contents offsite.

Technet24

9-41 VM Encryption Architecture

The VM encryption architecture consists of the following components:

- vCenter:
  - Communicates with a KMS and requests keys on behalf of ESXi hosts
- ESXi hosts:
  - Run encrypted VMs
- One of the following key management servers:
  - Standard key provider
  - Trusted key provider
  - vSphere Native Key Provider

9-42 Supported vSphere KMS Servers

- Standard key provider:
  - An external server that provides keys to services provided by a third-party security vendor:
- Trusted key provider:
  - vSphere Trust Authority (vTA)
- vSphere Native Key Provider
  - Introduced in vSphere 7.0 Update 2, you can use the built-in vSphere Native Key Provider to enable encryption technologies, such as virtual TPMs (vTPM)

Key Provider	External Key Server Required?	Quick Setup?	Works Only with vSphere?
Standard key provider	Yes	No	No
Trusted key provider	Yes	No	No
vSphere Native Key Provider	No	Yes	Yes

9-43 VM Encryption Prerequisites

To encrypt a VM, you must meet several prerequisites:

- A registered KMS that can provide vCenter with keys.
- An encryption storage policy.
  For example, you can use the built-in VM Encryption Policy.

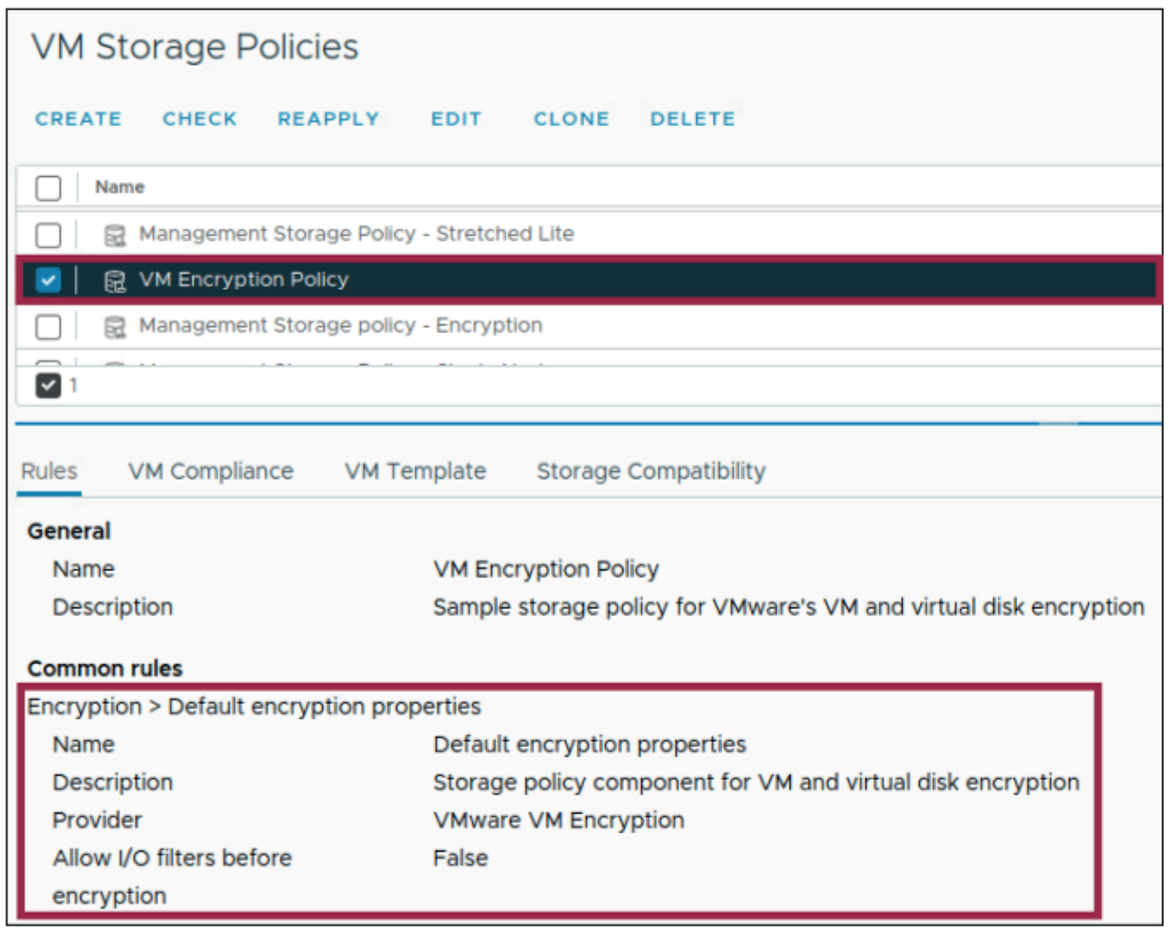

To encrypt VMs, you must create a storage policy with the encryption filter enabled.

You create a common rule in the storage policy that defines the encryption properties. The default encryption properties are appropriate in most cases. You create a custom policy only if you want to combine encryption with other features, such as caching or replication.

You must set the **Allow I/O filters before encryption** value to **True** to enable the use of a replication filter before the encryption filter. The replication filter sees plain text before it is encrypted.

Technet24

9-44 Enabling Encryption During VM Creation

If a new VM has an encryption storage policy applied to it, the VM is encrypted immediately when it is created.

You do not need to take additional steps to encrypt the new VM.

The encryption storage policy can be applied to any VM, regardless of both the operating system type and the storage on which the VM resides.

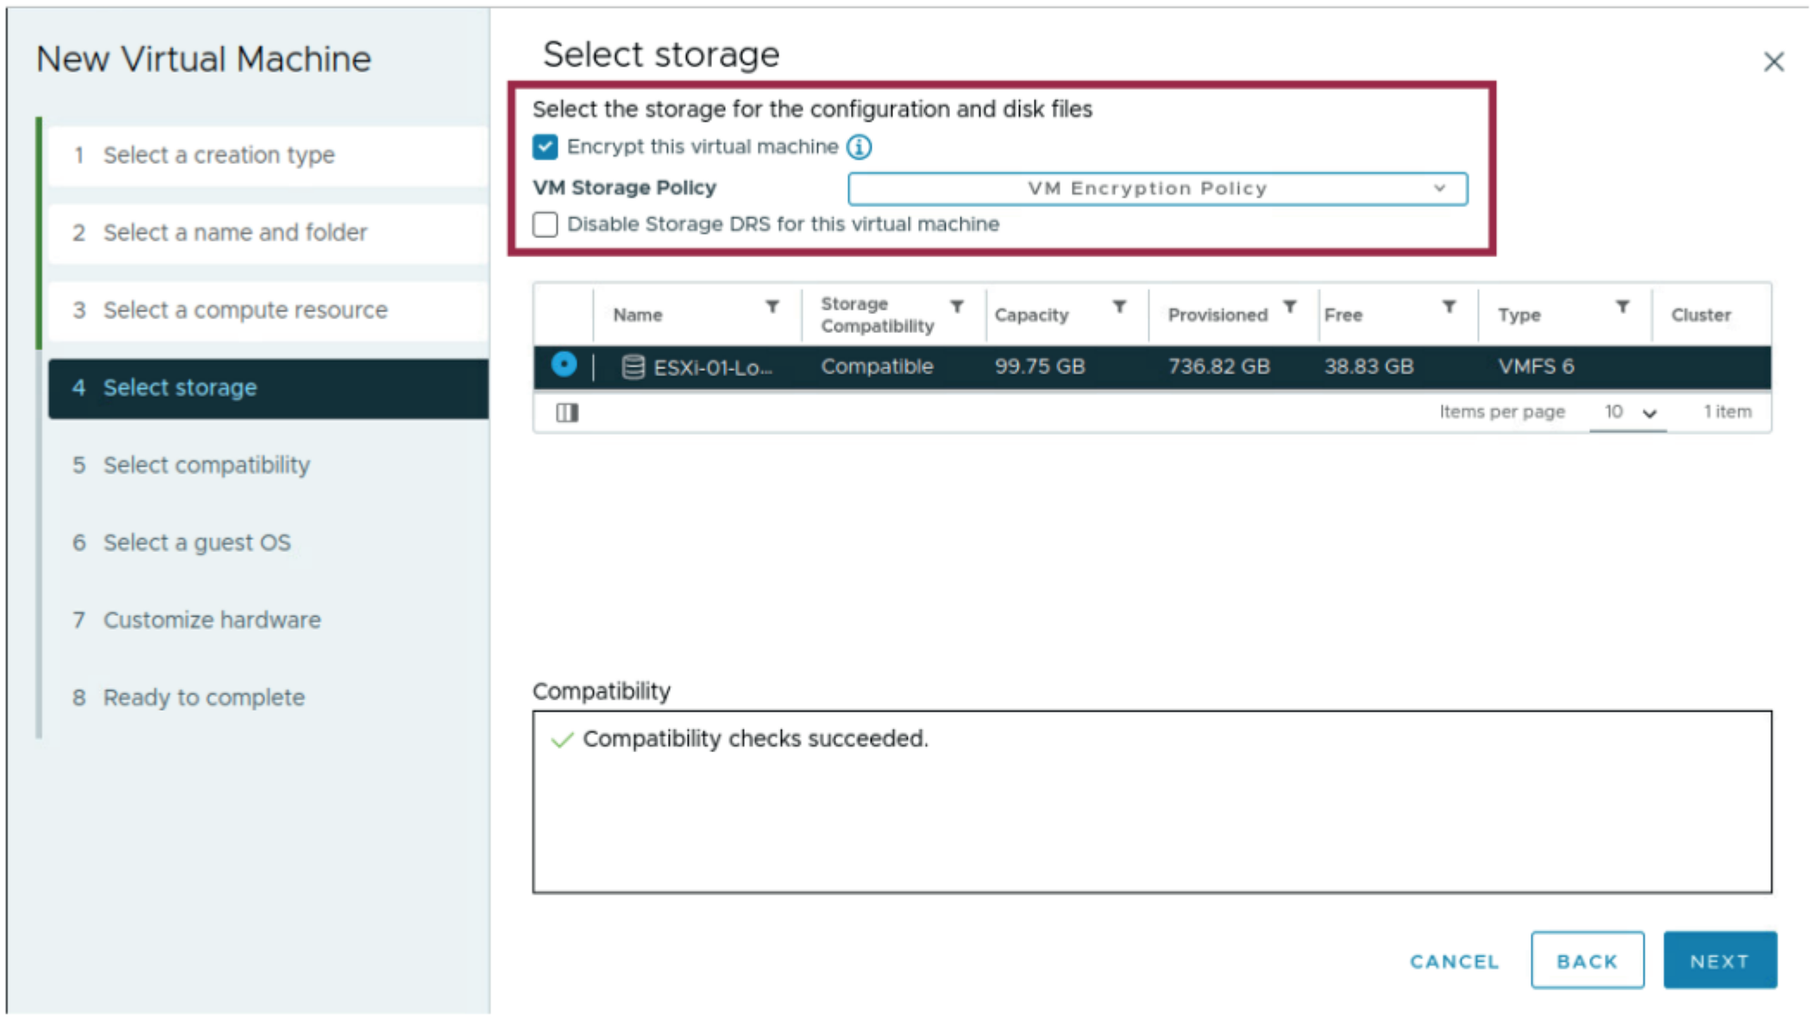

To encrypt a new VM, perform the following steps:

1. In the New Virtual Machine wizard, select the **Encrypt this virtual machine** check box to show only the encryption policies.
2. Select an encryption policy from the **VM Storage Policy** drop-down menu.

9-45 Encrypting Existing VMs

Encrypting an existing VM is slightly different from encrypting a new VM.

When a VM encryption storage policy is applied to an existing VM, vSphere automatically performs the following tasks:

- A duplicate disk is created with the encryption I/O filter applied.
- Data is copied from the unencrypted disk to the new disk, and the data is encrypted.
- After all the data is copied, the new encrypted VMDK file is attached to the VM.
- The old unencrypted disk is deleted.

This operation is performed on one disk at a time. You need at least the capacity of the largest disk available in the datastore to complete this operation.

This operation is allowed only when the VM is powered off.

For the data of an existing VM, a duplicate disk is created with the encryption I/O filter applied. Data is then copied from the unencrypted disk to the new encrypted disk, applying the encryption during the copy process. After all the data is copied, the new encrypted VMDK file is attached to the VM, and the old unencrypted disk is deleted.

The VM must be powered off.

9-46 Unlocking Encrypted VMs (1)

When powered-off, a VM goes into a locked state when its key encryption key (KEK) is unavailable.

A VM's KEK might be unavailable for the following reasons:

- The KEK is missing on the ESXi host.
- The KEK is missing on the KMS cluster.
- The KMS cluster encountered an error.

vCenter triggers an alarm when the VM goes into a locked state.

Technet24

9-47 Unlocking Encrypted VMs (2)

To unlock a VM, first ensure that the required key is available on the KMS cluster.

In the vSphere Client, click **Unlock VM** to unlock the encrypted VM:

- The Unlock Virtual Machine API sends the VM encryption key to the ESXi host.

9-48 About the KMS

To prepare the environment for VM encryption, you must set up the KMS.

The KMS has the following characteristics:

- Must be compatible with KMIP 1.1: KMIP is an industry-standard language for the management of security keys.
- Provides key management service for KMIP clients, such as vCenter
- Can be configured with a KMIP proxy server
- Is accessed over IP using IPv4 only, IPv6 only, or mixed mode (IPv6 and IPv4)

For vSphere, the KMS must be a server that communicates using the KMIP protocol. vCenter implements a KMIP client that can issue KMIP-formatted commands to request keys using specific identifiers. A KMIP server can return the key values to vCenter using a secure channel.

Having a KMS means that all the management tasks relating to keys can be centrally managed on a single server or server cluster. Multiple vCenter instances can access the same KMS cluster and can be granted access to the same keys. Having a centrally managed key store reduces the attack surface for your key management. Instead of keys existing in multiple locations, which you must secure, a single or limited number of servers are secured, which is easier and safer.

KMIP proxy servers can provide the key management capability of the KMS to a remote location. With KMIP proxy servers, you can keep your critical key server safe in your central data center while still allowing remote offices to use its services without compromising security.

Finally, using the KMIP requires only an IP-capable network, which is readily available in most data centers.

KMS is only needed if you're not using the NKP.

9-49 KMIP Client Certificate

The type of certificate used by KMIP client (vCenter) depends on the KMS server

- Some vendors accept self-generated certificates.
- Some vendors require that the KMS generate a trusted certificate for the client, which must be downloaded by the client.
- Other vendors accept only a client certificate provided by the KMS itself. Check with the KMS vendor for requirements

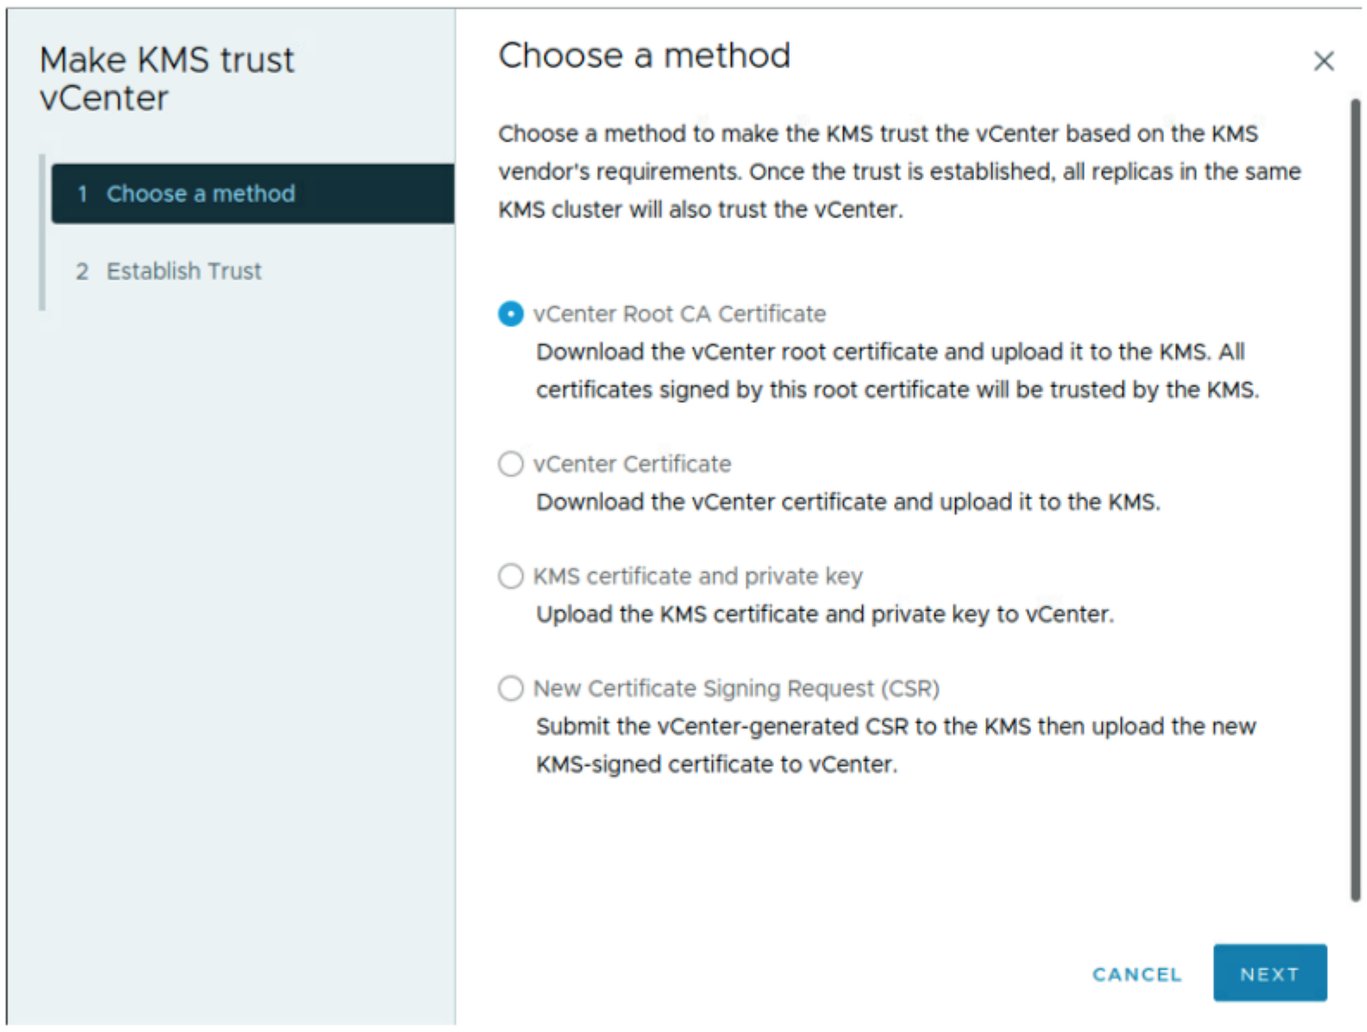

Different KMS vendors have varying requirements for the format of the client certificate that is used by vCenter.

Some vendors allow the use of the client's own certificate, such as a self-signed or CA-signed certificate. Other vendors require that the certificate used by the client be signed by the KMS server.

You can use the vSphere Client to provide certificate details based on the security requirements of the different types of implementations.

Technet24

9-50 Key Management Clusters

The KMS has a critical role in protecting VMs.

Make the KMS highly available by creating a key management cluster:

- A key management cluster is a group of KMIP-based key management servers that replicate keys to one another.
- A cluster must contain key management servers from one vendor.
- A cluster prevents single points of failure.
- One default cluster is allowed for each vCenter instance.
- The cluster is identified in vCenter with a Cluster Name value.

To protect against the potential failure of a KMS (also called a KMIP server), KMIP clusters are formed. These clusters provide some failover protection. vCenter keeps a list of key management servers in the cluster. If the first KMS on the list is not available in the cluster, vCenter tries to communicate with the next KMS on the list, and so on.

To form a KMIP cluster, the KMIP servers must be from a single vendor. See the vendor documentation for steps on how to create a KMIP cluster.

After the keys are replicated, a client can request the key from any of the servers participating in the cluster. So if one server goes down, another server can provide the key.

Although multiple KMIP clusters can be added to vCenter, one cluster must be identified as the default cluster. The default cluster is the cluster from which vCenter requests new keys. Keys can be retrieved from nondefault clusters if the cluster is identified. Clusters of KMIP servers are identified using a user-defined friendly name, populated in the vSphere Client.

9-51 Summary of Keys

The KMS provides keys to secure ESXi hosts and VMs.

Key Type	Description	Key Location	Created by
Host key	Used to encrypt core dumps	Stored in the ESXi host's memory and on the KMS	KMS
Data encryption key (DEK)	Used to encrypt the VM	Stored in the VM's configuration file (in an encrypted format)	ESXi host on which the VM was encrypted
Key encryption key (KEK)	Used to encrypt the DEK	Stored in memory of each ESXi host in the cluster and on the KMS	KMS

Technet24

9-52 Role of vCenter in VM Encryption

vCenter has several important functions in VM encryption:

- Stores KMS credential information
- Manages keys used for encryption:
  - Identifies keys by UUID
  - Retrieves keys from the KMS
- Pushes keys to ESXi hosts when required
- Manages permissions
- Implements and manages storage policies
- Records events for auditing purposes

vCenter itself cannot be encrypted.

vCenter stores the KMS login information, and it is the only vSphere component that can communicate with the KMS and retrieve keys. An ESXi host itself cannot request keys from the KMS. vCenter performs key management tasks on behalf of the ESXi hosts.

vCenter provides a mechanism whereby keys can be identified by a unique identifier, which can be stored in a VM's `VMX` file or `VMDK` file. vCenter can use these identifiers to request keys from the KMS server when required and then push these keys to the ESXi hosts.

vCenter provides the framework for managing permissions for cryptographic tasks. You can restrict access to critical operations to a subset of your administrators, while ensuring that your other administrators can continue with their daily tasks.

vCenter manages storage policies that define whether a VM is encrypted.

vCenter records events for auditing purposes. Event information includes the identity of the user who initiated the event.

9-53 KMS and vCenter Communication

To secure communication between the KMIP server (KMS) and the KMIP client (vCenter), trust must be established.

vCenter stores the private key information in the VMware Endpoint Certificate Store (VECS).

vCenter uses TLS 1.2 when communicating with the KMS.

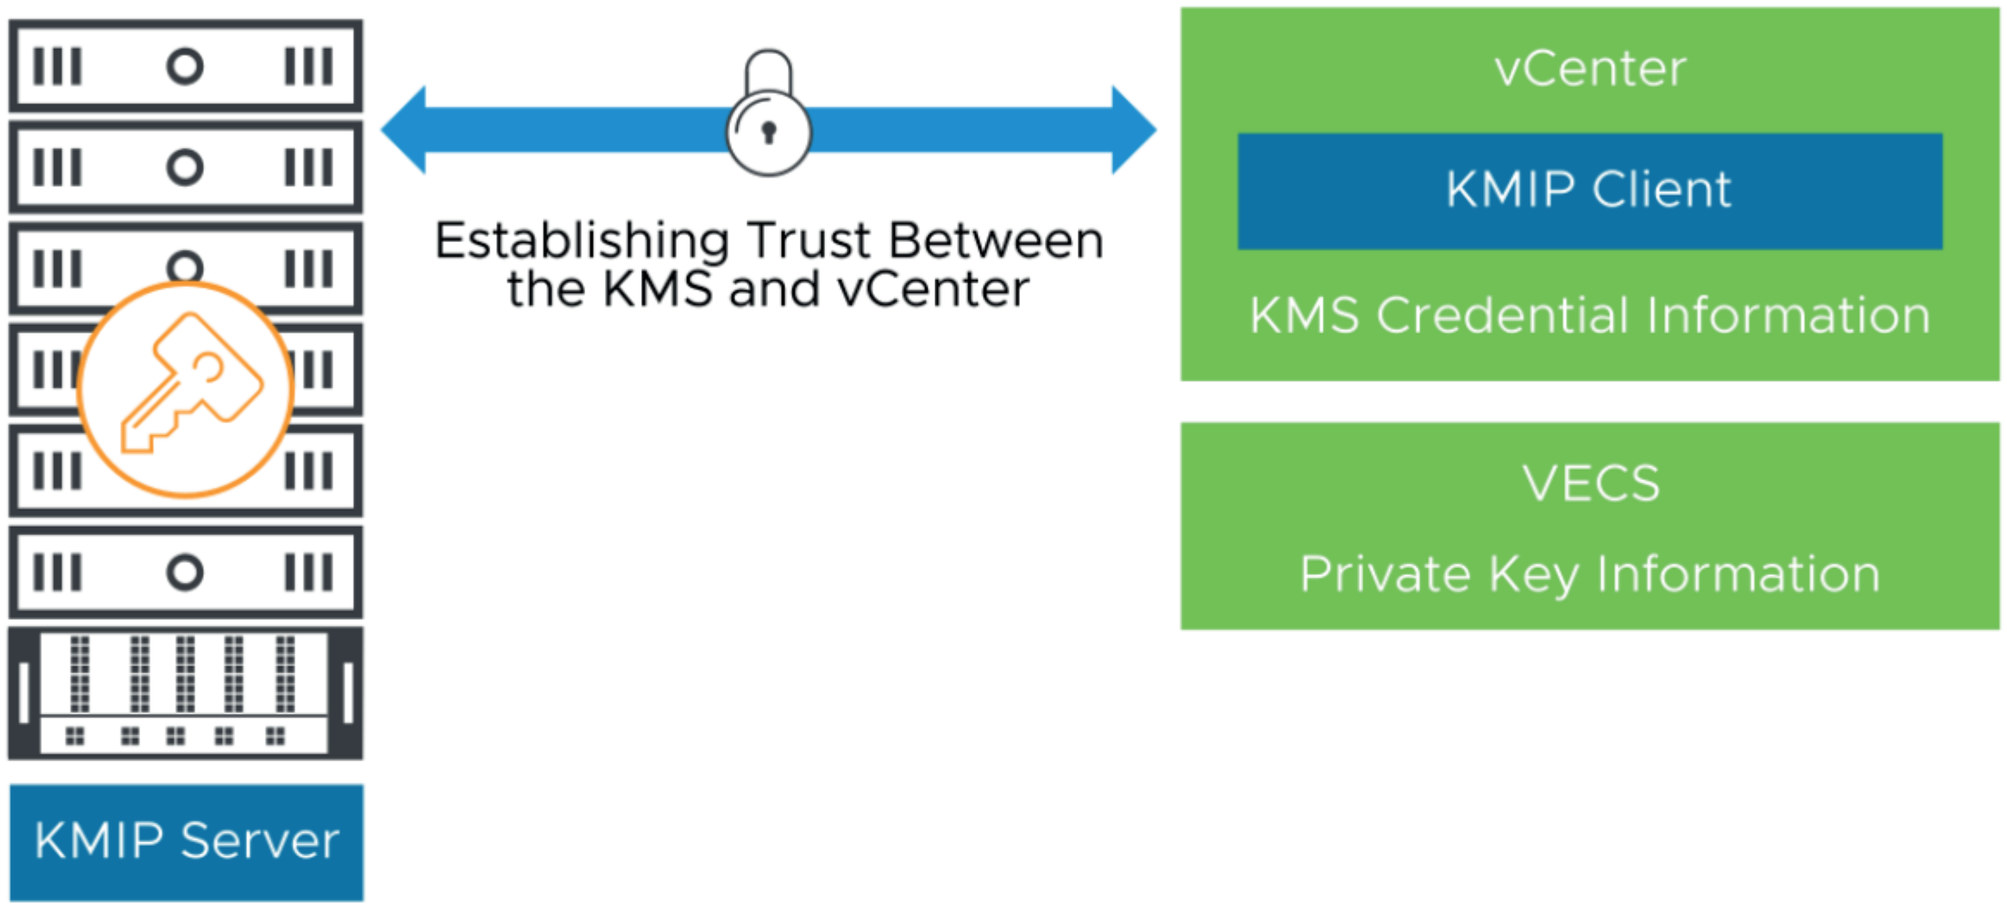

After trust is established between the KMS (KMIP server) and the vCenter (KMIP client), vCenter stores the KMS credential information. The VMware Endpoint Certificate Store (VECS) stores the SSL private key information.

Technet24

9-54 Ensuring the ESXi Host is Cryptographically Safe

An ESXi host ensures that it is cryptographically safe before it can manage encrypted VMs.

When a user performs an encryption task, eg creates an encrypted VM, the following tasks are performed automatically:

1 vCenter requests a key from the KMS for the cluster, known as the host key.

2 vCenter retrieves the host keys from the KMS and pushes a host key to each host in the cluster.

3 The host installs its host key in its memory key cache and enables core dump encryption. The host key is used to encrypt core dumps.

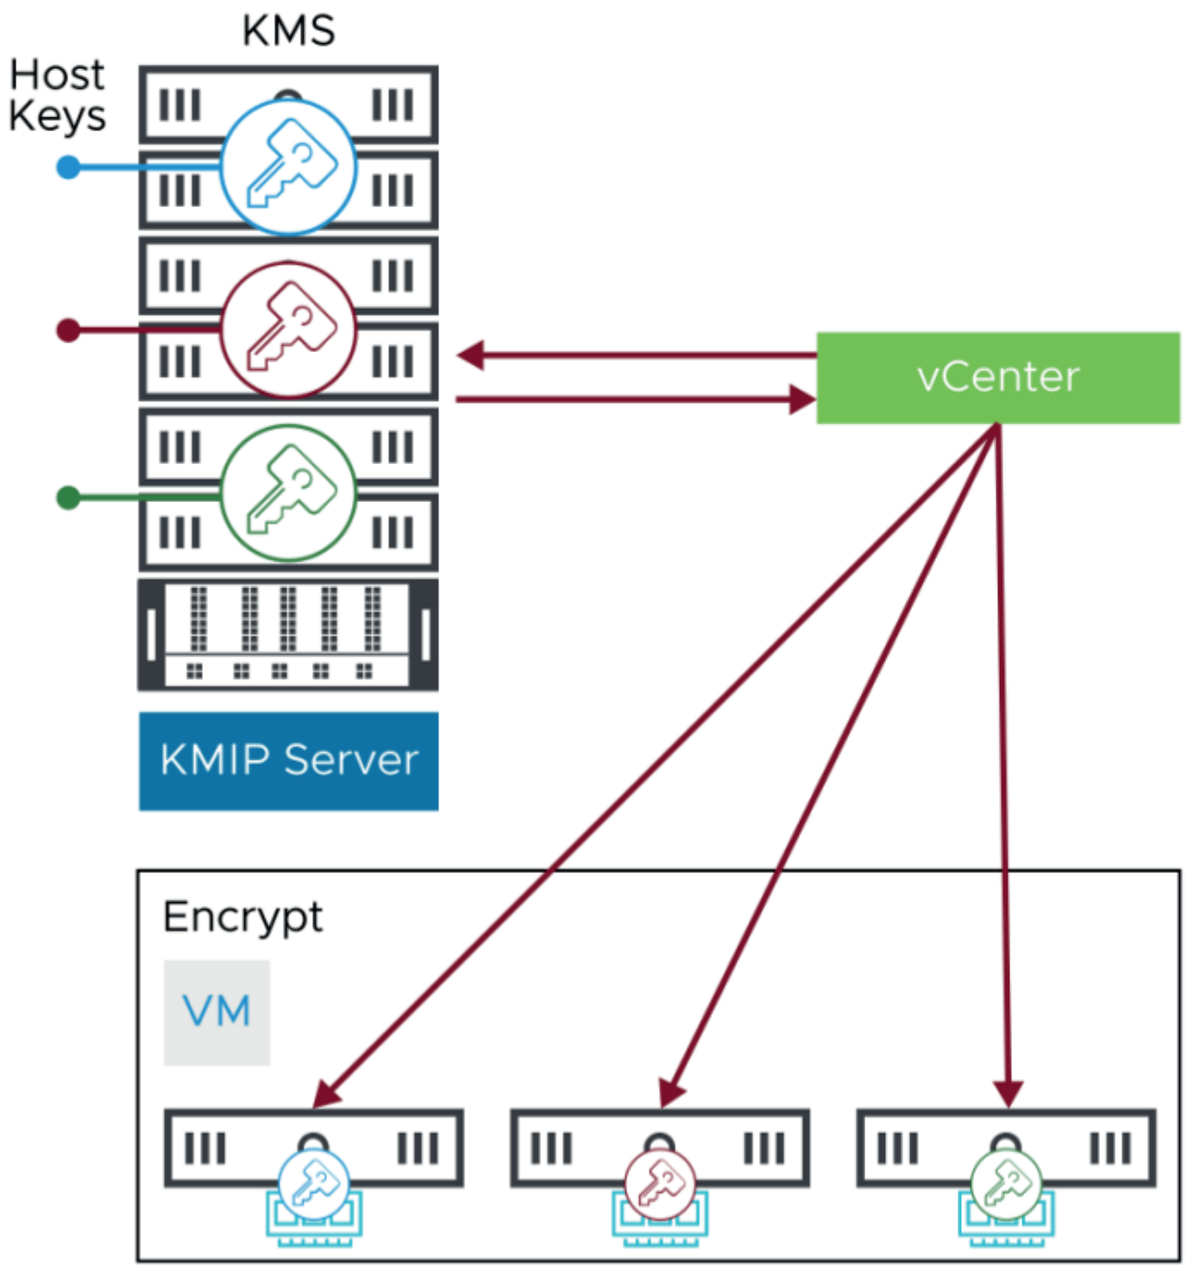

Users with the required privileges can perform cryptographic operations. These operations include creating encrypted VMs, encrypting existing VMs, and decrypting encrypted VMs.

When host encryption mode is enabled on one host in a cluster, it is also enabled for all other hosts in the cluster. Host encryption mode cannot be deactivated while hosts in the cluster are managing encrypted VMs.

9-55 VM Encryption Process Flow (1)

After the host is cryptographically safe, it can receive keys for the VM:

1. vCenter requests a VM KEK from the KMS.
2. The KMS sends the KEK ID to vCenter.
3. vCenter passes the KEK directly to the ESXi host.

   If the host belongs to a cluster, vCenter sends the KEK to each host in the cluster. The KEK is stored in the host's memory.

   The KEK itself is never stored on the vCenter system, only the key ID is.

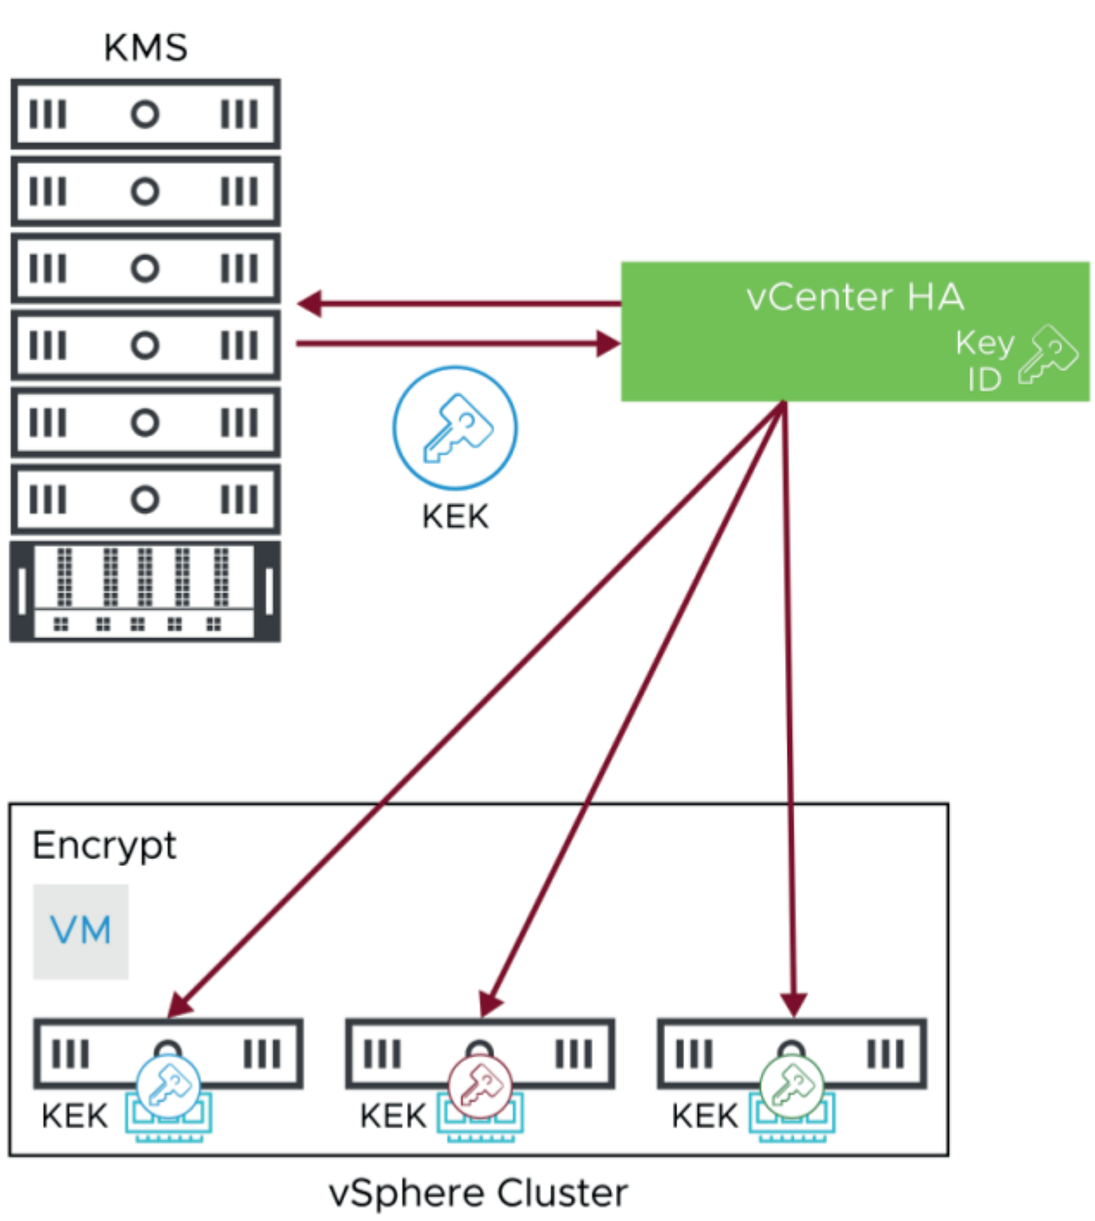

When a cryptographic operation occurs, vCenter must determine whether a VM key must be pushed to the ESXi hosts.

vCenter retrieves the key encryption key (KEK) from the KMS, but the KEK is never stored anywhere on vCenter, not even in its memory. The KEK remains in vCenter memory only while the key is in transit to the ESXi host. vCenter passes the KEK to the ESXi host that is performing the cryptographic operation.

If the host belongs to a cluster, the cluster is deemed the security boundary. In this case, vCenter gives the KEK to all hosts in the cluster so that any cluster operations, such as vSphere vMotion migrations and vSphere HA operations, can occur.

Technet24

9-56 VM Encryption Process Flow (2)

KEKs remain in ESXi memory until the next reboot. Additionally, the core dumps are encrypted.

4. The ESXi host generates a key called the data encryption key (DEK).
5. The DEK is used to encrypt the VM and its disks.
6. The ESXi host uses the KEK to encrypt the DEK.
7. The DEK is stored in the VM's configuration file in an encrypted format.

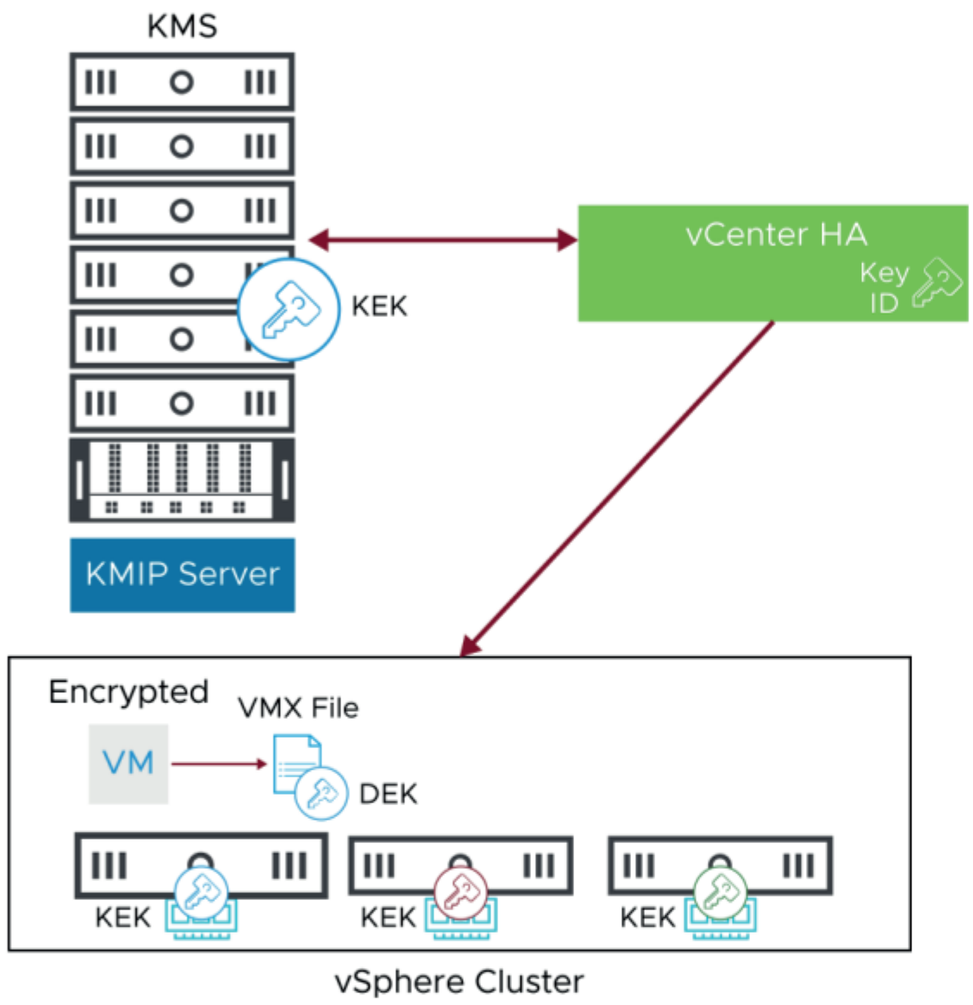

If an ESXi host reboots, a VM's KEK is no longer in the host's memory. vCenter requests the KEK with the corresponding key ID from the KMS and makes it available to the ESXi host when the host is running. The ESXi host can then decrypt the VM's DEK, as needed.

Because KEKs are stored in the host's memory and vSphere encrypts all core dumps (using the HEK) so that the KEK cannot be discovered

9-57 Activity: VM Keys (1)

Where does an ESXi host pull a VM's Key Encryption Key (KEK) from when required?

◊ ❒ Virtual Machine Disk (VMDK)

◊ ❒ Virtual Machine Configuration File (VMX)

◊ ❒ Key Management Server (KMS)

◊ ❒ Hypervisor Security Module (HSM)

9-58 Activity: VM Keys (2)

Where does an ESXi host pull a VM's Key Encryption Key (KEK) from when required?

◊ ❒ Virtual Machine Disk (VMDK)

◊ ❒ Virtual Machine Configuration File (VMX)

◊ ✓ Key Management Server (KMS) Administrator

◊ ❒ Hypervisor Security Module (HSM)

When an ESXi host needs to retrieve a virtual machine's Key Encryption Key (KEK), it obtains it from the Key Management Server (KMS). ESXi hosts cannot pull a VM's key from either the KMS or vCenter.

9-59 Managing VM Encryption

vCenter cryptographic privileges:

- The vCenter Administrator role has cryptographic privileges.
- However, not all administrators should be able to control encryption operations or access keys.
- vCenter provides a role called No Cryptography Administrator, for administrators who do not need encryption privileges.
- Administration and No Cryptography Administrator roles are built-in roles and can not be edited

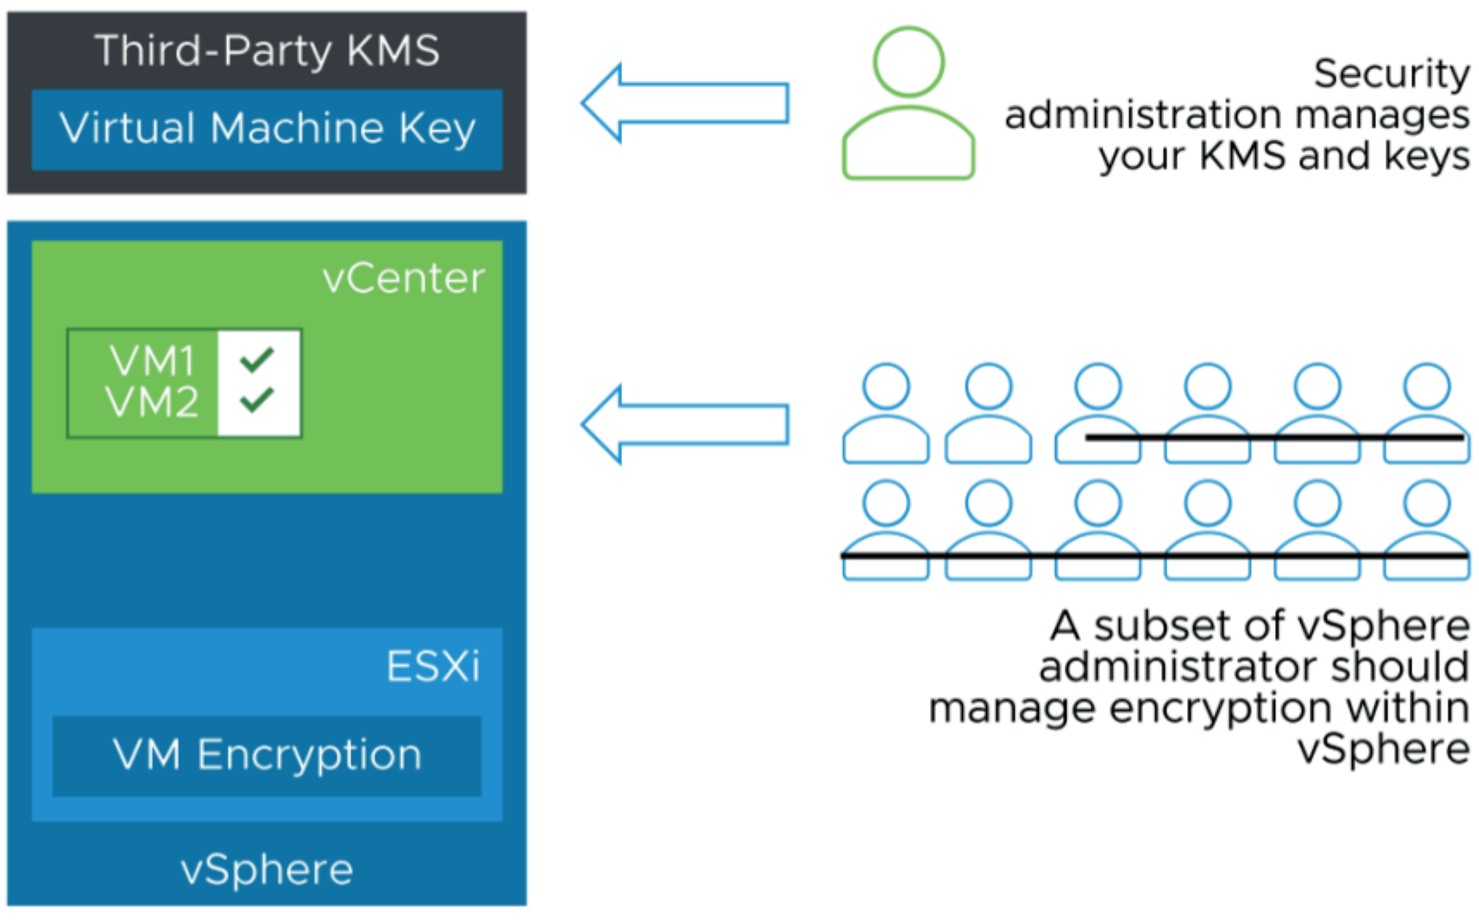

Technet24

9-60 vCenter Role: No Cryptography Administrator

This role does not include the following privileges:

- Cryptographic Operations
- Global Diagnostics
- Add host to cluster
- Add standalone host
- Manage user groups

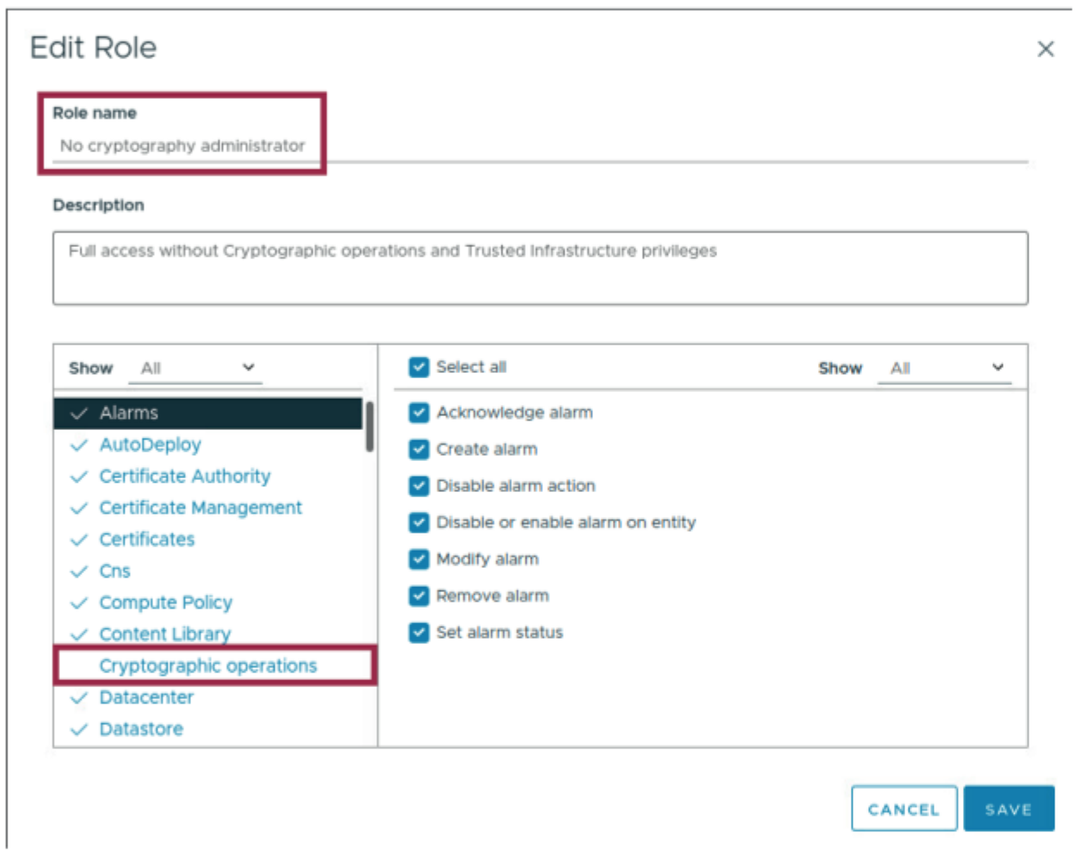

A predefined role called No Cryptography Administrator is provided. You can also create custom roles that include cryptographic privileges.

The No Cryptography Administrator role does not include the 'Add Host' privilege because that task might require the installation of a host key, making the task a cryptographic operation.

9-61 Activity: Cryptographic Privileges (1)

Which predefined vCenter role allows a user to perform cryptographic operations?

◊ ❒ Administrator

◊ ❒ Virtual Machine Power User

◊ ❒ Virtual Machine Console User

◊ ❒ No Cryptography Administrator

9-62 Activity: Cryptographic Privileges (2)

Which predefined vCenter role allows a user to perform cryptographic operations?

◊ ✓ Administrator

◊ ❒ Virtual Machine Power User

◊ ❒ Virtual Machine Console User

◊ ❒ No Cryptography Administrator

You can also create a custom role that has certain cryptographic privileges, but the only predefined role with cryptographic privileges is the Administrator role.

9-63 Cross vCenter Migration and Cloning of an Encrypted VM

In vSphere 7 and later, you can perform cryptographic operations across vCenter instances:

For cloning:

- When cloning an unencrypted VM, you can encrypt the destination VM.
- When cloning an encrypted VM, you can decrypt the destination VM.
- When cloning an encrypted VM, you can recrypt the destination VM.

For migrations:

- You can migrate an encrypted VM across vCenters when the same KMS is configured on both vCenter instances.
- You cannot decrypt the VM during the migration.
- You cannot recrypt the VM during the migration.

VMs can be powered on or powered off when cloning or migrating across vCenter instances.

During the migration operation, you cannot change keys and, therefore, recrypting the VM during the migration is not supported. Also, you cannot change the storage policy during migration, and, therefore, decrypting the VM during the migration is not supported.

Technet24

9-64 Performing a Shallow Recrypt of VMs with Snapshots

Introduced In vSphere 7, you can perform a shallow rekey of an encrypted VM with snapshots, while the VM is powered on or off.

You can encrypt only VMs with a single snapshot branch (multiple branches are not supported).

In the snapshot branch, one snapshot at a time is recrypted:

- If the recryption operation fails before recrypting all snapshots in the chain, some snapshots are encrypted with the new key and other snapshots use the old key.
- To fix this problem, perform the shallow recrypt again.

You use PowerCLI to perform the shallow rekey.

A deep recryption replaces all keys in a VM with new keys, from top to bottom. This operation requires rewriting all encrypted data in the VM.

A shallow recryption replaces only the top-level managed keys in a VM, namely, the KEK. Only a small amount of encrypted data must be rewritten.

9-65 Cloning an Encrypted VM with Instant Clone Technology

Introduced in vSphere 7, Instant Clone technology supports cloning an encrypted VM (the parent VM).

An instant clone is a child VM that shares virtual disks and memory with the parent VM in an ongoing manner.

By default, the encrypted instant clone VM inherits the same keys as the parent VM.

This cloning process has limitations:

- During the instant clone operation, the encrypted VM cannot be shallow or deep recrypted, or decrypted.
- After you create an instant clone of an encrypted VM, you cannot recrypt (shallow or deep) or decrypt the source VM or the destination VM.

You can create an instant clone of a VM by using PowerCLI or API.

You cannot change the storage policy or change encryption keys during the instant clone operation.

9-66 Cloning a VM: Encrypting or Decrypting the Destination VM

To encrypt or decrypt the destination VM, change the VM's storage policy.

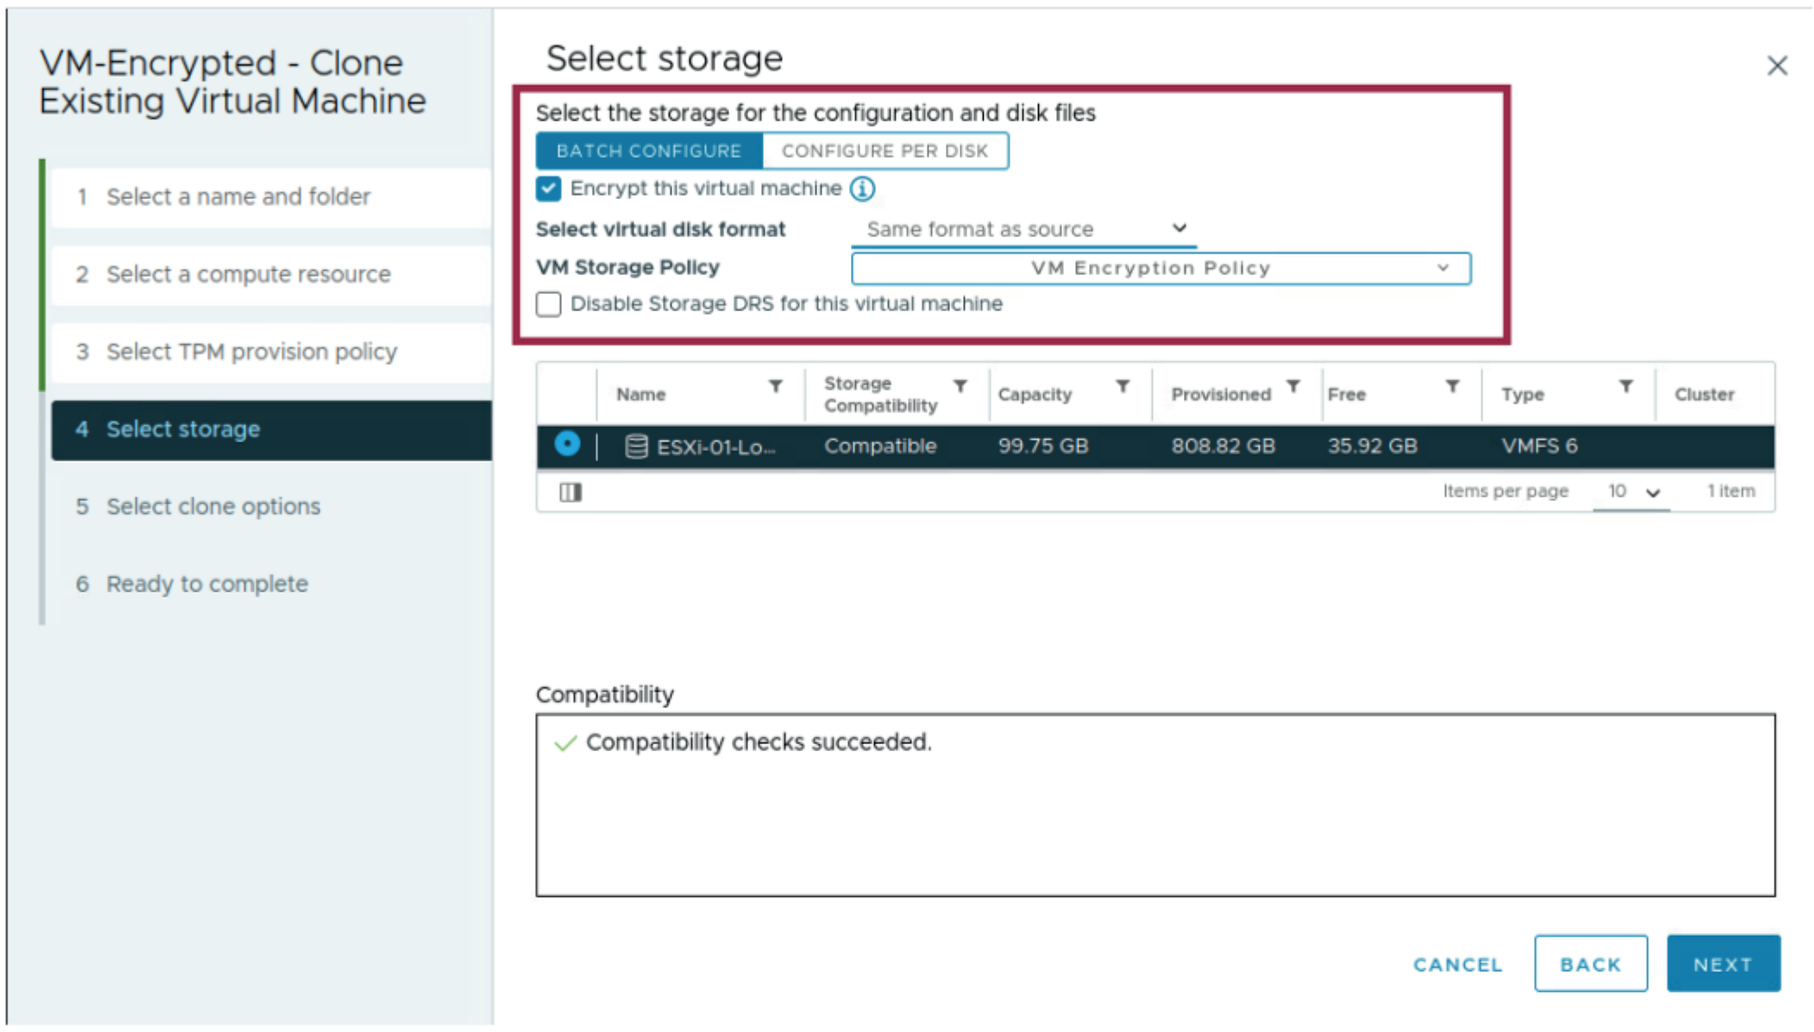

Decrypting a VM can be a security risk. Therefore, as a best practice, restrict most users from having the privilege to decrypt VMs.

9-67 System Requirements When Cloning a VM

For encrypting, decrypting, and recrypting VMs during cloning, the following requirements must be met:

- Establish a trusted connection with the KMS and select a default KMS.
- If cloning occurs across ESXi hosts, both source and destination ESXi hosts must be vSphere 7 or later.

Technet24

9-68 VM Encryption Operations When Cloning a VM

Introduced in vSphere 7, cryptographic operations are supported when you clone a VM. When cloning an unencrypted VM:

- You can encrypt the destination VM by changing the storage policy of the destination VM to the VM encryption policy (or custom encryption policy).
- You can decrypt the destination VM by changing the storage policy of the destination VM to a non-encryption policy.
- You can recrypt the destination VM by using PowerCLI to rekey VMs. You can perform either a shallow or a deep rekey.

A shallow rekey (or recrypt) changes the KEK associated with the VM. The data encryption key (DEK) is wrapped with the new KEK.

Deep rekey (or recrypt) changes the DEK and recrypts all data using the new DEK. This function requires temporary disk space and is slower than a shallow rekey because all data must be rewritten.

PowerCLI cmdlet for recrypt: – – <> – – <>

9-69 Backing Up Encrypted VMs (1)

You can back up encrypted VMs with a backup and recovery solution that uses vSphere Storage APIs - Data Protection (VAPD).

The backup and recovery solution might also provide its own encryption mechanism.

The ability to back up an encrypted VM depends on the transport mode used by the backup solution.

Transport Mode	Description
Hot-add mode	• A proxy VM residing on a host with direct access to the VM's disks being backed up can perform the reads from the virtual disk directly and perform a disk-to-disk copy of the backup data.
NBD/NBDSSL mode	• When no other transport is available, this mode is used. • The ESXi host performs the reads and then sends the data over a LAN, using the NFC protocol.

A transport mode defines the method that is used to transfer data to the backup server.

For more information about transport modes, see *Virtual Disk Development Kit (VDDK)* at https://developer.vmware.com/web/sdk.

Technet24

9-70 Backing Up Encrypted VMs (2)

◊ The backup server must first determine whether a VM is encrypted.

◊ The majority of backup products allow for VMs to be encrypted as they are backed up.

◊ If the VM is encrypted:

- SAN mode is not supported.
- NBD or NBD-SSL is available and can be used.

◊ If hot-add mode is used and the proxy machine is a VM:

- Hot-add mode is available but the proxy machine must also be encrypted.

◊ If hot-add mode is used and the proxy machine is a physical machine:

- Hot-add mode is available but the proxy machine must also be encrypted.

◊ The vCenter user registered as the Backup Appliance user must have cryptographic privileges.

◊ The backup data is stored on the backup server in plain text:

- Grant restore privileges only to trusted individuals.
- Set up a policy to re-encrypt a restored VM.

If the VMs encrypted, SAN mode is immediately ruled out as an available transport mode. The backup products that rely on this mode cannot back up encrypted VMs. The blocks are encrypted, and without vSphere assistance, the backup server cannot use the data.

If the backup proxy is a physical machine, hot-add mode is ruled out, and NBD or NBD-SSL mode is used. You rely on the ESXi host to open the disk, decrypt the data, and send the backup data in plain text to the backup server.

If the user with which you register the backup appliance to vCenter does not have cryptographic privileges, the disk cannot be decrypted to send the data to the backup server.

For any of the transport modes, the backup data is stored on the backup server in plain text. In this way, a VM can be restored as a plain-text machine. Privileges to restore VMs should be granted only to trusted individuals, most likely users with cryptographic privileges, so that the VM can be immediately re-encrypted after it is restored.

9-71 VM Encryption Events and Alarms

In vSphere 7 and later, certain logs warnings and error messages will help you provide information about the running state of vCenter and the VM encryption environment:

- Warning event:
  - Encrypted core dump exists.
- Critical events:
  - Insufficient disk space to perform the cryptographic operation
  - Key creation failed on the KMS cluster.
  - KMS server is unreachable.

These events help you find possible issues.

9-72 Review of Learner Objectives

- Describe the components of a VM encryption architecture
- Manage encrypted VMs
- Use encryption with instant clones
- Migrate Encrypted VMs
- List VM encryption events and alarms

Technet24

9-73 Lesson 4: Key Providers

9-74 Learner Objectives

- Recognize the role of the Key Provider in vCenter

9-75 Overview of Standard Key Provider

A Standard Key Provider can be used when you want vCenter to retrieve keys from a third-party key server.

The vCenter is responsible for distributing these keys to the necessary ESXi hosts within vSphere.

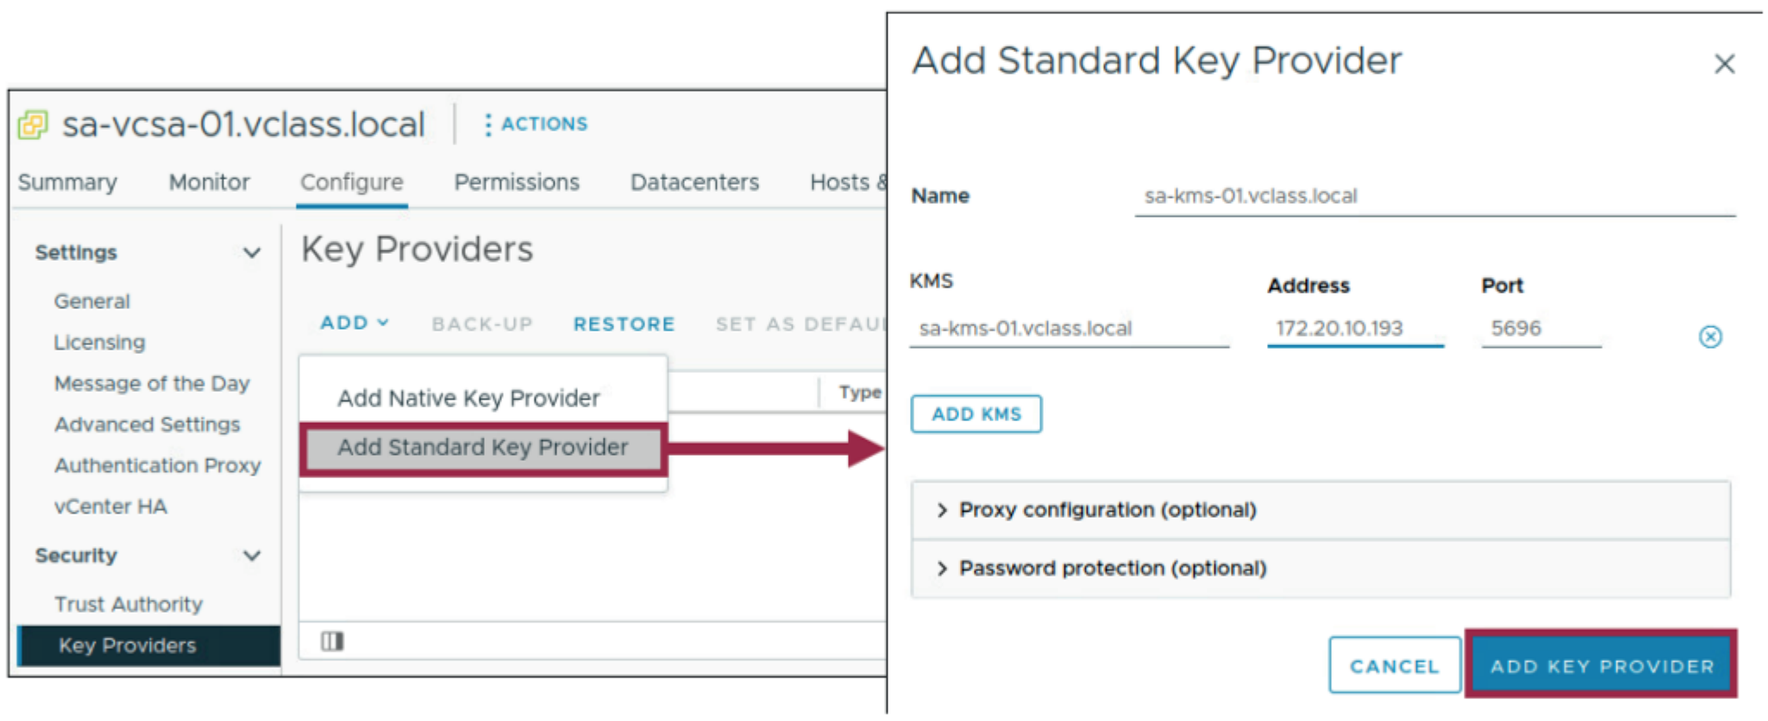

9-76 Available Options

Administrators can choose from the following when configuring a Standard Key Provider:

- Amazon Web Services Key Management Service
- Microsoft Azure Key Vault
- Microsoft Azure Dedicated HSM
- Google Cloud Key Management Service
- Thales Vormetric Data Security Manager
- External Key Management Systems

9-77 Implementation Considerations

Important factors to consider when implementing standard key providers in vSphere:

- Verify compatibility between vSphere versions and the selected standard key provider using the VMware Compatibility Guide - https://www.vmware.com/resources/compatibility/search.php
- Assess the security features and compliance certifications of the chosen key provider.
- Define key rotation and archival policies.

9-78 Overview of vSphere Native Key Provider

In vSphere 7.0 Update 2 and later, you can use the built-in vSphere Native Key Provider to enable encryption technologies, such as virtual TPMs (vTPM) and VM Encryption.

vSphere Native Key Provider only works with VMware infrastructure products and is not KMIP 1.1 compliant.

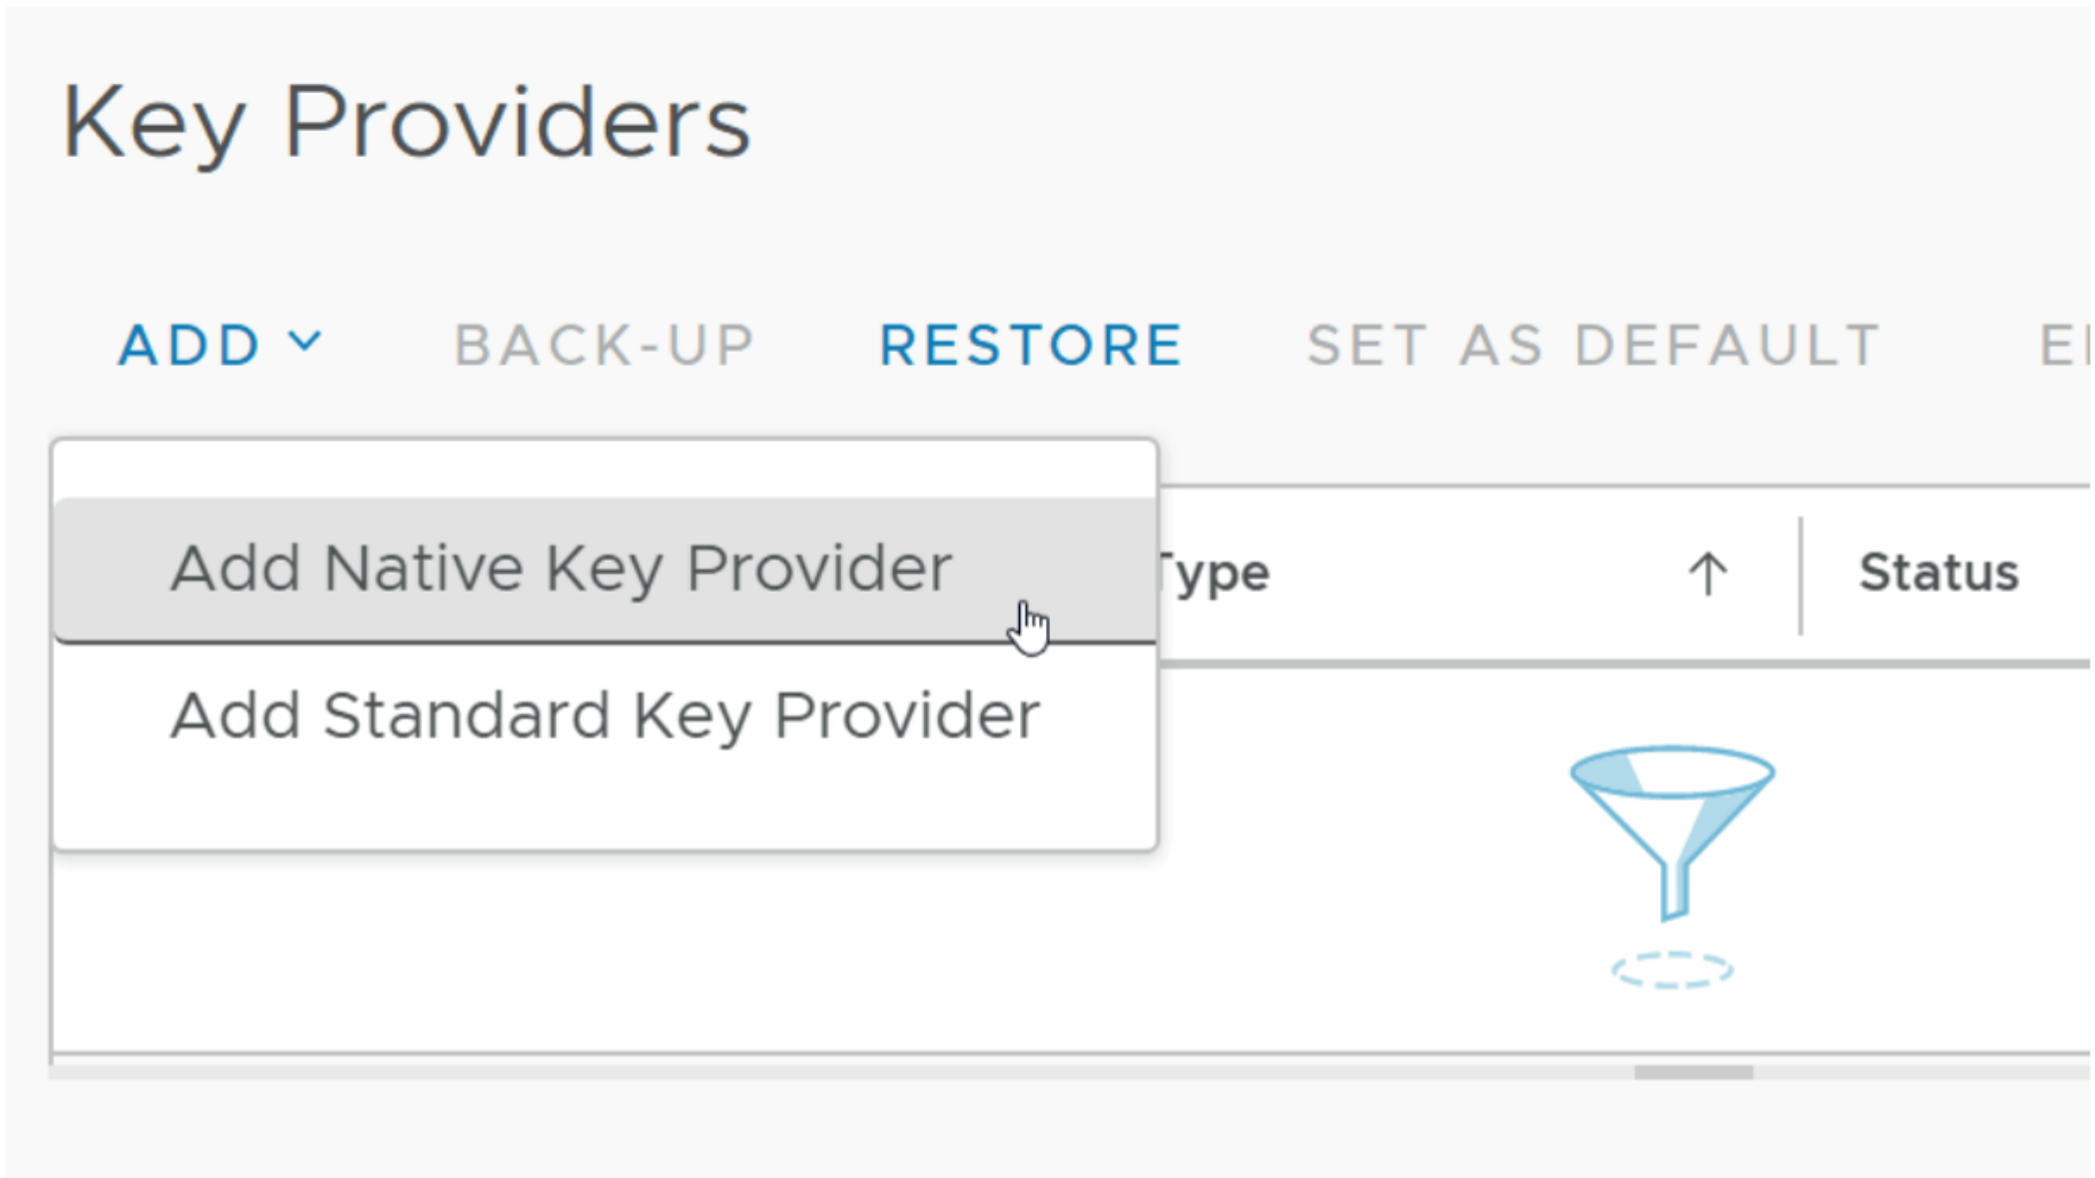

Currently, vSphere Native Key Provider does not support First Class Disk (FCD) encryption.

Technet24

9-79 Features of vSphere Native Key Provider

vSphere Native Key Provider features include the following:

- Enables the use of vTPMs, vSphere Virtual Machine Encryption, and vSAN Data at Rest Encryption, when you do not require or want an external key server
- Does not provide external interoperability, KMIP support, hardware security modules, or other features that a traditional, third-party, external key server can offer.
- Helps address the needs of organizations that either cannot use, or do not want not to use, an external key server.
- Improves data sanitization and system reuse practices by enabling earlier use of encryption technologies on media that is difficult to sanitize, such as flash and SSD.
- Works with multiple vCenters using an Enhanced Linked Mode.
- Works with vCenter High Availability configuration.

9-80 vSphere Native Key Provider Requirements

To use vSphere Native Key Provider, you must:

- Ensure that both the vCenter and ESXi hosts are running vSphere 7.0 Update 2 or later.
- Configure the ESXi hosts in a cluster.
- Configure the vCenter file-based backup and restore, and store the backups securely.

9-81 vSphere Native Key Provider vMotion and Cross-vCenter vMotion

vSphere Native Key Provider supports vMotion and Encrypted vMotion across ESXi hosts.

Cross-vCenter vMotion is supported if vSphere Native Key Provider is configured on the destination host.

9-82 vSphere Native Key Provider and Enhanced Linked Mode

You can configure a single vSphere Native Key Provider that is shareable across vCenter systems configured in an Enhanced Linked Mode configuration. The high-level steps in this scenario are:

1 Creating the vSphere Native Key Provider on one of the vCenter.

2 Backing up the Native Key Provider on the vCenter on which it was created

3 Exporting the Native Key Provider

4 Importing the Native Key Provider to the other vCenters in the Enhanced Link Mode configuration

9-83 Review of Learner Objectives

- Recognize the role of the Key Provider in vCenter

Technet24

9-84 Lesson 5: vSphere Trust Authority

9-85 Learner Objectives

- Describe vSphere Trust Authority
- List the use cases of vSphere Trust Authority
- Explain the general architecture of vSphere Trust Authority

9-86 About vSphere Trust Authority

vSphere Trust Authority secures a vSphere environment and its workloads:

- Uses a hardware root of trust to secure the environment
- Reports on the software running on a host
- Measures or attests software using a Trusted Platform Module (TPM)

A Trust Authority Administrator user must configure the following parameters:

- Permitted versions of ESXi
- Trusted TPM manufacturers or devices
  - This is important to protect against vulnerabilities which are discovered in a particular build of ESXi or a particular TPM implementation.

A trusted key provider is required. A valid attestation report is needed before access is granted to the key provider.

9-87 About vSphere Trust Authority Administrator

To enable vSphere Trust Authority, you must add a user to the vSphere TrustedAdmins group.

- This user becomes the Trust Authority Administrator.
- Use a separate user from the vCenter Administrator as your Trust Authority Administrator.
- You must add a Trust Authority Administrator for both the Trust Authority Cluster and the Trusted Cluster.
- Group membership is validated by ESXi using SAML token authentication.

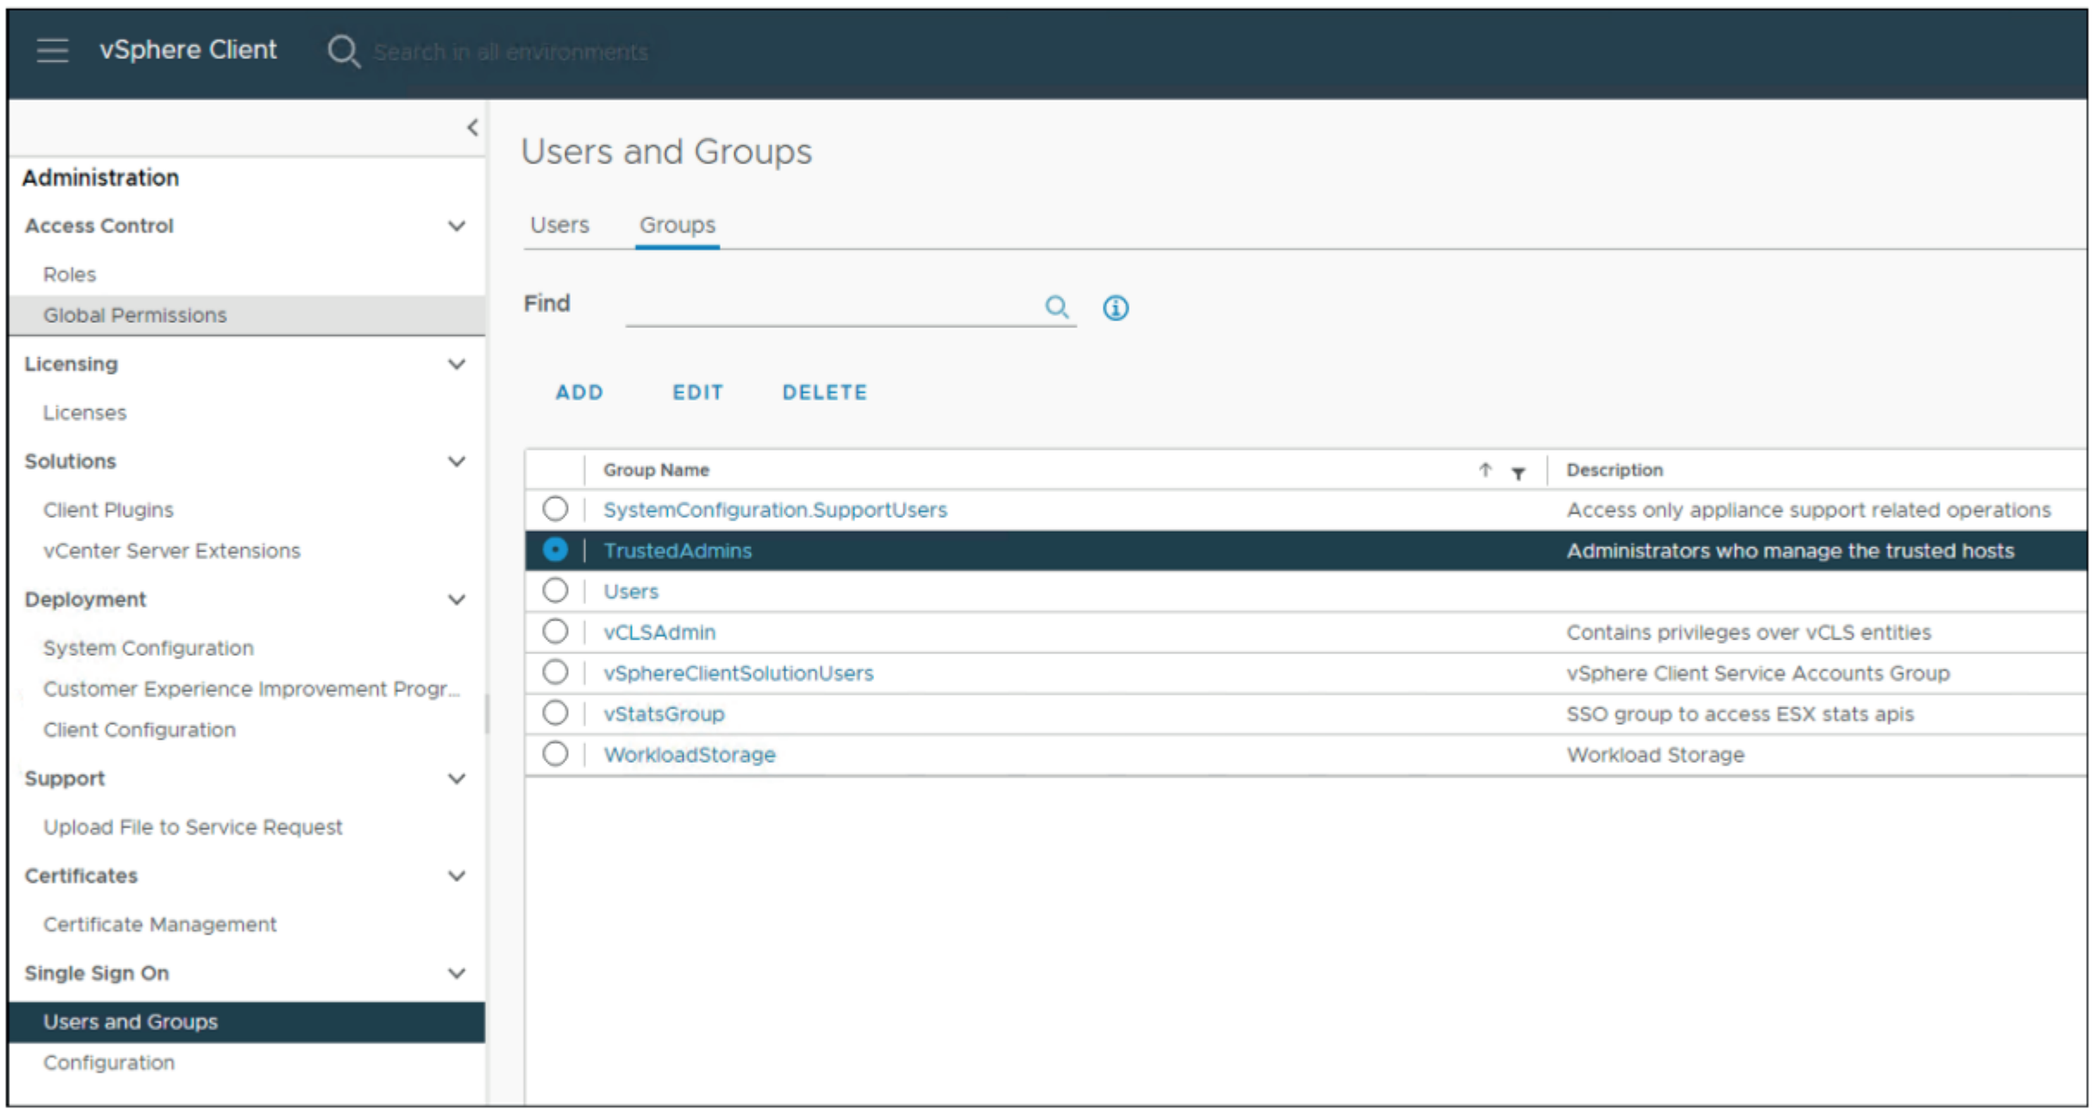

Technet24

9-88 Use Case for vSphere Trust Authority

You can use vSphere Trust Authority to perform security tasks:

- Provide ESXi hosts with a hardware root of trust and remote attestation capabilities.
- Perform cryptographic operations on VMs but with an enhanced level of encryption key management:
  - Key handoff is through the Trust Authority Cluster instead of vCenter.
  - Attestation is required for key access.
- Encrypt vCenter.

In vSphere 7 and later, you can create a trusted computing base, which consists of a secure, manageable set of ESXi hosts. vSphere Trust Authority implements a remote attestation service for the ESXi hosts that you want to trust. Furthermore, vSphere Trust Authority improves on TPM 2.0 attestation support (added to vSphere in the 6.7 release), to implement access restrictions on encryption keys and so better protect VM workload secrets.

In addition, vSphere Trust Authority allows only authorized Trust Authority administrators to configure vSphere Trust Authority services and configure Trust Authority hosts. The Trust Authority Administrator can be the same user as the vSphere administrator user, or a separate user.

9-89 vSphere Trust Authority Components

vSphere Trust Authority comprises a vSphere Trust Authority Cluster and a trusted key provider (KMS). Other ESXi clusters are attested against the Trust Authority Cluster.

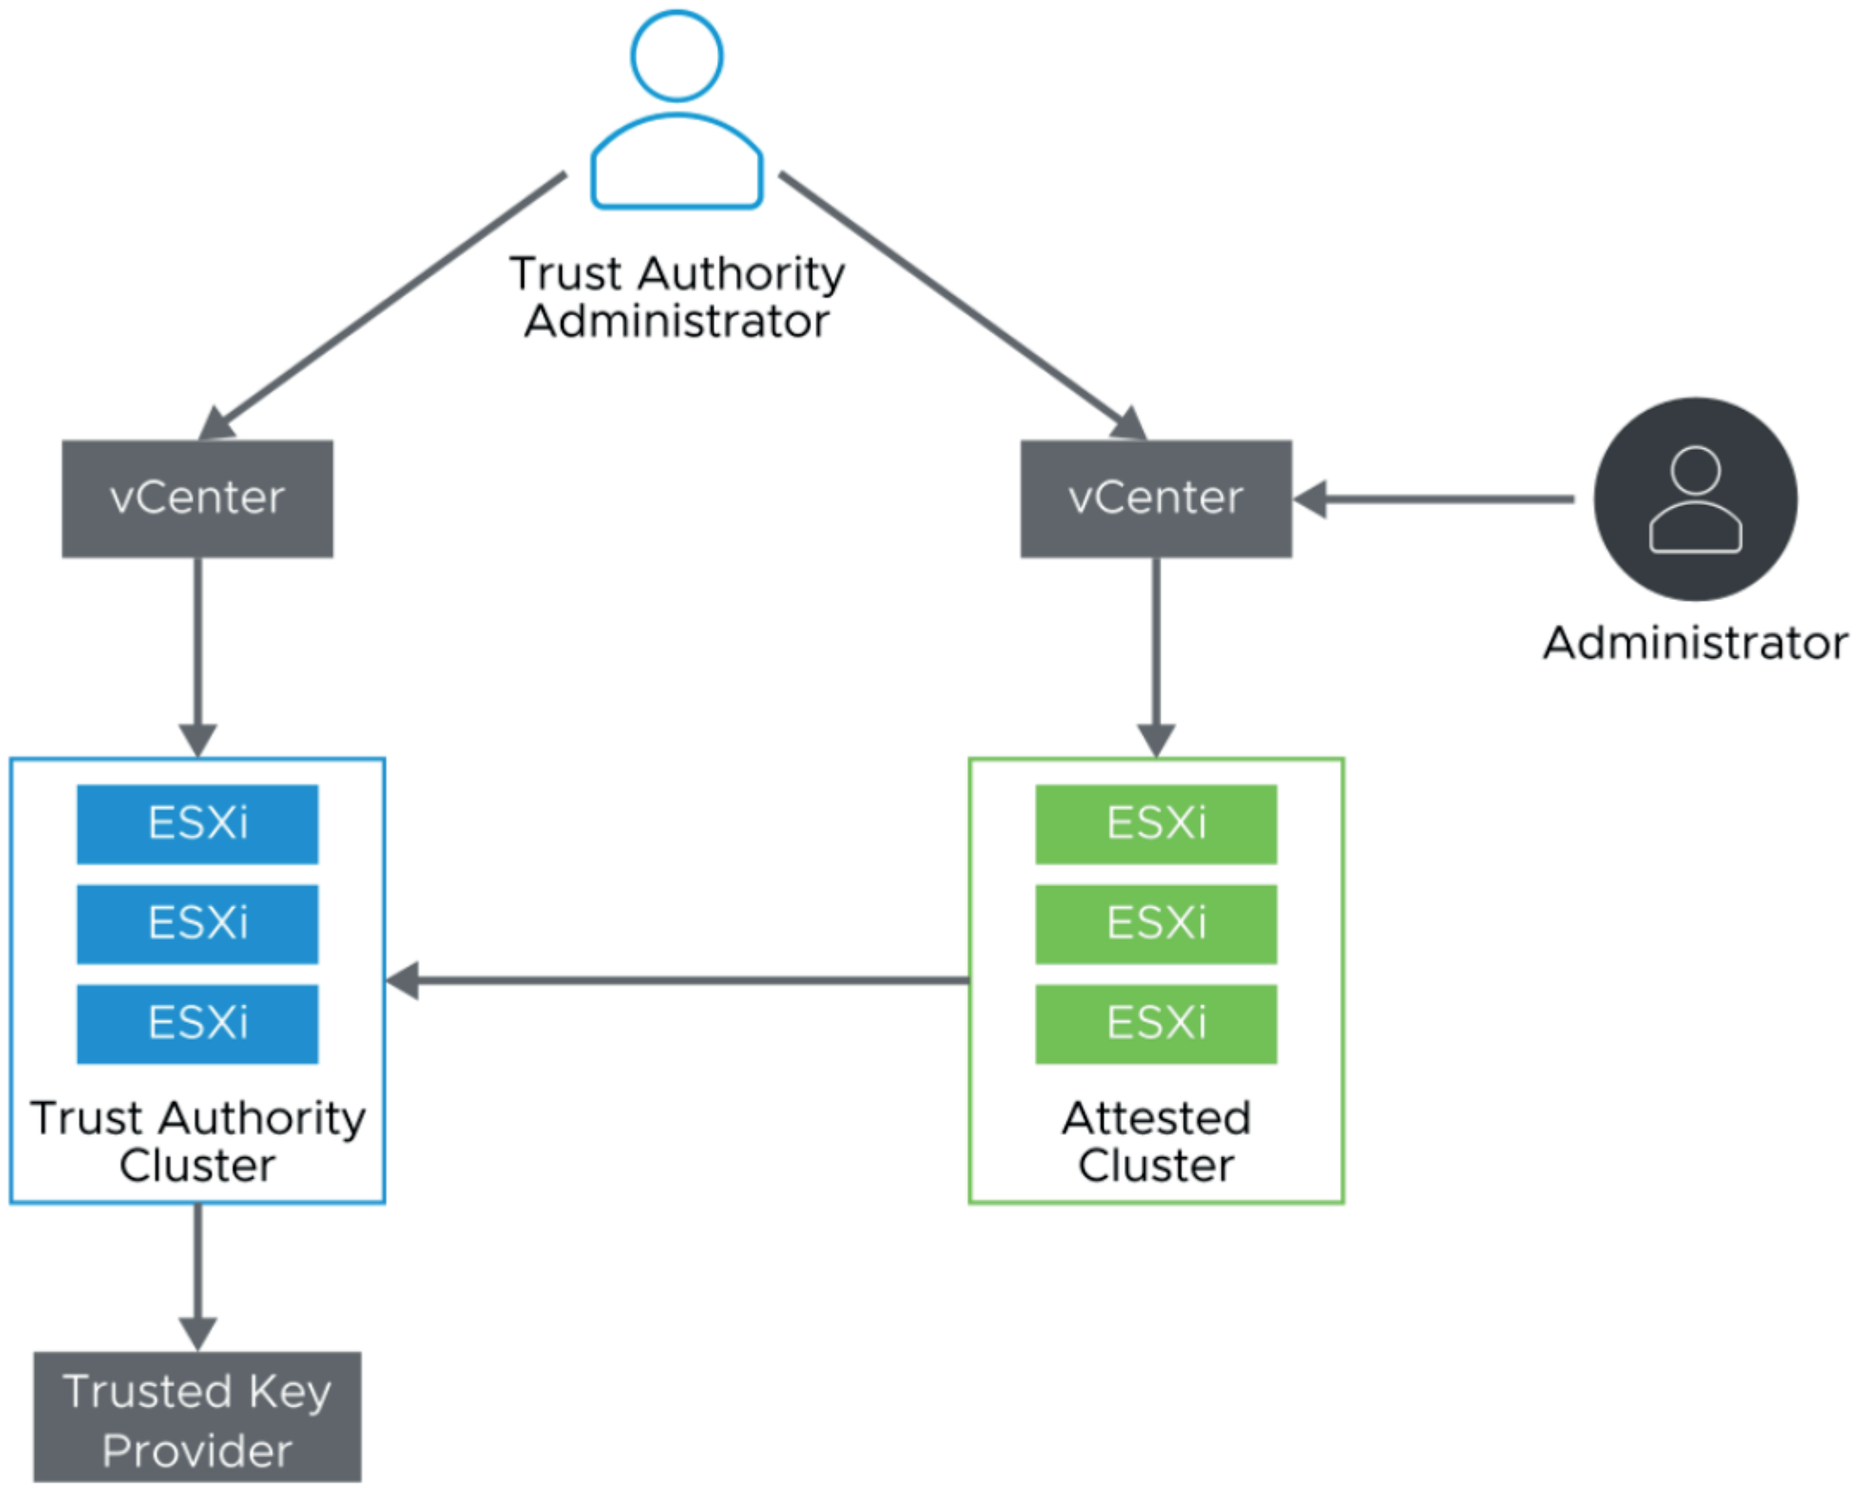

Several concepts apply to vSphere Trust Authority:

- Trust Authority Cluster: A set of restricted ESXi hosts with a known good configuration.
- Administered by a smaller number of trusted administrators.
- Trusted Key Provider (also called key management server [KMS]): A key provider that only the Trust Authority Cluster should know.
- Attested cluster: Also known as a Trusted Cluster. The attestation report for a cluster must be validated by the Trust Authority Cluster.

Technet24

9-90 ESXi Components

Services running on ESXi enabled vSphere Trust Authority:

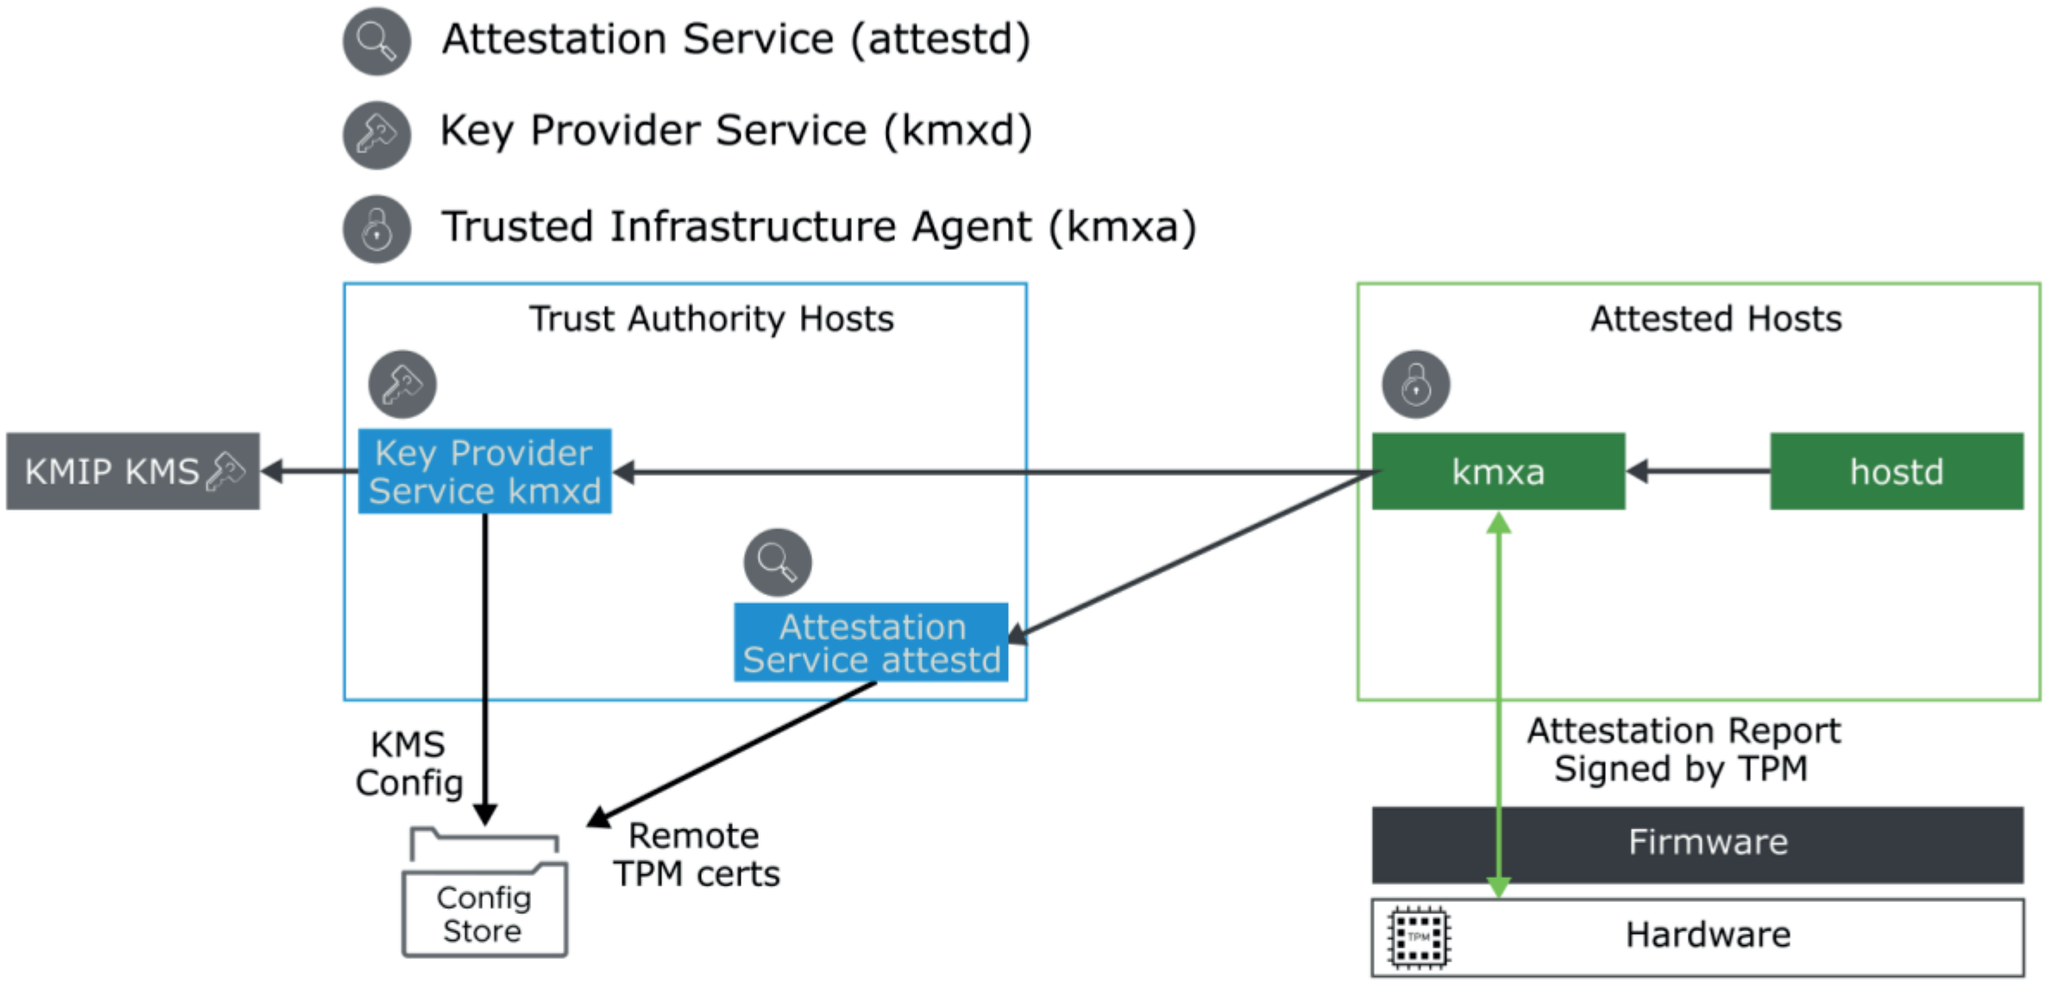

All vSphere Trust Authority services run behind the Reverse Proxy and API Forwarder. The Trusted Infrastructure Agent communicates over port 443 to the Trust Authority Hosts.

The following internal ports are used:

- Attestation Service : Port 7889
- Key Provider Service : Port 7888
- Trusted Infrastructure Agent : Port 7890

9-91 Attestation Service

The Attestation Service on the ESXi host is used to verify remote software and hardware:

- Receives an attestation report from the Trusted Infrastructure Agent of the attesting ESXi hosts
- Verifies that a TPM device is present on the attesting ESXi hosts
- Verifies the ESXi version of the attesting ESXi hosts
- Generates a signed document describing the binary and configuration state of the attesting ESXi hosts

The Attestation Service performs the following functions:

- Generates a signed document that contains assertions describing the binary and configuration state of the remote ESXi hosts in the Trusted Cluster.
- Attests the state of the ESXi hosts using a Trusted Platform Module (TPM) 2.0 chip for software measurement and reporting. The TPM measures the software stack and configuration data and reports it to the Attestation Service.
- Verifies that the software measurement signature can be attributed to a previously configured trusted TPM endorsement key (EK).
- Ensures that the software measurement matches one of a set of previously permitted ESXi images.
- Signs a SAML token that it issues to the ESXi host, providing the assertions about the identity, validity, and configuration of the ESXi host.

Technet24

9-92 Attestation Service Trust

The Attestation Service does not know what to trust.

A Trust Authority Administrator must configure the following items:

- Which versions of ESXi to trust
- Which TPM device manufacturers to trust
- Optionally, which specific TPM devices to trust

9-93 Key Provider Service

The Key Provider Service (`kmxd`) provides a method to encapsulate encryption key sources:

- The service supports KMIP-compliant key servers.
- Additional protocols can be added in the future.
- The service allows encryption key access only if a client presents a valid attestation report.

With the Key Provider Service, vCenter and ESXi hosts do not require direct key management server (KMS) credentials.

In vSphere Trust Authority, an ESXi host can access an encryption key by authenticating with the Key Provider Service rather than with vCenter (as in vSphere 6.7).

For the Key Provider Service to connect to a KMS, the Trust Authority Administrator must configure a trust setup. For most servers that are compliant with the Key Management Interoperability Protocol (KMIP), you configure a trust setup by creating client and server certificates.

9-94 Trusted Infrastructure Agent

Trusted Infrastructure Agent communicates with a Trusted Key Provider Service and Attestation Service :

- The Trusted Infrastructure Agent retrieves encryption keys from the Trusted Key Provider Service.
- The agent can be issued with encryption keys only if the ESXi host passes attestation.
- The ESXi host does not receive encryption keys from vCenter when using a Trusted Key Provider.

Technet24

9-95 vSphere Trust Authority Workflow

In the vSphere Trust Authority workflow, a Trust Authority Administrator performs the following tasks:

1. Exports the identity of the hosts to be attested
2. Enables Trusted Host Services on a Trust Authority Cluster
3. Imports the identity of the attesting hosts into the Trust Authority Cluster
4. Configures a trusted key provider
5. Exports the identity of the Trust Authority Cluster
6. Imports the Trust Authority Cluster identity into the attesting cluster
7. Configures the Trust Authority Cluster as the key provider for the attesting cluster

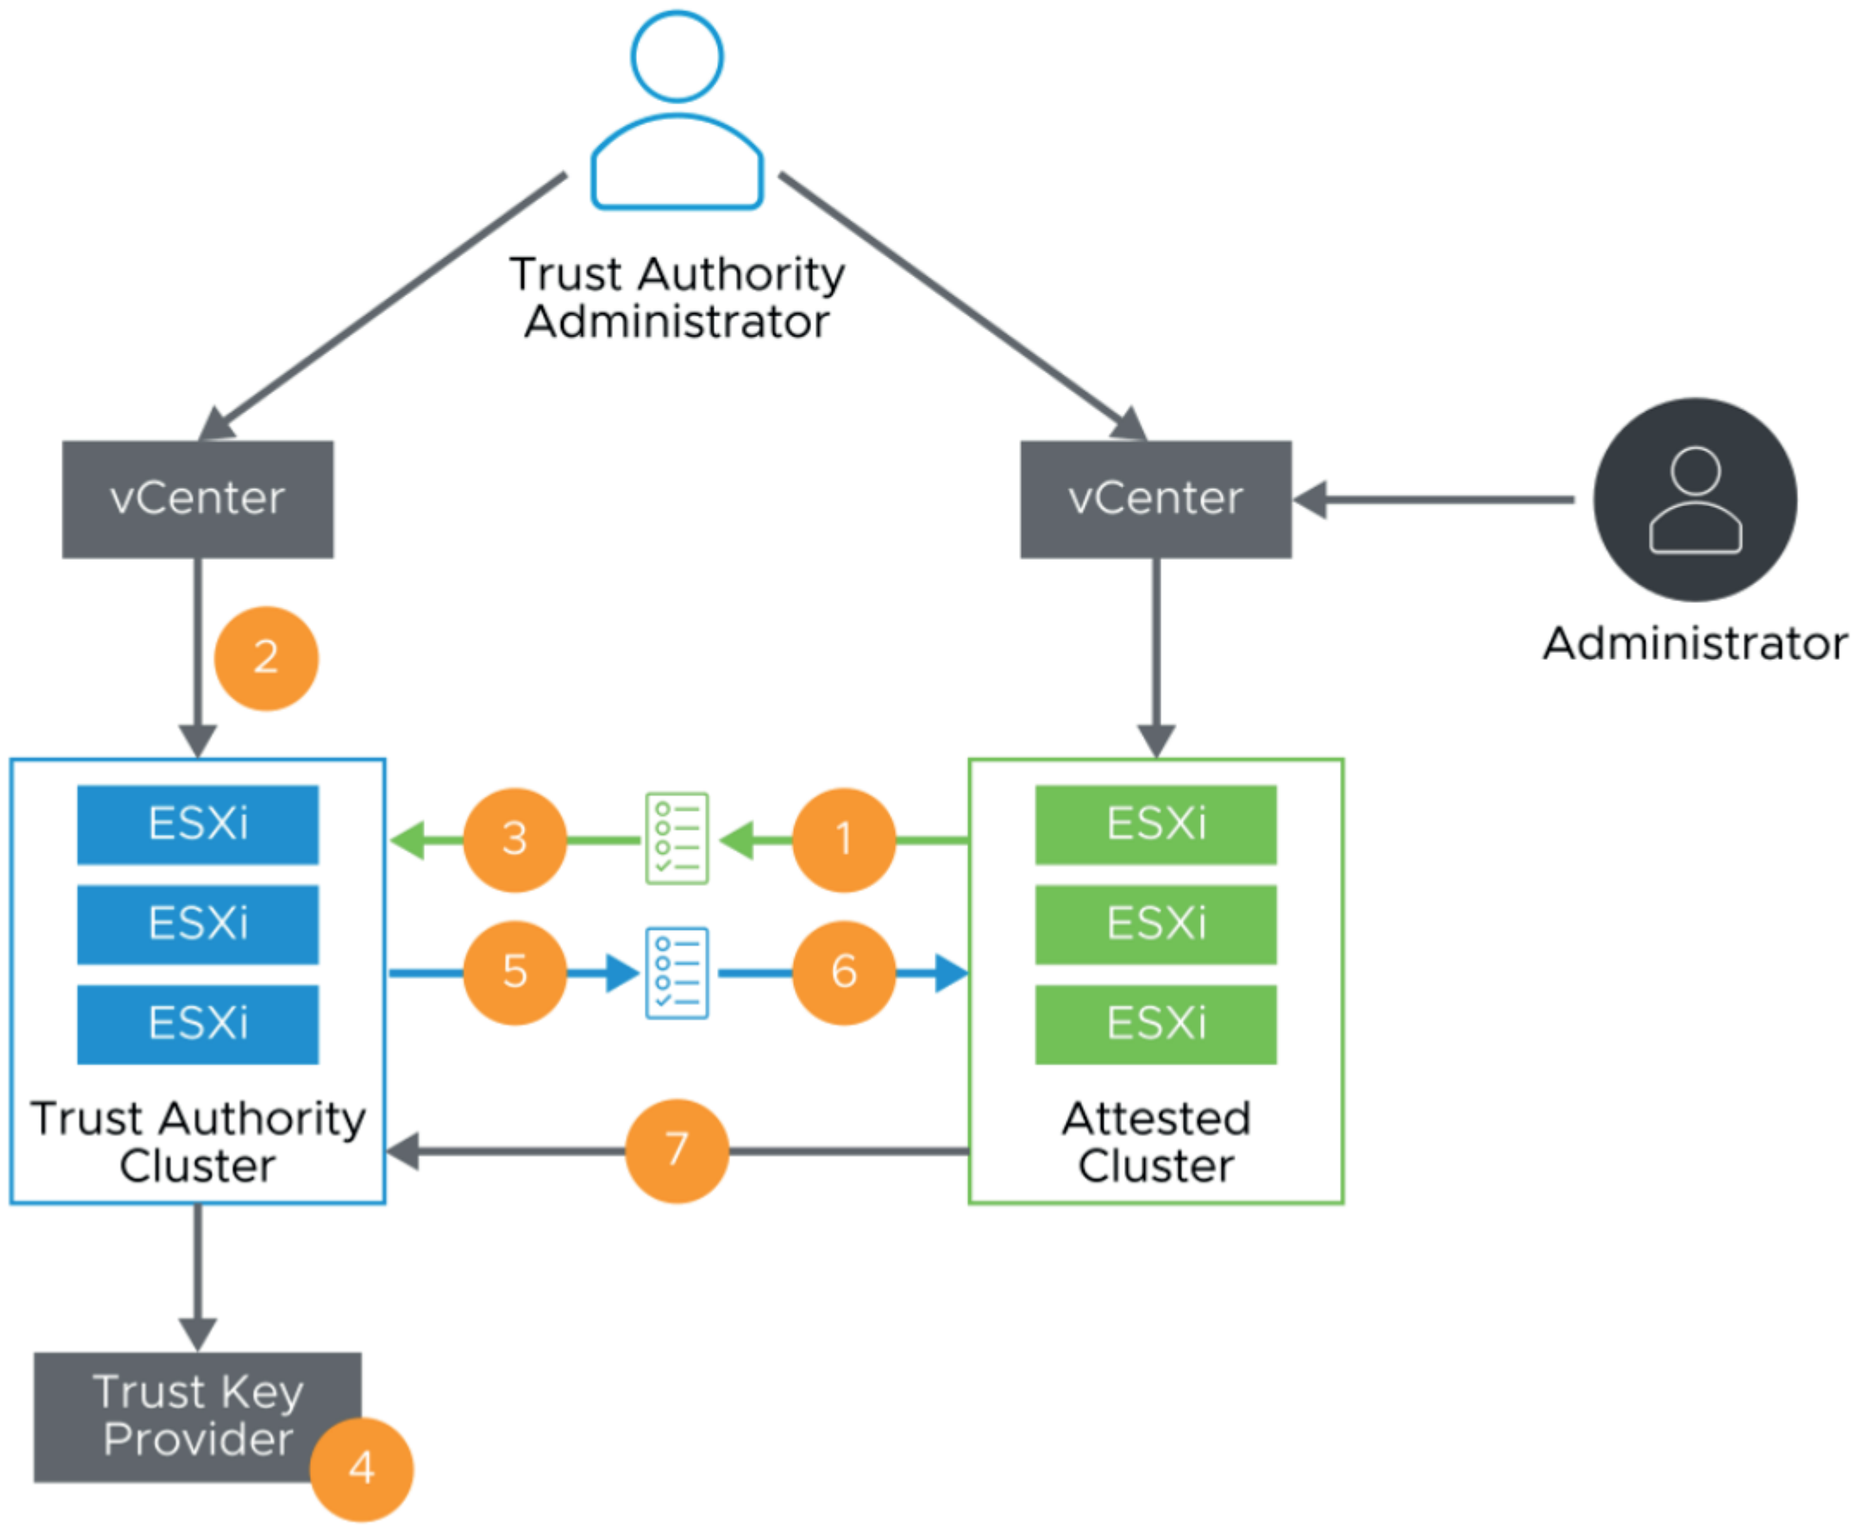

9-96 vCenter Topologies

The vSphere Trust Authority cluster must reside in a different vCenter inventory than the Trusted (Attested) clusters.

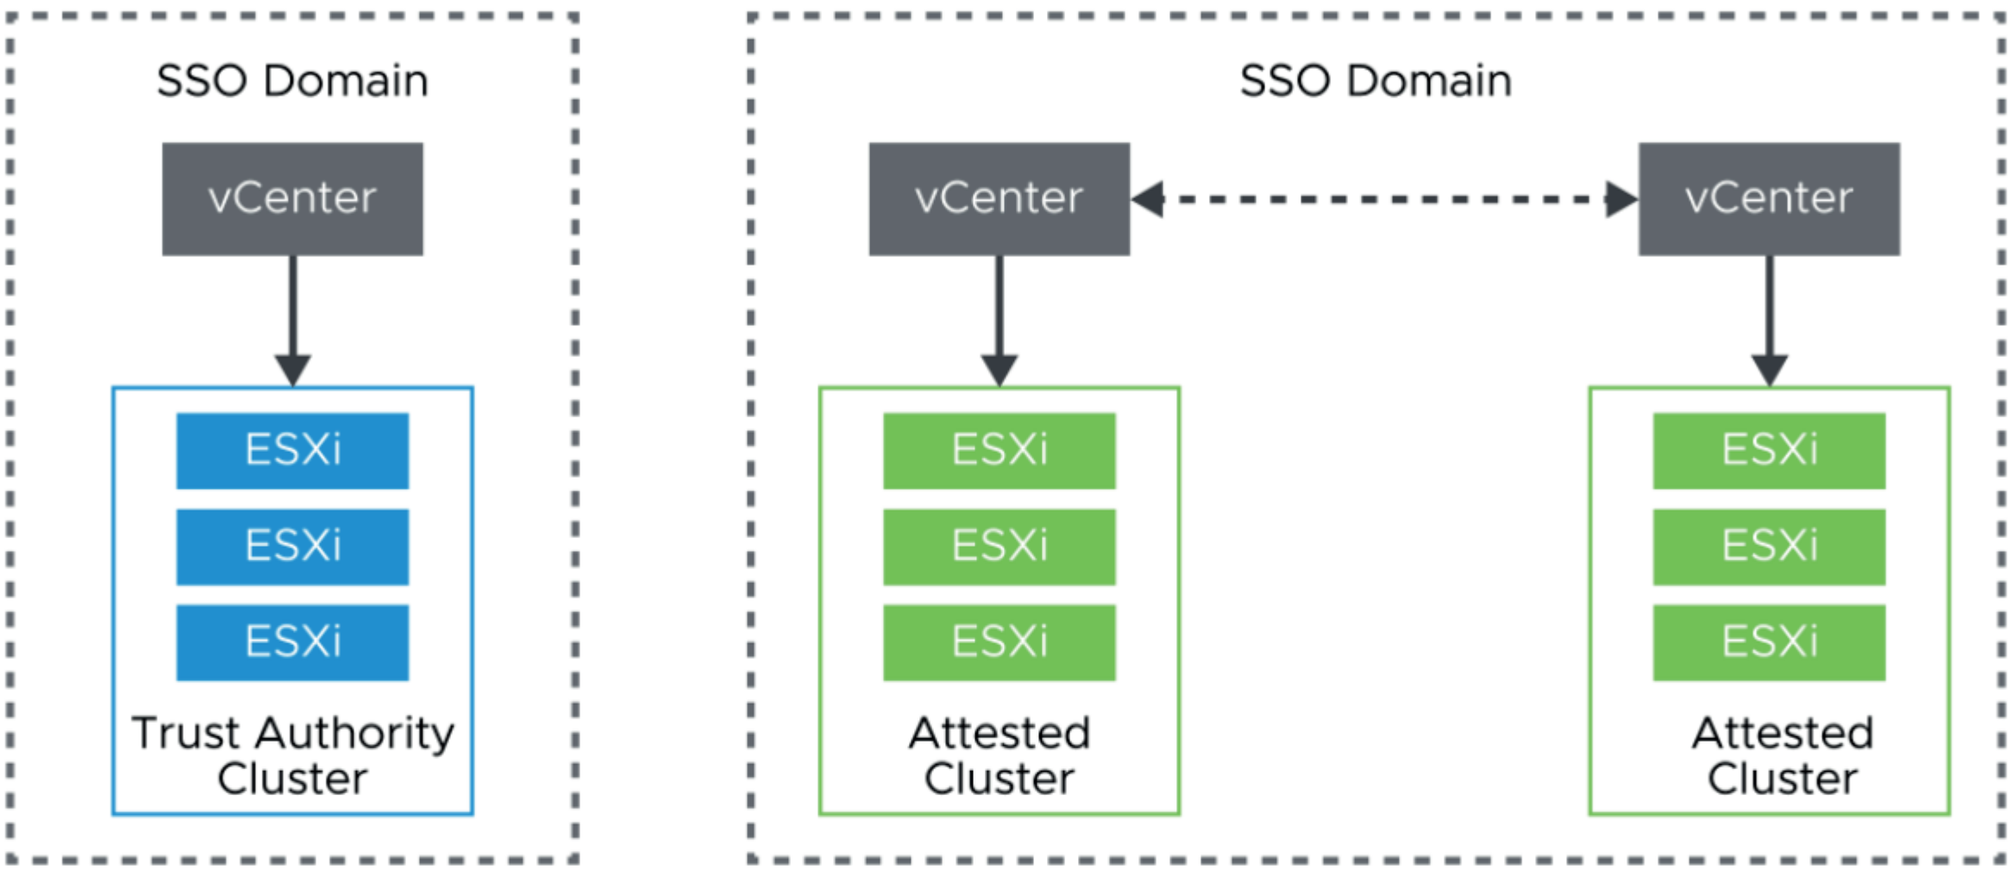

The ESXi hosting the vSphere Trust Authority cluster must be isolated from the vCenter inventory that they manage. Attested clusters can participate iin the same single sign-on (SSO) domain using Enhanced Linked Mode.

9-97 About Enabling vSphere Trust Authority

In vSphere 7 and later, you enable vSphere Trust Authority primarily through vSphere PowerCLI:

- vSphere PowerCLI 12.0 and later provides a new set of vSphere PowerCLI cmdlets.
- A user who runs the cmdlets must be a member of the TrustedAdmins single sign-on group.
- A user interface is planned for a future release.

Technet24

9-98 About VM Encryption with vSphere Trust Authority

The user experience of VM encryption does not change when it is used with vSphere Trust Authority:

- You can encrypt a VM by applying an encryption storage policy.
- An attested cluster uses a trusted key provider.
- You can mix a trusted key provider and a standard key provider in the same vSphere environment.
- When you add a virtual TPM device to a VM, the files in the VM home directory are encrypted with the trusted key provider. Virtual disks are not encrypted but can be if you apply an encryption storage policy to the VM.

9-99 VM Encryption with a Trusted Key Provider

A VM that is encrypted with a trusted key provider has a slightly different icon from a VM encrypted with a standard key provider. It also has the following characteristics:

- Can run only on ESXi hosts attested against the vSphere Trust Authority Cluster
- Cannot be migrated to an ESXi host that is not attested against the same vSphere Trust Authority Cluster

A VM encrypted with a Standard Key Provider can be migrated to an attested ESXi host, but it continues to use the standard key provider.

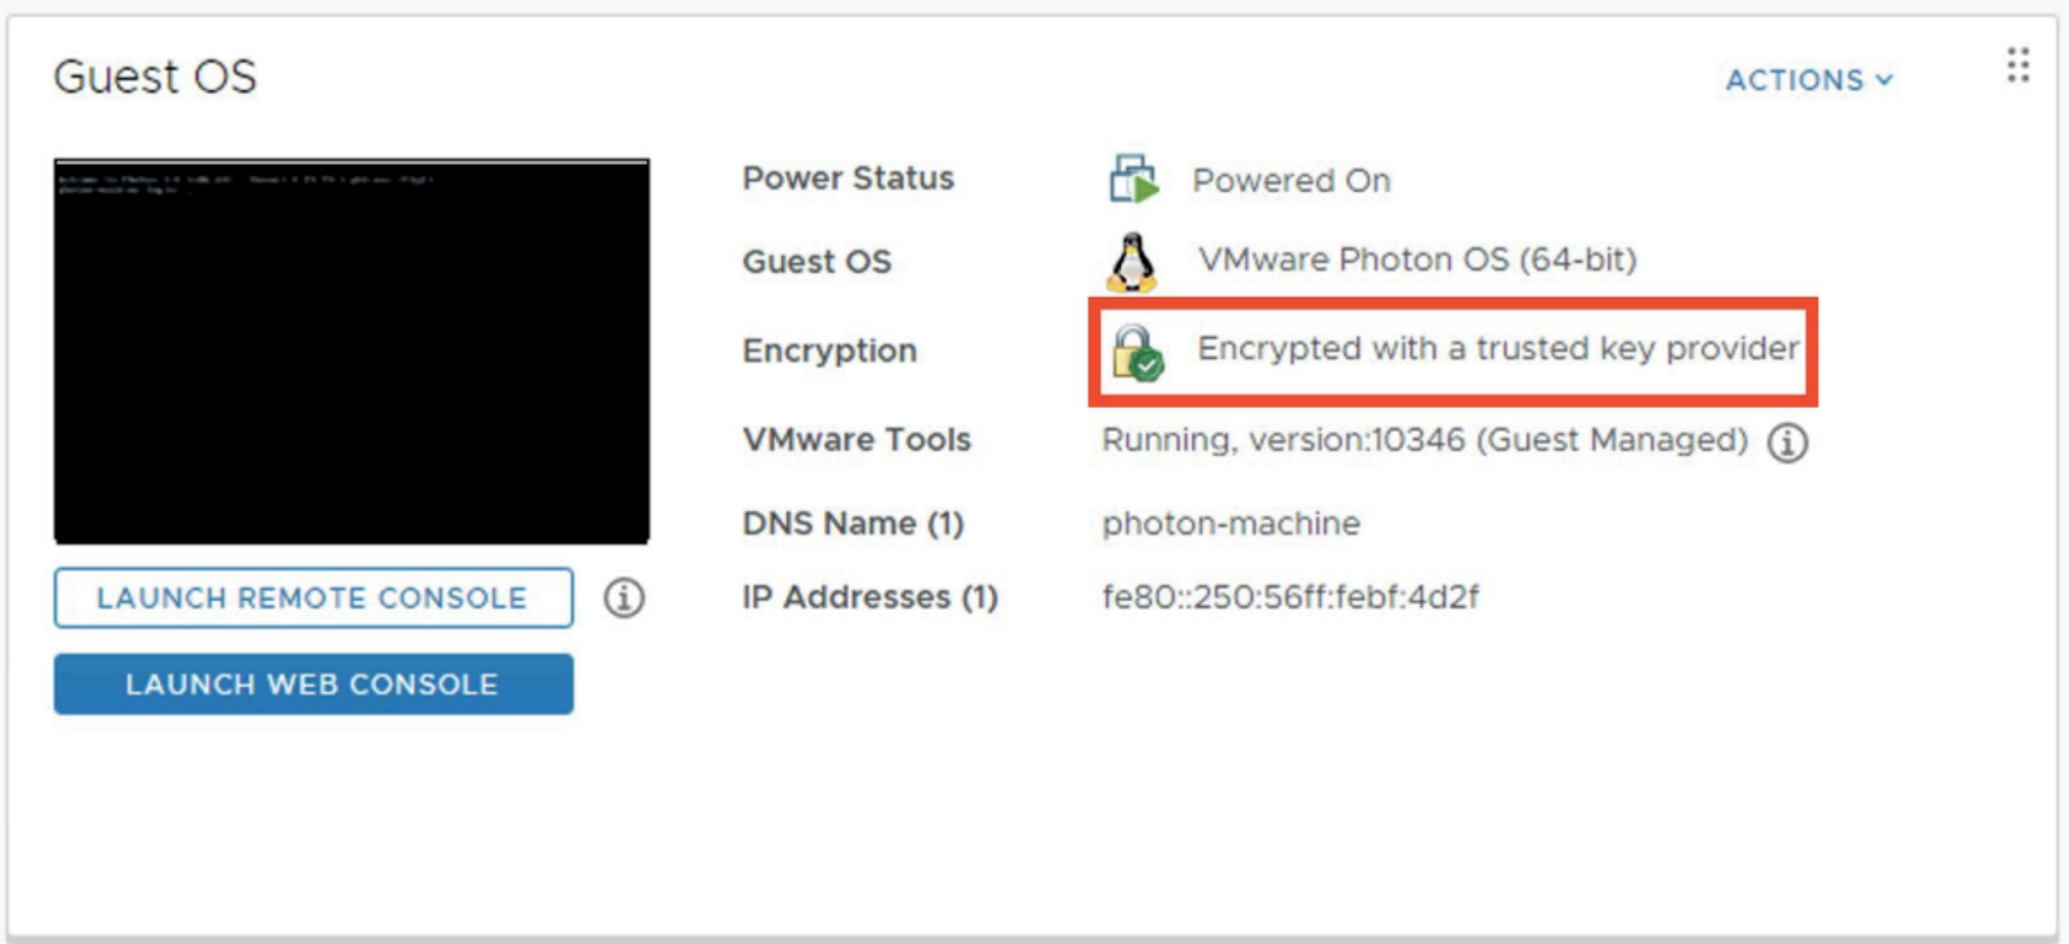

Technet24

9-100 Identifying VM Encryption Key Providers

In the vSphere Client, the VM summary indicates the type of key provider that is used for encryption.

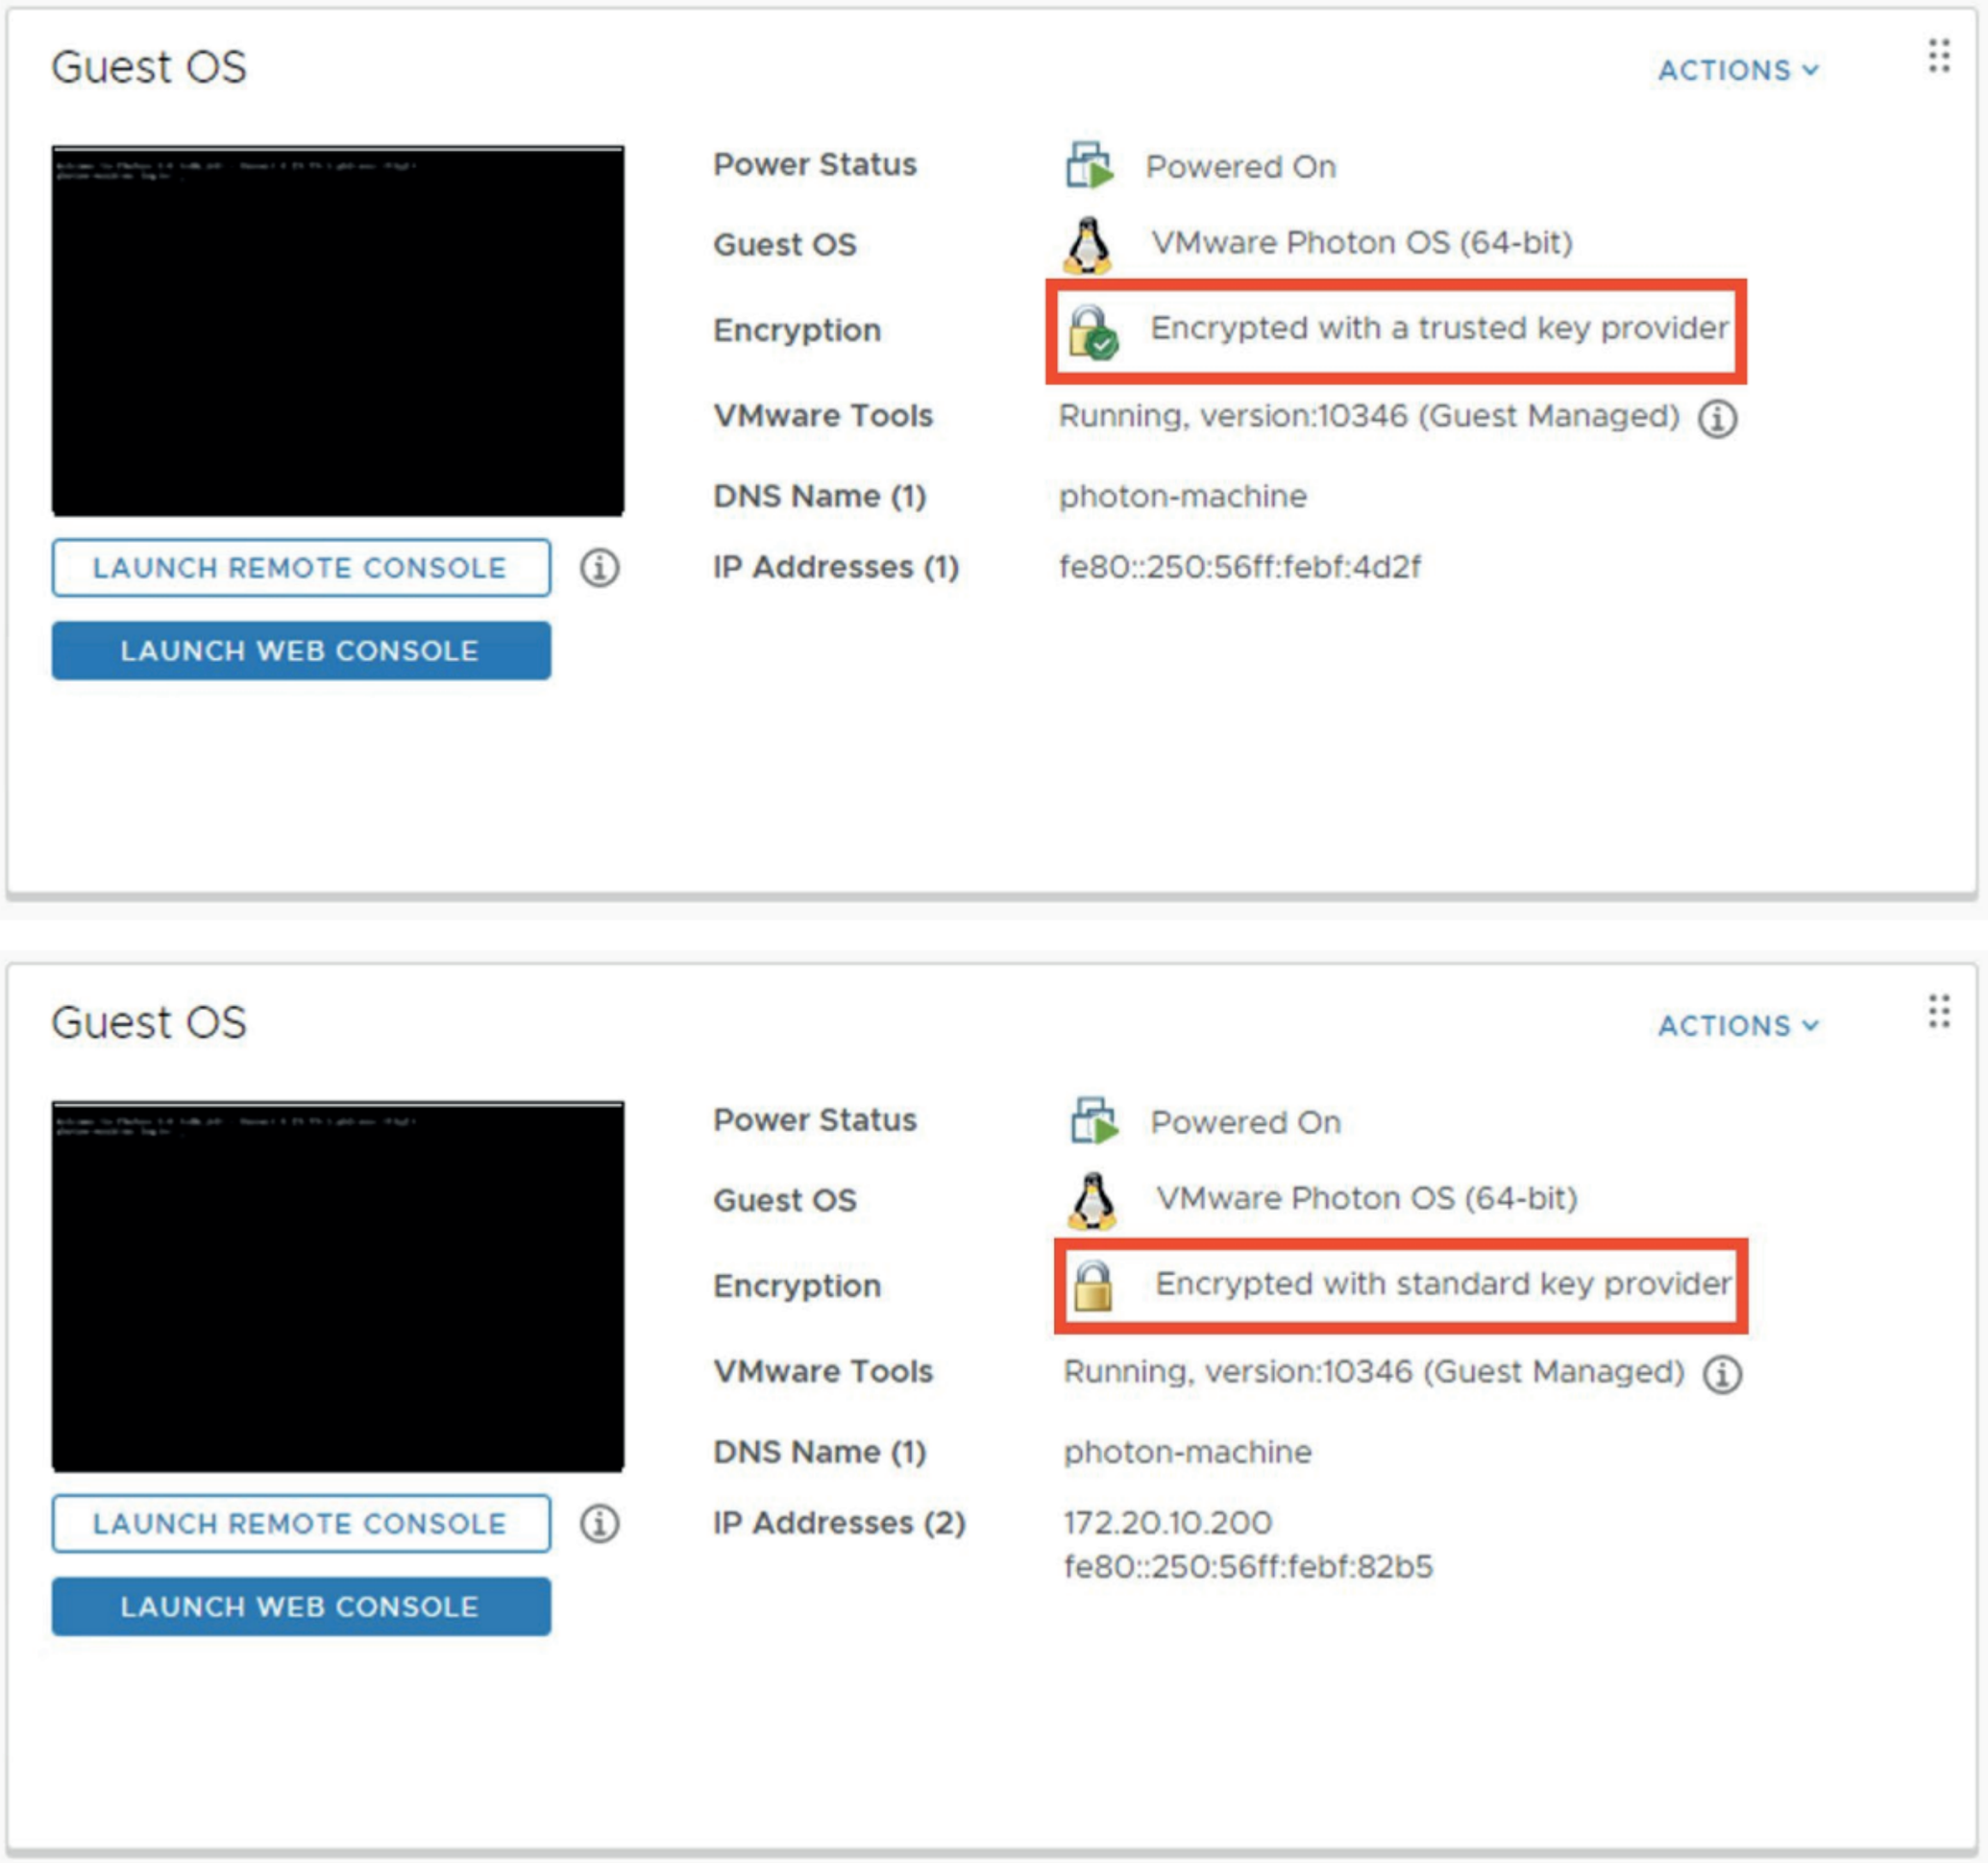

The `crypto-util dictionary describe /<path>/<vm>.vmx` command does not indicate whether the VM is encrypted with a standard or trusted key provider. However, the key provider name is present in the output so that a correlation can be made.

9-101 Migrating a VM Encrypted with a Trusted Key Provider (1)

A VM that is encrypted with a Trusted Key Provider cannot be migrated to an ESXi host that is not attested against the same vSphere Trust Authority Cluster.

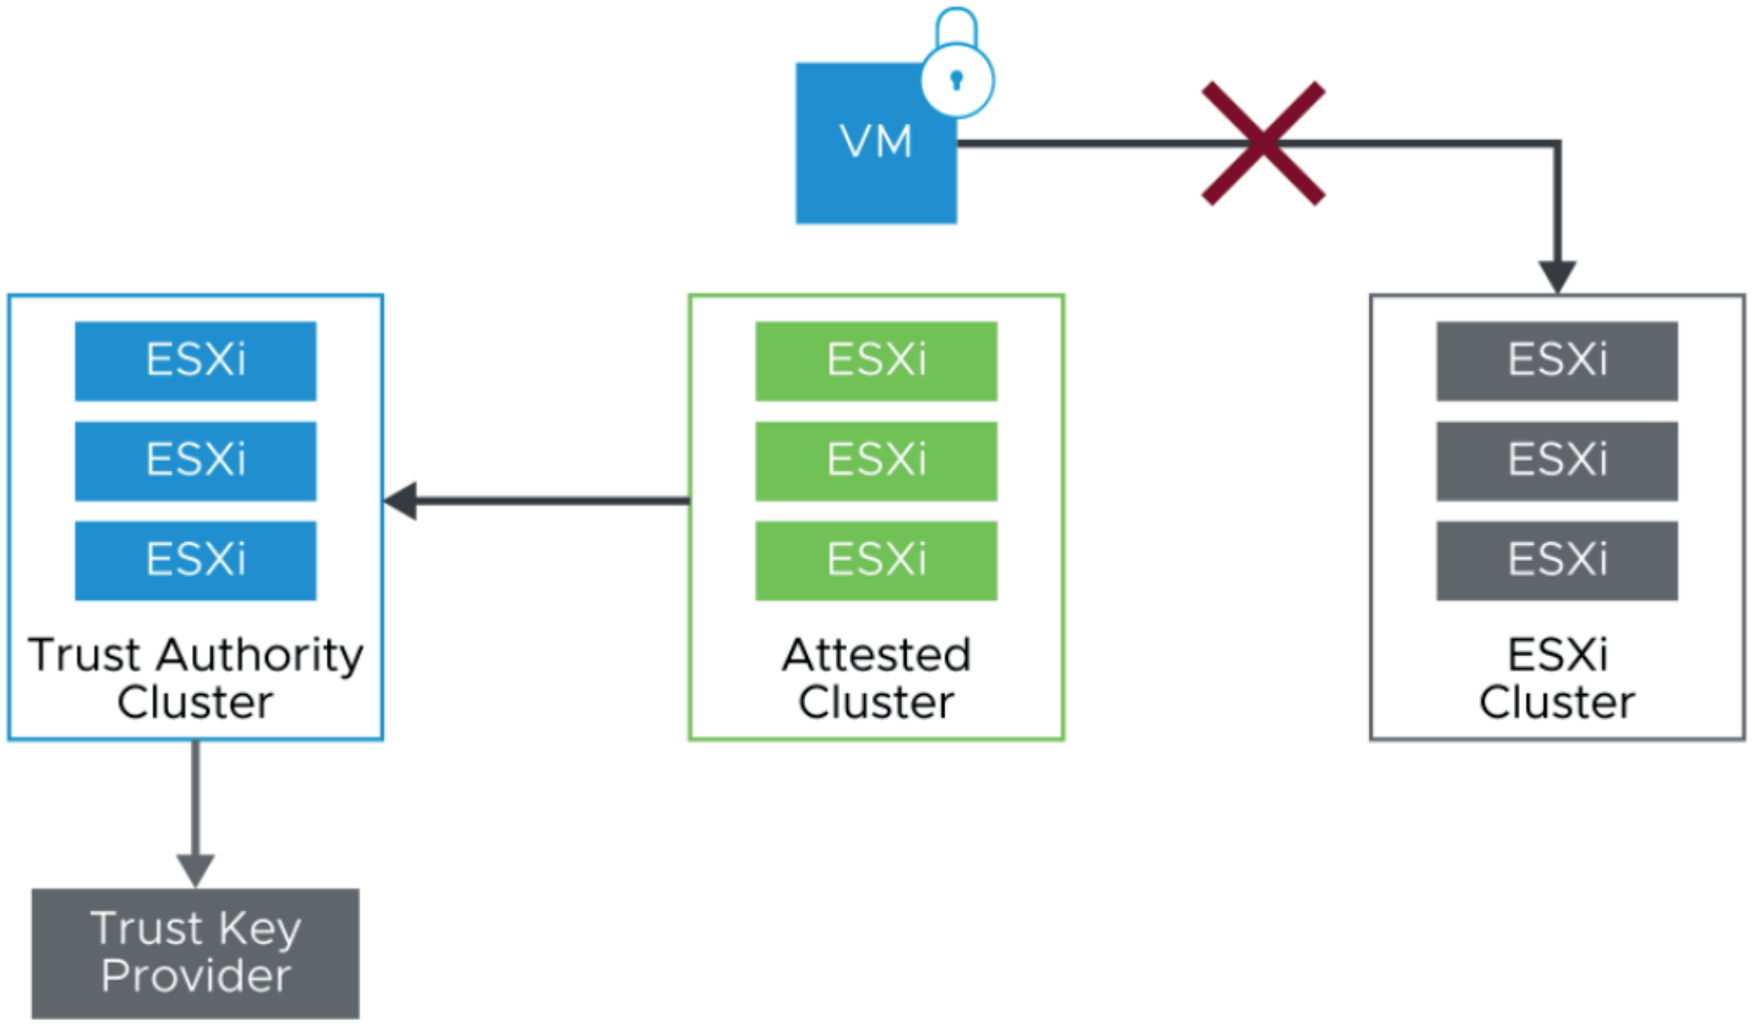

Technet24

9-102 Migrating a VM Encrypted with a Trusted Key Provider (2)

The vSphere Client displays a compatibility issue indicating that the destination host is not trusted and attested.

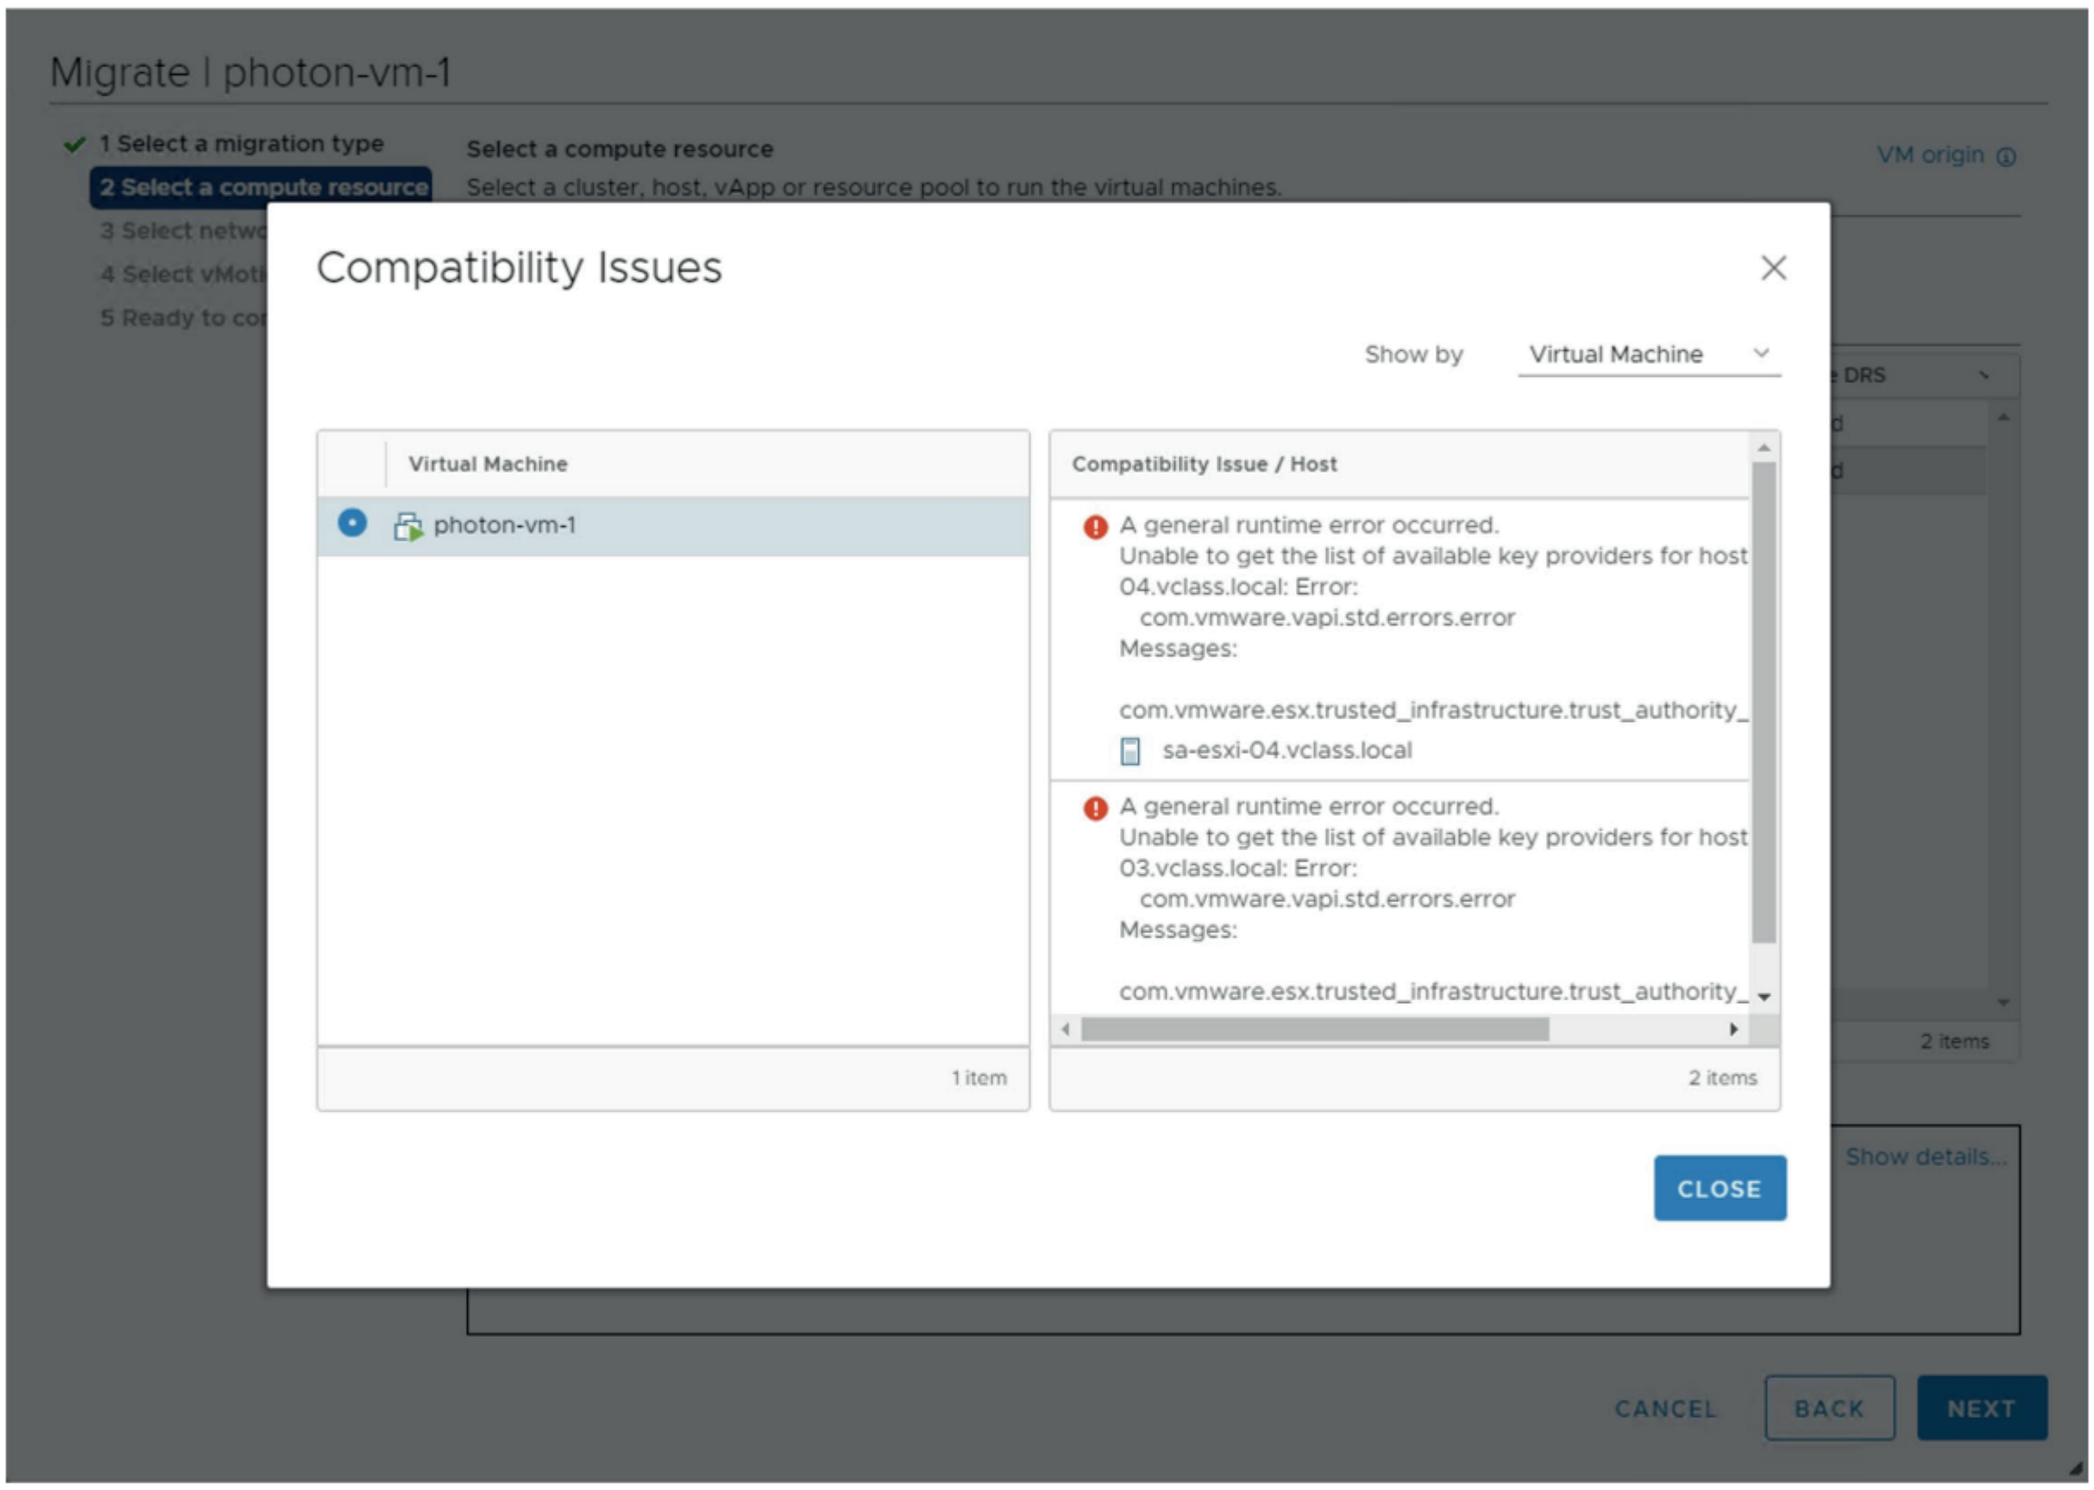

9-103 Lab 20: Configuring vCenter to work with an external KMS

Add a Key Management Server (KMS) to your vCenter from the vSphere Client:

1. Configure a KMS on vCenter
2. Establish Trust between KMS and vCenter

9-104 Lab 21: Creating an Encrypted Virtual Machine

Encrypt a new VM with a standard key provider:

1. Deploy an Encrypted VM
2. Confirm the VM is Encrypted with a Standard Key Provider

9-105 Review of Learner Objectives

- Describe vSphere Trust Authority
- List the use cases of vSphere Trust Authority
- Explain the general architecture of vSphere Trust Authority

9-106 Key Points

- Encrypted vSphere vMotion is available for all VMs, whether or not the VMs are encrypted.
- If you use VM encryption and the ESXi host fails, the resulting core dump is encrypted to protect customer data.
- To secure communication between the KMIP server (KMS) and the KMIP client (vCenter), trust must be established.
- VM encryption cannot operate without vCenter because vCenter communicates with the KMS and manages keys that are used for encryption.
- To enable vSphere Trust Authority, you must add a user to the vSphere TrustedAdmins group.

Questions?

Technet24

9-104 Lab 21: C
Machine

Encrypt a new VM with a stand

1. Deploy an Encrypted VM
2. Confirm the VM is Encrypte

9-105 Review c

- Describe vSphere Trust Au
- List the use cases of vSphe
- Explain the general architec

9-106 Key Poin

- Encrypted vSphere vMotio
- If you use VM encryption a protect customer data.
- To secure communication b trust must be established.
- VM encryption cannot ope KMS and manages keys tha
- To enable vSphere Trust A group.

Questions?

9-104 Lab 21: C
Machine

Encrypt a new VM with a stand

1. Deploy an Encrypted VM
2. Confirm the VM is Encrypte

9-105 Review c

- Describe vSphere Trust Au
- List the use cases of vSphe
- Explain the general architec

9-106 Key Poin

- Encrypted vSphere vMotio
- If you use VM encryption a protect customer data.
- To secure communication b trust must be established.
- VM encryption cannot ope KMS and manages keys tha
- To enable vSphere Trust A group.

Questions?