



Cisco HyperFlex Data Platform Administration Guide, Release 4.5

First Published: 2021-01-06

Last Modified: 2021-01-21

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Full Cisco Trademarks with Software License

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)



CONTENTS

Full Cisco Trademarks with Software License iii

CHAPTER 1

New and Changed Information for this Release 1

New and Changed Information for this Release 1

CHAPTER 2

HX Storage Cluster Overview 3

Cisco HX Data Platform Overview 3

Storage Cluster Physical Components Overview 3

HX Data Platform Capacity Overview 5

Understanding Capacity Savings 6

Storage Capacity Event Messages 7

HX Data Platform High Availability Overview 8

Storage Cluster Status 9

Operational Status Values 9

Resiliency Status Values 9

HX Data Platform Cluster Tolerated Failures 10

Data Replication Factor Settings 12

Cluster Access Policy 13

Responses to Storage Cluster Node Failures 13

HX Data Platform Ready Clones Overview 17

HX Native Snapshots Overview 17

CHAPTER 3

Logging in to HX Data Platform Interfaces 19

HyperFlex Cluster Interfaces Overview 19

Guidelines for HX Data Platform Login Credentials	20
HX Data Platform Names, Passwords, and Characters	21
AAA Authentication REST API	24
Logging into HX Connect	24
Logging into the Controller VM (steli) Command Line	25
Changing Storage Controller Password	27
Logging Into Cisco HX Data Platform Installer	27
Accessing the HX Data Platform REST APIs	28
Secure Admin Shell	29
Guidelines and Limitations	29
Information About Consent Token	29

CHAPTER 4

Monitoring HX Storage Clusters	31
Monitoring HyperFlex Clusters	31
Monitoring HyperFlex Clusters with HX Connect	31
Dashboard Page	32
Activity Page	33
System Information Overview Page	35
Nodes Page	38
Disks Page	39
Using the Cisco HX Data Platform Plug-in Interface	41
Cisco HX Data Platform Plug-in Integration with vSphere Web Client	41
Links Between the Cisco HX Data Platform Plug-in and the vSphere Interface	42
Cisco HX Data Platform Plug-in Tabs Overview	42
Monitoring Performance Charts	42
Storage Cluster Performance Chart	43
Hosts Performance Chart	43
Datastores Performance Chart	43
Datastore Trends Portlet	43
Customizing Performance Charts	44
Specify Performance Time Period	44
Specify Custom Range	45
Selecting Performance Charts	45
Audit Logging with HX Connect	46

Enabling Audit Logging	47
Configuring the Remote Syslog Server	48
Disabling Audit Logging	50
Deleting Audit Logging Server Configuration	50

CHAPTER 5

Managing HX Storage Clusters	51
Changing the Cluster Access Policy Level	51
Rebalancing the Cluster	51
Checking Cluster Rebalance and Self-Healing Status	52
Handling Out of Space Errors	53
Checking Cleaner Schedule	54
Moving the Storage Cluster from a Current vCenter Server to a New VCenter Server	54
Moving the Storage Cluster from a Current vCenter Server to a New VCenter Server	55
Unregistering a Storage Cluster from a vCenter Cluster	55
Unregistering and Removing EAM Extensions	56
Removing HX Data Platform Files from the vSphere Client	58
Verifying HX Cluster is Unregistered from vCenter	59
Registering a Storage Cluster with a New vCenter Cluster	59
Re-registering the vCenter using HX Connect	60
Renaming Clusters	61
Re-registering the vCenter using HX Connect	61
Replacing Self-Signed Certificate	62
Replacing Self-Signed Certificate with External CA Certificate on a vCenter Server	62
Replacing Self-Signed Certificate with External CA Certificate on an ESXi Host	63
Reregistering a HyperFlex cluster	64
Recreating a Self-Signed Certificate	64
Boost Mode	64
Configuring Boost Mode	65
Disabling Boost Mode	65
UEFI Secure Boot Mode	66
Enabling Secure Boot Mode	66

CHAPTER 6

Preparing for HX Storage Cluster Maintenance	69
Storage Cluster Maintenance Operations Overview	69

Serial vs. Parallel Operations	71
Checking Cluster Status	71
Setting a Beacon	71
Verify vMotion Configuration for HX Cluster	72
Maintenance Modes for Storage Cluster Nodes	73
Entering Cisco HyperFlex Maintenance Mode	74
Exiting Cisco HyperFlex Maintenance Mode	75
Creating a Backup Operation	76
Shut Down and Power Off the Cisco HX Storage Cluster	80
Power On and Start Up the Cisco HX Storage Cluster	82
Restoring the Configuration for a Fabric Interconnect	84
Configure PCI Passthrough After Changing vNIC or vHBAs	86

CHAPTER 7

Managing Encryption 87

Self-Encrypting Drives Overview	87
Verify if the HyperFlex Cluster Is Encryption Capable	88
Configuring Local Encryption Key	88
Modifying Local Encryption Key	89
Disabling Local Encryption Key	89
Secure Erase an Encrypted Disk	90
Remote Key Management	90
Configuring Remote Encryption Key	91
Generating Certificate Signing Requests	91
Configuring a Key Management Server Using CSRs (Certificate Signing Requests)	93
Generating Self-Signed Certificates	94
Configuring a key management server using SSCs (Self-Signed Certificates)	95
Restart Encryption	96

CHAPTER 8

Managing Datastores 97

Managing Datastores	97
Adding Datastores	98
Editing Datastores	99
Mounting Datastores	99
Unmounting Datastores	100

Deleting Datastores	101
Recovering from Partially Unmounted Datastores	101

CHAPTER 9

Managing Disks 103

Managing Disks in the Cluster	103
Disk Requirements	103
Replacing Self Encrypted Drives (SEDs)	105
Replacing SSDs	107
Replacing NVMe SSDs	108
Replacing Housekeeping SSDs	109
Replacing or Adding Hard Disk Drives	111

CHAPTER 10

Managing Nodes 113

Managing Nodes	113
Identify Node Maintenance Methods	115
Searching by DNS Address or Host Name	117
Changing ESXi Host Root Password	118
Reinstalling Node Software	118
Changing Node Identification Form in vCenter Cluster from IP to FQDN	119
Replacing Node Components	120
Removing a Node	122
Preparing to Remove a Node	123
Removing a Node from an Online Storage Cluster	125
Removing a Node from an Offline Storage Cluster	127
Removing a Compute Node	129
Deleting the Removed Node Data from a Disk and a Storage Controller VM	130
Replacing a Node	130
Replacing a Compute Node	134

CHAPTER 11

Expand Cisco HyperFlex System Clusters 135

Cluster Expansion Guidelines	135
ESXi Installation Guidelines	136
Prerequisites When Creating a Mixed M4/M5 Cluster	136
Mixed Cluster Expansion Guidelines - Cisco HX Release 4.5(x)	137

Steps During Mixed Cluster Expansion	137
Prerequisites for Adding a Converged (HX220c/HX240c) Node	138
Preparing a Converged Node	138
Adding a Converged Node to an Existing Cluster	139
Prerequisites for Adding a Compute-Only Node	146
Preparing a Compute-Only Node	147
Verify the HX Data Platform Installer	148
Apply an HX Profile on a Compute-only Node Using UCS Manager	148
Install VMware ESXi on Compute Nodes	148
Adding a Compute-Only Node to an Existing Cluster	149
Resolving Failure of Cluster Expansion	155
Enabling Logical Availability Zones	156
Requirements and Limitations for Logical Availability Zones	157
Expanding Cluster with Nodes Lesser than Zones	159

CHAPTER 12
Managing HX Controller VMs 161

Managing Storage Controller VMs	161
Powering On or Off Storage Controller VMs	161
Disabling HA VM Monitoring in HX Controller VMs	162

CHAPTER 13
Managing Ready Clones 165

HX Data Platform Ready Clones Overview	165
Benefits of HX Data Platform Ready Clones	166
Supported Base VMs	166
Ready Clone Requirements	166
Ready Clone Best Practices	167
Creating Ready Clones Using HX Connect	167
Creating Ready Clones Using the HX Data Platform Plug-In	169
Prepare to Customize HX Data Platform Ready Clones	170
Creating a Customization Specification for Linux in the vSphere Web Client	170
Create a Customization Specification for Windows in the vSphere Web Client	171
Configuring Ready Clones Using Customized Specifications	171
Managing Virtual Machine Networking	172

CHAPTER 14**Managing HX Native Snapshots 173**

- HX Native Snapshots Overview 173
- Benefits of HX Native Snapshots 174
- HX Native Snapshot Considerations 175
- HX Native Snapshots Best Practices 177
- Understanding SENTINEL Snapshots 178
- HX Native Snapshot Time Zones 178
- Creating HX Native Snapshots 179
- Scheduling HX Native Snapshots Overview 180
- Scheduling HX Native Snapshots 181
- Setting the Frequency of HX Native Scheduled Snapshots 182
- Deleting HX Native Snapshot Schedules 182
- Reverting to an HX Native Snapshot 183
- Deleting HX Native Snapshots 183

CHAPTER 15**Managing Clusters Running on Different HXDP Versions 185**

- Managing Clusters Running on Different HXDP Versions 185
- Scenario—Site A at HXDP 3.0 and Site B at HXDP 2.6 185
- Scenario—Site A at HXDP 2.6 and Site B at HXDP 3.0 187
- Functionality Limitations 188

CHAPTER 16**Managing Virtual Machine Disaster Recovery 189**

- HX Data Protection Snapshot Overview 189
- Replication and Recovery Considerations 190
- Replication Network and Pairing Considerations 192
- Storage Replication Adapter Overview 193
- Data Protection Terms 195
- Best Practices for Data Protection and Disaster Recovery 195
- Protecting Virtual Machines Overview 197
- Data Protection Workflow 197
- Configuring the Replication Network in HX Connect 199
- Test Local Replication Network 203
- Editing the Replication Network 203

Replication Pair Overview	204
Creating a Replication Pair	205
Test Remote Replication Network	207
Editing a Mapped Datastore Replication Pair	208
Deleting a Replication Pair	209
Creating a Protection Group	210
Editing Protection Groups	212
Deleting Protection Groups	212
Protecting Virtual Machines with an Existing Protection Group	212
Protecting Virtual Machines with a New Protection Group	214
Protecting Individual Virtual Machines	216
Unprotecting Virtual Machines	217
Disaster Recovery Overview	218
Configuring the Recovery Settings	218
Compatibility Matrix for Disaster Recovery Operations	220
Testing Virtual Machine Recovery	221
Recovering Virtual Machines	222
Recovering Virtual Machines in Protection Groups	224
Planned Migration	224
Migrating Virtual Machines in Protection Groups	226
Disaster Recovery and Re-protect	226
Protecting Virtual Machines After Disaster	228
Replication Maintenance Overview	229
Pausing Replication	230
Resuming Replication	230

CHAPTER 17
Managing Storage Replication Adapter 231

Storage Replication Adapter Overview	231
Configuring Storage Replication Adapter	232
System Requirements	232
Windows Site Recovery Manager	232
Installing a Site Recovery Manager for Windows	232
Installing a Storage Replication Adapter for Windows	232
Site Recovery Manager Virtual Appliance	233

Deploy the Site Recovery Manager Virtual Appliance	233
Installing a Photon OS Storage Replication Adapter	233
Setting up SRM Environment	234
Pairing Sites	234
Configuring Mapping	234
Refreshing Storage Replication Adapter	235
Adding an Array Manager Pair	235
Mapping Datastores on Cisco HyperFlex	236
Discovering Devices	236
Creating a Protection Group	237
Creating a Recovery Plan	237
Executing SRA Workflows	238
Testing a Recovery Plan	238
Testing the Recovery Plan Cleanup Task	239
Executing a Recovery	239
Executing a Reprotect Task	240
Troubleshooting SRA	240
Troubleshooting: Using Windows SRM Log Files	240
Troubleshooting: Using the SRM Appliance Log Files	241
Troubleshooting: Location of the VMware Virtual Center	241
Troubleshooting: Name of the VM	241
<hr/>	
CHAPTER 18	Managing Users 243
Managing Cisco HyperFlex Users Overview	243
User Management Terms	244
Audit Logs for AAA Accounting	245
Creating Cisco HX Data Platform RBAC Users	245
Assigning Users Privileges	246
<hr/>	
CHAPTER 19	Managing iSCSI 247
HyperFlex iSCSI Target Service Overview and Supported Use Cases	247
iSCSI Configuration Overview	247
iSCSI Network Page	248
Creating an iSCSI Network	248

Editing an iSCSI Network	250
Deleting an iSCSI Network	250
iSCSI Initiator Group	250
Creating an iSCSI Initiator Group	251
Editing an iSCSI Initiator Group	252
Deleting an iSCSI Initiator Group	252
Linking iSCSI Initiator Group with Targets	252
Unlinking an iSCSI Initiator Group	253
iSCSI Target Page	253
Creating an iSCSI Target	254
Editing an iSCSI Target	255
Deleting an iSCSI Target	255
Linking iSCSI Targets	255
Unlinking an iSCSI Target	256
iSCSI LUN Page	256
Creating an iSCSI LUN	257
Editing an iSCSI LUN	257
Deleting an iSCSI LUN	258
Configuring an iSCSI Initiator (Windows)	258
Configuring an iSCSI Initiator (Linux)	258
Cloning an iSCSI LUN	259
Limitations for the HX Windows Agent	260
Prerequisites for the HX Windows Agent	260
Installing HX Windows Agent for iSCSI Clone LUN	260
Installing HX Windows Agent (with Pre-Installed Dependencies) for iSCSI Clone LUN	261
Uninstalling HX Windows Agent for iSCSI Clone LUN	262
iSCSI HX Windows Agent Logs	262
Editing the Location of Service Logs	263
Accessing Cloned LUN(s) on Destination Target	264
CHAPTER 20	
Cisco HyperFlex HTML Plugin for VMware vCenter	265
Cisco HyperFlex vCenter Plugin	265
Cisco HyperFlex HTML5 Plugin for VMware vCenter	265
Cisco HyperFlex HTML5 Plugin Prerequisites	265

Install and Register the vCenter HTML5 Plugin	266
Verifying the Cisco HyperFlex HTML5 Plugin Installation from the vSphere Client	267
Uninstalling the HyperFlex HTML5 Plugin	268
Upgrading the HTML5 Plugin	268
Using the Cisco HyperFlex HTML5 Plugin	269
Navigating the HTML5 plugin	270
Cluster Management	271
Managing Users and Access to HX Clusters	271
Discover the Registered HX Cluster	271
View the HX Cluster Summary	272
Events	275
Alarms	276
HX Datastore Management	277
Managing Datastores	277
Create New Datastore	279
Edit Datastore	279
Mount or Unmount Datastore	279
Delete Datastore	280
vCenter: HyperFlex Plugin Embedded Actions	280
vCenter Server Action at Host and Clusters Level	280
vCenter Server Action at Virtual Machine Level	283
vCenter Server Action at the Storage Level	284
Cisco HyperFlex Flash Plugin for VMware vCenter	285
Cisco HyperFlex Flash Prerequisites	285
Launching HyperFlex Menu Options from vSphere	285
Checking Cluster Rebalance and Self-Healing Status from the Flash Plugin	285
Setting a Disk Beacon through the HX Flash Plugin	286



CHAPTER 1

New and Changed Information for this Release

- [New and Changed Information for this Release, on page 1](#)

New and Changed Information for this Release

The following table provides an overview of the significant changes to this guide for this current release. The table does not provide an exhaustive list of all changes made to this guide or of all new features in this release.

Feature	Description	Release/Date Added	Where Documented
Managing Storage Replication Adapter	Introduces the Site Recovery Manager Virtual Appliance	January 28, 2021	Storage Replication Adapter Overview, on page 231
Managing iSCSI	EAP Release	January 5, 2021	iSCSI Configuration Overview, on page 247
Secure Admin Shell	Improves the built-in security posture of the Controller VM by reducing its attack surface.	January 5, 2021	Secure Admin Shell, on page 29
UEFI Secure Boot	Provides an automated workflow for boot mode of converged and compute nodes in the cluster to Unified Extensible Firmware Interface (UEFI) Secure Boot.	January 5, 2021	UEFI Secure Boot Mode, on page 66
Re-register vCenter	Allows re-registering the vCenter after a vCenter upgrade and for additional scenarios.	January 5, 2021	Re-registering the vCenter using HX Connect, on page 60



CHAPTER 2

HX Storage Cluster Overview

- [Cisco HX Data Platform Overview, on page 3](#)
- [Storage Cluster Physical Components Overview, on page 3](#)
- [HX Data Platform Capacity Overview, on page 5](#)
- [HX Data Platform High Availability Overview, on page 8](#)
- [Storage Cluster Status, on page 9](#)
- [HX Data Platform Cluster Tolerated Failures, on page 10](#)
- [Responses to Storage Cluster Node Failures, on page 13](#)
- [HX Data Platform Ready Clones Overview, on page 17](#)
- [HX Native Snapshots Overview, on page 17](#)

Cisco HX Data Platform Overview

Cisco HyperFlex Data Platform (HX Data Platform) is a hyperconverged software appliance that transforms Cisco servers into a single pool of compute and storage resources. It eliminates the need for network storage and enables seamless interoperability between computing and storage in virtual environments. The Cisco HX Data Platform provides a highly fault-tolerant distributed storage system that preserves data integrity and optimizes performance for virtual machine (VM) storage workloads. In addition, native compression and deduplication reduce storage space occupied by the VMs and VM workloads.

Cisco HX Data Platform has many integrated components. These include: Cisco Fabric Interconnects (FIs), Cisco UCS Manager, Cisco HX specific servers, and Cisco compute only servers; VMware vSphere, ESXi servers, and vCenter; and the Cisco HX Data Platform Installer, controller VMs, HX Connect, vSphere HX Data Platform Plug-in, and `stcli` commands.

Cisco HX Data Platform is installed on a virtualized platform such as VMware vSphere. During installation, after specifying the Cisco HyperFlex HX Cluster name, and the HX Data Platform creates a hyperconverged storage cluster on each of the nodes. As your storage needs to increase and you add nodes in the HX cluster, the HX Data Platform balances the storage across the additional resources. Compute only nodes can be added to increase compute only resources to the storage cluster.

Storage Cluster Physical Components Overview

Cisco HyperFlex storage clusters contain the following objects. These objects are monitored by the HX Data Platform for the storage cluster. They can be added and removed from the HX storage cluster.

- **Converged nodes**—Converged nodes are the physical hardware on which the VM runs. They provide computing and storage resources such as disk space, memory, processing, power, and network I/O.

When a converged node is added to the storage cluster, a storage controller VM is installed. The HX Data Platform services are handled through the storage controller VM. Converged nodes add storage resources to your storage cluster through their associated drives.

Run the *Cluster Expansion* workflow from the HX Data Platform Installer to add converged nodes to your storage cluster. You can remove converged nodes using *stcli* commands.

- **Compute nodes**—Compute nodes add compute resource but not storage capacity to the storage cluster. They are used as a means to add compute resources, including CPU and memory. They do not need to have any caching (SSD) or storage (HDD) drives. Compute nodes are optional in a HX storage cluster.

When a compute node is added to the storage cluster, an agent controller VM is installed. The HX Data Platform services are handled through the agent controller VM.

Run the *Cluster Expansion* workflow from the HX Data Platform Installer to add compute nodes to your storage cluster. You can remove compute nodes using *stcli* commands.

- **Drives**—There are two types of drives that are required for any node in the storage cluster: Solid State Drive (SSD) and Hard Disk Drive (HDD). HDD typically provides the physical storage units associated with converged nodes. SSD typically supports management.

Adding HDD to existing converged nodes, also adds storage capacity to the storage cluster. When storage is added to a HX node in the storage cluster, an equal amount of storage must be added to every node in the storage cluster.

When disks are added or removed, the HX Data Platform rebalances the storage cluster to adjust for the change in storage resources.

Adding or removing disks on your converged nodes is not performed through the HX Data Platform. Before adding or removing disks, review the best practices. See the server hardware guides for specific instructions to add or remove disks in nodes.

NVMe Caching SSD's slot information is unavailable from HX-Connect for all AF server PIDs except for the All-NVMe server PIDs. Please refer to UCSM management console for NVMe SSD slot information.

- **Datastores**—Storage capacity and datastore capacity. This is the combined consumable physical storage available to the storage cluster through datastores, and managed by the HX Data Platform.

Datastores are logical containers that are used by the HX Data Platform to manage your storage use and storage resources.

Datastores are where the host places virtual disk files and other VM files. Datastores hide the specifics of physical storage devices and provide a uniform model for storing VM files.

HX Data Platform Capacity Overview

**Note**

Capacity addition in a cluster through the addition of disks or nodes can result in a rebalance. This background activity can cause interference with regular User IO on the cluster and increase the latency. You must note the time duration for the storage capacity at the time where performance impact can be tolerated. Also, this operation may be performed in urgent situations that may warrant capacity addition.

In the HX Data Platform the concept of capacity is applied to both datastores and storage clusters. Values are measured in base-2 (GiB/TiB), but for simplicity and consistency are labeled as GB or TB.

- **Cleaner**—A process run on all the storage cluster datastores. After it completes, all the storage cluster datastores total capacity should be in a similar range to the total storage cluster capacity, excluding the metadata. Datastore capacity listed typically will not match the HX storage cluster capacity. See the [Cisco HX Data Platform Command Line Interface Reference Guide](#) for information on the `cleaner` command.
- **Cluster capacity**—All the storage from all the disks on all the nodes in the storage cluster. This includes uncleaned data and the metadata overhead for each disk.

The total/used/free capacity of cluster is based on overall storage capacity and how much storage is used.

- **Condition**—When the HX Storage Cluster enters a space event state, the **Free Space Status** fields are displayed. The **Condition** field lists the space event state. The options are: **Warning**, **Critical**, and **Alert**.
- **Available Datastore capacity**—The amount of storage available for provisioning to datastores without over-provisioning. Generally, this is similar to the cleaned storage cluster capacity, but it is not an exact match. It does not include metadata or uncleaned data.

The provisioned/used/free capacity of each datastore is based on datastore (thin) provisioned capacity. Because the datastore is thin provisioned, the provisioned capacity (specified by the administrator when creating the datastore) can be well above the actual storage.

- **Free Capacity, storage cluster**—Same as available capacity. For the storage cluster, this is the difference between the amount available to the storage cluster and the amount used in the storage cluster.
- **Free capacity, datastore**—Same as available capacity. For all the storage cluster datastores, this is the difference between the amount provisioned to all the storage cluster datastores and the amount used on all the storage cluster datastores.

The amount used on the whole storage cluster is not included in this datastore calculation. Because datastores are frequently over provisioned, the free capacity can indicate a large availability on all the storage cluster datastores, while the storage cluster capacity can indicate a much lower availability.

- **Multiple users**—Can have different datastores with different provisioned capacities. At any point in time, users do not fully utilize their allocated datastore capacity. When allocating datastore capacity to multiple users, it is up to the administrator to ensure that each user's provisioned capacity is honored at all time.
- **Over-provisioning**—Occurs when the amount of storage capacity allocated to all the datastores exceeds the amount available to the storage cluster.

It is a common practice to initially over-provision. It allows administrators to allocate the capacity now and backfill the actual storage later.

The value is the difference between the usable capacity and provisioned capacity.

It displays zero (0) value, unless more space has been allocated than the maximum physical amount possible.

Review the over provisioned capacity and ensure that your system does not reach an out-of-space condition.

- **Provisioned**—Amount of capacity allowed to be used by and allocated to the storage cluster datastores.

The provisioned amount is not set aside for the sole use of the storage cluster datastores. Multiple datastores can be provisioned storage from the same storage capacity.

- **Space Needed**—When the HX Storage Cluster enters a space event state, the **Free Space Status** fields are displayed. **Space Needed** indicates the amount of storage that needs to be made available to clear the listed **Condition**.

- **Used**—Amount of storage capacity consumed by the listed storage cluster or datastore.

HX Data Platform internal meta-data uses 0.5% to 1% space. This might cause the HX Data Platform Plug-in or HX Connect to display a Used Storage value even if you have no data in your datastore.

Storage Used shows how much datastore space is occupied by virtual machine files, including configuration and log files, snapshots, and clones. When the virtual machine is running, the used storage space also includes swap files.

- **Usable Capacity**—Amount of storage in the storage cluster available for use to store data.

Understanding Capacity Savings

The Capacity portlet on the Summary tab displays the deduplication and compression savings provided by the storage cluster. For example, with 50% overall savings, a 6TB capacity storage cluster can actually store 9 TB of data.

The total storage capacity saved by the HX Data Platform system is a calculation of two elements:

- **Compression**—How much of the data is compressed.
- **Deduplication**—How much data is deduplicated. Deduplication is a method of reducing storage space by eliminating redundant data. It stores only one unique instance of the data.

Deduplication savings and compression savings are not simply added together. They are not independent operations. They are correlated using the following elements where essentially the number of unique bytes used for storage is reduced through deduplication. Then the deduplicated storage consumption is compressed to make even more storage available to the storage cluster.

Deduplication and compression savings are useful when working with VM clones.

If the savings is showing 0%, this indicates the storage cluster is new. The total ingested data to the storage cluster is insufficient to determine meaningful storage savings. Wait until sufficient data is written to the storage cluster.

For example:

1. Initial values

Given a VM of 100 GB that is cloned 2 times.

Total Unique Used Space (TUUS) = 100GB

Total Addressable Space (TAS) = 100x2 = 200 GB

Given, for this example:

Total Unique Bytes (TUB) = 25 GB

2. Deduplication savings

$$= (1 - \text{TUUS}/\text{TAS}) * 100$$

$$= (1 - 100\text{GB} / 200\text{GB}) * 100$$

$$= 50\%$$

3. Compression Savings

$$= (1 - \text{TUB}/\text{TUUS}) * 100$$

$$= (1 - 25\text{GB} / 100\text{GB}) * 100$$

$$= 75\%$$

4. Total savings calculated

$$= (1 - \text{TUB}/\text{TAS}) * 100$$

$$= (1 - 25\text{GB} / 200\text{GB}) * 100$$

$$= 87.5\%$$

Storage Capacity Event Messages

Cluster storage capacity includes all the storage from all the disks on all the nodes in the storage cluster. This available capacity is used to manage your data.

Calculating Cluster Capacity

A HyperFlex HX Data Platform cluster capacity is calculated as follows:

$$(((\text{capacity disk size in GB} * 10^9) / 1024^3) * \text{number of capacity disks per node} * \text{number of HyperFlex nodes} * 0.92) / \text{replication factor}$$

Divide the result by 1024 to get a value in TiB. The replication factor value is 3 if the HX cluster is set to RF=3, and the value is 2 if the HX cluster is set to RF=2. The 0.92 multiplier accounts for an 8% reservation set aside on each disk by the HX Data Platform software for various internal filesystem functions.

Calculation Example: <capacity disk size in GB> = 1200 for 1.2 TB disks <number of capacity disks per node> = 15 for an HX240c-M4SX model server <number of HyperFlex nodes> = 8
replication factor = 3

Result: $((1200 * 10^9) / 1024^3) * 15 * 8 * 0.92 / 3 = 41127.2049$ 41127.2049 / 1024 = 40.16 TiB



Note

This formula for calculating cluster capacity does not apply for Large Form Factor (LFF) clusters.

Error Messages

Error messages are issued if your data storage needs to consume high amounts of available capacity, the performance and health of your storage cluster are affected. The error messages are displayed in vCenter Alarms panels, HX Connect, and HX Data Platform Plug-in Alarms and Events pages.



Note When the warning or critical errors appear:

Add additional drives or nodes to expand capacity. Additionally, consider deleting unused virtual machines and snapshots. Performance is impacted until storage capacity is reduced.

- **SpaceWarningEvent** – Issues an error. This is a first level warning.

Cluster performance is impacted due to increased cleaner activity to reclaim the space as fast as possible. The effect on throughput and latency depend on the workload and how much read and writes are being performed.

Reduce the amount of storage capacity used to below the warning threshold, of 76% total HX Storage Cluster capacity.

- **SpaceAlertEvent** – Issues an error. Space capacity usage remains at error level.

This alert is issued after storage capacity has been reduced, but is still above the warning threshold.

Cluster performance is affected.

Continue to reduce the amount of storage capacity used, until it is below the warning threshold, of 80% total HX Storage Cluster capacity.

- **SpaceCriticalEvent** – Issues an error. This is a critical level warning.

Cluster is in a read only state.

Do not continue the storage cluster operations until you reduce the amount of storage capacity used to below this warning threshold, that is, 100% of the available disk space.

- **SpaceRecoveredEvent** - This is informational. The cluster capacity has returned to normal range.

Cluster storage space usage is back to normal.

HX Data Platform High Availability Overview

The HX Data Platform High Availability (HA) feature ensures that the storage cluster maintains at least two copies of all your data during normal operation with three or more fully functional nodes.

If nodes or disks in the storage cluster fail, the cluster's ability to function is affected. If more than one node fails or one node and disk(s) on a different node fail, it is called a *simultaneous failure*.

The number of nodes in the storage cluster, combined with the Data Replication Factor and Access Policy settings, determine the state of the storage cluster that results from node failures.



Note Before using the HX Data Platform HA feature, enable DRS and vMotion on the vSphere Web Client.

Storage Cluster Status

HX Data Platform storage cluster status information is available through HX Connect, the HX Data Platform Plug-in, and the storage controller VM `stcli` commands. Storage cluster status is described through resiliency and operational status values.

Storage cluster status is described through the following reported status elements:

- **Operational Status**—Describes the ability of the storage cluster to perform the functions storage management and storage cluster management of the cluster. Describes how well the storage cluster can perform operations.
- **Resiliency Status**—Describes the ability of the storage clusters to tolerate node failures within the storage cluster. Describes how well the storage cluster can handle disruptions.

The following settings take effect when the storage cluster transitions into particular operational and resiliency status states.

- **Data Replication Factor**—Sets the number of redundant data replicas.
- **Cluster Access Policy**—Sets the level of data protection and data loss.

Operational Status Values

Cluster Operational Status indicates the operational status of the storage cluster and the ability for the applications to perform I/O.

The Operational Status options are:

- **Online**—Cluster is ready for IO.
- **Offline**—Cluster is not ready for IO.
- **Out of space**—Either the entire cluster is out of space or one or more disks are out of space. In both cases, the cluster cannot accept write transactions, but can continue to display static cluster information.
- **Readonly**—Cluster cannot accept write transactions, but can continue to display static cluster information.
- **Unknown**—This is a transitional state while the cluster is coming online.

Other transitional states might be displayed during cluster upgrades and cluster creation.

Color coding and icons are used to indicate various status states. Click icons to display additional information such as reason messages that explain what is contributing to the current state.

Resiliency Status Values

Resiliency status is the data resiliency health status and ability of the storage cluster to tolerate failures.

Resiliency Status options are:

- **Healthy**—The cluster is healthy with respect to data and availability.
- **Warning**—Either the data or the cluster availability is being adversely affected.

- **Unknown**—This is a transitional state while the cluster is coming online.

Color coding and icons are used to indicate various status states. Click an icon to display additional information, such as reason messages that explain what is contributing to the current state.

HX Data Platform Cluster Tolerated Failures

If nodes or disks in the HX storage cluster fail, the cluster's ability to function is affected. If more than one node fails or one node and disk(s) on a different node fail, it is called a *simultaneous failure*.

How the number of node failures affect the storage cluster is dependent upon:

- **Number of nodes in the cluster**—The response by the storage cluster is different for clusters with 3 to 4 nodes and 5 or greater nodes.
- **Data Replication Factor**—Set during HX Data Platform installation and cannot be changed. The options are 2 or 3 redundant replicas of your data across the storage cluster.



Attention

Data Replication Factor of 3 is recommended.

- **Access Policy**—Can be changed from the default setting after the storage cluster is created. The options are strict for protecting against data loss, or lenient, to support longer storage cluster availability.

Cluster State with Number of Failed Nodes

The tables below list how the storage cluster functionality changes with the listed number of simultaneous node failures.

Table 1: Cluster State in 5+ Node Cluster with Number of Failed Nodes, HX Release 4.0(1x)

Replication Factor	Access Policy	Number of Failed Nodes		
		Read/Write	Read-Only	Shutdown
3	Lenient	2	--	3
3	Strict	1	2	3
2	Lenient	1	--	2
2	Strict	--	1	2

Table 2: Cluster State in 5+ Node Cluster with Number of Failed Nodes, HX Release 4.0(2x)

Replication Factor	Access Policy	Number of Failed Nodes	
		Read/Write	Read-Only
3	Lenient	2	--
3	Strict	1	2

Replication Factor	Access Policy	Number of Failed Nodes	
		Read/Write	Read-Only
2	Lenient	1	--
2	Strict	--	1

Table 3: Cluster State in 3 - 4 Node Clusters with Number of Failed Nodes HX Release 4.0(1x)

Replication Factor	Access Policy	Number of Failed Nodes		
		Read/Write	Read-Only	Shutdown
3	Lenient or Strict	1	--	2
2	Lenient	1	--	2
2	Strict	--	1	2

Table 4: Cluster State in 3 - 4 Node Clusters with Number of Failed Nodes HX Release 4.0(2x)

Replication Factor	Access Policy	Number of Failed Nodes	
		Read/Write	Read-Only
3	Lenient or Strict	1	--
2	Lenient	1	--
2	Strict	--	1

Cluster State with Number of Nodes with Failed Disks

The table below lists how the storage cluster functionality changes with the number of nodes that have one or more failed disks. Note that the node itself has not failed but disk(s) within the node have failed. **For example:** 2 indicates that there are 2 nodes that each have at least one failed disk.

There are two possible types of disks on the servers: SSDs and HDDs. When we talk about multiple disk failures in the table below, it's referring to the disks used for storage capacity. **For example:** If a cache SSD fails on one node and a capacity SSD or HDD fails on another node the storage cluster remains highly available, even with an Access Policy strict setting.

The table below lists the worst case scenario with the listed number of failed disks. This applies to any storage cluster 3 or more nodes. **For example:** A 3 node cluster with Replication Factor 3, while self-healing is in progress, only shuts down if there is a total of 3 simultaneous disk failures on 3 separate nodes.



Note HX storage clusters are capable of sustaining serial disk failures, (separate disk failures over time). The only requirement is that there is sufficient storage capacity available for support self-healing. The worst-case scenarios listed in this table only apply during the small window while HX is completing the automatic self-healing and rebalancing.

3+ Node Cluster with Number of Nodes with Failed Disks

Replication Factor	Access Policy	Failed Disks on Number of Different Nodes		
		Read/Write	Read Only	Shutdown
3	Lenient	2	--	3
3	Strict	1	2	3
2	Lenient	1	--	2
2	Strict	--	1	2

Replication Factor	Access Policy	Failed Disks on Number of Different Nodes	
		Read/Write	Read Only
3	Lenient	2	--
3	Strict	1	2
2	Lenient	1	--
2	Strict	--	1

Data Replication Factor Settings



Note Data Replication Factor cannot be changed after the storage cluster is configured.

Data Replication Factor is set when you configure the storage cluster. Data Replication Factor defines the number of redundant replicas of your data across the storage cluster. The options are 2 or 3 redundant replicas of your data.

- If you have hybrid servers (servers that contain both SSD and HDDs), then the default is 3.
- If you have all flash servers (servers that contain only SSDs), then you must explicitly select either 2 or 3 during HX Data Platform installation.

Choose a Data Replication Factor. The choices are:

- Data Replication Factor 3 — Keep three redundant replicas of the data. This consumes more storage resources, and ensures the maximum protection for your data in the event of node or disk failure.

Attention Data Replication Factor 3 is the recommended option.

- Data Replication Factor 2 — Keep two redundant replicas of the data. This consumes fewer storage resources, but reduces your data protection in the event of node or disk failure.

Cluster Access Policy

The Cluster Access Policy works with the Data Replication Factor to set levels of data protection and data loss prevention. There are two Cluster Access Policy options. The default is `lenient`. It is not configurable during installation, but can be changed after installation and initial storage cluster configuration.

- **Strict** - Applies policies to protect against data loss.

If nodes or disks in the storage cluster fail, the cluster's ability to function is affected. If more than one node fails or one node and disk(s) on a different node fail, it is called a simultaneous failure. The strict setting helps protect the data in event of simultaneous failures.

- **Lenient** - Applies policies to support longer storage cluster availability. This is the default.

Responses to Storage Cluster Node Failures

A storage cluster healing timeout is the length of time HX Connect or HX Data Platform Plug-in waits before automatically healing the storage cluster. If a disk fails, the healing timeout is 1 minute. If a node fails, the healing timeout is 2 hours. A node failure timeout takes priority if a disk and a node fail at same time or if a disk fails after node failure, but before the healing is finished.

When the cluster resiliency status is Warning, the HX Data Platform system supports the following storage cluster failures and responses.

Optionally, click the associated Cluster Status/Operational Status or Resiliency Status/Resiliency Health in HX Connect and HX Data Platform Plug-in, to display reason messages that explain what is contributing to the current state.

Cluster Size	Number of Simultaneous Failures	Entity Failed	Maintenance Action to Take
3 nodes	1	One node.	The storage cluster does not automatically heal. Replace the failed node to restore storage cluster health.

Cluster Size	Number of Simultaneous Failures	Entity Failed	Maintenance Action to Take
3 nodes	2	Two or more disks on two nodes are blacklisted or failed.	<ol style="list-style-type: none"> 1. If one cache SSD fails, the storage cluster does not automatically heal. 2. If one HDD fails or is removed, the disk is blacklisted immediately. The storage cluster automatically begins healing within a minute. 3. If more than one HDD fails, the system might not automatically restore storage cluster health. If the system is not restored, replace the faulty disk and restore the system by rebalancing the cluster.
4 nodes	1	One node.	<p>If the node does not recover in two hours, the storage cluster starts healing by rebalancing data on the remaining nodes.</p> <p>To recover the failed node immediately and fully restore the storage cluster:</p> <ol style="list-style-type: none"> 1. Check that the node is powered on and restart it if possible. You might need to replace the node. 2. Rebalance the cluster.

Cluster Size	Number of Simultaneous Failures	Entity Failed	Maintenance Action to Take
4 nodes	2	Two or more disks on two nodes.	<p>If two SSDs fail, the storage cluster does not automatically heal.</p> <p>If the disk does not recover in one minute, the storage cluster starts healing by rebalancing data on the remaining nodes.</p>
5+ nodes	2	Up to two nodes.	<p>If the node does not recover in two hours, the storage cluster starts healing by rebalancing data on the remaining nodes.</p> <p>To recover the failed node immediately and fully restore the storage cluster:</p> <ol style="list-style-type: none"> 1. Check that the node is powered on and restart it if possible. You might need to replace the node. 2. Rebalance the cluster. <p>If the storage cluster shuts down, see Troubleshooting, Two Nodes Fail Simultaneously Causes the Storage Cluster to Shutdown section.</p>
5+ nodes	2	Two nodes with two or more disk failures on each node.	The system automatically triggers a rebalance after a minute to restore storage cluster health.

Cluster Size	Number of Simultaneous Failures	Entity Failed	Maintenance Action to Take
5+ nodes	2	One node and One or more disks on a different node.	<p>If the disk does not recover in one minute, the storage cluster starts healing by rebalancing data on the remaining nodes.</p> <p>If the node does not recover in two hours, the storage cluster starts healing by rebalancing data on the remaining nodes.</p> <p>If a node in the storage cluster fails and a disk on a different node also fails, the storage cluster starts healing the failed disk (without touching the data on the failed node) in one minute. If the failed node does not come back up after two hours, the storage cluster starts healing the failed node as well.</p> <p>To recover the failed node immediately and fully restore the storage cluster:</p> <ol style="list-style-type: none"> 1. Check that the node is powered on and restart it if possible. You might need to replace the node. 2. Rebalance the cluster.

Review the table above and perform the action listed.

HX Data Platform Ready Clones Overview

HX Data Platform Ready Clones is a pioneer storage technology that enables you to rapidly create and customize multiple cloned VMs from a host VM. It enables you to create multiple copies of VMs that can then be used as standalone VMs.

A Ready Clone, similar to a standard clone, is a copy of an existing VM. The existing VM is called the host VM. When the cloning operation is complete, the Ready Clone is a separate guest VM.

Changes made to a Ready Clone do not affect the host VM. A Ready Clone's MAC address and UUID are different from that of the host VM.

Installing a guest operating system and applications can be time consuming. With Ready Clone, you can make many copies of a VM from a single installation and configuration process.

Clones are useful when you deploy many identical VMs to a group.

**Note**

The Hyper-V Host issues an Offloaded Data Transfer (ODX) request to create a full file copy. HyperFlex uses cloning technology to create an instantaneous copy as a response to the full copy request.

HX Native Snapshots Overview

HX native snapshots are a backup feature that saves versions (states) of VMs. VMs can be reverted back to a prior saved version using an HX native snapshot. A native snapshot is a reproduction of a VM that includes the state of the data on all VM disks and the VM powerstate (on, off, or suspended) at the time the native snapshot is taken. Taking a native snapshot to save the current state of a VM gives you the ability to revert back to the saved state.

The following methodologies are used in the administration of HX native Snapshots:

- The HX vSphere web client plug-in for FLEX mode facilitates the creation of on-demand HX native snapshots, as well as the creation of HX native snapshot schedules. The HX vSphere client plug-in for HTML 5 mode does not provide these capabilities.
- The vSphere “Manage Snapshots” function can revert to a specific HX native snapshot, or delete all snapshots.
- Cisco HyperFlex Connect can create on-demand HX native snapshots.
- The HyperFlex command line user interface can create HX native snapshots.
- HX REST APIs can create and delete HX native snapshots.

For additional information about VMware snapshots, see the VMware KB, Understanding virtual machine snapshots in VMware ESXi and ESX (1015180) at, http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1015180



CHAPTER 3

Logging in to HX Data Platform Interfaces

- [HyperFlex Cluster Interfaces Overview, on page 19](#)
- [AAA Authentication REST API, on page 24](#)
- [Logging into HX Connect, on page 24](#)
- [Logging into the Controller VM \(stcli\) Command Line, on page 25](#)
- [Logging Into Cisco HX Data Platform Installer, on page 27](#)
- [Accessing the HX Data Platform REST APIs, on page 28](#)
- [Secure Admin Shell, on page 29](#)

HyperFlex Cluster Interfaces Overview

Each HyperFlex interface provides access to information about and a means to perform actions upon the HX Storage Cluster. The HX Storage Cluster interfaces include:

- HX Connect—Monitoring, performance charts, and tasks for upgrade, encryption, replication, datastores, nodes, disks, and VM ready clones.
- HX Data Platform Plug-in—Monitoring, performance charts, and tasks for datastores, hosts (nodes), and disks.
- Admin Shell command line—Run HX Data Platform `hxcli` commands.
- HyperFlex Systems RESTful APIs—Enabling authentication, replication, encryption, monitoring, and management of HyperFlex Systems through an on-demand stateless protocol.

Additional interfaces include:

- HX Data Platform Installer—Installing HX Data Platform, deploying and expanding HX Storage Cluster, deploying stretched cluster, and deploying Hyper-V clusters.
- Cisco UCS Manager—Tasks for networking, storage, and storage access, and managing resources in the HX Storage Cluster.
- VMware vSphere Web Client and vSphere Client—Managing all the VMware ESXi servers in the vCenter cluster.
- VMware ESXi —Managing the individual ESXi host, providing host command line.

Guidelines for HX Data Platform Login Credentials

`hxcli` commands prompt for login credentials.

The Admin Shell password for the predefined users `admin` and `root` are specified during HX Data Platform installer. After installation you can change passwords through the `hxcli` command line.

When a user attempts to login with wrong credentials for 10 successive times, the account will be locked for two minutes. If the failed login attempts were made through SSH, the error message will not indicate that the account is locked. If the failed login attempts were made through HX Connect or REST API, the error message during the 10th attempt will indicate that the account is locked.

Component	Permission Level	Username	Password	Notes
HX Data Platform Installer VM	root	root	Cisco123	Important Systems ship with a default password of <code>Cisco123</code> that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.
HX Connect	administrator or read-only	User defined through vCenter.	User defined through vCenter.	
		Predefined <code>admin</code> or <code>root</code> users.	As specified during HX installation.	
Admin Shell		User defined during HX installation. Predefined <code>admin</code> or <code>root</code> users.	As specified during HX installation. Strong password required.	Must match across all nodes in storage cluster. Use the <code>hxcli</code> command when changing the password after installation.
vCenter	admin	<code>administrator@vsphere.local</code> default. SSO enabled. As configured, <code>MYDOMAIN\name</code> or <code>name@mydomain.com</code>	SSO enabled. As configured.	Ensure the vCenter credentials meet the vSphere 5.5 requirements if the ESX servers are at version 5.5. Read only users do not have access to HX Data Platform Plug-in.

Component	Permission Level	Username	Password	Notes
ESXi Server	root	SSO enabled. As configured.	SSO enabled. As configured.	Must match across all ESX servers in storage cluster.
Hypervisor	root	root	As specified during HX installation.	Use vCenter or <code>esxcli</code> command when changing the password after HX installation.
UCS Manager	admin	As configured.	As configured.	
Fabric Interconnect	admin	As configured.	As configured.	

HX Data Platform Names, Passwords, and Characters

Most printable and extended ASCII characters are acceptable for use in names and passwords. Certain characters are not allowed in HX Data Platform user names, passwords, virtual machine names, storage controller VM names, and datastore names. Folders and resource pools do not have character exceptions.

Passwords must contain a minimum of 10 characters, with at least 1 lowercase, 1 uppercase, 1 numeric, and 1 of the following characters:

ampersand (&), apostrophe ('), asterisk (*), at sign (@), back slash (\), colon (:), comma (,), dollar sign (\$), exclamation (!), forward slash (/), less than sign (<), more than sign (>), percent (%), pipe (|), pound (#), question mark (?), semi-colon (;)

When entering special characters, consider the shell being used. Different shells have different sensitive characters. If you have special characters in your names or passwords, place them in a single quote, 'speci@lword!'. It is not required to place passwords within single quotes in the HyperFlex Installer password form field.

HX Storage Cluster Name

HX cluster names cannot exceed 50 characters.

HX Storage Cluster Host Names

HX cluster host names cannot exceed 80 characters.

Virtual Machine and Datastore Names

Most characters used to create a virtual machine name, controller VM name, or datastore name are acceptable. Escaped characters are acceptable for virtual machine, controller VM names, or datastore names.

Maximum characters—Virtual machine names can have up to 80 characters.

Excluded characters—Do not use the following character in any user virtual machine name or datastore name for which you want to enable snapshots.

- accent grave (`)

Special characters—The following special characters are acceptable for user virtual machine or datastore names:

- ampersand (&), apostrophe ('), asterisk (*), at sign (@), back slash (\), circumflex (^), colon (:), comma (,), dollar sign (\$), dot (.), double quotation ("), equal sign (=), exclamation (!), forward slash (/), hyphen (-), left curly brace {}, left parentheses (), left square bracket ([), less than sign (<), more than sign (>), percent (%), pipe (|), plus sign (+), pound (#), question mark (?), right curly brace {}), right parentheses ()), right square bracket (]), semi-colon (;), tilde (~), underscore (_)

Username Requirements

Username can be specific to the HX Data Platform component and must meet UCS Manager username requirements.

UCS Manager username requirements.

- Number of characters: between 6 and 32 characters
- Must be unique within Cisco UCS Manager.
- Must start with an alphabetic character.
- Must have alphabetic characters (upper or lower case).
- Can have numeric characters. Cannot be all numeric characters.
- Special characters: Limited to underscore (_), dash (-), and dot (.)

Controller VM Password Requirements

The following rules apply to controller VM root and admin user passwords.



Note

General rule about passwords: Do not include them in a command string. Allow the command to prompt for the password.

- Minimum Length: 10
- Minimum 1 Uppercase
- Minimum 1 Lowercase
- Minimum 1 Digit
- Minimum 1 Special Character
- A maximum of 3 retry to set the new password

To change a controller VM password, always use the `stcli` command. Do not use another change password command, such as a Unix password command.

1. Login to the management controller VM.
2. Run the `stcli` command.

`stcli security password set [-h] [--user USER]`

The change is propagated to all the controller VMs in the HX cluster.

UCS Manager and ESX Password Format and Character Requirements

The following is a summary of format and character requirements for UCS Manager and VMware ESXi passwords. See the Cisco UCS Manager and VMware ESX documentation for additional information.

- **Characters classes:** lower case letters, upper case letters, numbers, special characters.

Passwords are case sensitive.

- **Character length:** Minimum 6, maximum 80

Minimum 6 characters required, if characters from all four character classes.

Minimum 7 characters required, if characters from at least three character classes.

Minimum 8 characters required, if characters from only one or two character classes.

- **Start and end characters:** An upper case letter at the beginning or a number at the end of the password do not count toward the total number of characters.

If password starts with uppercase letter, then 2 uppercase letters are required. If password ends with a digit, then 2 digits are required.

Examples that meet the requirements:

h#56Nu - 6 characters. 4 classes. No starting upper case letter. No ending number.

h5xj7Nu - 7 characters. 3 classes. No starting upper case letter. No ending number.

XhUwPcNu - 8 characters. 2 classes. No starting upper case letter. No ending number.

Xh#5*Nu - 6 characters counted. 4 characters classes. Starting upper case letter. No ending number.

h#5*Nu9 - 6 characters counted. 4 characters classes. No starting upper case letter. Ending number.

- **Consecutive characters:** Maximum 2. For example, hhh###555 is not acceptable.

Through vSphere SSO policy, this value is configurable.

- **Excluded characters:**

UCS Manager passwords cannot contain the escape (\) character.

ESX passwords cannot contain these characters.

- Cannot be the username or the reverse of the username.
- Cannot contain words found in the dictionary.
- Cannot contain the characters escape (\), dollar sign (\$), question mark (?), equal sign (=).

- **Dictionary words:**

Do not use any words that can be found in the dictionary.

AAA Authentication REST API

Cisco HyperFlex provides REST APIs to access resources in storage cluster. The AAA Authentication REST API provides a mechanism to authenticate a user and exchange the provided credentials for an Access Token. This access token can be used to invoke other REST API calls.

A rate limit is enforced on Authentication REST API (/auth): in a 15 minute window, /auth can be invoked (successfully) a maximum of 5 times. A user is allowed to create a maximum of 8 unrevoked tokens. Subsequent call to /auth will automatically revoke the oldest issued token to make room for the new token. A maximum of 16 unrevoked tokens can be present in system. In order to prevent brute-force attacks, after 10 consecutive failed authentication attempts, a user account is locked for a period of 120 seconds. Access Tokens issued are valid for 18 days (1555200 second).

**Note**

HxConnect makes use of /auth call for login purpose and the limit applies there also.

Logging into HX Connect

Cisco HyperFlex Connect provides an HTML5 based access to HX Storage Cluster monitoring, and replication, encryption, datastore, and virtual machine tasks.

About Sessions

Each login to HX Connect is a session. Sessions are the period of activity between time when you log into HX Connect and when you log out. Do not manually clear cookies in a browser during a session, because this also drops the session. Do not close the browser to close a session, though dropped, the session is still counted as an open session. Default session maximums include:

- 8 concurrent sessions per user
- 16 concurrent sessions across the HX Storage Cluster.

Before you begin

**Important**

- If you are a read-only user, you may not see all of the options described in the Help. To perform most actions in HX Connect, you must have administrative privileges.
- Ensure that the time on the vCenter and the controller VMs are in sync or near sync. If there is too large of a time skew between the vCenter time and the cluster time, AAA authentication will fail.

Step 1 Locate the HX Storage Cluster management IP address.

Use fully qualified domain name (FQDN) for the management IP address, rather than individual Storage Controller VM.

Step 2 Enter the HX Storage Cluster management IP address in a browser.

Step 3 Enter the HX Storage Cluster login credentials.

- **RBAC users**—Cisco HyperFlex Connect supports role-based access control (RBAC) login for:
 - **Administrator**—Users with administrator role have read and modify operations permissions. These users can modify the HX Storage Cluster
 - **Read only**—Users with read only role have read (view) permissions. They cannot make any changes to the HX Storage Cluster.

These users are created through vCenter. vCenter username format is: <name>@domain.local and specified in the User Principal Name Format (UPN). For example, administrator@vsphere.local. Do not add a prefix such as "ad:" to the username.

- **HX pre-defined users**—To login using the HX Data Platform predefined users `admin` or `root`, enter a prefix `local/`. For example: `local/root` or `local/admin`.

Actions performed with the `local/` login only affect the local cluster.

vCenter recognizes the session with HX Connect, therefore system messages that originate with vCenter might indicate the session user instead of `local/root`. For example, in Alarms, Acknowledged By might list `com.springpath.sysmgmt.domain-c7`.

Click the eye icon to view or hide the password field text. Sometimes this icon is obscured by other field elements. Click the eye icon area and the toggle function continues to work.

What to do next

- To refresh the HX Connect displayed content, click the refresh (circular) icon. If this does not refresh the page, the clear the cache and reload the browser.
- To logout of HX Connect, and properly close the session, select **User** menu (top right) > **Logout**.

Logging into the Controller VM (stcli) Command Line

All `stcli` command are divided into commands that read HX Cluster information and commands that modify the HX Cluster.

- **Modify commands**—Require administrator level permissions. Examples:

```
stcli cluster create
stcli datastore create
```

- **Read commands**—Permitted with administrator or read only level permissions. Examples:

```
stcli <cmd> -help
stcli cluster info
stcli datastore info
```

To execute HX Data Platform `stcli` commands, login to the HX Data Platform Storage Controller VM command line.

**Important**

Do not include passwords in command strings. Commands are frequently passed to the logs as plain text. Wait until the command prompts for the password. This applies to login commands as well as `stcli` commands.

You may login to the HX Data Platform command line interface in the Storage Controller VM in the following ways:

- From a command terminal
- From HX Connect Web CLI page

Only direct commands are supported through HX Connect.

- Direct commands—commands that complete in a single pass and do not require responses through the command line. Example direct command: `stcli cluster info`
- Indirect commands—multi-layered commands that require live response through the command line. Example interactive command: `stcli cluster reregister`

Step 1

Locate a controller VM DNS Name.

- Select a **VM > Summary > DNS Name**.
- From vSphere Web Client **Home > VMs and Templates > vCenter server > *datacenter* > ESX Agents > VVM**.
- Click through to the storage cluster list of controller VMs.

Step 2

From a browser, enter the DNS Name and `/cli` path.

- Enter the path.

Example

```
# cs002-stctlvm-a.eng.storvisor.com/cli
```

Assumed username: `admin`, password: defined during HX Cluster creation.

- Enter the password at the prompt.

Step 3

From a command line terminal using `ssh`.

Note Do not include the password in a `ssh` login string. The login is passed to the logs as plain text.

- Enter the `ssh` command string.
- Sometimes a certificate warning is displayed. Enter `yes` to ignore the warning and proceed.

```
-----
!!! ALERT !!!
This service is restricted to authorized users only.
All activities on this system are logged. Unauthorized
access will be reported.
-----
HyperFlex StorageController 2.5(1a)# exit
logout
Connection to 10.198.3.22 closed.]$ssh root@10.198.3.24
The authenticity of host '10.198.3.24 (10.198.3.24)' can't be established.
ECDSA key fingerprint is xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx.
Are you sure you want to continue connecting (yes/no)?
```

- c) Enter the password at the prompt.

```
# ssh admin@10.198.3.22
HyperFlex StorageController 2.5(1a)
admin@10.198.3.22's password:
```

Step 4 From HX Connect—Log in to HX Connect, select **Web CLI**.

Note Only non-interactive commands can be executed from the HX Connect Web CLI.

Changing Storage Controller Password

To reset the HyperFlex storage controller password post-installation, do the following.

Step 1 Log in to a storage controller VM.

Step 2 Change the Cisco HyperFlex storage controller password.

```
# stcli security password set
```

This command applies the change to all the controller VMs in the storage cluster.

Note If you add new compute nodes and try to reset the cluster password using the **stcli security password set** command, the converged nodes get updated, but the compute nodes may still have the default password. To change the compute node password, use the following procedure.

To change the password on compute nodes:

- a. Vmotion all the user VMs off the ESXi hosts.
- b. Launch the storage controller VM console from vCenter and log in as the root user.
- c. Run the **passwd** command to change the password.
- d. Log out and re-login to confirm that the password changed successfully.
- e. Run the **stcli node add -f** command to add the node back into the cluster.

Step 3 Type in the new password.

Step 4 Press **Enter**.

Logging Into Cisco HX Data Platform Installer

Next, you install the HX Data Platform software.



Note Before launching the Cisco HX Data Platform Installer, ensure that all the ESXi servers that are in the vCenter cluster that you plan to include in the storage cluster are in maintenance mode.

-
- Step 1** In a browser, enter the URL for the VM where HX Data Platform Installer is installed.
- You must have this address from the earlier section on **Deploying HX Data Platform Installer**. For example *http://10.64.4.254*
- Step 2** Enter the following credentials:
- **Username:** *root*
 - **Password** (Default): Cisco123
- Attention** Systems ship with a default password of `Cisco123` that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.
- Read the EULA. Click **I accept the terms and conditions**.
- Verify the product version listed in the lower right corner is correct. Click **Login**.
- Step 3** The HX Data Platform Installer Workflow page provides two options to navigate further.
- **Create Cluster** drop-down list—You can deploy a standard cluster, Stretched Cluster, or a Hyper-V cluster.
 - **Cluster Expansion**—You can provide the data to add converged nodes and compute nodes to an existing standard storage cluster.
-

Accessing the HX Data Platform REST APIs

Cisco HyperFlex HX-Series Systems provide a fully-contained, virtual server platform that combines all three layers of compute, storage, and network with the powerful Cisco HX Data Platform software tool resulting in a single point of connectivity for simplified management. Cisco HyperFlex Systems are modular systems designed to scale out by adding HX nodes under a single UCS management domain. The hyperconverged system provides a unified pool of resources based on your workload needs.

Cisco HyperFlex Systems RESTful APIs with HTTP verbs integrate with other third-party management and monitoring tools that can be configured to make HTTP calls. It enables authentication, replication, encryption, monitoring, and management of a HyperFlex system through an on-demand stateless protocol. The APIs allow for external applications to interface directly with the HyperFlex management plane.

These resources are accessed through URI or Uniform Resource Identifier and operations are performed on these resources using http verbs such as POST (create), GET (read), PUT (update), DELETE (delete).

The REST APIs are documented using swagger which can also generate client libraries in various languages such as python, JAVA, SCALA, and Javascript. Using libraries thus generated, you can create programs and scripts to consume HyperFlex resources.

HyperFlex also provides a built-in REST API access tool, the REST explorer. Use this tool to access HyperFlex resources in real time and observe responses. The REST explorer also generates CURL commands that can be run from command line.

-
- Step 1** Open a browser to the DevNet address <https://developer.cisco.com/docs/ucs-dev-center-hyperflex/>.

Step 2 Click **Login** and enter credentials, if needed.

Secure Admin Shell

Starting with Cisco HX Release 4.5(1a), limiting access provides the following:

- Controller VMs from outside the clusters through remote **root** access over SSH is disabled.
- Admin users have limited shell access with only restricted commands available. To know the allowed commands in the admin shell, execute **priv** and **help** or **?** commands.
- Access is only available through local **root** Consent Token process.
- Logging in to the root shell of a controller, for troubleshooting purposes, requires Cisco TAC to be involved.

Guidelines and Limitations

- Remote root access over ssh to any controller VM from outside the cluster is disabled. Only nodes part of the cluster can SSH as root to other nodes over the data network.
- If an ESX node is put in Maintenance Mode (MM) during or before consent token generation, the token will not be available on that SCVM and the sync utility will have to be invoked after the node exists MM and the SCVM is back online.
- If any root capable user exists in a release earlier than a HX release 4.5(1a) cluster than it needs to be deleted to upgrade to HX release 4.5(1a), otherwise the upgrade will not proceed.

Information About Consent Token

Consent Token is a security feature that is used to authenticate the system network administrator of an organization to access system shell with mutual consent from the administrator and Cisco Technical Assistance Centre (Cisco TAC).

In some debugging scenarios, the Cisco TAC engineer may have to collect certain debug information or perform live debug on a production system. In such cases, the Cisco TAC engineer will ask you (the network administrator) to access system shell on your device. Consent Token is a lock, unlock and re-lock mechanism that provides you with privileged, restricted, and secure access to the system shell.

For Secure Shell limited access, it is necessary for the network administrator and Cisco TAC to provide explicit consent. When logged in as admin, there is the option to run diagnostic commands as admin or request TAC assistance to request a **root** shell. **root** shell access is only intended to troubleshoot and fix issues within HyperFlex Data Platform.

Once TAC has completed the required troubleshooting, it's recommended to invalidate the consent token to disable the root access.



CHAPTER 4

Monitoring HX Storage Clusters

- [Monitoring HyperFlex Clusters, on page 31](#)
- [Monitoring HyperFlex Clusters with HX Connect, on page 31](#)
- [Using the Cisco HX Data Platform Plug-in Interface, on page 41](#)
- [Monitoring Performance Charts, on page 42](#)
- [Audit Logging with HX Connect, on page 46](#)

Monitoring HyperFlex Clusters

This chapter describes the monitoring content available through the following HX Storage Cluster interfaces:

- Cisco HX Connect
- Cisco HX Data Platform Plug-in
- Storage Controller VM command line

Monitoring HyperFlex Clusters with HX Connect

The Cisco HX Connect user interface provides a view of the Cisco HX storage cluster status, components, and features, such as encryption and replication.

Key monitoring pages include information about the local Cisco HX storage cluster:

- **Dashboard**—Overall Cisco HX storage cluster status.
- **Alarms, Events, Activity**—See the [Cisco HyperFlex Systems Troubleshooting Reference Guide](#) for details.
- **Performance**—Charts for IOPS, throughput, latency, and replication network bandwidth.
- **System Information**—HX storage cluster system-related information, including node and disk information, and access the HX maintenance mode.

See the [Cisco HyperFlex Systems Troubleshooting Reference Guide](#) for generating support bundles, [Storage Cluster Maintenance Operations Overview, on page 69](#) for entering and exiting maintenance mode, and [Setting a Beacon, on page 71](#) to set a node or disk beacon.
- **Datastores**—Status and tasks related to datastores.

- **Virtual Machines**—Status and tasks related to protecting virtual machines.

Additional Cisco HX Connect pages provide management access:

- **Encryption**—For data at rest disk and node encryption tasks.
- **Replication**—For disaster recovery VM protection tasks.

The **Upgrade** page provides access to HX Data Platform and Cisco UCS Manager firmware upgrade tasks.

Dashboard Page



Important

If you are a read-only user, you may not see all of the options available in the Help. To perform most actions in HyperFlex (HX) Connect, you must have administrative privileges.

Displays a status summary of your HX storage cluster. This is the first page that you see when you log in to Cisco HyperFlex Connect.

UI Element	Essential Information
Operational Status section	Provides the functional status of the HX storage cluster and application performance. Click Information (i) to access the HX storage cluster name and status data.
Cluster License Status section	Displays the following link when you log into the HX storage cluster for the first time or till the HX storage cluster license is registered: Cluster License not registered link—Appears when the HX storage cluster is not registered. To register a cluster license, click this link and provide product instance registration token in the Smart Software Licensing Product Registration screen. For more information on how to get a product instance registration token, refer the Registering a Cluster with Smart Licensing section in the Cisco HyperFlex Systems Installation Guide for Microsoft Hyper-V .
Resiliency Health section	Provides the data health status and ability of the HX storage cluster to tolerate failures. Click Information (i) to access the resiliency status, and replication and failure data.
Capacity section	Displays a breakdown of the total storage versus how much storage is used or free. Also displays the storage optimization, compression-savings, and deduplication percentages based on the data stored in the cluster.

UI Element	Essential Information
Nodes section	Displays the number of nodes in the HX storage cluster, and the division of converged versus compute nodes. Hovering over a node icon displays that node's name, IP address, node type, and an interactive display of disks with access to capacity, usage, serial number, and disk type data.
Performance section	Displays an HX storage cluster performance snapshot for a configurable amount of time, showing IOPS, throughput, and latency data. For full details, see Performance Page .
Cluster Time field	System date and time for the cluster.

Table Header Common Fields

Several tables in HX Connect provide one or more of the following three fields that affect the content displayed in the table.

UI Element	Essential Information
Refresh field and icon	The table automatically refreshes for dynamic updates to the HX Cluster. The timestamp indicates the last time the table was refreshed. Click the circular icon to refresh the content now.
Filter field	Display in the table only list items that match the entered filter text. The items listed in the current page of the table below are automatically filtered. Nested tables are not filtered. Type in the selection text in the Filter field. To empty the Filter field, click the x . To export content from other pages in the table, scroll to the bottom, click through the page numbers, and apply the filter.
Export menu	Save a copy of the current page of table data. The table content is downloaded to the local machine in the selected file type. If the listed items are filtered, the filtered subset list is exported. Click the down arrow to select an export file type. The file type options are: <code>cvs</code> , <code>xls</code> , and <code>doc</code> . To export content from other pages in the table, scroll to the bottom, click through the page numbers, and apply the export.

Activity Page

Displays a list of recent activity on the HX storage cluster allowing you to monitor the progress of VM operations, Cluster upgrade/expansion, enter/exit maintenance mode, and recovery jobs.

UI Element	Essential Information
Activity list	<p>Displays a list of recent tasks including the following details:</p> <ul style="list-style-type: none"> • ID • Description • VM power on/off/suspend status • Task status: <ul style="list-style-type: none"> • In Progress • Success • Failed <p>For failed VM-power operations, the Existing State and Required State fields are also included.</p> • Date and time stamp • Progress bar <p>An expanded list shows the task's step name and status.</p> <p>Click the circular icon to refresh the content and fetch recent activity. The page refreshes automatically every 2 minutes.</p>
Recovery list	<p>Displays progress of all recovery-related jobs (for example, migration, recovery, test recovery, re-protect) including the following details:</p> <ul style="list-style-type: none"> • ID • Description • Task status: <ul style="list-style-type: none"> • In Progress • Success • Failed • Date and time stamp • Progress bar <p>An expanded list shows the task's step name and status.</p> <p>Click the circular icon to refresh the content and fetch recent activity. The page refreshes automatically every 2 minutes.</p>
Expand All / Collapse All button	<p>Toggles the view of the job list to display top-level task information or task details.</p> <p>You can also expand and collapse individual tasks.</p>

System Information Overview Page

Displays HX storage cluster system-related information, including node and disk data, and provides access to HX maintenance mode.

HX Storage Cluster Configuration Data

Displays the basic configuration information for this HX storage cluster.

UI Element	Essential Information
HX storage cluster field	Name of the storage cluster.
Cluster License Status section	<p>Displays the Register Now link when you log into the HX storage cluster for the first time or till the HX storage cluster license is registered:</p> <p>Register Now link—To register a cluster license, click this link and provide product instance registration token in the Smart Software Licensing Product Registration screen. For more information on how to get a product instance registration token, refer the Registering a Cluster with Smart Licensing section in the Cisco HyperFlex Systems Installation Guide for VMware ESXi.</p> <p>Note To register a cluster license, you can also choose Register Now from the Actions drop-down field.</p>
License section	<ul style="list-style-type: none"> • License Type—Displays Evaluation, Edge, Standard, or Enterprise as the HX storage cluster license type. • License Status—Displays one of the following as the HX storage cluster license status: <ul style="list-style-type: none"> • In compliance • License expires in <n> days. Cluster not registered - Register Now. (This status appears only for Evaluation type license) • License expired. Cluster not registered - Register Now. (This status appears only for Evaluation type license) • Out of compliance - Insufficient license • Authentication expired—This status appears when HX is unable to communicate with Cisco Smart Software Manager or Smart Software Manager satellite for more than 90 days. <p>Note To refresh license certificate or renew license authorization, choose the respective options from the Actions drop-down field.</p>

UI Element	Essential Information
HX storage cluster status field	Provides functional status of the HX storage cluster. <ul style="list-style-type: none"> • Online—Cluster is ready. • Offline—Cluster is not ready. • Read Only—Cluster is out of space. • Unknown—Transitional state while the cluster is coming online.
vCenter link	Secure URL to the VMware vSphere associated with this HX storage cluster. Click the link to remotely access the vSphere Web Client.
Hypervisor field	Hypervisor version installed on this HX storage cluster.
HXDP Version field	Installer package version installed on this HX storage cluster.
Data Replication Factor field	Number of the redundant data replicas stored on this HX storage cluster.
Uptime field	Length of time this HX storage cluster has been online.
Total Capacity field	Overall storage size of this cluster.
Available Capacity field	Amount of free storage in this cluster.
DNS Server(s)	IP address for the DNS server(s) for this HX storage cluster.
NTP Server(s)	IP address for the NTP server(s) for this HX storage cluster.

Controller VM Access

Use **Actions** to access the controller VM using SSH as an administrator and perform actions such as **Enable Controller Access over SSH**, **Disable Controller Access over SSH** or register your license.



Note Actions to enable or disable SSH can only be performed by **domain** users, and not local users. Domain users are users in VC (ESXi) and AD (Hyper-V).

UI Element	Essential Information
Disable Controller Access over SSH	Secure Shell (SSH) is disabled by default.
Register License	Register your license.
Re-register vCenter	Re-register your license via vCenter

Disk View Options

Customize your Disk View display. Use the check box list to select and deselect the fields that appear in the Node Data section.

Disk View Legend

To display the Disk Legend icons and descriptions, click on **Disk View Legend**.

Node Data

Displays data about individual nodes in this HX storage cluster. To see this information in tabular format, go to the **Nodes** page.

UI Element	Essential Information
Node	Name of a node on this cluster.
Model	Physical hardware model number of this node.
Disks	Number of caching versus persistent disks in this node.
Node status	<ul style="list-style-type: none"> • Online • Offline • In Maintenance • Healthy • Warning
HXDP Version	HyperFlex Data Platform version installed on this cluster.
Type	<ul style="list-style-type: none"> • Hyperconverged • Compute
Hypervisor Status	<ul style="list-style-type: none"> • Online • Offline • In Maintenance • In Progress
Hypervisor Address	IP address for the management network to this HX storage cluster.
Disk Overview	<p>Graphic representation of the number of disks in use for each node, the usage type and number of empty slots.</p> <p>Note A disk outline with a red icon indicates a disk that is not recognized and requires a Catalog Upgrade.</p>

For nodes with disks, you can place your cursor over a disk to view an interactive display of information including the following.

Disks

UI Element	Essential Information
Slot Number	Location of the drive, for example Slot Number 2.
Type of Disk	<ul style="list-style-type: none"> • System • Cache • Persistent
Disk State	<ul style="list-style-type: none"> • Claimed • Available • Ignored • Blacklisted • Ok to Remove • Unknown
Locator LED	Activates a physical light on the host to help locate a disk; options are On and Off .
Capacity	Total disk size.
Used / Total Capacity (Persistent Disks only)	Amount of the disk used versus the total disk size.
Serial Number	Physical serial number of this disk.
Storage Usage (Persistent Disks only)	Percentage of disk storage used.
Version	Version of the disk drive.
Disk Drive Interface	The disk drive interface type, for example SAS or SATA.

Nodes Page

Displays data about all of the nodes in this HX storage cluster in a table. Each column can be used to sort the data.

UI Element	Essential Information
Enter HX Maintenance Mode button	<p>Select a node to access this button.</p> <p>Opens the Confirm HX Maintenance Mode dialog box.</p>
Exit HX Maintenance Mode button	<p>Select a node to access this button.</p> <p>After you complete any maintenance tasks, you must manually exit HX maintenance mode.</p>

UI Element	Essential Information
Node column	Name of a node in this HX storage cluster.
Hypervisor Address column	IP address for the management network of the Node referred in the Node column.
Hypervisor Status column	<ul style="list-style-type: none"> • Online—Node is available. • Offline—Node is not available. • In Maintenance—The running (and powered off) node is disconnected from the host. • In Progress—a backup job is in progress.
Controller Address column	IP address for the HX storage controller VM of the Node referred in the Node column.
Controller Status column	<ul style="list-style-type: none"> • Online—The connection between the VM and the disk is available. • Offline—The connection between the VM and the disk is not available. • In Maintenance—the connection between the VM and the disk is powered off from the host.
Model column	Physical hardware model number of this node.
Version column	HyperFlex Data Platform installer package version installed on this node.
Disks column	<p>Number of disks in the node.</p> <p>Click the number to open the Disks page filtered by the selected node name.</p>

Disks Page

Displays data about all of the disks in this HX storage cluster in a 7-column table. Each column can be used to sort the data.

UI Element	Essential Information
Node column	Name of the node where the disk resides.
Slot column	Location of the SED drive. This identifies the drive for maintenance procedures.
Capacity column	Total disk size.

UI Element	Essential Information	
Status column	<ul style="list-style-type: none"> • Claimed—State when a disk is recognized and in use. • Available—Initial state for a newly added, data-at-rest capable disk. Also, a transitional state when disks move into one of the other states. • Ignored—State when a disk is not being consumed by the cluster; for example, the HX controller VM system disk, a disk with other data (valid file system partitions), or a disk where the IO is failing. • Blacklisted—State when a disk is not being consumed by the cluster due to either a software error or an IO error. This could be a transitional state while the cluster attempts to repair the disk, if the disk is still available, before the state transitions to Repairing. • Ok To Remove—State when an SED disk was securely erased using the Secure Erase option and can safely be removed. • Repairing—State when a blacklisted disk is currently being repaired. • To Be Removed—State when a disk is scheduled for RMA. 	<p>The following states can be ignored:</p> <ul style="list-style-type: none"> • Invalid • Normal • Removed—State when an SED disk is removed after using the Secure Erase option. • Time out • Unknown
Encrypted column	<ul style="list-style-type: none"> • Enabled—Encryption is configured for this data-at-rest-capable disk. • Disabled—Encryption is not configured for this data-at-rest-capable disk. This occurs when a new disk is present, but the Key has not yet been applied. • Locked • Unknown 	
Type column	<ul style="list-style-type: none"> • Unknown • Rotational—Hybrid drive • Solid State—SSD drive 	

UI Element	Essential Information
Usage column	<ul style="list-style-type: none"> • Unknown • Cache • Persistent
Turn On Locator LED and Turn Off Locator LED radio buttons	<p>Select a disk to access the radio buttons.</p> <p>Activates or deactivates a physical light, or beacon, on the host to help locate the disk.</p>
(Optional) Secure erase button	<p>This button is visible only if your HX storage cluster is encrypted using local-key encryption.</p> <p>Select a disk to access the button.</p> <p>Enter the encryption key in use on the cluster, click Secure erase, and then click Yes, erase this disk to securely erase the local encryption key.</p>

Using the Cisco HX Data Platform Plug-in Interface

There are several Cisco HX Data Platform plug-in features that apply across the interface. These are described in the following topics.

Cisco HX Data Platform Plug-in Integration with vSphere Web Client

The Cisco HX Data Platform plug-in is tightly integrated with the VMware vSphere vCenter interface to provide a seamless data management experience. You can use either the vSphere Web Client or the vSphere Client vSphere vCenter interface. Most of the task examples in this guide refer to the vSphere Web Client interface.

The Cisco HX Data Platform plug-in can be accessed through the vSphere vCenter Inventory Lists. Select storage clusters to manage from the Cisco HX Data Platform plug-in. The Cisco HX Data Platform plug-in monitors and manages storage cluster specific objects such as datastores. vSphere monitors and manages objects in the storage cluster, such as ESX servers. Tasks overlap between the Cisco HX Data Platform plug-in and vSphere.



Important

The Cisco HX Data Platform Plug-in is not compatible with the VMware vSphere vCenter HTML5 interface. You cannot perform HX-related tasks such as HX Maintenance mode using the VMware vSphere vCenter HTML5 interface. Use the vSphere Web Client flash interface instead.



Note

HX 3.0 and previous versions supported options to view Support, Summary and Upgrade in vCenter. Starting with HX 3.5, only the Summary option is available.

Links Between the Cisco HX Data Platform Plug-in and the vSphere Interface

In the vSphere Web Client, both the Cisco HX Data Platform plug-in and vCenter provide information on component and cluster status. Selected tabs and panels provide direct links between Cisco HX Data Platform plug-in and vCenter information and actions.

Note that following a link from either the Cisco HX Data Platform plug-in or vCenter does not mean there is a single-click link to return to your starting location.

Cisco HX Data Platform Plug-in Tabs Overview

The Cisco HX Data Platform plug-in monitoring information and managing functions are distributed among three tabs. The following is a list of all the Cisco HX Data Platform plug-in tabs and panels that display Cisco HX Data Platform storage cluster status and provide options for storage cluster administrative tasks.

Summary tab contains a Summary area and a Portlets area. The Summary tab portlets are: Capacity, Performance, and Status.

Monitor tab has two sub tabs:

- Performance tab - Displays Latency, Throughput, and IOPs performance charts for Storage Clusters, Hosts, and Datacenters.
- Events tab - Displays a list Cisco HX Data Platform events and a detail panel for a selected event.

Manage tab has two sub tabs:

- Cluster tab - Describes storage clusters, hosts, disks, PSUs, and NICs. This includes: List of clusters and hosts, detail panels for any selected cluster or host, and additional sub tabs: Hosts, Disks, PSUs, and NICs.
- Datastores tab - Describes information about hosts from the datastore point of view. This includes: List of datastores and additional sub tabs for any selected datastore. The datastore sub tabs include: a Summary tab that includes portlets: Details, Trends, and Top VMs by Disk Usage, and a Hosts tab.

Monitoring Performance Charts

The Monitor Performance tab displays the read and write performance of the storage cluster, hosts, and datastores.

- Performance charts display a pictorial representation of the storage cluster, host, and datastore performance.
- The system updates the performance charts every 20 seconds.
- Hover your mouse over individual data points to view peak performance information and time-stamp.
- Light blue indicates write operations and dark blue indicates read operations.
- Gaps in the performance charts indicate time periods when data was not available. Gaps do not necessarily indicate a drop in performance.

Storage Cluster Performance Chart

You must use HX Connect or HX Plug-in to view storage capacity and not vCenter.

-
- Step 1** From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Monitor > Performance**.
- On the left there are three options you can choose to monitor: Storage Cluster, Hosts, and Datastores.
- Step 2** Click **Storage Cluster** to view the storage cluster performance tab.
- Step 3** Click **Hour, Day, Week, Month, Max, or Custom** option, to specify the time period in which you want to view storage cluster performance.
- Step 4** Click **IOPS, Throughput, Latency, and Show** check boxes to display selected performance and objects.
-

Hosts Performance Chart

-
- Step 1** From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Monitor > Performance**.
- On the left there are three options you can choose to monitor: Storage Cluster, Hosts, and/or Datastores.
- Step 2** Click **Hosts** to view the hosts performance tab.
- Step 3** Click **Hour, Day, Week, Month, Max, or Custom** option, to specify the time period in which you want to view the host performance.
- Step 4** Click **IOPS, Throughput, Latency, and Show** check boxes to display selected performance and objects.
- Step 5** Click **host** to exclude or view individual hosts. Compute nodes do not have storage cluster performance values.
-

Datastores Performance Chart

-
- Step 1** From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Monitor > Performance**.
- On the left there are three options you can chose to Monitor: Storage Cluster, Hosts, and Datastores.
- Step 2** Click **Datastores** to view the datastores performance tab.
- Step 3** Click **Hour, Day, Week, Month, Max, or Custom** option, to specify the time period in which you want to view the datastore performance.
- Step 4** Click **IOPS, Throughput, Latency, and Show** check boxes to display selected performance and objects.
-

Datastore Trends Portlet

The Datastore Trends portlet is a chart of the IO performance of the selected datastore.

- Step 1** From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Manage**.
- Step 2** Select a **datastore** from the table list. The **Summary** tab updates to display the information for the selected datastore.
- Step 3** Scrolls to view the **Trends** portlet.
- The tab displays IOPS plotted every 20 minutes.
- Hover your mouse over the peak values to obtain color-coded read IOPS and write IOPS.

Customizing Performance Charts

Procedure

	Command or Action	Purpose	
		Customize Item	Description
Step 1	Modify the performance charts to display all or some of the listed options.	Time period	Choose from hour, days, week, month, all, or custom. See Specifying Performance Time Period section in this chapter.
		Cluster objects	Choose from a list of storage clusters, hosts, or datastores.
		Chart type	Choose from IOPS, Throughput, or Latency.
		Show objects	Choose which listed object's data to display. See Selecting Performance Charts section in this chapter.

Specify Performance Time Period

- Step 1** From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Monitor > Performance**
- Step 2** Click one of the following tabs to specify the time period in which you want to view performance of the storage cluster, host, or datastore.

Parameter	Description
Hour	Displays performance in the past hour
Day	Displays performance in the past day

Parameter	Description
Week	Displays performance in the past week
Month	Displays performance in the past month
All	Displays the performance of the storage cluster since it was created
Custom	Select this tab and specify a custom range as described in Specifying Custom Range

Specify Custom Range

- Step 1** From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Monitor > Performance**
- Step 2** Click the **Custom** tab to display the Custom Range dialog box.
- Step 3** Choose a method, for the Custom Range dialog box:
- Click **Last**, type the number of minutes, hours, days, or months. Optionally, use the up or down arrow to increase or decrease the number.
 - Click the drop-down list to specify the minutes, hours, days, weeks, or months.
 - Click **From**, click the calendar icon, and select a date from which you want to start measuring the performance. Click the drop-down list to select a time.
 - Click **To**, click the calendar icon, and select a date up to which you want to start measuring the performance. Click the drop-down list to select a time.
- Step 4** Click **Apply** and then click **OK** to apply your configuration.

Selecting Performance Charts

You can select the performance charts to display for storage clusters, hosts, and datastores.

Select or deselect the check box corresponding to IOPS, Throughput, and Latency at the bottom of the tab to view specific information.

For example, to view only storage cluster IOPS performance:

- From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Monitor > Performance**.
- Click either **Storage Cluster, Hosts, or Datastores** chart set. In a Hosts table, compute nodes do not display IOPS, Throughput, or Latency values, as they do not provide storage to the storage cluster.
- Deselect chart options.

Field	Description
Chart types	Click the check box to select which charts and table columns to view or hide. Options are: <ul style="list-style-type: none"> • IOPS • Throughput • Latency
Show	For each storage cluster, hosts, and datastores, click the check boxes to select the specific object to include or exclude from the charts.
Read/Write	Indicates the color representation in the chart for the read and write values of each object.
Storage Cluster	Names of the storage clusters in the charts.
Hosts	Names of the hosts in the charts. This includes both converged nodes and compute nodes.
Datastores	Names of the datastores in the charts.
IOPS Read/Write	Latest data point for Input/Output Operations per Second.
Throughput Read/Write (Mbps)	Latest data point for the rate of data transfer in the storage cluster (measured in Mbps).
Latency Read/Write (msec)	Latest data point for the Latency that is a measure of how long it takes for a single I/O request to complete. It is the duration between issuing a request and receiving a response. Measured in msec.

Audit Logging with HX Connect

Audit logging implies storing all audit logs to a remote syslog server. Currently, each controller VM stores audit logs, but these logs are not stored indefinitely. The logs are overwritten based on the retention policy set for the controller VM. By configuring a remote syslog server to store audit logs, you can ensure that the logs are retained for a longer period of time.

Following are the audit logs that you can export to the remote syslog server:

- REST-related logs
 - /var/log/springpath/audit-rest.log
 - /var/log/springpath/hxmanager.log
 - /var/log/springpath/hx_device_connector.log

- /var/log/shell.log
- /var/log/springpath/stSSOMgr.log
- /var/log/springpath/stcli.log
- /var/log/springpath/hxcli.log
- /var/log/nginx/ssl-access.log

After you enable audit logging, these logs are exported to the remote syslog server. If the logs from the controller VM are not pushed to the remote syslog server, or if the remote syslog server is not reachable, an alarm is generated in the HX-Connect user interface. However, HX Connect does not monitor the disk space available on the remote syslog server. The HX Connect user interface will not display an alarm if the disk on the remote syslog server is full.



Attention

- Only an administrator user can enable audit logging.
- Logs from the compute-only nodes and witness nodes are not pushed to the remote syslog server.

After you enable audit logging, you can choose to either temporarily disable audit logging, or you can choose to delete the audit logging server configuration details.

Enabling Audit Logging

Before you begin

- Configure the remote syslog server. You must have the server details such as the server IP, the port number and certificate files to enable audit logging in HX-Connect.
- To configure an encrypted connection between the controller VM and the remote syslog server, you must generate a self-signed certificate or a CA-signed certificate and a private key for the syslog client in the controller VM.
- Configure the remote syslog server to categorize different types of logs into respective files.

Step 1 Choose **Settings > Audit Log Export Settings**.

Step 2 Check the **Enable audit log export to an external syslog server** check box.

Step 3 Complete the following details:

UI Element	Essential Information
Syslog Server	Enter the IP address of the syslog server.
Port	Enter the port number for the syslog server.

UI Element	Essential Information
Connection Type drop-down list	Choose TLS or TCP as the connection type. The default and recommended value is TLS. The TLS connection type is for encrypted transport over TLS. The TCP connection type is for unencrypted transport over TCP.
Client Certificate	Click Choose to search and locate a certificate file that must be stored on the controller VM. This certificate creates a TLS connection between the controller VM and the remote syslog server. A TLS connection ensures that the log files are encrypted. You must upload either a user-generated self-signed certificate or a CA-signed certificate.
Private Key	Click Choose to search and locate a generated private key file to be stored on the controller VM. This key creates a TLS connection between the controller VM and the remote syslog server. Choosing a certificate and private key for the syslog server ensures that the log files are encrypted. The certificate for the syslog server can either be a CA certificate or a self-signed certificate.
Are you using a self-signed certificate?	Check this check box if the syslog server uses a self-signed certificate. Click Choose to search and locate the self-signed certificate for the syslog server.

Step 4 Click **OK**.

Configuring the Remote Syslog Server

Prior to enabling audit logging, you must create a configuration file on the remote syslog server to categorize different log files into separate files. You could create a file titled `hx-audit.conf` in the `/etc/syslog-ng/conf.d` directory.

Following is a sample of the configuration file to establish an encrypted connection with the syslog server:

```
## Audit Logging Configuration ###
source demo_tls_src {
    tcp(ip(0.0.0.0) port(6515)
    tls(
        key-file("/etc/syslog-ng/CA/serverkey.pem")
        cert-file("/etc/syslog-ng/CA/servercert.pem")
        peer-verify(optional-untrusted)
    )
}; };

filter f_audit_rest { match("hx-audit-rest" value("MSGHDR")); };
filter f_device_conn { match("hx-device-connector" value("MSGHDR")); };
```

```

filter f_stssomgr { match("hx-stSSOMgr" value("MSGHDR")); };
filter f_ssl_access { match("hx-ssl-access" value("MSGHDR")); };
filter f_hxmanager { match("hx-manager" value("MSGHDR")); };
filter f_hx_shell { match("hx-shell" value("MSGHDR")); };
filter f_stcli { match("hx-stcli" value("MSGHDR")); };
filter f_hxcli { match("hx-cli" value("MSGHDR")); };

destination d_audit_rest { file("/var/log/syslog-ng/audit_rest.log"); };
destination d_device_conn { file("/var/log/syslog-ng/hx_device_connector.log"); };
destination d_stssomgr { file("/var/log/syslog-ng/stSSOMgr.log"); };
destination d_ssl_access { file("/var/log/syslog-ng/ssl_access.log"); };
destination d_hxmanager { file("/var/log/syslog-ng/hxmanager.log"); };
destination d_hx_shell { file("/var/log/syslog-ng/shell.log"); };
destination d_stcli { file("/var/log/syslog-ng/stcli.log"); };
destination d_hxcli { file("/var/log/syslog-ng/hxcli.log"); };

log { source(demo_tls_src); filter(f_audit_rest); destination(d_audit_rest); flags(final);
};
log { source(demo_tls_src); filter(f_device_conn); destination(d_device_conn);
flags(final); };
log { source(demo_tls_src); filter(f_stssomgr); destination(d_stssomgr); flags(final);
};
log { source(demo_tls_src); filter(f_ssl_access); destination(d_ssl_access); flags(final);
};
log { source(demo_tls_src); filter(f_hxmanager); destination(d_hxmanager); flags(final);
};
log { source(demo_tls_src); filter(f_hx_shell); destination(d_hx_shell); flags(final);
};
log { source(demo_tls_src); filter(f_stcli); destination(d_stcli); flags(final); };
log { source(demo_tls_src); filter(f_hxcli); destination(d_hxcli); flags(final); };

#####

```

Following is a sample of the configuration file to establish a TCP connection with the remote syslog server:

```

#####
## Audit Logging Configuration ##
source demo_tls_src {
    tcp(ip(0.0.0.0) port(6515)
    ); };

filter f_audit_rest { match("hx-audit-rest" value("MSGHDR")); };
filter f_device_conn { match("hx-device-connector" value("MSGHDR")); };
filter f_stssomgr { match("hx-stSSOMgr" value("MSGHDR")); };
filter f_ssl_access { match("hx-ssl-access" value("MSGHDR")); };
filter f_hxmanager { match("hx-manager" value("MSGHDR")); };
filter f_hx_shell { match("hx-shell" value("MSGHDR")); };
filter f_stcli { match("hx-stcli" value("MSGHDR")); };
filter f_hxcli { match("hx-cli" value("MSGHDR")); };

destination d_audit_rest { file("/var/log/syslog-ng/audit_rest.log"); };
destination d_device_conn { file("/var/log/syslog-ng/hx_device_connector.log"); };
destination d_stssomgr { file("/var/log/syslog-ng/stSSOMgr.log"); };
destination d_ssl_access { file("/var/log/syslog-ng/ssl_access.log"); };
destination d_hxmanager { file("/var/log/syslog-ng/hxmanager.log"); };
destination d_hx_shell { file("/var/log/syslog-ng/shell.log"); };
destination d_stcli { file("/var/log/syslog-ng/stcli.log"); };
destination d_hxcli { file("/var/log/syslog-ng/hxcli.log"); };

log { source(demo_tls_src); filter(f_audit_rest); destination(d_audit_rest); flags(final);
};
log { source(demo_tls_src); filter(f_device_conn); destination(d_device_conn);
flags(final); };
log { source(demo_tls_src); filter(f_stssomgr); destination(d_stssomgr); flags(final);
};

```

```

log { source(demo_tls_src); filter(f_ssl_access); destination(d_ssl_access); flags(final);
};
log { source(demo_tls_src); filter(f_hxmanager); destination(d_hxmanager); flags(final);
};
log { source(demo_tls_src); filter(f_hx_shell); destination(d_hx_shell); flags(final);
};
log { source(demo_tls_src); filter(f_stcli); destination(d_stcli); flags(final); };
log { source(demo_tls_src); filter(f_hxcli); destination(d_hxcli); flags(final); };

#####

```

Disabling Audit Logging

You can choose to temporarily disable audit logging. By doing so, the remote syslog server details such as the server IP and the port, that you previously configured are retained in the system. You need not enter the server details again when you re-enable audit logging at a later time. You will only need to upload the certificate and private key files to enable audit logging.

-
- Step 1** Choose **Settings > Audit Log Export Settings**.
 - Step 2** Clear the **Enable audit log export to an external syslog server** check box.
 - Step 3** Click **OK**.
- Audit logging is disabled.
-

Deleting Audit Logging Server Configuration

As an administrator, you can delete the remote syslog server configuration details from the system. When you do so, the system does not push server logs to the remote syslog server. To enable audit logging, you will have to provide the server details again.

-
- Step 1** Choose **Settings > Audit Log Export Settings**.
 - Step 2** Click **Delete**.
 - Step 3** In the **Confirm Delete** dialog box, click **Delete**.
- The remote syslog server details are deleted from the system.
-



CHAPTER 5

Managing HX Storage Clusters

- [Changing the Cluster Access Policy Level, on page 51](#)
- [Rebalancing the Cluster, on page 51](#)
- [Handling Out of Space Errors, on page 53](#)
- [Moving the Storage Cluster from a Current vCenter Server to a New VCenter Server, on page 54](#)
- [Renaming Clusters, on page 61](#)
- [Re-registering the vCenter using HX Connect, on page 61](#)
- [Replacing Self-Signed Certificate, on page 62](#)
- [Boost Mode, on page 64](#)
- [UEFI Secure Boot Mode, on page 66](#)

Changing the Cluster Access Policy Level

Step 1 The storage cluster must be in a healthy state prior to changing the Cluster Access Policy to strict.

Step 2 From the command line of a storage controller VM in the storage cluster, type:

```
# stcli cluster get-cluster-access-policy  
  
# stcli cluster set-cluster-access-policy --name {strict,lenient}
```

Rebalancing the Cluster

The storage cluster is rebalanced on a regular schedule. It is used to realign the distribution of stored data across changes in available storage and to restore storage cluster health. When a new node is added to the existing cluster, the added node(s) take on new writes as soon as it joins the existing cluster. The Cluster automatically rebalances if required (usually within 24 hours) and the new node may initially show less storage utilization than the existing converged nodes if the overall storage utilization is low. If the current storage utilization is high, and once the new node is added to the cluster, data is rebalanced onto the new node drives over a period of time.



Note Forcing a manual rebalance can cause interference with regular User IO on the cluster and increase the latency. Therefore, the HyperFlex system initiates a rebalance only when required in order to minimize performance penalties.

Verify rebalancing status from the storage controller VM.

a) Enter the following on the command line:

```
# stcli rebalance status
rebalanceStatus:
rebalanceState:
cluster_rebalance_ongoing
percentComplete: 10
rebalanceEnabled: True
```

b) Reenter the command line to confirm the process completes:

```
# stcli rebalance status
rebalanceStatus:
rebalanceState: cluster_rebalance_not_running
rebalanceEnabled: True
```

This sample indicates that `rebalance` is enabled, and ready to perform a rebalance, but is not currently rebalancing the storage cluster.

Checking Cluster Rebalance and Self-Healing Status

The storage cluster is rebalanced on a regular schedule and when the amount of available storage in the cluster changes. A rebalance is also triggered when there is a change in the amount of available storage. This is an automatic self-healing function.



Important Rebalance typically occurs only when a single disk usage exceeds 50% or cluster aggregate disk usage is greater than 50%.

You can check rebalance status through the HX Data Platform plug-in or through the storage controller VM command line.

Step 1 Check the rebalance status through HX Data Platform plug-in.

a) From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > *cluster* > Summary**.

The **Status** portlet lists the **Self-healing status**.

b) Expand the 'Resiliency Status' to see the 'Self-healing status' section. The Self-healing status field lists the rebalance activity or N/A, when rebalance is not currently active.

Step 2 Check the rebalance status through the storage controller VM command line.

- a) Login to a controller VM using `ssh`.
- b) From the controller VM command line, run the command.

```
# stcli rebalance status
```

The following output indicates that rebalance is not currently running on the storage cluster.

```
rebalanceStatus:
percentComplete: 0
rebalanceState: cluster_rebalance_not_running
rebalanceEnabled: True
```

The Recent Tasks tab in the HX Data Platform plug-in displays a status message.

Handling Out of Space Errors

If your system displays an Out of Space error, you can either add a node to increase free capacity or delete existing unused VMs to release space.

When there is an Out of Space condition, the VMs are unresponsive.



Note Do not delete storage controller VMs. Storage controller VM names have the prefix `stCtlVM`.

Step 1 To add a node, use the Expand Cluster feature of the HX Data Platform Installer.

Step 2 To delete unused VMs, complete the following:

- a) Determine which guest VMs you can delete. You can consider factors such as disk space used by the VM or naming conventions.
- b) Go to **vCenter > Virtual Machines** to display the virtual machines in the inventory.
- c) Double-click a VM that you want to delete.
- d) Select the **Summary > Answer Questions** to display a dialog box.
- e) Click the **Cancel** radio button and click **OK**.
- f) Power off the VM.
- g) Delete the VM.

Step 3 After the Out of Space condition is cleared, complete the following:

- a) Go to **vCenter > Virtual Machines** to display the VM in the inventory.
- b) Double-click a VM that you want to use.
- c) Select the **Summary > Answer Questions** to display a dialog box.
- d) Click the **Retry** radio button and click **OK**.

Checking Cleaner Schedule

The `stcli cleaner` command typically runs in the background continuously. `cleaner` goes into sleep mode when it is not needed and wakes when policy defined conditions are met. For example, if your storage cluster is experiencing ENOSPC condition, the cleaner automatically runs at High Priority.

Do not expand the cluster while the `cleaner` is running. Check the cleaner schedule or adjust the schedule, as needed.

Step 1 Login to any controller VM in the storage cluster. Run the listed commands from the controller VM command line.

Step 2 View the cleaner schedule.

```
# stcli cleaner get-schedule --id ID | --ip NAME
```

Parameter	Description
--id ID	ID of storage cluster node
--ip NAME	IP address of storage cluster node

Moving the Storage Cluster from a Current vCenter Server to a New VCenter Server

Before you begin

- If your HX Cluster is running HX Data Platform version older than 1.8(1c), upgrade before attempting to reregister to a new vCenter.
- Perform this task during a maintenance window.
- Ensure the cluster is healthy and upgrade state is OK and Healthy. You can view the state using the `stcli` command from the controller VM command line.

```
# stcli cluster info
```

Check response for:

```
upgradeState: ok
healthState: healthy
```

- Ensure vCenter must be up and running.
- Snapshot schedules are not moved with the storage cluster when you move the storage cluster between vCenter clusters.

Step 1 From the current vCenter, delete the cluster.

This is the vCenter cluster specified when the HX storage cluster was created.

- Step 2** On the new vCenter, create a new cluster using the same cluster name.
- Step 3** Add ESX hosts to new vCenter in the newly created cluster.
-

What to do next

Proceed to [Unregistering a Storage Cluster from a vCenter Cluster, on page 55](#).

Moving the Storage Cluster from a Current vCenter Server to a New VCenter Server

Before you begin

- If your HX Cluster is running HX Data Platform version older than 1.8(1c), upgrade before attempting to reregister to a new vCenter.
- Perform this task during a maintenance window.
- Ensure the cluster is healthy and upgrade state is OK and Healthy. You can view the state using the `stcli` command from the controller VM command line.

```
# stcli cluster info
```

Check response for:

```
upgradeState: ok  
healthState: healthy
```

- Ensure vCenter must be up and running.
 - Snapshot schedules are not moved with the storage cluster when you move the storage cluster between vCenter clusters.
-

- Step 1** From the current vCenter, delete the cluster.
- This is the vCenter cluster specified when the HX storage cluster was created.
- Step 2** On the new vCenter, create a new cluster using the same cluster name.
- Step 3** Add ESX hosts to new vCenter in the newly created cluster.
-

What to do next

Proceed to [Unregistering a Storage Cluster from a vCenter Cluster, on page 55](#).

Unregistering a Storage Cluster from a vCenter Cluster

This step is optional and not required. It is recommended to leave the HX Data Platform Plug-in registration alone in the old vCenter.

Before you begin

As part of the task to move a storage cluster from one vCenter server to another vCenter server, complete the steps in [Moving the Storage Cluster from a Current vCenter Server to a New VCenter Server, on page 54](#).

**Note**

- If multiple HX clusters are registered to the same vCenter, do not attempt this procedure until all HX clusters have been fully migrated to different vCenter. Running this procedure is disruptive to any existing HX clusters registered to the vCenter.

Step 1 Complete the steps in [Unregistering and Removing EAM Extensions, on page 56](#).

Note Newly deployed HX clusters starting with HyperFlex Release 4.0(1a) no longer leverage the vSphere ESX Agent Manager (EAM) for the HyperFlex Storage Controller VMs. HX clusters built prior to HX 4.0(1a) will continue to utilize EAM. If that cluster is migrated to a new vCenter, however, the EAM integration will not be configured.

This is the step that removes (unregisters) the HX cluster from the old vCenter server.

Also, if there are more ESX agencies than the number of HX clusters installed on the given vSphere server, it is likely there are stale EAM configurations that need cleanup.

Step 2 Complete the steps in [Removing HX Data Platform Files from the vSphere Client, on page 58](#).

Step 3 Complete the steps in [Verifying HX Cluster is Unregistered from vCenter, on page 59](#).

What to do next

Proceed to [Registering a Storage Cluster with a New vCenter Cluster, on page 59](#).

Unregistering and Removing EAM Extensions

If you have partially installed or uninstalled HX Data Platform, or unregistered a HX cluster where there are more agencies than the number of HX clusters installed on the given vSphere, sometimes a stale ESX Agent Manager (EAM) for the HX Data Platform extension remains. Remove stale extensions using the Managed Object Browser (MOB) extension manager.

Before you begin

- Download the vSphere ESX Agent Manager SDK, if you have not already done so.
- If multiple HX clusters are registered to the same vCenter, do not attempt this procedure until all HX clusters have been fully migrated to a different vCenter. Running this procedure is disruptive to any existing HX clusters registered to the vCenter.
- Remove the datacenter from your vSphere cluster.



Note Newly deployed HX clusters starting with HyperFlex Release 4.0 no longer leverage the vSphere ESX Agent Manager (EAM) for the HyperFlex Storage Controller VMs. HX clusters built prior to HX 4.0 will continue to utilize EAM. If that cluster is migrated to a new vCenter, however, the EAM integration will not be configured.

Step 1

Identify the HX cluster UUID.

Every agency has a field `cluster_domain_id` which refers to the underlying vSphere extension. This extension ID uses a Managed Object ID (moid).

If you have multiple HyperFlex clusters, ensure that you select the correct cluster ID to unregister.

From a storage controller VM command line, run the command:

```
# stcli cluster info | grep vCenterClusterId:
vCenterClusterId: domain-c26
```

Step 2

To unregister the storage cluster extension: Login to the vCenter server MOB extension manager

First unregister the HyperFlex cluster.

- a) In a browser, enter the path and command.

```
https://vcenter_server/mob/?moid=ExtensionManager
```

`vcenter_server` is the IP address of the vCenter where the storage cluster is currently registered.

- b) Enter administrator login credentials.

Step 3

Locate the HX storage cluster extensions with the cluster IDs. Scroll through the **Properties** > **extensionList** to locate the storage cluster extensions:

```
com.springpath.sysgmt.cluster_domain_id and com.springpath.sysgmt.uuid.cluster_domain_id.
```

Copy each of these strings into your clipboard. Exclude the double quotes (") on either end of string, if there are any.

Step 4

Unregister each storage cluster extension.

- a) From the Methods table click `UnregisterExtension`.
- b) In the **UnregisterExtension** popup, enter an extension key value, `com.springpath.sysgmt.cluster_domain_id`.

For example: `com.springpath.sysgmt.domain-26`

- c) Click **Invoke Method**.

Step 5

To remove stale EAM extensions: Login to the vCenter server MOB ESX agencies extension manager.

Second remove stale EAM extensions that were associated with the HyperFlex cluster.

- a) In a browser, enter the path and command.

```
https://vcenter_server/eam/mob/
```

`vcenter_server` is the IP address of the vCenter where the storage cluster is currently registered.

- b) Enter administrator login credentials.

Step 6

Locate the stale HX storage cluster ESX agency extensions with the cluster IDs.

- a) Scroll through the **Properties** > **agency** > **Value**.

- b) Click an agency value.
- c) In the **Agency** window, check the **Properties > solutionID > Value** extension. Verify has the correct *cluster_domain_id*.

For example: `com.springpath.sysgmt.domain-26`

Step 7 Remove stale ESX agency extensions.

- a) From the **Agency** window, **Methods** table select a method.

Stale ESX agencies can be removed using either the `destroyAgency` or `uninstall`.

- b) In the *method* popup, click **Invoke Method**.

Step 8 Refresh the **ExtensionManager** tab and verify that the **extensionList** entry does not include `com.springpath.sysgmt.cluster_domain_id` extensions.

Step 9 Restart the vSphere Client services.

The HX Data Platform extensions are removed when the vSphere Client services are restarted. Restarting the vSphere client service temporarily disables access to vCenter through the browser.

Removing HX Data Platform Files from the vSphere Client

This task is a step in unregistering a HX Storage Cluster from vCenter.

Remove the HX Data Platform files from the vSphere Client. Select a method.

Linux vCenter

- a) Login to the Linux vCenter server using `ssh` as a root user.
- b) Change to the folder containing the HX Data Platform Plug-in folder.

For vCenter 6.0

```
# cd /etc/vmware/vsphere-client/vc-packages/vsphere-client-serenity/
```

For vCenter 5.5

```
# cd /var/lib/just/vmware/vsphere-client/vc-packages/vsphere-client-serenity/
```

- c) Remove the HX Data Platform Plug-in folder and files.

```
# rm -rf com.springpath*
```

- d) Restart the vSphere Client.

```
# service vsphere-client restart
```

Windows vCenter

- a) Login to the Windows vCenter system command line using Remote Desktop Protocol (RDP).
- b) Change to the folder containing the HX Data Platform Plug-in folder.

```
# cd "%PROGRAMDATA%\VMware\VSphere Web Client\vc-packages\vsphere-client-serenity
```

- c) Remove the HX Data Platform Plug-in folder and files.

```
# rmdir /com.springpath*
```

- d) Open the Service screen.

```
# services.msc
```

- e) Restart the vSphere Web Client to logout of vCenter.

```
# serviceLogout
```

Verifying HX Cluster is Unregistered from vCenter

This task is a step in unregistering a HX Storage Cluster from vCenter.

Verify that the HX cluster is no longer on the old vCenter.

Before you begin

Complete the steps in:

- [Unregistering and Removing EAM Extensions, on page 56](#)
- [Removing HX Data Platform Files from the vSphere Client, on page 58](#)

- Step 1** Clear your cache before logging back into vCenter.
- Step 2** Log out of the old vCenter.
- Step 3** Log in again to the old vCenter and verify the HX Data Platform Plug-in has been removed.

Registering a Storage Cluster with a New vCenter Cluster

Before you begin

As part of the task to move a storage cluster from one vCenter server to another vCenter server, complete the steps in [Unregistering a Storage Cluster from a vCenter Cluster, on page 55](#).

- Step 1** Login to a controller VM.
- Step 2** Run the `stcli cluster reregister` command.

```
stcli cluster reregister [-h] --vcenter-datacenter NEWDATACENTER --vcenter-cluster NEWVCENTERCLUSTER
--vcenter-url NEWVCENTERURLIP [--vcenter-sso-url NEWVCENTERSSOURL] --vcenter-user
NEWVCENTERUSER
```

Apply additional listed options as needed.

Syntax Description	Option	Required or Optional	Description
	--vcenter-cluster NEWVCENTERCLUSTER	Required	Name of the new vCenter cluster.
	--vcenter-datacenter NEWDATACENTER	Required	Name of the new vCenter datacenter.

Option	Required or Optional	Description
--vcenter-sso-url NEWVCENTERSSOURL	Optional	URL of the new vCenter SSO server. This is inferred from --vcenter-url, if not specified.
--vcenter-url NEWVCENTERURLIP	Required	URL of the new vCenter, <vcentername>. Where <vcentername> can be IP or FQDN of new vCenter.
--vcenter-user NEWVCENTERUSER	Required	User name of the new vCenter administrator. Enter vCenter administrator password when prompted.

Example response:

```
Reregister StorFS cluster with a new vCenter ...
Enter NEW vCenter Administrator password:
Waiting for Cluster creation to finish ...
```

If, after your storage cluster is re-registered, your compute only nodes fail to register with EAM, or are not present in the EAM client, and not under the resource pool in vCenter, then run the command below to re-add the compute only nodes:

```
# stcli node add --node-ips <computeNodeIP> --controller-root-password <ctlvm-pwd> --esx-username
<esx-user> --esx-password <esx-pwd>
```

Contact TAC for assistance if required.

Step 3 Re-enter your snapshot schedules.

Snapshot schedules are not moved with the storage cluster when you move the storage cluster between vCenter clusters.

Re-registering the vCenter using HX Connect

You may need to re-register the vCenter in the following scenarios:

- The Controller VM certificate is changed.
- It is recommended to re-register vCenter extensions whenever a vCenter upgrade is performed.
- Re-registration is required when the extensions are manually removed due to misconfigurations.

Step 1 Navigate to **System Information > Actions** drop-down menu in the HX Connect UI. The Action drop-down menu is found on the top right of the **System Overview** tab window.

Step 2 From the **Actions** menu, select the **Re-register vCenter**.

Step 3 Select **Re-register vCenter** from the drop-down menu.

Step 4 Select the **Re-create vCenterExtensions** option if your CVM certificates are changed and if any extensions are mis-configured.

Note The **Re-create vCenterExtensions** option is only recommended when re-registering extensions after a major vCenter upgrade. If selected, enter the existing vCenter username and password.

Step 5 For the new vCenter registration, enter the new vCenter username, password, and vCenter URL for **Re-register vCenter**.

Note When **Re-register vCenter** is selected, you must remove the clusters from the old vCenter and recreate the new clusters and add hosts in the new vCenter.

Step 6 Enter the vCenter datacenter name.

Step 7 Enter the associated cluster name.

You have the option to enter a vCenter SSO URL.

Renaming Clusters

After you create a HX Data Platform storage cluster, you can rename it without disrupting any processes.



Note These steps apply to renaming the HX Cluster, not the vCenter cluster.

-
- Step 1** From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster** to rename.
- Step 2** Open the **Rename Cluster** dialog box. Either right-click on the storage cluster or click the **Actions** drop-down list at the top of the tab.
- Step 3** Select **Rename Cluster**.
- Step 4** Enter a new name for the storage cluster in the text field.
HX cluster names cannot exceed 50 characters.
- Step 5** Click **OK** to apply the new name.
-

Re-registering the vCenter using HX Connect

You may need to re-register the vCenter in the following scenarios:

- The Controller VM certificate is changed.
- It is recommended to re-register vCenter extensions whenever a vCenter upgrade is performed.
- Re-registration is required when the extensions are manually removed due to misconfigurations.

-
- Step 1** Navigate to **System Information > Actions** drop-down menu in the HX Connect UI. The Action drop-down menu is found on the top right of the **System Overview** tab window.
- Step 2** From the **Actions** menu, select the **Re-register vCenter**.
- Step 3** Select **Re-register vCenter** from the drop-down menu.

- Step 4** Select the **Re-create vCenterExtensions** option if your CVM certificates are changed and if any extensions are mis-configured.
- Note** The **Re-create vCenterExtensions** option is only recommended when re-registering extensions after a major vCenter upgrade. If selected, enter the existing vCenter username and password.
- Step 5** For the new vCenter registration, enter the new vCenter username, password, and vCenter URL for **Re-register vCenter**.
- Note** When **Re-register vCenter** is selected, you must remove the clusters from the old vCenter and recreate the new clusters and add hosts in the new vCenter.
- Step 6** Enter the vCenter datacenter name.
- Step 7** Enter the associated cluster name.
- You have the option to enter a vCenter SSO URL.
-

Replacing Self-Signed Certificate

Replacing Self-Signed Certificate with External CA Certificate on a vCenter Server

Set the certMgmt mode in vCenter to **Custom** to add the ESXi hosts with third party certificate to vCenter.

Note By default, the certMgmt mode is **vmsa**. In the default **vmsa** mode, you can add only the ESX host with self signed certificates. If you try to add an ESX with CA certificate to a vCenter, it will not allow you to add the ESX host unless CA certificate is replaced with self-signed certificate.

To update the certMgmt mode:

- Select the vCenter server that manages the hosts and click **Settings**.
- Click **Advanced Settings**, and click **Edit**.
- In the **Filter** box, enter **certmgmt** to display only certificate management keys.
- Change the value of **vpxd.certmgmt.mode** to **custom** and click **OK**.
- Restart the vCenter server service.

To restart services, enter the following link in a browser and then click **Enter**:

`https://<VC URL>:5480/ui/services`



Note The behavior of host addition in vCenter varies according to the certificate and certMgmt mode.

- When the host has self-signed certificate with the certMgmt mode set to the default value of **vmrsa** in vCenter:
 - Only ESX host with self-signed certificate can be added.
 - The addition of ESX with third party CA certificate is not allowed.
 - If you try to add an ESX to a vCenter after replacing the self-signed certificate with a third party CA certificate, the system will prompt you to replace third party CA certificate with self-signed certificate. You can add the ESX host after replacing CA certificate with self-signed certificate.
- When the host has self-signed certificate with the certMgmt mode set to **custom** in vCenter:
 - If you try to add an ESX to a vCenter after replacing the self-signed certificate with a third party CA certificate, the system throws an error: `ssl thumbprint mismatch and add host fails`. In this case, do the following to replace the third party CA certificate with the self-signed certificate:
 1. Place the host in the maintenance mode (MM mode).
 2. Replace the certified `ruicert` and `ruicert.key` files with the backed up previous key and certificate.
 3. Restart the `hostd` and `vxsa` service. The CA certificate comes up in the new node.
 4. Right-click and connect to vCenter. The host removes the CA certificate and gets replaced with self-signed certification in VMware.
- When the host has third party CA certificate with the certMgmt mode set to the default value of **vmrsa** in vCenter:
 - ESX host with self-signed certificate can be added.
 - The addition of ESX with third party CA certificate is not allowed.
- When the host has third party CA certificate with the certMgmt mode set to **custom** in vCenter:
 - ESX host with self-signed certificate cannot be added.
 - The self-signed certificate in ESX host needs to be replaced with a CA certificate of vCenter.

Replacing Self-Signed Certificate with External CA Certificate on an ESXi Host

Step 1 Generate the host certificate (`ruicert`) and key (`ruicert.key`) files and send the files to the certificate authority.

Note Ensure that a proper hostname or FQDN of the ESX host is provided while generating the `ruicert.key` and `ruicert` files.

Step 2 Replace the certified host certificate (`ruicert`) and key (`ruicert.key`) files in the `/etc/vmware/ssl` directory on each ESXi host after taking backup of the original host certificate (`ruicert`) and key (`ruicert.key`) files.

Note Replace host certificate (rui.crt) and key (rui.key) files in a rolling fashion by placing only one host in maintenance mode and then wait for the cluster to be healthy and then replace the certificates for the other nodes.

- a) Log in to the ESXi host from an SSH client with administrator privileges.
- b) Place the host in the maintenance mode (MM mode).
- c) Take a backup of the previous key and certificate to the rui.bak file in the /etc/vmware/ssl/ directory.
- d) Upload the new certified rui.crt and rui.key files to the /etc/vmware/ssl/ directory.
- e) Restart the hostd and vpxa service, and check the running status using the following commands:

```
/etc/init.d/hostd restart
/etc/init.d/vpxa restart
/etc/init.d/hostd status
/etc/init.d/vpxa status
```

- f) Reconnect the host to vCenter and exit the maintenance mode.

Note Repeat the same procedure on all the nodes. You can verify the certificate of each node by accessing it through web.

Reregistering a HyperFlex cluster

After adding all the hosts to the vCenter after replacing the certified files, reregister the HX cluster to the vCenter using the following command:

```
stcli cluster reregister
```

Recreating a Self-Signed Certificate

If you face any issue with the host certificate after replacing external CA certificate, you can recreate the self-signed certificate by executing the following procedure:

1. Log in to the ESXi host from an SSH client.
2. Delete the rui.key and rui.crt files from the /etc/vmware/ssl/ directory.
3. Recreate the self-signed certificate for the host using the following command:

```
/sbin/generate-certificates
```

4. Restart the hostd and vpxa service using the following commands:

```
/etc/init.d/hostd restart
/etc/init.d/vpxa restart
```

Boost Mode

Boost mode allows the Cisco HyperFlex cluster to deliver higher IOPs by increasing the storage controller VM CPU resources by 4 vCPU. Enabling Boost Mode takes additional CPU resources from user VM for the HX data platform, and should only be enabled in deployments where support has determined that the benefit of additional CPUs, outweighs the impact to the sizing of your deployment.

Configuring Boost Mode

Perform the following steps for each cluster you want to enable Boost Mode on:

Before you begin

Boost Mode Support is limited to the following configurations:

- Supported Hardware:
 - All NVMe
 - All Flash C240
 - All Flash C220
- Hypervisor: ESX only
- Boost Mode number of controller VM vCPUs:
 - All NVMe: 16
 - All Flash C240: 12
 - All Flash C220:12
- Cluster expansion requires you to apply Boost Mode to the new nodes.
- Boost Mode is supported in Cisco HX Release 4.0(2a) and later.
- Boost Mode should be enabled after support has determined that your deployment will benefit from the additional CPUs.

**Note**

CPU - number of physical cores must be equal to at least the new number of controller vCPUs. To verify the number of physical cores in the vSphere Client; Click **host** > **Configure** > **Hardware** > **Processors** > **Processor cores per socket**

-
- Step 1** From the vCenter, right-click one controller VM and **Shut Down Guest OS**.
 - Step 2** Increase the number of vCPUs by four (4). In the vSphere client, click **Edit Settings** for the VM and change the value of the CPU field in the first line.
 - Step 3** Click **OK** to apply the configuration changes.
 - Step 4** Power up the controller VM.
 - Step 5** Log in to HX Connect and wait for the cluster to become healthy.
 - Step 6** Repeat the process for each host (or node) in the cluster.
-

Disabling Boost Mode

To disable Boost Mode, perform the following steps:

-
- Step 1** From the vCenter, right-click one controller VM and **Shut Down Guest OS**.
- Step 2** Decrease the number controller VM vCPUs back to 12 for all-NVMe, or 8 for all flash C220, and all flash C240. In the vSphere client, click **Edit Settings** for the VM and change the value of the CPU field in the first line.
- Step 3** Click **OK** to apply the configuration changes.
- Step 4** Power up the controller VM.
- Step 5** Log in to HX Connect and wait for the cluster to become healthy.
- Step 6** Repeat the process for each host (or node) in the cluster.
-

UEFI Secure Boot Mode

Unified Extensible Firmware Interface (UEFI) is a specification that defines a software interface between an operating system and platform firmware. HX Data Platform uses UEFI to replace the BIOS firmware interfaces. This allows the BIOS to run in UEFI mode while still providing legacy support.

Starting with HX Data Platform release 4.5(1a), the hardening of Hypervisor (ESXi) boot security is simplified by providing an automated workflow that non-disruptively changes the boot mode of converged and compute nodes in the cluster to Unified Extensible Firmware Interface (UEFI) Secure Boot. The chain of trust is anchored by a hardware trust anchor (i.e. the Cisco Trust Anchor module) built-in to UCS rack and blade servers. UI and API-based queries of each node's secure boot status is supported so the cluster's security posture on-demand can be audited.

The following limitations apply to the UEFI boot mode:

- Support for Secure boot is available only for HyperFlex ESXi M3/M4/M5 Server.
- Cluster expansion supported on M2R1 nodes.
- VMware ESXi supported versions correspond to HX 4.5(1a) for 6.5, 6.7 and 7.0.
- Attestation of Secure Boot of ESXi hosts by vCenter is supported. This feature requires ESXi release 6.7 or later and TPM 2.0 modules on the converged or compute nodes. The TPM and TXT parameters, which are required to enable usage of the TPM module, are automatically configured in the course of enabling secure boot. No additional steps are required to use attestation.

Enabling Secure Boot Mode

- Enabling Secure Boot mode allows you to change the boot mode of your ESXi hosts from Legacy BIOS or UEFI (Non-Secure) to UEFI Secure Boot.
- Do not manually change the boot parameters in UCS Manager or Cisco IMC for a UCS server, which is part of an HyperFlex cluster. HyperFlex is not aware of and will not automatically remediate such changes.
- Use the Check Secure Boot Status (see step 4) action to audit the secure boot status of the cluster. If a node is found to be out-of-compliance, the Secure Boot mode upgrade type option is made available on the Upgrade tab and users can re-enable Secure Boot. Only those nodes that are out-of-compliance are rebooted and have their boot mode changed.

Before you begin

- Run **Check Secure Boot Status** to confirm whether Secure Boot is already enabled and then proceed accordingly. See step 4.
- Starting with HX release 4.5(1a), UEFI Secure Boot must be enabled as a separate Day 2 operation, after a refresh of HX release 4.5(1a) install or after upgrading an existing cluster to HX 4.5(1a).
- Validate that the cluster is in a ready state to enable Secure Boot by running a pre-flight validation.
- If your cluster has legacy, UEFI, and UEFI Secure booted nodes present in same cluster, secure boot operation will still get enabled on all nodes of the cluster and any expansions after that will be secure boot enabled.
- Option to Enable Secure Boot will be available only for ESXi clusters.
- Due to the ESXi hosts going in to rolling reboot to enable secure boot, plan the activity in a maintenance window.
- Enabling Secure Boot cannot be combined with other upgrade activities.
- If the Secure Boot is already enabled, **Enable Secure Boot** option is greyed out and no further action needed.

Step 1 From the HX Connect UI, navigate to **Upgrade > Select Upgrade Type**

Step 2 From the **Select Upgrade Type** tab, select the **Secure Boot mode** check box.

Note After Secure Boot is enabled, it cannot be disabled.

Step 3 Enter your vCenter and UCSM credentials: **Username** and **Admin password** and click **Upgrade**.

After enabling Secure Boot on the cluster, any new converged or compute nodes subsequently added, automatically have secure boot enabled. No manual intervention is required.

Step 4 To check the status of Secure Boot, navigat to **System Information > Actions** drop-down menu and select **Check Secure Boot Status**.

Note If all nodes are enabled, the **Secure Boot is enabled on all the nodes** message is displayed.

Step 5



CHAPTER 6

Preparing for HX Storage Cluster Maintenance

- [Storage Cluster Maintenance Operations Overview, on page 69](#)
- [Serial vs. Parallel Operations, on page 71](#)
- [Checking Cluster Status, on page 71](#)
- [Setting a Beacon, on page 71](#)
- [Verify vMotion Configuration for HX Cluster, on page 72](#)
- [Maintenance Modes for Storage Cluster Nodes, on page 73](#)
- [Entering Cisco HyperFlex Maintenance Mode, on page 74](#)
- [Exiting Cisco HyperFlex Maintenance Mode, on page 75](#)
- [Creating a Backup Operation, on page 76](#)
- [Shut Down and Power Off the Cisco HX Storage Cluster, on page 80](#)
- [Power On and Start Up the Cisco HX Storage Cluster, on page 82](#)
- [Restoring the Configuration for a Fabric Interconnect, on page 84](#)
- [Configure PCI Passthrough After Changing vNIC or vHBAs, on page 86](#)

Storage Cluster Maintenance Operations Overview

Maintaining the Cisco HyperFlex (HX) Data Platform storage cluster tasks affect both hardware and software components of the storage cluster. Storage cluster maintenance operations include adding or removing nodes and disks, and network maintenance.

Some steps in maintenance tasks are performed from the storage controller VM of a node in the storage cluster. Some commands issued on a storage controller VM affect all the nodes in the storage cluster.



Note **Three node storage clusters.** Contact Technical Assistance Center (TAC) for any task that requires removing or shutting down a node in a three node cluster. With any 3 node storage cluster, if one node fails or is removed, the cluster remains in an unhealthy state until a third node is added and joins the storage cluster.

Upgrading from vSphere 5.5 to 6.0. Before you upgrade either your ESX server or your vCenter server from 5.5 to 6.0, contact Technical Assistance Center (TAC).

Adding nodes. Nodes are added to the storage cluster through the Expand Cluster feature of the Cisco HX Data Platform Installer. All new nodes must meet the same system requirements as when you installed the Cisco HX Data Platform and created the initial storage cluster. For a complete list of requirements and steps for using the Expand Cluster feature, see the appropriate [Cisco HX Data Platform Install Guide](#).

Online vs Offline Maintenance

Depending upon the task, the storage cluster might need to be either online or offline. Typically maintenance tasks require that all nodes in the storage cluster are online.

When storage cluster maintenance is performed in an offline mode, this means the Cisco HX Data Platform is offline, however the storage controller VMs are up and Cisco HX Data Platform management is viewable through the `stcli` command line, HX Connect, and HX Data Platform Plug-in. The vSphere Web Client can report on the storage I/O layer. The `stcli cluster info` command returns that the overall storage cluster status is `offline`.

Pre-Maintenance Tasks

Before you perform maintenance on the storage cluster, ensure the following.

- Identify the maintenance task to be performed.
- All maintenance operations such as remove/replace resources are done during maintenance windows when the load on the system is low.
- The storage cluster is healthy and operational **before** the maintenance tasks.
- Identify disks using the HX Connect or HX Data Platform Plug-in Beacon options.

The HX Beacon option is not available for housekeeping 120GB SSDs. Physically check the server for the location of the housekeeping SSD.

- Check the list of maintenance tasks that cannot be performed in parallel. See [Serial vs. Parallel Operations, on page 71](#) for more information on these tasks.. You can perform only some tasks serially to each other.
- Ensure that SSH is enabled on all the ESX hosts.
- Put the ESX host into HX Maintenance Mode prior to performing a maintenance task on the host. The HX maintenance mode performs additional storage cluster specific steps compared to the vSphere provided ESX maintenance mode.

Post Maintenance Tasks

After the maintenance task is completed, the nodes need to exit Cisco HX Maintenance Mode and the storage cluster needs to be restarted. In addition, some changes to the Cisco HX storage cluster require additional post maintenance tasks. For example, if you change the vNICs or vHBAs, the PCI Passthrough needs to be reconfigured. For more information describing how to reconfigure the PCI Passthrough, see [Configure PCI Passthrough After Changing vNIC or vHBAs, on page 86](#).

Ensure the following:

- The ESX host is exited from Cisco HX maintenance mode after performing maintenance tasks on the host.
- The storage cluster is healthy and operational **after** any remove or replace tasks are completed.
- If vNICs or vHBAs have been added, removed, or replace on any ESX host in the Cisco HX storage cluster, reconfigure the PCI Passthrough.

Serial vs. Parallel Operations

Certain operations cannot be performed simultaneously. Ensure that you perform the following operations serially (not in parallel).

- Upgrade a storage cluster or a node.
- Create, re-create, or configure a storage cluster.
- Add or remove a node.
- Any node maintenance that requires a node be shutdown. This includes adding or removing disks or network interface cards (NICs).
- Start or shut down a storage cluster.
- Re-register a storage cluster with vCenter.

Checking Cluster Status

Step 1 Login to any controller VM in the storage cluster. Run the listed commands from the controller VM command line.

Step 2 Verify the storage cluster is healthy.

```
# stcli cluster info
```

Example response that indicates the storage cluster is online and healthy:

```
locale: English (United States)
state: online
upgradeState: ok
healthState: healthy
state: online
state: online
```

Step 3 Verify the number of node failures.

```
# stcli cluster storage-summary
```

Example response:

```
#of node failures tolerable to be > 0
```

Setting a Beacon

Beaconing is a method of turning on an LED to assist in locating and identifying a node (host) and a disk. Nodes have the beacon LED in the front near the power button and in the back. Disks have the beacon LED on the front face.

You set a node beacon through Cisco UCS Manager. You set a disk beacon through the Cisco HX Data Platform Plug-in or HX Connect user interface.

Step 1 Turn on and off a node beacon using UCS Manager.

- From the UCS Manager left panel, select **Equipment > Servers > server**.
- From the UCS Manager central panel, select **General > Turn on Locator LED**.
- After you locate the server, turn off the locator LED.

From the UCS Manager central panel, select **General > Turn off Locator LED**.

Step 2 Turn on and off a disk beacon using the Cisco HX Data Platform Plug-in.

- From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Manage**.
- From **Manage**, select **Cluster > cluster > host > Disks > disk**.
- Locate the physical location of the object and turn on the beacon.

From **Actions** drop-down list, select **Beacon ON**.

- After you locate the disk, turn off the beacon.

From **Actions** drop-down list, select **Beacon OFF**.

Step 3 Turn on and off a disk beacon using HX Connect.

- Log in to HX Connect.
- Select **System Information > Disks**.
- Select a node, and then click **Turn On Locator LED** or **Turn Off Locator LED**.

The beacon LED for all the disks on the selected node are toggled, except Housekeeping SSDs and cache NVMe SSDs. Housekeeping SSDs or cache NVMe SSDs do not have functioning LED beacons.

Verify vMotion Configuration for HX Cluster

Before you perform maintenance operations on the Cisco HyperFlex (HX) cluster, verify all nodes in the HX cluster are configured for vMotion. Confirm the following from your vSphere Web Client:

- Verify that the vMotion port group is configured with `vmnic3` and `vmnic7` in an active/standby configuration across all of the ESXi hosts in the cluster.
- Verify that a port group is configured for vMotion, and that the naming convention is EXACTLY the same across all ESXi hosts in the cluster.



Note The name is case-sensitive.

- Verify that you have assigned a static IP to each vMotion port group, and that the static IPs for each vMotion port group are in the same subnet.



Note The static IP address is defined as a VMKernel interface.

4. Verify that the vMotion port group has the vMotion option checked in the properties, and that no other port groups (such as management) have this option checked, on each ESXi host in the cluster.
5. Verify in the settings that the vMotion port group is set to 9000 MTU, (if you are using jumbo frames), and the VLAN ID matches the network configuration for the vMotion subnet.
6. Verify you can ping from the vMotion port group on one ESXi host to the vMotion IP on the other host.

Type `vmkping -I vmk2 -d -s 8972 <vMotion IP address of neighboring server>`

Maintenance Modes for Storage Cluster Nodes

Maintenance mode is applied to nodes in a cluster. It prepares the node for assorted maintenance tasks by migrating all VMs to other nodes before you decommission or shut the node down.

There are two types of maintenance modes.

- Cisco HX maintenance mode
- VMware ESX maintenance mode

Cisco HX Maintenance Mode

Cisco HX maintenance mode performs Cisco HX Data Platform specific functions in addition to the ESX maintenance mode. Be sure to select Cisco HX maintenance mode and not ESX maintenance mode for maintenance tasks performed on storage cluster nodes after initial storage cluster creation.

This mode is the preferred maintenance mode for performing selected tasks on individual nodes in the cluster. Including:

- Shutting down an individual host for maintenance, such as disk replacement.
- Upgrading selected software on a host, such as ESX Server version.

Cisco HX Maintenance Mode Considerations

- Ensure that SSH is enabled in ESX on all the nodes in the storage cluster prior to using Cisco HX Maintenance Mode.
- When Cisco HX Maintenance Mode is entered to enable performing tasks on an ESX host, be sure to exit Cisco HX Maintenance Mode after the tasks on the ESX host are completed.
- Cisco HX Maintenance Mode is applied to nodes in a healthy cluster only. If the cluster is unhealthy, for example too many nodes are down, or you are shutting down the cluster, use ESX Maintenance Mode.
- When nodes are added or removed from the cluster, the number of resources (controller VM, caching and persistent tier devices, etc) to serve the user IO changes. HXDP aims to use the available cluster resources to serve IO optimally. Each node is used to serve part of user IO as well as be responsible for doing internal bookkeeping activities.

When a node leaves (entering Maintenance Mode), the in-flight IO needs to failover to other nodes in the cluster. In addition to internal book-keeping resources and activities also need to failover. The time required for this is proportional to data and activities which were being served by the node. This results in additional latency for the in-flight user IO.

This is similar to the case where nodes come back from Maintenance Mode.

- See [Entering Cisco HyperFlex Maintenance Mode](#) and [Exiting Cisco HyperFlex Maintenance Mode](#), on page 75 for steps.

VMware ESX Maintenance Mode

This mode is used when you are installing Cisco HX Data Platform or applying cluster wide changes.

To enter or exit vSphere maintenance mode:

- Through the vCenter GUI, select the **host**, then from the right-click menu select **maintenance mode**.
- Through the ESX command line, use the `esx maintenance mode` command.

Entering Cisco HyperFlex Maintenance Mode

Using the Cisco HyperFlex (HX) Connect User Interface



Note

Maintenance Mode was introduced in Cisco HyperFlex Release 2.5(1a) and 2.5(1b).

1. Log in to Cisco HX Connect: `https://<cluster management ip>`.
2. In the menu, click **System Information**.
3. Click **Nodes**, and then click the row of the node you want to put in to maintenance mode.
4. Click **Enter HX Maintenance Mode**.
5. In the **Confirm HX Maintenance Mode** dialog box, click **Enter HX Maintenance Mode**.



Note

After you complete any maintenance tasks, you must manually exit HX maintenance mode.

Using the vSphere Web Client

1. Log in to the vSphere web client.
2. Go to **Home > Hosts and Clusters**.
3. Expand the **Datacenter** that contains the **HX Cluster**.
4. Expand the **HX Cluster** and select the node.
5. Right-click the node and select **Cisco HX Maintenance Mode > Enter HX Maintenance Mode**.

Using the Command-Line Interface

1. Log in to the storage controller cluster command line as a user with root privileges.

2. Move the node into HX Maintenance Mode.

- a. Identify the node ID and IP address.

```
# stcli node list --summary
```

- b. Enter the node into HX Maintenance Mode.

```
# stcli node maintenanceMode (--id ID | --ip IP Address) --mode enter
```

(see also `stcli node maintenanceMode --help`)

3. Log in to the ESXi command line of this node as a user with root privileges.
4. Verify that the node has entered HX Maintenance Mode.

```
# esxcli system maintenanceMode get
```

You can monitor the progress of the **Enter Maintenance Mode** task in vSphere Web Client, under the **Monitor > Tasks** tab.

If the operation fails, an error message displays. Try to fix the underlying problem and attempt to enter maintenance mode again.

Exiting Cisco HyperFlex Maintenance Mode

Using the Cisco HyperFlex (HX) Connect User Interface



Note

Maintenance Mode was introduced in Cisco HyperFlex Release 2.5(1a) and 2.5(1b).

1. Log in to HX Connect: <https://<cluster management ip>>.
2. In the menu, click **System Information**.
3. Click **Nodes**, and then click the row of the node you want to remove from maintenance mode.
4. Click **Exit HX Maintenance Mode**.

Using the vSphere Web Client

1. Log in to the vSphere web client.
2. Go to **Home > Hosts and Clusters**.
3. Expand the **Datacenter** that contains the **HX Cluster**.
4. Expand the **HX Cluster** and select the node.
5. Right-click the node and select **Cisco HX Maintenance Mode > Exit HX Maintenance Mode**.

Using the Command-Line Interface

1. Log in to the storage controller cluster command line as a user with root privileges.

2. Exit the node out of HX Maintenance Mode.

- a. Identify the node ID and IP address.

```
# stcli node list --summary
```

- b. Exit the node out of HX Maintenance Mode.

```
# stcli node maintenanceMode (--id ID | --ip IP Address) --mode exit
```

(see also `stcli node maintenanceMode --help`)

3. Log in to the ESXi command line of this node as a user with root privileges.
4. Verify that the node has exited HX Maintenance Mode.

```
# esxcli system maintenanceMode get
```

You can monitor the progress of the **Exit Maintenance Mode** task in vSphere Web Client, under the **Monitor > Tasks** tab.

If the operation fails, an error message displays. Try to fix the underlying problem and attempt to exit maintenance mode again.

Creating a Backup Operation

Before you shutdown your HX storage cluster, backup the configuration. Perform both the Full-State and All Configuration type backups with the Preserve Identities attribute.

Before you begin

1. Login to UCS Manager.
2. Obtain the backup server IPv4 address and authentication credentials.



Note All IP addresses must be IPv4. HyperFlex does not support IPv6 addresses.

- Step 1** In the **Navigation** pane, click **Admin**.
- Step 2** Click the **All** node.
- Step 3** In the **Work** pane, click the **General** tab.
- Step 4** In the **Actions** area, click **Backup Configuration**.
- Step 5** In the **Backup Configuration** dialog box, click **Create Backup Operation**.
- Step 6** In the **Create Backup Operation** dialog box, complete the following fields:

Name	Description
Admin State field	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • Enabled—Cisco UCS Manager runs the backup operation as soon as you click OK. • Disabled—Cisco UCS Manager does not run the backup operation when you click OK. If you select this option, all fields in the dialog box remain visible. However, you must manually run the backup from the Backup Configuration dialog box.
Type field	<p>The information saved in the backup configuration file. This can be one of the following:</p> <ul style="list-style-type: none"> • Full state—A binary file that includes a snapshot of the entire system. You can use the file generated from this backup to restore the system during disaster recovery. This file can restore or rebuild the configuration on the original fabric interconnect, or recreate the configuration on a different fabric interconnect. You cannot use this file for an import. <p>Note You can only use a full state backup file to restore a system that is running the same version as the system from which the backup file was exported.</p> <ul style="list-style-type: none"> • All configuration—An XML file that includes all system and logical configuration settings. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore. This file does not include passwords for locally authenticated users. • System configuration—An XML file that includes all system configuration settings such as usernames, roles, and locales. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore. • Logical configuration—An XML file that includes all logical configuration settings such as service profiles, VLANs, VSANs, pools, and policies. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore.

Name	Description
Preserve Identities check box	<p>This checkbox remains selected for All Configuration and System Configuration type of backup operation, and provides the following functionality:</p> <ul style="list-style-type: none"> • All Configuration—The backup file preserves all identities derived from pools, including vHBAs, WWPNs, WWNN, vNICs, MACs and UUIDs. Also, the identities for Chassis, FEX, Rack Servers, and user labels for Chassis, FEX, Rack Servers, IOMs and Blade Servers are preserved. <p>Note If this check box is not selected the identities will be reassigned and user labels will be lost after a restore.</p> <ul style="list-style-type: none"> • System Configuration—The backup file preserves identities for Chassis, FEX, Rack Servers, and user labels for Chassis, FEX, Rack Servers, IOMs and Blade Servers. <p>Note If this check box is not selected the identities will be reassigned and user labels will be lost after a restore.</p> <p>If this checkbox is selected for Logical Configuration type of backup operation, the backup file preserves all identities derived from pools, including vHBAs, WWPNs, WWNN, vNICs, MACs and UUIDs.</p> <p>Note If this check box is not selected the identities will be reassigned and user labels will be lost after a restore.</p>
Location of the Backup File field	<p>Where the backup file should be saved. This can be one of the following:</p> <ul style="list-style-type: none"> • Remote File System—The backup XML file is saved to a remote server. Cisco UCS Manager GUI displays the fields described below that allow you to specify the protocol, host, filename, username, and password for the remote system. • Local File System—The backup XML file is saved locally. <p>HTML-based Cisco UCS Manager GUI displays the Filename field. Enter a name for the backup file in <filename>.xml format. The file is downloaded and saved to a location depending on your browser settings.</p>

Name	Description
Protocol field	<p>The protocol to use when communicating with the remote server. This can be one of the following:</p> <ul style="list-style-type: none"> • FTP • TFTP • SCP • SFTP • USB A—The USB drive inserted into fabric interconnect A. This option is only available for certain system configurations. • USB B—The USB drive inserted into fabric interconnect B. This option is only available for certain system configurations.
Hostname field	<p>The hostname, IPv4 address of the location where the backup file is stored. This can be a server, storage array, local drive, or any read/write media that the fabric interconnect can access through the network.</p> <p>Note If you use a hostname rather than an IPv4 address, you must configure a DNS server. If the Cisco UCS domain is not registered with Cisco UCS Central or DNS management is set to local, configure a DNS server in Cisco UCS Manager. If the Cisco UCS domain is registered with Cisco UCS Central and DNS management is set to global, configure a DNS server in Cisco UCS Central.</p> <p>Note All IP addresses must be IPv4. HyperFlex does not support IPv6 addresses.</p>
Remote File field	The full path to the backup configuration file. This field can contain the filename as well as the path. If you omit the filename, the backup procedure assigns a name to the file.
User field	The username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP or USB.
Password field	<p>The password for the remote server username. This field does not apply if the protocol is TFTP or USB.</p> <p>Cisco UCS Manager does not store this password. Therefore, you do not need to enter this password unless you intend to enable and run the backup operation immediately.</p>

Step 7 Click **OK**.

Step 8 If Cisco UCS Manager displays a confirmation dialog box, click **OK**.

If you set the **Admin State** field to enabled, Cisco UCS Manager takes a snapshot of the configuration type that you selected and exports the file to the network location. The backup operation displays in the **Backup Operations** table in the **Backup Configuration** dialog box.

- Step 9** (Optional) To view the progress of the backup operation, do the following:
- If the operation does not display in the **Properties** area, click the operation in the **Backup Operations** table.
 - In the **Properties** area, click the down arrows on the **FSM Details** bar.
- The **FSM Details** area expands and displays the operation status.
- Step 10** Click **OK** to close the **Backup Configuration** dialog box.
- The backup operation continues to run until it is completed. To view the progress, re-open the **Backup Configuration** dialog box.

Shut Down and Power Off the Cisco HX Storage Cluster

Some storage cluster maintenance tasks require that the storage cluster be shut down. This is different than the storage cluster being in an offline state. It is also separate from shutting down a node in the storage cluster. Powering down the storage cluster affects all the physical components of the cluster.

- A **powered-off cluster** has all the physical components of the storage cluster removed from electrical power.
Very rarely would a storage cluster need to have all the components powered off. No regular maintenance or upgrade processes require that the entire storage cluster be completely powered off.
- A **shut-down cluster** has all storage cluster processes, including the working VMs, powered down. This does not include powering down the nodes in the cluster or shutting down the vCenter or FI cluster.
- An **offline cluster** is one of the storage cluster operational states. A storage cluster can be offline if there is an unknown or specific error, or if the storage cluster has been shutdown.

To shut down the Cisco HX storage cluster, perform the following steps:

Before you begin

- The storage cluster must be in a healthy state.
- Perform both the Full-State and All Configuration type backups with the Preserve Identities attribute. See [Creating a Backup Operation, on page 76](#).

- Step 1** Gracefully shut down all workload VMs on all the Cisco HX datastores.
Alternatively, use vMotion to migrate the workload VMs to another cluster.
- Note** Do not shut down or move the storage controller VMs (stCtlVMs).
- Step 2** Gracefully shut down the Cisco HX storage cluster.
- From any controller VM command line, run the command and wait for the shell prompt to return.
- Note** For clusters with a nested vCenter, performing an stcli cluster shutdown may have certain limitations. For more details, see [Known Constraints with vCenter Deployment](#).
- ```
stcli cluster shutdown
```

- b) Run the cluster information command. Confirm the storage cluster is offline.

```
stcli cluster info
```

In the command response text, check the cluster subsection and verify the `healthstate` is `unknown`.

This Cisco HX cluster shutdown procedure does not shut down the ESXi hosts.

If the maintenance or upgrade task does not require the physical components be powered off, exit these steps and proceed to *What to do next*:

**Step 3** To power off the **HX storage cluster**, complete Step 2 and Step 3, then complete the rest of the following steps.

**Step 4** On each storage cluster ESX host, shutdown the controller VM (`stctlvm`).

Choose a method:

Using vCenter VM Power Off

- a) From vCenter client, locate the controller VM on each ESX host.
- b) Right-click the controllerVM and select **Power > Power Off**.

This method performs a graceful guest VM shutdown.

Using vCenter ESX Agent Manager

- a) From vCenter client, open the ESX Agent Manager console.
- b) Locate the controller VM on each ESX host, and select **Power > Power Off**.

This method performs a graceful shutdown of agent VMs. The controller VM is an agent VM.

Using vCenter ESX Maintenance Mode

- a) From vCenter client, locate each ESX host.
- b) Right-click the ESX host and select **Maintenance Mode > Enter Maintenance Mode**.

This method performs a hard shutdown on every VM in the ESX host, including the controller VM.

**Step 5** Shutdown each storage cluster ESX host.

- a) From the vCenter client, locate the host.
- b) Right-click the host and select **Power > Shut Down**.

**Step 6** Power off the FIs, if this is needed for your maintenance task.

Cisco UCS FIs are designed for continuous operation. In a production environment, there is no need to shut down or reboot Fabric Interconnects. Therefore, there is no power button on UCS Fabric Interconnects.

**To power off Cisco UCS Fabric Interconnect**, pull the power cable manually. Alternatively, if you have the FI power cables connected to a smart PDUs, use the provided remote control to turn off the power from the electrical outlet.

- a) Verify all the storage cluster servers on the FI do not have a green power LED.
- b) Power off the secondary FI.
- c) Power off the primary FI.

---

The HX storage cluster is now safely powered off.

**What to do next**

1. Complete the task that required the storage cluster shutdown or power off. For example, an offline upgrade, physically moving the storage cluster, or performing maintenance on nodes.

- For upgrade tasks, see the [Cisco HyperFlex Systems Upgrade Guide](#).
- For hardware replacement tasks, see the server hardware guides.

Sometimes these tasks require that the host is shutdown. Follow the steps in the server hardware guides for migrating VMs, entering Cisco HX Maintenance Mode, and powering down the servers, as directed.

**Note**

Most hardware maintenance tasks do not require the Cisco HX cluster is shutdown.

2. To restart the Cisco HX storage cluster, proceed to [Power On and Start Up the Cisco HX Storage Cluster, on page 82](#).

## Power On and Start Up the Cisco HX Storage Cluster

The steps here are for use in restarting the Cisco HX storage cluster after a graceful shutdown and power off. Typically, this is performed after maintenance tasks are completed on the storage cluster.

**Before you begin**

Complete the steps in [Shut Down and Power Off the Cisco HX Storage Cluster, on page 80](#).

**Step 1**

Plug in to power up the FIs.

- a) Power on the primary FI. Wait until you can gain access to UCS Manager.
- b) Power on the secondary FI. Verify it is online in UCS Manager.

In some rare cases, you might need to reboot the Fabric Interconnects.

- a. Log in to each Fabric Interconnect using SSH.
- b. Issue the commands:

```
FI# connect local-mgmt
FI# reboot
```

**Step 2**

Connect all the ESX hosts to the FIs.

- a) Power on each node in the storage cluster if it does not power on automatically.

The node should automatically power on and boot into ESX. If any node does not, then connect to the UCS Manager and power up the servers (nodes) from UCS Manager.

- b) Verify each ESX host is up and associated with its respective service profile in UCS Manager.

**Step 3**

Verify all the ESXi hosts are network reachable.

Ping all the management addresses.

**Step 4** Exit each node from maintenance mode.

**Note** This is automatically completed by the `stcli cluster start` command.

**Step 5** If all the controller VMs are not automatically powered on, power on all the controller VMs (`stCtrlVM`) using one of the following methods:

Using vSphere Client

- From the vSphere Client, view a storage controller host.
- Right-click the `stCtrlVM` and select **Power > Power On**.
- Repeat for each host.

Using ESXi host command line

- Login to a host.
- Identify the VMID of the `stCtrlVM`.  

```
vim-cmd vmsvc/getallvms
```
- Using the VMID power on the controller VM.  

```
vim-cmd vmsvc/power.on VMID
```
- Repeat for each host.

**Step 6** Wait for all the controller VMs to boot and become network reachable. Then verify.

Ping the management addresses of each of the controller VMs.

**Step 7** Verify the storage cluster is ready to be restarted.

- SSH to any controller VM, run the command:  

```
stcli about
```
- If the command returns full storage cluster information, including build number, the storage cluster is ready to be started. Proceed to restarting the storage cluster.
- If the command does not return full storage cluster information, wait until all the services have started on the host.

**Step 8** Start the storage cluster.

From the command line of any controller VM, run the command.

```
stcli cluster start
```

Depending upon the maintenance or upgrade task performed while the HX cluster was shutdown, the nodes might be exited from HX maintenance mode or ESX maintenance mode. Ignore any error messages about an unknown host exception.

**Step 9** Wait until the storage cluster is online and returns to a healthy state.

- From any controller VM, run the command.  

```
stcli cluster info
```
- In the command response text, check the cluster subsection and verify the `healthstate` is `online`.

This could take up to 30 minutes, it could take less time depending upon the last known state.

**Step 10** Through vCenter, verify that ESX remounted the datastores.

Once the cluster is available, the datastores are automatically mounted and available.

If ESX does not recognize the datastores, from the ESX command line, run the command.

```
esxcfg-nas -r
```

**Step 11**

When the storage cluster is healthy and the datastores are remounted, power on the workload VMs.

Alternatively, use vMotion to migrate the workload VMs back to the storage cluster.

## Restoring the Configuration for a Fabric Interconnect

It is recommended that you use a full state backup file to restore a system that is running the same version as the system from which the backup file was exported. You can also use a full state backup to restore a system if they have the same release train. For example, you can use a full state backup taken from a system running Release 2.1(3a) to restore a system running Release 2.1(3f).

To avoid issues with VSAN or VLAN configuration, a backup should be restored on the fabric interconnect that was the primary fabric interconnect at the time of backup.

**Before you begin**

Collect the following information to restore the system configuration:

- Fabric interconnect management port IPv4 address and subnet mask.
- Default gateway IPv4 address.



**Note** All IP address must be IPv4. IPv6 addresses are not supported.

- Backup server IPv4 address and authentication credentials.
- Fully-qualified name of a Full State backup file



**Note** You must have access to a Full State configuration file to perform a system restore. You cannot perform a system restore with any other type of configuration or backup file.

**Step 1**

Connect to the console port.

**Step 2**

If the fabric interconnect is off, power on the fabric interconnect.

You will see the power on self-test message as the fabric interconnect boots.

**Step 3**

At the installation method prompt, enter **gui**.

**Step 4**

If the system cannot access a DHCP server, you may be prompted to enter the following information:

- IPv4 address for the management port on the fabric interconnect

- Subnet mask or prefix for the management port on the fabric interconnect
- IPv4 address for the default gateway assigned to the fabric interconnect

**Step 5** Copy the web link from the prompt into a web browser and go to the Cisco UCS Manager GUI launch page.

**Step 6** On the launch page, select **Express Setup**.

**Step 7** On the **Express Setup** page, select **Restore From Backup** and click **Submit**.

**Step 8** In the **Protocol** area of the **Cisco UCS Manager Initial Setup** page, select the protocol you want to use to upload the full state backup file:

- SCP
- TFTP
- FTP
- SFTP

**Step 9** In the **Server Information** area, complete the following fields:

| Name                    | Description                                                                                                                                                                                                                                                                         |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Server IP</b>        | The IPv4 address of the computer where the full state backup file is located. This can be a server, storage array, local drive, or any read/write media that the fabric interconnect can access through the network.                                                                |
| <b>Backup File Path</b> | <p>The file path where the full state backup file is located, including the folder names and filename.</p> <p><b>Note</b> You can only use a full state backup file to restore a system that is running the same version as the system from which the backup file was exported.</p> |
| <b>User ID</b>          | The username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP or USB.                                                                                                                                                        |
| <b>Password</b>         | The password for the remote server username. This field does not apply if the protocol is TFTP or USB.                                                                                                                                                                              |

**Step 10** Click **Submit**.

You can return to the console to watch the progress of the system restore.

The fabric interconnect logs in to the backup server, retrieves a copy of the specified full-state backup file, and restores the system configuration.

For a cluster configuration, you do not need to restore the secondary fabric interconnect. As soon as the secondary fabric interconnect reboots, Cisco UCS Manager synchronizes the configuration with the primary fabric interconnect.

# Configure PCI Passthrough After Changing vNIC or vHBAs

## Description

After vNIC or vHBA are manually added to a Cisco HyperFlex (HX) service profile or service profile template, the PCI devices are re-enumerated, and the VMware directpath I/O configuration is lost. When the service profile is changed, the host hardware is updated and the PCI passthrough must be reconfigured. Perform the following steps on each ESX host with a modified service profile.

Perform the following steps on the storage controller VM of the modified ESX host:

**Action: Update the vSphere Service Profile on the ESX Host**

- 
- Step 1** Put the ESX host into HX Maintenance mode.
- Step 2** Make or confirm the changes, such as adding hardware, in the Service Profile.
- Step 3** Reboot the ESX host.
- This host loses the direct path configuration.
- Step 4** Login to vCenter and select the DirectPath I/O Configuration page.
- From vCenter Client: Select the *ESX host* > **Configuration tab** > **Hardware pane** > **Advanced Settings** > **Edit**.
- From vCenter Web Client: From the **vCenter Inventory**, select **Resources** > **Hosts** > *ESX host* > **Manage** > **Settings** > **Hardware** > **PCI Devices** > **Edit**.
- Step 5** Select the LSI card for passthrough.
- From the DirectPath I/O Configuration page, select **Configure Passthrough**.
  - From the Mark devices for passthrough list, select the LSI card for the pass through.
  - Click **OK**.
- Step 6** Reboot the ESX host.
- Step 7** Re-map the PCI device to the HX storage controller VM (StCtlVM), by editing the storage controller VM settings.
- Locate and remove the unknown PCI Device.
- From vCenter Client: Right-click the *HX storage controller VM*, select **Edit Settings** > **PCI device 0** > **Remove** > **OK**.
- From vCenter Web Client: Right-click the *HX storage controller VM*, select **Edit Settings** > **Remove PCI device 0** > **OK**.
- Locate and re-add the LSI Logic PCI device.
- From vCenter Client: Right-click the *HX storage controller VM*, select **Edit Settings** > **Add** > **PCI Device** > **LSI Logic PCI device** > **OK**.
- From vCenter Web Client: Right-click the *HX storage controller VM*, select **Edit Settings** > **PCI Device** > **Add** > **LSI Logic PCI device** > **OK**.
- Step 8** Remove the ESX host from HX Maintenance mode.
- When the host is active again, the HX storage controller VM properly boots and rejoins the storage cluster.
-



## CHAPTER 7

# Managing Encryption

- [Self-Encrypting Drives Overview, on page 87](#)
- [Verify if the HyperFlex Cluster Is Encryption Capable, on page 88](#)
- [Configuring Local Encryption Key, on page 88](#)
- [Modifying Local Encryption Key, on page 89](#)
- [Disabling Local Encryption Key, on page 89](#)
- [Secure Erase an Encrypted Disk, on page 90](#)
- [Remote Key Management, on page 90](#)
- [Configuring Remote Encryption Key, on page 91](#)
- [Generating Certificate Signing Requests, on page 91](#)
- [Configuring a Key Management Server Using CSRs \(Certificate Signing Requests\), on page 93](#)
- [Generating Self-Signed Certificates, on page 94](#)
- [Configuring a key management server using SSCs \(Self-Signed Certificates\), on page 95](#)
- [Restart Encryption, on page 96](#)

## Self-Encrypting Drives Overview

Self-Encrypting Drives (SEDs) have special hardware that encrypts incoming data and decrypts outgoing data in real-time. The data on the disk is always stored in encrypted form. A media encryption key controls this encryption and decryption. This key is never stored in the processor or memory.

A security key, also known as Key-Encryption Key or an authentication passphrase, is used to encrypt the media encryption key. To enable SED, you must provide a security key. No key is required to fetch the data, if the disk is not locked.

Cisco HyperFlex Systems enables you to configure security keys locally or remotely. When you configure the key locally, you must remember the key. In case you forget the key, it cannot be retrieved, and the data is lost if the drive power cycles. You can configure the key remotely by using a key management server (also known as KMIP server). This method addresses the issues related to safe-keeping and retrieval of the keys in the local management.

The encryption and decryption for SEDs is done through the hardware. Thus, it does not affect the overall system performance. SEDs reduce the disk retirement and redeployment costs through instantaneous cryptographic erasure. Cryptographic erasure is done by changing the media encryption key. When the media encryption key of a disk is changed, the data on the disk cannot be decrypted, and is immediately rendered unusable.

An SED based cluster can have encryption enabled and disabled at will. You are free to move between the two states whenever you want. For more information, see the [HX-Hardening Guide](#).

## Verify if the HyperFlex Cluster Is Encryption Capable

### Verify Using the HX Data Platform Plug-in

1. From the HX Data Platform Plug-in, log in to vSphere Web Client.
2. Select **Global Inventory Lists > Cisco Hyperflex Systems > Cisco HX Data Platform > Cluster\_Name > Summary >** .
3. If the HyperFlex cluster has SED drives and is encryption capable, **Data At Rest Encryption-Capable** is listed at the top of the **Summary** tab.

### Verify Using the HX Connect User Interface

1. From the HX Connect UI, select **Encryption**.
2. If the HX cluster has SED drives and is encryption capable, **Data At Rest Encryption-Available** is listed on the **Encryption** page.

## Configuring Local Encryption Key

**Step 1** On the Cisco HyperFlex Connect Navigation Pane, choose **Encryption**.

**Step 2** On the Encryption Page, click **Configure encryption**.

**Step 3** Enter the following Cisco UCS Manager credentials.

| UI Element                  | Essential Information                                                                                |
|-----------------------------|------------------------------------------------------------------------------------------------------|
| UCS Manager host name field | Cisco UCS Manager cluster host name.<br>Enter an IP address or FQDN.<br><eng-fi12.eng.storvisor.com> |
| User name field             | <admin> username                                                                                     |
| Password field              | <admin> password                                                                                     |

Click **Next**.

**Step 4** To secure the HyperFlex cluster using an encryption key generated and stored locally, select **Local Key**.

Click **Next**.

**Step 5** Enter the **encryption key (passphrase)** for this cluster.

**Note** Enter exactly 32 alphanumeric characters.

**Step 6** Click **Enable Encryption**.

## Modifying Local Encryption Key

**Step 1** On the Cisco HyperFlex Connect Navigation Pane, choose **Encryption**.

**Step 2** On the Encryption Page, click **Re-key**.

**Step 3** Enter the following Cisco UCS Manager credentials.

| UI Element                  | Essential Information                |
|-----------------------------|--------------------------------------|
| UCS Manager host name field | For example, <i>10.193.211.120</i> . |
| User name field             | <admin> username.                    |
| Password field              | <admin> password.                    |

Click **Next**.

**Step 4** Enter the **Existing Encryption Key** and the **New Encryption Key** for the cluster.

**Note** Enter exactly 32 alphanumeric characters.

**Step 5** Click **Re-key**.

## Disabling Local Encryption Key

**Step 1** On the Cisco HyperFlex Connect Navigation Pane, choose **Encryption**.

**Step 2** On the Encryption Page, from the **Edit configuration** drop-down menu, choose **Disable encryption**.

**Step 3** Enter the following Cisco UCS Manager credentials.

| UI Element                  | Essential Information                                                                                |
|-----------------------------|------------------------------------------------------------------------------------------------------|
| UCS Manager host name field | Cisco UCS Manager cluster host name.<br>Enter an IP address or FQDN.<br><eng-fi12.eng.storvisor.com> |
| User name field             | <admin> username                                                                                     |
| Password field              | <admin> password                                                                                     |

Click **Next**.

**Step 4** To disable the encryption key on the cluster, enter the **encryption key** in use for the cluster.

- Step 5** Click **Disable encryption**.
- Step 6** To confirm disabling the encryption key on the cluster, in the **Disable encryption?** dialog box, click **Yes, disable encryption**.
- 

## Secure Erase an Encrypted Disk

---

- Step 1** On the Cisco HyperFlex Connect Navigation Pane, choose **System Information**.
- Step 2** From the **Disks** tab, select the **disk** from which you want to securely erase the local key.
- Step 3** Click the **Secure erase** button.
- Step 4** To securely erase the encrypted disk on the cluster, enter the encryption key in use on the cluster.
- Step 5** Click **Secure erase**.
- Step 6** In the **Erase this disk?** dialog box, click **Yes, erase this disk** to securely erase the encrypted disk.
- 

## Remote Key Management

The generic steps for remote KMIP certificate handling are as follows:

- If you are self-signing, specify local certificate authority in the configuration and get a root certificate.
- If you are using a trusted third-party CA, then specify that in the configuration and use their root certificate.
- Enter the root certificate in the HX encryption field that asks for the cluster key.
- Create an SSL server certificate and generate a Certificate Signing Request (CSR).
- Sign the CSR with whatever root certificate you are using.
- Update the KMIP server settings to use the client certificate.
- With the SSL certs and root CAs available, proceed with the KMIP service configuration specific to the vendor you have chosen.

### SafeNet Key Management

For details on managing encryption keys using a SafeNet key management server, see the [SafeNet Admin Guide](#).

### Vormetric Key Management

For details on managing encryption keys using a vormetric key management server, see the [Vormetric support portal](#) documentation downloads section.

# Configuring Remote Encryption Key

**Step 1** On the Cisco HyperFlex Connect navigation Pane, choose **Encryption**.

**Step 2** On the Encryption Page, click **Configure encryption**.

**Step 3** Enter the following Cisco UCS Manager credentials.

| UI Element                  | Essential Information                                                                                |
|-----------------------------|------------------------------------------------------------------------------------------------------|
| UCS Manager host name field | Cisco UCS Manager cluster host name.<br>Enter an IP address or FQDN.<br><eng-fi12.eng.storvisor.com> |
| User name field             | <admin> username                                                                                     |
| Password field              | <root> password                                                                                      |

Click **Next**.

**Step 4** To secure the HyperFlex cluster using a remote security key generated by the key management (KMIP) server, select **Key Management Server**.

You can configure a server with Self-Encrypting Drives in the cluster to use one of the following certificates.

- **Use certificate authority signed certificates**—Generate Certificate Signing Requests (CSRs) signed by an external certificate authority.
- **Use self-signed certificates**—Generate self-signed certificates.

Click **Next**.

**Step 5**

## What to do next

You can generate certificate signing requests or self-signed certificates.

# Generating Certificate Signing Requests

**Step 1** On the Cisco HyperFlex Connect navigation Pane, choose **Encryption**.

**Step 2** On the Encryption Page, click **Configure encryption**.

**Step 3** Enter the following Cisco UCS Manager credentials.

| UI Element                  | Essential Information                                                                                |
|-----------------------------|------------------------------------------------------------------------------------------------------|
| UCS Manager host name field | Cisco UCS Manager cluster host name.<br>Enter an IP address or FQDN.<br><eng-f112.eng.storvisor.com> |
| User name field             | <admin> username                                                                                     |
| Password field              | <admin> password                                                                                     |

Click **Next**.

**Step 4** Select **Key Management Server > Use certificate authority signed certificates**.

Click **Next**.

**Step 5** To generate the remote encryption key for configuring the key management (KMIP) server, complete the following details.

| UI Element                   | Essential Information                                                                                                 |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| Email address field          | <admin> email address.                                                                                                |
| Organization name field      | The organization requesting the certificate.<br>Enter up to 32 characters.                                            |
| Organization unit name field | The organizational unit.<br>Enter up to 64 characters.                                                                |
| Locality field               | The city or town in which the company requesting the certificate is headquartered.<br>Enter up to 32 characters.      |
| State field                  | The state or province in which the company requesting the certificate is headquartered.<br>Enter up to 32 characters. |
| Country field                | The country in which the company resides.<br>Enter two alphabetic characters in uppercase.                            |
| Valid for (days) field       | The validity period of the certificate.                                                                               |

**Step 6** To generate Certificate Signing Requests (CSRs) for all the HyperFlex nodes and download them, click **Generate certificates**.

**Step 7** Download the certificates to get them signed by a certificate authority. Click **Close**.

### What to do next

1. Upload the signed certificates.

2. Configure KMIP server (key management server).

# Configuring a Key Management Server Using CSRs (Certificate Signing Requests)

## Before you begin

Ensure that you have downloaded the generated CSRs on your local machine, signed it by a certificate authority and uploaded through the Cisco HX Data Platform UI for configuring the KMIP (key management) server.

- Step 1** On the Cisco HyperFlex Connect navigation Pane, choose **Encryption**.
- Step 2** On the Encryption Page, click **Continue configuration**.
- Step 3** From the **Continue configuration** drop-down list, select **Manage certificates** to upload the CSRs.
- Step 4** Enter the following Cisco UCS Manager credentials.

| UI Element                  | Essential Information                                                                                |
|-----------------------------|------------------------------------------------------------------------------------------------------|
| UCS Manager host name field | Cisco UCS Manager cluster host name.<br>Enter an IP address or FQDN.<br><eng-fi12.eng.storvisor.com> |
| User name field             | <admin> username                                                                                     |
| Password field              | <root> password                                                                                      |

Click **Next**.

- Step 5** Select **Upload certificate authority signed certificates**. Click **Next**.
- Step 6** Upload the CA signed certificate under **Upload new certificate**. Click **Upload**.
- Step 7** From the **Continue configuration** drop-down list select **Configure key management server** to configure the KMIP server.
- Step 8** Enter Cisco UCS Manager credentials to set up a primary key management server (KMIP) server and optionally a secondary KMIP server.

| UI Element                                       | Essential Information                                                                                               |
|--------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| Primary key management server field              | Enter the primary Key Management Server IP address.                                                                 |
| (Optional) Secondary key management server field | If you have a secondary key management server set up for redundancy, enter the details here.                        |
| Port number field                                | Enter the port number you wish to use for the key management servers.                                               |
| Public key field                                 | Enter the public root certificate of the certificate authority that you generated during KMIP server configuration. |

**Step 9** Click **Save** to encrypt the cluster with remotely managed keys.

### Example

## Generating Self-Signed Certificates

**Step 1** On the Cisco HyperFlex Connect navigation Pane, choose **Encryption**.

**Step 2** On the Encryption Page, click **Configure encryption**.

**Step 3** Enter the following Cisco UCS Manager credentials.

| UI Element                  | Essential Information                                                                                |
|-----------------------------|------------------------------------------------------------------------------------------------------|
| UCS Manager host name field | Cisco UCS Manager cluster host name.<br>Enter an IP address or FQDN.<br><eng-fi12.eng.storvisor.com> |
| User name field             | <admin> username                                                                                     |
| Password field              | <root> password                                                                                      |

Click **Next**.

**Step 4** Select **Key Management Server > Use self-signed certificates**.

Click **Next**.

**Step 5** To generate the remote encryption key for configuring the key management (KMIP) server, complete the following details.

| UI Element                   | Essential Information                                                                                                 |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| Email address field          | <admin> email address.                                                                                                |
| Organization name field      | The organization requesting the certificate.<br>Enter up to 32 characters.                                            |
| Organization unit name field | The organizational unit.<br>Enter up to 64 characters.                                                                |
| Locality field               | The city or town in which the company requesting the certificate is headquartered.<br>Enter up to 32 characters.      |
| State field                  | The state or province in which the company requesting the certificate is headquartered.<br>Enter up to 32 characters. |

| UI Element             | Essential Information                                                                      |
|------------------------|--------------------------------------------------------------------------------------------|
| Country field          | The country in which the company resides.<br>Enter two alphabetic characters in uppercase. |
| Valid for (days) field | The validity period of the certificate.                                                    |

**Step 6** To generate self-signed certificates for all the HyperFlex nodes and download them, click **Generate certificates**.

**Step 7** Upload the signed certificates and configure KMIP server (key management server).

## Configuring a key management server using SSCs (Self-Signed Certificates)

### Before you begin

Ensure that you have downloaded the generated SSCs on your local machine to configure the KMIP (key management) server.

**Step 1** On the Cisco HyperFlex Connect navigation Pane, choose **Encryption**.

**Step 2** On the Encryption Page, click **Edit configuration**.

**Step 3** From the **Edit configuration** drop-down list, select **Manage certificates**.

**Step 4** Enter the following Cisco UCS Manager credentials, to set up a primary key management (KMIP) server and optionally a secondary KMIP server.

| UI Element                  | Essential Information                                                                                             |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------|
| UCS Manager host name field | Cisco UCS Manager cluster host name.<br>Enter an IP address or FQDN.<br><i>&lt;eng-fi12.eng.storvisor.com&gt;</i> |
| User name field             | <i>&lt;admin&gt;</i> username                                                                                     |
| Password field              | <i>&lt;admin&gt;</i> password                                                                                     |

Click **Next**.

**Step 5** Enter the primary and secondary key management (KMIP) server credentials.

| UI Element                                       | Essential Information                                                                        |
|--------------------------------------------------|----------------------------------------------------------------------------------------------|
| Primary key management server field              | Enter the primary Key Management Server IP address.                                          |
| (Optional) Secondary key management server field | If you have a secondary key management server set up for redundancy, enter the details here. |

| UI Element        | Essential Information                                                                                               |
|-------------------|---------------------------------------------------------------------------------------------------------------------|
| Port number field | Enter the port number you wish to use for the key management servers.                                               |
| Public key field  | Enter the public root certificate of the certificate authority that you generated during KMIP server configuration. |

**Step 6** Click **Save** to encrypt the cluster with remotely managed keys.

---

## Restart Encryption

---

Enter Cisco UCS Manager credentials to restart configuring the key management server or local key, for securely encrypting the HyperFlex cluster.

| UI Element                  | Essential Information                                                                                |
|-----------------------------|------------------------------------------------------------------------------------------------------|
| UCS Manager host name field | Cisco UCS Manager cluster host name.<br>Enter an IP address or FQDN.<br><eng-fi12.eng.storvisor.com> |
| User name field             | <admin> username                                                                                     |
| Password field              | <admin> password                                                                                     |

---



## CHAPTER 8

# Managing Datastores

---

- [Managing Datastores, on page 97](#)
- [Adding Datastores, on page 98](#)
- [Editing Datastores, on page 99](#)
- [Mounting Datastores, on page 99](#)
- [Unmounting Datastores, on page 100](#)
- [Deleting Datastores, on page 101](#)
- [Recovering from Partially Unmounted Datastores, on page 101](#)

## Managing Datastores

Datastores are logical containers used by the HX Data Platform to manage your storage usage and storage resources. Datastores are where the host places virtual disk files and other VM files. Datastores hide the specifics of physical storage devices and provide a uniform model for storing VM files.

You can add, refresh the list, edit name and size, delete, mount and unmount datastores from either the HX Connect UI or the HX Data Platform Plug-in UI. You can only rename an unpaired datastore that is unmounted. Renaming the datastore is not supported by the vCenter administrator interface.



### Important

- Keep the number of datastores to as few as possible to avoid startup delay and to keep clone savings high.
- Configuring more than 10 datastores could result in excessive startup delay.
- HX Native Snapshots are not supported with multiple datastores.
- If using an M4 node, never use either the HyperFlex NFS or local Springpath datastore for ESXi logging or coredump partition. If using an M5 node, you can use any left over space in the HyperFlex NFS or local Springpath datastore for these purposes.
- When VMs have flat vmdk files, one with thin provisioned and one with thick provisioned, the total combined storage usage of all flat VMDK files as reported by the vCenter/ESXi and HX-Connect could be higher than the Datastore usage itself reported by vCenter & HX-Connect. This could be due to ESXi and vCenter space reporting for each VM files ignoring the "uniqueBytes" attributes sent by underlying NFS storage in Extended stats and attributes via VAAI APIs.

- 
- Step 1** Choose an interface.
- From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Manage > Datastores**.
  - From HX Connect, select **Datastores**.
- Step 2** Create a new or select an existing datastore, to view options.
- Create a new datastore
  - Refresh the datastore list
  - Edit the datastore name and size
  - Delete the datastore
  - Mount the datastore on the host
  - Unmount the datastore from the host
- 

## Adding Datastores

Datastores are logical containers, similar to file systems, that hide specifics of physical storage and provide a uniform model for storing VM files. You can also use datastores to store ISO images and VM templates.

- 
- Step 1** Choose an interface.
- From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Manage > Datastores**.
  - From HX Connect, select **Datastores**.
- Step 2** Select the create datastore.
- Step 3** Enter a name for the datastore. vSphere Web Client enforces a 42 character limit for the datastore name. Assign each datastore a unique name.
- Step 4** Specify the datastore size. Choose **GB** or **TB** from the drop-down list.
- Step 5** Specify the data blocksize. From HX Connect, choose **8K** or **4K**. Default is 8K. In the HX Data Platform Plug-in, the default is assumed. For VDI workloads, default is 4k.
- Step 6** Click **OK** to accept your changes or **Cancel** to cancel all changes.
- Step 7** Verify the datastore. Click the **Refresh** icon if needed to display your new datastore.
- From HX Data Platform Plug-in, Click the **Manage > Datastores > Hosts** tab to see the mount status of the new datastore.
- If you check the datastore through the vSphere Client application, **host > Configuration > Datastores**, the Drive Type is listed as `Unknown`. This is expected vSphere behavior, to list NFS datastores as Unknown.
-

# Editing Datastores

A HX Data Platform datastore can be modified using the edit (pencil) option. Edit options are: 1. Change the datastore name, or 2. Change the datastore storage allocation. That is, the size of the datastore.



---

**Note** Do not rename datastores with controller VMs.

---

- Step 1** Choose an interface.
- From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > *cluster* > Manage > Datastores**.
  - From HX Connect, select **Datastores**.
- Step 2** Select a *datastore*.
- Step 3** Unmount the datastore.
- If you are only resizing the datastore, you do not need to unmount the datastore. Skip this step.
- Step 4** Click the **Edit** (pencil icon) datastore.
- Step 5** Change the datastore name and/or size, as needed. Click **OK**.
- Step 6** Remount the datastore, if you previously unmounted it.
- 

# Mounting Datastores

**Prepare to mount a datastore.**

- No VM, template, snapshot, or CD/DVD image resides on the datastore. This is the most common error while unmounting.
- Storage I/O Control is disabled for the datastore.
- The datastore is not used for vSphere HA heartbeat.
- The datastore is not used to host RDM metadata files. RDM is not supported.
- The datastore is not used as a scratch location.



**Note** You cannot select an NFS datastore as a destination for the persistent scratch location on ESXi. If you select the HX datastore for the persistent scratch location, it will be removed after the ESXi host reloads.

For all M5 servers, M.2 boot SSD is automatically selected for use as scratch. This is configured out of the box on any new install.

For HX240M4 (non-SED), Intel SSD is used for persistent logs/scratch (same applies on 220M5/240M5, but on a different local SSD).

For HX220M4 and HX240M4 (SED), there is no location to store the scratch partition. So, the only option is to use syslog for persistent logging over the network.

### Mount a datastore.

- 
- Step 1** Choose an interface.
- From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Manage > Datastores**.
  - From HX Connect, select **Datastores**.
- Step 2** Select a *datastore*.
- Step 3** Click the **Mount**.
- Step 4** Confirm to mount the datastore, click **OK**.
- 

## Unmounting Datastores

### Prepare to unmount a datastore.

- No VM, template, snapshot, or CD/DVD image resides on the datastore. This is the most common error while unmounting.
- Storage I/O Control is disabled for the datastore.
- The datastore is not used for vSphere HA heartbeat.
- The datastore is not used to host RDM metadata files. RDM is not supported.
- The datastore is not used as a scratch location.

### Unmount a datastore.

- 
- Step 1** Choose an interface.
- From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Manage > Datastores**.

- From HX Connect, select **Datastores**.

**Step 2** Select a *datastore*.

**Step 3** Click the **Unmount**.

**Step 4** Confirm to unmount the datastore, click **OK**.

**Step 5** **If needed, recover from partial unmounts.**

- a) Go through the above checklist and unmount or delete through one of the UIs or CLI again.
- b) Use the UI or CLI to re-mount the datastore.

For additional information on recovering from partial unmounts, see [Recovering from Partially Unmounted Datastores](#), on page 101.

---

## Deleting Datastores

**Prepare to delete the datastores.**

- Power off all VMs.
- Close all open shells on the datastore mount point.
- Disable HA on the datastore.
- Close all applications that use the datastore.

**Delete datastores.**

---

**Step 1** Choose an interface.

- From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Manage > Datastores**.
- From HX Connect, select **Datastores**.

**Step 2** Select a *datastore*.

**Step 3** Click **Delete**.

**Step 4** Confirm to delete the datastore, click **OK**.

---

## Recovering from Partially Unmounted Datastores

When mounting, unmounting, or deleting datastores, sometimes a datastore can become partially unmounted. If this occurs, complete the following as needed.

---

**Step 1** Depending upon the task you are attempting, complete the items in Prepare to mount a datastore, Prepare to unmount a datastore, or Prepare to delete the datastores.

**Step 2** Retry to mount, unmount, or delete the datastore through the HX Connect or HX Data Platform Plug-in UI or CLI again.

**Step 3** If the datastore is not in the desired mount, unmount, or deleted state, complete the following.

- a) Ensure VMs are not running on the datastore.
- b) From ESX host, check to see if the HX Data Platform datastore is being used by VMware service, `storageRM`.

```
ls -ltra /vmfs/volumes/stfs-ds1/ | grep -i iorm
```

Sample response

```
-rwxr-xr-x 1 root root 16511 Jan 20 20:05 .iormstats.sf
drwxr-xr-x 1 root root 1125 Jan 20 20:06 .iorm.sf
```

- c) Check the `storagerm` status.

```
/etc/init.d/storagerm status
```

Sample response

```
storageRM is running
```

- d) Stop the `storagerm` service.

```
/etc/init.d/storagerm stop
```

Sample response

```
watchdog-storageRM: Terminating watchdog process with PID 34096
storageRM stopped
```

- e) Try to mount, unmount, or delete the datastore again.
  - f) This is one possible solution, if this doesn't resolve the issue, contact Technical Assistance Center (TAC).
-



## CHAPTER 9

# Managing Disks

---

- [Managing Disks in the Cluster, on page 103](#)
- [Disk Requirements, on page 103](#)
- [Replacing Self Encrypted Drives \(SEDs\), on page 105](#)
- [Replacing SSDs, on page 107](#)
- [Replacing NVMe SSDs, on page 108](#)
- [Replacing Housekeeping SSDs, on page 109](#)
- [Replacing or Adding Hard Disk Drives, on page 111](#)

## Managing Disks in the Cluster

Disks, SSDs or HDDs, might fail. If this occurs, you need to remove the failed disk and replace it. Follow the server hardware instructions for removing and replacing the disks in the host. The HX Data Platform identifies the SSD or HDD and incorporates it into the storage cluster.

To increase the datastore capacity of a storage cluster add the same size and type SSDs or HDDs to each converged node in the storage cluster. For hybrid servers, add hard disk drives (HDDs). For all flash servers, add SSDs.



### Note

When performing a hot-plug pull and replace on multiple drives from different vendors or of different types, pause for a few moments (30 seconds) between each action. Pull, pause for about 30 seconds and replace a drive, pause for 30 seconds. Then, pull, pause for 30 seconds and replace the next drive.

Sometimes, when a disk is removed it continues to be listed in cluster summary information. To refresh this, restart the HX cluster.

## Disk Requirements

The disk requirements vary between converged nodes and compute-only nodes. To increase the available CPU and memory capacity, you can expand the existing cluster with compute-only nodes as needed. These compute-only nodes provide no increase to storage performance or storage capacity.

Alternatively, adding converged nodes increase storage performance and storage capacity alongside CPU and memory resources.

Servers with only Solid-State Disks (SSDs) are All-Flash servers. Servers with both SSDs and Hard Disk Drives (HDDs) are hybrid servers.

The following applies to all the disks in a HyperFlex cluster:

- All the disks in the storage cluster must have the same amount of storage capacity. All the nodes in the storage cluster must have the same number of disks.
- All **SSDs** must support TRIM and have TRIM enabled.
- All **HDDs** can be either SATA or SAS type. All SAS disks in the storage cluster must be in a pass-through mode.
- Disk partitions must be removed from SSDs and HDDs. Disks with partitions are ignored and not added to your HX storage cluster.
- Optionally, you can remove or backup existing data on disks. All existing data on a provided disk is overwritten.




---

**Note** New factory servers are shipped with appropriate disk partition settings. Do not remove disk partitions from new factory servers.

---

- Only the disks ordered directly from Cisco are supported.
- On servers with Self Encrypting Drives (SED), both the cache and persistent storage (capacity) drives must be SED capable. These servers support Data at Rest Encryption (DARE).

In addition to the disks listed in the table below, all M4 converged nodes have 2 x 64-GB SD FlexFlash cards in a mirrored configuration with ESX installed. All M5 converged nodes have M.2 SATA SSD with ESXi installed.




---

**Note** Do not mix storage disks type or storage size on a server or across the storage cluster. Mixing storage disk types is not supported.

---

- When replacing cache or persistent disks, always use the same type and size as the original disk.
  - Do not mix any of the persistent drives. Use all HDD or SSD and the same size drives in a server.
  - Do not mix hybrid and All-Flash cache drive types. Use the hybrid cache device on hybrid servers and All-Flash cache devices on All-Flash servers.
  - Do not mix encrypted and non-encrypted drive types. Use SED hybrid or SED All-Flash drives. On SED servers, both the cache and persistent drives must be SED type.
  - All nodes must use same size and quantity of SSDs. Do not mix SSD types.
- 

Please refer to the corresponding server model spec sheet for details of drives capacities and number of drives supported on the different servers.

For information on compatible PIDs when performing an expansion of existing cluster, please refer to the [Cisco HyperFlex Drive Compatibility](#) document.

### Compute-Only Nodes

The following table lists the supported compute-only node configurations for compute-only functions. Storage on compute-only nodes is not included in the cache or capacity of storage clusters.



**Note** When adding compute nodes to your HyperFlex cluster, the compute-only service profile template automatically configures it for booting from an SD card. If you are using another form of boot media, update the local disk configuration policy. See the *Cisco UCS Manager Server Management Guide* for server-related policies.

| Supported Compute-Only Node Servers                                                                                                                                                                                                         | Supported Methods for Booting ESXi                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>• Cisco B200 M3/M4/M5</li> <li>• B260 M4</li> <li>• B420 M4</li> <li>• B460 M4</li> <li>• C240 M3/M4/M5</li> <li>• C220 M3/M4/M5</li> <li>• C460 M4</li> <li>• C480 M5</li> <li>• B480 M5</li> </ul> | <p>Choose any method.</p> <p><b>Important</b> Ensure that only one form of boot media is exposed to the server for ESXi installation. Post install, you may add in additional local or remote disks.</p> <p>USB boot is not supported for HX Compute-only nodes.</p> <ul style="list-style-type: none"> <li>• SD Cards in a mirrored configuration with ESXi installed.</li> <li>• Local drive HDD or SSD.</li> <li>• SAN boot.</li> <li>• M.2 SATA SSD Drive.</li> </ul> <p><b>Note</b> HW RAID M.2 (UCS-M2-HWRAID and HX-M2-HWRAID) is a supported boot configuration starting with HX Data Platform version 4.5(1a) and later.</p> |

## Replacing Self Encrypted Drives (SEDs)

Cisco HyperFlex Systems offers Data-At-Rest protection through Self-Encrypting Drives (SEDs) and Enterprise Key Management Support.

- Servers that are data at rest capable refer to servers with self encrypting drives.
- All servers in an encrypted HX Cluster must be data at rest capable.
- Encryption is configured on a HX Cluster, after the cluster is created, using HX Connect.
- Servers with self encrypting drives can be either solid state drive (SSD) or hybrid.

**Important**

To ensure the encrypted data remains secure, the data on the drive must be **securely erased** prior to removing the SED.

**Before you begin**

Determine if the encryption is applied to the HX Cluster.

- **Encryption not configured**—No encryption related prerequisite steps are required to remove or replace the SED. See [Replacing SSDs, on page 107](#) or [Replacing or Adding Hard Disk Drives, on page 111](#) and the hardware guide for your server.
- **Encryption is configured**—Ensure the following:
  1. If you are replacing the SED, obtain a Return to Manufacturer Authorization (RMA). Contact TAC.
  2. If you are using a local key for encryption, locate the key. You will be prompted to provide it.
  3. To prevent data loss, ensure the data on the disk is not the last primary copy of the data.  
If needed, add disks to the servers on the cluster. Initiate or wait until a rebalance completes.
  4. Complete the steps below before removing any SED.

- 
- Step 1** Ensure the HX Cluster is healthy.
- Step 2** Login to HX Connect.
- Step 3** Select **System Information** > **Disks** page.
- Step 4** Identify and verify the disk to remove.
- a. Use the Turn On Locator LED button.
  - b. Physically view the disks on the server.
  - c. Use the Turn Off Locator LED button.
- Step 5** Select the corresponding **Slot** row for the disk to be removed.
- Step 6** Click **Secure erase**. This button is available only after a disk is selected.
- Step 7** If you are using a local encryption key, enter the **Encryption Key** in the field and click **Secure erase**.  
If you are using a remote encryption server, no action is needed.
- Step 8** Confirm deleting the data on this disk, click **Yes, erase this disk**.  
**Warning** This deletes all your data from the disk.
- Step 9** Wait until the **Status** for the selected **Disk Slot** changes to **Ok To Remove**, then physically remove the disk as directed.
-

**What to do next**

**Note** Do not reuse a removed drive in a different server in this, or any other, HX Cluster. If you need to reuse the removed drive, contact TAC.

1. After securely erasing the data on the SED, proceed to the disk replacing tasks appropriate to the disk type: SSD or hybrid.

Check the **Type** column for the disk type.

- **Solid State** (SSDs)—See [Replacing SSDs, on page 107](#) and the hardware guide for your server.
- **Rotational** (hybrid drives)—See [Replacing or Adding Hard Disk Drives, on page 111](#) and the hardware guide for your server.

2. Check the status of removed and replaced SEDs.

When the SED is removed:

- **Status**—Remains **Ok To Remove**.
- **Encryption**—Changes from **Enabled** to **Unknown**.

When the SED is replaced, the new SED is automatically consumed by the HX Cluster. If encryption is not applied, the disk is listed the same as any other consumable disk. If encryption is applied, the security key is applied to the new disk.

- **Status**—Transitions from **Ignored** > **Claimed** > **Available**.
- **Encryption**—Transitions from **Disabled** > **Enabled** after the encryption key is applied.

## Replacing SSDs

The procedures for replacing an SSD vary depending upon the type of SSD. Identify the failed SSD and perform the associated steps.



**Note** Mixing storage disks type or size on a server or across the storage cluster is not supported.

- Use all HDD, or all 3.8 TB SSD, or all 960 GB SSD
- Use the hybrid cache device on hybrid servers and all flash cache devices on all flash servers.
- When replacing cache or persistent disks, always use the same type and size as the original disk.

### Step 1 Identify the failed SSD.

- For cache or persistent SSDs, perform a disk beacon check. See [Setting a Beacon, on page 71](#).

Only cache and persistent SSDs respond to the beacon request. NVMe cache SSDs and housekeeping SSDs do not respond to beacon requests.

- For cache NVMe SSDs, perform a physical check. These drives are in Drive Bay 1 of the HX servers.
- For housekeeping SSDs on HXAF240c or HX240c servers, perform a physical check at the back of the server.
- For housekeeping SSDs on HXAF220c or HX220c servers, perform a physical check at Drive Bay 2 of the server.

**Step 2** If the failed SSD is a housekeeping SSD, proceed based on the type of server.

- For HXAF240c M4 or HX240c M4 servers, contact Technical Assistance Center (TAC).
- For all other HX servers, see [Replacing Housekeeping SSDs, on page 109](#).

**Step 3** If a failed SSD is a cache or persistent SSD, proceed based on the type of disk.

- For NVMe SSDs, see [Replacing NVMe SSDs, on page 108](#).
- For all other SSDs, follow the instructions for removing and replacing a failed SSD in the host, per the server hardware guide.

After the cache or persistent drive is replaced, the HX Data Platform identifies the SDD and updates the storage cluster.

When disks are added to a node, the disks are immediately available for HX consumption.

**Step 4** To enable the Cisco UCS Manager to include new disks in the **UCS Manager > Equipment > Server > Inventory > Storage** tab, re-acknowledge the server node. This applies to cache and persistent disks.

**Note** Re-acknowledging a server is disruptive. Place the server into HX Maintenance Mode before doing so.

**Step 5** If you replaced an SSD, and see a message *Disk successfully scheduled for repair*, it means that the disk is present, but is still not functioning properly. Check that the disk has been added correctly per the server hardware guide procedures.

## Replacing NVMe SSDs

The procedures for replacing an SSD vary depending upon the type of SSD. This topic describes the steps for replacing NVMe cache SSDs.



**Note** Mixing storage disks type or size on a server or across the storage cluster is not supported. When replacing NVMe disks, always use the same type and size as the original disk.

### Before you begin

Ensure the following conditions are met when using NVMe SSDs in HX Cluster servers.

- NVMe SSDs are supported in HX240 and HX220 All-Flash servers.
- Replacing NVMe SSDs with an HGST SN200 disk requires HX Data Platform version 2.5.1a or later.

- NVMe SSDs are only allowed in slot 1 of the server. Other server slots do not detect NVMe SSDs.
- NVMe SSDs are only used for cache.
  - Using them for persistent storage is not supported.
  - Using them as the housekeeping drive is not supported.
  - Using them for hybrid servers is not supported.

- 
- Step 1** Confirm the failed disk is an NVMe cache SSD.
- Perform a physical check. These drives are in Drive Bay 1 of the HX servers. NVMe cache SSDs and housekeeping SSDs do not respond to beacon requests.
- If the failed SSD is not an NVMe SSD, see [Replacing SSDs, on page 107](#).
- Step 2** Put ESXi host into HX Maintenance Mode.
- Login to HX Connect.
  - Select **System Information > Nodes > node > Enter HX Maintenance Mode**.
- Step 3** Follow the instructions for removing and replacing a failed SSD in the host, per the server hardware guide.
- Note** When you remove an HGST NVMe disk, the controller VM will fail until you reinsert a disk of the same type into the same slot or reboot the host.
- After the cache or persistent drive is replaced, the HX Data Platform identifies the SDD and updates the storage cluster. When disks are added to a node, the disks are immediately available for HX consumption.
- Step 4** Reboot the ESXi host. This enables ESXi to discover the NVMe SSD.
- Step 5** Exit ESXi host from HX Maintenance Mode.
- Step 6** To enable the Cisco UCS Manager to include new disks in the **UCS Manager > Equipment > Server > Inventory > Storage** tab, re-acknowledge the server node. This applies to cache and persistent disks.
- Note** Re-acknowledging a server is disruptive. Place the server into HX Maintenance Mode before doing so.
- Step 7** If you replaced an SSD, and see a message *Disk successfully scheduled for repair*, it means that the disk is present, but is still not functioning properly. Check that the disk has been added correctly per the server hardware guide procedures.
- 

## Replacing Housekeeping SSDs



**Note** This procedure applies to HXAF220c M4, HX220c M4, HXAF220c M5, HX220c M5, HXAF240c M5, HX240c M5, servers only. To replace the housekeeping SSD on an HXAF240c M4 or HX240c M4 servers, contact Cisco TAC.

---

Identify the failed housekeeping SSD and perform the associated steps.

- 
- Step 1** Identify the failed housekeeping SSD.
- Physically check the SSD drives, as housekeeping drives are not listed through a beacon check.
- Step 2** Remove the SSD and replace with a new SSD of the same supported kind and size. Follow the steps in the server hardware guide.
- The server hardware guide describes the physical steps required to replace the SSD.
- Note** Before performing the hardware steps, enter the node into Cisco HX Maintenance Mode. After performing the hardware steps, exit the node from Cisco HX Maintenance Mode.
- Step 3** Using `SSH`, login into the storage controller VM of the affected node and run the following command.
- ```
# /usr/share/springpath/storfs-appliance/config-bootdev.sh -r -y
```
- This command consumes the new disk, adding it into the storage cluster.
- Sample response
- ```
Creating partition of size 65536 MB for /var/stv ...
Creating ext4 filesystem on /dev/sdg1 ...
Creating partition of size 24576 MB for /var/zookeeper ...
Creating ext4 filesystem on /dev/sdg2 ...
Model: ATA INTEL SSDSC2BB12 (scsi)
Disk /dev/sdg: 120034MB
Sector size (logical/physical): 512B/4096B
Partition Table: gpt
discovered. Rebooting in 60 seconds
```
- Step 4** Wait for the storage controller VM to automatically reboot.
- Step 5** When the storage controller VM completes its reboot, verify that partitions are created on the newly added SSD. Run the command.
- ```
# df -ah
```
- Sample response
- ```
.....
/dev/sdb1 63G 324M 60G 1%
/var/stv /dev/sdb2 24G 173M 23G 1% /var/zookeeper
```
- Step 6** Identify the HX Data Platform installer package version installed on the existing storage cluster.
- ```
# stcli cluster version
```
- The same version must be installed on all the storage cluster nodes. Run this command on the controller VM of any node in the storage cluster, but not the node with the new SSD.
- Step 7** Copy the HX Data Platform installer packages into the storage controller VM in `/tmp` folder.
- ```
cd /tmp
tar zxvf storfs-packages-<version>.tgz
```
- Step 8** Run the HX Data Platform installer deployment script.
- ```
# ./inst-packages.sh
# chmod 640 /usr/share/springpath/storfs-misc/springpath_security.properties
```

Note This is applicable for all affected nodes in the cluster.

Note You may need to restart the service (i.e., tomcat8) after the permission change.

For additional information on installing the HX Data Platform, see the appropriate [Cisco HX Data Platform Install Guide](#).

Step 9 After the package installation, HX Data Platform starts automatically. Check the status.

```
# status storfs
```

Sample response

```
storfs running
```

The node with the new SSD re-joins the existing cluster and the cluster returns to a healthy state.

Replacing or Adding Hard Disk Drives



Note Mixing storage disks type or size on a server or across the storage cluster is not supported.

- Use all HDD, or all 3.8 TB SSD, or all 960 GB SSD
- Use the hybrid cache device on hybrid servers and all flash cache devices on all flash servers.
- When replacing cache or persistent disks, always use the same type and size as the original disk.

Step 1 Refer to the hardware guide for your server and follow the directions for adding or replacing disks.

Step 2 Add HDDs of the same size to each node in the storage cluster.

Step 3 Add the HDDs to each node within a reasonable amount of time.

The storage starts being consumed by storage cluster immediately.

The vCenter Event log displays messages reflecting the changes to the nodes.

Note When disks are added to a node, the disks are immediately available for HX consumption although they will not be seen in the UCSM server node inventory. This includes cache and persistent disks. To include the disks in the **UCS Manager > Equipment > Server > Inventory > Storage** tab, re-acknowledge the server node.

Note Re-acknowledging a server is disruptive. Place the server into HX Maintenance Mode before doing so.



CHAPTER 10

Managing Nodes

- [Managing Nodes, on page 113](#)
- [Identify Node Maintenance Methods, on page 115](#)
- [Searching by DNS Address or Host Name, on page 117](#)
- [Changing ESXi Host Root Password, on page 118](#)
- [Reinstalling Node Software, on page 118](#)
- [Changing Node Identification Form in vCenter Cluster from IP to FQDN, on page 119](#)
- [Replacing Node Components, on page 120](#)
- [Removing a Node, on page 122](#)
- [Replacing a Node, on page 130](#)
- [Replacing a Compute Node, on page 134](#)

Managing Nodes

Nodes are initially added to a storage cluster using the Create Cluster feature of the HX Data Platform Installer. Nodes are added to an existing storage cluster using the Expand Cluster feature of the HX Data Platform Installer. When nodes are added or removed from the storage cluster, the HX Data Platform adjusts the storage cluster status accordingly.

- Tasks for node maintenance with a failed node.
 - The ESXi or HX software needs to be reinstalled.
 - A node component needs to be replaced.
 - The node needs to be replaced.
 - The node needs to be removed.
- Tasks for node maintenance with a non-failed node.
 - Putting the node into maintenance mode.
 - Changing the ESX password.



Note Though there are subtle differences, the terms **server**, **host**, and **node** are used interchangeably throughout the HyperFlex documentation. Generally a server is a physical unit that runs software dedicated to a specific purpose. A node is a server within a larger group, typically a software cluster or a rack of servers. Cisco hardware documentation tends to use the term node. A host is a server that is running the virtualization and/or HyperFlex storage software, as it is 'host' to virtual machines. VMware documentation tends to use the term host.

Step 1 Monitor the nodes in the cluster.

HX storage cluster, node, and node component status is monitored and reported to HX Connect, HX Data Platform Plug-in, vCenter UI, and assorted logs as Operational status (online, offline) and Resiliency (healthy, warning) status values.

Note Functional state distinctions contribute to, but are separate from, the storage cluster operational and resiliency status reported in the HX Connect and HX Data Platform Plug-in views. For each Data Replication Factor (2 or 3), Cluster Access Policy (lenient or strict), and given number of nodes in the storage cluster, the storage cluster shifts between Read and Write, Read Only, or Shutdown state, depending on the number of failed nodes or failed disks in nodes.

Note A replication factor of three is highly recommended for all environments except HyperFlex Edge. A replication factor of two has a lower level of availability and resiliency. The risk of outage due to component or node failures should be mitigated by having active and regular backups.

Step 2 Analyze the node failure and determine the action to take.

This frequently requires monitoring the node state through HX Connect, HX Data Platform Plug-in, vCenter, or ESXi; checking the server beacons; and collecting and analyzing logs.

Step 3 Complete the identified tasks.

- Reinstall or upgrade software.

For steps to reinstall ESXi or the HX Data Platform see [Cisco HyperFlex Systems Installation Guide for VMware ESXi](#). For steps to upgrade software, see the [Cisco HyperFlex Systems Upgrade Guide](#).

- Repair a component in the node.

Node components, such as solid state drives (SSD), hard disk drives (HDD), power supply units (PSU), and network interface cards (NIC) components are not configurable through HX Connect or HX Data Platform Plug-in, but the HX Data Platform monitors them and adjusts the storage cluster status when any of these items are disrupted, added, removed, or replaced.

The steps to add or remove disks, depends upon the type of disk. Field replaceable units (FRUs), such as PSUs and NICs are replaced following steps described in the server hardware guides.

- Replace a node in the cluster.

Replacing a node in a storage cluster typically requires TAC assistance. Provided the requirements are met, nodes can be replaced without TAC assistance while the storage cluster is online (5+ node clusters only) or offline (4+ node clusters). To replace a node in a 3 node cluster always requires TAC assistance. For more information, see [Removing a Node, on page 122](#).

- Remove a node from the cluster.

Note Removing the node must not reduce the number of available nodes below the minimum 3 nodes, as this makes the storage cluster unhealthy. To remove a node in a 3 node cluster always requires TAC assistance. You can remove a maximum of 2 nodes from an offline cluster. For more information, see [Replacing a Node, on page 130](#).

Identify Node Maintenance Methods

When performing maintenance tasks on nodes, some of these tasks are performed while the storage cluster is offline, others can be performed while the cluster is online and only require that the node is in HX maintenance mode.

- **Online tasks** - require that the storage cluster is healthy before the task begins.
- **Offline tasks** - require that the storage cluster will be shutdown.
If 2 or more nodes are down, then the storage cluster is automatically offline.
- **TAC assisted tasks** - typically require steps that are performed by the TAC representative.



Note There are several considerations to keep in mind before replacing a node. For more information, see [Replacing a Node, on page 130](#).

The following tables lists the methods available to perform the associated node maintenance task.

Repair Node Software

ESX and HX Data Platform software is installed on every node in the storage cluster. If it is determined after node failure analysis that either software item needs to be re-installed, see the [Cisco HyperFlex Systems Installation Guide for VMware ESXi](#). For steps to upgrade software, see the [Cisco HyperFlex Systems Upgrade Guide](#).

Repair Node Hardware

A reparable item on node fails. This includes FRUs and disks. Some node components require TAC assistance. Replacing a node's mother board, for example, requires TAC assistance.

No. Nodes in Cluster	No. Failed Nodes in Cluster	Method	Notes
3	1 or more	TAC assisted only node repair.	Node does not need to be removed to perform repair. Includes replacing disks on node.
4-8	1	Online or Offline node repair.	Node does not need to be removed to perform repair. Includes replacing disks on node.

Remove Node

A non-reparable item on node fails. Disks on the removed node are not reused in the storage cluster.

No. Nodes in Cluster	No. Failed Nodes in Cluster	Method	Notes
4	1	Offline node remove.	A 4 node cluster with 2 nodes down, requires TAC assistance.
5 or more	1	Online or Offline node remove.	
5 or more	2	Offline 2 node remove.	A 5 node cluster with 3 nodes down, requires TAC assistance.

Replace Node and Discard Storage

A non-reparable item on node fails. Disks on the removed node are not reused in the storage cluster.

No. Nodes in Cluster	No. Failed Nodes in Cluster	Method	Notes
3	1	TAC assisted only node replace.	TAC assisted node replacement required to return cluster to minimum 3 nodes. A 3 node cluster with 1 node down, requires TAC assistance.
4	1	Offline replace node. Not reusing the disks.	Use Expand cluster to add new nodes. All other nodes must be up and running. A 4 node cluster with 2 nodes down, requires TAC assistance.
5 or more	1	Online or offline replace node. Not reusing the disks.	Use Expand cluster to add new nodes. All other nodes must be up and running.
5 or more	2	Offline replace 1 or 2 nodes. Not reusing the disks.	Use Expand cluster to add new nodes. All other nodes must be up and running. Replacing up to 2 nodes is supported. Replacing 3 or more nodes requires TAC assistance.

Replace Node and Reuse Storage

A non-reparable item on node fails. Disks on the removed node are reused in the storage cluster.

No. Nodes in Cluster	No. Failed Nodes in Cluster	Method	Notes
3 or more	1 or more	TAC assisted only.	<p>TAC assisted node replacement required to return cluster to minimum 3 nodes.</p> <p>Note Reusing disks requires assigning old node UUID to new node. Disks UUIDs to node UUID relationship is fixed and cannot be reassigned. This is a TAC assisted task.</p>

Searching by DNS Address or Host Name

Sometimes for troubleshooting purposes it is useful to be able to search by the DNS server address or DNS server host name. This is an optional task.

Step 1

Assign DNS search addresses

- Login to the HX Data Platform Installer virtual machine. Use either `ssh` or the vSphere console interface.
- Edit `resolv.conf.d` file.

```
# vi /etc/resolvconf/resolv.conf.d/base
```

- Confirm the change.

```
# resolvconf -u
# cat /etc/resolv.conf
```

- Confirm the DNS server can be queried from either the IP address or the host name.

```
# nslookup ip_address
# nslookup newhostname
```

Step 2

Assign a DNS host name.

- Login to the HX Data Platform Installer virtual machine. Use either `ssh` or the vSphere console interface.
- Open the hosts file for editing.

```
# vi /etc/hosts
```

- Add the following line and save the file.

```
ip_address ubuntu newhostname
```

For each host *ip_address*, enter the host *newhostname*.

- Add the *newhostname* to *hostname*.

```
# hostname newhostname
```

Changing ESXi Host Root Password

You can change the default ESXi password for the following scenarios:

- During creation of a standard and stretch cluster (supports only converged nodes)
- During expansion of a standard cluster (supports both converged or compute node expansion)
- During Edge cluster creation



Note

In the above cases, the ESXi root password is secured as soon as installation is complete. In the event a subsequent password change is required, the procedure outlined below may be used after installation to manually change the root password.

As the ESXi comes up with the factory default password, you should change the password for security reasons. To change the default ESXi root password post-installation, do the following.



Note

If you have forgotten the ESXi root password, for password recovery please contact Cisco TAC.

Step 1 Log in to the ESXi host service control using SSH.

Step 2 Acquire root privileges.

```
su -
```

Step 3 Enter the current root password.

Step 4 Change the root password.

```
passwd root
```

Step 5 Enter the new password, and press **Enter**. Enter the password a second time for confirmation.

Note If the password entered the second time does not match, you must start over.

Reinstalling Node Software

To re-install software on a node that is a member of an existing storage cluster, contact TAC. This task must be performed with TAC assistance.

Step 1 Reinstall ESX following the directions from TAC.

Ensure the server meets the required hardware and configuration listed in Host ESX Server Setting Requirements. HX configuration settings are applied during the HX Data Platform process.

Step 2 Reinstall HX Data Platform, following the directions from TAC.

The HX Data Platform must always be re-installed after ESX is re-installed.

Changing Node Identification Form in vCenter Cluster from IP to FQDN

This task describes how to change how vCenter identifies the nodes in the cluster, from IP address to Fully Qualified Domain Name (FQDN).

Step 1 Schedule a maintenance window to perform this task.

Step 2 Ensure the storage cluster is healthy.

Check the storage cluster status through either HX Connect, HX Data Platform Plug-in, or from the `stcli cluster info` command on the storage controller VM.

Step 3 Lookup the FQDN for each ESXi host in the storage cluster.

a) From the ESXi host command line.

```
# cat /etc/hosts
```

In this example, the FQDN is `sjs-hx-3-esxi-01.sjs.local`.

```
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1        localhost.localdomain localhost
::1             localhost.localdomain localhost
172.16.67.157   sjs-hx-3-esxi-01.sjs.local sjs-hx-3-esxi-01
```

b) Repeat for each ESXi host in the storage cluster.

Step 4 Verify the FQDNs for each ESXi host are resolvable from vCenter, each other ESXi host, and the controller VMs.

a) From the vCenter command line.

```
# nslookup <fqdn_esx_host1>
# nslookup <fqdn_esx_host2>
# nslookup <fqdn_esx_host3>
...
```

b) Repeat for each ESXi host from an ESXi host.

c) Repeat for each ESXi host from each controller VM.

Step 5 If the FQDN name is not resolvable, then verify the DNS configuration on each ESXi host and each controller VM.

a) Check that the controller VMs have the correct IP address for the DNS server.

From a controller VM command line.

```
# stcli services dns show
10.192.0.31
```

a) Check the ESXi hosts have the same DNS configuration as the controller VMs.

From vCenter, select each ESXi host then **Configuration > DNS Servers**.

- Step 6** Locate and note the Datacenter Name and the Cluster Name.
- From vCenter client or web client, scroll through to see the Datacenter Name and Cluster Name. Write them down. They will be used in a later step.
- Step 7** Delete the **cluster** from vCenter.
- From vCenter, select **datacenter** > **cluster**. Right-click the **cluster** and select **Delete**.
- Note** Do not delete the **datacenter**.
- Step 8** Recreate the **cluster** in vCenter.
- From vCenter, right-click the **datacenter**. Select **New Cluster**.
 - Enter the exact same name for the **Cluster Name** as the cluster you deleted. This is the name you wrote down from Step 6.
- Step 9** Add ESXi hosts (nodes) to the **cluster** using the FQDN name. Perform these steps for all ESXi hosts.
- From vCenter, right-click the **datacenter** > **cluster**. Select **Add Host**.
 - Select an ESXi host using their FQDN.
 - Repeat for each ESXi host in the cluster.
- Step 10** Reregister the cluster with vCenter.
- ```
stcli cluster reregister
--vcenter-datacenter <datacenter_name>
--vcenter-cluster <hx_cluster_name>
--vcenter-url <FQDN_name>
--vcenter-user <vCenter_username>
--vcenter-password <vCenter_Password>
```
- The SSO URL is not required for HX version 1.8.1c or later. See [Registering a Storage Cluster with a New vCenter Cluster, on page 59](#) for additional information on reregistering a cluster.
- Step 11** Enable VMware cluster HA and DRS using the post install script:
- Login to the HX cluster IP as admin and run the command **# hx\_post\_install**.
  - Select Option 1 - "New/Existing Cluster" and input all login credentials
  - Type "y" if you want to enter a new license key
  - Type "y" to enable HA and DRS in the cluster
  - Select 'n' for all other options and exit the script.

## Replacing Node Components

Selected components on a node can be replaced. Some components can be replaced while the node is up and running. Replacing some components requires that the node be placed into a maintenance mode and shutdown. Refer to the hardware installation guide for your specific server for a complete list of field replaceable units (FRUs). Some components cannot be replaced or can only be replaced with TAC assistance. The following is a general list of components that can be replaced in a node.



**Note** When disks are removed, the disk UUIDs continue to be listed, even when not physically present. To reuse disks on another node in the same cluster see TAC for assistance.

- Components that do not require the node be shutdown. These are hot-swappable.
  - HDD data drives. Front bays  
See [Managing Disks](#) for the storage cluster tasks and the hardware installation guides for the hardware focused tasks. Both sets of tasks are required to replace this component.
  - SSD cache drive. Front bay 1  
See [Managing Disks](#) for the storage cluster tasks and the hardware installation guides for the hardware focused tasks. Both sets of tasks are required to replace this component.
  - Fan Modules  
See the hardware installation guides to replace this component.
  - Power Supplies  
See the hardware installation guides to replace this component.
- Components that do required the node be put into maintenance mode and shutdown.  
For all of the following components, see the hardware installation guides.

- Housekeeping SSD  
Both the storage cluster tasks, and hardware focused tasks are required to replace this component.
- RTC Battery on motherboard



**Note** The motherboard itself is not a replaceable component. You must purchase a battery from your local hardware store and replace it.

- DIMMS
- CPUs and Heatsinks
- Internal SD Card
- Internal USB Port
- Modular HBA Riser (HX 220c servers)
- Modular HBA Card
- PCIe Riser Assembly
- PCIe Card
- Trusted Platform Module
- mLOM Card

- RAID Controller
- Virtual Interface Card (VIC)
- Graphic Processing Unit (GPU)

## Removing a Node

Depending upon the node maintenance task, removing a node can be while the storage cluster is online or offline. Ensure you have completed the preparation steps before removing a node.



### Note

It is highly recommended that you work with your account team when removing a converged node in a storage cluster.

Do not reuse a removed converged node or its disks in the original cluster.

The affecting context is based on the number of converged nodes. The number of compute nodes does not affect the replacing a node workflow.

**Table 5: Removing a Node Workflows**

| Cluster Size   | Nodes Removed | Workflow                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 4 node cluster | 1             | <ol style="list-style-type: none"> <li>1. Cluster is healthy.</li> <li>2. Affected node in Maintenance mode.</li> <li>3. Shutdown the cluster (take cluster offline).<br/>Use the <code>stcli cluster shutdown</code> command.</li> <li>4. Remove the node.<br/>Use the <code>stcli node remove</code> command.</li> <li>5. Restart the cluster.<br/>Use the <code>stcli cluster start</code> command.</li> </ol> |
| 4 node cluster | 2 or more     | Workflow requires TAC assistance.                                                                                                                                                                                                                                                                                                                                                                                 |
| 5 node cluster | 1             | <ol style="list-style-type: none"> <li>1. Cluster is healthy.</li> <li>2. Affected node in Maintenance mode.</li> <li>3. Cluster remains online.</li> <li>4. Remove the node.<br/>Use the <code>stcli node remove</code> command.</li> </ol>                                                                                                                                                                      |

| Cluster Size   | Nodes Removed | Workflow                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 5 node cluster | 2             | <ol style="list-style-type: none"> <li>1. Cluster is healthy.</li> <li>2. Affected node in Maintenance mode.</li> <li>3. Shutdown the cluster (take cluster offline).<br/>Use the <code>stcli cluster shutdown</code> command.</li> <li>4. Remove the nodes.<br/>Use the <code>stcli node remove</code> command.<br/>Specify both nodes.</li> <li>5. Restart the cluster.<br/>Use the <code>stcli cluster start</code> command.</li> </ol> |
| 5 node cluster | 3 or more     | Workflow requires TAC assistance.                                                                                                                                                                                                                                                                                                                                                                                                          |

## Preparing to Remove a Node

Before you remove a node from a storage cluster, whether the cluster is online or offline, complete the following steps.



**Note** For all 3 node clusters, see TAC to assist with preparing, removing, and replacing a node.

**Step 1** Ensure the cluster is healthy.

```
stcli cluster info
```

Example response that indicates the storage cluster is online and healthy:

```
locale: English (United States)
state: online
upgradeState: ok
healthState: healthy
state: online
state: online
```

**Step 2** Ensure that SSH is enabled in ESX on all the nodes in the storage cluster.

**Step 3** Ensure that the Distributed Resource Scheduler (DRS) is enabled.

DRS migrates only powered-on VMs. If your network has powered-off VMs, you must manually migrate them to a node in the storage cluster that will not be removed.

**Note** If DRS is not available then manually move the Virtual Machines from the node.

**Step 4** Put the node to be removed into Cisco HX Maintenance mode. Choose a method: vSphere GUI or controller VM command line (CLI).

**GUI**

- a) From vSphere web client, select **Home > Hosts and Clusters > Hosts > host**.
- b) Right-click each host, scroll down the list, and select **Cisco HX Maintenance Mode > Enter HX Maintenance Mode**.

The vSphere Maintenance Mode option is at the top of the host right-click menu. Be sure to scroll to the bottom of the list to select Cisco HX Maintenance Mode.

## CLI

- a) On the ESX host, log in to a controller VM as a user with root privileges.
- b) Identify the node.

```
stcli node info

stNodes:

type: node
id: 689324b2-b30c-c440-a08e-5b37c7e0eefe
name: 192.168.92.144

type: node
id: 9314ac70-77aa-4345-8f35-7854f71a0d0c
name: 192.168.92.142

type: node
id: 9e6ba2e3-4bb6-214c-8723-019fa483a308
name: 192.168.92.141

type: node
id: 575ace48-1513-0b4f-bfe1-e6abd5ff6895
name: 192.168.92.143

```

Under `stNodes` section the `id` is listed for each node in the cluster.

- c) Move the ESX host into Maintenance mode.

```
stcli node maintenanceMode (--id ID | --ip NAME) --mode enter
```

(see also `stcli node maintenanceMode --help`)

## Step 5 Rebalance the storage cluster.

This ensures that all datastores associated with the node will be removed.

The rebalance command is used to realign the distribution of stored data across changes in available storage and to restore storage cluster health. If you add or remove a node in the storage cluster, you can manually initiate a storage cluster rebalance using the `stcli rebalance` command.

**Note** Rebalancing might take some time depending on the disk capacity used on the failed node or disk. Once rebalance is completed, confirm the cluster is healthy.

- a) Login to a controller VM in the storage cluster.
- b) From the controller VM command line, run the command:

```
stcli rebalance start --force
```

## Step 6 Open a command shell and login to the storage controller VM. For example, using `ssh`.

```
ssh root@controller_vm_ip
```

At the prompt, enter password, `Cisco123`.

### What to do next

Proceed to [Removing a Node](#). Choose the Online or Offline method per the condition of your storage cluster and the desired results listed in [Managing Nodes](#).

## Removing a Node from an Online Storage Cluster

Use the `stcli node remove` to clean up a deployment or remove a node from a storage cluster.

Removing a node from a storage cluster while the cluster remains online has slightly different requirements from removing a node while the cluster is offline.



**Note** You can remove multiple nodes in a series, so long as it is done one at a time and when the cluster is healthy between each successive node removal.

The affecting context is based on the number of converged nodes. The number of compute nodes does not affect the replacing a node workflow.

| Number of nodes in cluster                             | Method                                                                                                      |
|--------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| 3 node cluster                                         | See TAC to remove and replace the node.                                                                     |
| 4 node cluster                                         | Cluster must be offline. See <a href="#">Removing a Node from an Offline Storage Cluster, on page 127</a> . |
| 5 node cluster, removing 2 nodes                       | Cluster must be offline. See <a href="#">Removing a Node from an Offline Storage Cluster, on page 127</a> . |
| 5 node cluster, removing 1 node from a healthy cluster | Cluster can be online. Continue with the steps listed here.                                                 |



**Note** Do not remove the controller VM or other HX Data Platform components before you complete the steps in this task.

### Step 1

Complete the steps in [Preparing for Maintenance Operations](#) and [Preparing to Remove a Node](#). This includes:

- Ensure the cluster is healthy.

For 3 node clusters see TAC, as any node failure in a 3 node cluster means the cluster is not healthy.

- Ensure DRS is enabled or manually move the VMs from the node.
- Put the node being removed into HX maintenance mode.
- Rebalance the storage cluster.
- Login to the controller VM of a node that is not being removed.

- Step 2** Ensure the cluster is healthy.
- Step 3** Remove the desired node using the `stcli node remove` command.

`stcli node remove [-h] [--id-1 ID1 | --ip-1 NAME1] [--id-2 ID2 | --ip-2 NAME2]] [-f]`

| Syntax Description | Option              | Required or Optional | Description                                                                                                                                                                                                                 |
|--------------------|---------------------|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                    | <b>--id-1 ID1</b>   | One of set required. | A unique ID number for the storage cluster node. The ID is listed in the <code>stcli cluster info</code> command under the <code>stNode</code> field <code>id</code> .                                                      |
|                    | <b>--ip-1 NAME1</b> | One of set required. | IP address of storage cluster node. The IP is listed in the <code>stcli cluster info</code> command under the <code>stNode</code> field <code>name</code> .                                                                 |
|                    | <b>--id-2 ID2</b>   | Optional.            | A unique ID number for the storage cluster node. The ID is listed in the <code>stcli cluster info</code> command under the <code>stNode</code> field <code>id</code> .                                                      |
|                    | <b>--ip-2 NAME2</b> | Optional.            | IP address of storage cluster node. The IP is listed in the <code>stcli cluster info</code> command under the <code>stNode</code> field <code>name</code> .<br><br>The <code>--ip</code> option is currently not supported. |
|                    | <b>-f, --force</b>  | Optional.            | Forcibly remove storage cluster nodes.                                                                                                                                                                                      |

**Example:**

stNodes for a 5 node cluster:

```

type: node
id: 569c03dc-9af3-c646-8ac5-34b1f7e04b5c
name: example1

type: node
id: 0e0701a2-2452-8242-b6d4-bce8d29f8f17
name: example2

type: node
id: a2b43640-cf94-b042-a091-341358fdd3f4
name: example3

type: node
id: c2d43691-fab5-30b2-a092-741368dee3c4
name: example4

type: node
id: d2d43691-daf5-50c4-d096-941358fede374
name: example5
```

The `stcli node remove` command to remove nodes from the 5 node cluster are:

- To remove 1 node
  - `stcli node remove --ip-1 example5` or
  - `stcli node remove --id-1 d2d43691-daf5-50c4-d096-941358fede374`

After the `stcli node remove` command completes successfully, the system rebalances the storage cluster until the storage cluster state is **Healthy**. Do not perform any failure tests during this time. The storage cluster remains healthy.

Because the node is no longer in the storage cluster, you do not need to exit HX maintenance mode.

**Note** It is highly recommended that you work with TAC when removing a converged node in a storage cluster. Do not reuse a removed converged node or its disks in the original cluster.

**Note** If you want to reuse a removed node in another storage cluster, contact Technical Assistance Center (TAC). Additional steps are required to prepare the node for another storage cluster.

**Step 4** Confirm the node is removed from the storage cluster.

a) Check the storage cluster information.

```
stcli cluster info
```

b) Check the `ActiveNodes` entry in the response to verify the cluster has one less node.

**Step 5** Confirm all the node-associated datastores are removed.

**Note** If any node-associated datastores are listed, then manually unmount and delete those datastores.

**Step 6** Remove the host from the vCenter **Hosts and Cluster** view.

- Log in to vSphere Web Client Navigator. Navigate to **Host** in the vSphere Inventory.
- Right-click the host and select **Enter Maintenance Mode**. Click **Yes**.
- Right-click the host and select **All vCenter Actions > Remove from Inventory**. Click **Yes**.

**Step 7** Decommission the host from UCS Manager.

- Log in to UCS Manager. In the Navigation pane, click **Equipment**.
- Expand **Equipment > Chassis>Chassis Number>Servers**.
- Choose the HX server you want to decommission. In the work pane, click the **General** tab.
- In the **Actions** area, click **Server Maintenance**. In the **Maintenance** dialog box, click **Decommission**. Click **OK**.

## Removing a Node from an Offline Storage Cluster

Use the `stcli node remove` to clean up a deployment or remove a node from a storage cluster.



**Note** It is highly recommended that you work with TAC when removing a converged node in a storage cluster.

The affecting context is based on the number of converged nodes. The number of compute nodes does not affect the replacing a node workflow.

| Number of nodes in cluster       | Method                                  |
|----------------------------------|-----------------------------------------|
| 3 node cluster                   | See TAC to remove and replace the node. |
| 4 node cluster                   | Cluster must be offline.                |
| 5 node cluster, removing 2 nodes | Cluster must be offline.                |

| Number of nodes in cluster                             | Method                                                                                                   |
|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| 5 node cluster, removing 1 node from a healthy cluster | Cluster can be online. See <a href="#">Removing a Node from an Online Storage Cluster, on page 125</a> . |



**Note** Do not remove the controller VM or other HX Data Platform components before you complete the steps in this task.

You can remove a maximum of 2 nodes from an offline cluster.

**Step 1** Complete the steps in Preparing for Maintenance Operations and Preparing to Remove a Node. This includes:

a) Ensure the cluster is healthy.

For 3 node clusters see TAC, as any node failure in a 3 node cluster means the cluster is not healthy.

b) Ensure DRS is enabled or manually move the VMs from the node.

c) Rebalance the storage cluster.

d) Put the node being removed into HX maintenance mode.

e) Login to the controller VM of a node that is not being removed.

**Step 2** Prepare to shutdown, then shutdown the storage cluster.

This step is needed only for either of the following conditions:

- The cluster is less than 5 nodes.
- Removing 2 nodes from a 5 node cluster.

**Note** Do not use more than one node for removal if you have a cluster with RF=2.

a) Gracefully shutdown all resident VMs on all the HX datastores.

Optionally, vMotion the VMs.

b) Gracefully shutdown all VMs on non-HX datastores on HX storage cluster nodes, and unmount.

c) From any controller VM command line, issue the `stcli cluster shutdown` command.

```
stcli cluster shutdown
```

**Step 3** Remove the desired node using the `stcli node remove` command.

For example, you can specify the node to be removed by either IP address or domain name.

```
stcli node remove --ip-1 10.10.2.4 --ip-2 10.10.2.6
```

or

```
stcli node remove --name-1 esx.SVHOST144A.complab --name-2 esx.SVHOST144B.complab.lab
```

**Note** Enter the second IP address if you are removing a second node from a 5+ node storage cluster.

Response

```
Successfully removed node: EntityRef(type=3, id='', name='10.10.2.4' name='10.10.2.6')
```

This command unmounts all datastores, removes from the cluster ensemble, resets the EAM for this node, stops all services (stores, cluster management IP), and removes all firewall rules.

This command does not:

- Remove the node from vCenter. The node remains in vCenter.
- Remove the installed HX Data Platform elements, such as the controller VM.

Due to the node no longer being in the storage cluster, you do not need to exit HX maintenance mode.

**Note** If you want to reuse a removed node in another storage cluster, contact Technical Assistance Center (TAC). Additional steps are required to prepare the node for another storage cluster.

**Step 4** Restart the cluster.

```
stcli cluster start
```

**Step 5** To rebalance the storage cluster, run the rebalance command.

```
stcli rebalance start -f
```

**Note** Wait and confirm that rebalance has completed.

**Step 6** Confirm the node is removed from the storage cluster.

a) Check the storage cluster information.

```
stcli cluster info
```

b) Check the `ActiveNodes` entry in the response to verify the cluster has one less node.

**Step 7** Confirm all the node-associated datastores are removed.

**Note** If any node-associated datastores are listed, then manually unmount and delete those datastores.

---

## Removing a Compute Node

---

**Step 1** Migrate all the VMs from a compute node that needs to be removed.

**Step 2** Unmount the datastore from the compute node.

**Step 3** Check if the cluster is in the healthy state, by running the following command:

```
stcli cluster info --summary
```

**Step 4** Put ESXi host in the HX Maintenance mode.

**Step 5** Remove the compute node using the `stcli node remove` command, from CMIP (use the Cisco HX connect IP address as it is the cluster IP address).

```
stcli node remove --ip-1
```

Where, IP is the IP address of the node to be removed.

**Step 6** Remove any DVS from the ESXi host in vCenter, if there is a DVS.

**Step 7** Remove the ESXi host from vCenter.

- Step 8** Check if the cluster is in the healthy state, by running the following command:
- ```
stcli cluster info --summary
```
- Step 9** Clear stale entries in the compute node by logging out of Cisco HX Connect and then logging into Cisco HX Connect.
- Step 10** Disable and re-enable the High Availability (HA) and Distributed Resource Scheduler (DRS) services to reconfigure the services after node removal.

Deleting the Removed Node Data from a Disk and a Storage Controller VM

After removing a node from a storage cluster, you have to delete the removed node details from the disk and storage controller VM using the following procedure.



Warning Be aware that once the data is deleted you cannot recover it.

- Step 1** Destroy the cluster by running the following command:

```
run destroycluster -sxy
```

- Step 2** Remove the `stvbboot.cfg` configuration file from the `/etc/` folder.

- Step 3** Reboot the controller VM.

Note The reboot process takes few minutes.

- Step 4** After rebooting the controller VM, run the following command:

```
# for d in $(/bin/lsblk -dpn -e 1,2,7,11 | awk '{ print $1 }');do grep -qE "$d[0-9]" /proc/mounts && continue; dd if=/dev/zero of=$d bs=1M oflag=direct & done;
```

The data deletion action takes few hours. On completion of the drive data deletion, you will get the message: No space left on device. Ignore this message.

Replacing a Node

Replacing a node uses Expand Cluster to add the replacement node after removing the failed node. Replacing a node can be performed while the HX storage cluster is online or offline, provided the requirements are met. Replacing a converged node in a storage cluster typically requires TAC assistance.



Note It is highly recommended that you work with TAC when replacing a node in a storage cluster.

Conditions that require TAC assistance to replace a converged node.

- **3 node cluster**

All 3 node clusters require TAC assistance to replace a node. Replace the node during cluster maintenance.

• 4 node cluster

- Storage cluster is unhealthy.
- Storage cluster will become unhealthy if a node is removed.
- 2 or more nodes have failed.
- Disks on the replaced node will be reused.

When a node is added to a storage cluster, the HX Data Platform associates each disk UUID with the node UUID. This is a fixed relationship for the life of the storage cluster. Disks cannot be reassigned to nodes with different UUIDs. TAC will work with you to assign the old node's UUID to the new node to ensure the disk UUID to node UUID association.

- Storage cluster to remain online while replacing a node.

• 5 node cluster

- Storage cluster is unhealthy.
- Storage cluster will become unhealthy if a node is removed.
- 3 or more nodes have failed.
- Disks on the replaced node will be reused.

When a node is added to a storage cluster, the HX Data Platform associates each disk UUID with the node UUID. This is a fixed relationship for the life of the storage cluster. Disks cannot be reassigned to nodes with different UUIDs. TAC will work with you to assign the old node's UUID to the new node to ensure the disk UUID to node UUID association.

- Storage cluster to remain online while replacing 2 nodes.
- Storage cluster to remain online and the cluster was initially 3 or 4 nodes.

If your storage cluster was initially configured with either 3 or 4 nodes, adding nodes to make a total of 5 keeps your cluster as a 3+2 or 4+1 cluster. To keep the cluster online while replacing a node requires TAC assistance.

Workflows for Replacing a Node

The affecting context is based on the number of converged nodes. The number of compute nodes does not affect the replacing a node workflow.

Cluster Size	Nodes Replaced	Workflow
3 node cluster	1 or more	Workflow requires TAC assistance.

Cluster Size	Nodes Replaced	Workflow
4 node cluster	1	<ol style="list-style-type: none"> 1. Cluster is healthy. 2. Affected node in Cisco HX Maintenance mode. 3. Shutdown the cluster (take cluster offline). Use the <code>stcli cluster shutdown</code> command. 4. Remove the node. Use the <code>stcli node remove</code> command. 5. Restart the cluster. Use the <code>stcli cluster start</code> command. 6. Wait until cluster is online and healthy. 7. Use HX Installer > Expand Cluster to add replacement node. <p>Note Do not reuse the removed node or its disks in this or another cluster.</p>
4 node cluster	2 or more	Workflow requires TAC assistance.
5 node cluster	1	<ol style="list-style-type: none"> 1. Cluster is healthy. 2. Affected node in Cisco HX Maintenance mode. 3. Cluster remains online. 4. Remove the node. Use the <code>stcli node remove</code> command. 5. Restart the cluster. Use the <code>stcli cluster start</code> command. 6. Wait until cluster is online and healthy. 7. Use HX Installer > Expand Cluster to add replacement node. <p>Note Do not reuse the removed node or its disks in this or another cluster.</p>

Cluster Size	Nodes Replaced	Workflow
5 node cluster	2	<ol style="list-style-type: none"> 1. Cluster is healthy. 2. Affected node in Cisco HX Maintenance mode. 3. Shutdown the cluster (take cluster offline). Use the <code>stcli cluster shutdown</code> command. 4. Remove the nodes. Use the <code>stcli node remove</code> command. Specify both nodes. 5. Restart the cluster. Use the <code>stcli cluster start</code> command. 6. Wait until cluster is online and healthy. 7. Use HX Installer > Expand Cluster to add replacement node <p>Note Do not reuse the removed node or its disks in this or another cluster.</p>
5 node cluster	3 or more	Workflow requires TAC assistance.

Replacing a node and discarding the failed node's disks.

Step 1

Remove the old node. Follow the steps in the appropriate topic:

- *Removing a Node from an Online Storage Cluster*

Use this method only if the HX cluster was initially configured with at least 5 nodes and currently still has at least 5 nodes.

- *Removing a Node from an Offline Storage Cluster*

Use this method for all other non-TAC assisted node removal.

Note Even though you remove a node and its associated disks, the HX Data Platform remembers the disk UUIDs. When logs and reports are generated, messages indicate that the disks exist but cannot be found. Ignore these messages.

Step 2

Add the new node using the Expand option in the HX Data Platform Installer. See the *Cisco HyperFlex Systems Getting Started Guide*.

Replacing a Compute Node

If a compute node boot disk or blade is corrupted and the node needs to be replaced, perform the following steps:

1. Remove the compute node from the existing Hyper-V HyperFlex Cluster.
2. Reinstall OS and re-add the compute node into the cluster.



Note

Compute nodes are supported in HyperFlex release 3.5.2 and later releases.

This section provides the procedure for replacing a compute node that needs to be replaced due to faulty boot disk or blade.

Step 1 Use Hyper-V failover cluster manager and remove the bad compute node from the failover cluster manager.

Step 2 Clean up the computer object of the compute node from the Active Directory.

Note There is no need to clean up DNS entry of the compute node.

Step 3 Navigate to any controller VM and run the `remcomputenode.py` script to clean up the stale entries associated with the compute node.

The remove compute node Python script can be executed by providing either the UUID or host name of the compute node as an argument.

The following sample shows how to run the script with UUID of the compute node:

```
python remcomputenode.py -u C2581942-55D2-8021-B1B1-A117F396D671
```

The following sample shows how to run the script with host name of the compute node:

```
python remcomputenode.py -n node-hv1.cloud.local
```

Note Ensure that the following .egg files are available in the controller VM:

- `/usr/share/thrift-0.9.1.a-py2.7-linux-x86_64.egg`
- `/opt/springpath/storfs-mgmt-cli/stCli-1.0-py2.7.egg`

Step 4 Replace the faulty MB, compute blade, or boot disk.

Step 5 Run the compute node expansion workflow from the Installer VM.

- a) Install Windows 2016.
- b) On the **HX Data Platform Installer** page, select the **I know what I'm doing...** check box.
- c) Select the expansion workflow and complete the procedure.



CHAPTER 11

Expand Cisco HyperFlex System Clusters

- [Cluster Expansion Guidelines, on page 135](#)
- [Prerequisites When Creating a Mixed M4/M5 Cluster, on page 136](#)
- [Mixed Cluster Expansion Guidelines - Cisco HX Release 4.5\(x\), on page 137](#)
- [Steps During Mixed Cluster Expansion, on page 137](#)
- [Prerequisites for Adding a Converged \(HX220c/HX240c\) Node, on page 138](#)
- [Preparing a Converged Node, on page 138](#)
- [Adding a Converged Node to an Existing Cluster, on page 139](#)
- [Prerequisites for Adding a Compute-Only Node, on page 146](#)
- [Preparing a Compute-Only Node, on page 147](#)
- [Adding a Compute-Only Node to an Existing Cluster, on page 149](#)
- [Resolving Failure of Cluster Expansion, on page 155](#)
- [Enabling Logical Availability Zones, on page 156](#)

Cluster Expansion Guidelines

- If you have replication configured, put replication in pause mode before performing upgrade, expansion or cluster maintenance. After the upgrade, expansion or cluster maintenance is completed, then resume replication. Perform the pause and resume on any cluster that has replication configured to or from this local cluster.
- If you are using RESTful APIs to perform cluster expansion, sometimes the task may take longer than expected.
- ESXi installation is supported on SD cards for M4 converged nodes and M.2 SATA SSD for M5 converged nodes. For compute-only nodes, ESXi installation is supported for SD Cards, SAN boot, front SSD/HDD, or single M.2 SSD (using UCS-MSTOR-M2 controller). Installing ESXi on USB Flash is not supported for compute-only nodes



Note

HW RAID M.2 (UCS-M2-HWRAID and HX-M2-HWRAID) is a supported boot configuration starting with HX Data Platform version 4.5(1a) and later.

- You must click on the discovered cluster to proceed with expanding a standard ESX cluster in a 3.5.x or earlier release. Not doing so results in errors.

- Use only Admin credentials for the Controller VM during expansion work flow. Using any other credentials other than Admin may cause the expansion to fail.

ESXi Installation Guidelines

1. Modify boot policy for compute node.
To modify the template and boot policy for HyperFlex Stretched Cluster compute only node on M5 server:
 - a. Clone the template.
 - b. Uncheck the Flex flash from local boot policy, if the compute M5 node does not have flash cards.
 - c. Add the SAN boot with proper WWPN to the boot order.
2. Start the DPI expansion workflow.
3. When prompted, install ESXi using an ISO image.
4. Return to the DPI expansion workflow and complete the ESXi installation workflow.

**Note**

If the Hypervisor configuration fails with the SOL logging failure message, access the installer CLI through SSH with root and default password and configure the ESXi hypervisor. Then, run the advanced installer and check the **HX Storage Software** and **Expand Cluster** check boxes to proceed with the ESXi installation process.

Prerequisites When Creating a Mixed M4/M5 Cluster

The following tasks should be performed prior to beginning cluster expansion in a mixed cluster:

- **Hypercheck Health Check Utility**— Cisco recommends running this proactive health check utility on your HyperFlex cluster prior to upgrade. These checks provide early visibility into any areas that may need attention and will help ensure a seamless upgrade experience. For more information, see the [Hyperflex Health & Pre-Upgrade Check Tool](#) TechNote for full instructions on how to install and run Hypercheck.
- Upgrade UCS Manager to 3.2(1d) or later.
- Upgrade the existing HX cluster fully to 2.6(1a) or later.
- Download and deploy the matching 2.6(1a) or later HX Data Platform Installer to run the expansion workflow.
- Upgrade existing M4 server firmware to the latest 3.2(1) patch release.
- Upgrade vCenter to 6.5 or later. Without vCenter 6.5, Broadwell EVC mode cannot be enabled. Only vCenter upgrade is required. ESXi can remain on an older version subject to the VMware software interoperability matrix. Proceeding with EVC mode off is not supported and will cause operational issues in the future.

Mixed Cluster Expansion Guidelines - Cisco HX Release 4.5(x)

- Hypercheck Health Check Utility— Cisco recommends running this proactive health check utility on your HyperFlex cluster prior to upgrade. These checks provide early visibility into any areas that may need attention and help ensure a seamless **upgrade** experience. For more information, see the HyperFlexHealth & Pre-Upgrade Check Tool TechNote for full instructions on how to install and run Hypercheck.
- Expanding existing M4 cluster with M5 converged nodes is supported.
- Expanding existing M5 cluster with M4 converged nodes is not supported.
- Expanding existing mixed M4/M5 cluster with M4 or M5 converged nodes is supported.
- Adding any supported compute-only nodes is permitted with all M4, M5, and mixed M4/M5 clusters using the HX Data Platform 2.6 or later Installer. Some example combinations are listed here, many other combinations are possible.

Example combinations:

Expand mixed M4/M5 cluster with compute-only B200, C220, C240 M4/M5

Expand M4 cluster with compute-only B200 M5, C220 M5, C240M5

- Only expansion workflow is supported to create a mixed cluster. Initial cluster creation with mixed M4/M5 servers is not supported.
- All M5 servers must match the form factor (220/240), type (Hybrid/AF), security capability (Non-SED only) & disk configuration (QTY, capacity, and non-SED) of the existing M4 servers. For more information on drive compatibility, refer to the [Cisco Hyperflex Drive Compatibility](#) document.
 - HX220-M5 will use a maximum of 6 capacity disks (2 disk slots to remain empty) when mixed with HX220-M4.
- HX Edge, SED, LFF, Hyper-V, and Stretched Clusters do not support mixed M4 and M5 clusters.

Steps During Mixed Cluster Expansion

- During the validation steps, before expansion begins, an EVC check is performed. Follow the displayed guidance to manually enable EVC mode on the existing cluster at this time.



Caution

Failure to enable EVC at the time of the warning will require a complete shutdown of the storage cluster and all associated VMs at a later point in time. Do not skip this warning.

- Perform the EVC mode configuration in vCenter and then retry the validation.
- Cluster expansion will then validate a second time and then continue with the expansion.

Prerequisites for Adding a Converged (HX220c/HX240c) Node

A converged node can be added to a HyperFlex cluster after cluster creation. The storage on a converged node is automatically added to the cluster's storage capacity.

Before you start adding a converged node to an existing storage cluster, make sure that the following prerequisites are met.

- Ensure that the storage cluster state is healthy.
- Ensure that the new node meets the system requirements listed under **Installation Prerequisites**, including network and disk requirements.
- Ensure that the new node uses the same configuration as the other nodes in the storage cluster. This includes VLAN IDs and switch types (whether vSwitches), VLAN tagging with External Switch VLAN Tagging (EST), VLAN tagging with Virtual Switch Tagging (VST), or Virtual Distributed Switch.



Note If the storage cluster is in an out of space condition, when you add a new node, the system automatically rebalances the storage cluster. This is in addition to the rebalancing that is performed every 24 hours.

- Ensure that the node you add is of the same model (HX220 or HX240) type (Hybrid or All Flash), and disk configuration (SED or non-SED). In addition, ensure that the number of capacity disks matches the existing cluster nodes.
- To add a node that has a different CPU family from what is already in use in the HyperFlex cluster, enable EVC. For more details, see the *Setting up Clusters with Mixed CPUs* section in the *Cisco HyperFlex Systems Installation Guide for VMware ESXi*.
- Ensure that the software version on the node matches the Cisco HX Data Platform version, the ESXi version, and the vCenter version. To identify the software version, go to the Storage Cluster Summary tab in vCenter and check the HX Data Platform version in the top section. Upgrade if necessary.



Note If you upgraded the cluster, you must download and install a new installer VM, that matches the current version of HXDP running on the cluster.

- Ensure that the new node has at least one valid DNS and NTP server configured.
- If you are using SSO or Auto Support, ensure that the node is configured for SSO and SMTP services.
- Allow ICMP for ping between the HX Data Platform Installer and the existing cluster management IP address.

Preparing a Converged Node

Step 1 Connect the converged node to the hardware and the network of the existing storage cluster.

Step 2 Ensure that the HX node is a node prepared at factory.

Note Do not reuse a removed converged node or its disks in the original cluster.

Adding a Converged Node to an Existing Cluster



Note If you are using RESTful APIs to perform cluster expansion, the task may take longer than expected.

Step 1 Launch the Cisco HX Data Platform Installer.

- a) In your web browser, enter the IP address or the node name for the HX Data Platform Installer VM. Click **Accept** or **Continue** to bypass any SSL certificate errors. The Cisco HX Data Platform Installer login page appears. Verify the HX Data Platform Installer **Build ID** in the lower right corner of the login screen.
- b) In the login page, enter the following credentials:

Username: `root`

Password (Default): `Cisco123`

Note Systems ship with a default password of `Cisco123` that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.

- c) Read the EULA, check the **I accept the terms and conditions** checkbox, and click **Login**.

Step 2 On the **Workflow** page, select **Cluster Expansion**.

Step 3 On the **Credentials** page, complete the following fields.

To perform cluster creation, you can import a *JSON configuration* file with the required configuration data. The following two steps are optional if importing a JSON file, otherwise you can input data into the required fields manually.

Note For a first-time installation, contact your Cisco representative to procure the factory preinstallation JSON file.

- a. Click **Select a file** and choose your *JSON file* to load the configuration. Select **Use Configuration**.
- b. An **Overwrite Imported Values** dialog box displays if your imported values for Cisco UCS Manager are different. Select **Use Discovered Values**.

Field	Description
UCS Manager Credentials	
UCS Manager Host Name	UCS Manager FQDN or IP address. For example, <i>10.193.211.120</i> .
User Name	<admin> username.
Password	<admin> password.

Field	Description
vCenter Credentials	
vCenter Server	<p>vCenter server FQDN or IP address.</p> <p>For example, <i>10.193.211.120</i>.</p> <p>Note</p> <ul style="list-style-type: none"> • A vCenter server is required before the cluster can be made operational. • The vCenter address and credentials must have root level administrator permissions to the vCenter. • vCenter server input is optional if you are building a nested vCenter. See the Nested vCenter TechNote for more details.
User Name	<p><admin> username.</p> <p>For example, <i>administrator@vsphere.local</i>.</p>
Admin Password	<p><root> password.</p>
Hypervisor Credentials	
Admin User Name	<p><admin> username.</p> <p>This is root for factory nodes.</p>
Admin Password	<p><root> password.</p> <p>Default password is <i>Cisco123</i> for factory nodes.</p> <p>Note Systems ship with a default password of <i>Cisco123</i> that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.</p>

Step 4 Click **Continue**. A **Cluster Expand Configuration** page is displayed. Select the *HX Cluster* that you want to expand. If the HX cluster to be expanded is not found, or if loading the cluster takes time, enter the IP of the Cluster Management Address in the **Management IP Address** field.

Step 5 The **Server Selection** page displays a list of unassociated HX servers under the **Unassociated** tab, and the list of discovered servers under the **Associated** tab. Select the servers under the **Unassociated** tab to include in the HyperFlex cluster.

If HX servers do not appear in this list, check Cisco UCS Manager and ensure that they have been discovered.

For each server you can use the **Actions** drop-down list to set the following:

- **Launch KVM Console**—Choose this option to launch the KVM Console directly from the HX Data Platform Installer.

- **Disassociate Server**—Choose this option to remove a service profile from that server.

Note If there are no unassociated servers, the following error message is displayed:

No unassociated servers found. Please login to UCS Manager and ensure server ports are enabled.

The **Configure Server Ports** button allows you to discover any new HX nodes. Typically, the server ports are configured in Cisco UCS Manager before you start the configuration.

Step 6 Click **Continue**. The **UCSM Configuration** page appears.

Note If you imported a JSON file at the beginning, the **Credentials** page should be populated with the required configuration data from the preexisting HX cluster. This information must match your existing cluster configuration.

Step 7 In the **UCSM Configuration** page, complete the following fields for each network.

Field	Description
VLAN Configuration	
Note Use separate subnet and VLANs for each of the following networks.	
VLAN for Hypervisor and HyperFlex management	
VLAN Name	Name: hx-inband-mgmt
VLAN ID	Default VLAN ID: 3091
VLAN for HyperFlex storage traffic	
VLAN Name	Name: hx-storage-data
VLAN ID	Default VLAN ID: 3092
VLAN for VM vMotion	
VLAN Name	Name: hx-vmotion
VLAN ID	Default VLAN ID: 3093
VLAN for VM Network	
VLAN Name	Name: vm-network
VLAN ID(s)	Default VLAN ID: 3094 A comma-separated list of guest VLANs.
MAC Pool	
MAC Pool Prefix	Configure MAC Pool prefix by adding in two more hex characters (0-F). For example, 00:25:B5:A0.
'hx-ext-mgmt' IP Pool for Out-of-Band CIMC	

Field	Description
IP Blocks	The range of IP addresses designated for the HyperFlex nodes. This can be a comma-separated list of values for the guest VLANs. For example, <i>10.193.211.124-127, 10.193.211.158-163</i>
Subnet Mask	Set the subnet to the appropriate level to limit and control IP addresses. For example, <i>255.255.0.0</i> .
Gateway	IP address. For example, <i>10.193.0.1</i> .
iSCSI Storage	
Note This must be configured upfront if you want to use external storage at any point in the future.	
Enable iSCSI Storage checkbox	Check to configure iSCSI storage.
VLAN A Name	Name of the VLAN associated with the iSCSI vNIC, on the primary fabric interconnect (FI-A).
VLAN A ID	ID of the VLAN associated with the iSCSI vNIC, on the primary fabric interconnect (FI-A).
VLAN B Name	Name of the VLAN associated with the iSCSI vNIC, on the subordinate fabric interconnect (FI-B).
VLAN B ID	ID of the VLAN associated with the iSCSI vNIC, on the subordinate fabric interconnect (FI-A).
FC Storage	
Note This must be configured upfront if you want to use external storage at any point in the future.	
Enable FC Storage checkbox	Check to enable FC Storage.
WWxN Pool	A WWN pool that contains both WW node names and WW port names. For each fabric interconnect, a WWxN pool is created for WWPn and WWNN.
VSAN A Name	The name of the VSAN for the primary fabric interconnect (FI-A). By default, this is set to <i>hx-ext-storage-fc-a</i> .
VSAN A ID	The unique identifier assigned to the network for the primary fabric interconnect (FI-A). Caution Do not enter VSAN IDs that are currently used on the UCS or HyperFlex system. If you enter an existing VSAN ID in the installer which utilizes UCS zoning, zoning will be disabled in your existing environment for that VSAN ID.

Field	Description
VSAN B Name	The name of the VSAN for the subordinate fabric interconnect (FI-B). By default, this is set to <code>hx-ext-storage-fc-b</code> .
VSAN B ID	<p>The unique identifier assigned to the network for the subordinate fabric interconnect (FI-B).</p> <p>Caution Do not enter VSAN IDs that are currently used on the UCS or HyperFlex system. If you enter an existing VSAN ID in the installer which utilizes UCS zoning, zoning will be disabled in your existing environment for that VSAN ID.</p>
Advanced	
UCS Firmware Version	<p>Select the UCS firmware version to associate with the HX servers from the drop-down list. The UCS firmware version must match the UCSM version. See the latest Cisco HX Data Platform Release Notes for more details.</p> <p>For example, <i>3.2(1d)</i>.</p>
HyperFlex Cluster Name	The name applied to a group of HX Servers in a given cluster. This is a user-defined name. The HyperFlex cluster name adds a label to service profiles for easier identification.
Org Name	Displays a unique <i>Org Name</i> of the cluster, that ensures isolation of the HyperFlex environment from the rest of the UCS domain.

Note Review the VLAN, MAC pool, and IP address pool information in the **Configuration** pane. These VLAN IDs might be changed by your environment. By default, the Installer sets the VLANs as non-native. You must configure the upstream switches to accommodate the non-native VLANs by appropriately applying a trunk configuration.

Step 8

Click **Continue**. The **Hypervisor Configuration** page appears. Complete the following fields:

Attention You can skip the completion of the fields described in this step in case of a reinstall, and if ESXi networking has been completed.

Field	Description
Configure Common Hypervisor Settings	
Subnet Mask	<p>Set the subnet mask to the appropriate level to limit and control IP addresses.</p> <p>For example, <i>255.255.0.0</i>.</p>
Gateway	<p>IP address of gateway.</p> <p>For example, <i>10.193.0.1</i>.</p>

Field	Description
DNS Server(s)	<p>IP address for the DNS Server.</p> <p>If you do not have a DNS server, do not enter a hostname in any of the fields on the Cluster Configuration page of the HX Data Platform installer. Use only static IP addresses and hostnames for all ESXi hosts.</p> <p>Note If you are providing more than one DNS server, check carefully to ensure that both DNS servers are correctly entered, separated by a comma.</p>
Hypervisor Settings <p>Ensure to select Make IP Addresses and Hostnames Sequential, to make the IP addresses sequential.</p> <p>Note You can rearrange the servers using drag and drop.</p>	
Name	Server name.
Serial	Serial number of the server.
Static IP Address	Input static IP addresses and hostnames for all ESXi hosts.
Hostname	Do not leave the hostname fields empty.

Step 9 Click **Continue**. The **IP Addresses** page appears. You can add more compute or converged servers, by clicking **Add Compute Server** or **Add Converged Server**.

Ensure to select **Make IP Addresses Sequential**, to make the IP addresses sequential. For the IP addresses, specify if the network should belong to Data Network or Management Network.

For each HX node, complete the following fields for Hypervisor Management and Data IP addresses.

Field	Description
Management Hypervisor	Enter the static IP address that handles the Hypervisor management network connection between the ESXi host and the storage cluster.
Management Storage Controller	Enter the static IP address that handles the HX Data Platform storage controller VM management network connection between the storage controller VM and the storage cluster.
Data Hypervisor	Enter the static IP address that handles the Hypervisor data network connection between the ESXi host and the storage cluster.
Data Storage Controller	Enter the static IP address that handles the HX Data Platform storage controller VM data network connection between the storage controller VM and the storage cluster.

Field	Description
<p>When you enter IP addresses in the first row for Hypervisor (Management), Storage Controller VM (Management), Hypervisor (Data), and Storage Controller VM (Data) columns, the HX Data Platform Installer applies an incremental auto-fill to the node information for the rest of the nodes. The minimum number of nodes in the storage cluster is three. If you have more nodes, use the Add button to provide the address information.</p> <p>Note Compute-only nodes can be added only after the storage cluster is created.</p>	
Controller VM Password	<p>A default administrator username and password are applied to the controller VMs. The VMs are installed on all converged and compute-only nodes.</p> <p>Important</p> <ul style="list-style-type: none"> You cannot change the name of the controller VM or the controller VM's datastore. Use the same password for all controller VMs. The use of different passwords is not supported. Provide a complex password that includes 1 uppercase character, 1 digit, 1 special character, and a minimum of 10 characters in total. You can provide a user-defined password for the controller VMs and for the HX cluster to be created. For password character and format limitations, see the section on Guidelines for HX Data Platform Special Characters in the <i>Cisco HX Data Platform Management Guide</i>.
Advanced Configuration	
Jumbo frames Enable Jumbo Frames checkbox	<p>Check to set the MTU size for the storage data network on the host vSwitches and vNICs, and each storage controller VM.</p> <p>The default value is 9000.</p> <p>Note To set your MTU size to a value other than 9000, contact Cisco TAC.</p>
Disk Partitions Clean up Disk Partitions checkbox	<p>Check to remove all existing data and partitions from all nodes added to the storage cluster. You must backup any data that should be retained.</p> <p>Important Do not select this option for factory prepared systems. The disk partitions on factory prepared systems are properly configured. For manually prepared servers, select this option to delete existing data and partitions.</p>

Step 10 Click **Start**. A **Progress** page displays the progress of various configuration tasks.

Note If the vCenter cluster has EVC enabled, the deploy process fails with a message: The host needs to be manually added to vCenter. To successfully perform the deploy action, do the following:

- Log in to the ESXi host to be added in vSphere Client.
- Power off the controller VM.
- Add the host to the vCenter cluster in vSphere Web Client.
- In the HX Data Platform Installer, click **Retry Deploy**.

Step 11 When cluster expansion is complete, click **Launch HyperFlex Connect** to start managing your storage cluster.

Note When you add a node to an existing storage cluster, the cluster continues to have the same HA resiliency as the original storage cluster until auto-rebalancing takes place at the scheduled time.

Rebalancing is typically scheduled during a 24-hour period, either 2 hours after a node fails or if the storage cluster is out of space.

To rebalance the storage cluster before the scheduled time, perform the following steps to manually initiate rebalance storage cluster command.

- a. From a storage cluster controller VM command line, run the `# stcli rebalance start --force` command:
- b. To monitor rebalance status, run the `# stcli rebalance status` command.

Step 12 Create the required VM Network port groups and vMotion vmkernel interfaces using HyperFlex `hx_post_install` script or manually to match the other nodes in the cluster.

- a) SSH to HyperFlex cluster management IP.
- b) Log in as the admin user.
- c) Run the `hx_post_install` command.
- d) Follow the on-screen instructions, starting with vMotion and VM network creation. The other configuration steps are optional.

Step 13 After the new nodes are added to the storage cluster the High Availability (HA) services are reset so that HA can recognize the added nodes.

- a) Log in to vCenter.
- b) In the vSphere Web Client, navigate to the Host: **Home > vCenter > Inventory Lists > Hosts and Clusters > vCenter > Server > Datacenter > Cluster > Host**
- c) Select the new node.
- d) Right-click and select **Reconfigure for vSphere HA**.

Prerequisites for Adding a Compute-Only Node

A compute-only node can be added to a HyperFlex cluster, after cluster creation. It is added to provide extra compute resources. The Cisco UCS server does not need to have any caching or persistent drives as they do not contribute any storage capacity to the cluster.

Before you start adding a compute-only node, make sure that the following prerequisites are met.

- Ensure that the storage cluster state is healthy.
- Ensure that the new node meets the compute-only system requirements listed in *Installation Prerequisites*, including network and disk requirements.
- Install ESXi hypervisor after service profile association.
- Ensure that the new node uses the same configuration as the other nodes in the storage cluster. This includes VLAN IDs and switch types (whether vSwitches), VLAN tagging with External Switch VLAN Tagging (EST), VLAN tagging with Virtual Switch Tagging (VST), or Virtual Distributed Switch.



Note If the storage cluster is in an out of space condition, when you add a new node, the system automatically rebalances the storage cluster. This is in addition to the rebalancing that is performed every 24 hours.

- Enable EVC if the new node to be added has a different CPU family than what is already used in the HX cluster. For more details, see the *Setting up Clusters with Mixed CPUs* section in the *Cisco HyperFlex Systems Installation Guide for VMware ESXi*.
- Ensure that the software version on the node matches the Cisco HX Data Platform version, the ESXi version, and the vCenter version. To identify the software version, go to the **Storage Cluster Summary** tab in vCenter and check the *HX Data Platform version* in the top section. Upgrade if necessary.
- Ensure that the new node has at least one valid DNS and NTP server configured.
- If you are using SSO or Auto Support, ensure that the node is configured for SSO and SMTP services.
- ESXi installation is supported on SD cards for M4 converged nodes and M.2 SATA SSD for M5 converged nodes. For compute-only nodes, ESXi installation is supported for SD Cards, SAN boot, front SSD/HDD, or single M.2 SSD (using UCS-MSTOR-M2 controller). Installing ESXi on USB Flash is not supported for compute-only nodes.



Note HW RAID M.2 (UCS-M2-HWRAID and HX-M2-HWRAID) is a supported boot configuration starting with HX Data Platform version 4.5(1a) and later.

Preparing a Compute-Only Node

Step 1 Ensure that the server is a supported HX server and meets the requirements. For more details, see the *Host Requirements* section in the *Cisco HyperFlex Systems Installation Guide for VMware ESXi*.

Step 2 Log in to Cisco UCS Manager.

- a) Open a browser and enter the Cisco UCS Manager address for the fabric interconnect of the storage cluster network.
- b) Click the **Launch UCS Manager** button.
- c) If prompted, download, install, and accept Java.
- d) Log in with administrator credentials.

Username: admin

Password: <admin password>

- Step 3** Locate the server to ensure that the server has been added to the same FI domain as the storage cluster and is an approved compute-only model. Check the latest [Release Notes for Cisco HX Data Platform](#) for a full list of compatible Compute-only nodes.

Verify the HX Data Platform Installer

- Step 1** Verify that the HX Data Platform installer is installed on a node that can communicate with all the nodes in the storage cluster and compute nodes that are being added to the storage cluster.
- Step 2** If the HX Data Platform installer is not installed, see Deploy the HX Data Platform Installer.

Apply an HX Profile on a Compute-only Node Using UCS Manager

In Cisco UCS Manager the network policies are grouped into an HX profile. The HX installer handles automatic service profile association for compute-only nodes. Manual association is not required.

Once install begins, you should monitor compute-only node service profile association in UCS Manager. Wait until the server is fully associated before continuing on to install ESXi.

Install VMware ESXi on Compute Nodes



Important

Install VMware ESXi on each compute-only node.

Install a Cisco HX Data Platform supported version of ESXi. See the [Cisco HyperFlex Data Platform Release Notes](#) for a list of supported ESXi versions.

If the compute only node already has ESXi installed, it must be re-imaged with the Cisco HX Custom image.

Before you begin

Ensure the required hardware and network settings are met. For more details, see the *Installation Prerequisites* section in the *Cisco HyperFlex Systems Installation Guide for VMware ESXi*. Ensure the service profiles in the previous step have finished associating.

- Step 1** Download the *HX Custom Image for ESXi* from the Cisco.com download site for Cisco HyperFlex. See [Download Software](#).
- Select a networked location that can be accessed through Cisco UCS Manager.
- Step 2** Log in to Cisco UCS Manager.

- Step 3** Log in to the KVM console of the server through Cisco UCS Manager.
- In the Navigation Pane, click **Servers > Service Profiles > Sub-Organizations > *hx-cluster***.
 - Right click the *hx-cluster* and choose **KVM Console**.
- Step 4** Copy the *HX-Vmware.iso* image to the KVM path for the compute server.
- Example:**
- HX-Vmware-ESXi-60U3-5050593-Cisco-Custom-6.0.3.1.iso
- Step 5** From the KVM console session, select **Virtual Media > Map CD/DVD** and mount the *HX Custom Image for ESXi* image. If you do not see the **Map CD/DVD** option, first activate virtual devices.
- Select **Virtual Media > Activate Virtual Devices**.
- This opens in a pop-up window.
- Click **Accept the session > Apply**.
- Step 6** From the **Map CD/DVD** option, map to the location of the *HX-Vmware.iso* file.
- Select the *HX-Vmware.iso* file.
 - Select **Map Device**.
- There is a check mark indicating that the file is on a mapped location, once the process is complete. The mapped file's full name includes the ESXi build ID.
- Step 7** Reset the compute server.
- Click the **Reset** button on the KVM console. Click **OK** to confirm.
 - Select **Power Cycle**. Click **OK**.
- Step 8** Change the boot path to point to the *HX-Vmware.iso* file.
- Press **F6**.
 - From the **Enter boot selection** menu, use the arrow keys to highlight the *Cisco vKVM-Mapped vDVD1.22* option.
 - Press **Enter** to select.
- This launches the ESXi installer bootloader. Select one of the three compute-only node options based on desired boot type: SD Card, Local Disk, or Remote Disk. Type in **yes** (all lowercase) to confirm selection. The rest of the installation is automated. ESXi will reboot several times. It is normal to see warnings that automatically dismiss after a short wait period. Wait for the *ESXi DCUI* to fully appear, signaling the end of installation.
- Step 9** Repeat steps 3 to 8 for each Cisco HyperFlex server.
- Step 10** Once ESXi is fully installed, click **continue**. Then click **Retry Hypervisor Configuration** to complete the rest of the cluster expansion.

Adding a Compute-Only Node to an Existing Cluster

To add a HyperFlex compute-only node to an existing HyperFlex system cluster, complete the following steps.



Note If you are using RESTful APIs to perform cluster expansion, sometimes the task may take longer than expected.



Note After you add a compute-only node to an existing cluster, you must manually configure the vmk2 interface for vmotion.

Step 1

Launch the Cisco HX Data Platform Installer.

- In your web browser, enter the IP address or the node name for the HX Data Platform Installer VM. Click **Accept** or **Continue** to bypass any SSL certificate errors. The Cisco HX Data Platform Installer login page appears. Verify the HX Data Platform Installer **Build ID** in the lower right corner of the login screen.
- In the login page, enter the following credentials:

Username: root

Password (Default): Cisco123

Note Systems ship with a default password of Cisco123 that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.

- Read the EULA, check the **I accept the terms and conditions** checkbox, and click **Login**.

Step 2

On the **Workflow** page, select **Cluster Expansion**.

Step 3

On the **Credentials** page, complete the following fields.

To perform cluster creation, you can import a *JSON configuration* file with the required configuration data. The following two steps are optional if importing a JSON file, otherwise you can input data into the required fields manually.

- Note** For a first-time installation, contact your Cisco representative to procure the factory preinstallation JSON file.
- Click **Select a file** and choose your *JSON file* to load the configuration. Select **Use Configuration**.
 - An **Overwrite Imported Values** dialog box displays if your imported values for Cisco UCS Manager are different. Select **Use Discovered Values**.

Field	Description
UCS Manager Credentials	
UCS Manager Host Name	UCS Manager FQDN or IP address. For example, 10.193.211.120.
User Name	<admin> username.
Password	<admin> password.
vCenter Credentials	

Field	Description
vCenter Server	<p>vCenter server FQDN or IP address.</p> <p>For example, <i>10.193.211.120</i>.</p> <p>Note</p> <ul style="list-style-type: none"> • A vCenter server is required before the cluster can be made operational. • The vCenter address and credentials must have root level administrator permissions to the vCenter. • vCenter server input is optional if you are building a nested vCenter. See the Nested vCenter TechNote for more details.
User Name	<p><admin> username.</p> <p>For example, <i>administrator@vsphere.local</i>.</p>
Admin Password	<p><root> password.</p>
Hypervisor Credentials	
Admin User Name	<p><admin> username.</p> <p>This is root for factory nodes.</p>
Admin Password	<p><root> password.</p> <p>Default password is <i>Cisco123</i> for factory nodes.</p> <p>Note</p> <p>Systems ship with a default password of <i>Cisco123</i> that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.</p>

Step 4 Click **Continue**. A **Cluster Expand Configuration** page is displayed. Select the *HX Cluster* that you want to expand. If the HX cluster to be expanded is not found, or if loading the cluster takes time, enter the IP of the Cluster Management Address in the **Management IP Address field**.

Step 5 Click **Continue**. A **Server Selection** page is displayed. On the **Server Selection** page, the **Associated** tab lists all the HX servers that are already connected. Do not select them. On the **Unassociated** tab, select the servers you wish to add to the cluster.

Step 6 Click **Continue**. A **UCS Manager Configuration** page is displayed.

- Note**
- If you imported a JSON file at the beginning, the fields should be populated with the required configuration data from the pre-existing HX cluster. If the server to add into the cluster is not associated, select them and click **Continue**.
 - This information must match your existing cluster configuration. Ensure that you use the proper settings on this page.

Enter the following information for each network:

Field	Information
VLAN for ESXi and HyperFlex management	
VLAN Name	Name: hx-inband-mgmt
VLAN ID	Default VLAN ID: 3091
VLAN for HyperFlex Storage traffic	
VLAN Name	Name: hx-storage-data
VLAN ID	Default VLAN ID: 3092
VLAN for VM vMotion	
VLAN Name	Name: hx-vmotion
VLAN ID	Default VLAN ID: 3093
VLAN for VM Network	
VLAN Name	Name: hx-vm-network
VLAN ID	Default VLAN ID: 3094
MAC Pool	
MAC Pool Prefix	For example, 00:25:B5:99 (99 is default and can be set as per the user environment).
hx-ext-mgmt IP Pool for Out-of-Band CIMC	
IP Blocks	<p>The range of IP addresses designated for the HyperFlex nodes. This can be a comma-separated list of values for the guest VLANs.</p> <p>For example, 10.193.211.124-127, 10.193.211.158-163, 255.255.255.0.</p> <p>Enter IP address of gateway. For example, 10.193.0.1.</p>
Subnet	<p>Set the subnet to the appropriate level to limit and control IP addresses.</p> <p>For example, 255.255.0.0.</p>
Default gateway	<p>IP address.</p> <p>For example, 10.193.0.1.</p>
Advanced	

Field	Information
UCS Firmware version	Select the UCS firmware version to associate with the HX servers from the drop-down list. The UCS firmware version must match the UCSM version. See the latest Cisco HX Data Platform Release Notes for more details. For example, 3.2(1d).
HyperFlex Cluster Name	The name applied to a group of HX Servers in a given cluster. This is a user-defined name. The HyperFlex cluster name adds a label to service profiles for easier identification.

Note Review the VLAN, MAC pool, and IP address pool information in the Configuration pane. These VLAN IDs might be changed by your environment. By default, the Installer sets the VLANs as non-native. You must configure the upstream switches to accommodate the non-native VLANs by appropriately applying a trunk configuration.

Step 7

Click **Continue**. The **Hypervisor Configuration** page appears. Complete the following fields:

Attention You can skip the completion of the fields described in this step in case of a reinstall, and if ESXi networking has been completed.

Field	Description
Configure Common Hypervisor Settings	
Subnet Mask	Set the subnet mask to the appropriate level to limit and control IP addresses. For example, 255.255.0.0.
Gateway	IP address of gateway. For example, 10.193.0.1.
DNS Server(s)	IP address for the DNS Server. If you do not have a DNS server, do not enter a hostname in any of the fields on the Cluster Configuration page of the HX Data Platform installer. Use only static IP addresses and hostnames for all ESXi hosts. Note If you are providing more than one DNS server, check carefully to ensure that both DNS servers are correctly entered, separated by a comma.
Hypervisor Settings	
Ensure to select Make IP Addresses and Hostnames Sequential , to make the IP addresses sequential.	
Note You can rearrange the servers using drag and drop.	
Name	Server name.

Field	Description
Serial	Serial number of the server.
Static IP Address	Input static IP addresses and hostnames for all ESXi hosts.
Hostname	Do not leave the hostname fields empty.

Step 8

Click **Continue**. An **IP Addresses** page is displayed. Click **Add Compute-only Node** to add a new node.

If you are adding more than one compute-only node, select **Make IP Addresses Sequential**.

Field	Information
Management Hypervisor	Enter the static IP address that handles the Hypervisor management network connection between the ESXi host and storage cluster.
Management Storage Controller	None.
Data Hypervisor	Enter the static IP address that handles the Hypervisor data network connection between the ESXi host and the storage cluster.
Data Storage Controller	None.
Controller VM	Enter the default Admin username and password that were applied to controller VMs when they were installed on the existing HX Cluster. Note The name of the controller VM cannot be changed. Use the existing cluster password.

Step 9

Click **Start**. A **Progress** page displays the progress of various configuration tasks.

Note By default no user intervention is required if you are booting from FlexFlash (SD Card). However, if you are setting up your compute-only node to boot from a local disk, complete the following steps in Cisco UCS Manager:

- a. Click the service profile created by the HX Data Platform Installer.
For example, *blade-1(HX_Cluster_Name)*.
- b. On the **General** tab, click **Unbind from the Template**.
- c. In the working pane, click the **Storage** tab. Click the **Local Disk Configuration Policy** sub tab.
- d. In the Actions area, select **Change Local Disk Configuration Policy** > **Create Local Disk Configuration Policy**.
- e. Under **Create Local Disk Configuration Policy**, enter a name for the policy, and keep the rest as default. Click **Ok**.
- f. In the **Change Local Disk Configuration Policy** Actions area, select the newly created local disk configuration policy from the drop-down list. Click **Ok**.

- g. Now, go back to the HX Data Platform Installer UI and click **Continue**, and then click **Retry UCSM**

Compute Node Expansion - ESXi Installation Required

ESXi must be installed on all nodes being added at this point using the HX ESXi ISO on [cisco.com](https://www.cisco.com)

Using an existing installation of ESXi will cause installation to fail. Other ESXi ISOs other than the one posted on Cisco are not supported.

Once ESXi is installed, select Continue and then Retry to continue installation.
Full instructions can be found below.

If ESXi is already installed using the HX ESXi ISO wait for it to boot and then select Continue and Retry to continue installation.

 Instructions

 Launch UCS Manager

Continue

Configuration.

Note If the vCenter cluster has EVC enabled, the deploy process fails. The host needs to be manually added to vCenter. To successfully perform the deploy action, do the following:

- Log in to the ESXi host to be added in vSphere Client.
- Power off the controller VM.
- Add the host to the vCenter cluster in vSphere Web Client.
- In the HX installer, click **Retry Deploy**.

Step 10 When installation is complete, start managing your storage cluster by clicking **Launch HyperFlex Connect**.

Step 11 After the new nodes are added to the storage cluster, HA services are reset so that HA is able to recognize the added nodes.

- Log on to VMware vSphere Client.
- Select **Home > Hosts and Clusters > Datacenter > Cluster > Host**.
- Select the new node.
- Right-click and select **Reconfigure for vSphere HA**.

Step 12 After adding compute-only nodes to an existing cluster, you must manually configure the vmk2 interface for vmotion.

Resolving Failure of Cluster Expansion

If you receive an error dialog box and the storage cluster expansion doesn't complete, proceed with the resolution options listed below:

Step 1 **Edit Configuration** - Returns you to the Cluster Configuration page. You fix the issues listed in the validation page.

- Step 2** **Start Over** - Allows you to reverse the settings you applied by clearing progress table entries and you are returned to the Cluster Configuration page to restart a new deployment. See Technical Assistance Center (TAC).
- Step 3** **Continue** - Adds the node to the storage cluster in spite of the failure generating errors. See Technical Assistance Center (TAC).

Note Select the Continue button only if you understand the failures and are willing to accept the possibility of unpredictable behavior.

For more information about cleaning up a node for the purposes of redeploying HyperFlex, see the [HyperFlex Customer Cleanup Guides for FI and Edge](#).

Enabling Logical Availability Zones

Once you have successfully configured your HyperFlex cluster, refer to the commands below to enable or disable Logical Availability Zones (LAZ).



Important

Please be advised to wait for at least 10 seconds between successive invocations of LAZ disable and, LAZ enable operations in that order.

In case of fresh cluster installation, for best results, it is recommended that LAZ is enabled during cluster creation.

Command	Description
<code>stcli cluster get-zone</code>	Checks if zones are enabled.
<code>stcli cluster set-zone --zone 0</code>	Disables zones.
<code>stcli cluster set-zone --zone 1</code> <code>stcli rebalance start</code>	<p>(Recommended) Enables and creates zones (default number of zones)</p> <p>Important</p> <ul style="list-style-type: none"> • A cluster created without zoning enabled, will become zone compliant only after enabling zoning and successful completion of rebalance. • You must execute the rebalance start command after you enable and create zones. • Triggering rebalance activity may involve large scale data movements across several nodes in the cluster which may decrease the IO performance in the cluster.

Command	Description
<pre>stcli cluster set-zone --zone 1 --numzones <integer-value> stcli rebalance start</pre>	<p>Enables zones and creates a specific number of zones.</p> <p>Important You must execute the rebalance start command after you enable and create zones.</p>

Requirements and Limitations for Logical Availability Zones

Logical Availability Zones (LAZ) are subject to the following requirements and limitations:

- Only HyperFlex clusters with 8 nodes or more can be configured with Logical Availability Zones during the installation process.
- Logical Availability Zones can be enabled during the HyperFlex cluster installation, or it can be enabled using the command line interface at a later time. It is recommended to enable this feature during installation, in order to avoid a large migration and reorganization of data across the cluster, which would be necessary to comply with the data distribution rules if LAZ is turned on in a cluster already containing data.
- The number of zones can be manually specified as 3, 4, 5, or you can allow the installer to automatically choose, which is the recommended setting. When you let the installer decide the number of zones, the number of zones is decided based on the number of nodes in the cluster.

Here is a mapping table of number of nodes and resultant number of zones, with the assumption of RF as 3:

No. of Hyperconverged Nodes in the cluster	No. of Zones in the cluster
8	4
9	3
10	5
11	3
12	3
13	3
14	7
15	3
16	4
17	3
18	3
19	3
20	4

No. of Hyperconverged Nodes in the cluster	No. of Zones in the cluster
21	3
22	11
23	3
24	3
25	5
26	13
27	3
28	4
29	3
30	3
31	3
32	4

You can view the LAZ status and check which node is in which zone, by accessing **HyperFlex Connect** > **System information** and **HyperFlex Connect** > **Dashboard** respectively.

Also, you can view the LAZ details through CLI by running the following command:

```
stcli cluster get-zone
```

```
zones:
-----
pNodes:
-----
state: ready
name: 10.10.18.61
-----
state: ready
name: 10.10.18.59
-----
zoneId: 0000000057eebaab:0000000000000003
numNodes: 2
-----
pNodes:
-----
state: ready
name: 10.10.18.64
-----
state: ready
name: 10.10.18.65
-----
zoneId: 0000000057eebaab:0000000000000001
numNodes: 2
-----
pNodes:
-----
state: ready
name: 10.10.18.60
```

```

-----
state: ready
name: 10.10.18.63
-----
zoneId: 0000000057eebaab:00000000000000004
numNodes: 2
-----
pNodes:
-----
state: ready
name: 10.10.18.58
-----
state: ready
name: 10.10.18.62
-----
zoneId: 0000000057eebaab:00000000000000002
numNodes: 2
-----
isClusterZoneCompliant: True
zoneType: logical
isZoneEnabled: True
numZones: 4
AboutCluster Time : 08/22/2019 2:31:39 PM PDT

```

- The HyperFlex cluster determines which nodes participate in each zone, and this configuration cannot be modified.
- To maintain the most balanced consumption of space and data distribution, it is recommended that the number of nodes in a cluster are whole multiples of 3, 4, 5, or 7. For example, 8 nodes would evenly divide into 4 zones of 2 servers each, and 9 nodes would divide evenly into 3 zones of 3 servers each. Eleven nodes would create an unbalanced number of nodes across the zones, leading to unbalanced space consumption on the nodes.
- Expansion of a cluster should be done in multiples of the number of zones, when the cluster is operating with LAZ enabled. Expanding in such a way preserves a matched number of nodes in each zone, and prevents any unbalance of space consumption. For example, a cluster with 3 zones should be expanded by adding 3 more nodes, because adding only 1 or 2 nodes would lead to an imbalance, as would adding 4 nodes.

Expanding Cluster with Nodes Lesser than Zones

There are several requirements and guidelines to configuring Logical Availability Zones (LAZ) in a cluster. See [Requirements and Limitations for Logical Availability Zones, on page 157](#). If you choose to ignore these requirements, you can expand the cluster with the following procedure.

Step 1 Disable LAZ.

```
stcli cluster --set-zone 0
```

You must wait for a few minutes, roughly around 5 minutes for LAZ to be disabled.

Step 2 Add nodes to the cluster.

Step 3 After the nodes are added to the cluster, enable LAZ.

```
stcli cluster set-zone --zone 1
```

If the number of zones are not configured, 4 zones are created by default.

Step 4 Initiate rebalancing in the cluster.

```
stcli rebalance start
```

Note

- Rebalancing of nodes may decrease the performance of the cluster by about 30% to 40%.
- The time taken to rebalance the nodes in the cluster depends on the cluster size and the data within the cluster. You can use the **stcli rebalance status** command to review the progress of the rebalancing activity.

Step 5 Review the cluster configuration and determine if the cluster is zone compliant.

```
stcli cluster --get-zone
```



CHAPTER 12

Managing HX Controller VMs

- [Managing Storage Controller VMs, on page 161](#)
- [Powering On or Off Storage Controller VMs, on page 161](#)
- [Disabling HA VM Monitoring in HX Controller VMs, on page 162](#)

Managing Storage Controller VMs

Storage controller VMs provide critical functionality for the Cisco HX Distributed Data Platform. A storage controller VM is installed on every converged node in the storage cluster. The storage controller VMs provide a command line interface for running `stcli` commands on the storage cluster.



Note For HX220C-M4 server models, the storage controller VM is located on your SD cards and the datastore is a fixed size of 3.5Gb. The datastore is identified as Springpath<SN> and cannot be managed. If you notice an alert in vCenter regarding usage of the Springpath<SN> datastore, you can safely ignore it.

Powering On or Off Storage Controller VMs

You can power on or off VMs through the vSphere Web Client or through the ESX command line. This also applies to storage controller VMs, though generally the storage cluster operations handle powering on or off the storage controller VMs.

Step 1 Using the **vSphere Web Client** to power on or off a VM.

- a) Login to the vSphere Web Client.
- b) Locate the VM.

From the Navigator select, **Global Inventory Lists > Virtual Machines > vm**.

Storage controller VMs, have the prefix, `stCtlVM`.

- c) From the right-click or Actions menu select, **Power > Power On** or **Power > Power Off**.

Step 2 Using the **ESX command line** to power on or off a VM.

- a) Login to the command line for the ESX host for a VM.

- b) Locate the VM `vmid`.

This is specific to the ESX host. Run the command.

```
# vim-cmd vmsvc/getallvms
```

Sample response

```
Vmid  Name      File      Guest OS   Version  Annotation
1    stCtlVM-<vm_number> [SpringpathDS-<vm_number>] stCtlVM-<vm_number>/stCtlVM-<vm_number>.vmx
    ubuntu64Guest    vmx-11
3    Cisco HyperFlex Installer [test]  Cisco HyperFlex Installer/Cisco HyperFlex Installer.vmx
    ubuntu64Guest    vmx-09
Retrieved runtime info
Powered off
```

Storage controller VMs, have the prefix, `stCtlVM`.

- c) To power on a VM. Run the command specifying the VM to power on.

```
# vim-cmd vmsvc/power.on 1
```

- d) To power off a VM. Run the command specifying the VM to power off.

```
# vim-cmd vmsvc/power.off 1
```

Disabling HA VM Monitoring in HX Controller VMs

To avoid All Paths Down (APD) state in an HX cluster, use the vSphere Web Client to disable HA VM Monitoring for all the HX Controller VMs.

- Step 1** Login to the vSphere Web Client.
- Step 2** Select the HX cluster that you want to modify.
- Step 3** Select **Configure > VM Overrides** from the menu.
- Step 4** Click **Add**.

Add VM Override Sandbox window is displayed along with the list of VMs in vCenter.

- Step 5** Select all the available HX Controller VMs in the window.

Note The HX Controller VM names begin with `stCtlVM-`.

- Step 6** Click **Next**.
- Add VM Override** dialog box is displayed.

- Step 7** Locate the **vSphere HA - VM Monitoring** option and select the following:
- **Override** checkbox
 - **Disabled** from the drop-down list

- Step 8** Click **Finish** to apply the configuration changes.

HA VM Monitoring is disabled for all the HX controller VMs.



CHAPTER 13

Managing Ready Clones

- [HX Data Platform Ready Clones Overview, on page 165](#)
- [Benefits of HX Data Platform Ready Clones, on page 166](#)
- [Supported Base VMs, on page 166](#)
- [Ready Clone Requirements, on page 166](#)
- [Ready Clone Best Practices, on page 167](#)
- [Creating Ready Clones Using HX Connect, on page 167](#)
- [Creating Ready Clones Using the HX Data Platform Plug-In, on page 169](#)
- [Prepare to Customize HX Data Platform Ready Clones, on page 170](#)
- [Configuring Ready Clones Using Customized Specifications, on page 171](#)
- [Managing Virtual Machine Networking, on page 172](#)

HX Data Platform Ready Clones Overview

HX Data Platform Ready Clones is a pioneer storage technology that enables you to rapidly create and customize multiple cloned VMs from a host VM. It enables you to create multiple copies of VMs that can then be used as standalone VMs.

A Ready Clone, similar to a standard clone, is a copy of an existing VM. The existing VM is called the host VM. When the cloning operation is complete, the Ready Clone is a separate guest VM.

Changes made to a Ready Clone do not affect the host VM. A Ready Clone's MAC address and UUID are different from that of the host VM.

Installing a guest operating system and applications can be time consuming. With Ready Clone, you can make many copies of a VM from a single installation and configuration process.

Clones are useful when you deploy many identical VMs to a group.



Note

The Hyper-V Host issues an Offloaded Data Transfer (ODX) request to create a full file copy. HyperFlex uses cloning technology to create an instantaneous copy as a response to the full copy request.

Benefits of HX Data Platform Ready Clones

HX Data Platform Ready Clones provide the following benefits:

- **Create multiple clones of a VM at a time** - Simply right-click a VM and create multiple clones of the VM using the Ready Clones feature.
- **Rapid cloning** - HX Data Platform Ready Clones are extremely fast and more efficient than legacy cloning operations because they support VMware vSphere® Storage APIs – Array Integration (VAAI) data offloads and supported for powered on VMs. VAAI also called hardware acceleration or hardware offload APIs, are a set of APIs to enable communication between VMware vSphere ESXi hosts and storage devices. Use HX Data Platform Ready Clones to clone VMs in seconds instead of minutes.
- **Batch customization of guest VMs** - Use the HX Data Platform Customization Specification to instantly configure parameters such as IP address, host name, VM name for multiple guest VMs cloned from a host VM.
- **Automation of several steps to a one-click process** - The HX Data Platform Ready Clones feature automates the task to create each guest VM.
- **VDI deployment support** - Ready Clones are supported for desktop VMs on VDI deployments which are using VMware native technology.
- **Datastore access** - Ready Clone work on partially mounted/accessible datastores as long as the VM being cloned is on an accessible mountpoint.

Supported Base VMs

HX Data Platform supports:

- Base VMs stored on a HX Data Platform datastore
- Base VMs with HX Data Platform Snapshot. For Powered-on VMs, the Ready Clone workflow takes an HX Snapshot, and then uses the snapshot to create a clone after the clone is created. The same workflow happens when an HX Snapshot is removed.
- Maximum 2048 Ready Clones from one base VM
- Maximum 256 Ready Clones created in one batch at a time.
-

HX Data Platform does not support:

- Powered on base VMs with > 30 snapshots
- Powered on base VMs with Redo log snapshots

Ready Clone Requirements

- VMs must be within the HX Data Platform storage cluster. Non-HX Data Platform VMs are not supported.

- VMs must reside on a HX Data Platform datastore, VM folder, and resource pool.

Ready Clones fail for any VM that is not on a HX Data Platform datastore. This applies to Ready Clones on a VM level, VM folder level, or resource pool level.

- VMs can have only native snapshots. Ready Clones cannot be created from VMs with snapshots that have redo logs, (non-native snapshots).
- Use only the single vNIC customization template for Ready Clones.
- Beginning in Cisco HX Release 3.0, SSH does not need to be enabled in ESX on all the nodes in the storage cluster.

Ready Clone Best Practices

- Use the customization specification as a profile or a template.
- Ensure that properties that apply to the entire batch are in the customization specification.
- Obtain user-defined parameters from the HX Data Platform Ready Clone batch cloning work flow.
- Use patterns to derive per-clone identity settings such as the VM guest name.
- Ensure that the network administrator assigns static IP addresses for guest names and verify these addresses before cloning.
- You can create a batch of 1 through 256 at a given time. The HX Data Platform plug-in enables you to verify this.
- Do not create multiple batches of clones simultaneously on the same VM (when it is powered on or powered off) because it causes failures or displays incorrect information on the master task updates in the HX Data Platform plug-in.
- Only use the Hyper-V ReadyClone PowerShell script on a cluster node that is not in a paused state.

Hyper-V

Creating Ready Clones Using HX Connect

Use HX Data Platform Ready Clones feature to populate your cluster by creating multiple clones of a VM, each with different static IP addresses.

**Note**

If you click **Ready Clones** to clone a VM when the OVA deployment of that VM is in progress, you will get an error message. You can clone a VM only after the successful VM deployment.

Step 1 Login to HX Connect as an administrator.

Step 2 From **Virtual Machines** page, select a *virtual machine*, then click **Ready Clones**.

Step 3 Complete the **Ready Clone** dialog fields.

UI Element	Essential Information
Number of clones	Enter the number of Ready Clones that you want to create. You can create a batch of 1 through 256 clones at a given time.
Customization Specification	Optional field. Click the drop-down list and select a Customization Specification for the clone from the drop-down list (which includes the customization specifications available in vCenter). The system filters the customization specifications for the selected host virtual machine. For example, if the selected host virtual machine uses Windows OS for guest virtual machines, the drop-down list displays Windows OS customization specifications.
Resource Pool	Optional field. If you have resource pools defined in your HX Storage Cluster, you can select one to store the Ready Clones of the selected virtual machine.
VM Name Prefix	Enter a prefix for the guest virtual machine name. This prefix is added to the name of each Ready Clone created.
Starting clone number	Enter a clone number for the starting clone. Each Ready Clone must have a unique name, numbering is used to ensure a unique element in the name.
Increment clone numbers by	Enter a value using which the clone number in the guest virtual machine name must be increased, or leave the default value 1 as is. The system appends a number to the names of the virtual machine Ready Clones (such as clone1, clone2, and clone3). By default, the number starts from 1. You can change this value to any number.
Use same name for Guest Name	Select this check box to use the vCenter VM inventory name as the guest host virtual machine name. If you uncheck this box, a text box is enabled. Enter the name you want to use for the guest host virtual machine name.
Preview	After required fields are completed, HX Data Platform lists the proposed Ready Clones names. As you change the content in the required fields, the Clone Name and Guest Name fields update.
Power on VMs after cloning	Select this check box to turn the guest virtual machines on after the cloning process completes.

Step 4 Click **Clone**.

HX Data Platform creates the number of Ready Clones with the naming and location specified.

Creating Ready Clones Using the HX Data Platform Plug-In

If you use the VMware cloning operation, you can create only a single clone from a VM. This operation is manual and slower than batch processing multiple clones from a VM. For example, to create 20 clones of a VM, you must manually perform the clone operation over and over again.



Note Use HX Data Platform Ready Clones to create multiple clones of a VM in one click!

For example, you can create ten different clones with different static IP addresses from a Windows VM.

Step 1 From the vSphere Web Client Navigator, select **Global Inventory Lists > Virtual Machines**. This displays the list of VMs in vCenter.

Step 2 Select the VM to clone, and open the **Actions** menu. Either right-click the VM or click the **Actions** menu in the VM information portlet.

If needed, view the list of clusters and associated VMs to verify the VM is a storage cluster VM.

Step 3 Select **Cisco HX Data Platform > Ready Clones** to display the Ready Clones dialog box.

Step 4 Specify the following information in the Ready Clones dialog box:

Control	Description
Number of clones	Type the number of clones that you want to create. You can create a batch of 1 through 256 clones at a given time.
Customization Specification	Click the drop-down list and select a Customization Specification for the clone from the drop-down list (which includes the customization specifications available in vCenter). The system filters the customization specifications for the selected host VM. For example, if the selected host VM uses Windows OS for guest VMs, the drop-down list displays Windows OS customization specifications.
VM name prefix	Type a prefix for the guest VM name.
Starting clone number	Type a clone number for the starting clone.
Use same name for 'Guest Name'	Select this check box to use the vCenter VM inventory name as the guest host VM name. If you uncheck this box, a text box is displayed. Enter the name you want to use for the guest host VM name. The system displays the guest VM names in the Guest Name column in the dialog box. There is a similar option in the Customization Specification itself. This HX Data Platform Ready Clone batch customization process overrides the option that you specify in the Customization Specification option. <ul style="list-style-type: none"> If the Customization Specification contains a NIC or network adapter that uses a static gateway and static subnet and the guest name resolves to a static IP address, then the system assigns the network adapter the static IP address associated with the guest name. It also sets the storage cluster name or host name to the guest name specified.

Control	Description
	<ul style="list-style-type: none"> If the Customization Specification contains a NIC or network adapter that obtains the IP address using DHCP, then the system sets only the storage cluster name or host name to the guest name specified.
Increment clone number by	Type a value using which the clone number in the guest VM name must be increased, or leave the default value 1 as is. The system appends a number to the names of the VM clones (such as clone1, clone2, and clone3). By default, the number starts from 1. You can change this value to any number.
Power on VMs after cloning	Select this check box to turn the guest VMs on after the cloning process completes.

Step 5 Click **OK** to apply your configuration changes.

The vSphere Web Client Recent Tasks tab displays the status of the Ready Clones task. The system displays:

- Top-level progress with the initiator as the logged in vCenter user.
- Implementation work flows with the initiator as the logged in vCenter user and a HX Data Platform extension.
- As part of the Ready Clone workflow a temporary snapshot is listed in vCenter and HX Connect. This is listed as an extra powered off VM transiently, only while the Ready Clones are being created.

Prepare to Customize HX Data Platform Ready Clones

- Create a customization specification per the VMware documentation.
Apply the customization settings described in the following topics specific to either Linux or Windows VMs.
- Obtain the IP addresses from the administrator. For example, ten IP addresses 10.64.1.0 through 10.64.1.9.
- Gather information specific to your network such as the subnet mask for these IP addresses.
- Ensure that the base VM is valid (not disconnected, undergoing snapshots, or vMotion).
- Ensure that Guest Tools is installed on the base VM. Update it if necessary.
- Go to the VM Summary tab and verify that Guest Tools is working.

Creating a Customization Specification for Linux in the vSphere Web Client

Use the vSphere Web Client Guest Customization wizard to save guest operating system settings in a specification that you can apply when cloning virtual machines or deploying from templates.

Complete the wizard with the following considerations.

- You can use the HX Data Platform Ready Clones feature to overwrite the guest name that you specify in when you create the customization specification.

- HX Data Platform Ready Clones enable you to use patterns in the VM name or guest name.
- HX Data Platform supports only one NIC.
- Editing the NIC of a Customized Linux VM
 - You can use a fake IP address because the HX Data Platform Ready Clone customization process overwrites this address.
 - HX Data Platform Ready Clones resolve VM guest names to static IP addresses and sets them for the cloned VMs.

The customization specification you created is listed in the Customization Specification Manager. You can use it to customize virtual machine guest operating systems.

Create a Customization Specification for Windows in the vSphere Web Client

Use the vSphere Web Client Guest Customization wizard to save Windows guest operating system settings in a specification that you can apply when cloning virtual machines or deploying from templates.

**Note**

The default administrator password is not preserved for Windows Server 2008 after customization. During customization, the Windows Sysprep utility deletes and recreates the administrator account on Windows Server 2008. You must reset the administrator password when the virtual machine boots the first time after customization.

Complete the wizard with the following considerations:

- The operating system uses this name to identify itself on the network. On Linux systems, it is called the host name.
- HX Data Platform supports only one NIC.
- Editing the NIC of a Customized Windows VM

You can use a fake IP address because the HX Data Platform Ready Clone customization process overwrites it.

The customization specification you created is listed in the Customization Specification Manager. You can use it to customize virtual machine guest operating systems.

Configuring Ready Clones Using Customized Specifications

Use a customized specification to ensure IP addresses are applied correctly to the new VMs if you use static IP addresses.

For example, if you create a Windows server VM clone and you use DHCP, the guest VMs are automatically assigned new IP addresses. But, if you use static IP addresses, the IP address is not automatically replicated in the guest VM. To resolve this, configure HX Data Platform Ready Clones using a Customization Specification.

-
- Step 1** Obtain the valid DNS names and ensure that they resolve to valid IP addresses.
- For example, to provision a batch of 100 Windows VMs where the guest name is userwinvm1 to userwinvm100, check that userwinvm1 through userwinvm100 are valid IP addresses.
- Step 2** Install Guest VM tools on the source VM.
- Step 3** Clone the source VM using the Ready Clones feature. The cloned guest VMs obtain the identity of the source VM.
- Step 4** Use the Customization Specification to change the identity of all cloned VMs. You can configure parameters such as IP address, host name, and VM name.
-

Managing Virtual Machine Networking

After you have made changes to your storage cluster, you can ensure that the networking for the virtual machines on the nodes in the clusters is configured correctly. See the UCS Manager documentation for complete virtual machine networking information.

-
- Step 1** Verify the VLANs are configured correctly.
- See the VLANs chapter in the *Cisco UCS Manager Network Management Guide* at, [Cisco UCS Manager Network Guide](#).
- Step 2** Verify the vNICs are configured correctly.
- See the Configuring vNIC Templates topics in the *Cisco UCS Manager Network Management Guide* at, [Cisco UCS Manager Network Guide](#).
- Step 3** Verify the Virtual Port Groups are configured correctly.
- See the Add a Virtual Machine Port Group topic in the *VMware vSphere 6.0 Documentation* at, <http://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.networking.doc%2FGUID-004E2D69-1EE8-453E-A287-E9597A80C7DD.html>
-



CHAPTER 14

Managing HX Native Snapshots

- [HX Native Snapshots Overview, on page 173](#)
- [Benefits of HX Native Snapshots, on page 174](#)
- [HX Native Snapshot Considerations, on page 175](#)
- [HX Native Snapshots Best Practices, on page 177](#)
- [Understanding SENTINEL Snapshots, on page 178](#)
- [HX Native Snapshot Time Zones, on page 178](#)
- [Creating HX Native Snapshots, on page 179](#)
- [Scheduling HX Native Snapshots Overview, on page 180](#)
- [Scheduling HX Native Snapshots, on page 181](#)
- [Setting the Frequency of HX Native Scheduled Snapshots, on page 182](#)
- [Deleting HX Native Snapshot Schedules, on page 182](#)
- [Reverting to an HX Native Snapshot, on page 183](#)
- [Deleting HX Native Snapshots, on page 183](#)

HX Native Snapshots Overview

HX native snapshots are a backup feature that saves versions (states) of VMs. VMs can be reverted back to a prior saved version using an HX native snapshot. A native snapshot is a reproduction of a VM that includes the state of the data on all VM disks and the VM powerstate (on, off, or suspended) at the time the native snapshot is taken. Taking a native snapshot to save the current state of a VM gives you the ability to revert back to the saved state.

The following methodologies are used in the administration of HX native Snapshots:

- The HX vSphere web client plug-in for FLEX mode facilitates the creation of on-demand HX native snapshots, as well as the creation of HX native snapshot schedules. The HX vSphere client plug-in for HTML 5 mode does not provide these capabilities.
- The vSphere “Manage Snapshots” function can revert to a specific HX native snapshot, or delete all snapshots.
- Cisco HyperFlex Connect can create on-demand HX native snapshots.
- The HyperFlex command line user interface can create HX native snapshots.
- HX REST APIs can create and delete HX native snapshots.

For additional information about VMware snapshots, see the VMware KB, Understanding virtual machine snapshots in VMware ESXi and ESX (1015180) at, http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1015180

Benefits of HX Native Snapshots

HX native Snapshots provide the following benefits:

- **Revert Registered VMs** - If a VM is registered, whether powered-on or powered-off, HX native snapshots and VM snapshots, can be used to revert to an earlier point in time at which the snapshot was created.
- **High Performance** - The HX native snapshot process is fast because it does not incur I/O overhead.
- **VM Performance** - HX native snapshots do not degrade VM performance.
- **Crash-Consistent** - HX native snapshots are crash-consistent by default. I/O crash consistency is defined as maintaining the correct order of write operations to enable an application to restart properly from a crash.
- **Quiescence** - HX native snapshots can be created with the guest file system quiesced. The quiesce option is available when using Cisco HyperFlex Connect, the HyperFlex command line user interface, and HX REST APIs. VMware tools should be installed in the guest VM when creating HX native snapshots using the quiesce option.

Quiescing a file system is a process of bringing the on-disk data of a physical or virtual computer into a state suitable for backups. This process might include operations such as flushing dirty buffers from the operating system's in-memory cache to disk, or other higher-level application-specific tasks.

If your system displays quiesce errors, review the following VMware KB article [Troubleshooting Volume Shadow Copy \(VSS\) quiesce related issues \(1007696\)](#).

- **Scheduled Snapshots are Tolerant to Node Failures** - Scheduled snapshots are tolerant to administrative operations that require a node shutdown, such as HX maintenance mode and HX online upgrade.
Scheduled Snapshots are tolerant to failures in other HX clusters in multi cluster environments.
- **Individual or Grouped** - HX native snapshots can be created on a VM level, VM folder level, or resource pool level.
- **Granular Progress and Error Reporting** - Monitoring is performed at the task level for Resource Pool, Folder, and VM level HX native snapshots
- **Instantaneous Snapshot Delete** - Deletion of an HX native snapshot and consolidation is always instantaneous.
- **Parallel Batch Snapshots** - Support for up to 255 VMs in a Resource Pool or Folder for parallel batched snapshots.
- **VDI Deployment Support** - Scheduled HX native snapshots are supported for desktop VMs on VDI deployments which are using VMware native technology.
- **Datastore Access** - Snapshots work on partially mounted or accessible datastores as long as the VM being snapshot is on an accessible mountpoint.

HX Native Snapshot Considerations

Snapshot Parameters

- **HX native snapshots** - When creating the first HX native snapshot an HX SENTINEL snapshot is created prior to the HX native snapshot. The SENTINEL snapshot is a baseline snapshot which ensures that subsequent snapshots are HX native snapshots. For additional information about SENTINEL snapshots, see [Understanding SENTINEL Snapshots, on page 178](#). When a SENTINEL snapshot is present, creating additional snapshots using vSphere results in the creation of HX native snapshots.



Note There should be no VMware snapshots (non-native) present when creating native snapshots.

- **Maximum Number of Stored Snapshots** - VMware has a limitation of 31 snapshots per VM. The limitation total is equal to the sum of all VMware created snapshots, the HX SENTINEL snapshot, and HX native snapshots.

Beginning with HX Release 4.0 Snapshot operations beyond the number set in the `snapshot.maxSnapshots` property of the VM fail, with the following error message: `Snapshot operation cannot be performed.`

- **Scheduled Snapshots** - Do not have overlapping snapshots scheduled on VMs and their resource pools.

Snapshots During Upgrade Processes

- HX native snapshots are not supported while an upgrade of HX Data Platform, ESXi, or UCS is in progress.

VMs

- **HX Release 3.5 deployments** - Ensure that the VM is not attached to Graphic Processing Unit (GPU).
- **HX Release 3.5.2(a) and earlier**, all powered on VMs use synchronous consolidation (`asynConsolidate = false`) when taking HX native snapshots.
- Beginning with HX Release 3.5.2(b), all powered on VMs use asynchronous consolidation (`asynConsolidate = true`) when taking HX native snapshots. If the VM is powered off, the settings remain unchanged.
- **VM Hardware Version** - HX native Snapshots requires the use of VM hardware Version 9 or later. Using the most recent version is recommended.
- **Deleted VMs** - The life cycle of HX native snapshots, similar to VM snapshots, is tied to the virtual machine. If the VM is deleted (accidentally or intentionally), all associated HX native snapshots are deleted. HX native snapshots do not provide a mechanism to recover from a deleted VM. Use a backup solution to protect against VM deletion.
- **HX Data Platform storage controller VMs** - You cannot schedule HX native snapshots for storage controller VMs.

- **Non-HX Data Platform VMs** - HX native snapshots will fail for any VM that is not on a HX datastore. This applies to snapshots on a VM level, VM folder level, or resource pool level. To create an HX native snapshot, the VM must reside on a single HX datastore.
- **Suspended VMs** - Creating the first HX native snapshot and the HX SENTINEL snapshot on VMs in a suspended state is not supported.
- **VM Name** - The VM name must be unique per vCenter for taking an HX native snapshot.
- **Ready storage cluster** - To enable the creation of an HX native snapshot, the storage cluster must be healthy, have sufficient space, and be online. The datastore in which the VM resides must be accessible. The VM must be valid and not in a transient state, such as vMotioning.

Cluster with only 1 on-line Node Remaining

- HX native Snapshot is not supported on single (1-node) on-line node on a CBT Enabled VM in a powered-on state. Power-off the VM and take the SENTINEL snapshot, subsequent snapshots on powered-on VMs are supported.

vCenter

- **vMotion** - vMotion is supported on VMs with HX native snapshots.
- **Storage vMotion** - Storage vMotion is not supported on VMs with HX native snapshots. If the VM needs to be moved to a different datastore, delete the snapshots before running storage vMotion.

Naming

- **Duplicate names** - VMs or Resource Pools with duplicate names within the HX Data Platform vCenter are not supported and cause HX native snapshots to fail. This includes parents and children within nested resource pools and resource pools within different vCenter clusters.
- **Characters in names** - Special characters are not supported. Using special characters in names results in names appearing different than specified.

Disks and Datastores

- **VM Datastores** - Ensure that all the VMDKs belonging to a VM are on the same datastore prior to creating HX native snapshots.
 - HX native snapshots are not supported with multiple datastores.
- **Thick Disks** - If the source disk is thick, then the HX native snapshot of the VM will also be thick. Increase the datastore size to accommodate the snapshot if required.
- **Virtual Disk Types** - VMware supports a variety of virtual disk backing types. The most common is the FlatVer2 format. HX native snapshots are supported for this format.

There are other virtual disk formats such as Raw Device Mapping (RDM), SeSparse, VmfsSparse (Redlog format). VMs containing virtual disks of these formats are not supported with HX native snapshots.

Login Access

- **SSH** - Ensure that SSH is enabled in ESX on all the nodes in the storage cluster.

Limitations

Object	Maximum Number
HX native snapshots	30 per VM VMware limit is 31. One snapshot is consumed by the HX SENTINEL snapshot.
Datastores	48 per storage cluster
Maximum VMDK size	3 TB

HX Native Snapshots Best Practices



Important

Always use the HX native snapshots feature to create the first snapshot of a VM. This ensures that all subsequent snapshots are in native format.

- Do not use the VMware snapshot feature to create the first snapshot of a VM. VMware snapshots use redo log technology that result in degraded performance of the original VM. This performance degrades further with each additional snapshot.
 - If there are redo log snapshots that should be deleted, edit the `/etc/vmware/config` file and `setsnapshot.asyncConsolidate="TRUE"` on the ESXi host(s) where the redo log snapshots reside.

- HX native snapshots do not impact VM performance after the initial HX SENTINEL and HX native snapshot are created.

- Add all the VMDKs to the VM prior to creating the first HX native snapshot.

When VMDKs are added to a VM, additional HX SENTINEL snapshots are taken. Each additional HX SENTINEL consumes additional space.

For example, when adding 2 new VMDKs to an existing VM that has HX native snapshots, at the next scheduled HX native snapshot, 1 new HX SENTINEL is created. If it is necessary to add one or more additional VMDKs, review any existing HX native snapshot schedule retention plans and make sure that the total number of retained HX native snapshots plus any HX SENTINEL snapshots will not attempt to exceed a total value of 31.

- When creating large numbers of HX native snapshots consider the following:
 - Schedule HX native snapshots at a time the expected data traffic is low.
 - Use multiple resource pools or VM folders to group VMs rather than a single resource pool or VM folder. Then stagger the HX native snapshot schedules by group.

For example resourcePool1, schedule snapshots at :00, resourcePool2, schedule snapshots at :15, resourcePool3, schedule snapshots at :30.

- If vCenter running on a VM in the storage cluster, do not take an HX native snapshot of the vCenter VM. This is related to VMware KB, VMware VirtualCenter Server service fails due to a quiesced snapshot operation on the vCenter Server database virtual machine (2003674), at https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2003674

Understanding SENTINEL Snapshots

When creating the first HX native snapshot of a VM, a baseline snapshot called an HX SENTINEL snapshot is created as well as the requested HX native snapshot. The HX SENTINEL snapshot ensures that all subsequent of the VM snapshots are HX native snapshots

HX SENTINEL snapshots prevent reverted VMs from having VMware redo log-based virtual disks. Redo log-based virtual disks occur when an original snapshot is deleted and the VM is reverted to the second oldest snapshot.

An HX SENTINEL snapshot consumes one snapshot of the 31 total available per VMware limitation.

Using HX SENTINEL snapshots

- Do not delete the HX SENTINEL snapshot unless it is a required operation.
- Do not revert a VM to the HX SENTINEL snapshot.
- If the HX SENTINEL snapshot needs to be deleted for any reason, all snapshots should be deleted.

HX Native Snapshot Time Zones

There are three objects that display and affect the timestamp and schedule of snapshots:

- vSphere and vCenter use UTC time.
- vSphere Web client uses the browser timezone.
- The HX vSphere web client plug-in, HX storage cluster, and HX storage controller VM use the same time zone. This is enforced across the HX storage cluster. The time zone used by these entities is configurable. The default is UTC.

The HX storage controller VM time is used to set the schedule. The vSphere UTC time is used to create the HX native snapshots. The logs and timestamps vary depending upon the method used to view them.

When a schedule is created using the HX vSphere web client plug-in, the scheduled times are converted to UTC from the HX storage controller VM time zone. When you view the schedule through the vSphere web client Scheduled Tasks it displays the tasks in browser time zone.

When converted to the same timezone, they translate to the same time. For example: 5:30pm PST, 8:30PM EST, 1:30AM UTC are all the same time.

To have vSphere Scheduled Tasks tab display the same time for a scheduled snapshot that you create in the HX vSphere web client plug-in, set the storage controller VM to UTC.

To have scheduled snapshots run based on local time zone settings, set that time zone for the storage cluster. By default, the storage controller VM uses the UTC time zone set during HX Data Platform installation.

If the vSphere and the storage controller VM are not using the same time zone, the vSphere scheduled tasks tab might display a different time than the schedule in the HX vSphere web client plug-in schedule snapshot dialog.

When you configure an hourly snapshot, the snapshot schedule runs between a specific start time and end time. The vSphere Task window might display a status that a scheduled snapshot was completed outside the hourly end time based on the timezone

Identify and set the time zone used by the storage controller VM

1. From the storage controller VM command line, view the configured time zone.

```
$ stcli services timezone show
```

2. Change the storage cluster time zone.

```
$ stcli services timezone set --timezone timezone_code
```

See a time zone reference for time zone codes, such as

https://en.wikipedia.org/wiki/List_of_tz_database_time_zones

Creating HX Native Snapshots

This procedure requires the use of the HX vSphere web client plug-in for FLEX mode. The HX vSphere client plug-in for HTML 5 mode does not provide these capabilities.

To create HX native snapshots, perform the following steps:

Before you begin

Remove any redo log snapshots for VMs in the HX storage cluster. If this step is not completed, VMs might be stunned during snapshot consolidation.

Redo log snapshots are snapshots that are created through the VMware snapshot feature and not through the HX native snapshot feature. To edit the ESXi host configuration where the redo log snapshots reside,

1. Login to the ESXi host command line.
2. Locate and open for editing the file, `/etc/vmware/config`
3. Set the `snapshot.asyncConsolidate` parameter to `TRUE`.

```
snapshot.asyncConsolidate="TRUE"
```

-
- Step 1** From the vSphere Web Client Navigator, select the VM level, VM folder level, or resource pool level. For example, **vCenter Inventory Lists > Virtual Machines** to display the list of VMs in vCenter.
- Step 2** Select a storage cluster VM and open the **Actions** menu. Either right-click the VM or click the Actions menu in the VM information portlet.
- Note** Ensure there are no non-HX Data Platform datastores on the storage cluster resource pool or the snapshot will fail.
- Step 3** Select **Cisco HX Data Platform > Snapshot Now** to display the Snapshot dialog box.
- Step 4** Type a name for the snapshot in the snapshot dialog box.

Step 5 Type a description of the snapshot.

Step 6 Click **OK** to create an HX native snapshot.

The Recent Tasks tab displays the status message:

```
Create virtual machine native snapshot.
The first snapshot
```

Scheduling HX Native Snapshots Overview

Apply HX native snapshot schedules to storage cluster objects, such as VMs or resource pools.



Note

If you re-register the vCenter cluster, your HX native snapshot schedules are lost. If this happens, reconfigure HX native snapshot schedules.

When scheduling an HX native snapshots consider your back up requirements. For critical data, retain more frequent HX native snapshots. If there is a disaster, it is possible to restore recent HX native snapshots or create a custom real-time HX native snapshot. For less critical data, considering creating less frequent HX native snapshots.

HX native snapshot scheduling helps control backup costs. For each VM in a storage cluster, you can schedule hourly, daily, or weekly snapshots. The maximum frequency for any specific VM is once per hour. Hourly settings are available in 15 minute increments.

For example, HX native snapshots are taken each day, given the following settings.

VM 1 hourly snapshots to run at hour:15 minutes, between 10 PM and 1 AM.

VM 2 hourly snapshots to run at hour:30 minutes, between 8 PM and 12 AM.

VM 3 and 4 hourly snapshots to run at hour:45, between 6 AM and 8 AM.

VM 5 daily snapshot to run at 6:00 AM

Each day these HX native snapshots are taken. Notice that the last HX native snapshot is before the ending hour:00.

6:00 AM — VM 5

6:45 AM — VM 3, VM 4

7:45 AM — VM 3, VM 4

8:30 PM — VM2

9:30 PM — VM2

10:15 PM — VM1

10:30 PM — VM2

11:15 PM — VM1

11:30 PM — VM2

12:15 AM — VM1

To schedule an HX native snapshot every hour over 24 hours:

Before you begin

The HX vSphere web client plug-in for FLEX mode facilitates the creation of HX native snapshot schedules. The HX vSphere client plug-in for HTML 5 mode does not provide these capabilities.

Step 1 Set the start time

Step 2 Set the end time one hour before the start time.

Example:

hour: 15, start 4 PM, end 3 PM.

This takes an HX native snapshot at 4:15 PM, 5:15 PM, ... 12:15 AM, 1:15 AM ... 2:15 PM, 3:15 PM. Then restarts the 24 hour cycle.

Note The maximum number of HX native snapshots per VM is 31. One HX SENTINEL snapshot is also required. So, it is possible to take an hourly HX native snapshot and retain the most recent 30 HX native snapshots.

The HX native schedule snapshot displays the set time for the snapshot based on the current time zone setting for the storage controller VM. So, if an HX native snapshot was set at 7 pm PST and the storage controller VM time zone is changed to EST. The next time you open the HX native scheduler window, it automatically updates to 10 pm EST.

Scheduling HX Native Snapshots

Before you begin

The HX vSphere web client plug-in for FLEX mode facilitates the creation of HX native snapshot schedules. The HX vSphere client plug-in for HTML 5 mode does not provide these capabilities.

Step 1 From the vSphere Web Client Navigator Home page, select the VM or resource pool list.

For example, **vCenter Inventory Lists > Virtual Machines** to display the list of VMs in vCenter.

Step 2 Select a storage cluster VM or resource pool and open the **Actions** menu.

Either right-click the object or click the Actions menu.

Step 3 From the Actions menu, select **Cisco HX Data Platform > Schedule Snapshot** to display the Schedule Snapshot dialog box.

Step 4 Select the snapshot frequency.

Click the boxes for hourly, daily, and/or weekly frequency and set the starting days, times, and duration.

Step 5 Set the number of snapshots to retain.

When the maximum number is reached, older snapshots are removed as newer snapshots are created.

Step 6 Unselect existing scheduled items, as needed.

If a previous schedule existed, unselecting items deletes those items from the future schedule.

Step 7 Click **OK** to accept the schedule and close the dialog.

Setting the Frequency of HX Native Scheduled Snapshots

Create a snapshot every hour at specific times, daily at a specific time, or weekly on selected days and times.

Before you begin

The HX vSphere web client plug-in for FLEX mode facilitates the creation of HX native snapshot schedules. The HX vSphere client plug-in for HTML 5 mode does not provide these capabilities.

Open the **Schedule Snapshot** dialog box for a VM or resource pool.

-
- Step 1** From the Schedule Snapshot dialog box, select the **Enable Hourly Snapshot**, **Enable Daily Snapshot**, or **Enable Weekly Snapshot** check box.
- Step 2** Click the **Start at** drop-down list to select a start time. Select hour, minutes in 15 minute increments, and AM or PM.
- Step 3** For an hourly snapshot schedule, click the **Until** drop-down list to select an end time. Select hour, minutes in 15 minute increments, and AM or PM. Set the minute to the same value as the Start at time.
- The HX Data Platform plug-in creates a snapshot of the VM every hour between the start and end times.
- Step 4** Select the corresponding check box to specify **Days** of the week on which you want to take the snapshots.
- Step 5** Under **Retention**, either type a number or use the arrow button to specify the maximum number of copies to retain for each schedule.
-

Deleting HX Native Snapshot Schedules

Before you begin

The HX vSphere web client plug-in for FLEX mode facilitates the administration of HX native snapshot schedules. The HX vSphere client plug-in for HTML 5 mode does not provide these capabilities.

-
- Step 1** From the HX vSphere Web Client Navigator Home page, select the VM or resource pool list.
- For example, **vCenter Inventory Lists > Virtual Machines** to display the list of VMs in vCenter.
- Step 2** Select a storage cluster VM or resource pool and open the **Actions** menu.
- Either right-click the object or click the Actions menu.
- Step 3** From the Actions menu, select **Cisco HX Data Platform > Schedule Snapshot** to display the Schedule HX Native Snapshot dialog box.
- Step 4** Uncheck the scheduled options that are no longer required.
- Step 5** Click **OK** to accept the changes, this includes deleting previously configured schedules, and exit the dialog.
- Step 6** Confirm the schedule is deleted.

Select a storage cluster VM or resource pool. Click the HX vCenter tabs, **Manage > Scheduled Tasks**. The previous HX native snapshot schedule should not be listed.

Reverting to an HX Native Snapshot

Reverting a snapshot is returning a VM to a state stored in a snapshot. Reverting to a snapshot is performed on one VM at a time. This is not performed at the resource pool or VM folder level. Reverting snapshots is performed through the vCenter Snapshot Manager and not through the HX Data Platform plug-in.

Before you begin

Snapshots of the VM must exist.

- Step 1** From the vSphere Web Client Navigator Home page, select the VM level, VM folder level, or resource pool level. For example, **vCenter Inventory Lists > Virtual Machines** to display the list of VMs in vCenter.
- Step 2** Select a storage cluster VM and open the **Actions** menu. Either right-click the VM or click the Actions menu in the VM information portlet.
- Step 3** Select **Snapshots > Manage Snapshots** to open the vSphere Snapshot Manager.
- Step 4** Select a snapshot to revert to from the hierarchy of snapshots for the selected VM.
- Step 5** Click **Revert to > Yes > Close**.

The reverted VM is included in the list of VMs and powered off. In selected cases, a VM reverted from a VM snapshot is already powered on. See the following table for more details.

Table 6: VM Power State After Restoring a HX VM Snapshot

VM State When HX VM Snapshot is Taken	VM State After Restoration
Powered on (includes memory)	Reverts to the HX VM snapshot, and the VM is powered on and running.
Powered on (does not include memory)	Reverts to the HX VM snapshot, and the VM is powered off.
Powered off (does not include memory)	Reverts to the HX VM snapshot, and the VM is powered off.

- Step 6** If the reverted VM is powered off, then select the VM and power it on.

Deleting HX Native Snapshots

Deleting HX native snapshots is managed through the vSphere interface and not through the HX vSphere plug-in.

-
- Step 1** From the vSphere Web Client Navigator, select **VMs and Templates** > *vcenter_server* > **Snapshots** > *datacenter* > *vm*.
- Step 2** Right-click the *vm* and select **Snapshots** > **Manage Snapshots**.
- Step 3** Select an HX native snapshot and click **Delete**.

Note Delete the HX SENTINEL snapshot by using **Delete All** option only. Do not delete the HX SENTINEL snapshot individually.



CHAPTER 15

Managing Clusters Running on Different HXDP Versions

- [Managing Clusters Running on Different HXDP Versions, on page 185](#)
- [Scenario—Site A at HXDP 3.0 and Site B at HXDP 2.6, on page 185](#)
- [Scenario—Site A at HXDP 2.6 and Site B at HXDP 3.0, on page 187](#)
- [Functionality Limitations, on page 188](#)

Managing Clusters Running on Different HXDP Versions

Scenario—Site A at HXDP 3.0 and Site B at HXDP 2.6

The following terms and abbreviations are used:

- **Site A**—Source cluster
- **Site B**—Target cluster
- **dr_cleanup tool**—Contact Cisco TAC to obtain this tool available in the 3.0 internal support package.

Prerequisites

- Before upgrading make sure, there are no VMs or groups in **Recovered** or **Halted** state.
- If the VMs are in **Halted** State, recover and unprotect the VMs or groups.
- If the VMs are in **Recovered** state, then unprotect the VMs or groups.

Step	Site A	Site B	Result
1.	At HXDP version 2.6 or lower.	At HXDP version 2.6 or lower.	—

Step	Site A	Site B	Result
2.	Upgrade to HXDP 3.0.	—	<ul style="list-style-type: none"> • Ongoing replication continues. • Planned migration for VMs is not supported. • See Functionality Limitations, on page 188 for more details.
3.	Before upgrading Site B, if a disaster happens on Site A.	<ol style="list-style-type: none"> 1. Execute the command: <i>stcli dp peer forget</i> 2. Recover VMs that are required. 3. Run the dr_cleanup tool to delete all the VM information from the disaster recovery database. 	Workloads are now running on Site B.
4.	Restore Site A. After Site A is restored, do the following: <ol style="list-style-type: none"> 1. Execute the command: <i>stcli dp peer forget</i> 2. Run the dr_cleanup tool to delete all the VM information from the disaster recovery database. 	—	Sites are unpaired.
5.	—	Upgrade to HXDP 3.0.	—
6.	Pair the sites.	—	Site A and Site B can now be re-paired and workloads can be protected.

Scenario—Site A at HXDP 2.6 and Site B at HXDP 3.0

The following terms and abbreviations are used:

- **Site A**—Source cluster
- **Site B**—Target cluster
- **dr_cleanup tool**—Contact Cisco TAC to obtain this tool available in the 3.0 internal support package.

Prerequisites

- Before upgrading make sure, there are no VMs or groups in **Recovered** or **Halted** state.
- If the VMs are in **Halted** State, recover and unprotect the VMs or groups.
- If the VMs are in **Recovered** state, then unprotect the VMs or groups.

Step	Site A	Site B	Result
1.	At HXDP version 2.6 or lower.	At HXDP version 2.6 or lower.	—
2.	—	Upgrade to HXDP 3.0.	<ul style="list-style-type: none"> • Ongoing replication continues. • Planned migration for VMs is not supported. • See Functionality Limitations, on page 188 for more details.
3.	Before upgrading Site A, if a disaster happens on Site A.	<ol style="list-style-type: none"> 1. Execute the command: <i>stcli dp peer forget</i> 2. Recover VMs that are required. 3. Run the dr_cleanup tool to delete all the VM information from the disaster recovery database. 	<ul style="list-style-type: none"> • Not all recovery options are available. • See Functionality Limitations, on page 188 for more details. • Workloads are now running on Site B.

Step	Site A	Site B	Result
4.	Restore Site A. After Site A is restored, do the following: <ol style="list-style-type: none"> 1. Execute the command: <i>stcli dp peer forget</i> 2. Run the dr_cleanup tool to delete all the VM information from the disaster recovery database. 	—	Sites are unpaired.
5.	Upgrade Site A to HXDP 3.0.	—	—
6.	—	Pair the sites.	Site A and Site B can now be re-paired and workloads can be protected.

Functionality Limitations

Newer functionality in release 3.0 is supported **ONLY** when both the source and target clusters are on the same HXDP version. It can take a while during upgrade for both the source and target to be on the same version. Review the following functionality limitations:

- Planned migration for VMs is not supported when peer sites have mismatched versions, such as when the target cluster is on 2.6, and source cluster is on 3.0.
- When the source is upgraded, all the newer features in release 3.0, such as movein and moveout of group VMs, migration are blocked on the source cluster until the peer is upgraded.
- If only the target cluster is upgraded, in **HX Connect** UI, **Network Mapping** options in the **Recovery** dialog box will not available until the source cluster is upgraded.



CHAPTER 16

Managing Virtual Machine Disaster Recovery

- [HX Data Protection Snapshot Overview, on page 189](#)
- [Protecting Virtual Machines Overview, on page 197](#)
- [Disaster Recovery Overview, on page 218](#)
- [Replication Maintenance Overview, on page 229](#)

HX Data Protection Snapshot Overview

The HyperFlex Data Protection Snapshot (DP Snapshot) feature enables the protection of virtual machines from disaster by configuring the replication of running VMs between a pair of network connected clusters. Protected virtual machines running on one cluster replicate to the other cluster in the pair, and vice versa. The two paired clusters typically are located at a distance from each other, with each cluster serving as the disaster recovery site for virtual machines running on the other cluster.

Once protection is configured on a VM, the HX Data Platform periodically takes a DP Snapshot of the running VM on the local cluster and replicates (copies) the DP snapshot to the paired remote cluster. In the event of a disaster at the local cluster, the most recently replicated snapshot of each protected VM can be recovered on the remote cluster. Each cluster that serves as a disaster recovery site for another cluster must be sized with adequate spare resources so that upon a disaster, it can run the newly recovered VMs in addition to its normal workload.



Note

Only the most recently replicated DP snapshot is retained on the destination cluster. Retaining additional DP snapshots is not supported.

Each VM is individually protected by assigning it protection attributes, including the replication interval (schedule). The shorter the replication interval, the fresher the replicated snapshot data is likely to be. DP snapshot intervals can range from once every 5 minutes to once every 24 hours.

A protection group is a group of VMs that have a common DP snapshot schedule and quiescence parameter value

Setting up DP snapshots requires two existing clusters running HX Data Platform version 2.5 or later. Both clusters must be on the same HX Data Platform version. Use HyperFlex Connect to complete the setup

First, each cluster is set up for DP snapshot networking. Use HX Connect to provide a set of IP addresses to be used by local cluster nodes to replicate to the remote cluster. HX Connect creates VLANs through UCS Manager, for dedicated DP snapshot network use.



Note When this option is chosen in HX Connect, UCSM is configured only when both UCS Manager and fabric interconnect are associated with the HyperFlex cluster. When UCSM and FI are not present, you must enter the VLAN ID, and not select UCSM configuration in HX Connect.

The two clusters, and their corresponding existing relevant datastores must be explicitly paired. The pairing setup can be completed using HX Connect from one of the two clusters. This requires administrative credentials of the other cluster.

Virtual machines can be protected (or have their existing protection attributes modified) by using HX Connect at the cluster where they are currently active.

HX Connect can monitor the status of both incoming and outgoing replication activity on a cluster.

After a disaster, a protected VM can be recovered and run on the cluster that serves as the DP snapshot recovery site for that VM.

Replication and Recovery Considerations

The following is a list of considerations when configuring virtual machine replication and performing disaster recovery of virtual machines.



Note Cisco HX Data Replication is not supported with HyperFlex clusters that have Hardware Acceleration Cards installed.

- **Administrator**—All replication and recovery tasks, except monitoring, can only be performed with administrator privileges on the local cluster. For tasks involving a remote cluster, both the local and remote user must have administrator privileges and should be configured with the vCenter SSO on their respective clusters.
- **Minimum and Recommended Bandwidth**—Beginning with HX Release 4.0(2a) the minimum bandwidth can be configured to be 10 Mb for smaller size deployments. The replication network link should also be reliable and have sustained minimum symmetric bandwidth same as configured in a HyperFlex DR network. This should not be shared with any other applications on an Uplink or Downlink.
- **Maximum Latency**—Maximum latency supported is 75ms between two clusters.

If you are scheduling to run multiple replication jobs at the same time, for example 32 as maximum supported by DR, and your bandwidth (50Mbps) is low and latency (75ms) high, it is possible that some jobs will error out until bandwidth becomes available. If this situation occurs, increase bandwidth or reduce the concurrency by staggering the replications.

During this situation, unprotect operations can take longer than expected.

- **Network Loss**—When there is a packet loss in data transmission across two sites, protection and recovery operations will have unexpected results. The transmission should be reliable for these features to work as expected.
- **Storage Space**—Ensure that there is sufficient space on the remote cluster to support the DR replication schedule. The protected virtual machines are replicated (copied) to the remote cluster at every scheduled

interval. Though storage capacity methods are applied (deduplication and compression), each replicated virtual machine will consume some storage space.

Not having sufficient storage space on the remote cluster can cause the remote cluster to reach capacity usage maximums. If you see **Out of Space** errors, see [Handling Out of Space Errors, on page 53](#). Pause all replication schedules until you have appropriately adjusted the space available on the cluster. Always ensure that your cluster capacity consumption is below the space utilization warning threshold.

• **Supported Clusters**—Replication is supported between the following HyperFlex clusters:

- 1:1 replication between HX clusters running under fabric interconnects.
- 1:1 replication between All Flash and Hybrid HX cluster running under fabric interconnects.
- 1:1 replication between 3-Node and 4-Node HX Edge and another 3-Node and 4-Node HX Edge cluster.
- 1:1 replication between All Flash 3-Node and 4-Node Edge and Hybrid 3-Node and 4-Node HX Edge clusters.
- 1:1 replication between 3-Node and 4-Node HX Edge and an HX cluster running under fabric interconnects.



Note 1:1 replication with 2-Node HX Edge is not supported.

- **Rebooting Nodes**—Do not reboot any nodes in the HX Cluster during any restore, replication, or recovery operation.
- **Thin Provision**—Protected VMs are recovered with thin provisioned disks irrespective of how disks were specified in the originally protected VM.
- **Protection Group Limitations**
 - Do not add VMs with ISOs or floppies to protection groups.

Protected Virtual Machine Scalability

- Beginning with HX Release 4.0(2a):
 - The sum of protected VMs on all nodes should not exceed the maximum limit of 1500 protected VMs per cluster in a single direction configuration and 750 protected VMs in a bi-direction configuration.
 - The maximum number of VMs allowed in a protection group is 64.
 - A maximum of 100 protection groups are supported.
- **Non-HX Datastores**—Periodical replication fails on a protected a VM with storage on a non-HX datastore. To avoid the failure, unprotect the VM or remove non-HX storage from the VM.

Do not move protected VMs from HX datastores to non-HX datastores. If a VM is moved to a non-HX datastore through storage vMotion, unprotect the VM.
- **Templates**—Templates are not supported with disaster recovery.

• Protection and Recovery of Virtual Machines with Snapshots

When replication is enabled:

- **VM with no Snapshots**—The entire content of the VM is replicated.
- **VM with VMware redolog snapshots**—The entire content including the snapshot data is replicated. When a VM with redolog snapshots is recovered, all previous snapshots are preserved.
- **VM with HX native snapshots**—Only the latest data is replicated, and the HX native snapshot data is not replicated. When the VM is recovered, previous snapshots are not preserved.
- **Data Protection snapshots** are stored on the same datastore as the protected VMs. Manual deletion of DP snapshots is not supported. Deleting snapshot directories will compromise HX data protection and disaster recovery.



Caution

To avoid deleting DP snapshots manually, it is important to remember that VMware does not restrict operations on datastores by the Admin user. In any VMware environment, datastores are accessed by the Admin via vCenter browser or by logging into the ESX host. Because of this, the snapshot directory and contents are browseable and accessible to Admins.

Other points for consideration include:

- **Location of the VMware vCenter**—If you delete a VM from VMware vCenter that is located on a “Other DRO” datastore pair, a recovery plan for this datastore pair fails during recovery. To avoid this failure, first unprotect the VM using the following command on one of the controller VM

stcli dp vm delete --vmid <VM_ID>
- **Name of the VM**—If you rename a VM from the vCenter, HyperFlex recovers at the previous name folder but registers the VM with the new name on the recovery side. Following are some of the limitations to this situation:
 - VMware allows a VMDK located at any location to be attached to a VM. In such cases, Hyperflex recovers the VM inside the VM folder and not at a location mapped to the original location. Also, recovery can fail if the VMDK is explicitly referenced in the `virtualmachine name.vmx` file by its path. The data is recovered accurately but there could be problems with registering the VM to vCenter. Correct this error by updating the `virtualmachine name.vmx` file name with the new path.
 - If a VM is renamed and a VMDK is added subsequently, the new VMDK is created at `[sourceDs] newVm/newVm.vmdk`. Hyperflex recovers this VMDK with the earlier name. In such cases, recovery can fail if the VMDK is explicitly referenced in the `virtualmachine name.vmx` file by its path. The data is recovered accurately but there could be problems with registering the VM to the Virtual Center. Correct this error by updating the `virtualmachine name.vmx` file name with the new path.

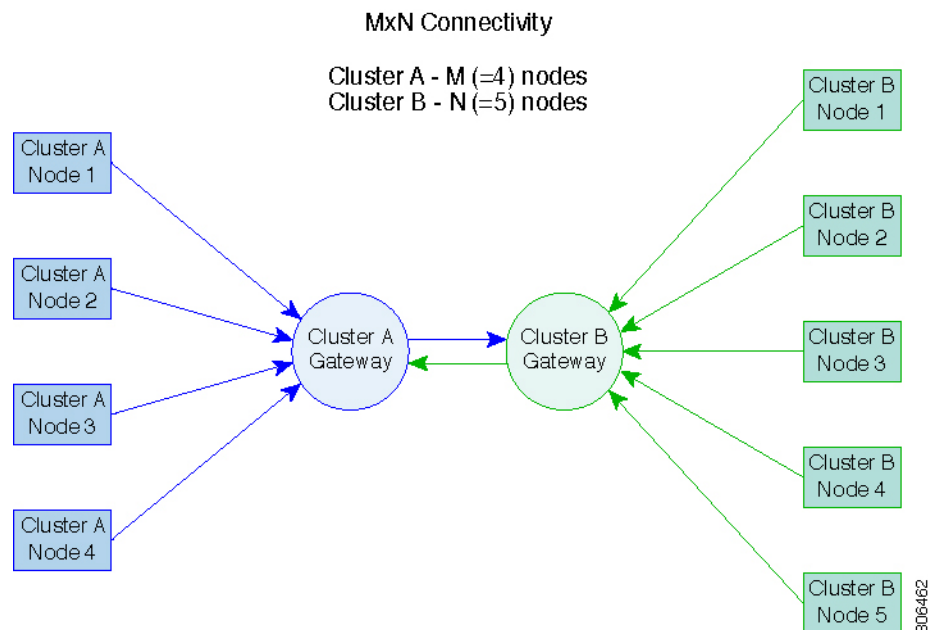
Replication Network and Pairing Considerations

A replication network must be established between clusters that are expected to use replication for DP snapshots. The Replication network is created to isolate inter-cluster replication traffic from other traffic within each cluster and site.

The following is a list of considerations when configuring replication network and pairing:

- To support efficient replication, all M nodes of cluster A have to communicate with all N nodes of cluster B, as illustrated in the *M x N connectivity between clusters* figure.
- To enable replication traffic between clusters to cross the site-boundary and traverse the internet, each node on Cluster A should be able to communicate with each node on Cluster B across the site boundary and the internet.
- The replication traffic must be isolated from other traffic within the cluster and the data center.
- To create this isolated replication network for inter-cluster traffic, complete these steps:
 - Create a replication network on each cluster.
 - Pair clusters to associate the clusters and establish M x N connectivity between the clusters.
- IP addresses, Subnet, VLAN, and Gateway are associated with the replication network of each cluster. Configure the corporate firewall and routers on both sites to allow communication between the clusters and the sites on TCP ports 9338,3049,9098,4049,4059.

M*N Connectivity Between Clusters



Storage Replication Adapter Overview

Storage Replication Adapter (SRA) for VMware vCenter Site Recovery Manager (SRM) is a storage vendor-specific plug-in for VMware vCenter server. The adapter enables communication between SRM and a storage controller at the Storage Virtual Machine (SVM) level as well as at the cluster level configuration. The adapter interacts with the SVM to discover replicated datastores.

For more information on installation and configuration of SRM, refer the following links as per the SRM release version:

- SRM 8.1 installation—<https://docs.vmware.com/en/Site-Recovery-Manager/8.1/srm-install-config-8-1.pdf>

- SRM 6.5 installation—<https://docs.vmware.com/en/Site-Recovery-Manager/6.5/srm-install-config-6-5.pdf>
- SRM 6.0 installation—<https://docs.vmware.com/en/Site-Recovery-Manager/6.0/srm-install-config-6-0.pdf>

You must install an appropriate SRA on the Site Recovery Manager Server hosts at both the protected and recovery sites. If you use more than one type of storage array, you must install the SRA for each type of array on both of the Site Recovery Manager Server hosts.

Before installing an SRA, ensure that SRM and JDK 8 or above version are installed on Windows machines at the protected and recovery sites.

To install an SRA, do the following:

1. Download SRA from the VMware site.

In the <https://my.vmware.com/web/vmware/downloads> page, locate VMware Site Recovery Manager and click **Download Product**. Click **Drivers & Tools**, expand **Storage Replication Adapters**, and click **Go to Downloads**.

2. Copy the Windows installer of SRA to SRM Windows machines at both the protected and recovery sites.
3. Double-click the installer.
4. Click **Next** on the Welcome page of the installer.
5. Accept the EULA and click **Next**.
6. Click **Finish**.



Note

The SRA is installed within the SRM program folder:

```
C:\Program Files\VMware\VMware vCenter Site Recovery Manager\storage\sra
```

After SRA installation, refer the following guide as per the SRM release version and do the SRM environment setup:

- SRM 8.1 configuration—<https://docs.vmware.com/en/Site-Recovery-Manager/8.1/srm-admin-8-1.pdf>
- SRM 6.5 configuration—<https://docs.vmware.com/en/Site-Recovery-Manager/6.5/srm-admin-6-5.pdf>
- SRM 6.0 configuration—<https://docs.vmware.com/en/Site-Recovery-Manager/6.0/srm-admin-6-0.pdf>

After configuration, SRM works with SRA to discover arrays and replicated and exported datastores, and to fail over or test failover datastores.

SRA enables SRM to execute the following workflows:

- Discovery of replicated storage
- Non-disruptive failover test recovery using a writable copy of replicated data
- Emergency or planned failover recovery
- Reverse replication after failover as part of failback
- Restore replication after failover as part of a production test

Data Protection Terms

Interval—Part of the replication schedule configuration, used to enforce how often the protected VMs DP snapshot must be taken and copied to the target cluster.

Local cluster—The cluster you are currently logged into through HX Connect, in a VM replication cluster pair. From the local cluster, you can configure replication protection for locally resident VMs. The VMs are then replicated to the paired remote cluster.

Migration—A routine system maintenance and management task where a recent replication DP snapshot copy of the VM becomes the working VM. The replication pair of source and target cluster do not change.

Primary cluster—An alternative name for the source cluster in VM disaster recovery.

Protected virtual machine—A VM that has replication configured. The protected VMs reside in a datastore on the local cluster of a replication pair. Protected VMs have a replication schedule configured either individually or by inclusion in a protection group.

Protection group—A means to apply the same replication configuration to a group of VMs.

Recovery process—The manual process to recover protected VMs in the event the source cluster fails or a disaster occurs.

Recovery test—A maintenance task that ensures the recovery process will be successful in the event of a disaster.

Remote cluster—One of a VM replication cluster pair. The remote cluster receives the replication snapshots from the Protected VMs in the local cluster.

Replication pair—Two clusters that together provide a remote cluster location for storing the replicated DP snapshots of local cluster VMs.

Clusters in a replication pair can be both a remote and local cluster. Both clusters in a replication pair can have resident VMs. Each cluster is local to its resident VMs. Each cluster is remote to the VMs that reside on the paired local cluster.

DP snapshot—Part of the replication protection mechanism. A type of snapshot taken of a protected VM, which is replicated from the local cluster to the remote cluster.

Secondary cluster—An alternative name for the target cluster in VM disaster recovery.

Source cluster—One of a VM replication cluster pair. The source cluster is where the protected VMs reside.

Target cluster—One of a VM replication cluster pair. The target cluster receives the replicated DP snapshots from the VMs of the source cluster. The target cluster is used to recover the VMs in the event of a disaster on the source cluster.

Best Practices for Data Protection and Disaster Recovery

The requirement for an effective data protection and disaster recovery strategy based on the environment being protected cannot be overstated. The solution should be designed and deployed to meet or exceed the business requirements for both Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO) of the production VMs. The following are some of the points that must be considered when designing this strategy:

- The number of Service Level Agreements (SLA) necessary to comply with various categories of production workloads that may include mission critical, business critical, and important VMs.

- Detailed constructs of each SLA that may include RPO, RTO, the number or recovery points retained, requirements for off-site copies of data, and any requirements for storing backup copies on different media types. There may be additional requirements that include the ability to recover to a different environment such as a different location, different hypervisor or different private/public cloud.
- An ongoing testing strategy for each SLA which serves to prove that the solution meets the business requirements it was designed for.

Note that backups and backup copies must be stored external to the HyperFlex cluster being protected. For example, backups performed to protect VMs on a HyperFlex cluster should not be saved to a backup repository or a disk library that is hosted on the same HyperFlex cluster.

The built-in HyperFlex data protection capabilities are generalized into the following categories:

- **Data Replication Factor**—Refers to the number of redundant copies of data within a HyperFlex cluster. A data replication factor of 2 or 3 can be configured during data platform installation and cannot be changed. The data replication factor benefit is that it contributes to the number of cluster tolerated failures. See the section titled, for additional information about the data replication factor.



Note

Data Replication Factor alone may not fulfill requirements for recovery in the highly unlikely event of a cluster failure, or an extended site outage. Also, the data replication factor does not facilitate point-in-time recovery, retention of multiple recovery points, or creation of point-in-time copies of data external to the cluster.

- **HX Native Snapshots**—Operates on an individual VM basis and enables saving versions of a VM over time. A maximum of 31 total snapshots, including the HX SENTINEL snapshot, can be retained.



Note

HX native snapshots alone may not fulfill requirements for recovery in the unlikely event of a cluster failure, or an extended site outage. Also, HX native snapshots do not facilitate the ability to create point-in-time copies of data external to the cluster. More importantly, unintentional deletion of a VM also deletes any HX native snapshots associated with the deleted VM.

- **Asynchronous Replication**—Also known as The HX Data Platform disaster recovery feature, it enables protection of virtual machines by replicating virtual machine DP snapshots between a pair of network connected HyperFlex clusters. Protected virtual machines running on one cluster replicate to the other cluster in the pair, and vice versa. The two paired clusters typically are located at a distance from each other, with each cluster serving as the disaster recovery site for virtual machines running on the other cluster.



Note

Asynchronous Replication alone may not fulfill requirements for recovery when multiple point-in-time copies need to be retained on the remote cluster. Only the most recent snapshot replica for a given VM is retained on the remote cluster. Also, asynchronous replication does not facilitate the ability to create point-in-time copies of data external to either cluster.

It is recommended to first understand the unique business requirements of the environment and then deploy a comprehensive data protection and disaster recovery solution to meet or exceed those requirements.

Protecting Virtual Machines Overview

To protect a virtual machine (VM), specify the following protection attributes:

- Replication interval, at which DP snapshots are created for replication.
- A start time (within the next 24 hours), which specifies the first-time replication is attempted for that VM.
- Specify if the DP snapshot should be taken with the VM quiesced or not. Proper use of the quiesce option requires that VMware Tools are installed on the VM or VMs being protected.

Protection attributes can be created and assigned to protection groups. To assign those protection attributes to VMs, they can be added to a protection group.

For example, there are three different SLAs: gold, silver, and bronze. Set up a protection group for each SLA, with replication intervals such as 5 or 15 minutes for gold, 4 hours for silver, and 24 hours for bronze. Most VMs can be protected by simply adding them to one of the three already created protection groups.

To protect VMs, you can choose from the following methods:



Note

When you select multiple VMs, you must add them to a protection group.

- **Independently**—Select one VM and configure protection. Set the replication schedule and the VMware quiesce option for the specific VM. Changes to the replication settings will only affect the independently protected VM. The VM is not included in a protection group.
- **Existing protection group**—Select one or more VMs and add them to an existing protection group. The schedule and VMware quiesce option settings are applied to all of the VMs in the protection group. If the protection group settings are changed, the changes are applied to all of the VMs in the protection group.
- **New protection group**—Select two or more VMs and choose to create a new protection group. Define the protection group name, schedule, and VMware quiesce option settings. These settings are applied to all the VMs in the protection group. If the protection group settings are changed, the changes are applied to all of the VMs in the protection group.

Data Protection Workflow

To protect VMs and their data using DP snapshots and replication, perform the following steps:

- Configure two clusters and pair them to each other, to support the replication network activity.
- Assign a replication schedule to the VMs to set the frequency (interval) for creating DP snapshots on the source cluster and replicate them to the target cluster. Replication schedules can be assigned to individual VMs and to protection groups.

Replication Workflow

1. Install HX Data Platform, create two clusters.
2. Create at least one datastore on each cluster.
3. Log in to HX Connect.
4. Before creating the replication network, verify the IP addresses, subnet mask, VLAN, gateway, and IP range to be used for the replication network. After the replication network is created, validate connectivity within the cluster over the replication network.
5. The default value of MTU is 1500. If the HyperFlex cluster uses OTV or other tunneling mechanisms, ensure choosing an MTU value which will work for inter-site or inter-cluster connectivity.
6. Configure cluster replication network on each cluster. The replication network information is unique to each cluster.

Specify the subnet, gateway, range of IP addresses, bandwidth limit for dedicated use by the replication network. HX Data Platform configures a VLAN through UCS Manager for both clusters.
7. An intra-cluster network test is performed to validate connectivity between the nodes in the cluster, after the replication network is configured. If the intra-cluster network test fails, the replication network configuration is rolled back. Reconfigure the replication network after resolving any issues.
8. Before creating the replication pair, ensure that you have updated the corporate network to support this pairing.
9. Create a replication pair from one cluster to the other, connecting the two clusters. After the replication pair is created, a test of the inter-cluster pair network is performed to validate bidirectional connectivity between the clusters. Set the datastore mapping(s) from both clusters.
10. Optionally, you can create protection groups.
 - Set the schedule. Each protection group must have one schedule.
 - Create multiple protection groups if you want to have various replication intervals (schedules) for different VMs. A VM can only belong to one protection group.
11. Select VMs to protect, as individual virtual machines or VMs assigned to protection groups.
12. Set protection, do the following:
 - a. Select one or more VMs. Click Protect.
 - b. From the Protect VM wizard, the options are:
 - Protect a single VM with an existing protection group.
 - Protect a single VM independently.

Set the schedule.
 - Protect multiple VMs with an existing protection group.
 - Protect multiple VMs with a new protection group.

Create new protection group and set schedule.

Configuring the Replication Network in HX Connect

Before a replication pair can be configured, the replication network has to be configured on both the local and remote cluster. Complete the configuration on the local cluster, then log in to the remote cluster and complete the configuration there.

Before you begin

Ensure that the following prerequisites are met, before configuring the replication network:

- A minimum of $N + 1$ IP addresses is required, where N is the number of converged nodes. An IP subnet spanning these new IP addresses, the gateway, and VLAN associated with this subnet is also required.
- To accommodate future cluster expansion, ensure that there are sufficient IP addresses in the subnet provided, for later use. Any new nodes in the expanded cluster would also need to be assigned IP addresses for replication. The subnet provided in the previous step should span the potentially new IP range as well.
- Additional IP-pool ranges can be added to the network later, however IP-pools already configured in the replication network cannot be modified.
- Make sure that the IP addresses to be used for the replication network are not already in use by other systems.
- Before creating the replication network, verify IP addresses, Subnet, VLAN, and Gateway to be used for the replication network.

Step 1 Log in to HX Connect as a user with administrator privileges.

Step 2 Select **Replication > Replication Configuration > Configure Network**.

Note You can only configure the replication network once. Once configured, you can edit the available IP addresses and the networking bandwidth.

Step 3 In the **Configure Replication Network** dialog box, under the tab, enter the network information.

UI Element	Essential Information
Select an existing VLAN radio button	Click this radio button to add an existing VLAN. If you manually configured a VLAN for use by the replication network through Cisco UCS Manager, enter that VLAN ID.
Create a new VLAN radio button	Click this radio button to create a new VLAN.

UI Element	Essential Information
VLAN ID field	<p>Click the up or down arrows to select a number for the VLAN ID or type a number in the field.</p> <p>This is separate from the HX Data Platform Management traffic network and Data traffic network.</p> <p>Important Be sure to use a different VLAN ID number for each HX Storage Cluster in the replication pair.</p> <p>Replication is between two HX Storage Clusters. Each HX Storage Cluster requires a VLAN dedicated to the replication network.</p> <p>For example, 3.</p> <p>When a value is added, the default VLAN Name is updated to include the additional identifier. The VLAN ID value does not affect a manually entered VLAN name.</p>
VLAN Name field	<p>This field is automatically populated with a default VLAN name when the Create a new VLAN radio button is selected. The VLAN ID is concatenated to the name.</p>
<p>For Stretched Cluster, provide Cisco UCS Manager credentials for primary and secondary FIs (site A and site B). For normal cluster, provide Cisco UCS Manager credential for single FI.</p>	
UCS Manager host IP or FQDN field	<p>Enter Cisco UCS Manager FQDN or IP address.</p> <p>For example, 10.193.211.120.</p>
Username field	Enter administrative username for Cisco UCS Manager.
Password field	Enter administrative password for Cisco UCS Manager.

Step 4 Click **Next**.

Step 5 In the **IP & Bandwidth Configuration** tab, set the network parameters and the replication bandwidth.

UI Element	Essential Information
Subnet field	<p>Enter the subnet for use by the replication network in network prefix notation. The subnet is separate from the HX Data Platform Management traffic network and Data traffic network.</p> <p>Format example: x.x.x.x/<number of bits> 1.1.1.1/20</p>
Gateway field	<p>Enter the gateway IP address for use by the replication network. The gateway is separate from the HX Data Platform Management traffic network and Data traffic network.</p> <p>For example, 1.2.3.4.</p> <p>Note The gateway IP address must be accessible even if the disaster recovery is being setup for a flat network.</p>

UI Element	Essential Information
IP Range field	<p>Enter a range of IP addresses for use by the replication network.</p> <ul style="list-style-type: none">• The minimum number of IP addresses required is the number of nodes in your HX Storage Cluster plus one more. <p>For example, if you have a 4 node HX Storage Cluster, enter a range of at least 5 IP addresses.</p> <ul style="list-style-type: none">• The from value must be lower than the to value. <p>For example, <i>From 10.10.10.20 To 10.10.10.30</i>.</p> <ul style="list-style-type: none">• If you plan to add nodes to your cluster, include sufficient number of IP addresses to cover any additional nodes. You can add IP addresses at any time. <p>Note The IP address range excludes compute-only nodes.</p>
Add IP Range button	Click to add the range of IP addresses entered in IP Range From and To fields.

UI Element	Essential Information
Set Replication Bandwidth Limit check box	<p>Enter the maximum network bandwidth that the replication network is allowed to consume for inbound and outbound traffic. Acceptable value is 33,000 Mbits/sec.</p> <p>The default value is <code>unlimited</code>, which sets the maximum network bandwidth to the total available to the network.</p> <p>The replication bandwidth is used to copy replication snapshots from this local HX Storage Cluster to the paired remote HX Storage Cluster.</p> <p>Note</p> <ul style="list-style-type: none"> • At lower bandwidth (typically, lesser than 50 Mbps), the replications of multiple VMs may exit without executing the replication process due to high data transfer rate. To overcome this issue, either increase the bandwidth or stagger VM replication schedule so that VMs do not replicate in the same window. • The bandwidth setting must be close to the link speed. The bandwidth setting for the clusters in the pair must be same. • The set bandwidth is applicable only for the incoming and outgoing traffic of the cluster to which the bandwidth is set to. For example, setting the bandwidth limit as 100Mb means that the 100Mb is set for incoming traffic and 100Mb is set for outgoing traffic. • The set bandwidth limit must not exceed the physical bandwidth. • The bandwidth configured must be same on both sites of the disaster recovery environment. • The allowed low bandwidth is 10Mb and the maximum latency supported with 10Mb is 75ms. If the initial replication of VMs fails due to lossy network or unstable HX clusters, the VM replication will be initiated again in the next schedule as a fresh replication job. In this case, you have to size the schedule accordingly to protect VMs.
Set non-default MTU check box	<p>Default MTU value is 1500.</p> <p>Select the check box to set a custom MTU size for the replication network. MTU can be set in the range 1024 to 1500.</p> <p>Note</p> <ul style="list-style-type: none"> • Ensure to use the same MTU value on both sides of the cluster. • After configuring the cluster, if the MTU value needs to be changed, the existing replication network configuration will need to be removed. The replication network can then be configured with the correct MTU value.

Note When you use an existing VLAN for replication network, the replication network configuration fails. You must add the self-created replication VLAN to the management vNIC templates in Cisco UCS Manager.

- Step 6** Click **Next**.
- Step 7** In the **Test Configuration** tab, check the replication network configuration.
- Step 8** Click **Configure**.

What to do next

- Be sure to configure the replication network on both HX Storage Clusters for the replication pair.
- After the replication network is created on the cluster, each converged node on the cluster would be configured with an IP address on the eth2 interface.
- Check for duplicate IP assignment using 'arping'.

For example: `arping -D -b -c2 -I ethX $replicationIP`` (replace ethX and \$replicationIP with actual values).

If there is a duplicate IP assignment, it is necessary to remove the replication network assignments.

Test Local Replication Network

To perform an intra-cluster replication network test, do the following:

-
- Step 1** Log in to HX Connect.
- Enter the HX Storage Cluster management IP address in a browser. Navigate to `https://<storage-cluster-management-ip>`.
 - Enter the administrative username and password.
 - Click **Login**.
- Step 2** In the Navigation pane, click **Replication**.
- Step 3** From the **Actions** drop-down list, select **Test Local Replication Network**.
- Step 4** Click **Run Test**.
- Step 5** On the **Activity** page, you can view the progress of the *Test Replication Network* task.
-

Editing the Replication Network

When you expand a HX Cluster that has replication configured, ensure that you have sufficient IP addresses available for the replication network. The replication network requires dedicated IP addresses, one for every node in the cluster plus one more. For example, in a 3 node cluster, four IP addresses are required. If you are adding one more nodes to the cluster, a minimum of five IP addresses are required.

To edit the replication network to add IP addresses, perform the following tasks:

-
- Step 1** Log in to HX Connect as administrator.
- Step 2** In the Navigation pane, Select **Replication**.
- Step 3** From the **Actions** drop-down list, select **Edit Replication Network**.

Step 4 In the **Edit Network Configuration** dialog box, you can edit the range of IPs to use and set the replication bandwidth limit for replication traffic. The replication network subnet and gateway are displayed for reference only and cannot be edited.

UI Element	Essential Information
Subnet field	The subnet that is configured for the replication network in network prefix notation. This value cannot be edited.
Gateway field	The gateway that is configured for the replication network. This is value cannot be edited.
IP Range field	<p>Enter a range of IP addresses for use by the replication network.</p> <ul style="list-style-type: none"> The minimum number of IP addresses required is the number of nodes in the HX Storage Cluster plus one more. <p>For example, if the HX Storage Cluster has 4 nodes, the IP Range must be at least 5 IP addresses.</p> <ul style="list-style-type: none"> The from value must be lower than the to value. <p>For example, <i>From 10.10.10.20 To 10.10.10.30</i>.</p> <ul style="list-style-type: none"> You can add IP addresses at any time. If you plan to add nodes to your cluster, include sufficient number of IP addresses to accommodate any additional nodes. <p>Note The IP address range excludes compute-only nodes.</p>
Add IP Range field	Click to add the range of IP addresses that are entered in IP Range From and To fields.
Set replication bandwidth limit check box (Optional)	<p>Enter the maximum network bandwidth that the replication network is allowed to consume for outbound traffic.</p> <p>Valid Range: 10 to 10,000. The default is <code>unlimited</code>, which sets the maximum network bandwidth to the total available replication network bandwidth.</p> <p>The replication bandwidth is used to copy DP snapshots from the local HX Storage Cluster to the paired remote HX Storage Cluster.</p>

Step 5 Click **Save Changes**.

The replication network is now updated. Any additional replication network IP addresses are available for new nodes should they be added to the storage cluster. Replication traffic adjusts to any changes made to the bandwidth limit.

Replication Pair Overview

Creating a replication cluster pair is a prerequisite for configuring VMs for replication. After two (2) HX clusters are paired, map the datastore on the remote cluster to a datastore on the local cluster.

Mapping datastore A on HX cluster 1 with a datastore B on HX cluster 2 enables any VM on HX cluster 1 that resides in datastore A and is configured for replication to be replicated to datastore B on HX cluster 2. Similarly, any VM on cluster 2 that resides in datastore B and is configured for replication to be replicated to datastore A on HX cluster 1.

Pairing is strictly 1-to-1. A cluster can be paired with no more than one other cluster.

Mapping is a strict 1-to-1 relationship. A datastore on a paired HX cluster can be mapped to no more than one datastore on the other HX cluster. Note that there can be multiple mapped datastores. For example, datastore A on HX cluster 1 mapped to datastore B on HX cluster 2, and datastore C on HX cluster 1 mapped to datastore D on HX cluster 2.

Creating a Replication Pair

The replication pair defines the two halves of the protection network. The HX cluster you are logged into is the local cluster, the first half of the pair. Through this dialog, you identify another HX cluster, the second half of the pair, the remote cluster. To configure the storage component, map the replication pair to datastores on each HX cluster. After the replication pair is configured, and at least one pair of datastores have been mapped, you can begin protecting virtual machines. See the **Virtual Machines** tab.



Important

When pairing clusters or mapping datastores with different HXDP versions, cluster pairing must be initiated on the 3.5 HX cluster and then the datastore mapping must be initiated from the 3.0 HX cluster.

- Cluster pairing must be initiated only from a 3.5 HX cluster to a 3.0 HX cluster.
- Datastore mapping must be initiated only from a 3.0 HX cluster to a 3.5 HX cluster.



Note

When pairing HX clusters, if you get the error: `Check your cluster status or logs for possible solutions` appears, check if the pairing is successful by running the following command:

```
stcli dp peer list
```

If the pairing is not successful, check the logs for solutions.

Before you begin

- Create a datastore on both the local and the remote cluster.
- Configure the replication network.

Step 1

From HX Connect, log in to either the local or remote HX cluster as a user with administrator privileges and do one of the following:

- Select **Replication > Pair Cluster** if you are doing cluster pairing for the first time.
- Select **Replication > Create Replication Pair**.

The **Create Replication Pair** option is enabled only when you delete an existing replication pair after unprotecting all the VMs and removing all the dependencies.

Step 2

Enter a **Name** for the replication pair and click **Next**.

Enter a name for the replication pairing between two HX Storage clusters. This name is set for both the local and remote cluster. The name cannot be changed.

Step 3 Enter the **Remote Connection** identification and click **Pair**.

UI Element	Essential Information
Management IP or FQDN field	Enter the cluster IP address or fully qualified domain name (FQDN) for the management network on the remote . For example: <i>10.10.10.10</i> .
User Name and Password fields	Enter vCenter single sign-on or cluster specific administrator credentials for the remote HX cluster.

HX Data Platform verifies the remote HX cluster and assigns the replication pair name.

Once the Test Cluster Pair job is successful, you can proceed to the next step. On the Activity page, you can view the progress of the Test Cluster Pair job.

Note Virtual machines to be protected must reside on one of the datastores in the replication pair.

Step 4 Click **Next**.

The **Create New Replication Pair** dialog box appears.

Step 5 To protect VMs using the HX Data Platform disaster recovery feature, click **Native Protection** and do the following:

- The **Local Datastore** column displays a list of the configured datastores on the local HX Storage cluster. Map one local datastore to one remote datastore.
- From the **Remote Datastore** pull-down menu, choose a datastore that needs to be paired with the local datastore.
- Click **Map Datastores**.

If you chose to cancel the datastore mapping by clicking **Cancel**, you can map the datastores later using **Map Datastores** that appears under **DATASTORE MAPPED** in the Replication dashboard.

To change the local datastore selection:

- From the **Remote Datastore** pull-down menu, choose **Do not map this datastore** to remove the mapping from the current local datastore.
- From the **Remote Datastore** pull-down menu, choose a datastore to be paired with the local datastore.

Note

- The virtual machines to be protected must be on the datastores you select. Moving virtual machines from the configured datastores for the replication pair, also removes protection from the virtual machines.
- Moving virtual machine to another paired datastore is supported. If the VMs are moved to unpaired datastore, the replication operation fails.

Note Once a local datastore is mapped to a remote datastore, the corresponding local datastore will not appear under **Other DRO Protection**.

Step 6 To protect VMs using SRM through disaster recovery orchestrator (DRO), click **Other DRO Protection** and do the following:

- The **Local Datastore** column displays a list of the unpaired configured datastores on the local HX cluster. Map one local datastore to one remote datastore.
- From the **Remote Datastore** pull-down menu, choose a datastore that need to be paired with the local datastore.

- c) From the **Direction** pull-down menu, choose **Incoming** or **Outgoing** as the direction of VM movement for the mapped datastores.
- d) From the **Protection Schedule** pull-down menu, choose the schedule for protecting all the VMs in the datastore.
- e) Click **Map Datastores**.

If you chose to cancel the datastore mapping by clicking **Cancel**, you can map the datastores later using **Map Datastores** that appears under **DATASTORE MAPPED** in the Replication dashboard.

Note If a new VM is added to the protected datastore, the newly added VM is also get protected by Cisco HyperFlex.

Note The replication pairs that are edited under **Other DRO Protection**, are exposed to SRM.

What to do next

To check the protection status of virtual machines, do one of the following:

- Click **Virtual Machines** in HX Connect. This displays the list of the virtual machines on the local cluster along with the protection status. If the VM is protected by SRM, the status is displayed as **Protected (by other DRO)**.



Note In the **Virtual Machine** page, the status of the VMs protected by SRM is displayed as **unprotected** until the completion of first auto protect cycle. Until then, the user is not recommended to manually protect those VMs.

- Click **Replication** in HX Connect.
- Click **Protection Group** under the **Local VMs** tab to view the VMs protected within a protection group. Click **Other DRO** under the **Local VMs** to view the VMs protected by SRM.
- Click **Replication** in HX Connect. Click **Replication Activity** to view the replication activity status of the protected VMs. If the VM is protected by SRM, the status is displayed as **Protected (by other DRO)**.

Test Remote Replication Network

To test the pairing between clusters in a remote replication network, do the following:

- Step 1** Log in to HX Connect.
 - a) Enter the HX Storage Cluster management IP address in a browser. Navigate to *https://<storage-cluster-management-ip>*.
 - b) Enter the administrative username and password.
 - c) Click **Login**.
- Step 2** In the Navigation pane, click **Replication**.
- Step 3** From the **Actions** drop-down list, select **Test Remote Replication Network**.

Field	Description
MTU Test Value	<p>Default MTU value is 1500. MTU can be set in the range 1024 to 1500.</p> <p>Note</p> <ul style="list-style-type: none"> After configuring the cluster, if the MTU value needs to be changed, the existing replication network configuration will need to be removed. The replication network can then be configured with the correct MTU value. Be sure to use the same MTU value on both sides of the cluster.

Step 4 Click **Run Test**.

Step 5 On the **Activity** page, you can view the progress of the *Replication Pair Network Check* task.

Editing a Mapped Datastore Replication Pair

Editing a replication pair is changing the datastores for the replication pair.

Step 1 Log in to HX Connect as an administrator.

Step 2 Select **Replication > Replication Pairs**.

Step 3 Select the replication pair that needs to be edited and click **Edit**.

The **Edit Replication Pair** dialog box appears.

Step 4 To protect VMs using the HX Data Platform disaster recovery feature, click **Native Protection** and do the following:

- The **Local Datastore** column displays a list of the configured datastores on the local HX Storage Cluster. Map one local datastore to one remote datastore.
- From the **Remote Datastore** pull-down menu, choose a datastore that needs to be paired with the local datastore.
- Click **Map Datastores**.

To change the local datastore selection:

- From the **Remote Datastore** pull-down menu, choose **Do not map this datastore** to remove the mapping from the current local datastore.
- From the **Remote Datastore** pull-down menu, choose a datastore to be paired with the local datastore.

Note Once a local datastore is mapped to a remote datastore, the corresponding local datastore will not appear under **Other DRO Protection**.

Step 5 To protect VMs using SRM through disaster recovery orchestrator (DRO), click **Other DRO Protection** and do the following:

- The **Local Datastore** column displays a list of the unpaired configured datastores on the local HX cluster. Map one local datastore to one remote datastore.
- From the **Remote Datastore** pull-down menu, choose a datastore that need to be paired with the local datastore.

- c) From the **Direction** pull-down menu, choose **Incoming** or **Outgoing** as the direction of VM movement for the mapped datastores.
- d) From the **Protection Schedule** pull-down menu, choose the schedule for protecting all the VMs in the datastore.
- e) Click **Map Datastores**.

Note New VMs added to the protected datastore are also protected.

Note The replication pairs that are edited under **Other DRO Protection**, are exposed to SRM.

What to do next

To check the protection status of virtual machines, do one of the following:

- Click **Virtual Machines** in HX Connect. This displays the list of the virtual machines on the local cluster along with the protection status. If the VM is protected by SRM, the status is displayed as **Protected (by other DRO)**.



Note In the **Virtual Machine** page, the status of the VMs protected by SRM is displayed as **unprotected** until the completion of first auto protect cycle. Until then, the user is not recommended to manually protect those VMs.

- Click **Replication** in HX Connect.
- Click **Protection Group** under the **Local VMs** tab to view the VMs protected within a protection group. Click **Other DRO** under the **Local VMs** to view the VMs protected by SRM.
- Click **Replication** in HX Connect. Click **Replication Activity** to view the replication activity status of the protected VMs. If the VM is protected by SRM, the status is displayed as **Protected (by other DRO)**.

Deleting a Replication Pair

Delete a replication pair on the local and remote clusters.

Select **Replication > Replication Pairs > Delete**.

Before you begin

On both the local and remote HX clusters, remove dependencies from the replication pair.

Log in to the local and the remote HX storage cluster and perform the following:

- Unprotect all virtual machines. Remove virtual machines from protection groups.
- Remove protection groups. If the protection group does not have a VM, deleting protection group is not required.

Step 1 Log in to HX Connect as an administrator.

Step 2 Unmap the datastores in the replication pair.

- a) Select **Replication > Replication Pairs > Edit**.

After the test cluster pair job is successful, you can proceed to the next step. You can view the progress of the Test Cluster Pair job on the Activity page.

- b) From the **Edit Replication Pair** dialog box, select **Do not map this datastore** from the **Remote Datastore** menu.

UI Element	Essential Information
Local Datastore column	List of the configured datastores on this cluster, the local HX clusters. Map one local datastore to one remote datastore.
Remote Datastore column	Pair the datastores between the HX clusters. <ol style="list-style-type: none"> 1. To change the local datastore selection, remove the mapping to the current local datastore. From the pull-down menu in the Remote Datastore column, select Do not map this datastore. 2. From the desired Local Datastore row, select a datastore from the Remote Datastore pull-down menu. This selects both the remote and local datastores in a single action.

- c) Ensure all the possible remote datastores are set to **Do not map this datastore**.

- d) Click **Finish**.

Step 3 Select **Replication > Replication Pairs > Delete**.

Step 4 Enter administrator credentials for the remote cluster and click **Delete**.

UI Element	Essential Information
User Name field	Enter the administrator user name for the remote HX Storage Cluster.
Password field	Enter the administrator password for the remote HX Storage Cluster.

Creating a Protection Group

A protection group is a group of VMs with the same replication schedule and VMware Tools quiescence settings.

Protection groups are created on the HX cluster that the administrative user is logged on to. Protection groups provide protection to the VMs that are members of a given protection group. If protection groups have protected virtual machines that replicate to the remote cluster, they are listed in HX Connect.



Note The administration of protection groups can only be performed from the local cluster where it was created.

Before you begin

Ensure that replication network and replication pair are configured.

- Step 1** Log in to HX Connect as an administrator.
- Step 2** Select **Replication > Protection Groups > Create Protection Group**.
- Step 3** Enter the information in the dialog fields.

UI Element	Essential Information
Protection Group Name field	Enter a name for the new protection group for this HX cluster. Protection groups are unique to each HX cluster. The name is referenced on the remote HX cluster, but not editable on the remote HX cluster. Multiple protection groups can be created on each HX cluster.
Protect virtual machines in this group every field	Select how often the virtual machines are to be replicated to the paired cluster. The pull-down menu options are: 5 minutes, 15 minutes, 30 minutes, 1 hour, 90 minutes, 2 hours, 4 hours, 8 hours, 12 hours, and 24 hours. The default value is 1 hour.
Start protecting the virtual machines immediately radio button	Select this radio button if you want the first replication to start immediately after you add the first virtual machine to this protection group.
Start protecting the virtual machines at radio button	Select this radio button if you want to set a specific time for the first replication operation to start. Before you start replication ensure: <ul style="list-style-type: none"> • At least one virtual machine is added to the protection group. • The scheduled start time is reached. To specify the protection operation start time: <ol style="list-style-type: none"> Check the Start protecting the virtual machines at radio button. Click in the time field and select an hour and minute. Then click out of the field. Cluster time zone and Current time on cluster are references to help you to choose the appropriate replication start time. Start time is based on the local cluster clock. For example: 10 hours, 3 minutes from now with Current time on cluster, 1:56:15PM, means that the first replication occurs at 11:59:00PM. The <i>hours, minutes from now</i> states when the first replication will occur. This is updated when you change the time field setting.
Use VMware Tools to quiesce the virtual machine check box	Select this check box to take quiesced DP snapshots. Leaving the check box in an unchecked state will take crash consistent DP snapshots. This only applies to virtual machines with VMware Tools installed.

- Step 4** Click **Create Protection Group**.

HX Data Platform adds the new group to the **Protection Groups** tab. This protection group is available to protect virtual machines on this cluster.

Step 5 Click the **Replication > Protection Groups** to view or edit the new protection group.

If the number of VMs is zero (0), add virtual machines to the new protection group to apply the replication schedule configured in the protection group.

Editing Protection Groups

Change the replication interval (schedule) for the virtual machines in the protection group. To edit the protection groups, perform the following steps:

Step 1 Login to HX Connect as an administrator.

Step 2 Select **Replication > Protection Groups > Edit Schedule**.

Step 3 Edit the information in the dialog fields.

UI Element	Essential Information
Protect virtual machines in this group every field	Use the pull-down list to select how often the virtual machines are to be replicated to the paired cluster. List values are: 5 minutes, 15 minutes, 30 minutes, 1 hour, 90 minutes, 2 hours, 4 hours, 8 hours, 12 hours, and 24 hours.
Use VMware Tools to quiesce the virtual machine check box	Select the check box to take quiesced DP snapshots. The checkbox is unchecked by default; leaving the check box unchecked, takes crash consistent DP snapshots. This only applies to virtual machines with VMware Tools installed.

Step 4 Click **Save Changes** to save the interval and VMware Tools quiescence settings for the protection group. View the Protection Groups tab to verify the interval frequency.

Deleting Protection Groups

Before you begin

Remove all virtual machines from the protection group.

Step 1 Select **Replication > Protection Groups > *protection_group_name***

Step 2 Click **Delete**. Click **Delete** in the verification pop-up.

Protecting Virtual Machines with an Existing Protection Group

This task describes how to protect multiple virtual machines using an existing protection group.

Using an **Existing protection group**—Select one or more virtual machines and add them to an existing protection group. The schedule and VMware quiesce option settings are applied to all the virtual machines in

the protection group. When the protection group settings are changed, the changes are applied to all the virtual machines in the protection group.

Before you begin

Replication network and replication pair configured.

Create protection group prior to adding the virtual machines.

Step 1 Log in to HX Connect with administrator privileges and select **Virtual Machines**.

This lists the virtual machines on the local HX cluster.

Step 2 Select one (1) or more unprotected VMs from the list.

Click in the virtual machine row to select it. As you click a virtual machine row, the corresponding virtual machine check box is selected.

Step 3 Click **Protect**.

The **Protect Virtual Machines** wizard, **Protection Group** page is displayed.

Step 4 Click the radio button, **Add to an existing protection group**

UI Element	Essential Information
Set the protection parameters table	Verify the selected virtual machine Name . Use the Storage Provisioned and Storage Used to check for sufficient storage resource availability on the remote HX cluster.
Add to an existing protection group radio button	Select an existing protection group from the pull-down list. The interval and schedule settings of the protection group are applied to the selected VM or VMs.
Create a new protection group radio button	Enter a name for the new protection group for this local cluster. Protection groups are unique to each cluster. The name is referenced on the remote cluster, but not editable on the remote cluster. You can create multiple protection groups on each cluster.

Step 5 Select a protection group from the pull-down list and click **Next**.

Be sure the protection group you choose has the schedule interval desired.

The **Protect Virtual Machines** wizard, **Summary** page is displayed.

Step 6 Confirm the information in the **Summary** page and click **Add to Protection Group**.

The selected VM or VMs are added to the protection group. View the **Replication** or **Virtual Machines** pages to confirm that the VM or VMs have been added to the protection group.

Protecting Virtual Machines with a New Protection Group

This task describes how to protect multiple virtual machines by creating a new protection group.

Using a **New protection group**—Select VMs and choose to create a new protection group. Define the protection group name, schedule, start time, and VMware quiesce option settings. These settings are applied to all the virtual machines in the protection group. When the protection group settings are changed, the changes are applied to all the virtual machines in the protection group.

Before you begin

Replication network and replication pair configured.

Step 1 Login to HX Connect with administrator privileges and select **Virtual Machines**.

This lists the virtual machines on the local HX cluster.

Step 2 Select one or more unprotected VMs from the list.

Click in the virtual machine row to select it. As you click a virtual machine row, the corresponding virtual machine checkbox is selected.

Step 3 Click **Protect**.

The **Protect Virtual Machines** wizard, **Protection Group** page is displayed.

Step 4 Click the radio button, **Create a new protection group**, add a name for the protection group, and click **Next**.

The **Protection Schedule Wizard Page** wizard page is displayed.

Step 5 Complete the schedule and VMware quiesce option, as needed, and click **Next**.

UI Element	Essential Information
Protect virtual machines in this group every field	Select how often the virtual machines are to be replicated to the paired cluster. The default value is every 1 hour.
Start protecting the virtual machines immediately radio button	Select this radio button if you want the first replication to start immediately after you add the first virtual machine to this protection group.

UI Element	Essential Information
<p>Start protecting the virtual machines at radio button</p>	<p>Select this radio button if you want to set a specific time for the first replication to start. To start replication requires:</p> <ul style="list-style-type: none"> • At least one virtual machine is added to the protection group. • The scheduled start time is reached. <p>To specify the protection operation start time:</p> <ol style="list-style-type: none"> a. Check the Start protecting the virtual machines at radio button. b. Click in the time field and select an hour and minute. Then click out of the field. <p>The <i>hours, minutes from now</i> states when the first replication will occur. This is updated when you change the time field setting.</p> <p>Cluster time zone and Current time on cluster are references to help you to choose the appropriate replication start time. Start time is based on the local cluster clock. For example:</p> <p>10 hours, 3 minutes from now with Current time on cluster, 1:56:15PM, means that the first replication occurs at 11:59:00PM.</p>
<p>Use VMware Tools to quiesce the virtual machine check box</p>	<p>Click the check box to take quiesced DP snapshots. An unchecked check box takes crash consistent DP snapshots. The check box is unchecked by default.</p> <p>This only applies to virtual machines with VMware Tools installed.</p>

The **Protect Virtual Machines** wizard, **Summary** page is displayed.

Step 6

Confirm the information in the **Summary** page and click **Add to Protection Group**.

Review the summary content to confirm the settings to apply to the selected virtual machines.

- Name of the protection group
- Number of virtual machines to protect
- Names of virtual machines
- Storage provisioned for each virtual machine
- Storage used (consumed) by each virtual machine

The selected VM or VMs are added to the protection group. View the **Replication** or **Virtual Machines** pages to verify that the VM(s) have been added to the protection group.

Protecting Individual Virtual Machines

This task describes how to protect a virtual machine (VM).

- **Independently**—Select one (1) VM and configure protection. Set the replication schedule and the VMware Tools quiesce option for the specific VM.

Changes to the replication settings only affect the independently protected VM. The VM is not a member of a protection group.

- **Existing protection group**—Select one or more VMs and add them to an existing protection group. The schedule and VMware Tools quiesce option settings are applied to all the VMs in the protection group. When the protection group settings are changed, the changes are applied to all VMs in the protection group.

Before you begin

Configure replication network and replication pair.

Step 1 Log in to HX Connect with administrator privileges and select **Virtual Machines**.

Step 2 Select one unprotected virtual machine from the list. Click in the virtual machine row to select it.

Click in the virtual machine row to select it. As you click a virtual machine row, the corresponding virtual machine check box is selected.

Step 3 Click **Protect**.

The **Protect Virtual Machine** dialog box is displayed.

Step 4 Complete the fields as needed.

UI Element	Essential Information
Add to an existing protection group radio button	<p>Select an existing protection group from the pull-down list.</p> <p>The interval and schedule settings of the protection group are applied to this virtual machine.</p> <p>No additional configuration is required, click Protect Virtual Machine.</p>
Protect this virtual machine independently radio button	<p>Enables the interval, schedule options, and VMware Tools quiescence option for defining protection for this VM.</p>
Protect this virtual machine every field	<p>Select from the pull-down list how often the virtual machines are to be replicated to the paired cluster.</p> <p>The list values are: 5 minutes, 15 minutes, 30 minutes, 1 hour, 90 minutes, 2 hours, 4 hours, 8 hours, 12 hours, and 24 hours.</p>

UI Element	Essential Information
Start protecting the virtual machines immediately radio button	Select this radio button if you want the first replication to start immediately after you add the first virtual machine to this protection group.
Start protecting the virtual machines at radio button	<p>Select this radio button if you want to set a specific time for the first replication to start. To start replication requires:</p> <ul style="list-style-type: none"> • At least one VM is added to the protection group. • The scheduled start time is reached. <p>To specify the protection operation start time:</p> <ol style="list-style-type: none"> Check the Start protecting the virtual machines at radio button. Click in the time field and select an hour and minute. Then click out of the field. <p>The <i>hours, minutes from now</i> states when the first replication will occur. This is updated when you change the time field setting.</p> <p>Cluster time zone and Current time on cluster are references to help you to choose the appropriate replication start time. Start time is based on the local cluster clock. For example:</p> <p>10 hours, 3 minutes from now with Current time on cluster, 1:56:15PM, means that the first replication occurs at 11:59:00PM.</p>
VMware Tools to quiesce the virtual machine check box	<p>To take quiesced DP snapshots, check the check box. The unchecked check box takes crash consistent DP snapshots. The check box is unchecked by default.</p> <p>This only applies to virtual machines with VMware Tools installed.</p>

Step 5 Click **Protect Virtual Machine**.

The VM status is updated in the **Virtual Machine** and the **Replication** page. Notice on the Replication page, that no Protection Group is listed for any VMs protected as individual VMs.

Replication is now enabled for this VM.

Unprotecting Virtual Machines



Note There is no need to unprotect VMs to pause replication for HX cluster activities. See [Pausing Replication, on page 230](#).

Step 1 Log in to HX Connect as an administrator.

Step 2 Select **Virtual Machines**.

Step 3 Select a protected virtual machine from the list. Click in the virtual machine row.

VMs can be unprotected one VM at a time.

Step 4 Click **Unprotect** and click to confirm.

The state changes for the virtual machine from **protected** to **unprotected**.

Disaster Recovery Overview

Disaster recovery is performed when the source site is not reachable and you want to failover the VMs and the protected groups to the target cluster. The process of recovery recovers the VM on the target cluster. Recovering virtual machines is restoring a most recent replication snapshot from the recovery (target) cluster.

Testing VM recovery—The ability to test recovery without breaking replication. It can bring up your VM workload on the target to verify the contents of the VM.

Recovering virtual machines—Restoring a most recent replication snapshot from the target (recovery) cluster. Once you start Recovery, all scheduled replication will be stopped.

Planned migration—Performing planned migration pauses the replication schedule, creates and replicates a DP snapshot, and recovers on the target. Ownership is switched from the source to the target, and resumes replication on the target that is now the new source

Disaster Recovery and Reprotect—Recovers the VM on the target, switches the ownership from the source to the target, and resumes replication on the target that is now the new source.

Protecting VMs after disaster—In the event of a disaster, you may lose the source site altogether. After the recovery is performed protect the recovered VMs to a newer cluster.

Configuring the Recovery Settings

The configuration of recovery settings allows defining global recovery parameters and mapping for resources across recovery sites. It is possible to configure the folder, network, or resource pool parameters to be used during recovery, test recovery and migrate operations. If the global recovery setting is not configured, explicitly mapping individual VMs at the time of recovery is required.

Step 1 Log in to HX Connect as administrator and do one of the following:

- a) If configuring recovery settings for the first time, select **Replication > Configure**.
- b) Select **Replication** and click **Actions** next to **Recovery Settings**. From the **Actions** drop-down list, choose **Edit Recovery Settings**.

Step 2 In the **Recovery Settings** dialog box, enter the following fields:

Field	Description
Virtual Machine Power State radio button	By default, the Off option is selected. The recovered VM is powered on as per the selected option.
Test Virtual Machine Name Prefix field	Enter a prefix that you want to add to the virtual machine after test recovery. The prefix helps to identify the type and context of the resources.

Field	Description
Notification Setting radio button	Choose Normal Mode to get a confirmation prompt with summary of configuration at the time of recovery, test recovery, or migration. Choose Silent Mode to not get a confirmation prompt. On choosing silent mode, a confirmation window appears with the description of default behavior of silent mode. If you agree to the default behavior of silent mode, click OK .
Resource Pool area	<p>Map the resources in the protected site to the resources in the remote site for the recovery, test recovery, and migrate configuration.</p> <p>Check the Same as Recovery Configuration check box to use the resource mapping of recovery configuration for the test recovery configuration.</p> <p>Click Add Rule to add one more resource pool mapping. Click the delete icon to remove a rule. To edit a rule, delete the rule and add the updated rule as a new rule.</p>
Folder area	<p>Map the folders in the protected site to the folders in the remote site for the recovery, test recovery, and migrate configuration.</p> <p>Check the Same as Recovery Configuration check box to use the folder mapping of recovery configuration for the test recovery configuration.</p> <p>Click Add Rule to add one more folder mapping. Click the delete icon to remove a rule. To edit a rule, delete the rule and add the updated rule as a new rule.</p>
Network area	<p>Map the network in the protected site to the network in the remote site for the recovery, test recovery, and migrate configuration.</p> <p>Check the Same as Recovery Configuration check box to use the network mapping of recovery configuration for the test recovery configuration.</p> <p>Click Add Rule to add one more network mapping. Click the delete icon to remove a rule. To edit a rule, delete the rule and add the updated rule as a new rule.</p>

Step 3

Click **Save**.

On successful saving of the recovery settings, the **RECOVERY SETTINGS** field on the **Replication** page displays one of the following status along with the notification setting mode:

- **Partially Configured**—This status is displayed when you have not set the recovery mapping for any of the resources or if any of the configured mapping is invalid.
- **Configured**—This status is displayed when all the recovery settings are configured and valid.

Note The **RECOVERY SETTINGS** field shows the last validated result. Once a rule is created for recovery, there is no automatic periodic validation. However, a validation job can be executed to check the validity of the rules existing in the recovery settings

The validation job summary in the **Activity** page directs the user to check the **Recovery Settings** page to view the validation result.

After configuring the recovery settings, validation of the recovery settings can be performed by choosing **Validate Recovery Settings** from the **Actions** drop-down list. The successful initiation of the recovery setting validation message is displayed. The **RECOVERY SETTINGS** field displays the time stamp of last validation. To monitor the progress of validation, click the **Activity** tab. In the normal notification setting mode, during recovery, test recovery, or migration of a virtual machine, the configured recovery settings are displayed.

It is possible to view the recovery configurations and edit them as required by checking the **Modify recovery configuration for current operation** check box. But, the recovery settings changes are applicable only for the current operation and the changes will not be updated to the global recovery settings.

Compatibility Matrix for Disaster Recovery Operations

The following compatibility matrix lists the DR operations that are supported when a cluster at HX Data Platform version 3.5(x) is paired with a cluster at HX Data Platform version 3.5(x) or 3.0(1x).

Feature	3.5(x) Paired With 3.5(x)	3.5(x) Paired With 3.0(1x)
Replication	✓	✓
Cluster Pairing	✓	✓
Datastore Mapping	✓	✓
Protection	✓	✓
Planned Migration (Single click using HX Connect)	✓	X
Planned Migration (Multi-step stcli or WebCLI and HX Connect for Recovery)	✓	✓
Test Recover using HX Connect	✓	✓
Recover using HX Connect	✓	✓
Re-protect using HX Connect	✓	X
Re-protect (Multi-step stcli or WebCLI)	✓	✓

Testing Virtual Machine Recovery

Testing VM recovery gives you the ability to test recovery without breaking replication. It can bring up your VM workload on the target to verify the contents of the VM.



Note

- Testing recovery does not disrupt the running clusters. The intent is to verify, in the event of an actual disaster, that the VMs are recoverable.
- Using the HX Connect user interface, to test VM recovery, you can run a maximum of 10 tasks in a sequence without waiting for the previously submitted task to complete.

Before you begin

Before you begin the test VM recovery process, ensure the following:

- The target cluster is up and in good health.
- The protected VMs completed a recent replication to the target cluster. These replicated VMs are stored as DP snapshots on the target clusters.



Important

Only one copy of the test recovered VM can be made at any point. If you need to have another test recovered VM, please delete the previously created test recovered VM.

To the test VM recovery process perform the following steps:

- Step 1** Log in to HX Connect on the target cluster as administrator.
- Step 2** Navigate to **Replication > Remote VMs Tab > *protected_ym***.
- Step 3** To test the recovery process, click the **Test Recovery** button.

Note When configuring recovery settings, the following fields are auto-populated.

UI Element	Essential Information
Resource Pool drop-down list	Select a location for the test VM to be stored.
Folders drop-down list	Select a location for the test VM to be stored, for example: <ul style="list-style-type: none"> • Discovered Virtual Machine • HX Test Recovery
Power On/Off radio button	By default, the Power ON option is selected. The recovered VM is powered on or left off after it is created as per the selected option.
VM Name field	Enter a new name for the created test VM.

UI Element	Essential Information
Test Networks radio button	<p>Select which HX Storage Cluster network to use for transferring the data from the replication snapshot.</p> <p>Network options for example:</p> <ul style="list-style-type: none"> • Storage Controller Data Network • Storage Controller Management Network • Storage Controller Replication Network • VM Network
Map Networks radio button	<p>Select to create a map between the source and the target cluster networks.</p> <ul style="list-style-type: none"> • Source Network—Network name at the source side on which the VM is connected. • Target Network—Select target network from the drop-down list, where the VM has to be connected.

Step 4 Click **Recover VM**.

Step 5 For VMs that are part of a protection group, perform a test recovery on each VM in the group.

Step 6 Verify the contents of the recovered VM.

Recovering Virtual Machines

Recovering VMs is restoring a most recent replication snapshot from the target (recovery) cluster.

**Attention**

- You may configure the folder, network, or resource pool parameters to be used during recovery, test recovery and migrate operations. If the global recovery setting is not configured, you will need to explicitly map individual VMs at the time of recovery.
- Recover VM is not supported between different vSphere versions. If the Target is at a lower version vSphere environment and does not support the hardware version of a protected VM on the primary, VM test recovery and recovery may fail. Cisco recommends to test recover each protected VM to validate the support on the target site.

Upgrade the target environment to enable recovery of protected VMs.

- When running recovery on VMs, you may specify explicit network mapping when recovering the VMs to avoid unintentional network connections to recovered VMs.

You can skip specifying network mapping in the following cases:

- If the source VMs use vSphere Standard Switches and if all ESXi hosts on the recovery side have standard switch networks with the same name.
- If the source VMs use vSphere Distributed Switch (vDS) port groups and if the recovery site has identically named vDS port groups.
- If you want to specify network mapping, ensure that both the name and the type of the VM network matches between the source and the target.
- When running recovery on virtual machines that are individually protected, or that are in different protection groups, the maximum number of concurrent recovery operations on a cluster is 20.

Before you begin

Ensure the following:

- The target cluster is up and in good health.
- The protected VMs completed a recent replication to the target cluster. Replicated VMs are stored as DP snapshots on the target clusters.

On the target cluster, perform the following to conduct disaster recovery.

Step 1 Log in to HX Connect as administrator.

Step 2 Select **Replication** > > **Remote VMs tab** > > *protected_vm* and click **Recover**.

Step 3 To recover the VM and build a new VM on the local cluster, click the **Recover VM** button.

Note When you configure recovery settings, the following fields are auto-populated.

UI Element	Essential Information
Resource Pool drop-down list	Select a location for the new VM to be stored.
Folders drop-down list	Select a location for the new VM to be stored.

UI Element	Essential Information
Power On/Off radio button	By default Power ON option is selected. The recovered VM is powered on or left off after it is created as per the selected option.
Map Networks	<p>Select to create a map between the source and target cluster networks.</p> <ul style="list-style-type: none"> • Source Network—Network name at the source side on which the VM is connected. • Target Network—Select target network from the drop-down list, where the VM has to be connected. <p>Network options for example:</p> <ul style="list-style-type: none"> • Storage Controller Data Network • Storage Controller Management Network • Storage Controller Replication Network • VM Network

Step 4 Click **Recover VM**.

Step 5 Wait for the recovery to complete. View the recovered VM in the target vCenter.

Recovering Virtual Machines in Protection Groups

Step 1 Select a *protected-vm* and click **Recover**.

All VMs will be moved from the protection group and the selected VMs will be recovered. Recovered VMs show protection status as *Recovered* and the remaining (protection group) VMs show protection status as *Recovering*. The protection group will go in *Recovered* state and is not reusable. You can delete it from the primary site.

Note Clicking **Recover** for a VM in a group puts it in a **Recovered** state (actual recovery happened), while the rest of the VMs in the standalone list are in **Ready for Recovery** state.

The recovered VMs are displayed in the **Standalone Protected VMs** subpane.

Step 2 Recover the remaining virtual machines from the **Standalone Protected VMs** subpane, which were a part of the protection group. See [Recovering Virtual Machines, on page 222](#) for more details.

Planned Migration

Performing planned migration pauses the replication schedule, replicates the most recent copy, recovers on the target, switches the ownership from the source to the target, and resumes replication on the target that is now the new source.

**Attention**

- This process cannot be rolled back.
- For a single vCenter deployment, the **Migrate** workflow performed entirely through the HX Connect UI is not supported. To perform a planned migration:
 1. Using the WebCLI, run the following command to prepare for failover on the source:

```
# stcli dp vm prepareFailover -vmid <VMID>
```

Result: The task ID is returned.

2. Log in to vSphere Web Client Navigator of the primary site and remove the VM from the primary site to unregister the VM.
Right-click on the virtual machine and select **All vCenter Actions > Remove from Inventory**.
3. Log in to HX Connect of the secondary site. Select **Replication > Remote VMs Tab > protected_vm**. Click **Migrate**.
4. After the Migrate task has completed successfully, log in to vSphere Web Client of the secondary site and manually register the VM.
 - a. Log in to vSphere Web Client Navigator. Select **Configuration > Storage**.
 - b. Right-click on the appropriate datastore and click **Browse Datastore**.
Navigate to the *virtualmachine name.vmx* file, right-click on the file and click **Add to Inventory**. Follow the wizard to manually register the VM.

- Step 1** Log in to HX connect of the target cluster. The target cluster is where the replicated DP snapshots were copied to.
- Step 2** On the target cluster, select **Replication > Remote VMs Tab > protected_vm**.
- Step 3** Click **Migrate**.

Note All the fields that are listed here are optional.

UI Element	Essential Information
Resource Pool drop-down list	Select a location for the new VM to be stored.
Folders drop-down list	Select a location for the new VM to be stored.
Power On/Off radio button	By default Power ON option is selected. The recovered VM is powered on or left off after it is created as per the selected option.

UI Element	Essential Information
Map Networks	<p>Select to create a map between the source and target cluster networks.</p> <ul style="list-style-type: none"> • Source Network—Network name at the source side on which the VM is connected. • Target Network—Select target network from the drop-down list, where the VM has to be connected. <p>Network options for example:</p> <ul style="list-style-type: none"> • Storage Controller Data Network • Storage Controller Management Network • Storage Controller Replication Network • VM Network

Step 4 Monitor the progress on the **Activity** page.

Migrating Virtual Machines in Protection Groups

Using the HX Connect user interface, to migrate VMs, you can run a maximum of 4 tasks in a sequence without waiting for the previously submitted task to complete.

Step 1 Select a *protected-vm* and click **Migrate**.

All the VMs are now moved out from the protection group and are displayed in the **Standalone Protected VMs** subpane. Only the selected VM is recovered.

Step 2 Migrate the remaining virtual machines from the **Standalone Protected VMs** subpane, which were a part of the protection group. See [Planned Migration, on page 224](#) for more details.

Disaster Recovery and Re-protect

Performing disaster recovery recovers the VM on the target, switches the ownership from the source to the target, and resumes replication on the target that is now the new source. Disaster recovery is typically done when disaster occurs, and when you want to reverse the direction of protection.

**Attention**

- This process cannot be rolled back.
- For a single vCenter deployment, the Disaster Recovery workflow performed entirely through the HX Connect UI is not supported. To perform Disaster Recovery and Re-protect:
 1. Log in to vSphere Web Client Navigator of the primary site and remove the VM from the primary site to unregister the VM.
Right-click on the virtual machine and select **All vCenter Actions > Remove from Inventory**.
 2. Log in to HX Connect of the secondary site. Select **Replication > Remote VMs Tab > protected_vm**. Click **Recover**.
 3. When the primary site comes back up, log in to HX Connect of the secondary site. Select **Replication > Remote VMs Tab > protected_vm**. Click **Re-protect**.
 4. After Re-protect has completed successfully, log in to vSphere Web Client of the secondary site and manually register the VM.
 - a. Log in to vSphere Web Client Navigator. Select **Configuration > Storage**.
 - b. Right-click on the appropriate datastore and click **Browse Datastore**.
Navigate to the *virtualmachine name.vmx* file, right-click on the file and click **Add to Inventory**. Follow the wizard to manually register the VM.
- Using the HX Connect user interface, you can run a maximum of 5 re-protect tasks in a sequence without waiting for the previously submitted task to complete.

Step 1 Log in to HX connect of the source and the target. The target cluster is where the replicated DP snapshots were copied to. The source cluster is the cluster where the VMs reside.

Step 2 Select a VM from the remote VM list. Execute the Recover VM workflow on this cluster workflow.

Note If both the target and source clusters are on the same vCenter, then unregister the VM on the source cluster. This ensures that vCenter no longer has a record of the VM and it stops managing the VM, but it retains the data for the VM.

Step 3 Select **Replication > > Remote VMs tab > > protected_vm** and click **Recover**.

Step 4 To recover on the target VM and build a new VM on the local cluster, click the **Recover VM** button.

Complete the following fields in the **Recover VM on this cluster** dialog box.

UI Element	Essential Information
Resource Pool drop-down list	Select a location for the new VM to be stored.
Folders drop-down list	Select a location for the new VM to be stored.
Power On/Off radio button	By default the Power ON option is selected. The recovered VM is powered on or left off after it is created as per the selected option.

UI Element	Essential Information
Map Networks	<p>Select to create a map between the source and target cluster networks.</p> <ul style="list-style-type: none"> • Source Network—Network name at the source side on which the VM is connected. • Target Network—Select target network from the drop-down list, where the VM has to be connected. <p>Network options for example:</p> <ul style="list-style-type: none"> • Storage Controller Data Network • Storage Controller Management Network • Storage Controller Replication Network • VM Network

Step 5 Click **Recover VM**.

Step 6 On the target cluster, select **Replication > Remote VMs Tab > *protected_vm***.

Step 7 Click **Re-protect**.

- Attention**
- If both the target cluster and source cluster are on the same vCenter, manually register the VM on the source cluster.
 - When the Re-protect task fails and in the HX Connect UI the **Re-protect** tab is not available, execute *stcli reverseprotect* to complete the Re-protect operation.

Protection status of the VM shows as **Protected**.

Step 8 After the original primary comes up, to migrate back to the primary do the following:

- On the target cluster, select **Replication > Remote VMs Tab > *protected_vm***.
- Click **Migrate** to unregister the target VM and transfer the VM ownership to the original primary. Protection status of the VM shows as **Protected**.

Protecting Virtual Machines After Disaster

In the event of a disaster, you may lose the source site altogether. After the recovery is performed, you may want to protect the recovered VMs to a newer cluster.

Step 1 Recover the Virtual Machines. Perform standalone recovery (Recovering VMs) or group recovery (Recovering VMs in protection groups). See [Recovering Virtual Machines, on page 222](#) for more details.

Step 2 To clear the existing pairing, run the following command in the HX Connect WebCLI:

```
stcli dp peer forget --all
```

Now the cluster is no longer paired to the original source.

Step 3 Unprotect all the local and remote VMs. See [Unprotecting Virtual Machines, on page 217](#) for more details.

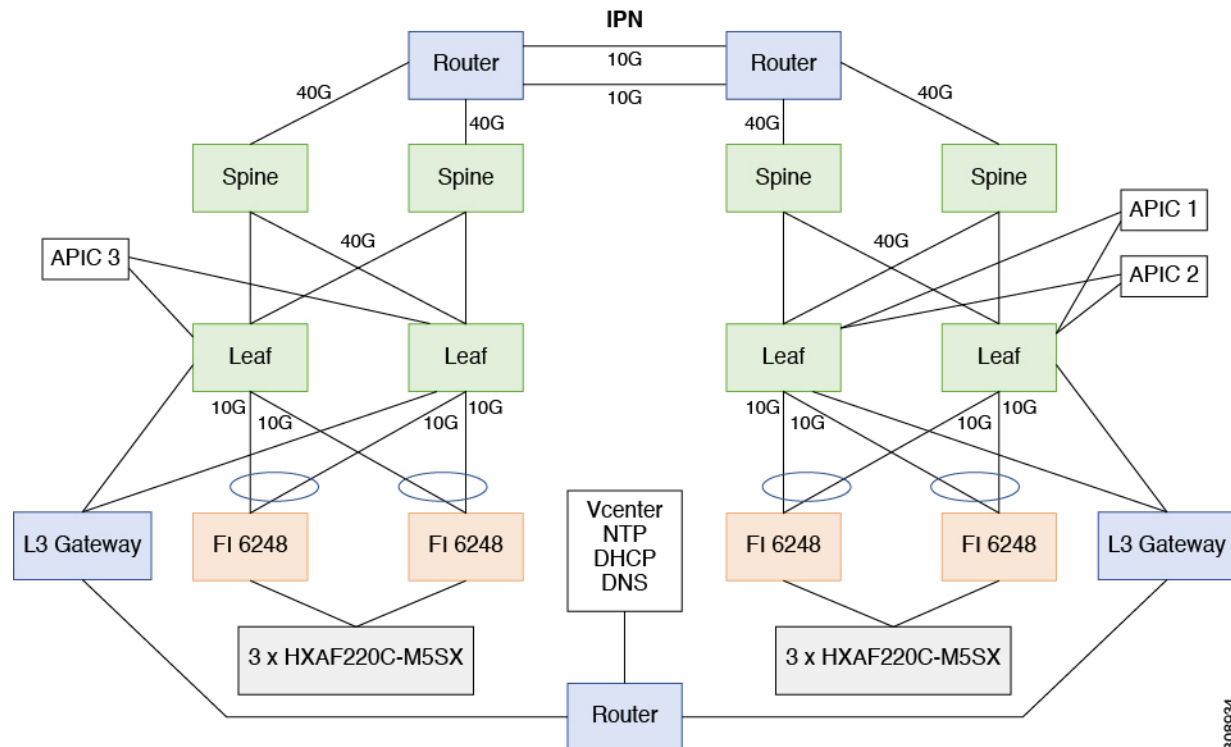
- Step 4** Pair to the new cluster. See the [Creating a Replication Pair, on page 205](#) section for more details.
- Step 5** Protect the virtual machines.

Replication Maintenance Overview

Replication, when configured, runs in the background according to defined schedules. Replication maintenance tasks include:

- **Testing recovery**—Testing if the recovery methods are working. See [Testing Virtual Machine Recovery, on page 221](#) for more details.
- **Pausing replication**—When preparing to upgrade the HX cluster and replication is configured, replication activity must be paused.
Use the `stcli dp schedule pause` command.
- **Resuming replication**—After HX cluster maintenance activities are complete, resume the replication schedule.
Use the `stcli dp schedule resume` command.
- **Migration**—The option to shift VMs from one source cluster to the replication paired target cluster, making the target cluster the new source cluster for the migrated VMs.

The following image illustrates which configuration is used for Disaster Recovery on HyperFlex if you are deploying in an ACI setup on a broad scale:



Pausing Replication

Before performing a storfs or platform upgrade, if replication is configured replication activity must be paused.

-
- | | |
|---------------|--|
| Step 1 | Log in to a Storage Controller VM. |
| Step 2 | From the command line, run the <code>stcli dp schedule pause</code> command. |
| Step 3 | Perform the upgrade task. |
| Step 4 | Resume the replication schedule. |
-

Resuming Replication

After successfully upgrading the HX Storage Cluster which had replication configured, do the following to resume the replication schedule.

Before you begin

Ensure that HX cluster replication is paused and that any maintenance or upgrade tasks are complete.

-
- | | |
|---------------|---|
| Step 1 | Log in to a Storage Controller VM. |
| Step 2 | From the command line, run the <code>stcli dp schedule resume</code> command. |
-

The previously configured replication schedule for all the protected virtual machines begins.



CHAPTER 17

Managing Storage Replication Adapter

- [Storage Replication Adapter Overview, on page 231](#)
- [Configuring Storage Replication Adapter, on page 232](#)
- [Executing SRA Workflows, on page 238](#)
- [Troubleshooting SRA, on page 240](#)

Storage Replication Adapter Overview

Storage Replication Adapter (SRA) for VMware vCenter Site Recovery Manager (SRM) is a storage vendor-specific or plugin for VMware vCenter server. The adapter enables communication between SRM and a storage controller at the Storage Virtual Machine (SVM) level as well as at the cluster level configuration. The adapter interacts with Cisco HyperFlex to discover replicated datastores.

SRA provides the replicated datastores the ability to perform test recovery and recovery that includes planned migration and disaster recovery.



Note This chapter covers the instructions to be executed in SRM Release 8.2. If you are using SRM Release 8.1, 6.5 or 6.1, refer the following guides:

- For SRM 8.1 installation—<https://docs.vmware.com/en/Site-Recovery-Manager/8.1/srm-install-config-8-1.pdf>
- For SRM 8.1 configuration—<https://docs.vmware.com/en/Site-Recovery-Manager/8.1/srm-admin-8-1.pdf>
- For SRM 6.5 installation—<https://docs.vmware.com/en/Site-Recovery-Manager/6.5/srm-admin-6-5.pdf>
- For SRM 6.1 installation—<https://docs.vmware.com/en/Site-Recovery-Manager/6.1/srm-admin-6-1.pdf>

To get started select your installation type:

- [Installing a Storage Replication Adapter for Windows, on page 232](#)
- [Installing a Photon OS Storage Replication Adapter, on page 233](#)

Configuring Storage Replication Adapter

System Requirements

Minimum requirements for SRM are:

Element	Minimum Supported Requirement
Processor	At least two 2.0 GHz or higher Intel or AMD x86 processors.
Memory	4 GB
Disk Storage	5 GB
Networking	1 GB is recommended for communication between SRM protected and recovery sites

For information on supported platforms and databases, see the VMware Compatibility Matrices for your release:

[Compatibility Matrices for VMware Site Recovery Manager 8.1](#)

[Compatibility Matrices for VMware Site Recovery Manager 8.2](#)

Windows Site Recovery Manager

Installing a Site Recovery Manager for Windows

You must install an SRM server instance at the protected site and also at the recovery site. For more information on the installation procedure, refer the *Installing Site Recovery Manager* chapter in the [Site Recovery Manager Installation and Configuration](#) guide.

After you install SRM, it remains in evaluation mode until you install an SRM license key. After the evaluation license expires, existing protection groups remain protected and you can recover them, but you cannot create new protection groups or add virtual machines to an existing protection group until you obtain and assign a valid SRM license key. You must obtain and assign SRM license keys after installing SRM. For more information on the licensing, refer the *Site Recovery Manager Licensing* section in the [Site Recovery Manager Installation and Configuration](#) guide.

Installing a Storage Replication Adapter for Windows

Before you begin

- Ensure that SRM and JDK 8 or above version are installed on Windows machines at the protected and recovery sites.
- Ensure that, you have obtained and assigned a valid SRM license key. For more information on the licensing, refer the *Site Recovery Manager Licensing* section in the [Site Recovery Manager Installation and Configuration](#) guide.

-
- Step 1** Download SRA from the VMware site.
- In the <https://my.vmware.com/web/vmware/downloads> page, locate **VMware Site Recovery Manager** and click **Download Product**. Click **Drivers & Tools**, expand **Storage Replication Adapters**, and click **Go to Downloads**.
- Step 2** Copy the Windows installer of SRA to SRM Windows machines at the protected and recovery sites.
- Step 3** Double-click the installer.
- Step 4** Click **Next** on the welcome page of the installer.
- Step 5** Accept the EULA and click **Next**.
- Step 6** Click **Finish**.
-

The SRA is installed within the SRM program folder:

```
C:\Program Files\VMware\VMware vCenter Site Recovery Manager\storage\sra
```

Site Recovery Manager Virtual Appliance

Deploy the Site Recovery Manager Virtual Appliance

The Site Recovery Manager Virtual Appliance is a preconfigured virtual machine that is optimized for running Site Recovery Manager and its associated services. You deploy the appliance on an ESXi host in your vSphere environment. refer the Deploy the Site Recovery Manager Virtual Appliance chapter in the <https://docs.vmware.com/en/Site-Recovery-Manager/8.2/srm-install-config-8-2.pdf>

After you install SRM, it remains in evaluation mode until you install an SRM license key. After the evaluation license expires, existing protection groups remain protected and you can recover them, but you cannot create new protection groups or add virtual machines to an existing protection group until you obtain and assign a valid SRM license key. You must obtain and assign SRM license keys after installing SRM. For more information on the licensing, refer the Site Recovery Manager Licensing section in <https://docs.vmware.com/en/Site-Recovery-Manager/8.2/srm-install-config-8-2.pdf>

Installing a Photon OS Storage Replication Adapter

-
- Step 1** Download SRA from the VMware site.
- In the <https://my.vmware.com/web/vmware/downloads> page, locate **VMware Site Recovery Manager** and click **Download Product**.
- Step 2** Download Drivers and Tools from the VMware site.
- Click **Drivers & Tools**, expand **Storage Replication Adapters**, and click **Go to Downloads**.
- Step 3** Log in to the Site Recovery Manager Appliance Management Interface as admin.
- Click on **Storage Replication Adapter** tab.
- Step 4** Click **NEW ADAPTOR** button and upload the SRA.
-

Setting up SRM Environment

This section covers the following procedures for both the Windows and Photon SRA:

- [Pairing Sites, on page 234](#)
- [Configuring Mapping, on page 234](#)
- [Refreshing Storage Replication Adapter, on page 235](#)
- [Adding an Array Manager Pair, on page 235](#)
- [Mapping Datastores on Cisco HyperFlex, on page 236](#)
- [Discovering Devices, on page 236](#)
- [Creating a Protection Group, on page 237](#)
- [Creating a Recovery Plan, on page 237](#)

Pairing Sites

Before you can use SRM, you must connect the SRM server instances on the protected and recovery sites. This is known as site pairing.

-
- Step 1** Log in to your protected site vCenter server using vSphere Web Client.
- Step 2** In the home page, click **Site Recovery** plug-in and select sites.
- Step 3** Enter the platform service controller address of the SRM remote site (recovery site). Leave the port number as default if you don't have any custom port for platform services controller (PSC). Click **Next**.
- The system displays the vCenter server instance in which SRM server is registered on the remote site.
- Step 4** Enter the vCenter server single sign-on (SSO) credentials and click **Finish**.
- The connection process proceeds automatically after authentication with the correct privileges established.
- For more information, refer the *Site Recovery Manager Licensing* section in the [Site Recovery Manager Installation and Configuration Guide](#).
-

When the site pairing is established between protected and recovery site, you will be able to see the remote site under **Sites** in the SRM interface. It displays the information about the primary site and paired site.

Configuring Mapping

Mappings allow you to specify how SRM maps virtual machine resources on the protected site to resources on the recovery site. During recovery, when virtual machines start on the recovery site, the virtual machines use the resources on the recovery site that you specify in the mappings.

-
- Step 1** Log in to the protected site vCenter using vSphere Web Client.
- Step 2** Click **SRM** plugin. Select the protected site and click **Configure Inventory Mappings**.
- Step 3** Click the following options and do the respective mappings:

Option	Action
Create resource mappings	<p>Map resource pools, standalone hosts, vApps, or clusters on the protected site to resource pools, standalone hosts, vApps, or clusters on the recovery site. You can map any type of resource on one site to any type of resource on the other site.</p> <p>Note You cannot map individual hosts that are part of clusters to other resource objects.</p>
Create folder mappings	Map datacenters or virtual machine folders on the protected site to datacenters or virtual machine folders on the recovery site.
Create network mappings	Map networks on the protected site to networks on the recovery site.
Configure placeholder datastores	<p>Select the datastore from the list to be used as placeholder datastore for the production cluster.</p> <p>Important You have to configure placeholder datastores on recovery site as well.</p>

For more information on inventory mappings, refer [VMware Site Recovery Manager Administration Guide](#)

Refreshing Storage Replication Adapter

- Step 1** Log in to vSphere Web Client.
- Step 2** In the home page, click the **Site Recovery** plugin.
- Step 3** Click **Sites** and select the protected site.
- Step 4** Click the **Monitor** tab.
- Step 5** Click **Refresh**.
- The Refresh icon is available next to **Add Array Manager**.

Execute the same procedure for the recovery site to refresh SRA in the recovery site.

What to do next

After refresh, the status of the site is set to **OK** with a green tick. This action refreshes SRA information, allowing SRM to discover the SRAs.

Adding an Array Manager Pair

- Step 1** Log in to vSphere Web Client.

- Step 2** Click **SRM** plugin. Select the protected site and click **Array Based Replication**.
- Step 3** Click **Add (+)**.
- The **Add Array Manager** wizard appears.
- Step 4** In the **Add Array Manager** wizard, do the following:
- Click **Add a pair of array managers** to add array managers for both the protected and recovery sites at one shot. You will be prompted to add two array managers, one for each site. Click **Next**.
 - Select the pair of sites (the protected site and recovery site). Verify the displayed site information and click **Next**.
The system automatically detects the type of SRAs installed on both the sites and statuses of SRAs.
 - Click **Next**.
 - Enter the display name, IP address, and credentials for the protected site array manager. Click **Next**.
 - Enter the display name, IP address, and credentials for the recovery site array manager. Click **Next**.
 - Check to enable the array pairs and click **Next**.
 - Review the specified information for array manager and click **Finish**.

Mapping Datastores on Cisco HyperFlex

Before you begin

Ensure that a replication pair is created from protected cluster to recovery cluster.

- Step 1** Log in to HX Connect as an administrator.
- Step 2** Select **Replication > Replication Pairs**.
- Step 3** Select the replication pair of protected and recovery sites, and click **Edit**.
- Step 4** Click the **Other DRO Protection** tab and do the following:
- The **Local Datastore** column displays a list of the unpaired configured datastores on the local HX cluster. Map one local datastore to one remote datastore.
 - From the **Remote Datastore** pull-down menu, choose a datastore that need to be paired with the local datastore.
 - From the **Direction** pull-down menu, choose **Incoming** or **Outgoing** as the direction of VM movement for the mapped datastores.
 - From the **Protection Schedule** pull-down menu, choose the schedule for protecting all the VMs in the datastore.
 - Click **Map Datastores**.
- Step 5** Click **Map Datastores**.

Discovering Devices

- Step 1** Log in to vSphere Web Client.
- Step 2** In the home page, click **Site Recovery** plugin.
- Step 3** Expand **Inventories** and click **Array Based Replication**.
- Step 4** Select the protected array.
- Step 5** Click **Manage**.

Step 6 Click blue refresh icon.

This actions runs the discover devices on the selected array pair.

Creating a Protection Group

- Step 1** Log in to vSphere Web Client.
 - Step 2** In the home page, click **Site Recovery** plugin.
 - Step 3** Expand **Inventories** and click **Protection Groups**.
 - Step 4** Click **Objects**.
 - Step 5** Click **Create Protection Group**.
 - Step 6** Enter the name and location for the protection group.
 - Step 7** Choose incoming or outgoing as the direction of protection.
 - Step 8** Choose **Datastore groups** from the **Protected Group type** drop-down list. Click **Next**.
 - Step 9** Select the datastore groups from the list to be grouped to use for the protection group.
All the virtual machines stored on these datastore groups will be recovered together as part of the protection group.
 - Step 10** Click **Next**.
 - Step 11** Review all the information, which you have selected to create the SRM protection group and click **Finish**.
-

Creating a Recovery Plan

- Step 1** In the vSphere Web Client home page, click **Site Recovery** plugin.
 - Step 2** Expand **Inventories** and click **Recovery Plans**.
 - Step 3** Click **Objects**.
 - Step 4** Click **Create Recovery Plan**.
 - Step 5** Enter the name and description for the recovery plan and select the location for the recovery plan. Click **Next**.
 - Step 6** Select the recovery site in which the VMs will be recovered as part of the recovery plan. Click **Next**.
 - Step 7** Select the list of protection groups which will be included as part of the recovery plan. Click **Next**.
 - Step 8** Specify the test networks to use while running tests of the recovery plan.
You can select the test networks if you have created for the test purposes or you can leave the default auto created test network called isolated network.
 - Step 9** Click **Next**.
 - Step 10** Review all the information, which you have selected to create the recovery plan and click **Finish**.
-

Executing SRA Workflows

After SRA configuration, SRM works with SRA to discover arrays and replicated and exported datastores, and to failover or test failover datastores. The workflows described here apply to both the Photon OS based SRA and the Windows based SRA.

SRA enables SRM to execute the following workflows:

- Discovery of replicated storage
- Non-disruptive failover test recovery using a writable copy of replicated data
- Emergency or planned failover recovery
- Reverse replication after failover as part of failback

**Note**

After failover, the source datastore is unmounted on HX data platform, due to which, the datastore to VM mapping returns empty list. So, we cannot unpair the datastores. To unpair the datastore, you have to bring the VMs on the source datastore to the Active state by invoking the Reprotect operation.

This section covers the following topics:

- [Testing a Recovery Plan, on page 238](#)
- [Testing the Recovery Plan Cleanup Task, on page 239](#)
- [Executing a Recovery, on page 239](#)
- [Executing a Reprotect Task, on page 240](#)

Testing a Recovery Plan

Step 1 In the vSphere Web Client home page, click **Site Recovery** plugin.

Step 2 Expand **Inventories** and click **Recovery Plans**.

Step 3 Select the recovery plan and click **Monitor**.

Step 4 Click the green play button to run the test recovery.

Caution Be cautious in selecting the test recovery option.

The recovery process runs in the test mode to recover the virtual machines in a test environment on the recovery site. You have the storage option to choose whether to replicate the recent changes to the recovery site or not. This process may take several minutes to complete the replication of recent changes to the recovery site. According to your requirement, you can choose to run the test recovery with recent changes or not.

Step 5 Click **Next**.

Step 6 Review the settings and click **Finish** to initiate the SRM test recovery.

The series of recovery steps involved as part of this recovery plans are displayed with its progress percentage.

Note Once the plan is executed, the datastore on the recovery site will be unmounted and it will stay unmounted. During this time, you can ignore alerts on the HyperFlex recovery cluster as this is expected behavior.

On completion of the test recovery, the status of the recovery plan is changed to **Test Complete** and status of each recovery steps is changed to **Success**. Also, the virtual machines that are part of the protection group as part of this recovery plan are powered on in both protected and recovery site without causing any interruption to the production virtual machines.

After successful execution of test recovery, you can run a recovery plan cleanup to return virtual machines to their initial state at which it was before running the recovery plan test. The recovery plan status is reset to **Ready** state. You can also run a forced cleanup when you experience errors during execution of a recovery plan. For more information, see [Testing the Recovery Plan Cleanup Task, on page 239](#).

Testing the Recovery Plan Cleanup Task

You can run a recovery plan cleanup task to return virtual machines to their initial state at which it was before running the recovery plan test. Also, you can run a forced cleanup when you experience errors during execution of a recovery plan.

Before you begin

A recovery plan must be run.

-
- Step 1** In the vSphere Web Client home page, click **Site Recovery** plugin.
 - Step 2** Expand **Inventories** and click **Recovery Plans**.
 - Step 3** Select the recovery plan and click **Monitor**.
 - Step 4** Click the brush icon to cleanup the test recovery.
Running a cleanup operation on this plan will remove the test environment and reset the plan to ready state.
 - Step 5** Click **Next** after receiving the cleanup confirmation message.
 - Step 6** Review the clean information and click **Finish** to start the recovery test cleanup task.
-

You can monitor the progress of the test recovery cleanup process. The recovery plan cleanup task returns virtual machines to their initial state as it was before running the recovery plan test. After cleanup, the virtual machines recovered as part of the test recovery in the recovered site are in the powered off state. The recovery plan status is reset to **Ready** state. The recovery plan is ready to test again.

Executing a Recovery

-
- Step 1** In the vSphere Web Client home page, click **Site Recovery** plugin.
 - Step 2** Expand **Inventories** and click **Recovery Plans**.
 - Step 3** Select the recovery plan and click **Monitor**.
 - Step 4** Click the red play button to run the recovery plan.
 - Step 5** Choose **Planned Migration** or **Disaster recovery** as per your requirement. Click **Next**.

Step 6 Review the settings set for initiating the SRM recovery and click **Finish**.

The series of recovery steps involved in the recovery plan, are displayed with the percentage of progress.

On completion of recovery process, the status of the recovery plan is set to **Recovery Complete** and the status of each recovery steps is set to **Success**. The virtual machines that are part of the protection group of the recovery plan are powered on in both the protected and recovery sites, without interrupting the production virtual machines.

Executing a Reprotect Task

After completing the recovery, you can reprotect the recovery plan.

-
- Step 1** In the vSphere Web Client home page, click **Site Recovery** plugin.
 - Step 2** Expand **Inventories** and click **Recovery Plans**.
 - Step 3** Select the recovery plan and click **Monitor**.
 - Step 4** Click **Reprotect**.
 - Step 5** Review the settings set for initiating the SRM reprotect and click **Finish**.
-

The progress of each reprotect steps is displayed. On completion of reprotection, the status of the recovery plan is set to **Reprotect Successful** and the status of each reprotect steps is set to **Success**. The virtual machines are protected, and the test recovery and recovery options are enabled.



Note

When you execute the Reprotect task, the datastore replication direction does not change to reflect reverse replication. We recommend you to view the VM direction at **Replication > Replication Pairs > Other DRO Protection** to confirm replication direction.

Troubleshooting SRA

Troubleshooting: Using Windows SRM Log Files

To identify the cause of any problems that you may encounter during the day-to-day running of SRM, you have to collect and review SRM log files. To get VMware support, you have to send the log files to them.

The VMware SRM logs from SRM server are available at: `C:\ProgramData\VMware\VMware vCenter Site Recovery Manager\Logs`

SRA specific logs are available under the SRA folder at: `C:\ProgramData\VMware\VMware vCenter Site Recovery Manager\Logs\SRAs\`

To download SRM log files from the SRM interface, follow this procedure:

1. In the vSphere Web Client home page, click **Site Recovery > Sites**.

2. Select a site for which the log file has to be collected.
3. From the **Actions** menu, select **Export SRM Log**.
Alternatively, right-click the site and select **Export SRM Log**.
4. In the **Export SRM Log** wizard, click **Generate Log**.
5. Click **Download Log**.

Troubleshooting: Using the SRM Appliance Log Files

To identify the cause of any problems that you may encounter during the day-to-day running of SRM, you must collect and review SRM log files. To get VMware support, you have to send the log files to them. You have to SSH to SRM Appliance to get the log file with root login

The VMware SRM logs from the SRM Appliance are available at: `/var/log/vmware/srm`

SRA specific logs are available under the SRA folder at: `/var/log/vmware/srm/SRAs`

To download SRM log files from the SRM interface, follow this procedure:

1. In the vSphere Web Client home page, click **Site Recovery > Sites**.
2. Select a site for which the log file has to be collected.
3. From the **Actions** menu, select **Export SRM Log**.
Alternatively, right-click the site and select **Export SRM Log**.
4. In the **Export SRM Log** wizard, click **Generate Log**.
5. Click **Download Log**.

Troubleshooting: Location of the VMware Virtual Center

If you delete a VM from VMware vCenter that is located on a **Other DRO** datastore pair, a recovery plan for this datastore pair fails during recovery. To avoid this failure, first unprotect the VM using the following command on one of the controller VM:

```
stcli dp vm delete --vmid <VM_ID>
```

Troubleshooting: Name of the VM

If you rename a VM from the vCenter, HyperFlex recovers at the previous name folder but registers the VM with the new name on the recovery side. Following are some of the limitations to this situation:

- VMware allows a VMDK located at any location to be attached to a VM. In such cases, Hyperflex recovers the VM inside the VM folder and not at a location mapped to the original location. Also, recovery can fail if the VMDK is explicitly referenced in the `virtualmachine name.vmx` file by its path. The data is recovered accurately but there could be problems with registering the VM to vCenter. Correct this error by updating the `virtualmachine name.vmx` file name with the new path.
- If a VM is renamed and a VMDK is added subsequently, the new VMDK is created at `[sourceDs] newVm/newVm.vmdk`. Hyperflex recovers this VMDK with the earlier name. In such cases, recovery

can fail if the VMDK is explicitly referenced in the `virtualmachine name.vmx` file by its path. The data is recovered accurately but there could be problems with registering the VM to the Virtual Center. Correct this error by updating the `virtualmachine name.vmx` file name with the new path.



CHAPTER 18

Managing Users

- [Managing Cisco HyperFlex Users Overview, on page 243](#)
- [Creating Cisco HX Data Platform RBAC Users, on page 245](#)
- [Assigning Users Privileges, on page 246](#)

Managing Cisco HyperFlex Users Overview

The user types allowed to perform actions on or view content in the HX Data Platform, include:

- **admin**—A predefined user included with Cisco HX Data Platform. The password is set during HX Cluster creation. Same password is applied to `root`. This user has read and modify permissions.
- **root**—A predefined user included with Cisco HX Data Platform. The password is set during HX Cluster creation. Same password is applied to `admin`. This user has read and modify permissions.
- **administrator**—A created Cisco HX Data Platform user. This user is created through vCenter and assigned the RBAC role, `administrator`. This user has read and modify permissions. The password is set during user creation.
- **read-only**—A created Cisco HX Data Platform user. This user is created through vCenter and assigned the RBAC role, `read-only`. This user only has read permissions. The password is set during user creation.

HX Interface	admin	root	hx_admin	hx_readonly
HX Data Platform Installer	Required	Optional	Not valid	Not valid
HX Connect	Can perform most HX tasks. local/ prefix required for login. Example: local/admin	Not valid	Can perform most HX tasks. A preferred user.	Can only view monitoring information. Cannot perform HX tasks. A preferred user.

HX Interface	admin	root	hx_admin	hx_readonly
Storage Controller VM with <code>stcli</code> command line	Can perform most HX tasks.	Can perform most HX tasks.	vc- prefix required for login. Example: <code>vc-hx_admin</code>	Cannot perform HX tasks. vc- prefix required for login. Example: <code>vc-hx_readonly</code>
HX Data Platform Plug-in through vCenter	Can perform most HX tasks.	Can perform most HX tasks.	Can perform most HX tasks. A vCenter SSO user.	Can only view vCenter information. Cannot view HX Data Platform Plug-in. A vCenter SSO user.
HX REST API	Can perform most HX tasks. <code>local/ prefix</code> required for login. Example: <code>local/admin</code>	Can perform most HX tasks. <code>local/ prefix</code> required for login. Example: <code>local/root</code>	Not valid	Not valid

User Management Terms

- **Authentication**—For login credentials. These processes verify user credentials for a named user, usually based on a username and password. Authentication generally verifies user credentials and associates a session with the authenticated user.
- **Authorization**—For access permissions. These processes allow a user/client application to perform some action, such as create, read, update, or delete a managed entity or execute a program, based on the user's identity. Authorization defines what an authenticated user is allowed to do on the server.
- **Accounting**—For tracking user actions. These processes perform record-keeping and track user activities including login sessions and command executions. The information is stored in logs. These logs are included in the support bundle that can be generated through Cisco HX Connect or other Cisco HX Data Platform interface.
- **Identity**—Individuals are provisioned with identities that are assigned roles with granted permissions.
- **Permission**—Settings given to roles to use the Resource. It is the link between roles, resource and the function exposed by the resource. For example, Datastore is a resource and a modifying role is granted permission to mount the datastore, while a read only role can only view that the datastore exists.
- **Privilege**—The link between Identity and the application. It is used in the context of specific interaction with the application. Examples: Power On a Virtual Machine, Create a Datastore, or Rename a datastore.
- **Resource**—The entire Cisco HX Platform, whose functionality and management controls are exposed over HTTP using GET, POST, PUT, DELETE, HEAD and other HTTP verbs. Datastores, Disks, Controller Nodes, Cluster Attributes, are all resources that are exposed to client applications using REST API.

- **Role**—Defines an authority level. An application function may be performed by one or more roles. Examples: Administrator, Virtual Machine Administrator, or Resource Pool Administrator. Role is assigned to a given Identity.

Audit Logs for AAA Accounting

To support AAA accounting, Cisco HX Data Platform implements audit logs of user activity. These logs are included in the generated support bundle.

See the [Cisco HyperFlex Systems Troubleshooting Guide](#) for information on generating the support bundles through HX Data Platform interfaces, including Cisco HX Connect.

- **stMgrAudit.log**—Contains audit records of `stcli` activity.

Sample entry. Note the keyword, `Audit`.

```
2017-03-27-22:10:02.528 [pool-1-thread-1] INFO Audit - 2017-03-27-03.10.02 127.0.0.1
--> 127.0.0.1 POST /stmgr 200 : root 27ms
```

This file contains other information as well. To filter for audit events, use a script to filter for the word, `Audit`.

- **audit.log**—Contains audit records for REST API activity.

Sample entry. Note the user name, `administrator@vsphere.local`

```
2017-03-29-01:47:28.779 - 127.0.0.1 -> 127.0.0.1 - GET /rest/clusters 200;
administrator@vsphere.local 454ms
```

Creating Cisco HX Data Platform RBAC Users

Cisco HX Data Platform supports two users: Administrator and Read Only. New users are created for the HX Data Platform through the VMware vCenter interface.

Before you begin

Creating users requires Administrator privileges.

-
- Step 1** Log in to vSphere Web Client as a vCenter administrator.
 - Step 2** From **Navigator Home, Administration > Users and Groups > Users**.
 - Step 3** Click **Add (+)** icon to add a user. Then complete the **New User** information and click **OK**.

Specify a user name and password for the new user.

For passwords, do not use escape character (\), dollar sign (\$), question mark (?), equal sign (=). In user names, the only special characters allowed are underscore (_), dash (-), dot (.). For more information on user name and password requirements, see [HX Data Platform Names, Passwords, and Characters, on page 21](#).

What to do next

Add the user to an RBAC role group. See [Assigning Users Privileges, on page 246](#).

Assigning Users Privileges

Privileges are assigned to users through the RBAC roles in vCenter. To assign privileges, add users to either the Administrator or Read-only group.

Before you begin

Create the user.

-
- Step 1** From the Cisco vSphere Web Client, select **Navigator Home > Administration > Global Permissions > Manage**.
- Step 2** Click **Add (+)** icon to assign roles.
- Step 3** Select an **Assigned Role**.
- In the **Global Permission Root - Add Permission** dialog box, select from the **Assigned Role** drop down menu. Choose one:
- **Administrator**
 - **Read only**
- Step 4** In the **Users and Groups** area, click **Add**.
- Step 5** In the **Select Users/Groups** dialog box, select the *user_name* and click **Add**.
- Step 6** Click **Check names** button, to verify the user name.
- Step 7** Then click **OK** to close out of each dialog box.
-



CHAPTER 19

Managing iSCSI

- [HyperFlex iSCSI Target Service Overview and Supported Use Cases, on page 247](#)
- [iSCSI Configuration Overview, on page 247](#)
- [iSCSI Network Page, on page 248](#)
- [iSCSI Initiator Group, on page 250](#)
- [iSCSI Target Page, on page 253](#)
- [iSCSI LUN Page, on page 256](#)
- [Configuring an iSCSI Initiator \(Windows\), on page 258](#)
- [Configuring an iSCSI Initiator \(Linux\), on page 258](#)
- [Cloning an iSCSI LUN, on page 259](#)

HyperFlex iSCSI Target Service Overview and Supported Use Cases

The HyperFlex iSCSI Target Service is available in HyperFlex 4.5(1a). Supported use cases for the the HX iSCSI Target Service are as follows:

- Support failover clustering, such as Microsoft Failover Clusters, for applications that require highly available shared storage, including Microsoft SQL Server and Hyper-V
- Present block storage to applications running inside or outside the HyperFlex cluster, such as Oracle Database or Oracle RAC deployments on external compute hosts
- Support Microsoft Exchange deployment over iSCSI
- Provision persistent volumes for Kubernetes through the HyperFlex Container Storage Interface for Kubernetes

Initiators are currently supported on Windows Server 2016, Windows Server 2019, Ubuntu 18.04 and 20.04, Oracle Linux 7, Red Hat Enterprise Linux 7.

iSCSI Configuration Overview

The process to configure the iSCSI Target Service is as follows:

- [Creating an iSCSI Network](#)

- [Creating an iSCSI Initiator Group](#)
- [Creating an iSCSI Target](#)
- [Linking iSCSI Initiator Group with Targets](#)
- [Creating an iSCSI LUN](#)
- [Configuring an iSCSI Initiator \(Windows\)](#)
- [Cloning an iSCSI LUN](#)

iSCSI Network Page

Displays configuration information for your iSCSI network.

iSCSI Network Data

If you have already configured your iSCSI network, basic network information is displayed.

Table 7: iSCSI Network Data

UI Element	Description
iSCSI Configuration	Status of iSCSI network configuration. If configured, status is indicated as "Network Configured".
LUNs	Number of LUNs created.
Capacity Used	Capacity used (in GB).
Cluster Capacity Utilization	Percentage of cluster capacity used by LUNs, Other and Free.
Targets	List and configuration information for iSCSI Targets.
Initiator Groups	List and configuration information for iSCSI Initiator Groups.

Creating an iSCSI Network

To create your iSCSI Network:

1. Enter the information required for the iSCSI Network: **iSCSI Network**, **Subnet**, **Gateway**, **IP Range**, **iSCSI Storage IP**, **Set non-default MTU**, and **VLAN Configuration**. Click **Next** when done.
2. Click **Configure**; or click **Cancel** to cancel your changes.
3. Click **Confirm** to confirm that you want to create your iSCSI Network; or click **Cancel** to cancel your changes.



Note After you confirm creation of your iSCSI network, you can not change some iSCSI network parameters without TAC assistance.

iSCSI Network Configuration Data

The following information is required to create your iSCSI Network.

Table 8: iSCSI Network Configuration Data

UI Element	Description
iSCSI Network	Configuration information for the iSCSI Network is displayed.
Subnet	Enter a valid Subnet
Gateway	Enter a valid Gateway
IP Range	Enter a valid IP Range
iSCSI Storage IP	Enter a valid IP Address for iSCSI Storage
Set non default MTU	<p>Checkbox to enable setting MTU (Message Transport Unit) manually. The MTU defines the maximum size of a network frame that you can send in a single data transmission across the network. The default MTU size is 9000.</p> <p>To disable Jumbo Frames, click the Set non default MTU checkbox, and then use the pull-down to change the value to 1500.</p> <p>Note If any of the initiators are crossing a router, the router will need to allow Jumbo Frames.</p>

UI Element	Description
VLAN Configuration	<p>Click in the checkbox to Create a new VLAN (Recommended) or Select an existing VLAN.</p> <p>To Create a new VLAN, you need to specify the following: VLAN ID, VLAN Name, UCS Manager host IP or FQDN, Username (Username for authentication with UCS), Password (Password for authentication with UCS).</p> <p>To Select an existing VLAN, you need to specify the VLAN ID.</p> <p>Note To configure the VLAN manually in UCS-M, use the create VLANs menu option. In the Create VLANs window, leave the checkboxes as is. In the vNIC templates for HX, attach the VLAN to "vNIC Template storage-data-a" and "vNIC Template storage-data-b". This configuration is non-disruptive.</p>

Editing an iSCSI Network

To edit your iSCSI network configuration:

1. Go to the **Actions** menu, and select **Edit Network**. The Edit Network for <cluster_name> window appears.
2. Add an IP Range for the iSCSI Network.
3. Click **Save Changes**; or click **Cancel** to cancel your changes.

Deleting an iSCSI Network

The option to delete your iSCSI network and retain objects you have configured is available using the "hxcli iscsi network delete" command or API. For more information, see the [CLI Reference Guide](#).

iSCSI Initiator Group

Displays configuration information for iSCSI Initiator Groups.

iSCSI Initiator Group Data

The following information is required to create an iSCSI Initiator Group.

Table 9: iSCSI Initiator Group Data

UI Element	Description
Name	Name for the Initiator Group.
Initiator IQN	iSCSI Qualified Name (IQN) for your Initiator. The IQN format takes the form <code>iqn.yyyy-mm.naming-authority:unique name.</code>

Creating an iSCSI Initiator Group

To create an iSCSI Initiator Group:

Before you begin

Before creating an iSCSI Initiator Group, you should have already created an [Creating an iSCSI Network](#).

Step 1 Go to the **Initiator Groups** tab, and click on **Create**.

The **Create Initiator Group** window appears.

Step 2 Enter a name for your Initiator Group in the **Name** field.

Step 3 Enter the **Initiator IQN** in the field provided. If you do not know the Initiator IQN, you can fetch it from the server:

- For Windows, login to the Windows machine. Go to Server Manager and click on **iSCSI Initiator**. Go to the **Configuration** tab. The Initiator IQN is identified in the **Initiator Name** field.
- For Linux, enter the command: `"sudo cat /etc/iscsi/initiatorname.iscsi"`. The Initiator IQN is identified in the **Initiator Name** field.

```
Cisco-Ubuntu:~$ sudo cat /etc/iscsi/initiatorname.iscsi
## DO NOT EDIT OR REMOVE THIS FILE!
## If you remove this file, the iSCSI daemon will not start.
## If you change the InitiatorName, existing access control lists
## may reject this initiator. The InitiatorName must be unique
## for each iSCSI initiator. Do NOT duplicate iSCSI InitiatorNames.
InitiatorName=iqn.2020-08.local.hx.green:Ubuntu-1
Cisco-Ubuntu:~$
```

The initiator name is in the file and can be changed with the appropriate rights.

Step 4 Click **Add Initiator**.

Note To allow access to initiators outside of the iSCSI VLAN subnet, use the `hxcli iscsi allowlist` command. For example:

```
hxcli iscsi allowlist add --ips 192.168.101.3
```

For more information, see the [CLI Guide, 4.5](#).

Step 5 Repeat the above steps to add multiple IQNs to participate in the same Initiator Group.

Step 6 Click **Create Initiator Group**.

The Initiator Group appears in the **Initiator Groups** tab.

Editing an iSCSI Initiator Group

To edit an iSCSI Initiator Group:

-
- Step 1** Click on the **Edit** (pencil) icon next to the **Create** button.
The **Edit Initiator Group** window appears.
- Step 2** Edit data for the iSCSI Initiator Group.
- Step 3** Click **Save Changes**; or click **Cancel** to cancel your changes.
-

Deleting an iSCSI Initiator Group

To delete an iSCSI Initiator Group:



Note You cannot delete an iSCSI Initiator Group when linked with a Target.

-
- Step 1** Click on the **Delete** (X) icon next to the **Create Initiator Group** button. The **Delete Initiator Group** window appears.
- Step 2** Click **Delete**; or **Cancel** to cancel your changes.
-

Linking iSCSI Initiator Group with Targets

To link an iSCSI Initiator Group:

Before you begin

Verify you have [Creating an iSCSI Initiator Group](#) and [Creating an iSCSI Target](#).

-
- Step 1** Go to the **Targets** tab, and select the name of the Target you want to link an Initiator Group.
- Step 2** From the **Linked Initiator Groups** tab, click in the checkbox to **Link**.
The **Link Initiator Groups** window appears.
- Step 3** Select the Initiator Group(s) to link to the Target.
- Step 4** Click **Link Initiator Groups**.
-

What to do next

After you have linked Initiator Groups with Targets;

- [Creating an iSCSI LUN](#)
- [Configuring an iSCSI Initiator \(Windows\)](#)
- [Cloning an iSCSI LUN](#)

Unlinking an iSCSI Initiator Group

To unlink an iSCSI Initiator Group:

-
- Step 1** Go to the **Targets** tab, and select the name of the Target you want to unlink an Initiator Group.
- Step 2** From the **Link Initiator Groups** tab, click in the checkbox to select the Initiator Group(s) to unlink to the Target.
- Step 3** Click on the **Unlink Initiator Group** button.
- Step 4** Click **Unlink Initiator Group(s)** to proceed or **Cancel**.
-

iSCSI Target Page

Displays configuration information for Targets.

Target Data

The following information is required to create a Target.

Table 10: Target Creation Data

UI Element	Description
Name	Name for the target.
Enable CHAP authentication	Click to enable CHAP authentication.
Username	Username for CHAP authentication
Secret	Secret for CHAP authentication

After a Target is created, the following information appears.

Table 11: Target Information Data

UI Element	Description
IQN	Target iSCSI Qualified Name for your Initiator.
Active Initiators Count	Total number of active Initiators.

UI Element	Description
Total LUN Capacity	Total storage capacity of LUN used and available (in Tb and in Gb).
CHAP authentication	Indicates whether you have Enabled CHAP authentication.
LUNs	Tab allowing you to create, edit, clone, delete and view LUNs for the Target.
Linked Initiator Groups	Tab allowing you to create and view LUNs for the Target.

Creating an iSCSI Target

To create an iSCSI target:

Before you begin

Before creating an iSCSI Target, you should have already created an [Creating an iSCSI Network](#). It is also recommended that you create [Creating an iSCSI Initiator Group](#).

-
- Step 1** Go to the **Targets** tab, and click on **Create**. The **Create Target** window appears.
- Step 2** Enter a name for the target in the **Target name** field.
- Step 3** (Optional) If you want to enable CHAP authentication, click in the **Enable CHAP authentication** checkbox. Fields for Username and Secret appear. Enter the Username and Secret.
- Note** For Windows, the Secret must be 12-16 characters.
- Note** CHAP based authentication is not supported during iSCSI Discovery phase.
- Step 4** Click **Create Target**.
-

The Target appears in the Targets tab.

What to do next

After you have created an iSCSI Target:

- [Linking iSCSI Initiator Group with Targets](#)
- [Creating an iSCSI LUN](#)
- [Configuring an iSCSI Initiator \(Windows\)](#)

Editing an iSCSI Target

To edit an iSCSI Target:

-
- Step 1** Click on the **Edit** (pencil) icon next to the **Create** button.
The **Edit Target** window appears.
 - Step 2** Edit data for the Target.
 - Step 3** Click **Save Changes**; or click **Cancel** to cancel your changes.
-

Deleting an iSCSI Target

To delete a Target:

-
- Step 1** Click on the Delete (X) icon next to the Create Target button.
The Delete Target window appears.
 - Step 2** Click **Delete Target**; or click **Cancel** to cancel your changes.
- Note** You can not delete a Target if there is an Initiator Group linked to the target.
- Note** You can not delete a Target if LUNs are created for it. If you still want to delete the Target, you must first delete all LUNs created for it.
-

Linking iSCSI Targets

To link an iSCSI Target:

Before you begin

Before linking an iSCSI Target, you need to create and configure [Creating an iSCSI Initiator Group](#). If you have not yet created an Initiator Group, you will need to do so before linking to an iSCSI Target. For more information, see [Creating an iSCSI Initiator Group, on page 251](#).

-
- Step 1** Go to the **Initiator Groups** tab, and select the name of the Initiator Group you want to link to a Target.
 - Step 2** From the **Linked Targets** tab, click on the **Link** button.
The Link Target window appears.
 - Step 3** Select the Target(s) to link to the Initiator Group(s).
 - Step 4** Click **Link Target**.
-

The linked targets you selected appear in the Linked Targets tab.

What to do next

After you have linked iSCSI Targets.

- [Creating an iSCSI LUN](#)
- [Configuring an iSCSI Initiator \(Windows\)](#)
- [Cloning an iSCSI LUN](#)

Unlinking an iSCSI Target

To unlink an iSCSI Target:

-
- Step 1** Go to the **Initiator Groups** tab, and select the name of the Initiator Group you want to unlink targets.
- Step 2** From the **Linked Targets** tab, click in the checkbox to select the Target(s) that you want to unlink.
- Step 3** Click on the **Unlink Target** button.
- Step 4** Click **Unlink**; or click **Cancel** to cancel your changes.
-

iSCSI LUN Page

Displays configuration information for iSCSI LUN.

iSCSI LUN Data

The following information is required to create a LUN.

Table 12: iSCSI LUN Data

UI Element	Description
Name	Name of the LUN.
Size	Total capacity size of the LUN (in GB).

After a LUN is created, the following information appears.

Table 13: LUN Information Data

UI Element	Description
Name	Name of the LUN.
LUN ID	Unique ID for the LUN.
Serial No.	Serial Number of the LUN.
Size	Total capacity size of the LUN (in Gb).

UI Element	Description
Used	Total capacity of the LUN used (in Gb).
Available	Total capacity of the LUN available (in Gb).

Creating an iSCSI LUN

To create an iSCSI LUN:

-
- Step 1** Go to the **Targets** tab, and select the name of the Target you want to create a LUN.
 - Step 2** Click in the **Create LUN** checkbox.
The **Create LUN** window appears.
 - Step 3** Enter a name for the LUN in the **LUN** field.
 - Step 4** Enter the size and unit of the LUN in the **Size** field.
 - Step 5** Click **Create LUN**.
-

The LUN appears in the LUNs tab for the Target.



Note There is a limit on the number of LUNs that you can expose per target. On Linux systems, the limit is 255 LUNs per target. On Windows systems, the limit is 254 LUNs per target.

What to do next

After you have created an iSCSI LUN, you can [Configuring an iSCSI Initiator \(Windows\)](#)

Editing an iSCSI LUN

To edit an iSCSI LUN:

-
- Step 1** Go to the **Targets** tab, and select the name of the Target you want to edit LUNs.
 - Step 2** From the LUNs tab, click in the checkbox to select the LUN you want to edit.
 - Step 3** Click on the **Edit** icon next to the Create LUN button. The Edit LUN window appears.
 - Step 4** Edit data for the LUN.
 - Step 5** Click **Edit LUN**; or click **Cancel** to cancel your changes.
-

Deleting an iSCSI LUN

To delete an iSCSI LUN:

-
- Step 1** Go to the **Targets** tab, and select the name of the Target you want to delete LUNs.
 - Step 2** From the LUNs tab, click in the checkbox to select the LUN you want to delete.
 - Step 3** Click on the **Delete** (X) icon next to the Clone LUN button. The **Delete LUN** window appears.
 - Step 4** Click **Delete**; or click **Cancel** to cancel your changes.
-

Configuring an iSCSI Initiator (Windows)

This procedure describes how to configure a Windows machine as an iSCSI Initiator. This must be done to initialize, online and create a volume on the iSCSI LUN prior to cloning the LUN.

-
- Step 1** Login to the Windows machine that you want to configure as an iSCSI Initiator.
 - Step 2** Go to **Server Manager**, and click on **iSCSI Initiator**. Click **Yes** to continue.
 - Step 3** Enter the Target's Hostname or IP address in the Target tab, and then click on **Quick Connect**.
Discovered Targets appear as "HX Cluster IP(CIP)" and display the Target IQN and status of the Targets.
 - Step 4** Click **Done**.
 - Step 5** Select the target, and then click **Connect** on the Target tab.
 - Step 6** Click **Advanced**.
 - Step 7** Click to "Enable CHAP log on", and then specify the username and password for HyperFlex. Click **OK**.
If there are no issues with the configuration, the status is updated to indicate "Connected".
 - Step 8** Verify that the iSCSI LUN(s) is attached in the Disk Management tool.
-

You can now initialize, online and create a volume on the iSCSI LUN(s).

Configuring an iSCSI Initiator (Linux)

This procedure describes how to configure an iSCSI Initiator on Linux. This must be done to initialize, online and create a volume on the iSCSI LUN prior to cloning the LUN.

-
- Step 1** Verify that the `iscsiadm` command is present.
 - Step 2** Run the following command: `sudo apt-get install open-iscsi`.
 - Step 3** Discovery the target. You can do so by running the following command: `sudo iscsiadm -m discovery -t sendtargets -p <HX iSCSI CIP>`.

Step 4 Login to the target. You can do so by running the following command: `sudo iscsiadm -m discovery -t sendtargets -p <HX iSCSI CIP> -l.`

Note With `lsblk -scsi`, you can verify which device is the target. You can now partition it with `gdisk` and format the HX iSCSI drive.

Cloning an iSCSI LUN

If you require the LUN to be application consistent, you can clone the LUN. Cloning LUNs is supported on Windows and Linux hosts. However, the application consistent LUN clones are only supported on Windows via VSS.

To clone an iSCSI LUN on a Windows machine, you must first install HX Windows Agent. For more information, see [Installing HX Windows Agent for iSCSI Clone LUN, on page 260](#).

To clone an iSCSI LUN:

Before you begin

Install the HX Windows Agent on each initiator. For more information on installing the HX Windows Agent, see [Installing HX Windows Agent for iSCSI Clone LUN, on page 260](#).

-
- Step 1** Go to the **Targets** tab, and select the name of the Target you want to clone LUNs.
- Step 2** From the LUNs tab, click in the checkbox(es) to select the LUN(s) you want to clone.
- Step 3** Click on the **Clone LUN** icon next to the Edit LUN button. The Clone LUNs window appears.
- Step 4** Click in the **Application Consistent** checkbox to enable. Specify the administrator account (local or AD) username and password for the Windows machine to verify and authenticate the VSS user.
- Note** If you do not enable the Application Consistent checkbox, the iSCSI Clone LUN(s) will be Crash Consistent.
- Step 5** Enter a name for the new destination target in the **New destination target name** field.
- Step 6** If you want to enable CHAP authentication, select the check box to **Enable CHAP Authentication**. For each source LUN(s), enter the destination LUN(s) name in the **Destination LUN Name** field.
- Note** CHAP based authentication is not supported during iSCSI Discovery phase.
- Step 7** Click **Clone**; or click **Cancel** to cancel your changes.
- Note** The HX Windows Agent will not identify LUN(s) which are involved for other applications for example, SQL Server, Exchange. You must select the required LUN(s) for cloning.
-

What to do next

To access cloned LUN(s), link the destination Target with an Initiator Group, and then discover the LUN(s) by refreshing the iSCSI target from the Initiator window. Use the `HxWindowsAgentUtils.exe` to change the disk/volume properties for the destination LUN(s).

Limitations for the HX Windows Agent

The following limitations apply for the HX Windows Agent:

- iSCSI LUN(s) should be added using an in-guest Initiator in Windows virtual machines deployed on the vCenter Server\ESXi Host.
- Shared Disk on the Windows machine not managed by Microsoft Failover Cluster is not supported for iSCSI Clone LUN operations.
- From HX Connect, you can clone the iSCSI LUN on a new destination target only.
- For application consistent iSCSI Clone LUN, make sure to select all the LUNs which are used by the application (for example, SQL Server, Oracle, etc.)
- Clone of Cluster Shared Volume(CSV) and clustered disk/standalone disk combination is not supported. SCV LUN(s) only or combination of standalone/clustered disk is supported for the clone operation.
- The HyperFlex VSS Hardware Provider service can only be invoked by the HyperFlex Windows Agent. If third-party backup vendor tries to clone the HX LUN, “Microsoft Software Shadow Copy” Provider will be invoked.
- Upgrade and patching of the HX Windows Agent is supported.
- Upgrade of HyperFlex will not upgrade the HX Windows Agent. You must manually upgrade the HX Windows Agent.
- If you restart or power off the HyperFlex cluster node, the iSCSI Clone LUN operation fails.

Prerequisites for the HX Windows Agent

To run the HX Windows Agent, you must meet the following prerequisites:

- Windows 2016 or later with basic configuration of the host machine.
- Ports 10152 and 9347 must be open on the Windows Server
- Ports 10151 and 9347 must be open on the HX Controller VMs
- Ensure that the “Hyperflex Windows Agent” and “Hyperflex VSS Hardware Provider” services are installed on the Windows machine.
- Ensure that the “Hyperflex Windows Agent” is in running state.
- Ensure that both the services - “Hyperflex Windows Agent” and “Hyperflex VSS Hardware Provider” are installed on all Windows machines that are exposed as an iSCSI Initiator.
- Provide Administrator or AD credentials while triggering the Application Consistent iSCSI Clone LUN workflow.
- Ensure that the source LUN you are going to clone is discovered and the volume with volume label is created on the Windows machine.

Installing HX Windows Agent for iSCSI Clone LUN

This procedure describes how to install the HX Windows Agent for iSCSI Clone LUN.

Before you begin

Ensure that you are running Microsoft Windows Server 2016 or later with basic configuration on the VM/Bare Metal. The Installer adds rules to allow inbound communication over 10152 port. If you are using any third-party firewall or antivirus software however, you will need to ensure that port 10152 is open.

-
- | | |
|---------------|--|
| Step 1 | Login to the Windows machine using Administrator or AD credentials. |
| Step 2 | Run the Windows HX Agent install executable by double-clicking on the file:
<code>HxWindowsAgentIscsiClone-v4.5.1a-39020.exe</code> . |
| Step 3 | Click Next . |
| Step 4 | Click to select “I accept the terms in license agreement” and click Next . |
| Step 5 | Confirm <Program File>\Cisco\HxWindowsAgent as the location of the installation directory, then click Next . |
| Step 6 | Click Install . |
| Step 7 | Click Finish . |
-

This installs the HyperFlex Windows Agent and HyperFlex VSS Hardware Provider services. Other installation notes are as follows:

- You can view the HyperFlex Windows Agent and HyperFlex VSS Hardware Provider services in Windows services. The HyperFlex Windows Agent should be in running state and the HyperFlex VSS Hardware Provider is in stopped state. The HyperFlex VSS Hardware Provider service is started when you request to clone or backup a LUN.
- You can view MSI installation and other install details in %appdata%\HxWinAgentMsiInstall.log.
- In the installation directory, you may notice some dependent dll's. Do not delete or update these dependencies. If deleted, you will need to use the repair option in the installer to restore.
- View service logs in C:\HxWindowsAgent\Logs\ HxAgentService_<DateTime>.log. The Windows Registry is in location : HKEY_LOCAL_MACHINE\SOFTWARE\HyperFlex. This entry will have the location of Service logs.
- To verify the Agent version: Go to the Installation directory, right-click on HxWinAgentService.exe and select **Properties**. Go to the **Details** tab and verify the Product Version.
- You can view installation and operation events on the service in the Event viewer with Source as “HxVssHardwareProvider” and “HxWindowsAgent”. Inbound rules are set with the name “Hx Windows Agent” on port 10152 and enabled for all IP Addresses in Windows firewall .

Installing HX Windows Agent (with Pre-Installed Dependencies) for iSCSI Clone LUN

This procedure describes how to install the HX Windows Agent to enable iSCSI Clone LUN for HyperFlex, if there are dependencies - for example, if Microsoft .NET framework 4.5 (version : 4.5.50709 or above), Microsoft Visual C++ 2017 Redistributable (x64) (version: 14.10.25017 or above) , Microsoft Visual C++ 2017 Redistributable (x86) (version: 14.10.25017 or above) programs are already installed on the Windows machine.

-
- Step 1** Login to the Windows machine using Administrator or AD credentials.
- Step 2** Double-click on the file: HxWindowsAgentIscsiClone-v4.5.1a-39020.msi
- Note** You cannot use HxWindowsAgentIscsiClone-v4.5.1a-39020.msi to uninstall the agent.
- Step 3** Click **Next**.
- Step 4** Click to select “I accept the terms in license agreement” and then click **Next**.
- Step 5** Select the Installation directory. The default location is <Program File>\Cisco\HxWindowsAgent.
- Step 6** Click **Install**.
- Step 7** Click **Finish**.
-

Uninstalling HX Windows Agent for iSCSI Clone LUN

This procedure describes how to uninstall the HX Windows Agent for iSCSI Clone LUN.

-
- Step 1** Login to the Windows machine using Administrator or AD credentials.
- Step 2** Double click on the file: HxWindowsAgentIscsiClone-v4.5.1a-39020.exe.
- Step 3** Click **Next**.
- Step 4** Select Remove and Click **Next**.
- Step 5** Click **Remove**.
- Step 6** Click **Finish**.
-

- When done, the uninstaller removes “Hyperflex Windows Agent” and “Hyperflex VSS Hardware Provider” services.
- The uninstaller deletes inbound rules with the name “Hx Windows Agent” having port 10152 from Windows Firewall.
- The uninstaller does not remove Microsoft .NET framework 4.5 (version : 4.5.50709 or above), Microsoft Visual C++ 2017 Redistributable (x64) (version: 14.10.25017 or above), Microsoft Visual C++ 2017 Redistributable (x86) (version: 14.10.25017 or above) programs.
- The uninstaller does not delete files and folders from the location: C:\HxWindowsAgent\Logs\HxAgentService_<DateTime>.log and Installation directory\HxCAInstallLogMsi.txt”.
- HKEY_LOCAL_MACHINE\SOFTWARE\HyperFlex registry entry is retained.

iSCSI HX Windows Agent Logs

The following logs are available for the HX Windows Agent.

Table 14: HX Windows Agent Logs

Log File	Description
hxCloneSvcMgr.log	Located in: /var/log/springpath
hxApplicationConsistentSvcMgr.log	Located in: /var/log/springpath
HxAgentService_<DateTime>.log	Located in: %SystemDrive%\HxWindowsAgent\Logs. This file contains Windows Agent logs specific to the VSS requestor.
HxVSSProvider_<DateTime>.log	Located in: %SystemDrive%\HxWindowsAgent\Logs. This file contains Windows Agent logs specific to the VSS hardware provider.
HxWinAgentMsiInstall.log	Located in: %appdata%. This file contains Windows Agent Installation logs.

Transferring HX Windows Agent Logs

To transfer helpful support information including the HXWindows Agent logs from a Windows machine to the Controller VM, run the following command from the Controller VM machine:

bash-4.2# hxWindowsAgentLogging



Note This command accepts input parameters including the Windows IP, username and password required to pull the logs.

The logs are transferred to the following location on the Controller VM machine:

: /var/log/springpath/<WindowsIP> in a file called "HXLogs.zip". The HXLogs.zip file includes the following information: HX Windows Agent logs, Disk Details from the HxDiskInfo.log, and system information from HxSystem.log.

Editing the Location of Service Logs

This procedure describes how to change the location of service logs.

-
- Step 1** Login to the Windows machine using Administrator or AD credentials.
 - Step 2** Open the Registry Editor.
 - Step 3** Open HKEY_LOCAL_MACHINE\SOFTWARE\HyperFlex.
 - Step 4** Right-click on the **Data** field of TargetDirectory and select **Modify**.
 - Step 5** Edit the value for the location of the log file.
 - Step 6** Restart the Hx Windows Agent service from Windows services.
-

Accessing Cloned LUN(s) on Destination Target

Follow the below procedure to discover the Destination LUN(s) on the Destination Target using the iSCSI Initiator window. You can also use the `HxWindowsAgentUtils.exe` located in the HX Windows Agent install directory.

-
- Step 1** Go to `diskmgmt.msc` and right-click on the required Disk <Disk ID> as "Online".
- Step 2** Open a command prompt as "Administrator". Run `diskpart.exe`.
- Step 3** Run the command: `List Disk`.
- Step 4** Run the command: `Select Disk <Disk ID>` (Select the corresponding disk accordingly).
- Step 5** Run the command: `Detail Disk`.
- Step 6** If the disk attribute "Read-only" is "Yes", set as "No": "attributes disk clear readonly"
- Step 7** Select Volume <Volume ID> (select volume shown as part of detail disk)
- Step 8** Run the following commands:

- `attributes volume clear READONLY`
- `attributes volume clear SHADOWCOPY`
- `attributes volume clear NODEFAULTDRIVELETTER`
- `attributes volume clear HIDDEN`

You should now have a disk with volume accessible and the volume label with read/write permissions.



CHAPTER 20

Cisco HyperFlex HTML Plugin for VMware vCenter

- [Cisco HyperFlex vCenter Plugin, on page 265](#)
- [Cisco HyperFlex HTML5 Plugin for VMware vCenter , on page 265](#)
- [Cisco HyperFlex Flash Plugin for VMware vCenter , on page 285](#)

Cisco HyperFlex vCenter Plugin

Cisco HyperFlex offers the flexibility of monitoring and managing content via Cisco HX Connect, Cisco HyperFlex Plugin, or Storage Controller VM command line. This chapter describes managing and monitoring content through the Cisco HyperFlex Plugin.

The Cisco HyperFlex vCenter Plugin is integrated with the vSphere Web Client and supports all of the HX Data Platform post-installation management and monitoring functions. Access the Cisco HyperFlex vCenter Plugin directly through the vSphere Web Client Navigator.

Cisco HyperFlex HTML5 Plugin for VMware vCenter

This section describes how to monitor and manage your HyperFlex clusters from the VMware vCenter using the Cisco HyperFlex HTML5 plugin.

Cisco HyperFlex HTML5 Plugin Prerequisites

The following hardware and software prerequisites apply to the Cisco HyperFlex HTML5 Plugin:

- Cisco HyperFlex HTML5 plugin for VMware vCenter support was introduced in Cisco HX Release 4.0(2a) and later on vCenter 6.5U2 and later.
- Browser compatibility: The Cisco HyperFlex HTML plugin works with Chrome, Firefox and IE.
- Cisco HyperFlex HTML5 plug-in 2.0.0 is the minimum version supported. If the running version is 1.0.1, upgrade to the latest version.
- Administrative Privileges are required for managing users and roles.



Note **Cisco HyperFlex Release 4.5(x)** - Perform install, uninstall, and upgrade operations by running `install_vc_plugin` on your shell.

Install and Register the vCenter HTML5 Plugin

Install the Cisco HyperFlex HTML5 plugin with the VMware vSphere web client. The plugin installation process requires the following information:

- vCenter Server FQDN/IP
- vCenter server Username and Password
- HX Controller VM Password

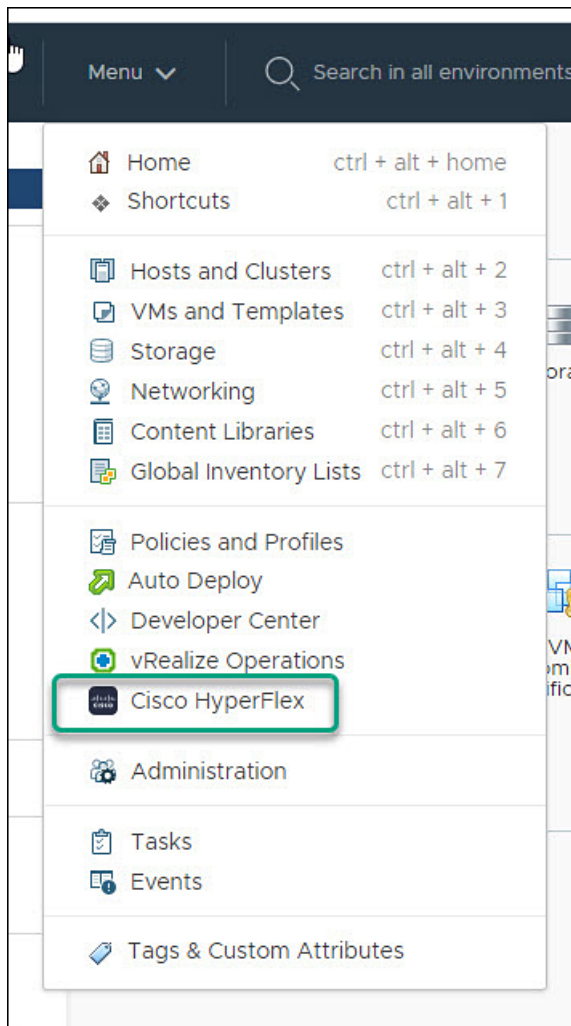
Table 15: CLI Arguments

Option	Required or Optional	Description
-h, --help	Optional	Shows the help message relative to the listed command and exits.
-u, --unregister	Optional	Unregister Cisco HyperFlex vCenter plugin.
-s, --show	Optional	Displays the HTML5 vCenter plugin details.
-v, --verbose	Optional	Make the operation more verbose.

Before you begin

- Check and confirm the HTTP (port 80) and HTTPS (port 443) connectivity between vCenter and Controller VMs.

-
- Step 1** Download the Cisco HyperFlex HTML plugin for VMware vCenter from the [Cisco Software Download](#) site.
- Step 2** Copy the `HyperFlex-VC-HTML-Plugin-2.0.x.zip` file into a temporary directory in one of the controller VMs and unzip.
- Step 3** Execute `install_vc_plugin` on your shell and enter:
- vCenter FQDN/IP address
 - Administrator username and password of vCenter server
 - Controller VM root password
- Step 4** Log on to vCenter and a blue banner message appears to confirm that the new plugin is installed.
- Step 5** Logout and login again to vCenter to see the Cisco HyperFlex menus for HTML5 plugin.



Verifying the Cisco HyperFlex HTML5 Plugin Installation from the vSphere Client

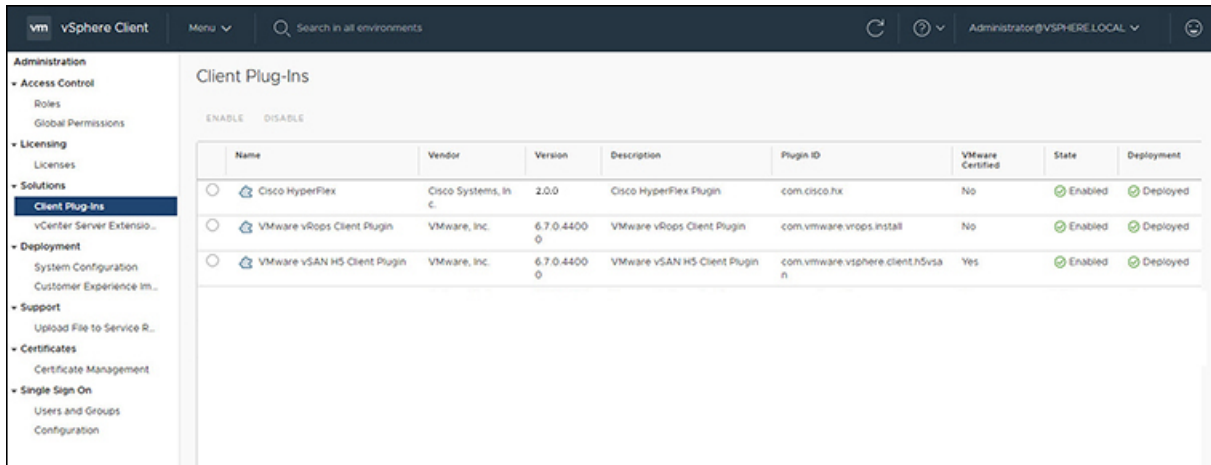
To verify Cisco HyperFlex plugin installation from the vSphere Client UI.

Before you begin

The HTML5 plugin should be installed on vCenter server.

Launch the vSphere client, select **Menu > Administration > Solutions > Client Plug-Ins**

Uninstalling the HyperFlex HTML5 Plugin



Uninstalling the HyperFlex HTML5 Plugin

To uninstall the HX Data Platform HTML5 Plugin, perform the following steps.

Step 1 Execute the uninstall command `install_vc_plugin` on the shell and enter the following credentials:

- vCenter FQDN/IP address
- Administrator username and password for the vCenter server

Step 2 Restart vSphere UI service of vCenter server.

Upgrading the HTML5 Plugin

When you want to upgrade to the latest HTML plugin, download the Cisco HyperFlex HTML plugin for VMware vCenter from the [Cisco Software Download](#) site.

Before you begin

Use this task only if the version of the HTML Plugin installed on the vCenter server is before 2.0.x.

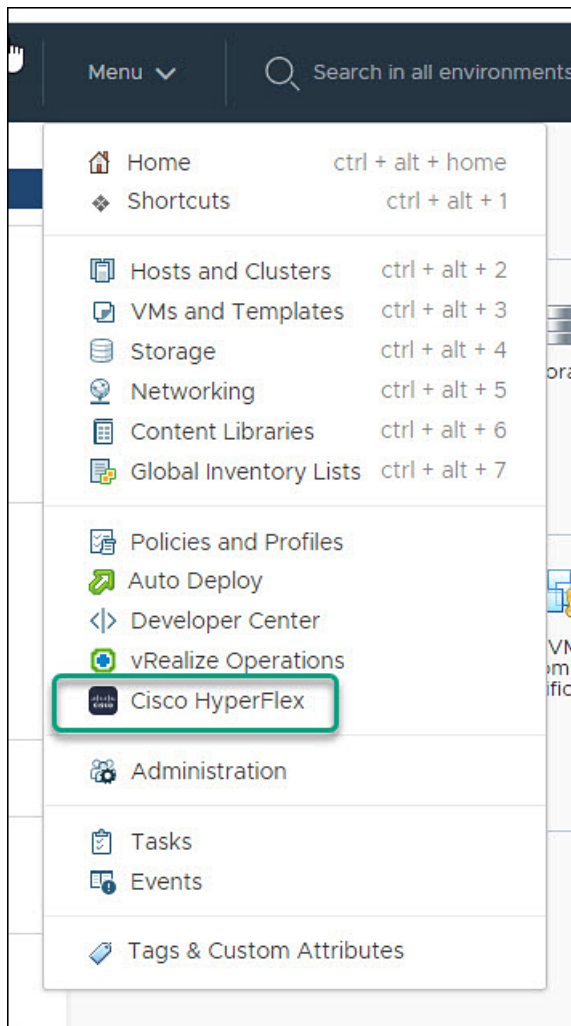
Step 1 Copy the `HyperFlex-VC-HTML-Plugin-2.0.x.zip` file into a temporary directory in one of the controller VMs and unzip.

Step 2 Execute `install_vc_plugin` on your shell and enter:

- vCenter FQDN/IP address
- Administrator username and password of vCenter server

Step 3 Select **Y** to continue the Upgrade process with controller VM password.

Step 4 Logout and login again into vCenter to see Cisco HyperFlex listed in the vCenter menus.



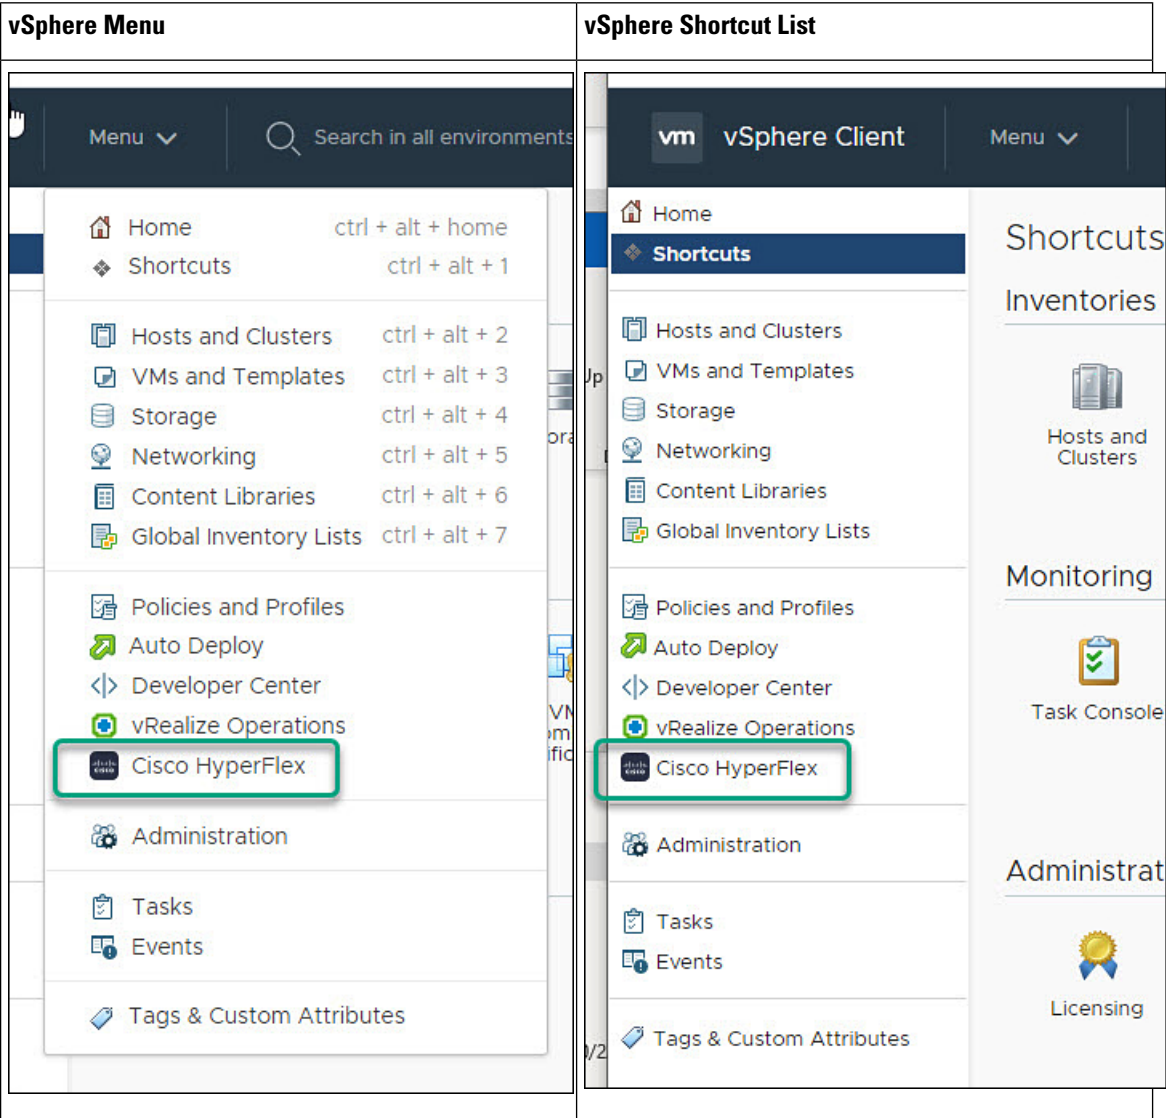
Using the Cisco HyperFlex HTML5 Plugin

The plugin interface allows you to:



- Discover the Registered HX Cluster
- View HX Cluster Summary
- Events and Alarms
- HX Datastore Management
- Create HX Snapshots and clones at the virtual machine level
- Manage users and access to HX clusters
- Cross-launch HX Connect for upgrade
- Embedded vCenter server actions at the Host and Clusters level



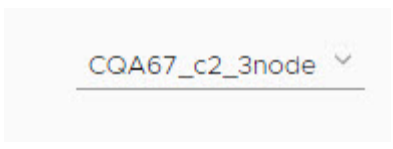

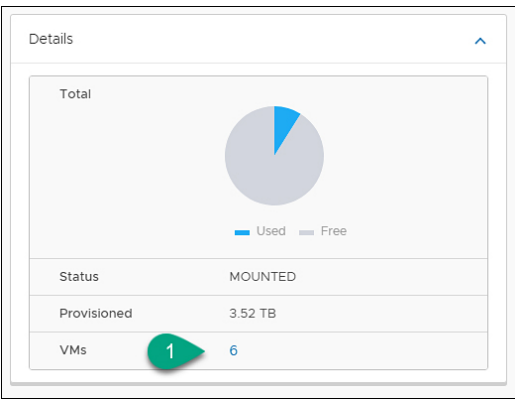
Navigating the HTML5 plugin

Accessing the Cisco HyperFlex HTML5 Plugin is easily accessed from the vSphere Menu or the Shortcuts list.



The Cisco HyperFlex HTML5 Plugin has functionality that is common throughout the plugin. This section describes the icons and their usage.

Icon	Usage
	Cisco HyperFlex Plugin. When installed, this icon is found on the Menu and the Shortcuts list.
	Refresh the view.

Icon	Usage
	Filter the content seen in the browser.
	Expand or collapse the contents.
	Navigate between clusters.
	Use the VC Cluster button to jump from the HyperFlex Events or Alarms view to vCenter Events or Alarm Page.
	Clicking on a VM count (number) takes the user directly to the Datastore page which list all VMs for that datastore.

Cluster Management

Managing Users and Access to HX Clusters

The vCenter plugin requires the user to have administrator privileges. You can create a user and assign administrator role to that user from **Permissions** tab on cluster level.

To manage users and access to HX clusters, assign the **No Access** Role to all the clusters for that user.



Note Administrative Privileges are required for managing users and roles.

Discover the Registered HX Cluster

To discover your HX clusters and map the vSphere managed objects in your deployment perform the following steps:

-
- Step 1** Login to the vSphere web client.
- Step 2** Select **Menu > Cisco HyperFlex**
- Step 3** Click **Rescan** to refresh the list of HX clusters displayed.
Registered clusters are displayed in HyperFlex Cluster table along with a summary of the cluster details.
- Step 4** If you have added new HX Cluster(s) to the vCenter server and they are not appearing in the cluster list, Click the **Rescan** icon on top of the cluster list grid to reload the cluster list from HyperFlex.
-

View the HX Cluster Summary

To view a summary of the HX Clusters in your deployment perform the following steps:

-
- Step 1** Log in to the vSphere web client.
- Step 2** Select **Menu > Cisco HyperFlex**
- Step 3** Click the discovered HX cluster name to view its summary.
- Step 4** Click on **Summary** to view details about Total Nodes, Datastores, HyperFlex Version, Model, vCenter Cluster, ESXi Version and Uptime.

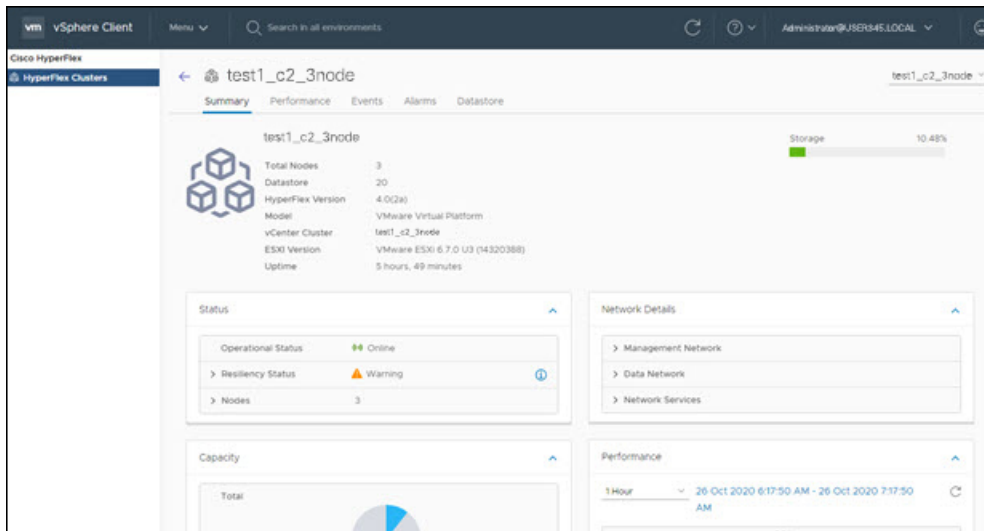


Table 16: Cluster Summary View Details

Field Name	Additional Information
Total Nodes	Total number of nodes
Datastore	Datastore connected to the cluster
HyperFlex Version	Version of HyperFlex on the cluster
Model	Model name
vCenter Cluster	vCenter cluster name

Field Name	Additional Information
ESXi Version	ESXi Version
Uptime	Length of time that the cluster has been up and running

- a) The summary view includes four portlets with additional details about the cluster: Status, Network Details, Capacity and Performance.

Use the arrows to collapse and expand the portlet contents.

Table 17: Status Portlet

Field Name	Additional Information
Operational Status	Online or Offline
Resiliency Status	<p>Warning or Healthy</p> <p>Click the arrow to collapse or expand additional Resiliency Status details:</p> <ul style="list-style-type: none"> • Host(s) failure tolerance - Number of host failures tolerated • Replication Factor- Number of copies • Creation time- Cluster creation time • Persistent Device failures tolerable- Number of device failures tolerated • Caching Device failures tolerable-Number of caching device failures tolerated
Nodes	<p>Number of nodes in the cluster.</p> <p>Click the arrow to collapse or expand additional Nodes details:</p> <ul style="list-style-type: none"> • Node Type • Version

Table 18: Capacity Portlet

Field Name	Additional Information
Total	Used and Free Capacity expressed as a percentage
Usable Capacity	Amount of usable capacity
Used	Used Capacity
Free	Free Capacity

Field Name	Additional Information
Saving	Total amount space saved Click the arrow to collapse or expand the details about the saved space with Compression and Deduplication. The data is expressed as a percentage.

Table 19: Network Details Portlet

Field Name	Additional Information
Management Network	Management Network details Click the arrow to view the following Management Network details: <ul style="list-style-type: none"> • Management IP address / FQDN • VLAN • Subnet Mask • Default Gateway
Data Network	Data network details Click the arrow to view the following Data Network details: <ul style="list-style-type: none"> • Data IP address / FQDN • VLAN • Subnet Mask • Default Gateway
Network Services	Network Services details Click the arrow to view the following Network Services details: <ul style="list-style-type: none"> • DNS Server(s) • NTP Server(s)

Table 20: Performance Portlet

Field Name	Additional Information
General Usage	<ul style="list-style-type: none"> • Click on the Time Interval list to select the length of time viewed in the performance chart. • Hover over the chart line to display totals for a specific time. • Click the refresh arrow to refresh the view. • To change the timezone, click the current time interval, complete the Time Range pop-up, and click OK. The time seen reflects the browser time.

Field Name	Additional Information
IOPS	Display IOPS performance chart
Throughput	Display Throughput performance chart
Latency	Display Latency performance chart

Events

To view events specific to the Cluster, Node, Host, VM or Disk, perform the following steps:

- Step 1** Log in to the vSphere client.
- Step 2** Select **Menu > Cisco HyperFlex**.
- Step 3** Click the Cluster name that you want to view.
- Step 4** Using the Cluster Summary Tabs, Click **Events**.
The Events Detail view appears.

Table 21: Event Details

Field Name	Additional Information
Description	Text description of the event
Event Name	Event name
Entity Type	Entity affected. Values include: All, Cluster, Node, Virtual Machine, and Disk
Severity	Event severity level. Valid values include: All, Info, Warning, Error, and Critical
DateTime	Date and time the event occurred

- Step 5** (Optional) Use filters to limit the results that appear in the Events Table.

Filter	Filter Options
Entity Type	All, Cluster, Node, Virtual Machine, and Disk
Severity	All, Info, Warning, Error, and Critical
Filter	Type a keyword in the Filter option to filter the table contents seen in the browser

- Step 6** In the list of events, Click on the event name that you want more information about.
The details appear below the Events table. Details include:

- Description

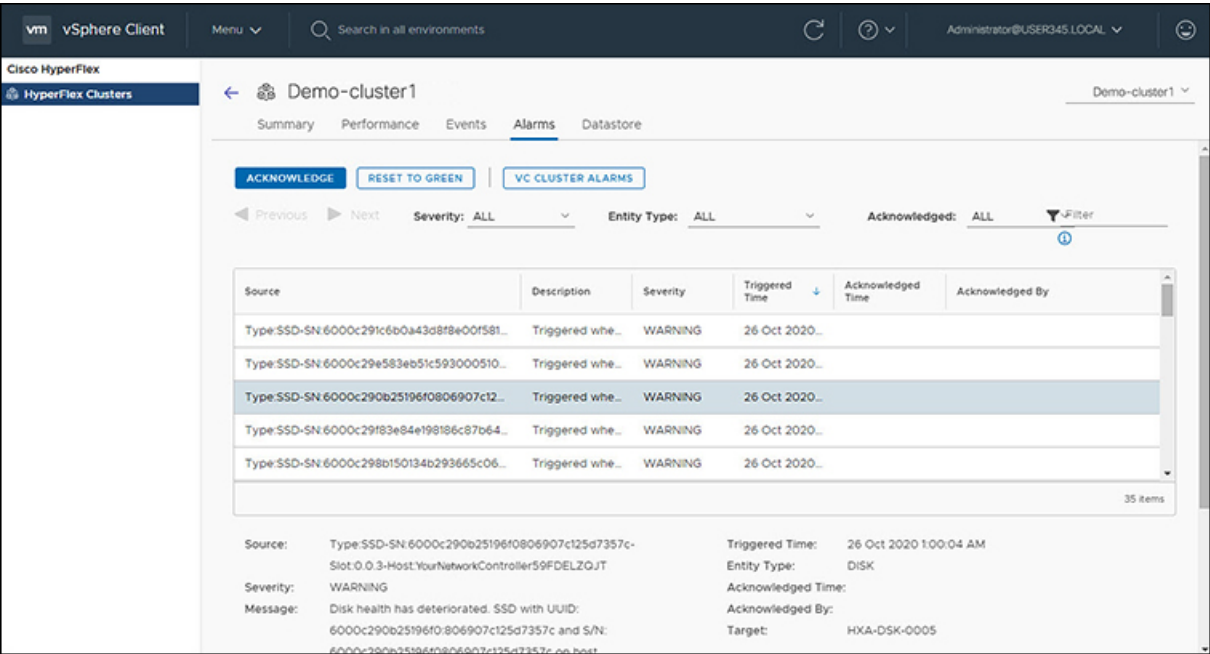
- Event Name
- Entity Type
- Severity
- DateTime

Alarms

To view alarms specific to the Cluster, Host, and VMs, perform the following steps:



Note Acknowledged alarms on HX Connect or HTML plugin does not acknowledge the equivalent vCenter alarm.



- Step 1** Log in to the vSphere client.
- Step 2** Select **Menu > Cisco HyperFlex**
- Step 3** Click the Cluster name that you want to view.
- Step 4** Using the Cluster Summary Tabs, Click **Alarms**. The Alarms Detail view appears.

Table 22: Alarm Details

Field Name	Additional Information
Source	Text description of the alarm

Field Name	Additional Information
Description	Alarm name
Severity	Alarm severity level. Valid values include: All, Info, Warning, and Error
Triggered Time	Date and time the alarm occurred Use the arrow to sort and re-sort the table results
Acknowledged Time	Time when the Alarm was acknowledged
Acknowledged by	Auto-enters who acknowledged the alarm

Step 5 (Optional) Use filters to limit the results that appear in the Alarms Table.

Filter	Filter Options
Severity	All, Info, Warning, and Error
Entity Type	All, Cluster, Node, Virtual Machine, Disk and Datastore
Acknowledged	All, True, and False
Filter	Type a keyword in the Filter option to filter the table contents seen in the browser

Step 6 Click the **Acknowledged** button to acknowledge that the alarm has been seen.
Clicking the **Acknowledged** button auto-enters who acknowledged the alarm in the **Acknowledged by** field.

Step 7 Click the **Reset To Green** button to remove the alarm from the list.

HX Datastore Management

Managing Datastores

The Datastore page allows users to view datastore details, create, edit, mount, unmount or delete datastores on a Cluster.

Step 1 Log in to the vSphere client.

Step 2 Select **Menu > Cisco HyperFlex**

Step 3 Click the cluster that you want.

Step 4 Click **Datastore**.

The Datastore Detail Table appears. Use the arrows to advance to the next or previous page, and first or last page of datastores.

Table 23: Datastore Table Details

Field Name	Additional Information
Name	Datastore name

Field Name	Additional Information
Mount Summary	Mounted or Unmounted
Mount Status	Normal
Provisioned	Amount of space provisioned
Used	Amount of space used
Free	Amount of space available

Step 5 Click on a Datastore name in the table to view additional details about the Datastore. The Details and Trends portlets appear below the table.

Table 24: Details Portlet

Field Name	Additional Information
Total	Graph of space provisioned and used
Status	Mounted or Unmounted
Provisioned	Amount of space provisioned
VMs	Number of VMs created on the datastore Clicking on a VM count (number) takes the user directly to the Datastore page which list all VMs for that datastore.

Table 25: Trends Portlet

Field Name	Additional Information
General Usage: <ul style="list-style-type: none"> Click on the Time Interval list to select the length of time viewed in the performance chart. Hover over the chart line to display totals for a specific time. Click the refresh arrow to refresh the view. To change the timezone, click the current time interval, complete the Time Range pop-up, and click OK. The time seen reflects the browser time. 	
IOPS	Display IOPS performance chart
Throughput	Display Throughput performance chart
Latency	Display Latency performance chart

Create New Datastore

To create a new datastore:

Before you begin

Login to the vSphere client and Select **Menu** > **Cisco HyperFlex**.

-
- Step 1** Click the cluster that you want.
- Step 2** Click **Datastore**.
- Step 3** Click the **Create** button. The Create New Datastore window appears.
- a) Enter the Datastore Name.
 - b) Enter the Size, and select GB or TB.
 - c) Select the Block Size, Select 4K or 8K.
- Step 4** Click **OK**. A new Datastore is created and added to the Datastore table list.
- Step 5** If the new Datastore does not appear in the list, click the **Refresh** arrow and recheck the list.
-

Edit Datastore

To edit an existing datastore:



Note You can only change the name of a datastore after it has been Unmounted.

Before you begin

Login to the vSphere client and Select **Menu** > **Cisco HyperFlex**.

-
- Step 1** Click the cluster that you want.
- Step 2** Click **Datastore** and select Datastore for Edit action.
- Step 3** Click the **Edit** button. The Edit Datastore window appears.
- Step 4** Edit the Datastore details.
- Step 5** Click **OK** to save your changes. The Datastore information is updated.
-

Mount or Unmount Datastore

The Mount and Unmount buttons are active based on the current status of the datastore. A Mounted datastore offers the option to Unmount the datastore, while an Unmounted datastore offers the option to Mount the datastore. To Mount or Unmount a datastore:

Before you begin

- Remove any VMs created or registered to the datastore before starting the Unmount action.
- Login to the vSphere client and Select **Menu** > **Cisco HyperFlex**.

-
- Step 1** Click the cluster that you want.
- Step 2** Click **Datastore** and select Datastore for Mount (Unmount) action.
- Step 3** Click the **Mount(Unmount)** button.
- The Mount (Unmount) Datastore window appears with a confirmation question "Do you want to mount (Unmount) the datastore?"
- Step 4** Click **OK** to continue with the Mount (Unmount) action, or click **Cancel** to exit the Mount (or Unmount) Datastore window. The datastore status is changed from Mounted to Unmounted or Unmounted to Mounted.
-

Delete Datastore

To delete a datastore:

Before you begin

- Remove any VMs created or registered to the datastore and unmount the datastore before starting the Delete Datastore action.
- Login to the vSphere client and Select **Menu > Cisco HyperFlex**.

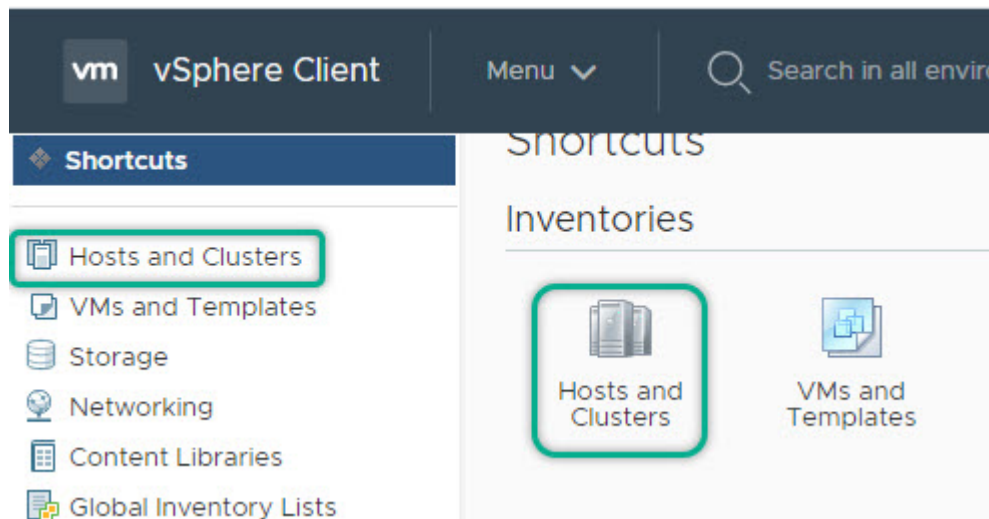
-
- Step 1** Click the cluster that you want.
- Step 2** Click **Datastore** and select Datastore for Delete action.
- Step 3** Click the **Delete** button.
- The Delete Datastore window appears with the confirmation question "Do you want to delete the datastore?"
- Step 4** Click **OK** to continue with the delete action, or **Cancel** to exit the Delete Datastore window. The selected Datastore is deleted from the Datastore table list.
-

vCenter: HyperFlex Plugin Embedded Actions

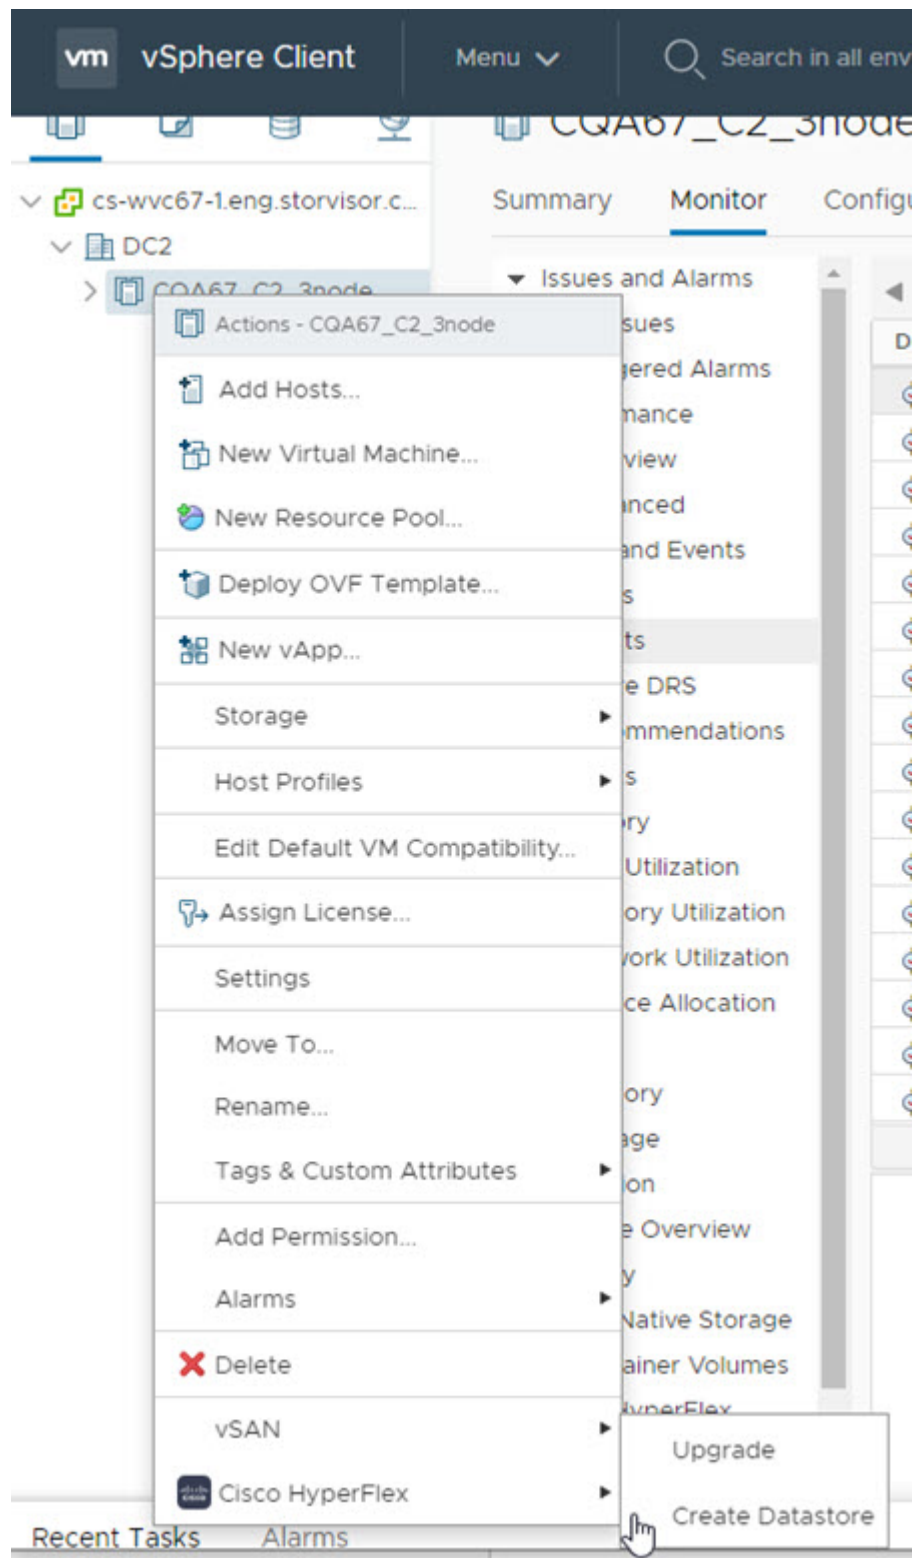
vCenter Server Action at Host and Clusters Level

From the **Host and Cluster** level in the vCenter server users have the ability to create new datastores or upgrade HyperFlex.

- Access the **Hosts & Clusters** either from the vSphere Menu or the Shortcut link.



- Right click on the cluster and select **Cisco HyperFlex > Upgrade**.
 - **Upgrade** Launches HyperFlex Connect and takes the user directly to the upgrade page to complete the upgrade process.
- Right click on the cluster and select **Cisco HyperFlex > Create Datastore**. The Create New Datastore window appears.



1. Enter the Datastore Name

2. Enter the Size, and select GB or TB.
3. Select the Block Size, Select 4K or 8K.
4. Click **OK**.

Viewing Cisco HyperFlex HTML5 Plugin Portlets from the vSphere Web UI Cluster

To view Cisco HyperFlex HTML5 Status, Network Details, and Capacity Plugin Portlets from the vSphere web UI perform the following steps:

1. Access **Hosts & Clusters** either from the vSphere Menu or the Shortcut link.
2. Click on the cluster name and select the **Summary** tab. The summary page appears.
3. Scroll down to locate the Cisco HyperFlex Plugin Portlets

To view Cisco HyperFlex HTML5 Storage Capacity and Performance Portlets from the vSphere web UI perform the following steps:

1. Access the **Hosts & Clusters** either from the vSphere Menu or the Shortcut link.
2. Click on the cluster name and select the **Monitor** tab. The Monitor page appears.
3. Scroll down the Monitor page left navigation panel and locate **Cisco HyperFlex**.
4. Click on **Storage Capacity** or **Performance** to display the related HyperFlex HTML5 Plugin chart.

Related Topics

[View the HX Cluster Summary](#), on page 272

[View Cluster and Datastore Performance Charts](#)

vCenter Server Action at Virtual Machine Level

Access the **Hosts & Clusters** either from the vSphere Menu or the Shortcut link.

- Right click on the virtual machine. Select **Cisco HyperFlex > ReadyClones**.
 - The **Ready Clones** window appears. Complete the General Settings Fields:
 - **Number of Clones**- Valid entry 1-256
 - **Customization specifications** - If configured, select from the list
 - **Resource Pool** - If configured, select from the list

Complete the Clone Name Fields:

- **Power on VMs after cloning**- Check box
- **Name of VM Prefix** - Type VM prefix
- **Starting clone number**- The default is 1
- **Increment clone number by** - The default is 1
- **Use same name for Guest Name**- Uncheck to provide a guest name

Click **OK** to complete the action. When the VM clone is complete, it will appear in the VM list.

- Right click on the virtual machine. Select **Cisco HyperFlex > Snapshot Now**.

The **Take VM Native Snapshot** window appears. Complete the following fields:

- **Name**- Snapshot name
- **Description** - Description of the snapshot
- **Quiesce guest file system** - Check box.

Click **OK** to create a VM Snapshot. You will see the snapshot task active in the background. After the Snapshot is complete it will be listed in the Snapshot Manager.

vCenter Server Action at the Storage Level

From the Datastore level users have the ability to edit or delete existing datastores.

- Access the **Datastores** either from the vSphere Menu or the Shortcut link.
- Right click on the datastore name.
 - Click **Cisco HyperFlex** and select **Edit Datastore** to make the Edit Datastore window appear.
 1. Edit the Datastore details.
 2. Click **OK** to save your changes.
 - Click **Cisco HyperFlex** and select **Delete Datastore** to make the Delete Datastore window appear.



Note Remove any VMs created or registered to the datastore and unmount the datastore before starting the Delete Datastore action.

1. Click the **Delete** button.

The Delete Datastore window appears with the confirmation question "Do you want to delete the datastore?"

2. Click **OK** to continue with the delete action, or **Cancel** to exit the Delete Datastore window.

Unmount Datastores from the vSphere Interface



Note Remove any VMs created or registered to the datastore before starting the Unmount Datastore action.

- Right click on the name.
 - Select **Unmount Datastore**. The Unmount Datastore window appears.
 1. Use the checkbox to select the node(s) you want to unmount.

2. Click **OK** to save your changes.

Related Topics

[Mount or Unmount Datastore](#), on page 279

Cisco HyperFlex Flash Plugin for VMware vCenter

Cisco HyperFlex Flash plugin for VMware vCenter, enables virtualization administrator to manage and monitor the Cisco HyperFlex physical infrastructure.

Cisco HyperFlex Flash Prerequisites

The following hardware and software prerequisites apply to the Cisco HyperFlex Flash Plugin:



Note

The Cisco HyperFlex Flash Plugin is the original plugin. It is recommended that users upgrade to the [Cisco HyperFlex HTML5 Plugin for VMware vCenter](#), on page 265.

- Cisco HyperFlex Flash plugin for VMware vCenter support was introduced in Cisco HX Release 4.0(2a) on vCenter 6.7.
- Browser compatibility: The Cisco HyperFlex Flash Plugin works with all browsers where vSphere Client and HyperFlex Connect work.

Launching HyperFlex Menu Options from vSphere

You can launch HyperFlex Connect from the vSphere Client UI in a new browser tab and perform management actions in the HyperFlex UI.

- To launch HyperFlex Connect Ready Clone in a new browser tab; in the vSphere client at Virtual Machine level, select **Actions** > **Cisco HyperFlex Systems** > **Ready Clone**
- To launch HyperFlex Connect Dashboard in a new browser tab; in the vSphere client at HX Cluster level, select **Actions** > **Cisco HyperFlex Systems** > **Summary**

Checking Cluster Rebalance and Self-Healing Status from the Flash Plugin

The storage cluster is rebalanced on a regular schedule and when the amount of available storage in the cluster changes. A rebalance is also triggered when there is a change in the amount of available storage. This is an automatic self-healing function.



Important

Rebalance typically occurs only when a single disk usage exceeds 50% or cluster aggregate disk usage is greater than 50%.

You can check rebalance status through the HX Data Platform plugin or through the storage controller VM command line. For directions on using the VM command line, see [Checking Cluster Rebalance and Self-Healing Status, on page 52](#).

Check the rebalance status through HX Data Platform plugin.

- a) From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Summary**.

The **Status** portlet lists the **Self-healing status**.

- b) Expand the **Resiliency Status** to see the **Self-healing status** section. The Self-healing status field lists the rebalance activity or N/A, when rebalance is not currently active.

The **Recent Tasks** tab in the HX Data Platform plugin displays a status message.

Setting a Disk Beacon through the HX Flash Plugin

Beaconing is a method of turning on an LED to assist in locating and identifying a node (host) and a disk. Nodes have the beacon LED in the front near the power button and in the back. Disks have the beacon LED on the front face.

You set a node beacon through Cisco UCS Manager. You set a disk beacon through the Cisco HX Data Platform plugin or HX Connect user interface.

Turn on and off a disk beacon using the HX Data Platform Plugin.

- a) From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Manage**.
- b) From **Manage**, select **Cluster > cluster > host > Disks > disk**.
- c) Locate the physical location of the object and turn on the beacon.

From **Actions** drop-down list, select **Beacon ON**.

- d) After you locate the disk, turn off the beacon.

From **Actions** drop-down list, select **Beacon OFF**

The beacon LED for all the disks on the selected node are toggled, except Housekeeping SSDs and cache NVMe SSDs. Housekeeping SSDs or cache NVMe SSDs do not have functioning LED beacons.
