

Cisco HyperFlex Data Platform Administration Guide, Release 5.5

First Published: 2023-08-22

Last Modified: 2025-01-24

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

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)

© 2023–2025 Cisco Systems, Inc. All rights reserved.



CONTENTS

Full Cisco Trademarks with Software License ?

PREFACE **Communications, Services, Bias-free Language, and Additional Information** xv

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 10

HX Data Platform Cluster Tolerated Failures 10

Data Replication Factor Settings 11

Cluster Access Policy 12

Responses to Storage Cluster Node Failures 12

HX Data Platform Ready Clones Overview 14

HX Native Snapshots Overview 15

CHAPTER 3 **Logging in to HX Data Platform Interfaces** 17

HyperFlex Cluster Interfaces Overview 17

Guidelines for HX Data Platform Login Credentials 18
HX Data Platform Names, Passwords, and Characters 19
AAA Authentication REST API 22
Logging into HX Connect 22
Logging into the Controller VM (hxcli) Command Line 23
Changing Storage Controller Password 25
Logging Into Cisco HX Data Platform Installer 25
Recovering the root password for the SCVM 26
Recovering the admin password for the SCVM 27
Accessing the HX Data Platform REST APIs 28
Secure Admin Shell 29
Guidelines and Limitations 29
Information About Consent Token 29
Diag User Overview 30

CHAPTER 4 **Monitoring HX Storage Clusters** 33
Monitoring HyperFlex Clusters 33
License Compliance and Feature Functionality 33
Monitoring HyperFlex Clusters with HX Connect 34
Dashboard Page 34
Activity Page 36
System Information Overview Page 38
Nodes Page 42
Disks Page 43
Audit Logging with HX Connect 45
Enabling Audit Logging 46
Configuring the Remote Syslog Server 47
Disabling Audit Logging 48
Deleting Audit Logging Server Configuration 49

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
Moving the Storage Cluster from a Current vCenter Server to a New VCenter Server 54
Unregistering a Storage Cluster from a vCenter Cluster 55
Unregistering and Removing EAM Extensions 55
Removing HX Data Platform Files from the vSphere Client 57
Verifying HX Cluster is Unregistered from vCenter 58
Registering a Storage Cluster with a New vCenter Cluster 58
Renaming Clusters 59
Replacing Self-Signed Certificate 60
Replacing Self-Signed Certificate with External CA Certificate on a vCenter Server 60
Replacing Self-Signed Certificate with External CA Certificate on an ESXi Host 61
Reregistering a HyperFlex cluster 62
Recreating a Self-Signed Certificate 62
Boost Mode 63
Configuring Boost Mode 63
Disabling Boost Mode 64
UEFI Secure Boot Mode 65
Enabling Secure Boot Mode 65
hx_edge.py Script Fails after ESXi Redeploy on Auto Secure Booted Node 66
Catalog Update 67
Catalog Update: HX Installer 68
Catalog Update: Cluster Creation using HX Installer 68
Catalog Update: Cluster Expansion using HX Installer 68
Catalog Update from the HX Installer Settings 69
Catalog Update: HX Connect 70
Cluster Catalog Upgrade using HX Connect 70
Catalog Update: Intersight 70
Catalog Upgrade using Intersight 70

CHAPTER 6 **Preparing for HX Storage Cluster Maintenance** 71
Storage Cluster Maintenance Operations Overview 71
Serial vs. Parallel Operations 73
Checking Cluster Status 73
Setting a Beacon 73

Verify vMotion Configuration for HX Cluster 74
Maintenance Modes for Storage Cluster Nodes 75
Entering Cisco HyperFlex Maintenance Mode 76
Exiting HXDP Maintenance Mode 77
Creating a Backup Operation 78
Shut Down and Power Off the Cisco HX Storage Cluster 82
Power On and Start Up the Cisco HX Storage Cluster 84
Restoring the Configuration for a Fabric Interconnect 86
Configure PCI Passthrough After Changing vNIC or vHBAs 88

CHAPTER 7 **Managing Encryption** 91
SED Encryption 91
Self-Encrypting Drives Overview 91
Verify if the HyperFlex Cluster Is Encryption Capable 91
Configuring Local Encryption Key 92
Modifying Local Encryption Key 93
Disabling Local Encryption Key 93
Secure Erase an Encrypted Disk 94
Remote Key Management 94
Configuring Remote Encryption Key 95
Generating Certificate Signing Requests 95
Configuring a Key Management Server Using CSRs (Certificate Signing Requests) 97
Generating Self-Signed Certificates 98
Configuring a key management server using SSCs (Self-Signed Certificates) 99
Restart Encryption 100
HyperFlex Software Encryption 100
Enabling HyperFlex Software Encryption Workflow 100
HyperFlex Software Encryption Guidelines and Limitations 101
Install HX Software Encryption Package: Clusters with 1 - 12 Nodes 101
Install HX Software Encryption Package: Clusters with 13+ Nodes 102
Backup Encryption Key of HyperFlex Software Encryption 103
Secure Disk Erase for HyperFlex Software Encryption 104

CHAPTER 8 **Managing Datastores** 107

Managing Datastores 107
Adding Datastores 108
Editing Datastores 109
Unmounting Datastores 110
Deleting Datastores 111
Encryption Support for Datastores 111
Recovering from Partially Unmounted Datastores 112

CHAPTER 9 **Managing Disks** 115
Managing Disks in the Cluster 115
Disk Requirements 116
Replacing SSDs 118
Replacing NVMe SSDs 119
Hot Swap NVME Drives in M5 and M6 Servers 120
Replacing Housekeeping SSDs for Cisco HX Release 5.0(2b) and Later 121
Replacing Self Encrypted Drives (SEDs) 123
Replacing or Adding Hard Disk Drives 125

CHAPTER 10 **Managing Nodes** 127
Managing Nodes 127
Identify Node Maintenance Methods 129
Searching by DNS Address or Host Name 131
Changing ESXi Host Root Password 132
Reinstalling Node Software 133
Changing Node Identification Form in vCenter Cluster from IP to FQDN 134
Replacing Node Components 135
Removing a Node 137
Preparing to Remove a Node 137
Removing a Node from an Online Storage Cluster 140
Removing a Node from an Offline Storage Cluster 143
Removing a Compute Node 146
Reuse a Previously Removed Node Within the Same Cluster 146

CHAPTER 11 **Expand Cisco HyperFlex System Clusters** 149

Cluster Expansion Guidelines 149
ESXi Installation Guidelines 150
Prerequisites When Expanding M5/M6 Clusters 151
Mixed Cluster Expansion Guidelines - Cisco HX Release 5.5(x) 151
Steps During Mixed Cluster Expansion 152
Prerequisites for Adding a Converged Node 152
Preparing a Converged Node 153
Adding a Converged Node to an Existing Cluster 153
Prerequisites for Adding a Compute-Only Node 158
Preparing a Compute-Only Node 159
Verify the HX Data Platform Installer 160
Apply an HX Profile on a Compute-only Node Using UCS Manager 160
Install VMware ESXi on Compute Nodes 160
Adding a Compute-Only Node to an Existing Cluster 162
Resolving Failure of Cluster Expansion 166
Logical Availability Zones 167

CHAPTER 12 **Managing HX Controller VMs** 171
Managing Storage Controller VMs 171
Powering On or Off Storage Controller VMs 171
Disabling HA VM Monitoring in HX Controller VMs 172

CHAPTER 13 **Managing Ready Clones** 175
HX Data Platform Ready Clones Overview 175
Benefits of HX Data Platform Ready Clones 175
Supported Base VMs 176
Ready Clone Requirements 176
Ready Clone Best Practices 177
Creating Ready Clones Using HX Connect 177
Creating Ready Clones Using the HX Data Platform Plug-In 179
Prepare to Customize HX Data Platform Ready Clones 180
Creating a Customization Specification for Linux in the vSphere Web Client 181
Create a Customization Specification for Windows in the vSphere Web Client 181
Configuring Ready Clones Using Customized Specifications 182

Managing Virtual Machine Networking 182

CHAPTER 14 **Managing HX Native Snapshots** 183

HX Native Snapshots Overview 183

Benefits of HX Native Snapshots 184

HX Native Snapshot Considerations 185

HX Native Snapshots Best Practices 188

HX Native Snapshot Time Zones 189

Creating HX Native Snapshots 190

HX Native Snapshots using ESXi 7.0 U2 191

Scheduling HX Native Snapshots Overview 191

Scheduling HX Native Snapshots 192

Setting the Frequency of HX Native Scheduled Snapshots 193

Deleting HX Native Snapshot Schedules 193

Reverting to an HX Native Snapshot 194

Deleting HX Native Snapshots 195

CHAPTER 15 **Managing Virtual Machine Disaster Recovery** 197

HX Disaster Recovery Overview 197

Replication and Disaster Recovery Requirements and Considerations 198

Admin Role Requirements 199

Networking Requirements 199

Cluster Requirements 203

Replication Network and Pairing Requirements 205

Replication and Disaster Recovery Virtual Machine Considerations 207

Storage Replication Adapter Overview 209

Data Protection Terms 210

Best Practices for Data Protection and Disaster Recovery 211

Protecting Virtual Machines Overview 212

Data Protection Workflow 213

Configuring the Replication Network in HX Connect 214

Test Local Replication Network 218

Editing the Replication Network 218

Replication Pair Overview 220

Creating a Replication Pair 220
Test Remote Replication Network 223
Editing a Mapped Datastore Replication Pair 224
Removing a Peer Cluster 225
Deleting a Replication Pair 225
Creating a Protection Group 227
Quiescence Overview 228
Editing Protection Groups 229
Deleting Protection Groups 230
Protecting Virtual Machines with an Existing Protection Group 230
Protecting Virtual Machines with a New Protection Group 231
Protecting Individual Virtual Machines 233
Unprotecting Virtual Machines 234
Disaster Recovery Overview 235
Configuring the Recovery Settings 235
Compatibility for Disaster Recovery Operations 237
Testing Virtual Machine Recovery 237
Recovering Virtual Machines 239
Recovering Virtual Machines in Protection Groups 241
Planned Migration 241
Planned Migration for a Single vCenter Deployment 242
Migrating Virtual Machines in Protection Groups 243
Disaster Recovery and Re-protect 243
Protecting Virtual Machines After Disaster 245
Removing Protection from an Auto-Protected Cluster VM 246
Replication Maintenance Overview 247
Pausing Replication 247
Resuming Replication 248
Replication Page 248
Local Virtual Machines Page 252
Remote Virtual Machines Page 255
Prepare to Protect Virtual Machines Alert 257
Configure or Edit Replication Network Dialog Box 258
Prepare Group Recovery Dialog Box 262

Recover VM on This Cluster Dialog Box 262
Test Recovery Parameters Dialog Box 262
Protected Virtual Machines Tab 263
Protection Groups 266
Replication Pairs Tab 268
Recovery Settings Dialog Box 276

CHAPTER 16 **Managing Storage Replication Adapter** 279
Storage Replication Adapter Overview 279

CHAPTER 17 **Managing Users** 281
Managing Cisco HyperFlex Users Overview 281
User Management Terms 282
Audit Logs for AAA Accounting 283
Creating Cisco HX Data Platform RBAC Users 283
Assigning Users Privileges 284

CHAPTER 18 **Managing iSCSI** 285
HyperFlex iSCSI Target Service Overview and Supported Use Cases 285
HyperFlex iSCSI Best Practices 286
iSCSI Configuration Overview 286
iSCSI Scale and Support 286
iSCSI Network Page 287
Creating an iSCSI Network 287
Editing an iSCSI Network 289
Deleting an iSCSI Network 289
iSCSI Initiator Group 289
Creating an iSCSI Initiator Group 290
Editing an iSCSI Initiator Group 291
Deleting an iSCSI Initiator Group 291
Linking iSCSI Initiator Group with Targets 291
Unlinking an iSCSI Initiator Group 292
iSCSI Target Page 292
Creating an iSCSI Target 293

Editing an iSCSI Target 294
Deleting an iSCSI Target 294
Linking iSCSI Targets 295
Unlinking an iSCSI Target 295
iSCSI LUN Page 296
Creating an iSCSI LUN 296
Editing an iSCSI LUN 297
Deleting an iSCSI LUN 298
Configuring an iSCSI Initiator (Windows) 298
Configuring an iSCSI Initiator (Linux) 299
Cloning an iSCSI LUN 299
Limitations for the HX Windows Agent 300
Prerequisites for the HX Windows Agent 301
Installing HX Windows Agent for iSCSI Clone LUN 301
Installing HX Windows Agent (with Pre-Installed Dependencies) for iSCSI Clone LUN 302
Uninstalling HX Windows Agent for iSCSI Clone LUN 303
iSCSI HX Windows Agent Logs 303
Editing the Location of Service Logs 304
Accessing Cloned LUN(s) on Destination Target 305

CHAPTER 19 **Cisco HyperFlex HTML Plugin for VMware vCenter** 307
Cisco HyperFlex vCenter Plugin 307
Cisco HyperFlex HTML5 Plugin for VMware vCenter 308
Cisco HyperFlex HTML5 Plugin Prerequisites 308
Install and Register the vCenter HTML5 Plugin 308
Verifying the Cisco HyperFlex HTML5 Plugin Installation from the vSphere Client 310
Uninstalling the HyperFlex HTML5 Plugin 311
Upgrading the HTML5 Plugin 311
Using the Cisco HyperFlex HTML5 Plugin 312
Navigating the HTML5 plugin 313
Cluster Management 315
Managing Users and Access to HX Clusters 315
Discover the Registered HX Cluster 315
Rename Cluster 316

View the HX Cluster Summary 316
View Cluster and Datastore Performance Charts 322
Disks 324
Nodes 325
Network 327
iSCSI 330
HX Datastore Management 335
VMs 340
Events 343
Alarms 344
Tasks 346
vCenter: HyperFlex Plugin Embedded Actions 347
vCenter Server Actions at the Host and Cluster Level 347
Create New Datastore 347
Enter or Exit Maintenance Mode 349
View HTML5 Plugin Portlets from the Summary Tab 350
View HTML5 Plugin Portlets from the Monitor Tab 351
View iSCSI and Datastore Summary from the Configure Tab 352
vCenter Server Actions at Virtual Machine Level 354
Snapshot Now 354
ReadyClones 355
Schedule Snapshot 357
vCenter Server Actions at the Storage Level 359
Edit Datastore 359
Delete Datastore 360
Cisco HyperFlex Remote Plugin for VMware vCenter 360
Install, Register and Upgrade the Remote Plugin 362
Remote Plugin Installation and Registration 362
Uninstalling and Unregistering the HyperFlex Remote Plugin from vCenter 364
Upgrade the Remote Plugin Application 3.0.0 using CLI 365
Encryption Support 365
Remote Plugin Encryption Support 365
Generate Support Bundles 366
Plugin Support Bundle Generation 366

APPENDIX A **Appendix** 367

Creating VLANs for HX Servers 367

Creating MAC Address Pools 368

Configure the vSwitches 370

Migrating vMotion Networks to Virtual Distributed Switches (VDS) or Cisco Nexus 1000v (N1Kv) 371



Communications, Services, Bias-free Language, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business impact you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco Marketplace.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Documentation Feedback

To provide feedback about Cisco technical documentation, use the feedback form available in the right pane of every online document.

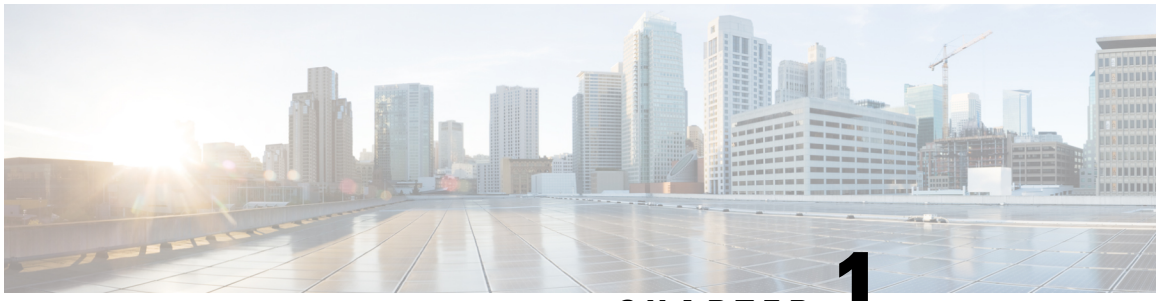
Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

Bias-Free Language

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

Communications, Services, Bias-free Language, and Additional Information



CHAPTER 1

New and Changed Information for this Release

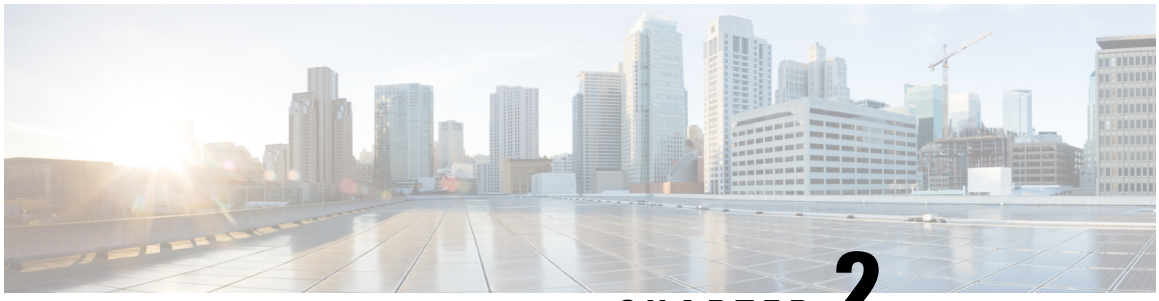
• New and Changed Information for this Release, on page 1

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
Updated vCLS modes	Updated Shut Down and Start Up procedure for using HX shared storage	5.5(1a)	Shut Down and Power Off the Cisco HX Storage Cluster, on page 82 Power On and Start Up the Cisco HX Storage Cluster, on page 84
HX Remote Plugin for VMware vCenter	Introduced the HTML remote plugin	5.5(1a)	Cisco HyperFlex Remote Plugin for VMware vCenter, on page 360
ESXi Support update	Starting with 5.5(1a) ESXi versions starting 6.5, 6.7, and 7.0.U1 are not supported.	5.5(1a)	HX Native Snapshots Overview, on page 15
New guide for HyperFlex Release 5.5(1a)	-	5.5(1a)	This guide.

New and Changed Information for this Release



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 12
- HX Data Platform Ready Clones Overview, on page 14
- HX Native Snapshots Overview, on page 15

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 `hxcli` 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 *hxcli* 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 *hxcli* 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 - TUUS/TAS) * 100
```

```
= (1 - 100GB / 200GB) *100
```

```
= 50%
```

3. Compression Savings

```
= (1 - TUB/TUUS) * 100
```

```
= (1 - 25GB / 100GB) * 100
```

```
= 75%
```

4. Total savings calculated

```
= (1 - TUB/TAS) * 100
```

```
= (1 - 25GB / 200GB) * 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:

(((<capacity disk size in GB> * 10^9) / 1024^3) * <number of capacity disks per node> * <number of HyperFlex nodes> * 0.92) / 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-M6SX 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 The event and alarm details provided on vCenter and HX Connect are not always a 1:1 relationship. When reviewing messages in HX Connect, it is a best practice to also review the events and tasks in vCenter.

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 `hxcli` 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 indicated 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. Production clusters should always use RF3. RF2 should be reserved for use in labs and demos.



Important Production clusters should be set to Data Replication Factor 3.

- **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 4.5(x) and later

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

Table 2: Cluster State in 3 - 4 Node Clusters with Number of Failed Nodes HX Release 4.5(x) and later.

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
3	Lenient	2	--
3	Strict	1	2
2	Lenient	1	--
2	Strict	--	1

Data Replication Factor Settings



Important 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.

Procedure

Choose a Data Replication Factor. The choices are:

- Data Replication Factor 3 — (**Recommended Usage:**All production environments) 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.
- Data Replication Factor 2 — (**Recommended Usage:**Non production labs and demos) 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.

Procedure

Review the table and perform the appropriate action.

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.
3 nodes	2	Two or more disks on two nodes are blocklisted or failed.	**a.** If one cache SSD fails, the storage cluster does not automatically heal. **b.** If one HDD fails or is removed, the disk is blocklisted immediately. The storage cluster automatically begins healing within a minute. **c.** 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.	If the node does not recover in two hours, the storage cluster starts healing by rebalancing data on the remaining nodes. To recover the failed node immediately and fully restore the storage cluster: **a.** Check that the node is powered on and restart it if possible. You might need to replace the node. **b.** Rebalance the cluster.
4 nodes	2	Two or more disks on two nodes.	If two SSDs fail, the storage cluster does not automatically heal. If the disk does not recover in one minute, the storage cluster starts healing by rebalancing data on the remaining nodes.

Cluster Size	Number of Simultaneous Failures	Entity Failed	Maintenance Action to Take
5+ nodes	2	Up to two nodes.	If the node does not recover in two hours, the storage cluster starts healing by rebalancing data on the remaining nodes. To recover the failed node immediately and fully restore the storage cluster: **a.** Check that the node is powered on and restart it if possible. You might need to replace the node. **b.** Rebalance the cluster. If the storage cluster shuts down, see Troubleshooting, Two Nodes Fail Simultaneously Causes the Storage Cluster to Shutdown section.
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.
5+ nodes	2	One node and One or more disks on a different node.	If the disk does not recover in **one minute**, the storage cluster starts healing by rebalancing data on the remaining nodes. If the node does not recover in **two hours**, the storage cluster starts healing by rebalancing data on the remaining nodes. 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. To recover the failed node immediately and fully restore the storage cluster: **a.** Check that the node is powered on and restart it if possible. You might need to replace the node. **b.** Rebalance the cluster.

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.

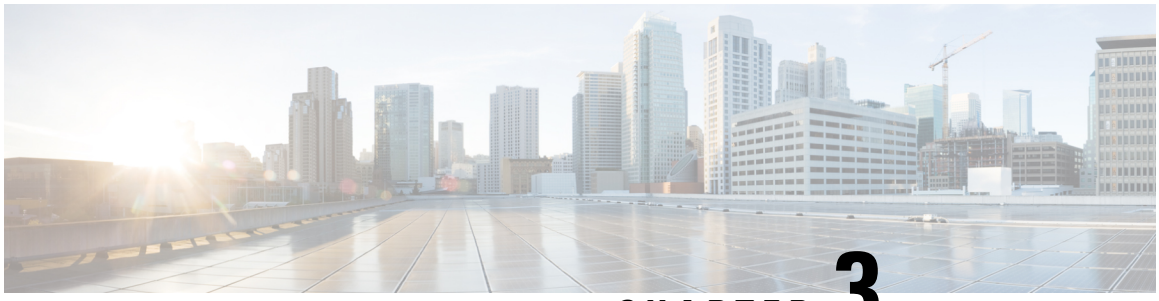
HX Native Snapshots Overview

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:

- Support for HX native Snapshot in the vSphere client plug-in for HTML 5 was introduced in plugin version 2.0.0. For more information, see Snapshot Now, on page 354.
- Support for Schedule Snapshot the vSphere client plug-in for HTML 5 was introduced in plugin version 2.1.0. For more information, see Schedule Snapshot, on page 357
- The vSphere “Manage Snapshots” function can revert to a specific HX native snapshot, or delete all snapshots.
- Cisco HyperFlex Connect can create on-demand and schedule HX native snapshots.
- The HyperFlex command line user interface can create HX native snapshots.
- HX REST APIs can create and delete HX native snapshots.
- Significant changes in Cisco HXDP Release 5.5(x) and later:
  - ESXi versions 6.5, 6.7 and 7.0 U1 are not supported.
  - VMware VAAI snapshot workflow is used instead of the Sentinel Snapshot Create workflow.

For additional information about VMware snapshots, see the "Overview of virtual machine snapshots in vSphere (KB 1015180)" on the VMware Customer Connect site.



CHAPTER 3

Logging in to HX Data Platform Interfaces

- HyperFlex Cluster Interfaces Overview, on page 17
- AAA Authentication REST API, on page 22
- Logging into HX Connect, on page 22
- Logging into the Controller VM (hxcli) Command Line, on page 23
- Logging Into Cisco HX Data Platform Installer, on page 25
- Recovering the root password for the SCVM, on page 26
- Recovering the admin password for the SCVM, on page 27
- Accessing the HX Data Platform REST APIs, on page 28
- Secure Admin Shell, on page 29
- Diag User Overview, on page 30

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.
- For the most accurate read of performance, refer to the HX Connect Cluster Level performance charts. The other charts may not present the complete picture due to the manner in which storage is distributed in HyperFlex and consumed in the VMs via the datastores.

Additional interfaces include:

- HX Data Platform Installer—Installing HX Data Platform, deploying and expanding HX Storage Cluster, and deploying stretched 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 Cisco123 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 `admin` or `root` users.	As specified during HX installation.	
Admin Shell		User defined during HX installation. Predefined `admin` user.	As specified during HX installation. Strong password required.	Must match across all nodes in storage cluster. Support for SSH to the secure admin shell is limited to the user admin. Use the `hxcli` command when changing the password after installation.

Component	Permission Level	Username	Password	Notes
vCenter	admin	administrator@vsphere.local default. SSO enabled. As configured, MYDOMAIN\name or name@mydomain.com	SSO enabled. As configured.	Read only users do not have access to HX Data Platform Plug-in.
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 `esxcli` 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

Usernames 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 `hxcli` command. Do not use another change password command, such as a Unix password command.

1. Log into the management controller VM.
2. Run the `hxcli` command.

   hxcli 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 RES 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.

Procedure

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 (hxcli) Command Line

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

- Modify commands—Require administrator level permissions. Examples:

```
hxcli cluster create
```

```
hxcli datastore create
```

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

```
hxcli <cmd> -help

hxcli cluster info

hxcli datastore info
```

To execute HX Data Platform `hxcli` commands, log into 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 `hxcli` commands.

You may log into 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: `hxcli cluster info`
  - Indirect commands—multi-layered commands that require live response through the command line. Example interactive command: `hxcli cluster reregister`

Procedure

Step 1 Locate a controller VM DNS Name.

a. Select a **VM** > **Summary** > **DNS Name**.

b. From vSphere Web Client **Home** > **VMs and Templates** > **vCenter server** > ***datacenter*** > **ESX Agents** > **VVM**.

c. Click through to the storage cluster list of controller VMs.

Step 2 From a browser, enter the DNS Name and `/cli` path.

a) Enter the path.

Example

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

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

b) 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.

a) Enter the `ssh` command string.
b) 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 admin@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 into 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.

Procedure

Step 1 Log into a storage controller VM.
Step 2 Change the Cisco HyperFlex storage controller password.

```
# hxcli security password set
```

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

Step 3 Type 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.

Procedure

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, or a Stretched cluster.
- **Cluster Expansion**—You can provide the data to add converged nodes and compute nodes to an existing standard storage cluster.

Recovering the root password for the SCVM

The only option to perform a root password recovery is using Linux single user mode.

Procedure

Contact Cisco TAC to complete this process.

Recovering the admin password for the SCVM

For HX 4.5(2c) and HX 5.0(2x) and later, you can recover the Storage Controller VM (SCVM) Admin password, by using SSH from the ESXi host with the RSA key and running the **recover-password** command. You will need to contact TAC to complete this process.

Before you begin

Contact TAC to support the Consent Token workflow.

Procedure

Step 1 Log in to the ESXi host using SSH.

Step 2 For ESXi 7.0 and 8.0, SSH to the Storage Controller VM for which the password has to be recovered, from ESXi using the **host_ecdsa_key** command.

Example:

```
[root@ucsblr625:~] ssh admin@`/opt/hxtools/bin/getstctlvmip.sh "Storage
 Controller Data Network"` -i /etc/ssh/ssh_host_ecdsa_key
The authenticity of host '10.21.24.89 (10.21.24.89)' can't be established.
ECDSA key fingerprint is SHA256:OkA9czzcL7I5fYbfLNtSI+D+Ng5dYp15qk/9C1cQzzk.
This key is not known by any other names

Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Warning: Permanently added '10.21.24.89' (ECDSA) to the list of known hosts.
 HyperFlex StorageController 5.5(1a)

This is a Restricted shell.
Type '?' or 'help' to get the list of allowed commands.
```

Step 3 Run the `recover-password` command. A prompt appears requesting Consent Token.

Note
Contact TAC to help provide the Consent Token.

a) Enter Option 1 to Generate Challenge.
b) Copy the Consent Token.
c) Enter Option 2 to Accept Response.
d) Enter the Constent Token.
e) Enter the new password for admin.
f) Re-enter the new password for admin.

Example

```
admin:~$ recover-password
Consent token is needed to reset password. Do you want to continue?(y/[n]):
y
-------------------------------------
1.  Generate Challenge
2.  Accept Response
3.  Exit
-------------------------------------
Enter Option:
```

```
1
Generating Challenge.....................................
Challenge String (Please copy everything between the asterisk lines exclusively):
**************************************BEGIN TOKEN*****************************************
2g9HLgAAAQEBAAQAAAABAgAEAAAAAQMACL7HPAX+PhhABAAQo9ijSGjCx+Kj+Nk1YrwKlQUABAAAAGQGAAlIeXBl
cmZsZXgHAAxIeXBlcmZsZXhfQ1QIAAlIWVBFUkZMRVgJACBhNzAxY2VhMGZlOGVjMDQ2NDllMGZhODVhODIyYTY2NA==
**************************************END TOKEN*******************************************
-----------------------------------
1.  Generate Challenge
2.  Accept Response
3.  Exit
-----------------------------------
Enter Option:
2
Starting background timer of 30 mins
Please input the response when you are ready:
Gu4aPQAAAQEBAAQAAAABAgAEAAAAAQMBYnlQdnRGY1NiNkhtOUlyanlDQVJic0ZXYnp3MVpzdm1pcVh3ZzJLSlZZSVl
yeXBydU9oejVQWkVXdlcvWWdFci8NCnBrVFVpS1d0dVRlczZ6TkdITXl0T3dNaFhaT2lrM3pKL1M5cDJqR0xxcGFOY1
Ruc05SVFNybCtQeGwvK1Z1b1gNCjBHYVVxcExXdUhtUUc0UG9ZU2FBL0lwelRFYzlaRmFNeUFmYUdkOThMSmliZnl2UF
c2d0tNY1FCM3lPWmRjU1ENCklGeWZJTVpKL1RWd1lOaERZT001dXQveHZxUU1HN1hTbjdXb2R4Wng2NVNqVktWK2lId
FMyZzdxZUIzc3R2TEgNCld1VWNYS3lWdFdOaXRiaHBvWUIwT1J0N2l3dHlrSkcyWldWbnk4KzZIUUNJbW9xdnFoSU91S
kk4aElsWWNNaUENCnlEbEpkQ0wwcHVObSswNVVyTWM0M1E9PQ==
Response Signature Verified successfully !
Response processed successfully.
Consent token workflow is successful, allowing password reset.
Enter the new password for admin:
Re-enter the new password for admin:
Changing password for admin...
Password changed successfully for user admin.
```

After using the **recover-password** command to change the password, passwords will no longer be synced on all nodes. You will need to use **hxcli security password set** to change and sync the password again on all nodes.

Step 4 To sync the password on all nodes, run the **hxcli security password set** command from any node, and enter the new password.

Example

```
admin:~$ hxcli security password set
Enter new password for user admin:
Re-enter new password for user admin:
admin:~$
```

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.

Procedure

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 into 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 a root capable user exists in an HX Release 4.0(x) or earlier cluster, delete it before starting an upgrade to HX Release 4.5(1a). If the root capable user is not removed, 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.

Diag User Overview

Starting with HX 5.0(2a), a new "diag" user for the HyperFlex command line interface, HX Shell, is introduced. This account is a local user account with escalated privileges designed for troubleshooting. Log into HX Shell remains restricted to the "admin" user account, and you must switch-user (su) to the "diag" user by providing the diag user password and passing a CAPTCHA test. When using the "diag" user, please note the following:

- Has more relaxed privileges than the admin user, but is more restricted than the root user
- Uses bash as the default shell, easing limitations of the lshell
- You can only access it by running '**su diag**' from admin shell. Direct ssh to diag is blocked.
- After entering the password for diag, a CAPTCHA test appears. You will need to enter the correct CAPTCHA to enter the diag shell.
- Write permission is limited to a pre-defined set of files for the diag user

Any command that would cause changes to the system software are blocked for the "diag" user. The default list of blocked commands include:

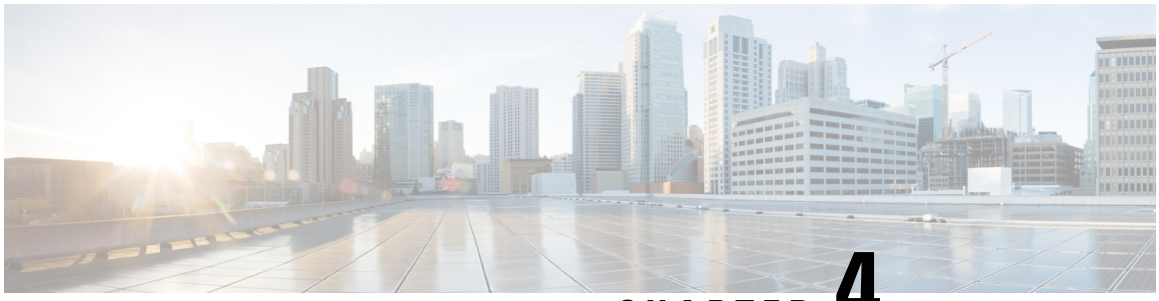
- **sudo**
- **apt-get**
- **li**
- **dpkg**
- **apt**
- **easy_install**
- **setfacl**
- **adduser**
- **deluser**
- **userdel**
- **groupadd**
- **groupdel**

- **addgroup**
- **delgroup**

The following is sample output for the **diag user** command.

```
This is a Restricted shell.
Type '?' or 'help' to get the list of allowed commands.
hxshell:~$ su diag
Password:
 ____          ____          _  _ _                        _
| ___|    _   |___ \        | \ | (_)_ __  ___    _    / |
|___ \   _| |_   __) | _____ | \| | | '_ \ / _ \ _| |_  | |
 ___) | |_  _| / __/ |_____| | |\ | | | | | __/ |_  _| | |
|____/    |_|  |_____|      |_| \_|_|_| |_|\___|  |_|  |_|


Enter the output of above expression: -1
Valid captcha
diag#
```



CHAPTER 4

Monitoring HX Storage Clusters

- Monitoring HyperFlex Clusters, on page 33
- License Compliance and Feature Functionality , on page 33
- Monitoring HyperFlex Clusters with HX Connect, on page 34
- Audit Logging with HX Connect, on page 45

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

License Compliance and Feature Functionality

Beginning with Cisco HXDP Release 5.0(2a), full feature functionality and configuration changes require a valid Cisco HyperFlex Software License. HX Connect users with expired or insufficient licenses at the end of the evaluation or the grace period after the license compliance date, view a prominent countdown banner that alerts the user to the license compliance need and provides a link to the license renewal page until the license expiration is remedied.

In the event a license passes both the license expiration date and the grace period countdown, the current configurations will operate as expected with limited information. Renewing the license allows a user to resume full feature functionality, and make configuration changes. For details and examples of the banners, see the License Compliance and Feature Functionality section of the Cisco HyperFlex Systems Ordering and Licensing Guide.

To review the Cisco End User Agreement (Cisco EULA), see https://www.cisco.com/c/en/us/about/legal/cloud-and-software/end_user_license_agreement.html

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 HXDP Maintenance Mode.

  See the Cisco HyperFlex Systems Troubleshooting Reference Guide for generating support bundles, Storage Cluster Maintenance Operations Overview, on page 71 for entering and exiting maintenance mode, and Setting a Beacon, on page 73 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 into Cisco HyperFlex Connect.

UI Element	Essential Information
Operational Status section	Provides the functional status of the HX storage cluster and application performance. Click **Information** (ⓘ) to access the HX storage cluster name and status data.

UI Element	Essential Information
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 VMware ESXi. Beginning with HXDP Release 5.0(2a), HX Connect users with expired or insufficient licenses will be unable to access certain features or have limited feature functionality, for more information see License Compliance and Feature Functionality .
Resiliency Health section	Provides the data health status and ability of the HX storage cluster to tolerate failures. Click **Information** (ⓘ) 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.
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.
VMs section	Displays the total number of VMs in the cluster as well as the breakdown of VMs by status (Powered on/off, Suspended, VMs with Snapshots and VMs with Snapshot Schedules).
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: `cvs`, `xls`, and `doc`. 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	Displays a list of recent tasks including the following details: • ID • Description • VM power on/off/suspend status • Task status: • **In Progress** • **Success** • **Failed** For failed VM-power operations, the **Existing State** and **Required State** fields are also included. • Date and time stamp • Progress bar An expanded list shows the task's step name and status. Click the circular icon to refresh the content and fetch recent activity. The page refreshes automatically every 2 minutes.
Recovery list	Displays progress of all recovery-related jobs (for example, migration, recovery, test recovery, re-protect) including the following details: • ID • Description • Task status: • **In Progress** • **Success** • **Failed** • Date and time stamp • Progress bar An expanded list shows the task's step name and status. Click the circular icon to refresh the content and fetch recent activity. The page refreshes automatically every 2 minutes.
Expand All / Collapse All button	Toggles the view of the job list to display top-level task information or task details. You can also expand and collapse individual tasks.

The following table specifies which Snapshot operations create an HX Task in the Activity Page.

Table 3: Snapshot Operations that create an HX Task in the Activity Page

Operation	HX Task Creation in Activity Page
Ready Clone from HX plugin	No HX task created.
Ready Clone from HX Connect	HX task added to the Activity page.
Scheduled Snapshot task creation from HX Plugin	No HX task created.
Scheduled Snapshot task creation from HX Connect	HX task added to the Activity page.
Snapshot creation from Schedule Snapshot	HX task added to the Activity page.
Snapshot now from HX Plugin	No HX task created.
Snapshot now from HX Connect	HX task added to the Activity page.

System Information Overview Page

Displays HX storage cluster system-related information, including node and disk data, and provides access to HXDP 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	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: **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. **Note** To register a cluster license, you can also choose **Register Now** from the **Actions** drop-down field.

UI Element	Essential Information

License section	• **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:
	Beginning with HXDP Release 5.0(2a), HX Connect users with expired or insufficient licenses will be unable to access certain features or have limited feature functionality, for more information see the License Compliance and Feature Functionality
	  • 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.
	Note
	To refresh license certificate or renew license authorization, choose the respective options from the **Actions** drop-down field.

HX storage cluster status field	Provides functional status of the HX storage cluster. • **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.

UI Element	Essential Information
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).

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

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.

UI Element	Essential Information
Node status	• **Online** • **Offline** • **In Maintenance** • **Healthy** • **Warning**
HXDP Version	HyperFlex Data Platform version installed on this cluster.
Type	• **Hyperconverged** • **Compute**
Hypervisor Status	• **Online** • **Offline** • **In Maintenance** • **In Progress**
Hypervisor Address	IP address for the management network to this HX storage cluster.
Disk Overview	Graphic representation of the number of disks in use for each node, the usage type and number of empty slots. **Note** A disk outline with a red icon indicates a disk that is not recognized and requires a Catalog Upgrade.

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	• System • Cache • Persistent

UI Element	Essential Information
Disk State	• **Claimed** • **Available** • **Ignored** • **Blocked** • **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 HXDP Maintenance Mode button	Select a node to access this button. Opens the **Confirm HXDP Maintenance Mode** dialog box.
Exit HXDP Maintenance Mode button	Select a node to access this button. After you complete any maintenance tasks, you must manually exit HXDP Maintenance Mode.
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.

UI Element	Essential Information
Hypervisor Status column	• **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	• **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	Number of disks in the node. Click the number to open the **Disks** page filtered by the selected node name.

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	• **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. • **Blocked**—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 blocked disk is currently being repaired. • **To Be Removed**—State when a disk is scheduled for RMA.	The following states can be ignored: • **Invalid** • **Normal** • **Removed**—State when an SED disk is removed after using the **Secure Erase** option. • **Time out** • **Unknown**
Encrypted column	• **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	• **Unknown** • **Rotational**—Hybrid drive • **Solid State**—SSD drive	

UI Element	Essential Information
Usage column	• **Unknown** • **Cache** • **Persistent**
Turn On Locator LED and **Turn Off Locator LED** radio buttons	Select a disk to access the radio buttons. Activates or deactivates a physical light, or beacon, on the host to help locate the disk.
(Optional) **Secure erase** button	This button in visible only if your HX storage cluster is encrypted using local-key encryption. Select a disk to access the button. 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.

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/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 sylog 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.

Procedure

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.
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.

UI Element	Essential Information
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_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_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_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_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_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_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.

Procedure

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.

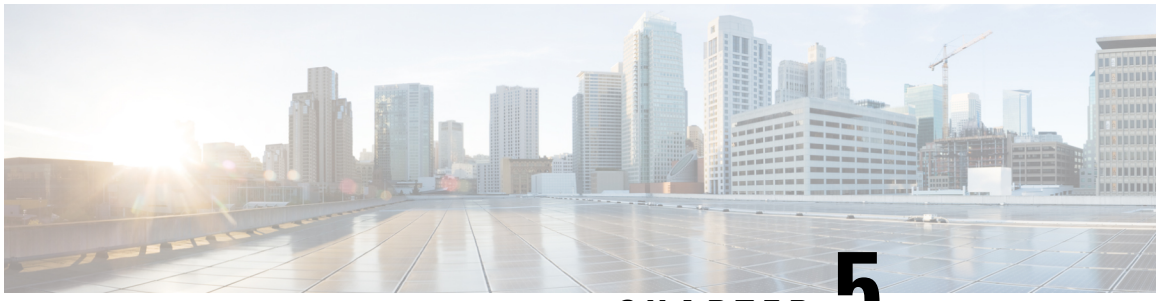
Procedure

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
- Unregistering a Storage Cluster from a vCenter Cluster, on page 55
- Renaming Clusters, on page 59
- Replacing Self-Signed Certificate, on page 60
- Boost Mode, on page 63
- UEFI Secure Boot Mode, on page 65
- hx_edge.py Script Fails after ESXi Redeploy on Auto Secure Booted Node, on page 66
- Catalog Update, on page 67

Changing the Cluster Access Policy Level

Procedure

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.

Restriction The following workflow should only be performed by Cisco TAC. If you have the need to manually rebalance a cluster, contact TAC for assistance.

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.

Procedure

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.

Procedure

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) Log into 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`.

Procedure

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**.

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

Before you begin

- 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:

  ```
  Resiliency Health: 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.

Procedure

Step 1 From the current vCenter, delete the cluster.

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

Caution
Distributed Virtual Switch (DVS) users: Deleting a cluster when using a DVS in the cluster is not recommended.

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.

Procedure

Step 1 Complete the steps in Removing HX Data Platform Files from the vSphere Client, on page 57.

Step 2 Complete the steps in Verifying HX Cluster is Unregistered from vCenter, on page 58.

What to do next

Proceed to Registering a Storage Cluster with a New vCenter Cluster, on page 58.

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.

Procedure

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:

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

Step 2 To unregister the storage cluster extension: Log into 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.sysmgmt.cluster_domain_id` and `com.springpath.sysmgmt.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: Log into 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. For additional information, see the *VMware KB, Stopping, starting, or restarting VMware vCenter Server Appliance 6.0 services (2109887)*article on the VMwear Customer connect site.

Removing HX Data Platform Files from the vSphere Client

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

Procedure

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

Linux vCenter

a) Log into 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) Log into 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: Removing HX Data Platform Files from the vSphere Client, on page 57

Procedure

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

Before attempting to register the HyperFlex cluster to vCenter, you must disable ESXi Lockdown mode on all ESXi hosts, and ensure SSH service is enabled and running.

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.

Procedure

Step 1 Log into a controller VM.
Step 2 Run the `stcli cluster reregister` command.

Example:

```
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.
--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.

Step 4 (Optional) Once registration is successful, re-enable ESXi Lockdown mode if you disabled it prior to registering the HyperFlex cluster to vCenter.

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.

Procedure

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.

Replacing Self-Signed Certificate

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

Procedure

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:

a) Select the vCenter server that manages the hosts and click **Settings**.
b) Click **Advanced Settings**, and click **Edit**.
c) In the **Filter** box, enter **certmgmt** to display only certificate management keys.
d) Change the value of **vpxd.certmgmt.mode** to **custom** and click **OK**.
e) 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 **vmsa** 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 rui.crt and rui.key files with the backed up previous key and certificate.
    3. Restart the hostd and vpxa 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 **vmsa** 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

Procedure

Step 1 Generate the host certificate (rui.crt) and key (rui.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 rui.key and rui.crt files.

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

```
hxcli license register
```



Note Before attempting to register the HyperFlex cluster to vCenter, you must disable ESXi Lockdown mode on all ESXi hosts, and ensure SSH service is enabled and running. Once registration is successful, you may re-enable Lockdown mode.

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 into 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. For more information about the CPUs supported by Boost Mode, see the Cisco HyperFlex Spec Sheets.

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 C245
  - All Flash C240
  - All Flash C225
  - All Flash C220
- Hypervisor: ESX only
- Boost Mode number of controller VM vCPUs:
  - All NVMe: 16
  - All Flash C245: 12
  - All Flash C240: 12
  - All Flash C225: 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**

Procedure

Step 1 From the vCenter, right-click one controller VM and **Shut Down Guest OS**.

Step 2 Increase the number controller VM vCPUs to 16 for all-NVMe, or 12 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.

Note
Boost Mode number of controller VM vCPUs:

- All NVMe: 16
- All Flash C245: 12
- All Flash C240: 12
- All Flash C225: 12
- All Flash C220:12

Step 3 Click **OK** to apply the configuration changes.

Step 4 Power up the controller VM.

Step 5 Log into 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:

Procedure

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 into 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:

- For HX Edge clusters, UEFI secure boot should only be enabled on HX Edge clusters running Cisco IMC version 4.1(2a) and later. If secure boot is enabled on earlier Cisco IMC versions, secure boot will need to be temporarily disabled during firmware updates.
- Support for Secure boot is available only for HyperFlex ESXi M5/M6 Server.
- 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.
- All the factory prepped M.2 RAID edge nodes run HXDP server firmware version 4.1(2a) or later. If a customer downgrades in field or retrofits existing setup and tries to bring up a cluster with M.2RAID nodes with HXDP server firmware version earlier than 4.1(2a), then the install may fail with the error `UEFI boot parameters cannot be configured for Legacy boot mode`. The HXDP server firmware must be upgraded to version 4.1(2a) or later and then re-try the install.

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.
- Startingwith 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.
- In the case that **Enable Secure Boot** workflow fails, then from the vCenter, confirm whether the host is still in **Maintenance Mode**. If so, then exit **Maintenance Mode** before retrying the **Enable Secure Boot** workflow.

Procedure

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, navigate 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.

hx_edge.py Script Fails after ESXi Redeploy on Auto Secure Booted Node

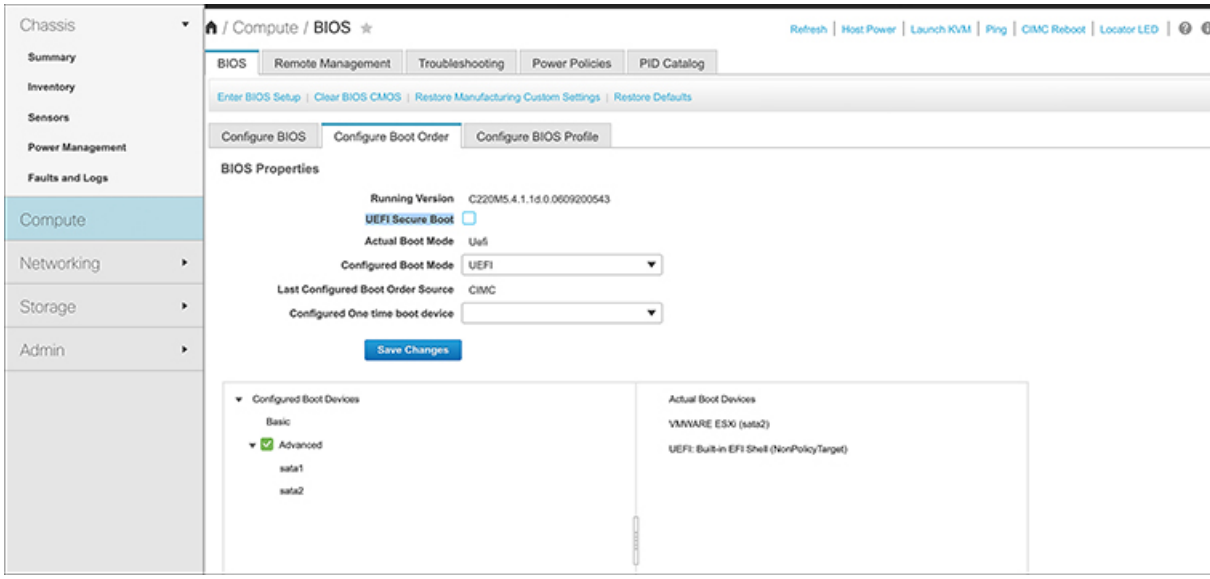
Description

Cisco UCS tools are not supported on HyperFlex Edge configuration with classic installer. The HX Edge configuration with classic installer uses the SOL for IP configuration and needs confirmation on whether the ESXi login console has been recieved.

To enable SOL on a Secure Boot node, perform the following steps:

Action

1. Disable boot security by unchecking the **UEFI Secure Boot** check box on the **Compute** > **Configure Boot Order** page.



2. Power cycle the server.
3. Install ESXi.
4. Run the `hx_edge.py` script.

Catalog Update

Compatibility Catalog Update feature was introduced in Cisco HyperFlex Release 4.5(1a) for ESXi.

Catalog Update provides the ability to update the catalog version across a cluster during cluster creation, expansion or hot adds of a newer model drives without needing to upgrade the HXDP version running on the cluster.

- Clearly identify drives that are unsupported by the current catalog.
- Reduces the overhead when adding a new drive model to a cluster node by eliminating the need to upgrade HXDP.
- Supported on HX Installer, HX Connect, and Intersight.
- Catalog is updated online and without impact to the running cluster.

Guidelines and Limitations

- Before adding a new drive, review the HyperFlex Release notes to confirm that the current HXDP version supports the new drive model.

- Catalog Update does not guarantee a drive is supported. Hardware issues and HXDP versions may contribute to a drive being unrecognized to HXDP.
- Do not use the Catalog Upgrade for drives that require HXDP tuning for custom settings, such as a higher drive capacity point; These require a full HXDP upgrade.
- Downgrading the catalog bundle to an earlier version is not supported.

Catalog Update: HX Installer

Catalog Update: Cluster Creation using HX Installer

Perform the following steps to upgrade the catalog during cluster creation using the HX VM based installer (OVA).

Before you begin

- Download the catalog bundle from CCO https://software.cisco.com/download/home/286305544/type/286305994/.

Procedure

Step 1 Log in to the HX Data Platform Installer.

Step 2 Follow the **Create Cluster** workflow for a standard cluster.

Step 3 On the Server Selection page, the installer validates the drive supportability and if unsupported drive(s) is found, the unsupported drive(s) are identified and the **Upgrade Catalog** button appears.

Step 4 Click the **Upgrade Catalog** button. The Upgrade Catalog window appears.

Note
The window displays the catalog version in use.

Step 5 Upload the Catalog file saved locally. Drag and drop the file into the target or click the target to browse to the file location. The upload operation is complete.

Step 6 Click **Upgrade** to complete the upgrade or **Close** to exit the Catalog Upgrade window.

The drive supportability check runs again after the catalog is upgraded. When all drives have compatible catalogs a green success banner appears.

Catalog Update: Cluster Expansion using HX Installer

When expanding a cluster, the Compatibility Catalog feature identifies if the installer-catalog is lower than the cluster-catalog and performs a drive supportability validation. Perform the following steps to upgrade the catalog during cluster expansion using the HX VM based installer (OVA).

Before you begin

- Download the catalog bundle from CCO https://software.cisco.com/download/home/286305544/type/286305994/.

Procedure

Step 1 Log into the HX Data Platform Installer.

Step 2 Follow the **Expand Cluster** workflow for a standard cluster.

Step 3 On the Server Selection page, the installer validates the drive supportability and if unsupported drive(s) is found, the unsupported drive(s) are identified and the **Upgrade Catalog** button appears.

Step 4 Click the **Upgrade Catalog** button. The Upgrade Catalog window appears.

Note
The window displays the catalog version in use.

Step 5 Upload the Catalog file saved locally. Drag and drop the file into the target or click the target to browse to the file location. The upload operation is complete.

Step 6 Click **Upgrade** to complete the upgrade or **Close** to exit the Catalog Upgrade window.

The drive supportability check runs again after the catalog is upgraded. When all drives have compatible catalogs a green success banner appears.

To view the current catalog version after the Catalog Upgrade is complete, navigate to the upgrade page in HX Connect for the running cluster.

Catalog Update from the HX Installer Settings

Perform the following steps to out-of-band catalog upgrade for the HX VM Installer (OVA):

Procedure

Step 1 Log into the HX Data Platform Installer.

Step 2 Click the **Settings** gear icon on any page.

Step 3 Click the **Upgrade Catalog** button. The Upgrade Catalog window appears.

Note
The window displays the catalog version in use.

Step 4 Upload the catalog file saved locally. Drag and drop the file into the target or click the target to browse to the file location. The upload operation is complete.

Step 5 Click **Upgrade** to complete the upgrade or **Close** to exit the Catalog Upgrade window.

To view the current catalog version after the Catalog Upgrade is complete, return to the upgrade catalog window by clicking the **Settings Icon** > **Upgrade Catalog**.

Catalog Update: HX Connect

Cluster Catalog Upgrade using HX Connect

When a new disk is not recognized in HX Connect it can be an indication that the catalog requires an update. Perform the following steps to upgrade the catalog using HX Connect.

Before you begin

- Download the catalog bundle from CCO https://software.cisco.com/download/home/286305544/type/286305994/.



Note The cluster catalog is automatically upgraded when you upgrade HXDP version. No manual update of catalog is needed when upgrading HXDP to a version that already contains an updated catalog.

Procedure

Step 1 Click on the **Upgrade** tab in HX Connect.

Step 2 Check the **HX Data Platform** box on the Select Upgrade tab.

Note
Combining a Catalog Upgrade with any other type of upgrade is not supported.

Step 3 Upload the locally saved Catalog file. Drag and drop the file into the target or click the target to browse to the file location. The Catalog file upload operation is complete.

Step 4 Click **Upgrade** to complete the upgrade.

a) To monitor progress of the upgrade tasks on an HX storage cluster, click the **Activity** page in HX Connect.

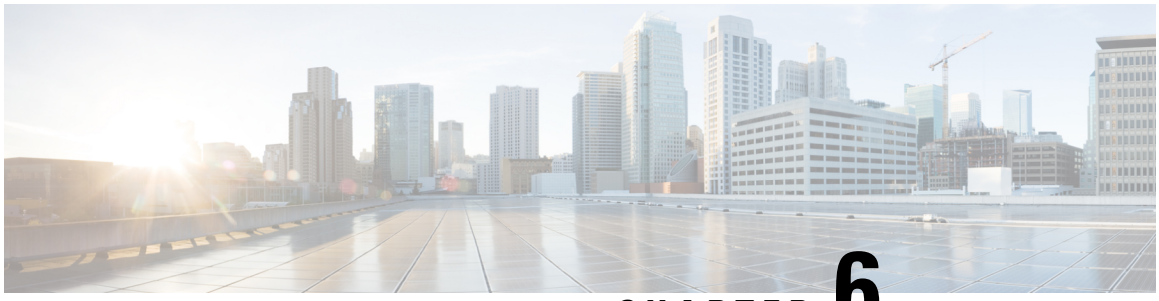
Step 5 Click on the **System Information** page and verify that all disks have been claimed by HXDP and are in use.

Catalog Update: Intersight

Catalog Upgrade using Intersight

Unlike the HX installer VM, the Intersight HX installer is kept up to date with the latest compatibility catalog automatically. Cisco releases updates to the Intersight HX Installer regularly and includes any catalog updates as part of that standard process.

Similarly, Intersight connected clusters are automatically updated to the latest catalog version without the need for manual download and upload through HX Connect. To receive these automatic updates, ensure the HyperFlex cluster is connected to Intersight.



CHAPTER 6

Preparing for HX Storage Cluster Maintenance

- Storage Cluster Maintenance Operations Overview, on page 71
- Serial vs. Parallel Operations, on page 73
- Checking Cluster Status, on page 73
- Setting a Beacon, on page 73
- Verify vMotion Configuration for HX Cluster, on page 74
- Maintenance Modes for Storage Cluster Nodes, on page 75
- Entering Cisco HyperFlex Maintenance Mode, on page 76
- Exiting HXDP Maintenance Mode, on page 77
- Creating a Backup Operation, on page 78
- Shut Down and Power Off the Cisco HX Storage Cluster, on page 82
- Power On and Start Up the Cisco HX Storage Cluster, on page 84
- Restoring the Configuration for a Fabric Interconnect, on page 86
- Configure PCI Passthrough After Changing vNIC or vHBAs, on page 88

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.

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 73 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 HXDP Maintenance Mode prior to performing a maintenance task on the host. The HXDP Maintenance Mode performs additional storage cluster specific steps compared to the vSphere provided Host Maintenance Mode.

Post Maintenance Tasks

After the maintenance task is completed, the nodes need to exit HXDP Maintenance Mode and the storage cluster needs to be healthy. 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 88.

Ensure the following:

- The ESX host is exited from HXDP 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

Procedure

Step 1 Log into 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.

```
# hxcli 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.

Procedure

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

a) From the UCS Manager left panel, select **Equipment** > **Servers** > ***server***.
b) From the UCS Manager central panel, select **General** > **Turn on Locator LED**.
c) 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.

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**

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

a) Log into HX Connect.
b) Select **System Information** > **Disks**.
c) 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:



Remember Some VMs not supported by vMotion should be shut down, since it will hold the nodes from going into maintenance mode.

1. 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.
2. 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.

3. 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.

- HXDP Maintenance Mode
- Host Maintenance Mode

HXDP Maintenance Mode

HXDP Maintenance Mode performs Cisco HX Data Platform specific functions in addition to the Host Maintenance Mode. Be sure to select HXDP Maintenance Mode and not Host 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.

HXDP Maintenance Mode Considerations

- Ensure that SSH is enabled in ESX on all the nodes in the storage cluster prior to using HXDP Maintenance Mode.
- When HXDP Maintenance Mode is entered to enable performing tasks on an ESX host, be sure to exit HXDP Maintenance Mode after the tasks on the ESX host are completed.

- HXDP 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 Host 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 HXDP Maintenance Mode, on page 77 for steps.

Host 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 `esxcli system maintenanceMode` command.

Entering Cisco HyperFlex Maintenance Mode

Using the Cisco HyperFlex (HX) Connect User Interface

1. Log into 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 HXDP Maintenance Mode**.
5. In the **Confirm HXDP Maintenance Mode** dialog box, click **Enter HXDP Maintenance Mode**.



Note After you complete any maintenance tasks, you must manually exit HXDP Maintenance Mode.

Using the vSphere Web Client

1. Log into 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 **HXDP Maintenance Mode** > **Enter HXDP Maintenance Mode**.

Using the Command-Line Interface

1. Log into the storage controller cluster command line as a user with root privileges.
2. Move the node into HXDP Maintenance Mode.
   a. Identify the node ID and IP address.

      ```
      # hxcli node list --summary
      ```

   b. Enter the node into HXDP Maintenance Mode.

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

      (see also `hxcli node maintenanceMode --help`)
3. Log into the ESXi command line of this node as a user with root privileges.
4. Verify that the node has entered HXDP 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 HXDP Maintenance Mode

Using the Cisco HyperFlex (HX) Connect User Interface

1. Log into 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 HXDP Maintenance Mode**.

Using the vSphere Web Client

1. Log into 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 **HXDP Maintenance Mode** > **Exit HXDP Maintenance Mode**.

Using the Command-Line Interface

1. Log into the storage controller cluster command line as a user with root privileges.
2. Exit the node out of HXDP Maintenance Mode.
   a. Identify the node ID and IP address.

      ```
      # hxcli node list --summary
      ```

   b. Exit the node out of HXDP Maintenance Mode.

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

      (see also `stcli node maintenanceMode --help`)
3. Log into the ESXi command line of this node as a user with root privileges.
4. Verify that the node has exited HXDP 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. Log into 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.

Procedure

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	This can be one of the following: • **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	The information saved in the backup configuration file. This can be one of the following: • **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. **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. It is also important that the bundles from which the backup was taken remain present in the Cisco UCS Manager and should not be deleted. • **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	This checkbox remains selected for **All Configuration** and **System Configuration** type of backup operation, and provides the following functionality: • **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. **Note** If this check box is not selected the identities will be reassigned and user labels will be lost after a restore. • **System Configuration**—The backup file preserves identities for Chassis, FEX, Rack Servers, and user labels for Chassis, FEX, Rack Servers, IOMs and Blade Servers. **Note** If this check box is not selected the identities will be reassigned and user labels will be lost after a restore. 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. **Note** If this check box is not selected the identities will be reassigned and user labels will be lost after a restore.
Location of the Backup File field	Where the backup file should be saved. This can be one of the following: • **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. Java-based Cisco UCS Manager GUI displays the **Filename** field with an associated **Browse** button that let you specify the name and location for the backup file. **Note** Once you click **OK**, the location cannot be changed. 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.

Name	Description
Protocol field	The protocol to use when communicating with the remote server. This can be one of the following: • **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	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. **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. **Note** All IP addresses must be IPv4. HyperFlex does not support IPv6 addresses.
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 into the remote server. This field does not apply if the protocol is TFTP.
Password field	The password for the remote server username. This field does not apply if the protocol is TFTP. 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.

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:

a) If the operation does not display in the **Properties** area, click the operation in the **Backup Operations** table.
b) 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.
- Prior to shutdown, verify that the HyperFlex cluster has one reachable external NTP and DNS resource configured that resides outside the HyperFlex.
- Perform both the Full-State and All Configuration type backups with the Preserve Identities attribute. See Creating a Backup Operation, on page 78.

Procedure

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 From the**vSphere Cluster Services** drop down list, select **General**. The **Edit vCLS Mode** dialog box appears.

Step 3 Select **Retreat Mode**.

Step 4 Click **OK**.

Step 5 Gracefully shut down the Cisco HX storage cluster.

a) 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.

```
# hxcli 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 6 On each storage cluster ESX host, shutdown the controller VM (`stCtlVM`).

Choose a method:

Using vCenter Shut Down Guest OS

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

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** > **Shut Down Guest OS.**.

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

Using vCenter Host 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 7 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 8 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 HXDP 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 84.

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 82.

Procedure

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 into 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 **`hxcli cluster start`** command.

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

Using vSphere Client

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

Using ESXi host command line

a) Log into a host.
b) Identify the VMID of the stCtlVM.

```
# vim-cmd vmsvc/getallvms
```

c) Using the VMID power on the controller VM.

```
# vim-cmd vmsvc/power.on VMID
```

d) 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.

a) SSH to any controller VM, run the command:

```
# hxcli about
```

b) 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.
c) 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.

```
# hxcli cluster start
```

Depending upon the maintenance or upgrade task performed while the HX cluster was shutdown, the nodes might be exited from HXDP Maintenance Mode or Host 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.

a) From any controller VM, run the command.

```
# hxcli cluster info
```

b) 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 From the **vSphere Cluster Services** drop down list, select **General**. The **Edit vCLS Mode** dialog box appears.

Step 11 Select **System Managed**.

Step 12 Click **OK**.

Step 13 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 14 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 4.5(1a) to restore a system running Release 4.5(2a).

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.

Procedure

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, 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
Server IP	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.
Backup File Path	The file path where the full state backup file is located, including the folder names and filename. **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.
User ID	The username the system should use to log into the remote server. This field does not apply if the protocol is TFTP.

Name	Description
Password	The password for the remote server username. This field does not apply if the protocol is TFTP.

Step 10 Click **Submit**.

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

The fabric interconnect logs into 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

Procedure

Step 1 Put the ESX host into HXDP 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 Log into 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.

a) From the DirectPath I/O Configuration page, select **Configure Passthrough**.
b) From the Mark devices for passthrough list, select the LSI card for the pass through.
c) 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.

a) 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**.

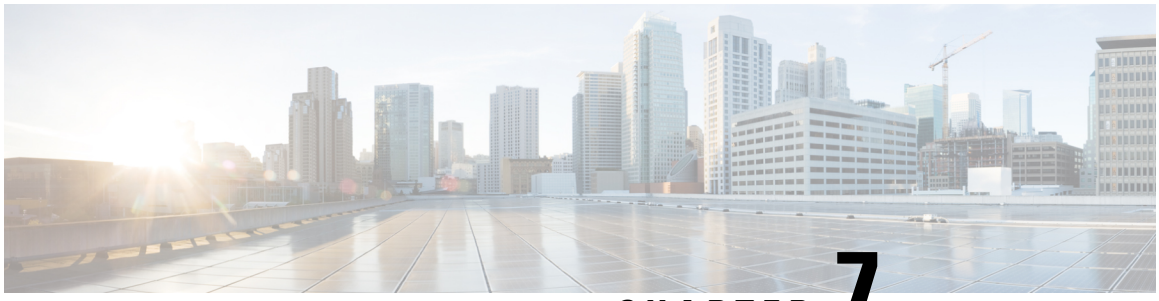
b) Locate and re-add the LSI Logic PCI device.

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 HXDP Maintenance Mode.

When the host is active again, the HX storage controller VM properly boots and rejoins the storage cluster.



CHAPTER 7

Managing Encryption

- SED Encryption, on page 91
- HyperFlex Software Encryption, on page 100

SED Encryption

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 into 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

Procedure

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. Enter an IP address or FQDN. <*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

Procedure

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, *10.193.211.120*.
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

Procedure

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. Enter an IP address or FQDN. <*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

Procedure

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

Procedure

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. Enter an IP address or FQDN. <*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

Procedure

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. Enter an IP address or FQDN. <*eng-fi12.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. Enter up to 32 characters.
Organization unit name field	The organizational unit. Enter up to 64 characters.
Locality field	The city or town in which the company requesting the certificate is headquartered. Enter up to 32 characters.
State field	The state or province in which the company requesting the certificate is headquartered. Enter up to 32 characters.
Country field	The country in which the company resides. 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.

Procedure

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. Enter an IP address or FQDN. <*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.

UI Element	Essential Information
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.

Generating Self-Signed Certificates

Procedure

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. Enter an IP address or FQDN. <*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. Enter up to 32 characters.
Organization unit name field	The organizational unit. Enter up to 64 characters.
Locality field	The city or town in which the company requesting the certificate is headquartered. Enter up to 32 characters.

UI Element	Essential Information
State field	The state or province in which the company requesting the certificate is headquartered. Enter up to 32 characters.
Country field	The country in which the company resides. 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.

Procedure

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. Enter an IP address or FQDN. <*eng-fi12.eng.storvisor.com*>
User name field	<*admin*> username
Password field	<*admin*> 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.
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

Procedure

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. Enter an IP address or FQDN. <*eng-fi12.eng.storvisor.com*>
User name field	<*admin*> username
Password field	<*admin*> password

HyperFlex Software Encryption

Enabling HyperFlex Software Encryption Workflow

The following table summarizes the enabling HyperFlex Software Encryption workflow:

Step	Description	Reference
1.	Download the HyperFlex Software Encryption package from Cisco License Central.	Cisco Software Central

Step	Description	Reference
2.	Log into the management CIP to install the package on all the controller VMs in the cluster.	Run the command `priv install-package`.
3.	Install the encryption package.	See Install HX Software Encryption Package: Clusters with 13+ Nodes, on page 102.
4.	Follow the enable procedure on Intersight.	Intersight HyperFlex Software Encryption
5.	Verify your cluster is encrypted.	Run the command `hxcli encryption info`.

Note If your cluster has VMware EVC enabled, make sure that the EVC baseline supports nodes with Advanced Encryption Standards New Instructions (AES-NI). If your current EVC baseline does not support AES-NI, change the EVC settings before enabling Software Encryption.

HyperFlex Software Encryption Guidelines and Limitations

Review these guidelines before enabling HyperFlex Software Encryption:

- HyperFlex Software Encryption can be enabled when all HyperFlex Cluster Nodes are running HXDP Release 5.0(1b) and later.
- Support for HyperFlex Stretch Clusters with HyperFlex Software Encryption was introduced in HXDP Release 5.0(2a).
- SED HyperFlex configurations are not supported with HyperFlex Software Encryption.
- HyperFlex Software Encryption is supported only with Vmware ESXi HyperFlex configurations.
- AES-NI enablement is required to install HyperFlex Software Encryption packages on the HX Cluster.
- HyperFlex Software Encryption can not be enabled on existing Datastores.
- HyperFlex Software Encryption can only be enabled on newly created Datastores.
- Once HyperFlex Software Encryption is enabled for a cluster/datastore, it cannot be disabled for the cluster or datastore.
- Once HyperFlex Software Encryption is enabled for a cluster, administrators can create either encrypted or non-encrypted datastores.

Install HX Software Encryption Package: Clusters with 1 - 12 Nodes

Before you begin

Download the HyperFlex Software Encryption package from Cisco License Central, see Cisco Software Central.

Note The HyperFlex Software Encryption package is licensed by its own software PID, which is in addition to the HyperFlex Data Platform and Intersight software licenses. For more information, refer to the Cisco HyperFlex Systems Ordering and Licensing Guide.

Procedure

Step 1 SFTP the encryption package to the HyperFlex cip node (the node holding the cluster management IP). Use the admin account for the **username/password** and a file transfer application, such as winscp. This should upload the package to the /tmp or /home/admin directory.

Step 2 To install the package on all available nodes of a cluster, SSH to the cip node and use the **priv install-package --cluster** option.

Example:

```
priv install-package --cluster --path /tmp/storfs-se-core_<latest version>_x86_64.deb.gz
```

Note
Make sure that all nodes are up and not in maintenance mode when using the `--cluster` option to install the encryption package.

What to do next

Go to Intersight HyperFlex Software Encryption to enable encryption on your cluster

Install HX Software Encryption Package: Clusters with 13+ Nodes

Before you begin

Download the HyperFlex Software Encryption package from Cisco License Central, see Cisco Software Central.

Note The HyperFlex Software Encryption package is licensed by its own software PID, which is in addition to the HyperFlex Data Platform and Intersight software licenses. For more information, refer to the Cisco HyperFlex Systems Ordering and Licensing Guide.

Procedure

Step 1 SFTP the encryption package to each HyperFlex node. Use the admin account for the **username/password** and a file transfer application, such as winscp. This should upload the package to the /tmp or /home/admin directory.

Step 2 For clusters with more than 12 nodes, SSH to each node and use the `priv install-package --local` option.

Example:

```
priv install-package --local --path /home/admin/<package-filename>
```

Note
Do not shut down the cluster before proceeding to the next step, enabling HyperFlex Software Encryption. If you shut down the cluster and restart it, you will need to re-install the encryption package.

What to do next

Go to Intersight HyperFlex Software Encryption to enable encryption on your cluster.

Backup Encryption Key of HyperFlex Software Encryption

Encryption keys are stored in multiple copies in a distributed fashion in the cluster. To safeguard against catastrophic failures impacting the entire cluster, it is recommended to create an out-of-band backup of the encryption key to protect against data loss.



Note It is recommended to backup DEK after HX Software Encryption is enabled and after every Rekey. A preivously backed-up DEK cannot be restored after you perform a Rekey on your cluster.

To restore encrypted DEK configuration to the cluster when lost or corrupted from a previously saved back-up, contact TAC.

Procedure

Step 1 Run the `hxcli encryption backup-keys -f <path to file name>` command.

Note
Filename path should start with **`/home/admin/`**.

Step 2 Enter a passphrase after prompted when the command is executed.

After all the password rules are passed the command completes successfully saving the file in encrypted format.

Note
The passphrase should be a minimum of 8 characters in length and should contain at least 1 match of small case characters, at least 1 match of upper case characters, at least 1 match of numerical, and at least 1 match of special characters (one of !@#$%^&*()_+{}?).

Secure Disk Erase for HyperFlex Software Encryption

A software-based disk erase utility that provides the option to do a basic (Mode '0') and standard (Mode '1'/Mode '2') sanitization of the disk. The categorizations are primarily based on the areas of the disk that get sanitized, number of overwrite cycles and patterns on the drive as part of data erasure.

Consider the followings before performing the secure erase operation:

- `secure disk erase` is destructive and irreversible and improper use can lead to data loss.
- `secure disk erase` utility by default checks whether the selected disk contains any last copy of data. This check should not be bypassed.
- `secure disk erase` can be a time-consuming operation depending on the mode of sanitization and the size of the drive.
- You can trigger the secure erase operation from admin mode.
- More than one disk can be sanitized in parallel.

Limitations:

- Boot Disk/Housekeeping disks are not allowed to be secure erased.
- Once a disk is secure erased, the disk can not be re-introduced in the same cluster.
- `secure disk erase` is not supported on SED drives.
- When `secure disk erase` is in-progress, you cannot perform erase on the same disk until it is completed.

Procedure

Step 1 Run the `secure_disk_erase` command and specify the absolute path of the target disk.

Example:

```
-d DISK_PATH, --disk-path DISK_PATH
```

Step 2 Select from different modes of erase:

Example for basic (default) mode erase (i.e. mode '0'):

Example:

```
admin:~$ secure_disk_erase -d /dev/sdh -m 1

 THIS UTILITY WILL IRRECOVERABLY ERASE DATA FROM DRIVE.PROCEED WITH CAUTION.
 All data (including storfs) from the disk /dev/sdh will be destroyed, proceed [Y/N]:y
 Successfully removed the disk from the system: '/dev/sdh'
 Starting erase operation for disk '/dev/sdh'
     SEAGATE   ST1200MM0009      CN03   peripheral_type: disk [0x0]
       << supports protection information>>
       Unit serial number: WFK25FY70000C917H4GQ
       LU name: 5000c500a762ca2b

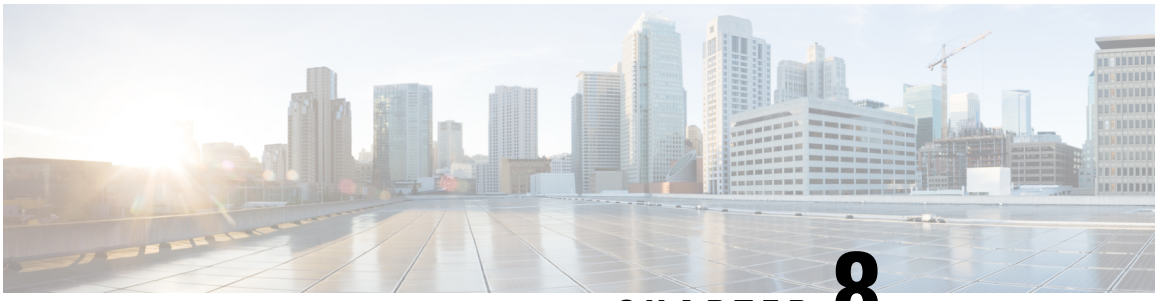
 Successfully triggered secure erase operation for the disk: '/dev/sdh'
 Please use following command to track the erase progress:
  secure_disk_erase -d /dev/sdh --progress
```

Step 3 Check erase progress using the following command:

Example:

```
admin:~$ secure_disk_erase -d /dev/sdh --progress
 Fetching the secure erase progress:
 Progress indication: 80.15% don
```

Step 4 After the erase process is completed, physically remove the erased drive from the node.



CHAPTER 8

Managing Datastores

- Managing Datastores, on page 107
- Adding Datastores, on page 108
- Editing Datastores, on page 109
- Unmounting Datastores, on page 110
- Deleting Datastores, on page 111
- Encryption Support for Datastores, on page 111
- Recovering from Partially Unmounted Datastores, on page 112

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 HX Datastore from the vCenter administration interface is not supported, and should not be done.

Before starting, review the following support notes:

Important

- Do not rename an HX datastore from vCenter. The datastore names shown in HX Connect or Intersight and in the ESXi host datastore (that appears in vCenter) must be identical including case sensitive. If they are not identical, some operations such as expansion, mount/unmount of a datastore will be impacted.
- Enabling encryption on your cluster is only possible during the datastore creation procedure. Encryption cannot be disabled for a datastore once enabled.
- HX Native Snapshots are not supported with multiple datastores.
- If using an M5/M6 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.
- For Vmware ESXi environments, ensure Storage I/O is disabled for all HyperFlex datastores in the vCenter. This setting is on a per datastore setting, and enabling this can cause unexpected performance impacts.

Procedure

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.

Procedure

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 To encrypt your datastore, click the **Software Encryption** check box.

For more information on enabling Software Encryption on your cluster, see Enabling HyperFlex Software Encryption Workflow, on page 100.

Step 7 Click **OK** to accept your changes or **Cancel** to cancel all changes.

Step 8 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 Starting with HX Release 5.0(2a), decreasing the size of an existing datastore is not supported. When you attempt to reduce the size of a datastore in 5.0(2a) or later release, the following error appears: Reducing datastore size is not allowed to prevent data loss. If the datastore is new, you can delete and recreate it with the correct size.



Note Do not rename datastores with controller VMs.

Procedure

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.

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.

Procedure

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 112.

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.

Procedure

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**.

Encryption Support for Datastores

To enable remote plugin Enryption, perform the following steps. For more information on enabling Software Encryption on your cluster, see Enabling HyperFlex Software Encryption Workflow.

Procedure

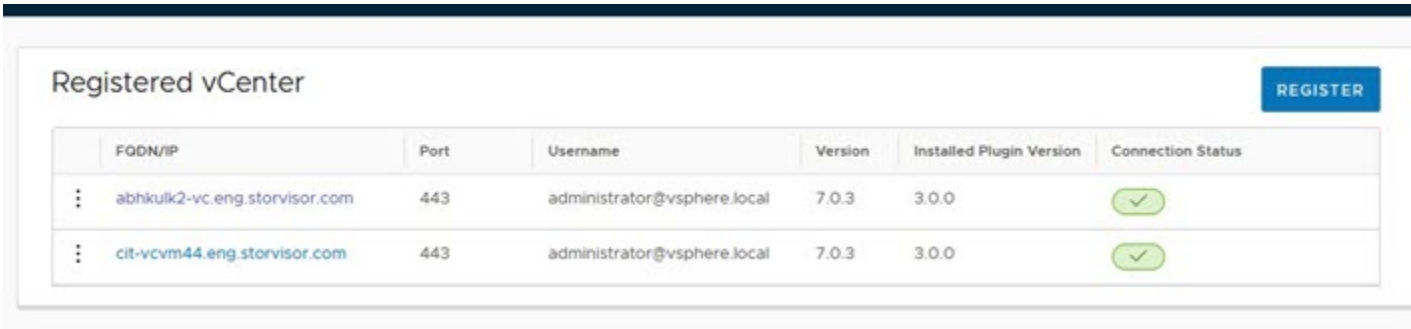
Step 1 Select the cluster you want to encrypt.

Step 2 Click **Datastore**.

Step 3 Click the **Create** button. The Create New Datastore window appears.

a) Type the Datastore Name.
b) Type the Size, and select GB or TB.
c) Select the Block Size, Select 4K or 8K
d) Check the **Software Encryption** check box

Step 4 Click **OK**. A new Datastore is created and added to the Datastore table list.



If the new Datastore does not appear in the list, click the **Refresh** arrow and recheck the list.

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.

Procedure

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 desire 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
```

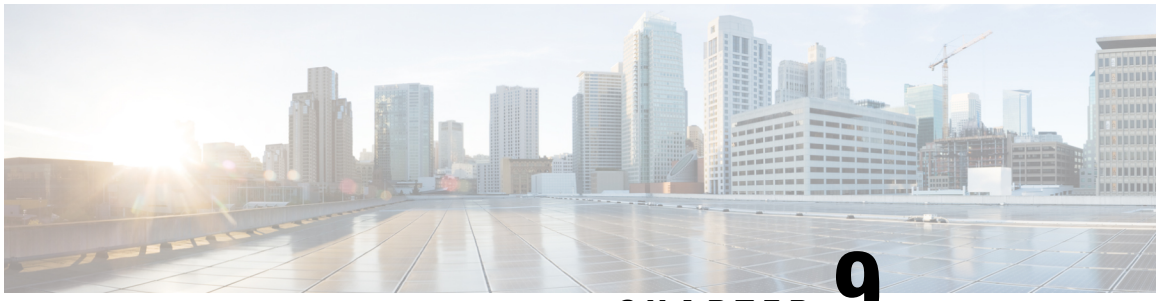
d) Stop the `storagerm` service.

```
# /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 115
- Disk Requirements, on page 116
- Replacing SSDs, on page 118
- Replacing NVMe SSDs, on page 119
- Replacing Housekeeping SSDs for Cisco HX Release 5.0(2b) and Later, on page 121
- Replacing Self Encrypted Drives (SEDs), on page 123
- Replacing or Adding Hard Disk Drives, on page 125

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.

Note Removing a functional drive from one HX cluster and installing it into another HX cluster is not supported.

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.
- Moving operational disks between servers within same cluster or moving them into expansion nodes within the same active cluster is not supported.
- 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 the event you see an error about unsupported drives or catalog upgrade, see the Catalog Update.
- 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.
- 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. After a successful rebalance the Cluster Flag Resiliency Status shows as Healthy.

In addition to the disks listed in the table below, all M5/M6 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
• Cisco B200 M5/M6 • C240 M5/M6 • C220 M5/M6 • C480 M5 • B480 M5	Choose any method. **Important** 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. USB boot is not supported for HX Compute-only nodes. • SD Cards in a mirrored configuration with ESXi installed. • Local drive HDD or SSD. • SAN boot. • M.2 SATA SSD Drive. **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.

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.

Procedure

Step 1 Identify the failed SSD.

- For cache or persistent SSDs, perform a disk beacon check. See Setting a Beacon, on page 73.

  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 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 119.
- 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 3 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 HXDP Maintenance Mode before doing so.

Step 4 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 disk types 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 and All-NVMe servers.
- Hot Swap NVME Drives in M5 ad M6 Servers is supported in HX Release 4.5(1a) and later.
- Replacing NVMe SSDs with an HGST SN200 disk requires HX Data Platform Release 2.5(1a) or later.
- For All-Flash nodes, NVMe SSDs are only allowed in slot 1 of the server. Other server slots do not detect NVMe SSDs.
- For All-Flash nodes, NVMe SSDs are only used for cache.

Note You can not use NVMe SSDs as the capacity or housekeeping drive(s) in All-Flash nodes.

- For M5 servers: If you are replacing an NVMe cache drive with a non-NVMe drive (or vice versa, if you are replacing a non-NVMe cache drive with an NVMe drive), you must replace the cable with a different SAS cable (for example, UCSC-RNVME-240M5 = HXAF240c M5 Rear NVMe cable (1) or UCSC-RSAS-C240M5 = C240 Rear UCSC-RAID-M5 SAS cbl(1)). This is required to ensure that the drive is discovered properly.



Note For M6 servers: you cannot replace an NVMe cache drive with a non-NVMe cache drive, due to the placement of the slots which are in the front.

Procedure

Step 1 Confirm the failed disk is an NVMe cache SSD.

Perform a physical check. NVMe cache SSDs and housekeeping SSDs do not respond to beacon requests.

If the failed SSD is not an NVMe SSD, see the Replacing SSD section of this guide.

Step 2 Put ESXi host into HXDP Maintenance Mode.

a) Log into HX Connect.
b) Select **System Information** > **Nodes** > ***node*** > **Enter HXDP 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 HXDP 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 HXDP 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.

Hot Swap NVME Drives in M5 and M6 Servers

Beginning with Cisco HyperFlex Release 4.5(1a), M5 and M6 servers with the VMD Enable bios option being active may hotswap NVME drives in new installs, and upgrades which have a combined HX+ UCS upgrade

performed. VMD enabled is set in the bios which allows NVME drives to be hot swappable without requiring HXDP Maintenance Mode or a reboot of ESXi.

To verify that VMD is enabled:

Before you begin

Ensure the Replacing NVMe SSDs, on page 119 conditions are met when using NVMe SSDs in HX Cluster servers.

Procedure

Step 1 In the Navigation pane, click **Servers**

Step 2 Go to **Policies** > **Root** > **BIOS Polices**

Step 3 Expand **root** > **Sub-Organizations** > ***your org***

Step 4 Select **hx-bios-af** (for M5) or **hx-bios-m6-af** (for m6)

Step 5 Click on **info**

Step 6 The BIOS Policy window apppears. Select **Advanced tab** > **LOM** > **PCIe Slots**

Step 7 Scroll down to see the **VMD Enable** settings and verify it is set to **Enabled**.

Replacing Housekeeping SSDs for Cisco HX Release 5.0(2b) and Later

Note This procedure applies to HXAF220c M5, HX220c M5, HXAF240c M5, HX240c M5, HXAF220c M6, HX220c M6, HXAF240c M6, and HX240c M6 servers only.

Identify the failed housekeeping SSD and perform the associated steps.

Procedure

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 HXDP Maintenance Mode. After performing the hardware steps, exit the node from HXDP Maintenance Mode.

Step 3 Using SSH, log into the storage controller VM of the cip node (any other working node) and run the following command to create **bootdev** partitions.

```
priv createBootdevPartitions --target 10.20.24.69
```

Sample response

```
hxshell:~$ priv createBootdevPartitions --target 10.20.24.69
Enter the root password:
create Bootdev Partitions initiated on 10.20.24.69
```

Note
The target should be the storage controller VM IP of the affected node.

This command reboots the affected node.

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.

```
# hxcli 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 SFTP the HX Data Platform installer packages into the storage controller VM of the affected node using the admin account for **user name**/**password** and a file transfer application, such as **winscp**. This should upload the package to the /tmp directory. Untar the package after copying it to the /tmp directory.

```
# tar -zxvf storfs-packages-<version>.tgz
```

Step 8 Using SSH, log into the storage controller VM of the cip node (any other working) node and run the following command:

```
priv housekeeping-preinstall --target 10.20.24.69
```

Sample response:

```
hxshell:~$ priv housekeeping-preinstall --target 10.20.24.69
      Enter root password :
      Copied secure files
```

Note
This step copies the secure files of /etc/springpath/secure/* folder from another working controller machine into the affected node.

Step 9 Run the following command on the storage controller VM of the cip node (any other working node) to install the HX Data Platform installer packages.

```
priv housekeeping-inst-packages --target 10.20.24.69
```

Sample response:

```
hxshell:~$ priv housekeeping-inst-packages --target 10.20.24.69
         Enter root password :
  Installed packages successfully
```

The package installation takes about 10 to 15 minutes.

Step 10 Enter the following command on the storage controller VM of the cip node (any other working node) to perform the post-install tasks..

```
priv housekeeping-postinstall --target 10.20.24.69
```

Sample response:

```
hxshell:~$ priv housekeeping-postinstall --target 10.20.24.69
        Enter root password :
        Successfully done post install tasks
       Successfully installed SE core package on 10.20.24.69 (optional only when Software Encryption
 is enabled on the cluster
```

For post-installation tasks, take the following steps:

a) Install the SE core package (optional if SE is enabled on the cluster).
b) Reboot the CVM.

This step reboots the affected node. Wait for the storage controller VM to automatically reboot.

Step 11 To confirm that the **cip-monitor** and **storfs** are in running status, run the `priv service cip-monitor status` and the `priv service storfs status` commands.

Example:

```
hxshell:~$ priv service cip-monitor status
       cip-monitor start/running, process 18251

       hxshell:~$ priv service storfs status
       storfs start/running, process 22057
```

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 118 or Replacing or Adding Hard Disk Drives, on page 125 and the hardware guide for your server.
- **Encryption is configured—**Ensure the following:
  - If you are replacing the SED, obtain a Return to Manufacturer Authorization (RMA). Contact TAC.
  - If you are using a local key for encryption, locate the key. You will be prompted to provide it.
  - Complete the steps below before removing any SED.

Procedure

Step 1 Ensure the HX Cluster is healthy.

Step 2 Log into HX Cluster .

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 118 and the hardware guide for your server.
   - **Rotational** (hybrid drives)—See Replacing or Adding Hard Disk Drives, on page 125 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 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.

Procedure

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.

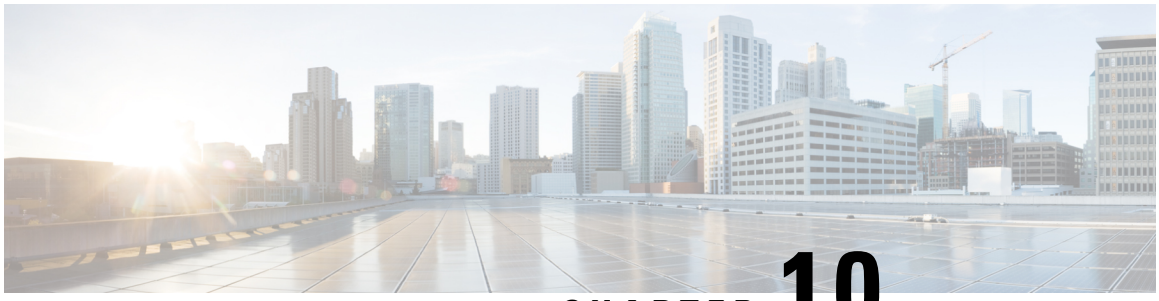
The storage starts being consumed by storage cluster immediately.

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 **Equipment** > **Manager** > **UCS** > **Equipment** > **Server** > **Inventory** > **Storage** tab, re-acknowledge the server node.

Note
Re-acknowledging a server is disruptive. Place the server into HXDP Maintenance Mode before doing so.



CHAPTER 10

Managing Nodes

- Managing Nodes, on page 127
- Identify Node Maintenance Methods, on page 129
- Searching by DNS Address or Host Name, on page 131
- Changing ESXi Host Root Password, on page 132
- Reinstalling Node Software, on page 133
- Changing Node Identification Form in vCenter Cluster from IP to FQDN, on page 134
- Replacing Node Components, on page 135
- Removing a Node, on page 137

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.

Procedure

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.

- To replace a node in the cluster, review the Replacing Node Components, on page 135.

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 remove a node from the cluster, review the Removing a Node from an Online Storage Cluster, on page 140 or Removing a Node from an Offline Storage Cluster, on page 143

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.

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 HXDP 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.

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. **How to Manage a Secure ESXi Configuration**[1]	TAC assisted node replacement required to return cluster to minimum 3 nodes. **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.

[1] **How to Manage a Secure ESXi Configuration:** This task applies only to an ESXi host that has a Trusted Platform Module (TPM). In general, you list the contents of the secure ESXi configuration recovery key to create a backup or as part of rotating recovery keys.

1. Run the following command in ESXi .

   ```
   esxcli system settings encryption recovery list
   ```

2. Save the output in a secure, remote location as a backup in case you must recover the secure configuration.

   For Example,

   ```
   [root@host1] esxcli system settings encryption recovery list
   Recovery ID                              Key
   ---------------------------------------  ---
   {2DDD5424-7F3F-406A-8DA8-D62630F6C8BC}
   478269-039194-473926-430939-686855-231401-642208-184477-602511-225586-551660-586542-338394-092578-687140-267425
   ```

Perform the recovery manually. Do not perform the recovery as part of an installation or upgrade script.

1. If the TPM fails, move the disk (containing the boot bank) to another host with a TPM.
2. Start the ESXi host.
3. When the ESXi installer window appears, press Shift+O to edit the boot options.
4. To recover the configuration, at the command prompt (from the ESXi host command line), add the following boot option to any existing boot options.

   ```
   encryptionRecoveryKey=recovery_key
   ```

   The secure ESXi configuration is recovered, and the ESXi host starts.

5. To ensure the change is saved permanently, enter the following command.

   ```
   /sbin/auto-backup.sh
   ```

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.

Procedure

Step 1 Assign DNS search addresses

a) Log into the HX Data Platform Installer virtual machine. Use either `ssh` or the vSphere console interface.

b) Edit `resolv.conf.d` file.

```
# vi /etc/resolvconf/resolv.conf.d/base
```

c) Confirm the change.

```
# resolvconf -u
# cat /etc/resolv.conf
```

d) 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.

a) Log into the HX Data Platform Installer virtual machine. Use either `ssh` or the vSphere console interface.

b) Open the hosts file for editing.

```
# vi /etc/hosts
```

c) Add the following line and save the file.

ip_address ubuntu *newhostname*

For each host *ip_address*, enter the host *newhostname*.

a) 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.

Procedure

Step 1 Log into 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.

Procedure

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).

Procedure

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.

a) From vCenter, right-click the **datacenter**. Select **New Cluster**.
b) 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.

a) From vCenter, right-click the **datacenter** > **cluster**. Select **Add Host**.
b) Select an ESXi host using their FQDN.
c) 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 58 for additional information on reregistering a cluster.

Step 11 Enable VMware cluster HA and DRS using the post install script:

a) Log into the HX cluster IP as admin and run the command **# hx_post_install** .
b) Select Option 1 - "New/Existing Cluster" and input all login credentials
c) Type "y" if you want to enter a new license key
d) Type "y" to enable HA and DRS in the cluster
e) 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 than 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

Removing a node is supported on the following cluster types:

Table 4: Cluster Types that Support Node Removal

Cluster Type	Converged	Compute
Standard	Yes	Yes
Stretch	No	Yes
Edge	Yes (see Note) **Note** Removing a node (compute or converged) is supported only on a Edge clusters with more than 3 nodes. For Edge clusters with 4 nodes, you will need to follow the offline node removal process. For Edge clusters with 5 or more nodes, the online and offline node removal process is supported. The online node removal method is however **recommended**.	

Depending upon the number of nodes in a cluster, you can remove a node when the cluster is either online or you need to make the cluster offline. Before you do so, you must first ensure that you have completed the required preparation steps.

The affecting context is based on the number of converged nodes. The number of compute nodes does not affect the process to remove a node.

You can only remove 1 converged node at any time.

For clusters with 4 converged nodes, follow the offline node removal process. For clusters with 5 converged nodes or more, follow the online node removal process.

Note Removing a converged node from a 3-node cluster is not supported

Note If you remove a node when the cluster is offline, you cannot add the node back to the cluster.

Prior to removing a node or nodes for HyperFlex clusters with Logical Availability Zones (LAZ) configured, LAZ must be disabled.

If LAZ is utilized in the HyperFlex cluster, then the number of remaining nodes must be in a balanced configuration that supports LAZ per the LAZ Guidelines and Considerations prior to reenabling LAZ.

Preparing to Remove a Node

Before you remove a node from a storage cluster, complete the following steps.

Procedure

Step 1 Ensure the cluster is healthy.

```
# stcli cluster info
```

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

```
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 Make a note of zkEnsemble. It contains the data IP of controller VMs (CVM).

Example:

```
admin:~$ cat /etc/springpath/storfs.cfg | grep -i ense
crmZKEnsemble=10.104.18.37:2181,10.104.18.38:2181,10.104.18.39:2181, 10.104.18.40:2181,
10.104.18.41:2181
```

If the node which was removed was ucs-308 whose CVM data IP is 10.104.18.40, then that CVMs data IP should no longer appear after you run the above command after the node removal.

Step 5 Put the node to be removed into Maintenance mode. Choose a method: vSphere GUI, controller VM command line (CLI) or HyperFlex Connect System Information panel:

GUI

a) Right-click each host, scroll down the list, and select **Maintenance Mode** > **Enter 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 Maintenance Mode.

b) In HX Connect, from the **MANAGE** > **System Information panel's Node tab**, select the node, and then click on the button to **Enter HXDP Maintenance Mode**.

CLI

a) Log in to a controller VM as an admin user.
b) Run **stcli cluster info** and look for `stNodes: section`

```
# stcli cluster info

stNodes:
    ----------------------------------------
    type: node
    id: 689324b2-b30c-c440-a08e-5b37c7e0eefe
    name: ucs-305
```

```
----------------------------------------
type: node
id: 9314ac70-77aa-4345-8f35-7854f71a0d0c
name: ucs-306
----------------------------------------
type: node
id: 9e6ba2e3-4bb6-214c-8723-019fa483a308
name: ucs-307
----------------------------------------
type: node
id: 575ace48-1513-0b4f-bfe1-e6abd5ff6895
name: ucs-308
----------------------------------------
type: node
id: 875ebe8-1617-0d4c-afe1-c6aed4ff6732
name: ucs-309
```

Under the `stNodes` section, the `id` is listed for each node in the cluster. Find the node id or name you need to remove.

c) Move the ESX host into Maintenance mode.

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

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

For example, to remove node ucs-308:

Example:

```
stcli node maintenanceMode –id 575ace48-1513-0b4f-bfe1-e6abd5ff6895
or
stcli node maintenanceMode –ip 10.104.18.40
```

Step 6 Wait for 2 hours, monitor healing info in stcli cluster storage-summary. You should wait until you see "Storage cluster is healthy." as shows in the following example:

Example:

```
admin:$ stcli cluster storage-summary | grep -i heali -A 8
healingInfo:
    inProgress: False
resiliencyInfo:
    messages:
        ----------------------------------------
        Storage node 10.104.18.40 is unavailable.
        ----------------------------------------
        Storage cluster is healthy.
        ----------------------------------------
```

Before the healing starts you will see following:

```
admin:$ date; stcli cluster storage-summary | grep -i heali -A 8
Thu Sep 30 12:33:57 PDT 2021
healingInfo:
inProgress: False
resiliencyInfo:
messages:
----------------------------------------
Storage cluster is unhealthy   .
----------------------------------------
Storage node 10.104.18.40 is unavailable   .
----------------------------------------
```

After 2 hours + you will see following:

```
admin:$ stcli cluster storage-summary | grep -i heali -A 8
healingInfo:
```

```
messages:
Space rebalanc   ing in progress, 83 % completed.
inProgress: True
percentComplete: 83
estimatedCompletionTimeInSeconds: 211
resiliencyInfo:
messages:
```

What to do next

Proceed to Removing a Node. Choose the Online or Offline method based on the number of nodes in your storage cluster.

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.

Note You can remove multiple nodes in a series, as long it is done one at a time and when the cluster is healthy between each successive node removal. You must also have followed the steps required to prepare to remove a node. For more information, see Preparing to Remove a Node, on page 137.

Note Prior to removing a node or nodes for HyperFlex clusters with Logical Availability Zones (LAZ) configured, LAZ must be disabled.

If LAZ is utilized in the HyperFlex cluster, then the number of remaining nodes must be in a balanced configuration that supports LAZ per the LAZ Guidelines and Considerations prior to reenabling LAZ.

Note Do not remove the controller VM or other HX Data Platform components before you complete the steps in this task.

Procedure

Step 1 Run the **stcli cluster info** command and look for `stNodes: section` to find the node which needs to be removed.. This information is also available when you put the node in maintenance mode.

Example:

```
    ----------------------------------------
stNodes:
type: node
id: 689324b2-b30c-c440-a08e-5b37c7e0eefe
name: ucs305

type: node
id: 9314ac70-77aa-4345-8f35-7854f71a0d0c
```

```
name: ucs306

type: node
id: 9e6ba2e3-4bb6-214c-8723-019fa483a308
name: ucs307

type: node
id: 575ace48-1513-0b4f-bfe1-e6abd5ff6895
name: ucs308

type: node
id: 875ebe8-1617-0d4c-af
name: ucs 309
```

The **stcli node remove** command to remove nodes from the 5-node cluster are:

- **stcli node remove –-ip-1 ucs 308** or
- **stcli node remove –-id-1 575ace48-1513-0b4f-bfe1-e6abd5ff6895** or
- **stcli node remove --ip-1 "vmk0 IP" -f**

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 HXDP 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 2 Confirm that the node is removed from the storage cluster.

a) Check the storage cluster information.

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

b) Check the `ActiveNodes` entry in the response to verify the cluster has one less node.

c) Check that the node which was removed is not part of Ensemble. For example:

Example:

```
admin:~$ cat /etc/springpath/storfs.cfg | grep -i ense
crmZKEnsemble=10.104.18.37:2181,10.104.18.38:2181,10.104.18.39:2181, 10.104.18.40:2181,
10.104.18.41:2181
```

For example, if the node which was removed was ucs-308 whose CVM data IP is 10.104.18.40, then that CVMs data IP should no longer appear after you run the above command after the node removal as seen above.

.

If there are more than 5 nodes and the removed node was part of ensemble, then the new node ip appears in the `crmZKEnsemble`. For example, if the cluster initially has 7 nodes (10.104.18.37 to 10.104.18.43), and `crmZKEnsemble` has 10.104.18.37:2181,10.104.18.38:2181,10.104.18.39:2181, 10.104.18.40:2181, 10.104.18.41:2181, then after removal of 10.104.18.40, `crmZKEnsemble` has either:

10.104.18.37:2181,10.104.18.38:2181,10.104.18.39:2181, 10.104.18.42:2181, 10.104.18.41:2181, or:

10.104.18.37:2181,10.104.18.38:2181,10.104.18.39:2181, 10.104.18.43:2181, 10.104.18.41:2181

Step 3 Verify that disks from the removed node no longer appear by running the `hxcli disk list` command:

```
admin:~$ hxcli disk list --no-loader
+-----------+------------------+------+----------+---------+------------+-------------+
| NODE NAME | HYPERVSIOR STATUS | SLOT | CAPACITY | STATUS | TYPE | USAGE |
+-----------+------------------+------+----------+---------+------------+-------------+
| ucs305 | ONLINE | 1 | 111.8 GB | Claimed | Solidstate | System |
| ucs305 | ONLINE | 2 | 894.3 GB | Claimed | Solidstate | Caching |
| ucs305 | ONLINE | 3 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 4 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 5 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 6 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 7 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE |8|0B| Unknown | | |
| ucs306 | ONLINE | 1 | 111.8 GB | Claimed | Solidstate | System |
| ucs306 | ONLINE | 2 | 894.3 GB | Claimed | Solidstate | Caching |
| ucs306 | ONLINE | 3 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 4 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 5 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 6 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 7 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE |8|0B| Unknown | | |
| ucs307 | ONLINE | 1 | 111.8 GB | Claimed | Solidstate | System |
| ucs307 | ONLINE | 2 | 894.3 GB | Claimed | Solidstate | Caching |
| ucs307 | ONLINE | 3 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 4 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 5 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 6 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 7 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE |8|0B| Unknown | | |
| ucs309 | ONLINE | 1 | 111.8 GB | Claimed | Solidstate | System |
| ucs309 | ONLINE | 2 | 894.3 GB | Claimed | Solidstate | Caching |
| ucs309 | ONLINE | 3 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs309 | ONLINE | 4 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs309 | ONLINE | 5 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs309 | ONLINE | 6 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs309 | ONLINE | 7 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs309 | ONLINE |8|0B| Unknown | | |
+-----------+------------------+------+----------+---------+------------+-------------+
```

For example, if you removed ucs-308, then its disks no longer appear.

Step 4 Remove the host from the vCenter **Hosts and Cluster** view.

a) Log in to vSphere Web Client Navigator. Navigate to **Host** in the vSphere Inventory.

b) Right-click on the host and select **All vCenter Actions** > **Remove from Inventory**. Click **Yes**.

Step 5 Confirm that all the node-associated datastores are removed. For example, run the following command in ESXi:

```
[root@ucs308:~] esxcfg-nas -l
ds4 is 169.254.226.1:ds4 from 6894152532647392862-8789011882433394379 mounted available
ds3 is 169.254.226.1:ds3 from 6894152532647392862-8789011882433394379 mounted available
ds2 is 169.254.226.1:ds2 from 6894152532647392862-8789011882433394379 mounted available
ds5 is 169.254.226.1:ds5 from 6894152532647392862-8789011882433394379 mounted available
ds1 is 169.254.226.1:ds1 from 6894152532647392862-8789011882433394379 mounted available
```

Note
If any node-associated datastores are listed, then unmount and delete those datastores manually.

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.

Number of nodes in cluster	Number of failed nodes in cluster	Method	Notes
3	1	See TAC to remove and replace the node.	-
4	1	Offline node remove. **Note** If "*storage cluster manager is not configured*" error is seen in hxconnect or as part of stcli cluster storage-summary | grep -i heali -A 8 command, Contact TAC for assistance.	A 4 node cluster with 2 nodes down, requires TAC assistance.
5 or more	1	Cluster can be online.	Online mode is recommended.
5 or more	2	Cluster must be offline.	A 5 node cluster with 3 nodes down, requires TAC assistance.

Note Do not remove the controller VM or other HX Data Platform components before you complete the steps in this task.

Procedure

Step 1 Follow the process to prepare for removing a node. For more information, see Preparing to Remove a Node, on page 137.

Step 2 (For 4-node clusters only) Prepare to shutdown, then shutdown the storage cluster.

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 Run the **stcli cluster info** command and look for `stNodes: section` to find the node which needs to be removed.. This information is also available when you put the node in maintenance mode.

```
Example:
    ----------------------------------------
    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: d2d43691-daf5-50c4-d096-941358fede374
    name: example5
```

Step 4 Remove the desired node using the `stcli node remove` command.

For example:

To remove 1 node

- **stcli node remove –ip-1 example5** or
- **stcli node remove –id-1 d2d43691-daf5-50c4-d096-941358fede374** or
- **stcli node remove --ip-1 "vmk0 IP" -f**

```
Response:
    Successfully removed node: EntityRef(type=3, id='', name='10.10.2.4')
```

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. This command also does not 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 HXDP 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 5 Restart the cluster.

```
# hxcli cluster start
```

Step 6 Confirm that the node is removed from the storage cluster once the cluster is up.

a) Check the storage cluster information.

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

b) Check the `ActiveNodes` entry in the response to verify the cluster has one less node.

c) Check that the node which was removed is not part of Ensemble.

```
Example:
admin:~$ cat /etc/springpath/storfs.cfg | grep -i ense
crmZKEnsemble=10.104.18.37:2181,10.104.18.38:2181,10.104.18.39:2181, 10.104.18.40:2181,
10.104.18.41:2181
```

For example, if you removed 10.104.18.40, note that it no longer appears.

d) If the node remove action is successful for a cluster with 5 or fewer nodes, the ensemble may contain 4 participant nodes, and 1 observer node. The only participant nodes are updated in stMgr.cfg; the observer node details are not updated.

Example:

Ensemble after node remove server.0=10.107.16.111:2888:3888:participant;10.107.16.111:2181

server.1=10.107.16.107:2888:3888:participant;10.107.16.107:2181

server.2=10.107.16.108:2888:3888:participant;10.107.16.108:2181

server.4=10.107.16.110:2888:3888:participant;10.107.16.110:2181

server.5=10.107.16.109:2888:3888:observer;10.107.16.109:2181

version=10000558b

Step 7 Verify that disks from the removed node no longer appear by running the `hxcli disk list` command:

```
admin:~$ hxcli disk list --no-loader
+-----------+------------------+------+----------+---------+------------+-------------+
| NODE NAME | HYPERVSIOR STATUS | SLOT | CAPACITY | STATUS | TYPE | USAGE |
+-----------+------------------+------+----------+---------+------------+-------------+
| ucs305 | ONLINE | 1 | 111.8 GB | Claimed | Solidstate | System |
| ucs305 | ONLINE | 2 | 894.3 GB | Claimed | Solidstate | Caching |
| ucs305 | ONLINE | 3 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 4 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 5 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 6 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 7 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 8 | 0 B | Unknown | | |
| ucs306 | ONLINE | 1 | 111.8 GB | Claimed | Solidstate | System |
| ucs306 | ONLINE | 2 | 894.3 GB | Claimed | Solidstate | Caching |
| ucs306 | ONLINE | 3 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 4 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 5 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 6 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 7 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 8 | 0 B | Unknown | | |
| ucs307 | ONLINE | 1 | 111.8 GB | Claimed | Solidstate | System |
| ucs307 | ONLINE | 2 | 894.3 GB | Claimed | Solidstate | Caching |
| ucs307 | ONLINE | 3 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 4 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 5 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 6 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 7 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 8 | 0 B | Unknown | | |
+-----------+------------------+------+----------+---------+------------+-------------+
```

For example, if you removed ucs-308, then its disks no longer appear.

Step 8 Remove the host from the vCenter **Hosts and Cluster** view.

a) Log in to vSphere Web Client Navigator. Navigate to **Host** in the vSphere Inventory.

b) Right-click on the host and select **All vCenter Actions** > **Remove from Inventory**. Click **Yes**.

Step 9 Confirm that all the node-associated datastores are removed. For example, run the following command in ESXi:

```
[root@ucs308:~] esxcfg-nas -l
ds4 is 169.254.226.1:ds4 from 6894152532647392862-8789011882433394379 mounted available
ds3 is 169.254.226.1:ds3 from 6894152532647392862-8789011882433394379 mounted available
ds2 is 169.254.226.1:ds2 from 6894152532647392862-8789011882433394379 mounted available
```

```
ds5 is 169.254.226.1:ds5 from 6894152532647392862-8789011882433394379 mounted available
ds1 is 169.254.226.1:ds1 from 6894152532647392862-8789011882433394379 mounted available
```

Note
If any node-associated datastores are listed, then unmount and delete those datastores manually.

Removing a Compute Node

Procedure

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 HXDP 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 --id-1
```

Or

```
stcli node remove --ip-1
```

Or

```
stcli node remove --ip-1 "vmk0 IP" -f
```

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 If compute node virtnode entry still exists in `stcli cluster info` output, perform the following:

Restart the stMgr management service using `priv service stMgr restart` on the SCVMs.

Step 10 Clear stale entries in the compute node by logging out of Cisco HX Connect and then logging into Cisco HX Connect.

Step 11 Disable and re-enable the High Availability (HA) and Distributed Resource Scheduler (DRS) services to reconfigure the services after node removal.

Reuse a Previously Removed Node Within the Same Cluster

To reuse a node within the same cluster that was previously removed, perform the following steps:

Before you begin

- Nodes must be removed on a cluster running HXDP 4.5(2b) or later in order to be reused in the same cluster.
- The HX Node must be removed only while the cluster is online. If the node is removed while the cluster is offline, then the node may not be re-used within the same cluster.
- The HX Node removal must be performed with the stcli node command listed within the administration guide. If the node is not properly removed from the cluster then it may not be reused within the same cluster.
- For a 4 node cluster, use Removing a Node from an Offline Storage Cluster, on page 143 only. Node reuse is not supported for 4 node clusters.

Procedure

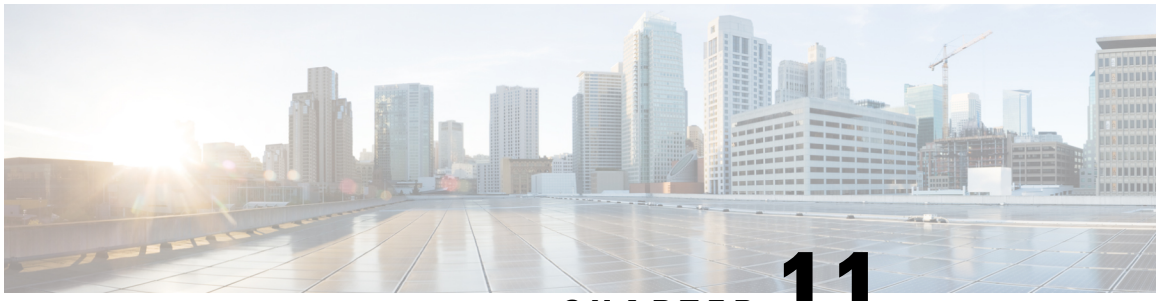
Step 1 Remove the ESXi host from the vCenter Cluster inventory.

Step 2 Re-install ESXi with the same version that matches the rest of the HX Cluster.

Step 3 Delete only the UCS Service Profile from UCS Manager that was previously associated with this removed node.

Step 4 Use the HX Installer of the same version and run an Expansion Workflow.

Note
Be sure to check the **Clear Disk Partitions** box during the expansion workflow.



CHAPTER 11

Expand Cisco HyperFlex System Clusters

- Cluster Expansion Guidelines, on page 149
- Prerequisites When Expanding M5/M6 Clusters, on page 151
- Mixed Cluster Expansion Guidelines - Cisco HX Release 5.5(x) , on page 151
- Steps During Mixed Cluster Expansion, on page 152
- Prerequisites for Adding a Converged Node, on page 152
- Preparing a Converged Node, on page 153
- Adding a Converged Node to an Existing Cluster, on page 153
- Prerequisites for Adding a Compute-Only Node, on page 158
- Preparing a Compute-Only Node, on page 159
- Adding a Compute-Only Node to an Existing Cluster, on page 162
- Resolving Failure of Cluster Expansion, on page 166
- Logical Availability Zones, on page 167

Cluster Expansion Guidelines

Please review these guidelines before expanding your cluster.

Note If you have LAZ configured (enabled by default for clusters of size 8 or more), please review Logical Availability Zones, on page 167 prior to moving ahead with expansion.

- **Non Pre-configured Cisco HyperFlex Systems:** The Cisco HyperFlex System must have VMware ESXi installed before starting the actual Cisco HyperFlex Installation. In the event your system does not have VMware ESXi preinstalled, perform the tasks in the Cisco HyperFlex Systems Customized Installation Method chaper of the Cisco HyperFlex Systems Installation Guide for VMware ESXi guide for your release.
- 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 M.2 SATA SSD for M5/M6 converged nodes. For compute-only nodes, ESXi installation is supported for 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.

  HW RAID M.2 (UCS-M2-HWRAID and HX-M2-HWRAID) is a supported boot configuration starting with HX Data Platform release 4.5(1a) and later.

- You must click on the discovered cluster to proceed with expanding a standard ESX cluster. Not doing so results in errors.
- Use only Admin credentials for the Controller VM during expansion workflow. Using any other credentials other than Admin may cause the expansion to fail.
- In the event you see an error about unsupported drives or catalog upgrade, see the Compatibility Catalog.
- Starting with HX Release 5.0(1b) and later, you can expand ESXi based 10/25 GbE HyperFlex Edge clusters with 2 nodes via Intersight.

  Please refer to the Intersight documentation for all requirements: Cluster Expansion Requirements.

- Starting with HX Release 5.0(2b) you can not add new nodes with 375G WL cache drives to an existing cluster with nodes that have 1.6TB cache drives.
- Moving operational disks between servers within same cluster or moving them into expansion nodes within the same active cluster is not supported.

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/M6 server:

   a. Clone the template.

   b. Uncheck the Flex flash from local boot policy, if the compute M5/M6 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 Expanding M5/M6 Clusters

Prior to beginning cluster expansion in M5/M6 clusters, perform the following tasks:

- **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 the HX cluster and UCS Manager to the appropriate recommended release for your deployment. For more information, see the Cisco HyperFlex Recommended Software Release and Requirements Guide.
- Download and deploy the matching HX Data Platform Installer (release should be same as cluster) release to run the expansion workflow.

Mixed Cluster Expansion Guidelines - Cisco HX Release 5.5(x)

General Guidelines:

- HX240c M6 is not able to use the additional slots if combined in a cluster with M5 nodes.
- All servers must match the form factor (220/240), type (Hybrid/AF/NVME), security capability (Non-SED only) and disk configuration (QTY, capacity, and non-SED) across the cluster.

Mixed Cluster Expansion Options: Supported

- Expanding existing M5 clusters with M6 converged nodes is supported.
- Expanding existing mixed M5/M6 cluster with M5 or M6 converged nodes is supported.
- Only expansion workflow is supported to create a mixed cluster (Initial cluster creation with mixed M5/M6 servers is not supported).
- Adding any supported compute-only nodes is permitted with all M5, M6 and mixed M5/M6 clusters using the HX Data Platform 5.0 or later Installer. Some example combinations are listed here, many other combinations are possible.

Mixed Cluster Expansion Options: Not Supported

- Expanding existing M6 cluster with M5 converged nodes is NOT supported.
- Mixing Intel and AMD M6 is not supported.
- HX Edge does not support mixed M5 and M6 clusters.
- Initial cluster creation with mixed M5/M6 servers is not supported.

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 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, All Flash or NVME), 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 release, 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 release 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 release 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

Procedure

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.

Procedure

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, *10.193.211.120*.
User Name	<*admin*> username.
Password	<*admin*> password.
vCenter Credentials	
vCenter Server	vCenter server FQDN or IP address. For example, *10.193.211.120*. **Note** • 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	<*admin*> username. For example, *administrator@vsphere.local*.
Admin Password	<*root*> password.
Hypervisor Credentials	
Admin User Name	<*admin*> username. This is **root** for factory nodes.

Field	Description
Admin Password	<*root*> password. Default password is `Cisco123` for factory nodes. **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.

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 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*.

Field	Description
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.
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**. 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
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. **Note** Compute-only nodes can be added only after the storage cluster is created.	
Advanced Configuration	
Jumbo frames **Enable Jumbo Frames** checkbox	Check to set the MTU size for the storage data network on the host vSwitches and vNICs, and each storage controller VM. The default value is 9000. **Note** To set your MTU size to a value other than 9000, contact Cisco TAC.
Disk Partitions **Clean up Disk Partitions** checkbox	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. **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.

Step 9 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 into the ESXi host to be added in vSphere Client.
- Power off the controller VM.
- Add the host to the vCenter cluster in vSphere Client.
- In the HX Data Platform Installer, click **Retry Deploy**.

Step 10 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.

Step 11 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 12 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 into 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

You can add a compute-only node 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.
- 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 release on the node matches the Cisco HX Data Platform release, the ESXi release and the vCenter release. To identify the software release, 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.
- Compute-only nodes are deployed with automatic detection and configuration of disk and boot policies based on the boot hardware.

  Starting with HX Data Platform release 4.5(1a) and later, compute-only nodes are deployed with automatic detection and configuration of disk and boot policies based on the inventoried boot hardware. Users cannot directly select the UCSM policies. Instead, the boot device is automatically determined based on the first acceptable boot media discovered in the server. The tables below show the priority order for M5/M6 generation servers. Reading from top to bottom, the first entry that is a match based on the

inventoried hardware are selected automatically during cluster expansion. For example, when expanding with a B200 compute node with a single M.2 boot SSD, the second rule in the table below is a match and used for SPT association.

If the server is booted using a mechanism not listed (such a SAN boot), the catch-all policy of **anyld** is selected and administrators may subsequently modify the UCSM policies and profiles as needed to boot the server.

Table 5: Priority for M6

Priority for M6			
Priority	**SPT Name**	**Boot Device**	**Number of disks**
1	compute-nodes-m6-m2r1	M6 - M.2 - 2 Disks	2
2	compute-nodes-m6-m2sd	M6 - M.2 - 1 Disk	1
3	compute-nodes-m6-ldr1	MegaRAID Controller	2
4	compute-nodes-m6-anyld	M6 - Generic	Any

Table 6: Priority for M5

Priority for M5			
Priority	**SPT Name**	**Boot Device**	**Number of disks**
1	compute-nodes-m5-m2r1	M.2 Raid	2
2	compute-nodes-m5-m2pch	PCH/Non-RAID M.2	1
3	compute-nodes-m5-sd	FlexFlash	2
4	compute-nodes-m5-ldr1	MegaRAID	2
5	compute-nodes-m5-sd	FlexFlash	1
6	compute-nodes-m5-anyld	Any other config	Any

Preparing a Compute-Only Node

Procedure

Step 1 Ensure that the server is a supported HX server and meets the requirements. For more details, see the Host Requirements section if the Cisco HyperFlex Installation guide for your release..

Step 2 Log into 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. Review the Cisco HyperFlex Software Requirements and Recommendations document for the list of compatible compute-only nodes.

Verify the HX Data Platform Installer

Procedure

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.

Procedure

Once the install beings, 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 HyperFlex Data Platform supported release 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.

Procedure

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 into Cisco UCS Manager.

Step 3 Log into the KVM console of the server through Cisco UCS Manager.

a) In the Navigation Pane, click **Servers** > **Service Profiles** > **Sub-Organizations** > ***hx-cluster***.
b) 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-ESXi-7.0U3-20328353-Cisco-Custom-7.3.0.10-install-only.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.

a) Select **Virtual Media** > **Activate Virtual Devices**.

This opens in a pop-up window.

b) Click **Accept the session** > **Apply**.

Step 6 From the **Map CD/DVD** option, map to the location of the *HX-Vmware.iso* file.

a) Select the *HX-Vmware.iso* file.
b) 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.

a) Click the **Reset** button on the KVM console. Click **OK** to confirm.
b) Select **Power Cycle**. Click **OK**.

Step 8 Change the boot path to point to the *HX-Vmware.iso* file.

a) Press **F6**.
b) From the **Enter boot selection** menu, use the arrow keys to highlight the ***Cisco vKVM-Mapped vDVD1.22*** option.
c) 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.

Procedure

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 expansion, 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

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, *10.193.211.120*.
User Name	<*admin*> username.
Password	<*admin*> password.
vCenter Credentials	
vCenter Server	vCenter server FQDN or IP address. For example, *10.193.211.120*. **Note** • 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	<*admin*> username. For example, *administrator@vsphere.local*.
Admin Password	<*root*> password.
Hypervisor Credentials	
Admin User Name	<*admin*> username. This is **root** for factory nodes.
Admin Password	<*root*> password. Default password is `Cisco123` for factory nodes. **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.

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 **(M6 Servers Only)** 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**. 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.
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 7 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.

Field	Information
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 8 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**.

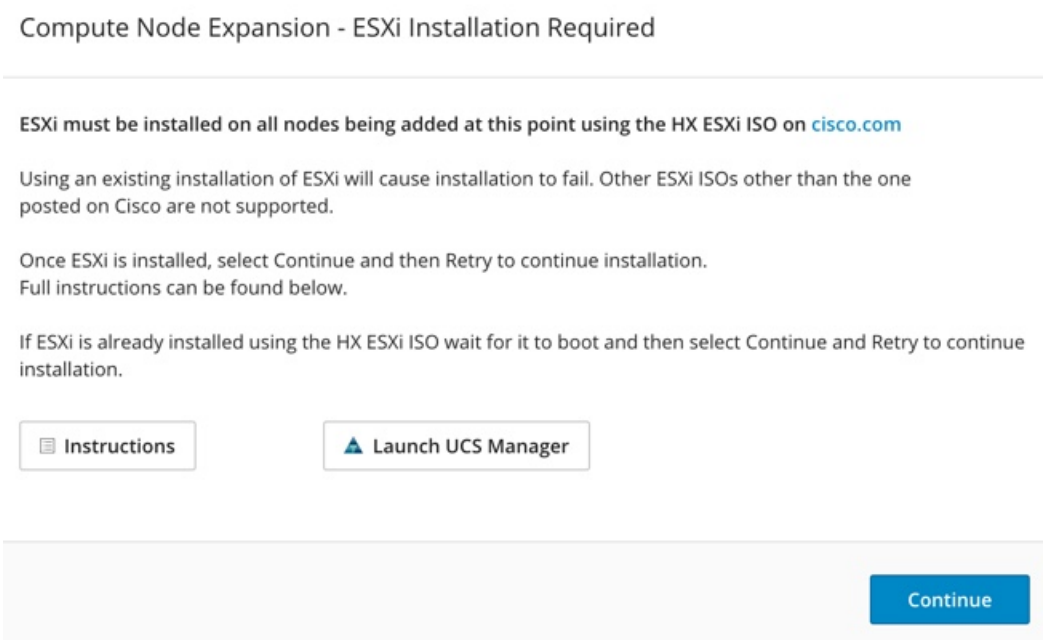
c. In the working pane, click the **Storage** tab. Click the **Local Disk Configuration Policy** sub tab.

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 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:

a) Log into the ESXi host to be added in vSphere Client.
b) Power off the controller VM.
c) Add the host to the vCenter cluster in vSphere Web Client.
d) In the HX installer, click **Retry Deploy**.

Step 9 When installation is complete, start managing your storage cluster by clicking **Launch HyperFlex Connect**.

Step 10 After the new nodes are added to the storage cluster, HA services are reset so that HA is able to recognize the added nodes.

a) Log on to VMware vSphere Client.
b) Select **Home** > **Hosts and Clusters** > **Datacenter** > **Cluster** > **Host**.
c) Select the new node.
d) Right-click and select **Reconfigure for vSphere HA**.

Step 11 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:

Procedure

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.

Logical Availability Zones

The Logical Availability Zones (LAZ) feature groups cluster storage nodes in fixed number pools of nodes which enable higher resiliency. The number of zones that can be set automatically or selected manually based on cluster parameters, such as replication factor and cluster size. LAZ is enabled by default on HyperFlex clusters with 8 or more storage nodes. The feature remains enabled through the life cycle of the cluster unless explicitly disabled either at install time or post installation.

Advantages of Logical Availability Zones

Reducing the failure of large clusters in a distributed system is the primary advantage of enabling LAZ on install. In any distributed storage system, when the number of resources in the cluster grow, so does the failure risk. Multiple simultaneous failures could result in permanent data unavailability.

LAZ helps reduce risk of multiple simultaneous component and node failures from causing a catastrophic failure. It does this by grouping resources based on some basic constraints, you can improve the availability from 20% up to 70% in comparison to the same cluster without LAZ. The amount of improvement depends on the cluster replication factor (RF) as well as the number of zones configured. In principle, a cluster with fewer zones and a higher replication factor provides optimal results. Additionally, LAZ saves time by performing maintenance tasks on multiple resources grouped in the same zone, an option not possible in clusters without LAZ.

It is recommended that LAZ be enabled during the HyperFlex cluster installation. Enabling LAZ during install provides optimal cluster performance and data availability. With the guidance of support, LAZ can be enabled or disabled at a later time using the command line interface (CLI). Review the LAZ guidelines before disabling.

Specifying the Number of Zones and Optimizing Balance

The number of zones is set automatically by default and recommended. 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.

To maintain the most balanced consumption of space and data distribution, it is recommended that the number of nodes in a cluster are whole multiple of number of zones, which is either 3, 4, or 5. 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. Users with a need may manually specify 3, 4 or 5 zones.

LAZ Guidelines and Considerations

- HyperFlex clusters determine which nodes participate in each zone. This configuration cannot be modified.
- When changing the number of resources, add or remove an equal number of resources from each configured zone.
- **Cluster Expansion:** Perform expansions in the same increment number of nodes as zones in order to maintain a balanced zone. A balanced zone is when the number of nodes and zones added during install or expansion (or a permanent failure of nodes from zone(s) occurs) are equal. For example, a cluster with 12 nodes and 4 zones is a balanced zone. In this case, it is recommended to add 4 nodes during expansion.
- **Imbalanced Zones:** Zones may become imbalanced when the number of nodes and zones added during install or expansion (or permanent failure of nodes from zone(s)) are not equal. Imbalanced zones can lead to non-optimal performance and are **not** recommended. For example, a cluster with 11 nodes and 4 zones will have 3 nodes per zone except the last zone. In this case, you need to add 1 node to make it balanced. The new node is added automatically to the last zone.
- **Disabling and Re-enabling LAZ:** You can disable and enable LAZ dynamically. It is not recommended to disable and re-enable LAZ in the same cluster with a different number of zones. Doing so could result in an excessive amount of movement and reorganization of data across the cluster - to comply with existing data distribution rules if LAZ is turned on in a cluster already containing data. This can result in the cluster becoming no longer zone compliant for example, if the cluster usage is already greater than 25%.

Viewing LAZ Status and Connections

- To view LAZ information from the HX Connect dashboard, log into HX Connect and use the **System information** and **HyperFlex Connect** > **Dashboard** menu.
- You can also view LAZ details through CLI by running the `stcli cluster get-zone` command. The following is sample output from the `stcli cluster get-zone` 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:0000000000000004
    numNodes: 2
    ----------------------------------------
    pNodes:
        ----------------------------------------
        state: ready
        name: 10.10.18.58
        ----------------------------------------
        state: ready
        name: 10.10.18.62
        ----------------------------------------
    zoneId: 0000000057eebaab:0000000000000002
    numNodes: 2
    ----------------------------------------
isClusterZoneCompliant: True
zoneType: logical
isZoneEnabled: True
numZones: 4
AboutCluster Time : 08/22/2019 2:31:39 PM PDT
```

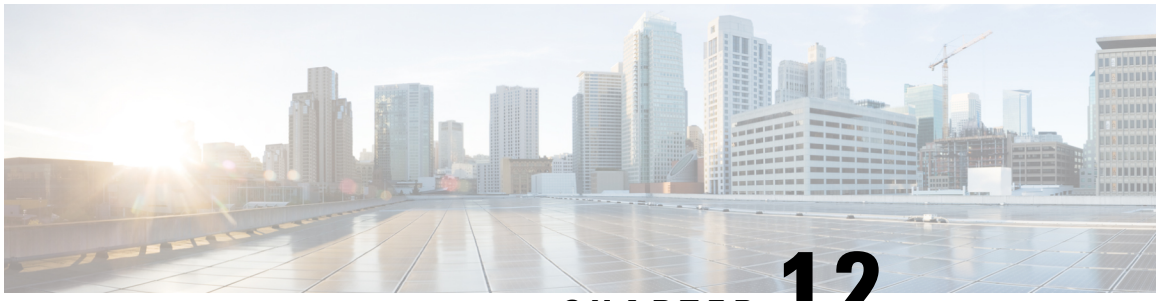
LAZ Related Commands

The following STCLI commands are used for LAZ operations. For more information on CLI commands, see the Cisco HyperFlex Data Platform CLI Guide.

Please be advised to wait at least 10 seconds between successive invocations of LAZ disable and LAZ enable operations in that order.

Command	Description
stcli cluster get-zone	Gets the zone details. This option is used to check if the zone is enabled.
stcli cluster set-zone --zone 0	Enables or Disables zones.

Command	Description
stcli cluster set-zone --zone 1 **stcli rebalance start**	**(Recommended)** Enables and creates zones (default number of zones) **Important** You must execute the **rebalance start** command after you enable and create zones. A cluster created without zoning enabled, will become zone compliant only after enabling zoning and successful completion of rebalance. **Warning** Rebalance is a critical background service. Disabling the service may lead to unexpected behavior including loss of cluster resiliency. Support for this command is limited to Cisco Tech support only. General use is not supported. Triggering rebalance activity may involve large scale data movements across several nodes in the cluster which may decrease the IO performance in the cluster.
stcli cluster set-zone --zone 1 --numzones **<integer-value>** **stcli rebalance start**	Enables zones and creates a specific number of zones. **Important** The number of zones can only be 3, 4, or 5. **Important** You must execute the **rebalance start** command after you enable and create zones. **Warning** Rebalance is a critical background service. Disabling the service may lead to unexpected behavior including loss of cluster resiliency. Support for this command is limited to Cisco Tech support only. General use is not supported.



CHAPTER 12

Managing HX Controller VMs

- Managing Storage Controller VMs, on page 171
- Powering On or Off Storage Controller VMs, on page 171
- Disabling HA VM Monitoring in HX Controller VMs, on page 172

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 `hxcli` commands on the storage cluster.

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.

Procedure

Step 1 **Using the vSphere Web Client** to power on or off a VM.

a) Log into 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** > **Shut Down Guest OS**.

Step 2 **Using the ESX command line** to power on or off a VM.

a) Log into 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.shutdown 1
```

These options will perform a relinquish action for a graceful shutdown versus a hard shutdown which is not desired.

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.

Procedure

Step 1 Log into 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**.

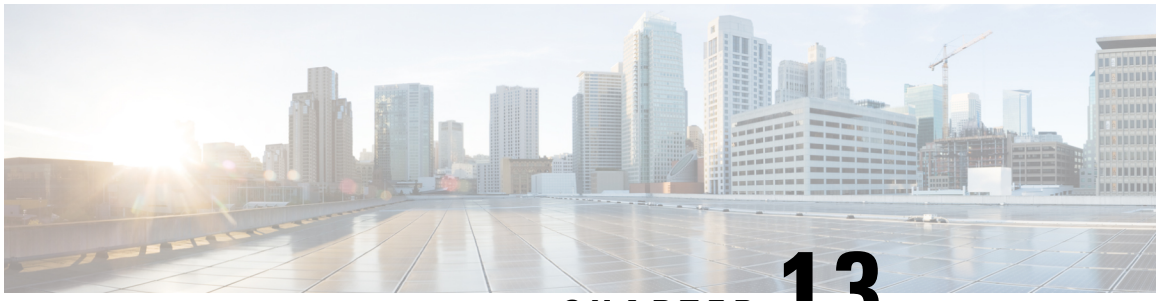
Add VM Override dialog box is displayed.

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 175
- Benefits of HX Data Platform Ready Clones, on page 175
- Supported Base VMs, on page 176
- Ready Clone Requirements, on page 176
- Ready Clone Best Practices, on page 177
- Creating Ready Clones Using HX Connect, on page 177
- Creating Ready Clones Using the HX Data Platform Plug-In, on page 179
- Prepare to Customize HX Data Platform Ready Clones, on page 180
- Configuring Ready Clones Using Customized Specifications, on page 182
- Managing Virtual Machine Networking, on page 182

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.

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 Snapshots is removed.

Note For sentinel based HX snapshot, sentinel snapshots are not automatically deleted after ready clone operation. See the HX Native Snapshots Overview, on page 15 for implications of using sentinel based HX snapshots.

- Storage vMotion is not supported on VMs with HX native snapshots.
- 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.
- VM migration is supported provided that the HX Connect Ready Clones operation has not been performed. If the VM needs to be moved to a different datastore, delete the snapshots first.

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.

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.

Procedure

Step 1 Log into 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. **Note** The VM Name Prefix which is used to name a Ready Clone, must contain only letters, numbers, and the hyphen (-) character. The name must start with a letter and cannot contain only digits or hyphen.
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.

Procedure

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. **Note** The VM Name Prefix which is used to name a Ready Clone, must contain only letters, numbers, and the hyphen (-) character. The name must start with a letter and cannot contain only digits or hyphen.
Starting clone number	Type a clone number for the starting clone.

Control	Description
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. • 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. • If the Customization Specification contains a NIC or network adapter that obtains the IP address using DHCP, then the systems 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.

Procedure

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.

Procedure

Step 1 Verify the VLANs are configured correctly.

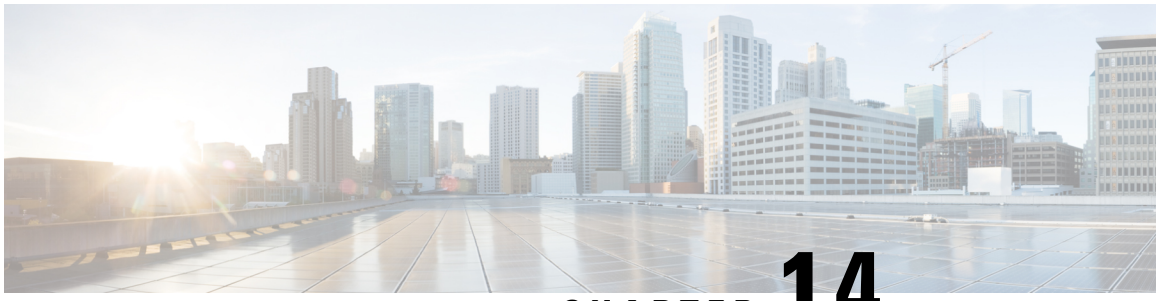
See the VLANs chapter in the *Cisco UCS Manager Network Management Guide* at, Cisco UCS Manager Network Guide.

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 7.0 Documentation* on the vmware site.



CHAPTER 14

Managing HX Native Snapshots

- HX Native Snapshots Overview, on page 183
- Benefits of HX Native Snapshots, on page 184
- HX Native Snapshot Considerations, on page 185
- HX Native Snapshots Best Practices, on page 188
- HX Native Snapshot Time Zones, on page 189
- Creating HX Native Snapshots, on page 190
- HX Native Snapshots using ESXi 7.0 U2, on page 191
- Scheduling HX Native Snapshots Overview, on page 191
- Scheduling HX Native Snapshots, on page 192
- Setting the Frequency of HX Native Scheduled Snapshots, on page 193
- Deleting HX Native Snapshot Schedules, on page 193
- Reverting to an HX Native Snapshot, on page 194
- Deleting HX Native Snapshots, on page 195

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:

- Support for HX native Snapshot in the vSphere client plug-in for HTML 5 was introduced in plugin version 2.0.0. For more information, see Snapshot Now, on page 354.
- Support for Schedule Snapshot the vSphere client plug-in for HTML 5 was introduced in plugin version 2.1.0. For more information, see Schedule Snapshot, on page 357
- The vSphere “Manage Snapshots” function can revert to a specific HX native snapshot, or delete all snapshots.
- Cisco HyperFlex Connect can create on-demand and schedule HX native snapshots.
- The HyperFlex command line user interface can create HX native snapshots.
- HX REST APIs can create and delete HX native snapshots.

- Significant changes in Cisco HXDP Release 5.5(x) and later:
  - ESXi versions 6.5, 6.7 and 7.0 U1 are not supported.
  - VMware VAAI snapshot workflow is used instead of the Sentinel Snapshot Create workflow.

For additional information about VMware snapshots, see the "Overview of virtual machine snapshots in vSphere (KB 1015180)" on the VMware Customer Connect site.

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.

  Improved performance and reliability of Quiesced Snapshot beginning with HyperFlex Release 4.5(2a) and VMware ESXi 7.0 U2.

  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.

  Quiesce based snapshot is not supported on the Windows2008R2 operating system. The operating system is End-of-Life. Please review the VMware documentation for the current list of supported operation systems. VMware Compatibility Guide

  If your system displays quiesce errors, review the *Troubleshooting Volume Shadow Copy (VSS) quiesce related issues (1007696)* article on the VMware Customer Connect site.
- **Scheduled Snapshots are Tolerant to Node Failures -** Scheduled snapshots are tolerant to administrative operations that require a node shutdown, such as HXDP Maintenance Mode and HX online upgrade.

  Scheduled Snapshots are tolerant to failures in other HX clusters in multi-cluster environments.
- **Granular Progress and Error Reporting -** Monitoring is performed at the task level for VM level HX native snapshots.
- **Instantaneous Snapshot Delete -** Deletion of an HX native snapshot and consolidation is always instantaneous.
- **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 snapshotted is on an accessible mount-point.

HX Native Snapshot Considerations

Snapshot Parameters



Attention Beginning with HX Release 4.5(2a) and VMware ESXi 7.0 U2 Sentinel snapshots are not applicable.

- **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. 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.

HX native snapshot does not support for VMs with VMware Fault Tolerance enabled.

- **HX Snapshots compatibility with VMware VAIO -** Creating HX Snapshots with VMware VAIO configured is not supported. Attempting to create HX snapshots will power the VM off. HX Snapshots cannot coexist with Virtual Machines enabled with vSphere APIs for IO Filtering (VAIO). The VAIO framework might be used by backup solutions to enable Continuous Data Protection (CDP) for Virtual Machines. To use the backup solutions with CDP, delete any existing HX snapshots before enabling the CDP functionality.

  To determine if your product uses VMware VAIO framework, review the VMware site list of qualified vendors in the VMware Compatibility Guide.

- **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.

Performance

- **VMware vSphere Storage APIs Array Integration (VAAI) -** To attain the best HX snapshot performance and functionality, upgrade to the ESXi 7.0 U2 or later.

Snapshots During Upgrade Processes

- HX native snapshots are not supported while an upgrade of HX Data Platform, ESXi, or UCS is in progress.

VMs

Table 7: Release Specific VM Considerations

Release	Consideration
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.(2b)	All powered on VMs use asynchronous consolidation (asyncConsolidate = true) when taking HX native snapshots. If the VM is powered off, the settings remain unchanged.
Beginning with HX Release 4.5(2a) using VMware ESXi 7.0 U2 and later	Consolidation time is no longer proportional to the I/O load on the Virtual Machine. Sentinel snapshots are no longer created.

The following considerations apply to all supported releases:

- **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 an HX datastore. HX native snapshots for VMs spanning HX datastores requires HXDP version 4.5(2a) or later with VMware ESXi version 7.0 U2 or later.
- **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.
- **Maximum Snapshot Name Length -**80 characters.

Disks and 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.



Note When creating a new virtual machine disk on the HyperFlex datastore and you want to enable the creation of thick-provisioned disks, there is no option to create a thick-provisioned disk. This is a known issue in VMware. For more information, see Creating VMDK with NFS-backed storage does not allow thick-provisioning with vendor plugin.



Note ESXi cannot distinguish between thick provision lazy zeroed and thick provision eager zeroed virtual disks on NFS datastores. When you use NFS datastores, the vSphere client allows you to create virtual disks in Thick Provision Lazy Zeroed (zeroedthick) or Thick Provision Eager Zeroed (eagerzeroedthick) format. However, when you check the disk type on the Virtual Machine Properties dialog box, the Disk Provisioning section always shows Thick Provision Eager Zeroed as the disk format (no matter which format you selected during the disk creation).

- **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

- When creating large numbers of HX native snapshots consider the following:
  - Schedule HX native snapshots at a time the expected data traffic is low.
  - Stagger HX native snapshot schedules such that a large numbers of VMs are not scheduled to be snapshotted at the same time.
- If vCenter running on a VM in the storage cluster, do not take an HX native snapshot of the vCenter VM. For additional informaion see the *VMware VirtualCenter Server service fails due to a quiesced snapshot operation on the vCenter Server database virtual machine (2003674)* article on the VMware site.

HX Native Snapshots Best Practices HX Release 4.5(1x) and Earlier

Significant updates were introduced in HX Release 4.5(2a) and ESXi 7.0 U2. The following recommendations only apply to users using a release that was introduced before this release.



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 set `snapshot.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.

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 client (HTML5) uses the browser timezone.
- The HX vSphere client (HTML5) 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 client (HTML5) 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 client (HTML5) 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 client (HTML5) 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 client (HTML5) 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.

   ```
   $ hxcli services timezone show
   ```

2. Change the storage cluster time zone.

   ```
   $ hxcli services timezone set --timezone timezone_code
   ```

Related Topics

Schedule Snapshot, on page 357

Creating HX Native Snapshots

To create HX native snapshots, perform the following steps:

Before you begin

Remove any redolog 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. Log into the ESXi host command line.
2. Locate and open the `/etc/vmware/config` file for editing.
3. Set the `snapshot.asyncConsolidate` parameter to `TRUE`.

   ```
   snapshot.asyncConsolidate="TRUE"
   ```

Procedure

Step 1 From the vSphere client (HTML5) Navigator display the list of VMs in vCenter at the VM level. Display the VM list with one of the following methods, **Hosts and Clusters**, **VMs and Templates** , **Storage**, **Networking**, or **Global Inventory Lists**

Example:

Global Inventory Lists > **VMs**

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** > **Take Snapshot** to display the Snapshot dialog box.

Step 4 Complete the dialog box

Table 8: Take Snapshot Dialog Box

Field	Description and Usage Notes
Name	Type the Snapshot name. Maximum snapshot name length: 80 characters.
Description	Type a description of the snapshot.
Snapshot option check box	Use the check boxes to select **Snapshot the virtual machine's memory** or **Quiesce guest file system (Needs VMware Tools installed)**

Step 5 Click **OK** to create an HX native snapshot.

The Recent Tasks tab displays the status message:

```
Create virtual machine native snapshot.
 The first snapshot
```

Related Topics

Snapshot Now, on page 354

HX Native Snapshots using ESXi 7.0 U2

Creating snapshots using ESXi 7.0 U2 introduces the following enhancements:

- Sentinel snapshots are not created.
- Supports VAAI offload for all snapshots of a VM by automatically configuring VM attribute`snapshot.alwaysAllowNative=TRUE`.
- Improved performance, reliability and functionality.
- Supports snapshot of VM on spanned datastores.
- Automatically identifies and deletes sentinels when no longer needed.

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:

Procedure

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.

Related Topics

Schedule Snapshot, on page 357

Scheduling HX Native Snapshots

Procedure

Step 1 From the vSphere client (HTML5) 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

Open the **Schedule Snapshot** dialog box for a VM or resource pool.

Procedure

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.

Related Topics

Schedule Snapshot, on page 357

Deleting HX Native Snapshot Schedules

Procedure

Step 1 From the HX vSphere client (HTML5), 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. 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.

Procedure

Step 1 From the vSphere client (HTML5), 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 9: 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.

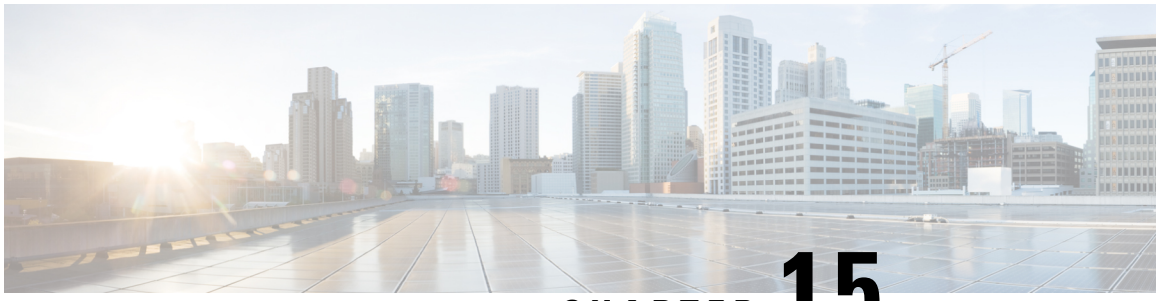
Procedure

Step 1 From the vSphere client (HTML5), 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 Virtual Machine Disaster Recovery

- HX Disaster Recovery Overview, on page 197
- Protecting Virtual Machines Overview, on page 212
- Disaster Recovery Overview, on page 235
- Replication Maintenance Overview, on page 247

HX Disaster Recovery Overview

HyperFlex DR 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 Data Protection (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 , quiescence parameter value, and a common start time.

Setting up DP snapshots requires two existing clusters running a currently supported HX Data Platform Software Release. Both clusters must be on the same HX Data Platform version. Use HyperFlex Connect to complete the setup.

First, set up a local replication network for each cluster. 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 local replication 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 Disaster Recovery Requirements and Considerations

Note Documentation for the N:1 DR for HyperFlex feature is located in the Intersight Help Center. The URL is https://www.intersight.com/help/saas/resources/replication_for_cisco_hyperFlex_clusters.

The following is a list of requirements and considerations when configuring virtual machine replication and performing disaster recovery of virtual machines:

- **Datastore Unmapped Behavior**:

  HXDP Release 5.5(2a) and later: Unmap Other DRO datastore pair is supported when VMs are in an Active (protected) state on either of the site. VMs in a state other than Active (for example Recovered, Recovery_Failed, and Migrate_Failed) need to be moved to the Active (Protected) state for unmap other DRO datastore pair support.

- **Other DRO Datastore Delete Operations**:

  The following guidelines apply to HXDP Release 5.5(2a) and later:

  - Delete operations are supported after upgrading to HXDP Release 5.5(2a)
  - Adding and editing schedules is not supported.
  - Datastore map and unmap is not supported if any DRO datastore unmap happens in parallel.

- **All Protected VMs Tab Usage**: Do not to perform any DR actions available from **All Protected VMs** tab for SRM (OtherDRO) VMs.

- **Post HXDP Release 5.5(2a) and later Upgrade Action:**

  When HyperFlex creates a SRA/SRM pair and protects VMs, SRM creates placeholder VMs on opposite site with same name to reserve the resources for the VM during migration and recovery. Cisco recommends

that after upgrading to HXDP 6.0(1a) you delete the placeholder VMs. Failing to complete this action will result in DR workflows failing due to the VM with same name exists on the opposite site.



Caution **Operations via SRM**: Even though VMs in SRM the environment are available for performing Planned Migration or Disaster Recovery. Cisco does not recommend users perform any kind of operations via SRM.

Admin Role Requirements

You can perform all replication and recovery tasks with administrator privileges on the local cluster. For tasks involving a remote cluster, both the local and remote user must have administrative privileges. You can configure administrative privileges with vCenter SSO on the respective clusters.

Networking Requirements

The replication network should be reliable and have a sustained minimum symmetric bandwidth that is the same as the bandwidth configured in the HyperFlex replication network. Do not share the network with any other applications or traffic on an uplink or downlink. Other requirements are as follows:

Table 10: Networking Requirements

Requirement	Description
Minimum and Recommended Bandwidth	The minimum supported bandwidth is 10 Mbps. Recommended bandwidth is half of the network link bandwidth available for replication. For example, if the network link bandwidth available is 100 Mbps, you should configure the replication bandwidth to be 50 Mbps.
Adaptive Bandwidth Control	Replication network variability may cause network bandwidth to vary and may include the introduction of network errors. Adaptive Bandwidth Control for replication will dynamically adjust the replication speed to scale down if errors are detected and scale up to the configured replication bandwidth limit when the errors are cleared. **Note** Adaptive Bandwidth Control is only in effect when the replication network bandwidth limit is enabled and configured as a non-zero number. Adaptive Bandwidth Control is disabled when the replication network bandwidth limit is not enabled (default). Enabling the replication bandwidth limit requires entering a bandwidth value in the range of 10 to 100,000 Mbit/s. It is recommended that the Administrator always configure a replication network bandwidth limit on both clusters, rather than use the default setting.

Requirement	Description

Measuring Available Replication Network Bandwidth	You can measure the bandwidth of a HyperFlex replication network between two sites by using the iperf utility. In preparation for using the iperf utility, configure the local replication networks on both HyperFlex clusters. After you have configured the local replication networks, you can then pair the HyperFlex clusters. Once you have paired the HyperFlex clusters, map one of the local datastores to a datastore on the remote HyperFlex cluster.
	• Deploy user VMs on both of the mapped datastores. Configure the user VMs with the same network and gateway as used by the respective HyperFlex replication network. The VM can be Ubuntu 16.04 to match the Linux distribution of the HyperFlex storage controller VMs.
	Note
	These VMs are only intended for the purpose of testing network bandwidth. After you have completed testing, you can delete them. There is no need to protect these VMs.
	• Install the iperf utility on both user VMs by running the command:
	`apt get install iperf`
	• Run the iperf server on the user VM deployed on site B by running the command:
	`iperf -s`
	`Example Output:`
	`------------------------------------------------------------`
	`Server listening on TCP port 5001`
	`TCP window size: 85.3 KByte (default)`
	`------------------------------------------------------------`
	Port 5001 should be open between the sites.

Requirement	Description

Measuring Available Replication Network Bandwidth (continued)	• Run the following iperf command on the user VM on site A
	
	`iperf -c <server ip> -i <interval in secs> -t <time in seconds>`
	
	`Example Output:`
	`Client connecting to a.b.c.d TCP port 5001`
	`TCP window size: 85.0 KByte (default)`
	`------------------------------------------------------------`
	`local w.x.y.z port 47642 connected with a.b.c.d port 5001`
	`0.0-10.0 sec  44.8 MBytes  37.6 Mbits/sec`
	`10.0-20.0 sec  222 MBytes  187 Mbits/sec`
	`20.0-30.0 sec  312 MBytes  261 Mbits/sec`
	`30.0-40.0 sec  311 MBytes  261 Mbits/sec`
	`40.0-50.0 sec  312 MBytes  262 Mbits/sec`
	`50.0-60.0 sec  311 MBytes  261 Mbits/sec`
	`60.0-70.0 sec  312 MBytes  262 Mbits/sec`
	`70.0-80.0 sec  312 MBytes  262 Mbits/sec`
	`80.0-90.0 sec  311 MBytes  261 Mbits/sec`
	`90.0-100.0 sec  312 MBytes  262 Mbits/sec`
	`100.0-110.0 sec  311 MBytes  261 Mbits/sec`
	
	Note
	Conduct testing in both directions, from the first paired cluster to the second paired cluster, and then from the second paired cluster to the first paired cluster. If this is a shared link with other applications, perform testing at the time the replication schedules are planned to run. When this link is shared, the available bandwidth for replication could be impacted and may result in congestion on the replication network which may result in packet drops. The HyperFlex replication engine monitors packet drops and throttles the replication traffic if required.

Requirement	Description

Maximum Latency	The maximum replication network latency supported is 75 ms between two paired clusters. There are conditions where it is possible that some replication jobs will encounter error conditions and fail. For example, this may occur when multiple replication jobs execute simultaneously with low network bandwidth and high latency. If this situation occurs, increase the replication network bandwidth, or reduce job concurrency by staggering the number of concurrent replication jobs. If this situation persists, VM unprotect operations may take longer than expected.
	Measuring Replication Network Latency
	You can measure the average replication network latency by running a ping command on any of the storage controller VMs on site A and site B.
	From site A, execute ping command as performed in the following example:
	`ping -I eth2 "Repl IP of any ctlvm on site B" -c 100`
	`Example Output:`
	`100 packets transmitted, 100 received, 0% packet loss, time 101243ms`
	`rtt min/avg/max/mdev = 0.112/0.152/0.275/0.031 ms`
	The average latency should be 75 ms or less.
	Note
	Execute the ping command in both directions, from site A to site B, and from site B to site A.

Network Ports	The comprehensive list of ports required for HyperFlex component communication is located in appendix A of the HX Data Platform Security Hardening Guide. The port/protocl requirements (as of Version 4.5.2a rev 3 dated September 2021) for HyperFlex replication are: ICMP, TCP: 9338, 9339, 3049, 4049, 4059, 9098, 8889, and 9350. **Testing Network Ports** Internal to the HyperFlex cluster, firewall entries are made on the source and destination storage controller VMs during the site pairing operation to allow the HX data platform access to the systems bi-directionally. You need to allow this traffic on WAN routers for each HX node replication IP address and management CIP address. When you configure the local replication network on a HyperFlex cluster, you can manually perform a **Test Local Replication Network** action to test connectivity across the replication IP addresses of each storage controller VM on the local cluster. This test includes port connectivity and firewall checks. When the two clusters have been paired, you can manually perform a **Test Remote Replication Network** action to test connectivity between each local storage controller VM and each remote storage controller VM. Port connectivity and firewall checks are performed. You can also use the Linux “netcat” utility as an additional option to check port connectivity.

Requirement	Description
Network Loss	Reliable transmission of data enables replication between two paired clusters to function optimally. Packet loss in data transmission between two paired clusters may degrade replication performance. **Diagnosing Dropped Packets** There are two cases where packet loss may occur - network congestion and transient network device errors. If dropped packets occur on a replication network due to network congestion the HyperFlex cluster replication engine automatically throttles back replication bandwidth. Throttling replication network bandwidth reduces network congestion and results in the reduction of dropped packets. In extreme cases, replication bandwidth throttling may result in replication jobs taking longer to complete than anticipated. Dropped packets that occur on a replication network due to transient network device errors may cause replication failures that occur randomly or at specific times of the day. Dropped packets are not reported in the HX Connect user interface. Occurrences of packet drop are logged in the HyperFlex storage controller logs. Users that experience noticeable replication job elongation or other failures can contact support for further assistance.

Cluster Requirements

When configuring virtual machine replication and performing disaster recovery of virtual machines, please ensure that the following cluster requirements are met:

Storage Space Requirements

Ensure that there is sufficient space on both clusters to accommodate the retention and processing of replicated DP snapshots. Each protected VM will result in the creation and subsequent replication of a DP snapshot based on the configured schedule interval. The most recent successfully replicated DP snapshot is retained on the destination HyperFlex cluster. Note that for every protected VM, there is a maximum of two DP snapshots present on the source cluster and two DP snapshots present on the destination cluster. This approach facilitates efficient difference-based replication and also assures that the most recent successfully replicated DP snapshot is available for recovery in the event that a newer DP snapshot fails to successfully complete the replication process. Although storage capacity reduction methods are applied, including deduplication and compression, each replicated virtual machine consumes some storage space.

- **Space Consumed by Protected VMs with Redolog Snapshots**—When protecting a VM that also has VMware redolog snapshots, the entire content of the VM is replicated. The entire content includes the VM as well as any retained VMware redolog snapshot(s). This results in increased storage space utilization on both of the paired HyperFlex clusters. When a greater number of redolog snapshots are retained, storage space consumption will also increase.
- **Space Consumed by Protected VMs with HX Native Snapshots**—When protecting a VM that also has HX native snapshots, only the latest VM data is replicated. Retained HX native snapshot data is not replicated. Typically, there is no need to account for space consumed by HX native snapshots on a replication destination HyperFlex cluster.

- **Space Consumed by Deleted VMs**—Deleting a protected VM will not cause space to be reclaimed on the paired HyperFlex cluster datastore. The most recent successfully replicated DP snapshot will be retained to protect the VM from accidental deletion. In order to reclaim space consumed by protected VMs, the VMs must first be unprotected. When a VM is unprotected, the associated DP snapshots are deleted on both paired HyperFlex clusters.
- **Space Consumption Calculations**—The amount of predicted space consumption in addition to the size of a protected VM can be expressed as:

  VM change rate times the number of DP snapshots retained

  The number of DP snapshots retained equals two (2). When a protected VM has VMware redolog snapshots the calculation will be skewed based on the number of retained snapshots.

  Space calculations should also consider that when a protected VM fails over or is migrated to the paired site, the calculations for the source and target can be reversed.
- **Difference Based Replication and Full Copy Replication**—In a typical replication data protection lifecycle, a full copy of a protected VM is replicated in the form of a DP snapshot only once. This full copy replication job occurs when a VM is initially protected. After the initial replication job completes, subsequent replication jobs take advantage of efficient differencing-based technology to replicate only new and changed data.

  You cannot use difference-based technology in the following known corner cases:

  - A protected VM also has HX native snapshots. If the VM is reverted back to a retained HX native snapshot, the next scheduled protection job will perform a full copy replication job instead of a difference-based replication job. An additional full copy worth of space needs to be budgeted on both of the paired clusters.
  - A protected VM undergoes storage vMotion and is migrated to a different datastore. If the destination datastore is mapped to a datastore on the paired cluster, the next scheduled protection job will perform a full copy replication job instead of a difference-based replication job. An additional full copy worth of space needs to be budgeted on both of the paired clusters.
  - A protected VM has a DP snapshot that was taken in conjunction with a replication job. Subsequent to this, an initial HX native snapshot is created that also creates an HX Sentinel snapshot. The next scheduled protection job will perform a full copy replication job instead of a difference-based replication job. An additional full copy worth of space needs to be budgeted on both of the paired clusters.
  - When a protected VM DP snapshot is taken during an HX native snapshot workflow that created intermediate delta disks, the next scheduled protection job will perform a full copy replication job instead of a difference-based replication job. An additional full copy worth of space needs to be budgeted on both of the paired clusters.
  - When a new VMDK is added to an already protected VM, that specific VMDK will be full-copied once.

  Not having sufficient storage space on the remote cluster can cause the remote cluster to reach capacity usage maximums. If you note any **Out of Space** errors, refer to Handling Out of Space Errors for more information. Pause all replication schedules until space available on the cluster has been properly adjusted. Always ensure that cluster capacity consumption is below the space utilization warning threshold.

Supported Configurations

Supported configurations for native replication (NRDR 1:1) are: 2N/3N/4N Edge, FI, and DC-no-FI based clusters to 2N/3N/4N Edge, FI, and DC-no-FI based clusters, including stretched clusters, all managed through HX Connect.

HyperFlex hardware acceleration cards (PID: HX-PCIE-OFFLOAD-1) are supported with native replication beginning with HX 4.5(1a). You must enable HX Hardware Acceleration on both of the paired HyperFlex clusters.

Rebooting Nodes

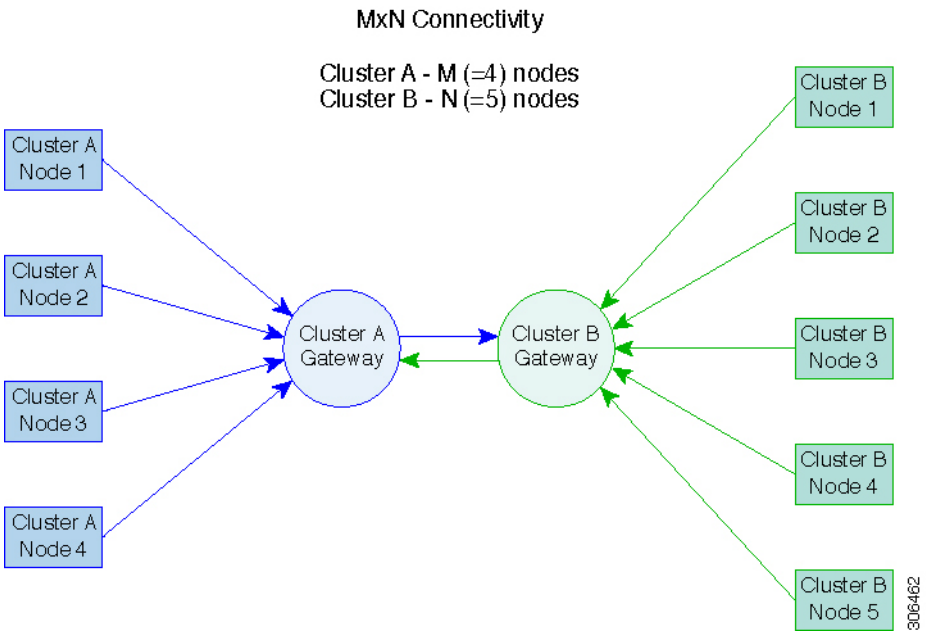
Do not reboot any nodes in an HX cluster during a restore, replication, or recovery operation. Note that node reboot operation may occur as part of an upgrade process. You should pause the replication scheduler prior to an upgrade, and then resume it after the upgrade has completed.

Replication Network and Pairing Requirements

You must establish a replication network between HyperFlex clusters that uses replication for Data Protection (DP) snapshots. The replication network is created to isolate inter-cluster replication traffic from other traffic within each cluster and site. Please also consider the following:

Table 11: Replication Network and Pairing Requirements

Component	Description
HX Data Platform Version	Ensure that the HyperFlex clusters that are going to be paired for replication are running the same HX data platform software version. Note that the use of different HX data platform versions is only supported during HX data platform upgrades. In this scenario, one of the paired HyperFlex clusters may be running a different version of HX data platform software for the period of time until both of the paired clusters have been upgraded. Ensure that you upgrade both of the paired clusters to the same HX data platform version within the shortest possible time frame based on site specific constraints. Also note that a maximum of one major HX data platform release version difference is permitted when upgrading paired clusters. Additionally, the changing of any replication configuration parameter is not supported when the paired clusters are not both running the same HX data platform version during an upgrade.
Node Status	Ensure that all HyperFlex cluster nodes are online and fully operational prior to the creation of the local replication networks and performing the site pairing process.

Component	Description
Node Communication Requirements	Requirements are as follows: • To support efficient replication, ensure that all M nodes of cluster A can communicate with all N nodes of cluster B, as illustrated in the M x N connectivity between clusters graphic. • To enable replication traffic between clusters to cross the site-boundary and traverse the internet, ensure that each node on Cluster A can communicate with each node on Cluster B across the site boundary and the internet. • Isolate the replication traffic from other traffic within the cluster and the data center. For more information, see the graphic below.
M*N Connectivity Between Clusters 	
Node Failure	In the highly unlikely and rare event of a node failure, there may be an impact to replication. As an example, replication jobs in progress will stop if the node which has the replication CIP address enters an inoperative state. At the point in time when the replication CIP address is claimed by another node in the cluster, the replication job will automatically resume. Similarly, if a recovery job was running on the node with replication CIP address and the node failed, the job would fail. The replication CIP address would subsequently be claimed by another node in the cluster. Retry the operation upon noting the failure.
vCenter Recommendations	Ensure that each of the two paired HyperFlex clusters is managed by a unique vCenter instance. Also ensure that vCenter is deployed in a different fault domain for availability during disaster recovery scenarios.

Replication and Disaster Recovery Virtual Machine Considerations

The following are considerations for VMs:

Table 12: Virtual Machine Considerations

Consideration	Description
Thin Provisioning	Protected VMs are recovered with thin provisioned disks irrespective of how disks were specified in the originally protected VM.
VM Device Limitations	Do not protect VMs with connected ISO images or floppies as individually protected VMs, or within a protection group. You can set any configured CD or DVD drive to “Client Device” with the “Connected” state disabled. There is no need to delete the device from the VM configuration. If there is a need to temporarily mount an ISO image, you can unprotect the VM and then protect it again once you have set the CD or DVD drive to “Client Device” and then disconnected.
Protected Virtual Machine Scalability	Beginning with HX Release 4.5(1a): • The sum of protected VMs on all nodes should not exceed the maximum limit of 2000 protected VMs per cluster in a single direction configuration or 1000 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 before using vMotion.
Templates	Templates are not supported with disaster recovery. Do not attempt to protect a template.
Recovery of Virtual Machines with Snapshots	When recovering a protected VM that has VMware redolog snapshots, the VM is recovered and all previous snapshots redolog snapshots are preserved.

Consideration	Description
Data Protection Snapshots	Replicated DP snapshots are stored on the mapped datastore on the paired cluster. You should not perform a manual deletion of DP snapshots as this is not supported. Deleting snapshot directories or individual files will compromise HX data protection and disaster recovery. **Note** To avoid deleting DP snapshots manually, it is important to remember that VMware does not restrict operations on datastores by the administrative user. In any VMware environment, datastores are accessed by an administrative user via the vCenter browser or by logging into the ESXi host. Because of this, the snapshot directory and contents are browsable and accessible to administrators.
VMware SRM – Purposeful VM Deletion	Supported through HXDP Release 5.5(1a). Not supported in HXDP Release 5.5(2x). If a VM is deleted from VMware vCenter and the VM is located in an "Other DRO" datastore pair, the SRM recovery plan for this datastore pair will fail during recovery. To avoid this scenario, first unprotect the VM using the following command on one of the HyperFlex storage controller VMs: stcli dp vm delete --vmid <VM_ID>
VMware SRM – Purposeful VM Deletion	Not supported in HXDP Release 5.5(2x) and later.
VM Naming	If a protected VM is renamed within vCenter, HyperFlex recovers at the previous name folder but registers the VM with the new name on the recovery side cluster. 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.
HyperFlex Software Encryption	Software encryption must be enabled on clusters in both paired datastores to be able to protect VMs on encrypted datastores.

Storage Replication Adapter Overview

The Storage Replication Adapter Feature is not supported in HXDP 5.5(2a) and later.

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, HX Data Platform Cluster Tolerated Failures, on page 10 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.
- VMware Guest Tool for quiesce snapshot in Disaster Recovery is supported. Install the most recent VMware Guest Tool Service or verify that the existing service is current.



Note Using third-party guest tool (open-vm-tools) usage is allowed.

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 with 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 into 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. Starting with Cisco HyperFlex Release 5.0(2a) the MTU field is editable.
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 into 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.

Procedure

Step 1 Log into 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 **VLAN Configuration** 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. **Note** If you are configuring replication network on edge cluster, do not use the **Create VLAN** option. Use the existing VLAN option and follow the same procedure.
VLAN ID field	Click the up or down arrows to select a number for the VLAN ID or type a number in the field. This is separate from the HX Data Platform Management traffic network and Data traffic network. **Important** Be sure to use a different VLAN ID number for each HX Storage Cluster in the replication pair. Replication is between two HX Storage clusters. Each HX Storage cluster requires a VLAN dedicated to the replication network. For example, *3*. 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.
VLAN Name field	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.
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.	

UI Element	Essential Information
UCS Manager host IP or FQDN field	Enter Cisco UCS Manager FQDN or IP address. For example, *10.193.211.120*.
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	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.
	`Format example:`
	`x.x.x.x/<number of bits>`
	`1.1.1.1/20`

Gateway field	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. For example, *1.2.3.4*. **Note** The gateway IP address must be accessible even if the disaster recovery is being setup for a flat network.
IP Range field	Enter a range of IP addresses for use by the replication network. • The minimum number of IP addresses required is the number of nodes in your HX Storage Cluster plus one more. For example, if you have a 4 node HX Storage Cluster, enter a range of at least 5 IP addresses. • The **from** value must be lower than the **to** value. For example, *From 10.10.10.20 To 10.10.10.30*. • 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. **Note** The IP address range excludes compute-only nodes.
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	Click the check box to enable setting the replication bandwidth limit. Enter the maximum network bandwidth that the replication network is allowed to consume for inbound and outbound traffic. This a value in the range of 10 to 100,000 Mbps. Failure to enable replication bandwidth limiting disables Adaptive Bandwidth Control. This is not recommended as replication network variability may cause bandwidth-related replication errors. The replication bandwidth is used to copy replication snapshots from this local HX Storage Cluster to the paired remote HX Storage Cluster. **Note** • 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	Default MTU value is 1500. Select the check box to set a custom MTU size for the replication network. MTU can be set in the range 1024 to 1500. **Note** • Use the same MTU value on both of the paired HX clusters. • Starting with HXDP Release 5.0(2a) and later, you can edit the MTU value after configuring the cluster. With older versions of HXDP, 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:

Procedure

Step 1 Log into 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 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 mininum of five IP addresses are required.

To edit the replication network to add IP addresses, perform the following tasks:

Procedure

Step 1 Log into 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

Replication Network Subnet field	Subnet for the replication network. The subnet that is configured for the replication network in network prefix notation. This value cannot be edited.
	`Format example:`
	`p.q.r.s/<length>`
	`209.165.201.0/27`

Gateway field	The gateway that is configured for the replication network. This is value cannot be edited.
IP Range field	Enter a range of IP addresses for use by the replication network. • The minimum number of IP addresses required is the number of nodes in the HX Storage Cluster plus one more. For example, if the HX Storage Cluster has 4 nodes, the IP Range must be at least 5 IP addresses. • The **from** value must be lower than the **to** value. For example, *From 10.10.10.20 To 10.10.10.30*. • 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. **Note** The IP address range excludes compute-only nodes.
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)	Enter the maximum network bandwidth that the replication network is allowed to consume for outbound traffic. Valid Range: 10 to 10,000. The default is `unlimited`, which sets the maximum network bandwidth to the total available replication network bandwidth. The replication bandwidth is used to copy DP snapshots from the local HX Storage Cluster to the paired remote HX Storage Cluster.

UI Element	Essential Information
Set non-default MTU check box	Default MTU value is 1500. Select the check box to set a custom MTU size for the replication network. MTU can be set in the range 1024 to 1500. **Note** • Use the same MTU value on both of the paired HX clusters. • Starting with HXDP Release 5.0(2a) and later, you can edit the MTU value after configuring the cluster. With older versions of HXDP, the existing replication network configuration will need to be removed. The replication network can then be configured with the correct MTU value.

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 cluster1 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. 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. Below are the prerequisites and steps to create a replication pair.

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.
- Create an encrypted datastore on the remote cluster to protect the encrypted datastore on local site.

Note Software encryption must be enabled on clusters in both paired datastores to be able to protect VMs on encrypted datastores.

- Configure the replication network.

Procedure

Step 1 From HX Connect, log into either the local or remote HX cluster as a user with administrator privileges and do one of the following:

a) Select **Replication** > **Pair Cluster** if you are doing cluster pairing for the first time.
b) 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: *10.10.10.10*.
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:

a) 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.
b) From the **Remote Datastore** pull-down menu, choose a datastore that needs to be paired with the local datastore.
c) 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:

a. From the **Remote Datastore** pull-down menu, choose **Do not map this datastore** to remove the mapping from the current local datastore.

b. 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:

a) 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.
b) 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:

Procedure

Step 1 Log into 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	Default MTU value is 1500. MTU can be set in the range 1024 to 1500. **Note** • Starting with HXDP versions 5.0(2a) and later, you can edit the MTU value after configuring the cluster. With older versions of HXDP, the existing replication network configuration will need to be removed. The replication network can then be configured with the correct MTU value.

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.



Note Datastores with same encryption property can be mapped.

Procedure

Step 1 Log into 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:

a) 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.

b) From the **Remote Datastore** pull-down menu, choose a datastore that needs to be paired with the local datastore.

c) Click **Map Datastores**.

To change the local datastore selection:

a. From the **Remote Datastore** pull-down menu, choose **Do not map this datastore** to remove the mapping from the current local datastore.

b. 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:

a) 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.

b) 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)**.

Removing a Peer Cluster

The preferred method for removing a paring relation for any reason is via HxConnect. In the event that it's necessary to unpair the clusters using the **stcli dp peer delete** command. The **stcli dp peer delete** command is a 2-cluster operation and removes pairing from both clusters.

In a situation where Cluster A and B were paired, and cluster B is permanently down, or unavailable for an extended period of time, it may be necessary to remove the pairing relation on cluster A the proper solution is to use the **stcli dp peer forget --pair-name** on cluster A.

To remove a peer cluster using the **stcli dp peer delete**:

Procedure

Run the **stcli dp peer delete** on one of the clusters in a pair to ensure that the pairing relation is removed from both clusters in the pair.

When successful, both the clusters are available for fresh configuration of data protection.

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 into 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.

Procedure

Step 1 Log into 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. **Note** The lock/unlock icon next to the datastore name indicates whether the dateastore encryption is enable or disabled: • Locked icon: encryption enabled • Unlocked icon: encryption disabled If encrypted local datastores are selected then only encrypted remote datastore information is displayed.
Remote Datastore column	Pair the datastores between the HX clusters. 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.
- Install the most recent VMware Guest Tool Service or verify that the existing service is current.

Procedure

Step 1 Log into 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.

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	Select this radio button if you want to set a specific time for the first replication operation to start. Before you start replication ensure: • At least one virtual machine is added to the protection group. • The scheduled start time is reached. To specify the protection operation start time: **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. **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 *hours, minutes from now* 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.

Quiescence Overview

Quiescence 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.

HX Data Protection (DP) 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 (UI), and HX

REST APIs. VMware tools should be installed in the guest VM when creating HX DP snapshots using the **quiesce** option. For information on VMware, go to the VMware Website for the following:

- VMware Compatibility Guide.
- VMware Tools Documentation
- Virtual Machine Tools, Version, and Status.
- VMware Guest Operating System Installation Guide

HXDP Software Release 5.0(2a) and earlier supports the following guest states:

- guestToolsCurrent
- guestToolsUnmanaged

When the quiesce data protection snapshot fails, the **DataProtectionVmError** occurs which prompts an HX event and an HX alarm.

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:

Procedure

Step 1 Log into 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.

Procedure

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.

Procedure

Step 1 Log into 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.

UI Element	Essential Information
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 theremote 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 oir 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.

Procedure

Step 1 Log into 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.
Start protecting the virtual machines at radio button	Select this radio button if you want to set a specific time for the first replication to start. To start replication requires: • At least one virtual machine is added to the protection group. • The scheduled start time is reached. To specify the protection operation start time: **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. The *hours, minutes from now* states when the first replication will occur. This is updated when you change the time field setting. **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.
Use VMware Tools to quiesce the virtual machine check box	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. This only applies to virtual machines with VMware Tools installed.

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.

Procedure

Step 1 Log into 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	Select an existing protection group from the pull-down list. The interval and schedule settings of the protection group are applied to this virtual machine. No additional configuration is required, click **Protect Virtual Machine**.
Protect this virtual machine independently radio button	Enables the interval, schedule options, and VMware Tools quiescence option for defining protection for this VM.

UI Element	Essential Information
Protect this virtual machine every field	Select from the pull-down list how often the virtual machines are to be replicated to the paired cluster. 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.
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 to start. To start replication requires: • At least one VM is added to the protection group. • The scheduled start time is reached. To specify the protection operation start time: **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. The *hours, minutes from now* states when the first replication will occur. This is updated when you change the time field setting. **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.
VMware Tools to quiesce the virtual machine check box	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. This only applies to virtual machines with VMware Tools installed.

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 247.

Procedure

Step 1 Log into 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.

Software encryption must be enabled on clusters in both paired datastores to be able to protect VMs on encrypted datastores.

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.

Procedure

Step 1 Log into 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.
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	Map the resources in the protected site to the resources in the remote site for the recovery, test recovery, and migrate configuration. Check the **Same as Recovery Configuration** check box to use the resource mapping of recovery configuration for the test recovery configuration. 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.
Folder area	Map the folders in the protected site to the folders in the remote site for the recovery, test recovery, and migrate configuration. Check the **Same as Recovery Configuration** check box to use the folder mapping of recovery configuration for the test recovery configuration. 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.
Network area	Map the network in the protected site to the network in the remote site for the recovery, test recovery, and migrate configuration. Check the **Same as Recovery Configuration** check box to use the network mapping of recovery configuration for the test recovery configuration. 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.

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 for Disaster Recovery Operations

As previously stated in the Replication Network and Pairing Requirements section, the use of different HX data platform versions is only supported during HX data platform upgrades. During the period of time until both of the paired clusters have been upgraded, the changing of any replication configuration parameter is not supported. The test-recover, recover, re-protect, and planned migration operations are expected to function during the period of time until both of the paired clusters have been upgraded. In some cases, the use of the command line user interface may be required to complete the re-protect and planned migration operations.

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:

Procedure

Step 1 Log into HX Connect on the target cluster as administrator.

Step 2 Navigate to **Replication** > **Remote VMs Tab** > ***protected_vm***.

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: • 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.
Test Networks radio button	Select which HX Storage Cluster network to use for transferring the data from the replication snapshot. Network options for example: • Storage Controller Data Network • Storage Controller Management Network • Storage Controller Replication Network • VM Network

UI Element	Essential Information
Map Networks radio button	Select to create a map between the source and the target cluster networks. • 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.
- Cancelling a recovery process (rolling back) is not supported. Attempting to cancel a recovery process changes all VMs in an unprotected ‘ready to recovery’ state.
- 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.

Procedure

Step 1 Log into 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.
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	Select to create a map between the source and target cluster networks. • 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. Network options for example: • 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

Procedure

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 239 for more details.

Planned Migration

Performing a 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.

To perform a planned migration, take the following steps:



Attention This process cannot be rolled back.

Procedure

Step 1 Log into 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.

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	Select to create a map between the source and target cluster networks. • 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. Network options for example: • Storage Controller Data Network • Storage Controller Management Network • Storage Controller Replication Network • VM Network

Step 4 Monitor the progress on the **Activity** page.

Low Bandwidth and Temporary Packet Loss - In the event VM migration operation fails with an error message that contains "THRIFT_EAGAIN (timed out)", retry the VM Migration. The timeout error is due to temporary network congestion caused by bandwidth saturation or underlying network packet loss.

Planned Migration for a Single vCenter Deployment

To perform a planned migration for a single vCenter deployment, take the following steps:



Attention This process cannot be rolled back.

Procedure

Step 1 Using the Web CLI, run the following command to prepare for failover on the source:

```
# stcli dp vm prepareFailover --vmid <VMID>
```

Note
You can use the `stcli dp vm list --brief` command to determine the VMID of a protected VM.

The task ID is returned.

Step 2 Log into 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**.

Step 3 Log into HX Connect of the secondary site. Select **Replication** > **Remote VMs Tab** > ***protected_vm***. Click **Migrate**.

Step 4 After the Migrate task has completed successfully, log into vSphere Web Client of the secondary site and manually register the VM.

a) Log into 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.

Low Bandwidth and Temporary Packet Loss - In the event VM migration operation fails with an error message that contains "THRIFT_EAGAIN (timed out)", retry the VM Migration. The timeout error is due to temporary network congestion caused by bandwidth saturation or underlying network packet loss.

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.

Procedure

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 241 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.

  1. Log into 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 into HX Connect of the secondary site. Select **Replication** > **Remote VMs Tab** > ***protected_vm***. Click **Recover**.

  3. When the primary site comes back up, log into HX Connect of the secondary site. Select **Replication** > **Remote VMs Tab** > ***unprotected***. Click **Re-protect**.

  4. After Re-protect has completed successfully, log into vSphere Web Client of the secondary site and manually register the VM.

     a. Log into vSphere Web Client Navigator. Select **Configuration** > **Storage**.

     b. Right-click on the appropriate data store 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.

Procedure

Step 1 Log into 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 >** > ***unprotected*** 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	Select to create a map between the source and target cluster networks. • 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. Network options for example: • 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** > ***unprotected***.

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:

a) On the target cluster, select **Replication** > **Remote VMs Tab** > ***unprotected***.

b) 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.

Important Usage Guideance: Starting with Cisco HyperFlex Release 5.0(2b) users should review the following use case before proceeding.

The **stcli dp peer forget --pair-name** operation is a single cluster operation and only affects the cluster where the command is executed. The **stcli dp peer delete** is a 2-cluster operation and removes pairing from both clusters.

In a situation where Cluster A and B were paired, and cluster B is permanently down, or unavailable for an extended period of time, it may be necessary to remove the pairing relation on cluster A the proper solution is to use the **stcli dp peer forget --pair-name** on cluster A.

Procedure

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 239 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 234 for more details.

Step 4 Use STCLI to clean-up the protection group data.

```
Remove Protection group (if any)
stcli dp group list
stcli dp group delete --groupid <groupUUID>
```

Note
GroupUUID is the vmGroupEr.id from the group list command.

Group delete is not supported in HX connect for remote cluster. Use stcli.

Step 5 (Optional) If needed, use the **stcli drnetwork cleanup** command to change the DR network. For more information see the Cisco hyperFlex Data Platform CLI Guide for your release.

Step 6 Pair to the new cluster. See the Creating a Replication Pair, on page 220section for more details.

Step 7 Protect the virtual machines.

Removing Protection from an Auto-Protected Cluster VM

In the event that the vCLS vms are not showing in Virtual machines page, but still they are getting auto protected, you can perform the following steps to remove protection from the auto-protected cluster VM.

Before you begin

- VSphere Cluster Services (vCLS) VM should not reside on Backup or DR/SRM datastores.
- Do not place vCLS VMs on HX datastores that are intended for 1:1 DR or N:1 Backup functionality.

Procedure

Step 1 Unprotect the VM using the `stcli dp vm delete <vmid>` cli.

Step 2 Use VCenter to Storage VMotion the VM to a different datastore.

Replication Maintenance Overview

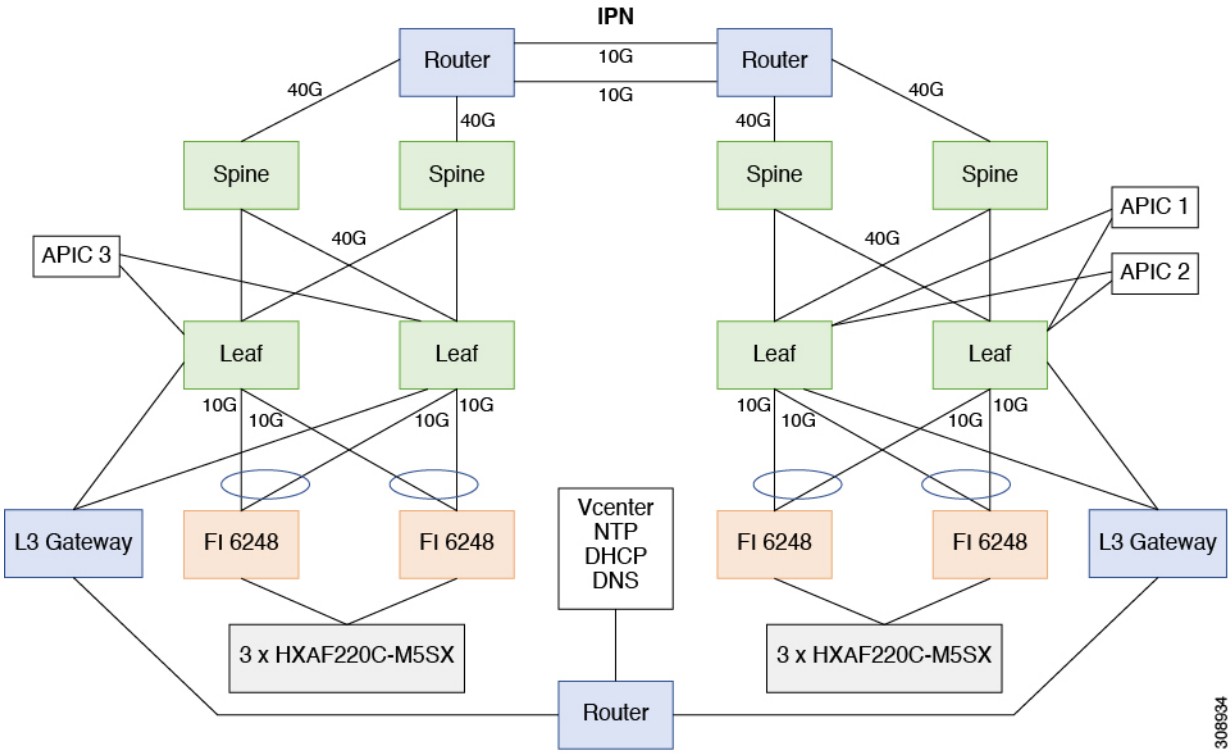
Replication, when configured, runs in the background according to defined schedules. Replication maintenance tasks include:

- **Testing recovery—**Testing if the recovery methods are working. See Testing Virtual Machine Recovery, on page 237 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.

Procedure

Step 1 Log into a Storage Controller VM.
Step 2 From the command line, run the `stcli dp schedule pause` 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.

Procedure

Step 1 Log into a Storage Controller VM.
Step 2 From the command line, run the `stcli dp schedule resume` command.

The previously configured replication schedule for all the protected virtual machines begins.

Replication Page

Displays summary information and links to detailed information related to Replication Configuration, Local Protection, and Remote Protection.

Replication Configuration Ribbon

UI Element	Essential Information
REPLICATION CONFIGURATION field	Displays the state of the replication network configuration. • **Replication network not configured**—The replication network has not been configured. Click **Configure** to begin. • **Network configured**—The replication network is configured. Click **Edit** to adjust replication network IP ranges or bandwidth limit.

UI Element	Essential Information
BANDWIDTH LIMIT field	Displays the configured bandwidth allowed for transmitting incoming and outgoing replication data. • **Blank**—The replication network is not configure. • **# Mbps**—Configured setting in Mega bits per second (Mbps). • **Maximum—**The default setting. Allows the replication network to use the total available network bandwidth. Click **Edit** to change the bandwidth limit.
Bandwidth chart	Displays the bandwidth being consumed for data being replicated between this HX Storage Cluster and the paired HX Storage Cluster. Vertical axis is bandwidth and horizontal axis is time. For more details, see the Performance Page.
Actions drop-down list	Click to create or edit the replication network for this HX Storage Cluster and to test the replication network. • **Test Local Replication Network**—<define> • **Edit Replication Network**—Edit IP Range and set replication bandwidth limit.

Recovery Settings Ribbon

UI Element	Essential Information
Cluster Pairing field	Displays the name of the cluster pair. • Click **Pair Cluster** that appears when cluster pairing is not done, to initiate cluster pairing. • Click **Create Replication Pair** to initiate cluster pairing. The **Create Replication Pair** option is displayed only when you delete an existing replication pair after unprotecting all the VMs and removing all the dependencies.
DATASTORE MAPPED field	Displays the number of mapped datastores. • Click **Map Datastore Pairs** that appears when datastore mapping is not done, to map one local datastore to one remote datastore.
RECOVERY SETTINGS field	Displays the state of the recovery settings configuration. • Click **Configure** that appears when recovery settings configuration is not done, to configure the recovery settings to return the network to a known working state during recovery or test recovery.

UI Element	Essential Information
Actions drop-down list	Choose any one of the action to perform specific operation on replication network, recovery settings, and datastore mapping. • **Test Remote Replication Network**—To test the pairing between clusters in a remote replication network. • **Validate Recovery Settings**—To validate the configured recovery settings. • **Edit Recovery Settings**—To edit the recovery settings configuration. • **Edit Datastore Mapping**—To edit the mapping between local and remote datastores.

Local/Remote Protection Summary Ribbons

UI Element	Essential Information
VMs field	Displays the total number of virtual machines configured for protection on the Local cluster or on the Remote cluster. Displays the details for individual virtual machines and virtual machines in protection groups. Click the field to display the list of protected virtual machines in the **Local VMs** or **Remote VMs** tab.
Protected field	Displays the total number of virtual machines that have a replication snapshot created. Click the field to display the list of protected virtual machines in the **Local VMs** or **Remote VMs** tab.
Exceeds Interval field	Displays the number of replications that took longer than the configured interval to complete. For example, if a virtual machine has an interval of every 15 minutes and replicating its snapshot from the local to the remote cluster took 20 minutes, the replication exceeded the interval. Click the field to display the list of virtual machines with exceeded interval in the **Local VMs** or **Remote VMs** tab.
Current Replication Failures field	Displays the current number of replications that did not complete. Click the field to display the list of virtual machines with failed replication in the **Local VMs** or **Remote VMs** tab.
Protection Groups field	Displays the total number of protection groups configured for this HX Storage Cluster. Click the field to display the list of protection groups and their associated VMs in the **Protection Groups** section under the **Local VMs** or **Remote VMs** tab.

The **Replication** page table provides four tabs: **Local VMs**, **Remote VMs**, **Replication Activity**, and **Replication Pairs**. Each of these tabs provide replication protection configuration options.

Replication Activity Tab

UI Element	Essential Information
Virtual Machine column	Name of the virtual machine protected by replication in the HX Storage Cluster.
Remote Cluster column	Name of the corresponding remote cluster associated with the protected virtual machine. This is the recovery cluster for the listed virtual machine.
Status column	Displays the current status of the virtual machine protection on this cluster: • **Success**—The scheduled replication of the virtual machine and its data to the remote cluster is completed. • **Starting**—Replication task is starting. • **In progress**—The replication task is proceeding. • **Failed**—The scheduled replication task did not complete. • **Deleted**—Replication task is deleted. • **Paused**—Replication task is paused.
Start Time column	Displays the timestamp for the most recently started replication process.
End Time column	Displays the timestamp for the most recently completed replication process.
Protection Group column	If the associated virtual machine belongs to a protection group, the protection group name is listed. If it does not have a protection group, the field displays **None**.
Direction column	The direction of the replicated virtual machine. The direction is relative to the local cluster. The cluster you are logged into is always the local cluster. The options are: • **Incoming**—The virtual machine resides on the remote cluster. It is replicated from the remote cluster to the local cluster. • **Outgoing**—The virtual machine resides on the local cluster. It is replicated from the local cluster to the remote cluster.
Data Transferred column	The size (in Bytes) of the virtual machine that is replicated. When the replication is in progress, the amount completed is listed. When the replication is complete, the amount of data transferred is listed in Bytes.

Replication Pairs Tab

UI Element	Essential Information
Name column	Name of this local cluster.
Remote Cluster column	Hostname and IP address of the remote cluster.
Remote Cluster Status column	Status of the remote cluster. Options include: Online, Offline
VMs Outgoing column	The number of virtual machines configured for replication to the remote HX Storage Cluster from this local HX Storage Cluster. Includes the number of protection groups on this local cluster. Click the field to display the VM replications on the **Virtual Machines** page.
Replications Outgoing column	The number of virtual machines being replicated, transferring their data, to the remote HX Storage Cluster from this local HX Storage Cluster.
VMs Incoming column	The number of virtual machines configured for replication from the remote HX Storage Cluster to this local HX Storage Cluster. Includes the number of protection groups on the remote cluster. Click the field to display the VM replications on the **Virtual Machines** page.
Replications Incoming column	The number of virtual machines being replicated, transferring their data, from the remote HX Storage Cluster to this local HX Storage Cluster.
Mapped Datastores Pairs column	The number of datastores used for replication on the local cluster. Click the field to display the list of datastores on the **Datastores** page.
Create Replication Pair button	This button is only available when a replication pair is not configured for this local cluster. Click the button and complete the **Create Replication Pair** dialog box.
Edit button	Select the replication pair and click **Edit** to change the local or remote datastores to use for replication. Complete the **Edit Replication Pair** dialog box.
Delete button	Select the replication pair and click **Delete**. Confirm the action. Perform this operation when you want to remove the pairing of this local cluster from the remote cluster. **Note** All VMs on both clusters lose their replication configuration. To apply protection to the VMs, you must complete all the protection steps, including creating a new replication pair.

Local Virtual Machines Page

Displays detailed information related to local virtual machines.

UI Element	Essential Information
Protection Groups sub table	**+ Create Group** button—Opens **Create Protection Group** dialog box. Lists protection groups created on the local cluster. You can filter the virtual machines by **All Protected VMs** or **Standalone Protected VMs**. Displays the following protection group data: • Group name • Number of VMs in the group • Status of VMs: Protected, Recovered, Recovering, Recovery Failed • Replication interval time, tooltip shows the time of last replication • To edit the group schedule, click the pen icon. To delete the protection group, use the trash icon.
Pause button	Pause outgoing replication stops all ongoing and new virtual machines from being protected to the target site.
Virtual Machine Name column	Lists the names of the virtual machines protected by replication in the HX Storage Cluster.

UI Element	Essential Information
Protection Status column	Displays the current status of the virtual machine protected on this cluster: • **Recovering**—The virtual machine is restoring from a replication snapshot on the remote cluster. VM State—Prepare Failover Started, Prepare Failover Completed • **Recovery Failed**—The virtual machine failed to restore from a replication snapshot on the remote cluster. VM State—Prepare Failover Failed, Failover Failed • **Recovered**—The virtual machine was recently restored from a replication snapshot on the remote cluster. VM State—Failover Completed • **Protecting**—Reverse protect started for that virtual machine. VM State—Prepare Reverse Protect Started, Prepare Reverse Protect Completed, Reverse Protect Started • **Protection Failed**—Reverse protect failed for the virtual machine. VM State—Prepare Reverse Protect Failed, Reverse Protect Failed • **Protected**—The virtual machine has at least one snapshot available for recovery. VM State—Success • **Active**—Protection is configured, but no snapshot is available. VM State—Active • **Exceed Interval**—The last replication process took longer than the configured interval to complete.
Last Protection Time column	Displays the timestamp for the most recently completed replication process.
Direction column	Displays the direction of the replicated virtual machine, relative to the local cluster. The cluster you are logged into is always the local cluster. • **Incoming**—The virtual machine resides on the remote cluster. It is replicated from the remote cluster to the local cluster. • **Outgoing**—The virtual machine resides on the local cluster. It is replicated from the local cluster to the remote cluster.
Protection Group column	If the associated virtual machine belongs to a protection group, the protection group name is listed. If it does not have a protection group, the field displays **None**.

UI Element	Essential Information
Interval column	Displays the length of time between the start of each replication. Select an Interval time sufficient to complete each replication. For example, an Interval time of `Every 1 hour` means that a replication of the VM is started every hour.
Edit Schedule button	Select an individually protected VM and click **Edit Schedule** to modify the replication interval.
Remove from Group button	Select one or more VMs from the same protection group and click **Remove from Group** to remove the selected VMs from the group. The selected VMs continue to be protected individually with the same replication schedule as the protection group. Click **Remove from Protection Group** to confirm.
Add to Group button	Click to add virtual machines that are protected to a group. VM schedule is changed to be a group schedule.
Unprotect button	Select an individually protected VM and click **Unprotect** to remove replication protection from the VM. Unprotecting prevents a replication snapshot from starting. Click **Unprotect** to confirm. The VM is removed from the list. **Note** Unprotecting removes protection for the selected VMs. To protect the VMs, you must reapply replication configuration.

Remote Virtual Machines Page

Displays detailed information that is related to remote virtual machines.

UI Element	Essential Information
Protection Groups sub table	**+ Create Group** button—Opens Create Protection Group dialog box. Lists protection groups that are created on the remote cluster. You can filter the virtual machines by **All Protected VMs** or **Standalone Protected VMs**. Displays the following protection group data: • Group name • Number of VMs in the group • Status of VMs: Protected, Recovered, Recovering, Recovery Failed • Replication interval time, tooltip shows the time of last replication.

UI Element	Essential Information
Virtual Machine Name column	Displays the name of the virtual machine that is protected by replication in the HX Storage Cluster.
Protection Status column	Displays the current status of the virtual machine protection on this cluster: • **Recovering**—The virtual machine is restoring from a replication snapshot on the remote cluster. VM State—Prepare Failover Started, Prepare Failover Completed • **Recovery Failed**—The virtual machine failed to restore from a replication snapshot on the remote cluster. VM State—Prepare Failover Failed, Failover Failed • **Recovered**—The virtual machine was recently restored from a replication snapshot on the remote cluster. VM State—Failover Completed • **Protecting**—Reverse protect started for that virtual machine. VM State—Prepare Reverse Protect Started, Prepare Reverse Protect Completed, Reverse Protect Started • **Protection Failed**—Reverse protect failed for the virtual machine. VM State—Prepare Reverse Protect Failed, Reverse Protect Failed • **Protected**—The virtual machine has at least one snapshot available for recovery. VM State—Success • **Active**—Protection is configured, but no snapshot is available. VM State—Active • **Exceed Interval**—The last replication process took longer than the configured interval to complete.
Last Protection Time column	Displays the timestamp for the most recently completed replication process.
Direction column	Displays the direction of the replicated virtual machine, relative to the local cluster. The cluster that you are logged into is always the local cluster. • **Incoming**—The virtual machine resides on the remote cluster. It is replicated from the remote cluster to the local cluster. • **Outgoing**—The virtual machine resides on the local cluster. It is replicated from the local cluster to the remote cluster.
Protection Group column	If the associated virtual machine belongs to a protection group, the protection group name is listed. If it does not have a protection group, the field displays **None**.

UI Element	Essential Information
Interval column	Displays the length of time between start of each replication. Select an interval time sufficient to complete each replication. For example, an interval time of `Every 1 hour` means that a replication of the VM is started every hour.
Recover button	Select a VM and click **Recover**, to take the most recent replication snapshot of the VM and build a new VM on the local cluster. Ensure that the VM on the remote cluster is unavailable. Complete this step after **Unprotect**, for recovering a VM.
Migrate button	Select a VM and click **Migrate**, to migrate the protected VM from the source to the target.
Unprotect button	Select an individually protected VM and click **Unprotect** to remove replication protection from the VM. Unprotecting prevents a replication snapshot from starting. Complete this step prior to **Recover on Cluster**, for recovering a VM.
Re-protect button	Select an individually unprotected VM and click **Re-protect** to reprotect the VM.
Test Recovery button	Select a VM and click **Test Recovery**, to take the most recent replication snapshot of the VM and builds a new VM on the local cluster.

Prepare to Protect Virtual Machines Alert

The replication network and a replication pair must be configured before virtual machines can be protected.

To perform the following tasks, you must be logged in as a user with administrator privileges.

1. Create a datastore on the local and the remote storage cluster. For each cluster, on the **Datastores** tab, click the **Create Datastore** button.

   Create one or more datastores on the local cluster, then log into the remote cluster and create datastores there.

2. Configure the replication network on both the local and remote cluster. For each cluster, on the **Replication** tab, click the **Configure** button.

   Complete the configuration on the local cluster, then log into the remote cluster and complete the configuration there.

3. Configure the replication pair between the local and remote storage clusters. Select **Replication** > **Pair Cluster**.

   Map datastores between the local and the remote storage clusters. Mapping datastores can be completed from either the local or the remote storage cluster.

Configure or Edit Replication Network Dialog Box

Configure Replication Network



Note To perform this task, you must be logged in as a user with administrator privileges.

You are required to configure a replication network and a replication pair before protecting virtual machines.

Configure the replication network both on the local and remote cluster. Configure replication network on the local cluster first, then log into the remote cluster and complete the configuration.

1. Select **Replication** > **Configure Network**.
2. Complete the following fields in the **VLAN Configuration** tab:

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. **Note** If you are configuring replication network on edge cluster, do not use the **Create VLAN** option. Use the existing VLAN option and follow the same procedure.
VLAN ID field	Click the up or down arrows to select a number for the VLAN ID or type a number in the field. This is separate from the HX Data Platform Management traffic network and Data traffic network. **Important** Be sure to use a different VLAN ID number for each HX Storage Cluster in the replication pair. Replication is between two HX Storage clusters. Each HX Storage cluster requires a VLAN dedicated to the replication network. For example, *3*. 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.
VLAN Name field	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.
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.	

UI Element	Essential Information
UCS Manager host IP or FQDN field	Enter Cisco UCS Manager FQDN or IP address. For example, *10.193.211.120*.
Username field	Enter administrative username for Cisco UCS Manager.
Password field	Enter administrative password for Cisco UCS Manager.

Click **Next**.

3. Complete the following fields in the **IP & Bandwidth Configuration** tab:

IP & Bandwidth Configuration Tab

UI Element	Essential Information

Replication Network Subnet field	Enter the subnet for use by the replication network in network prefix notation. This is separate from the HX Data Platform Management traffic network and Data traffic network.
	`Format example:`
	`p.q.r.s/<length>`
	`209.165.201.0/27`

Gateway field	Enter the gateway for use by the replication network. This is separate from the HX Data Platform Management traffic network and Data traffic network. For example, *1.2.3.4*.
IP Range field	Enter a range of IP addresses for use by the replication network. • The minimum number of IP addresses required is the number of nodes in your HX Storage Cluster plus one more. For example, if you have a 4 node HX Storage Cluster, enter a range of at least 5 IP addresses. • The **from** value must be lower than the **to** value. For example, *From 10.10.10.20 To 10.10.10.30*. • 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.
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	Enter the maximum network bandwidth that the replication network is allowed to consume for outbound traffic. Acceptable value is between 10 and 10,000. The default value is `unlimited`, which sets the maximum network bandwidth to the total available to the network. The replication bandwidth is used to copy replication snapshots from this local HX Storage Cluster to the paired remote HX Storage Cluster.

4. Click **Configure**.

Edit Replication Network



Note To perform this task, you must be logged in as a user with administrator privileges.

Add available IP addresses to the configured replication network. There must be one IP address for every node in the storage cluster, plus one more for management. If you expand your storage cluster, available IP addresses are consumed.

1. Select **Replication** > **Actions** ***drop-down list*** > **Edit Configuration**.
2. 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, gateway, and VLAN ID are displayed for reference only and cannot be edited.

Edit Network Configuration Dialog Box

UI Element	Essential Information

Replication Network Subnet field	Subnet for the replication network. The subnet that is configured for the replication network in network prefix notation. This value cannot be edited.
	`Format example:`
	`p.q.r.s/<length>`
	`209.165.201.0/27`

Gateway field	The gateway that is configured for the replication network. This is value cannot be edited.

UI Element	Essential Information
IP Range field	Enter a range of IP addresses for use by the replication network. • The minimum number of IP addresses required is the number of nodes in the HX Storage Cluster plus one more. For example, if the HX Storage Cluster has 4 nodes, the IP Range must be at least 5 IP addresses. • The **from** value must be lower than the **to** value. For example, *From 10.10.10.20 To 10.10.10.30*. • 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. **Note** The IP address range excludes compute-only nodes.
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)	Enter the maximum network bandwidth that the replication network is allowed to consume for outbound traffic. Valid Range: 10 to 10,000. The default is `unlimited`, which sets the maximum network bandwidth to the total available replication network bandwidth. The replication bandwidth is used to copy DP snapshots from the local HX Storage Cluster to the paired remote HX Storage Cluster.
Set non-default MTU check box	Default MTU value is 1500. Select the check box to set a custom MTU size for the replication network. MTU can be set in the range 1024 to 1500. **Note** • Use the same MTU value on both of the paired HX clusters. • Starting with HXDP Release 5.0(2a) and later, you can edit the MTU value after configuring the cluster. With older versions of HXDP, the existing replication network configuration will need to be removed. The replication network can then be configured with the correct MTU value.

3. Click **Save Changes**.

The replication network is now updated. Added IP addresses are available for new nodes when they are added to the storage cluster. Replication traffic adjusts to any changes made to the bandwidth limit.

Prepare Group Recovery Dialog Box

Caution Complete this action only in the event of a disaster.

Prepare group recovery stops the replication schedule for all the Virtual Machines in the protection group. After the replication schedule is stopped for all the VMs, proceed to the Standalone VM tab and recover each VM.

Recover VM on This Cluster Dialog Box

To recover the VM and build a new VM on the local cluster, click the **Recover VM** button.

Note All the fields 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	Choose if the recovered VM must be powered on or left powered off after it is created.
Map Networks	Select to create a map between the source and target cluster networks. • Source Network—the network on the cluster with the VM replication snapshot. • Target Network—the network on the cluster where the new VM is created. Network options include: • Storage Controller Data Network • Storage Controller Management Network • Storage Controller Replication Network • VM Network

Click **Recover VM**.

Test Recovery Parameters Dialog Box

To test the recovery process, click the **Recover VM** button.

Note All the fields listed here are optional.

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: • Discovered Virtual Machine • ESX Agents • HX Test Recovery
Power On/Off radio button	Click a button. The recovered VM is powered on or left off after it is created.
VM Name field	Enter a new name for the created test VM.
Test Networks radio button	Select which HX Storage Cluster network to use for transferring the data from the replication snapshot. Network options include: • Storage Controller Data Network • Storage Controller Management Network • Storage Controller Replication Network • VM Network
Map Networks radio button	Select to create a map between the source and target cluster networks. Source—the cluster with the VM replication snapshot. Target—the cluster where the test VM is created.

Click **Recover VM**.

Protected Virtual Machines Tab

Displays the virtual machine protection status, you can edit the protection schedule and unprotect virtual machines. To select virtual machines for protection see the **Virtual Machines** page.

Protected Virtual Machines Actions

UI Element	Essential Information
Edit Schedule button	Change the replication interval or VMware Tools quiesce setting for the replication of the selected virtual machine. Select the virtual machine and click **Edit Schedule**.

UI Element	Essential Information
Unprotect button	To unprotect a virtual machine: **1.** Select the **Replication** > **Protected Virtual Machines** tab. **2.** Select one or more virtual machines that resides on the local cluster with outgoing protection configured. Independently protected virtual machines must be selected one at a time. Multiple virtual machines selected must be in the same protection group. **3.** Select the virtual machine name and click **Unprotect**. **4.** Repeat for virtual machines in another protection group or independently protected. To move a virtual machine from one protection group to another: **1.** Unprotect the virtual machine. From the **Replication** > **Protected Virtual Machines** tab, select the virtual machine and click **Unprotect**. This removes all protection for the virtual machine. **2.** Re-protect the virtual machine selecting the new protection group. From **Virtual Machines**, select the virtual machine and click **Protect**.

Protected Virtual Machines Table

UI Element	Essential Information
# selected column	The number of virtual machine checkboxes selected from the table. Actions performed are applied to all virtual machines selected.
Virtual Machine Name column	Name of the virtual machine protected by replication in the HX Storage Cluster.

UI Element	Essential Information

Protection Status column	The most recent protection action on the virtual machine protection. The arrow on the status indicates the direction of the data transmission.
	
	The directional arrows indicate data transmission:
	• **Left to Right**—From the local cluster to the remote cluster.
	• **Right to Left**—From the remote cluster to the local cluster.
	
	The protection status options are:
	• **Active**—The virtual machine is configured for replication and replication occurs per the defined interval. Additional information might be listed.
	  • **Protected**—The virtual machine has a replication schedule.
	  • **Paused**—The replication schedule for the virtual machine is temporarily stopped. This is used during cluster maintenance.
	  • **Invalid**—An error in the virtual machine replication settings.
	  • **In Progress**—A scheduled replication for the virtual machine is proceeding.
	  • **Error**—A replication task for this virtual machine did not complete.
	  • **Deleted**—A replication snapshot was deleted from the remote cluster.
	  • **None**—No replication scheduled for this virtual machine.
	• **Exceeds Interval—**The last replication process took longer than the configured interval to complete.
	• **Halted—**The virtual machine replication schedule is stopped. This prevents a potentially corrupted virtual machine (that is in a state of disaster recovery) from replicating to the remote cluster.
	• **Recovered**—The virtual machine was recently restored from a replication snapshot on the remote cluster.

Last Protection Time column	Timestamp for when the most recent virtual machine replication process started.
Direction column	The direction of the replicated virtual machine. The direction is relative to the local cluster. The cluster you are logged into is always the local cluster. The options are: • **Incoming**—The virtual machine resides on the remote cluster. It is replicated from the remote cluster to the local cluster. • **Outgoing**—The virtual machine resides on the local cluster. It is replicated from the local cluster to the remote cluster.

UI Element	Essential Information
Protection Group column	If the associated virtual machine belongs to a protection group, the protection group name is listed. If it does not have a protection group, the field lists -.
Interval column	The configured interval setting for replicating the virtual machine. To change this, select the virtual machine row and click **Edit Schedule**.

Edit Protected Virtual Machine Schedule Dialog Box

Change the replication interval or VMware Tools quiesce setting for the replication of the selected virtual machine.

Select **Replication** > **Protected Virtual Machines** > **Edit Schedule**.

UI Element	Essential Information
Protect this virtual machine every field	Select how often the virtual machines are to be replicated to the paired cluster. Default is every 1 hour. The pull-down menu options are: 15 minutes, 30 minutes, 1 hour, 90 minutes, 2 hours, 4 hours, 8 hours, 12 hours, 24 hours
Use VMware Tools to quiesce the virtual machine check box	To have HX Data Platform quiesce the virtual machines before taking the replication snapshot, click this checkbox. This only applies to virtual machines with VMware Tools installed.

Click **Save Changes**.

HX Data Platform updates the interval and VMware Tools quiesce setting for the protection group. See the **Protection Groups** tab to view the new interval frequency.

Protection Groups

Create Protection Group Dialog Box

Select **Replication** > **Protection Groups** > **+ New Group**.

Create Protection Group Dialog Box

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.

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	Select this radio button if you want to set a specific time for the first replication operation to start. Before you start replication ensure: • At least one virtual machine is added to the protection group. • The scheduled start time is reached. To specify the protection operation start time: 1. Check the **Start protecting the virtual machines at** radio button. 2. 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 *hours, minutes from now* 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.

Click **Create Protection Group**.

HX Data Platform adds the new group to the **Protection Groups** tab. Notice that the number of VMs is zero, (0). You must add virtual machines to this new protection group to apply the replication schedule set in this protection group.

Edit Protection Group Schedule Dialog Box

Change the replication interval for the virtual machines in the protection group.

Select **Replication** > **Protection Groups** > **Edit Schedule**.

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.

UI Element	Essential Information
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.

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.

Add to Protection Group Dialog Box

UI Element	Essential Information
Add to an existing protection group drop-down list	Select a protection group, click to add virtual machines that are protected to a protection group.

Click **Save Changes**.

Replication Pairs Tab

From the Replication Pairs tab, you can create, edit, or delete replication pairs by selecting datastores on the local and remote clusters, and view the replication pair status. You can also expand the replication pair to view the list of virtual machines protected by this replication pair.

The replication pair defines the two halves of the protection network. The HX Storage Cluster you are logged into is the local cluster, the first half of the pair. When you configure a replication pair, you identify another HX Storage Cluster, the second half of the pair. To ensure the storage component, map the replication pair to datastores on each HX Storage Cluster. When the replication pair is configured, you can then protect virtual machines. See the **Virtual Machines** tab.

Replication Pair Actions

UI Element	Essential Information
Create Replication Pair button	Establish the connection between the local and remote storage clusters. **Prerequisites:** Create datastores on both the local and remote clusters. Configure the replication network on both the local and remote clusters. Click **Create Replication Pair** and complete the wizard.
Edit button	Change the datastores assigned the replication pair name. Select the replication pair and click **Edit**.
Delete button	Remove the replication pair between the local and remote cluster. **Prerequisites:** Remove all the dependencies: Remove protection from all virtual machines. Remove datastore mapping. Select the replication pair and click **Delete**.

Replication Pairs Table

UI Element	Essential Information
Name column	Name of the replication pair for this cluster.
Remote Cluster column	Name of the remote cluster in this replication pair.
Remote Cluster Status column	Displays the current status of the remote cluster. This is different that the general cluster status. Options are: • **Online** • **Offline** • **Upgrading** • **Out of Space** • **Shutdown** • **Unknown**
VMs Outgoing column	Number of virtual machines protected and number of protection groups on the local cluster. Click number to display the outgoing **Local VMs**.
Replications Outgoing column	Number of replication snapshots of the protected virtual machines being replicated from the local cluster to the remote cluster.
VMs Incoming column	Number of virtual machines protected and number of protection groups on the remote cluster. Click number to display the outgoing **Remote VMs**.
Replications Incoming column	Number of snapshots of the protected virtual machines being replicated from the remote cluster to the local cluster.
Mapped Datastore Pairs column	Number of datastores mapped to this replication pair. Click number to display the **Datastores** page.

Replication Pairs Detail Table

Click the replication pair **Name** to view the details table.

UI Element	Essential Information
Virtual Machine Name column	Name of the virtual machine protected by replication in the HX Storage Cluster.

UI Element	Essential Information

Protection Status column	The most recent protection action on the virtual machine protection. The arrow on the status indicates the direction of the data transmission.
	
	The directional arrows indicate data transmission:
	
	• **Left to Right**—From the local cluster to the remote cluster.
	
	• **Right to Left**—From the remote cluster to the local cluster.
	
	The protection status options are:
	
	• **Active**—The virtual machine is configured for replication and replication occurs per the defined interval. Additional information might be listed.
	
	  • **Protected**—The virtual machine has a replication schedule.
	
	  • **Paused**—The replication schedule for the virtual machine is temporarily stopped. This is used during cluster maintenance.
	
	  • **Invalid**—An error in the virtual machine replication settings.
	
	  • **In Progress**—A scheduled replication for the virtual machine is proceeding.
	
	  • **Error**—A replication task for this virtual machine did not complete.
	
	  • **Deleted**—A replication snapshot was deleted from the remote cluster.
	
	  • **None**—No replication scheduled for this virtual machine.
	
	• **Exceeds Interval—**The last replication process took longer than the configured interval to complete.
	
	• **Halted—**The virtual machine replication schedule is stopped. Halting the replication schedule prevents a potentially corrupted virtual machine (that is in a state of disaster recovery) from replicating to the remote cluster.
	
	• **Recovered**—The virtual machine was recently restored from a replication snapshot on the remote cluster.

Last Protection Time column	Timestamp for when the most recent virtual machine replication process started.

UI Element	Essential Information
Direction column	The direction of the replicated virtual machine. The direction is relative to the local cluster. The cluster you are logged into is always the local cluster. The options are: • **Incoming**—The virtual machine resides on the remote cluster. It is replicated from the remote cluster to the local cluster. • **Outgoing**—The virtual machine resides on the local cluster. It is replicated from the local cluster to the remote cluster.
Protection Group column	If the associated virtual machine belongs to a protection group, the protection group name is listed. If it does not have a protection group, the field displays **None**.
Interval column	The configured interval setting for replicating the virtual machine. To change the Interval, select the virtual machine row and click **Edit Schedule**.

Create New Replication Pair Wizard

The replication pair defines the two halves of the protection network. The HX Storage Cluster you are logged into is the local cluster, the first half of the pair. When you configure a replication pair, identify another HX Storage Cluster, the second half of the pair. To ensure the storage component, create the replication pair and map the datastores on the first half to the second half of the pair. After the replication pair is configured and datastores are mapped, you can begin protecting virtual machines. See the **Virtual Machines** tab.

Prerequisites

- Must be using HXDP Release 5.5(1a) or earlier to enable DRO Protection.
- Create a datastore on both the local and remote cluster.
- Configure the replication network.

Start the Replication Pair Wizard

Log in to either the local or remote 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.

Name Page

UI Element	Essential Information
Replication Pair Name field	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.

Click **Next**.

Remote Connection Page

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: *10.10.10.10*.
User Name and **Password** fields	Enter vCenter single sign-on or cluster specific administrator credentials for the remote HX cluster.

Click **Pair**.

HX Data Platform verifies the remote HX Storage Cluster and assigns the replication pair name.



Note Virtual machines to be protected must reside on one of the datastores in the replication pair.

Create New Replication Page > Map Datastores : Native Protection



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, replication schedule fails.

To protect VMs using the HX Data Platform disaster recovery feature, click **Native Protection** and do the following:

UI Element	Essential Information
Local Datastore column	List of the configured datastores on this cluster, the local HX Storage Cluster. Map one local datastore to one remote datastore.
Remote Datastore column	Pair the datastores between the HX Storage Clusters. 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.

Click **Map Datastore**.

Create New Replication Page > Map Datastores : Other DRO Protection

Prerequisites

- Must be using HXDP Release 5.5(1a) or earlier.

To protect VMs using SRM through disaster recovery orchestrator (DRO), click **Other DRO Protection** and do the following:

UI Element	Essential Information
Local Datastore column	List of the configured datastores on this cluster, the local HX Storage Cluster. Map one local datastore to one remote datastore.
Remote Datastore column	Pair the datastores between the HX Storage Clusters. 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.
Direction column	Choose **Incoming** or **Outgoing** as the direction of VM movement for the mapped datastore pairs.
Protection Schedule column	Choose the shedule for protecting all the VMs in the datastore.

Click **Map Datastore**.

Note

- The VMs in the datastores that are under other DRO, are protected by SRM.
- If a new VM is added to the datastore protected by other DRO, the newly added VM is automatically protected by Cisco HyperFlex. If a VM is added to the datastore protected using native DRO, you have to protect the VM.

  The replication pairs that are edited under **Other DRO Protection**, are exposed to SRM.

Edit Replication Pair Dialog Box

Change the Replication Pair Datastores

Change the datastores used for a replication pair on the local and remote clusters. The replication pair name cannot be changed, once created.

Note Changing the datastores used in a replication pair removes protection from all virtual machines on both the local and remote clusters.

This task requires a user with administrator privileges.

1. Unprotect all protected virtual machines, including those virtual machines protected independently or through a protection group. Perform the Unprotect operation on both the local and remote clusters.

   Select **Replication** > **Local VMs** > ***virtual_machine*** > **Unprotect**.

2. Select **Replication** > **Replication Pairs** > ***replication_pair*** > **Edit**.

   To edit the replication pair protected by the HX Data Platform disaster recovery feature, click the **Native Protection** tab and do the following:

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. **Note** The lock/unlock icon next to the datastore name indicates whether the dateastore encryption is enable or disabled: • Locked icon: encryption enabled • Unlocked icon: encryption disabled If encrypted local datastores are selected then only encrypted remote datastore information is displayed.
Remote Datastore column	Pair the datastores between the HX clusters. **a.** 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**. **b.** 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.

To protect VMs using SRM through disaster recovery orchestrator (DRO)[2], click the **Other DRO Protection** tab and do the following:

UI Element	Essential Information
Local Datastore column	List of the configured datastores on this cluster, the local HX cluster. Map one local datastore to one remote datastore. **Note** The lock/unlock icon next to the datastore name indicates whether the dateastore encryption is enable or disabled: • Locked icon: encryption enabled • Unlocked icon: encryption disabled If encrypted local datastores are selected then only encrypted remote datastore information is displayed.

[2] Must be using HXDP Release 5.5(1a) or earlier to enable Other DRO Protection.

UI Element	Essential Information
Remote Datastore column	Pair the datastores between the HX clusters. **a.** 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**. **b.** 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.
Direction column	Choose **Incoming** or **Outgoing** as the direction of VM movement for the mapped datastore pairs.
Protection Schedule column	Choose the shedule for protecting all the VMs in the datastore.

3. Click **Finish**.
4. Re-protect your virtual machines. Select **Virtual Machines** > **virtual_machines** > **Protect**.

Test Pair Cluster Network Dialog Box

UI Element	Essential Information
MTU field	Default is 1500. Enter MTU of the replication network to run the test. • Starting with HXDP Release 5.0(2a) and later, you can edit the MTU value after configuring the cluster. With older versions of HXDP, the existing replication network configuration will need to be removed. The replication network can then be configured with the correct MTU value.

Click **Run Test**, to test cluster pairing between the clusters in the remote replication network.

Delete Replication Pair Dialog Box

Prerequisites to Delete Replication Pair

Remove dependencies from the replication pair. Complete the prerequisites on both the local and remote clusters.

1. Unprotect all protected virtual machines. This includes those virtual machines protected independently or through a protection group. Perform this on both the local and remote clusters.

   Select **Replication** > **Protected Virtual Machines** > ***virtual_machine*** > **Unprotect**.
2. Remove datastore mappings from either the local or remote cluster.
   a. Select **Replication** > **Replication Pairs** > ***replication_pair*** > **Edit**.
   b. From **Remote Datastore** pull-down menu, select **Do not map this datastore**.

c. Click **Finish**.

Delete Replication Pair

Delete a replication pair on the local and remote clusters.

This task requires a user with administrator privileges.

1. Select **Replication** > **Replication Pairs** > ***replication_pair*** > **Delete**.
2. Complete the **Delete Replication Pair** dialog box.

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.

3. Confirm to delete the replication pair, click **Delete**.

Recovery Settings Dialog Box

Edit Recovery Settings



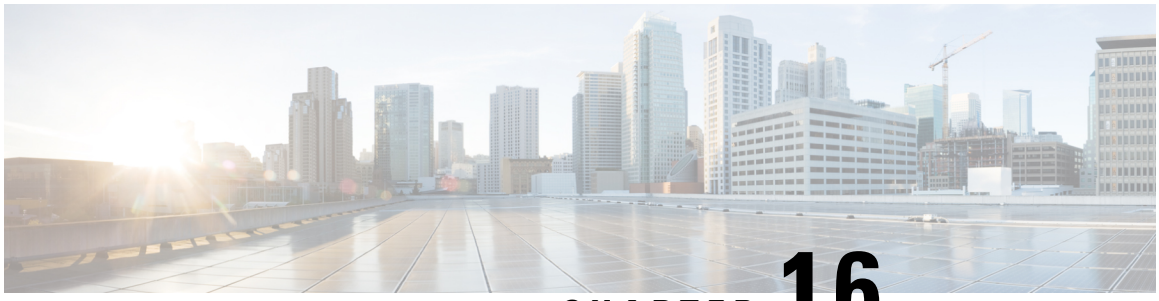
Note To perform this task, you must be logged in as a user with administrator privileges.

1. Complete the fields in the **Edit Network Configuration** Dialog Box.

UI Element	Essential Information
Virtual Machine Power State radio button	Specify the power state for the resource for when the network returns to a known working state.
Test Virtual Machine Name Prefix field	(Optional) Use a common prefix to help identify the type and context of the resource.
Notification Setting radio button	Select the type of notification prompt sent after a recovery event. • 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.

UI Element	Essential Information
Recovery Mappings fields	Define global recovery parameters and mapping for resources across recovery sites by the folder, network, or resource pool parameters to be used during recovery and test recovery operations. Click the parameter type to reveal the configuration fields. Complete the following: **Recovery Configuration** • **Rule**- the number of recovery rules configured. This value cannot be edited. • **Local** drop-down list • **Remote** drop-down list **Test Recovery Configuration** • **Same as Recovery Configuration** check box • **Local** drop-down list • **Remote** drop-down list
Add Rule button	Click to add an additional rule. The default value is 0.
Trash icon	Click the Trash icon to delete a rule.

2. Click **Save Changes**.



CHAPTER 16

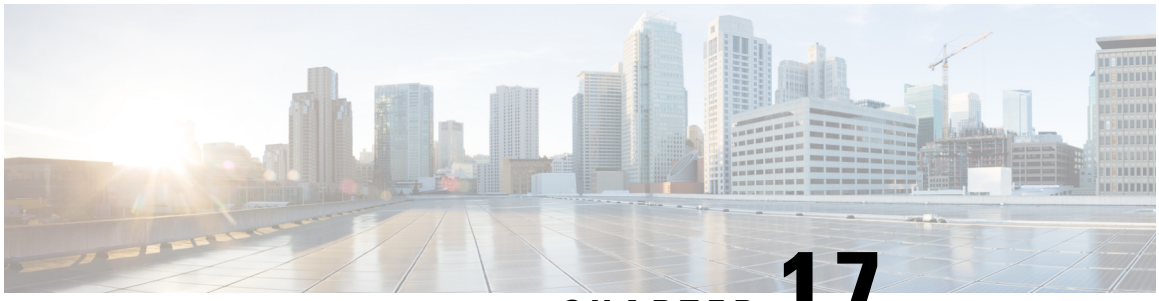
Managing Storage Replication Adapter

- Storage Replication Adapter Overview, on page 279

Storage Replication Adapter Overview

Restriction The Storage Replication Adapter feature is supported through HXDP Release 5.5(1x).

HXDP Release 5.5(1x) users that need SRA admin content should refer to the Managing Storage Replication Adapter chapter in the HyperFlex Admin Guide for HXDP Release 5.0(x).



CHAPTER 17

Managing Users

- Managing Cisco HyperFlex Users Overview, on page 281
- Creating Cisco HX Data Platform RBAC Users, on page 283
- Assigning Users Privileges, on page 284

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.
- ***diag***—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.

HX Interface	admin	root	*hx_admin*	*hx_readonly*	*diag*
HX Data Platform Installer	Required	Not valid	Not valid	Not valid	Not valid
HX Connect	Can perform most HX tasks.	Not valid	Can perform most HX tasks. A preferred user.	Can only view monitoring information. Cannot perform HX tasks. A preferred user.	Not valid

HX Interface	admin	root	*hx_admin*	*hx_readonly*	*diag*
Storage Controller VM with `hxcli` command line	Can perform most HX tasks.	Can perform most HX tasks.	`vc-` prefix required for login. Example: `vc-hx_admin`	Cannot perform HX tasks. `vc-` prefix required for login. Example: `vc-hx_readonly`	Can perform most HX tasks
HX Data Platform Plug-in through vCenter	Can perform most HX tasks.	Not valid	Not valid	Not valid	Not valid
HX REST API	Can perform most HX tasks.	Not valid	Not valid	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.

Procedure

Step 1 Log into 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 19 .

What to do next

Add the user to an RBAC role group. See Assigning Users Privileges, on page 284.

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.

Procedure

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:

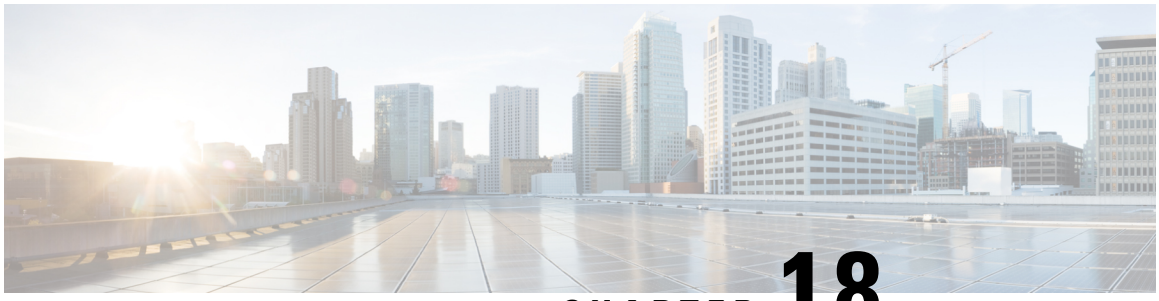
- **Administrator**
- **Read only**

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 18

Managing iSCSI

Note The iSCSI features are supported in Cisco HyperFlex Release 4.5(x) and later.

- HyperFlex iSCSI Target Service Overview and Supported Use Cases, on page 285
- HyperFlex iSCSI Best Practices , on page 286
- iSCSI Configuration Overview, on page 286
- iSCSI Network Page, on page 287
- iSCSI Initiator Group, on page 289
- iSCSI Target Page, on page 292
- iSCSI LUN Page, on page 296
- Configuring an iSCSI Initiator (Windows), on page 298
- Configuring an iSCSI Initiator (Linux), on page 299
- Cloning an iSCSI LUN, on page 299

HyperFlex iSCSI Target Service Overview and Supported Use Cases

The HyperFlex iSCSI Target Service is introduced 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.
- 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
- Support for HX iSCSI support to Edge and DC-No-FI was introduced in Cisco HX Release 5.0(2a).

Initiators are currently supported on Windows Server 2016, Windows Server 2019, Ubuntu 18.04 and 20.04, Oracle Linux 8.2, Oracle Linux 7, Red Hat Enterprise Linux 8.2, Red Hat Enterprise Linux 7.

Note The HyperFlex iSCSI Target Service is not supported on Stretched clusters.

HyperFlex iSCSI Best Practices

When enabling iSCSI in HyperFlex, it is a best practice to also enable Boost Mode. Boost Mode allows the Cisco HyperFlex Cluster to deliver higher IOPs by increasing the storage controller VM CPU resources by 4 vCPU, and mitigate the performance impact of iSCSI. For more information about enabling or configuring Boost Mode, see Boost Mode, on page 63.

iSCSI Configuration Overview

The process to configure the iSCSI Target Service is as follows:

- Create the ISCSI Network in HX Connect.
- Create an Initiator Group.
- Create an ISCSI Target.
- Link the Initiator Group with Targets.
- Create an iSCSI LUN.
- Configure the iSCSI Initiator.
- Clone the LUN.

iSCSI Scale and Support

The following describes iSCSI support recommendations. The noted values are based on what Cisco has tested, and what provides optimal performance. Using these as your "Maximum supported" guideline is strongly encouraged.

Table 13: iSCSI Scale and Support Recommendations

Scale Item	HXDP Support
iSCSI LUNs per HyperFlex cluster	32,640
iSCSI Initiator Groups per HyperFlex cluster	128
iSCSI targets per HyperFlex cluster	128
iSCSI LUNs per target	255
Maximum iSCSI LUN size	64 TB
Maximum number of iSCSI sessions per controller VM	64

Scale Item	HXDP Support
iSCSI IO initiator queue depth per iSCSI session	256
iSCSI IO target side queue depth per controller	2,048

iSCSI Network Page

Displays configuration information for your iSCSI network.

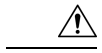
iSCSI Network Data

If you have already configured your iSCSI network, basic network information is displayed.

Table 14: 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



Caution An ip address range is required for the new VLAN configuration. This range should not already exist in the cluster. Not adhering to this requirement can result in a cluster outage.

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 cannot change some iSCSI network parameters without TAC assistance.

Note If you configured the hx-storage-data network with 1500 MTU (no Jumbo frames), but you want to utilize Jumbo frames (as recommended for the iSCSI network), you will need to manually edit the hx-storage-data network vswitch on all ESXi hosts in the HyperFlex cluster to 9000 MTU.

iSCSI Network Configuration Data

The following information is required to create your iSCSI Network.

Table 15: 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 (this range should not include the iSCSI Storage IP)
iSCSI Storage IP	Enter a valid IP Address for iSCSI Storage (this IP address should not be an address used in the IP Range field)
Set non default MTU	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. To disable Jumbo Frames, click the **Set non default MTU** checkbox, and then use the pull-down to change the value to 1500. **Note** If any of the initiators are crossing a router, the router will need to allow Jumbo Frames.

UI Element	Description
VLAN Configuration	Click in the checkbox to **Create a new VLAN**(Recommended) or **Select an existing VLAN**. 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). To Select an existing VLAN, you need to specify the VLAN ID. **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.

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

You may delete and reconfigure an iSCSI network configuration while preserving the iSCSI LUNs, initiator groups and targets on the cluster.

To delete an iSCSI network configuration:

1. Go to the **Actions** menu, and select **Delete Network**. The Delete Network Configuration window appears prompting "Are you sure you want to delete Network Configuration?".
2. Click **Confirm**; or click **Cancel** to cancel your changes.



Note The option to delete an iSCSI network and retain objects you have configured is also 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 16: 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 `iqn.yyyy-mm.naming-authority:unique name`.

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 iSCSI Network using HX Connect.

Procedure

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:

a) For Windows, log into 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.

b) 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 for your release.

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:

Procedure

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.

Procedure

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 created an Initiator Group and iSCSI Targets.

Procedure

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;

- Create an iSCSI LUN
- Configure an iSCSI Initiator
- Clone LUNs

Unlinking an iSCSI Initiator Group

To unlink an iSCSI Initiator Group:

Procedure

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 17: Target Creation Data

UI Element	Description
Name	Name for the target.
Enable CHAP authentication	Click to enable CHAP authentication. **Note** HXDP supports one-way CHAP authentication.
Username	Username for CHAP authentication
Secret	Secret for CHAP authentication

After a Target is created, the following information appears.

Table 18: Target Information Data

UI Element	Description
IQN	Target iSCSI Qualified Name for your Initiator.
Active Initiators Count	Total number of active Initiators.
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 iSCSI Network using HX Connect. It is also recommended that you create Initiator Groups.

Procedure

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.

HXDP supports one-way CHAP authentication.

Note
When using CHAP and booting from SAN (iSCSI LUN), you must set the CHAP user/password at the top initiator area in UCS (for example, in UCS manager, navigate to Boot Policy > Set iSCSI Boot Parameters > Authentication Profile).

Step 4 Click **Create Target**.

The Target appears in the Targets tab.

What to do next

After you have created an iSCSI Target:

- Link the Initiator Group with Targets
- Create an iSCSI LUN
- Configure an iSCSI Initiator

Editing an iSCSI Target

To edit an iSCSI Target:

Procedure

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:

Procedure

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 cannot delete a Target if there is an Initiator Group linked to the target.

Note
You cannot 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 Initiator Groups. 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 290.

Procedure

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.

- Create an iSCSI LUN
- Configure an iSCSI Initiator
- Clone LUNs

Unlinking an iSCSI Target

To unlink an iSCSI Target:

Procedure

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 19: iSCSI LUN Data

UI Element	Description
Name	Name of the LUN.
Size	Total capacity size of the LUN (in GB). **Note** The maximum LUN size is 64TB.

After a LUN is created, the following information appears.

Table 20: 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).
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:

Procedure

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.

Note
The maximum LUN size is 64 TB.

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.



Note You can create volumes that are protected by CHAP. The limit that you can create with one storage class for each target is 255 volumes (Persistent Volume Claims).

What to do next

After you have created an ISCSI LUN, you can Configure an iSCSI Initiator

.

Editing an iSCSI LUN

To edit an iSCSI LUN:

Procedure

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.

Note
The maximum LUN size is 64 TB.

Step 5 Click **Edit LUN**; or click **Cancel** to cancel your changes.

Deleting an iSCSI LUN

To delete an iSCSI LUN:

Procedure

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.

Note HXDP supports one-way CHAP authentication. Two-way CHAP authentication is not supported.

ESXi host cannot be set as an initiator for HX iSCSI.

Procedure

Step 1 Log into 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.

Note ESXi host cannot be set as an initiator for HX iSCSI.

Procedure

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 Log into 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 301.

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 301.

Procedure

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.

Procedure

Step 1 Log into the Windows machine using Administrator or AD credentials.

Step 2 Run the Windows HX Agent install executable by double-clicking on the file: `HxWindowsAgentIscsiClone-v4.5.1a-39020.exe`.

Note
Agent logs and files extracted from `HxWindowsAgentIscsiClone-v4.5.1a-39020.exe` appear with build number `4.5.1a.38547` in properties of the file. This is a version display issue with no impact on functionality and can be disregarded.

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 installatoin 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.

Procedure

Step 1 Log into 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.

Note
Agent logs and files extracted from `HxWindowsAgentIscsiClone-v4.5.1a-39020.exe` appear with build number `4.5.1a.38547` in properties of the file. This is a version display issue with no impact on functionality and can be disregarded.

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.

Procedure

Step 1 Log into the Windows machine using Administrator or AD credentials.

Step 2 Double click on the file: `HxWindowsAgentIscsiClone-v4.5.1a-39020.exe`.

Note
Agent logs and files extracted from `HxWindowsAgentIscsiClone-v4.5.1a-39020.exe` appear with build number `4.5.1a.39020` in properties of the file. This is a version display issue with no impact on functionality and can be disregarded.

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 later), Microsoft Visual C++ 2017 Redistributable (x64) (version: 14.10.25017 or later), Microsoft Visual C++ 2017 Redistributable (x86) (version: 14.10.25017 or later) 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 21: 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.

Procedure

Step 1 Log into 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.

Procedure

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`.

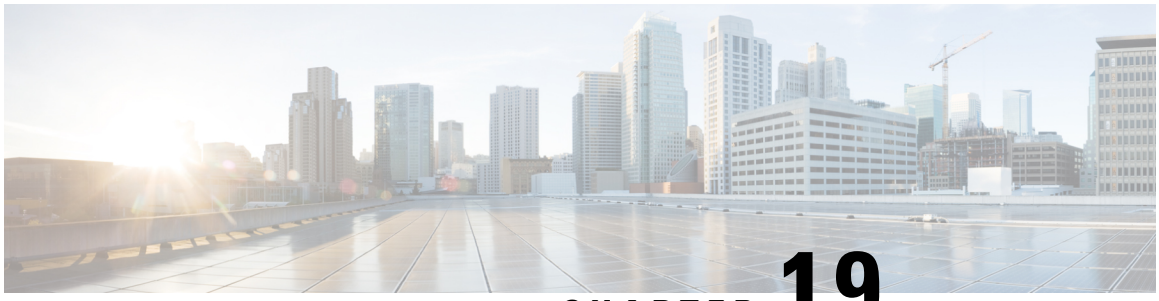
Step 6 If the disk attribute "Read-only" is "Yes", set as "No": "attributes disk clear readonly"

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 19

Cisco HyperFlex HTML Plugin for VMware vCenter

- Cisco HyperFlex vCenter Plugin, on page 307
- Cisco HyperFlex HTML5 Plugin for VMware vCenter , on page 308
- vCenter: HyperFlex Plugin Embedded Actions, on page 347
- Cisco HyperFlex Remote Plugin for VMware vCenter, on page 360
- Install, Register and Upgrade the Remote Plugin, on page 362
- Encryption Support, on page 365
- Generate Support Bundles, on page 366

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. The vCenter plugin is designed to ease the HyperFlex administration tasks without a need to visit to each HyperFlex clusters. This chapter describes the usage of Local vCenter Plugin and Remote vCenter Plugin.

Note VMware local plugin architecture support is limited to vSphere versions 6.5, 6.7, and 7.0. For more information, see the *vSphere Client Local plugins are deprecated (87880)* article on the VMware site.

To get started, use the following table to determine the plugin type for your deployment. Click the link in the Plugin Type column to advance to the full documentation.

vCenter Plugin Version	Cisco HyperFlex Release	vCenter/ESXi version	Plugin Type
3.0.0 and later	4.0(2f) and later	6.7 U3 and later	Cisco HyperFlex Remote Plugin for VMware vCenter, on page 360[3]
2.0.0, 2.1.0 and 2.2.0	4.0(2e) and earlier	Up to 7.0 U3	Local Plugin

[3] In alignment with the VMware deprecation of local plugins starting with vSphere 8.0, remote plugin is the only supported version for HX 4.0(2f) and later

Cisco HyperFlex HTML5 Plugin for VMware vCenter

The Cisco HyperFlex Local 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.

This section describes how to monitor and manage your HyperFlex clusters from the VMware vCenter using the Cisco HyperFlex HTML5 plugin versions 2.0.0, 2.1.0 and 2.2.0.

Cisco HyperFlex HTML5 Plugin Prerequisites

The following hardware and software prerequisites apply to the Cisco HyperFlex HTML5 Plugin:

- **Browser compatibility:** The Cisco HyperFlex HTML plugin works with Chrome, Firefox and IE.
- Administrative Privileges are required for managing users and roles.
- The installation workflow is the same for single and linked mode vCenter instances.
- Beginning with HX Release 5.0(1a) full HTML5 plugin feature functionality requires the license status to be In-compliance.
- HXDP Release 5.0(x) and later releases do not support Cisco HyperFlex Flash Plugin (the original plugin).
- Cisco HyperFlex HTML5 plugin for VMware vCenter support was introduced in Cisco HX Release 4.0(2a) and vCenter 6.5U2.
- Cisco HyperFlex HTML5 plug-in 2.2.0 is the minimum version supported. If the running version is 2.1.0 or 1.0.1, upgrade to the latest version.
- HTML Plugin v2.2 supports vCenter Linked Mode.

Install and Register the vCenter HTML5 Plugin

Install the Cisco HyperFlex HTML5 plugin with the VMware vSphere web client. During the plugin installation process, enter the HX Storage controller VM admin password and vCenter Username & Password information:

Table 22: 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.
- For deployments using Cisco HX Release 5.0 and later, review the Secure Admin Shell feature.
- HTML-Plugin v2.2 supports vCenter Linked Mode.
- The installation workflow is the same for single and linked mode vCenter instances.

Procedure

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-2.2.0.zip` file into a temporary directory in one of the controller VMs and unzip.

a) The file transfer may be completed by using sftp cli or any file transfer app such as winscp or filezilla.

To use sftp transfer via a file transfer app copy the file to the /tmp folder on SCVM, using HX admin account.

b) SSH to that SCVM and login with admin account.

c) Change to the `/tmp` directory using the command `"cd /tmp"`.

Example:

```
"cd /tmp"
```

d) Unzip the plugin file `HyperFlex-VC-HTML-Plugin-2.2.0.zip` using the command unzip.

Example:

```
unzip HyperFlex-VC-HTML-Plugin-2.2.0.zip
```

Step 3 Execute the command `install_vc_plugin` on your shell and enter:

- vCenter FQDN/IP address
- Administrator username and password of vCenter server
- Storage Controller VM admin password

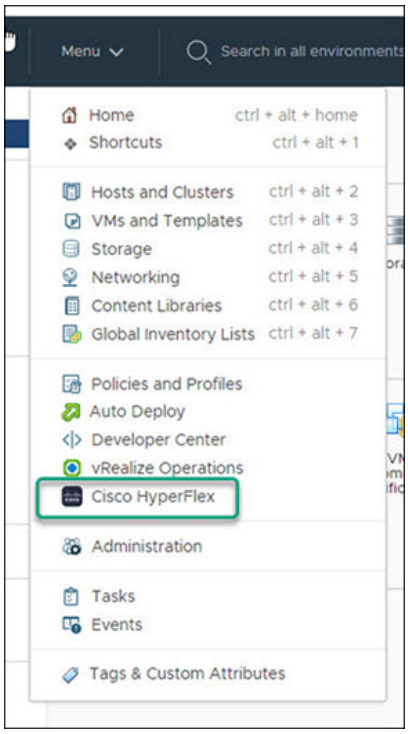
Note
Beginning with Cisco HX 4.5(x) with secure shell, the default Storage Controller VM root password is the same as the Storage Controller VM admin password.

If prompted for the Storage Controller VM root password, the default root password is the first password assigned to the Controller VM during setup.

- Storage Controller VM admin password

Step 4 Log on to vCenter and a blue banner message appears to confirm that the new plugin is installed.

Step 5 Log out and log in again to vCenter to see the Cisco HyperFlex menus for HTML5 plugin.



Verifying the Cisco HyperFlex HTML5 Plugin Installation from the vSphere Client

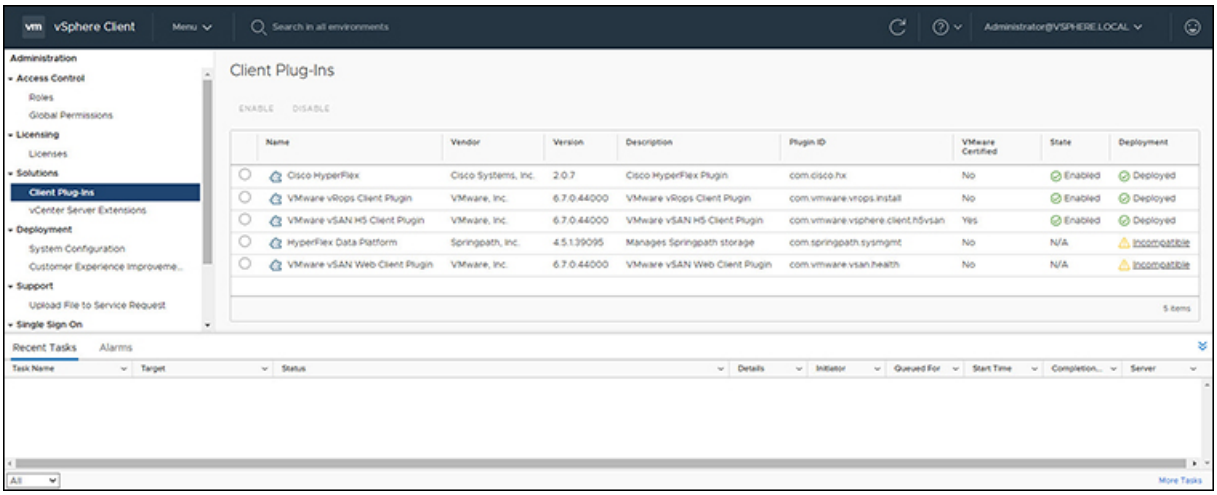
To verify Cisco HyperFlex plugin installation from the vSphere Client UI.

Before you begin

The HTML5 plugin should be installed on vCenter server.

Procedure

Launch the vSphere client, select **Menu** > **Administration** > **Solutions** > **Client Plug-Ins**



Uninstalling the HyperFlex HTML5 Plugin

To uninstall the HX Data Platform HTML5 Plugin, perform the following steps.

Procedure

Step 1 Execute the uninstall command `install_vc_plugin -u` 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 2.2.x.

Procedure

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-2.2.x.zip` file into a temporary directory in one of the controller VMs and unzip.

a) The file transfer may be completed by using sftp cli or any file transfer app such as winscp or filezilla.

   To use sftp transfer via a file transfer app copy the file to the /tmp folder on SCVM, using HX admin account.

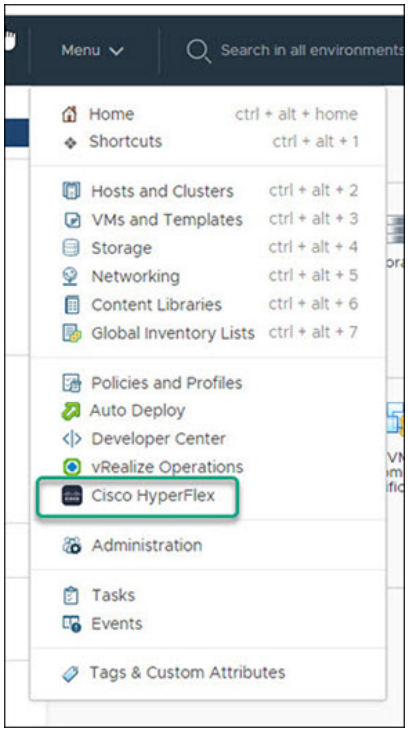
b) SSH to that SCVM and login with admin account.
c) Change to the /tmp directory "cd /tmp"
d) Unzip the plugin file `HyperFlex-VC-HTML-Plugin-2.2.x.zip`

Step 3 Execute `install_vc_plugin` command on your shell and enter:

- vCenter FQDN/IP address
- Administrator username and password of vCenter server

Step 4 Select **Y** to continue the Upgrade process with controller root and admin password.

Step 5 Logout and log in again into vCenter to see Cisco HyperFlex listed in the vCenter menus.



Using the Cisco HyperFlex HTML5 Plugin

The following table defines feature support by plugin version:

Table 23: HTML5 Local Plugin Feature Support

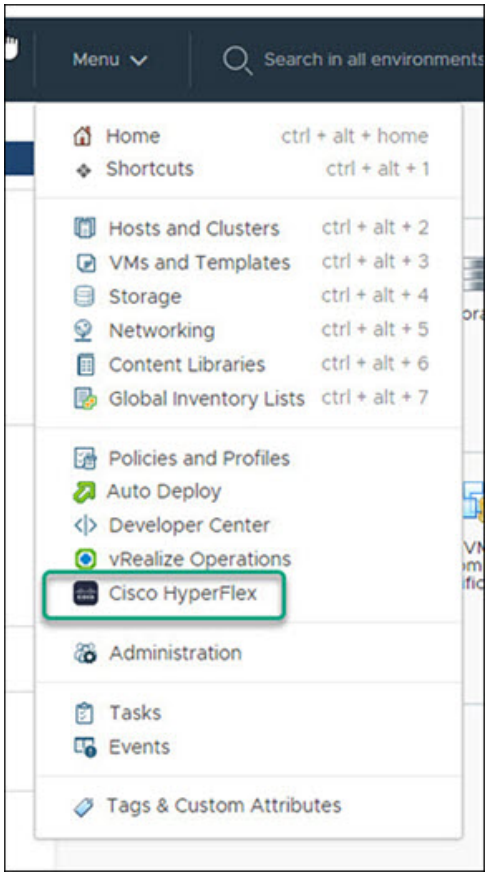
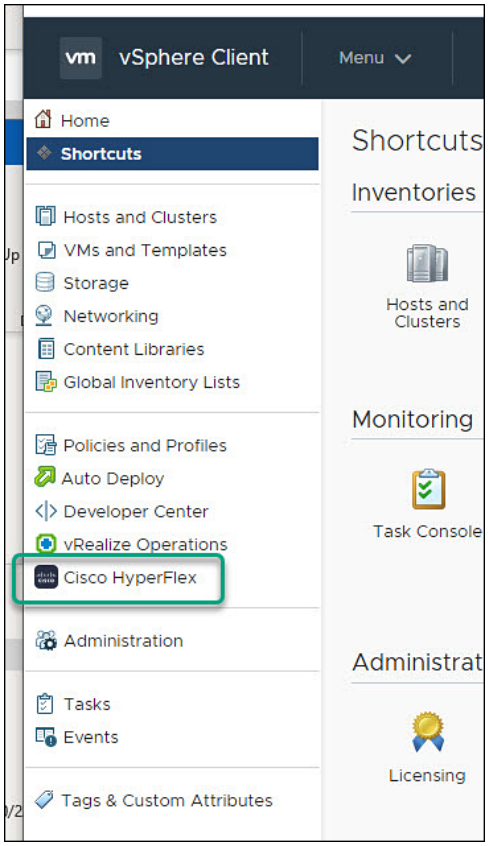
Feature	Plugin Version 2.0.0	Plugin Version 2.1.0	Plugin Version 2.2.0
Discover the Registered HX Cluster	✓	✓	✓

Rename Clusters[4]	-	✓	✓
View HX Cluster Summary	✓	✓	✓
View Cluster and Datastore Performance Charts	✓	✓	✓
Disks View	✓	✓	✓
Nodes View	✓	✓	✓
HX Datastore Management	✓	✓	✓
VM Summary and Top VM Consumers	✓	✓	✓
Network Management	-	✓	✓
iSCSI Management[5]	-	✓	✓
Events and Alarms	✓	✓	✓
Manage Tasks	-	✓	✓
HX Snapshots and clones at the virtual machine level	-	✓	✓
Schedule Snapshot[6]	-	✓	✓
Manage users and access to HX clusters	✓	✓	✓
Cross-launch HX Connect for upgrade	✓	✓	✓
Embedded vCenter server actions at the Host and Clusters level	✓	✓	✓
HTML 5 License Status[7]	-	-	✓
Linked Mode	-	-	✓

[4] Requires HXDP Release 4.5(x) or later.
[5] Requires HXDP Release 4.5(x) or later.
[6] Requires HXDP Release 4.5(x) or later.
[7] Requires HXDP Release 5.0(x) or later.

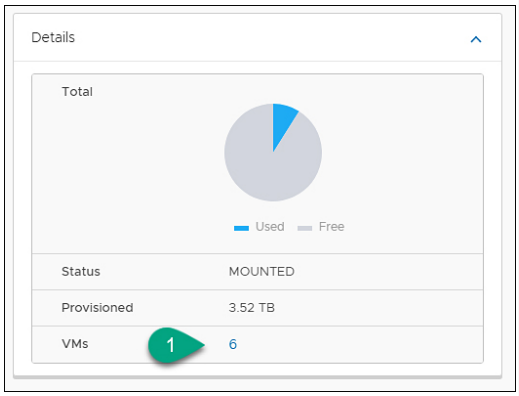
Navigating the HTML5 plugin

Accessing the Cisco HyperFlex HTML5 Plugin is easily accessed from the vSphere Menu or the Shortcuts list.

vSphere Menu	vSphere Shortcut 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. **Note** The cluster list uses dynamic function loading, the Scanning Clusters icon indicates when the cluster list is complete.
	Indicates that the cluster table is still populating. The icon disappears when the cluster list is complete.
	Filter the content seen in the browser.

Icon	Usage
	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:

Procedure

Step 1 Log into 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. The **Scanning Clusters** icon indicates that the cluster table is still populating. The icon disappears when the cluster list is complete.

Rename Cluster

The rename cluster was introduced in HX Release 4.5. To rename a cluster, perform the following steps:



Procedure

Step 1 Log into the vSphere web client.

Step 2 Select **Menu** > **Cisco HyperFlex**
The HyperFlex Clusters List appears.

Step 3 Click on the row of the cluster that you want to rename.
The **Rename** button appears for supported clusters.

Note
The rename cluster feature is supported on HXDP Release 4.5 and later.

Step 4 Click the **Rename** button.
The Rename Cluster window appears.

Step 5 Type the new name on the **Cluster Name:** line.

Step 6 Click **OK** to confirm the name change.

View the HX Cluster Summary

To view a summary of the HX Clusters in your deployment perform the following steps:

Procedure

Step 1 Log into 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 Release, Model, vCenter Cluster, ESXi Version and Uptime.

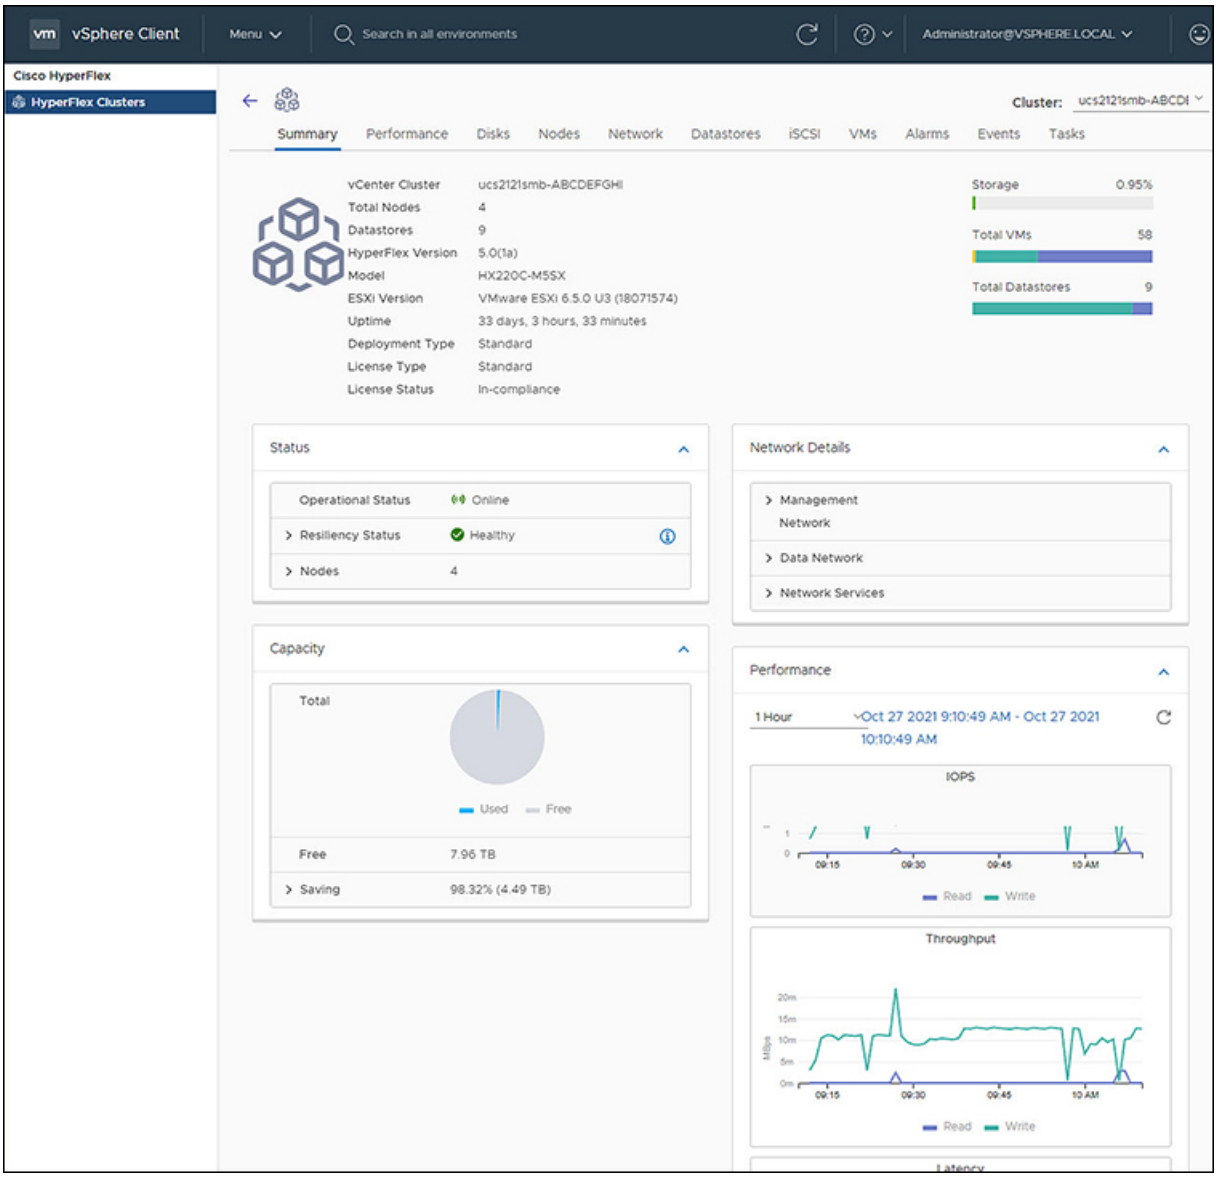


Table 24: Cluster Summary View Details

Field Name	Additional Information
vCenter Cluster	Name of the vCenter cluster
Total Nodes	Total number of nodes
Datastore	Datastore connected to the cluster
HyperFlex Release	Version of HyperFlex on the cluster
Model	Model name
ESXi Version	ESXi Version

Field Name	Additional Information
Uptime	Length of time that the cluster has been up and running
Cluster Type	Type of cluster
Deployment Type	Type of cluster deployment. Valid options are Standard and Edge.
License Type[8]	Type of License. Valid options include: Evaluation, Standard, and Enterprise. **Note** New users enjoy a 90-day grace period to register the license. After 90-days a "License is not in compliance" appears and product features are limited.
License Status[9]	License status. Status include: In-compliance, Out of Compliance, and License expires in *x* days, Cluster not registered with Cisco Licensing. Cluster not registered with Cisco Licensing.
Storage Capacity Bar	A graphical representation of the percentage of total storage used. Hover over the bar to view the amount of storage used.
Total VMs Bar	A graphical representation of the total number of VMs in the cluster.
Total Datastore Bar	Total number of datastores connected to the cluster. Hover over the bar to view the number of datastores mounted and unmounted.

[8] Added in HX Release 5.0(x)
[9] Added in HX Release 5.0(x)

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 25: Status Portlet

Field Name	Additional Information
Operational Status	Online or Offline
Resiliency Status	Warning or Healthy Click the arrow to collapse or expand additional Resiliency Status details: • **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

Field Name	Additional Information
Nodes	Number of nodes in the cluster. Click the arrow to collapse or expand additional Nodes details: • Node Type • Version

Table 26: Capacity Portlet

Field Name	Additional Information
Total	Used and Free Capacity expressed as a percentage
Total Capacity	Amount of usable capacity
Used	Used Capacity
Free	Free Capacity
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 27: Network Details Portlet

Field Name	Additional Information
Management Network	Management Network details Click the arrow to view the following Management Network details: • Management IP address / FQDN • VLAN • Default Gateway
Data Network	Data network details Click the arrow to view the following Data Network details: • Data IP address / FQDN • VLAN • Default Gateway
Network Services	Network Services details Click the arrow to view the following Network Services details: • DNS Server(s) • NTP Server(s)

Table 28: Performance Portlet

Field Name	Additional Information
General Usage • Performance charts are visible when the License Status is **In-compliance**[10]. • 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. • The Scanning Cluster icon indicates that the cluster table is still populating. The icon disappears when the cluster list is complete. • 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

[10] Supported in HX Release 5.0(1a) and later.

Register your license

New users have a 90-day grace period to register their license. During the 90-days you have full access to all feature functionality. To continue using the complete set of features, perform the following steps to register your license using the in-product link.

Before you begin

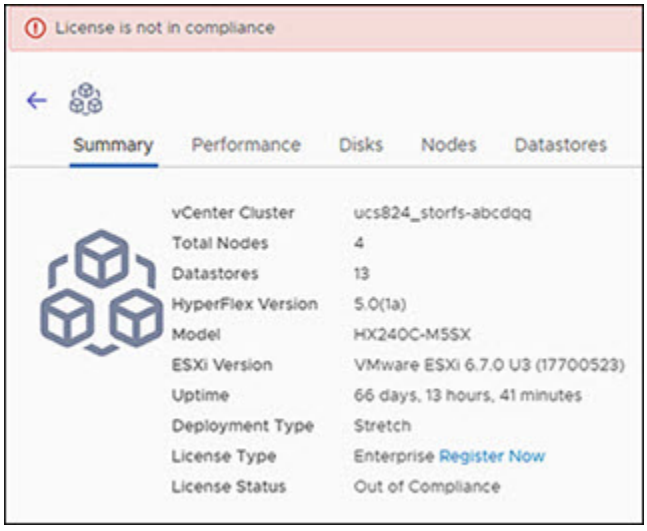
Beginning with HX Release 5.0(1a) full HTML plugin feature functionallity requires the license status to be In-compliance. Verify the your License type and Status on the Summary Page, if you need to register your license, complete this task.

License Compliance Examples

Figure 1: License In-Compliance:



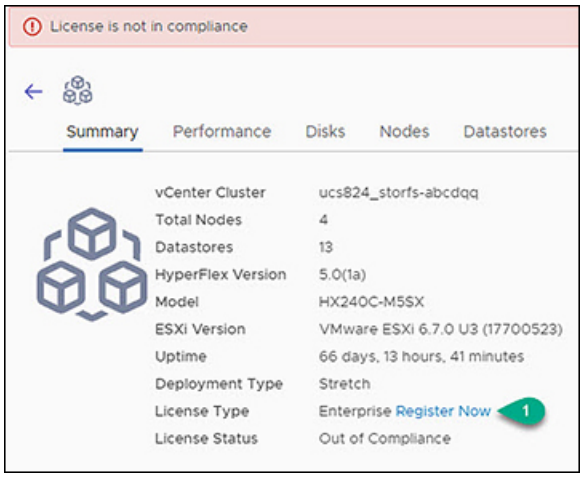
Figure 2: License Out of Compliance:



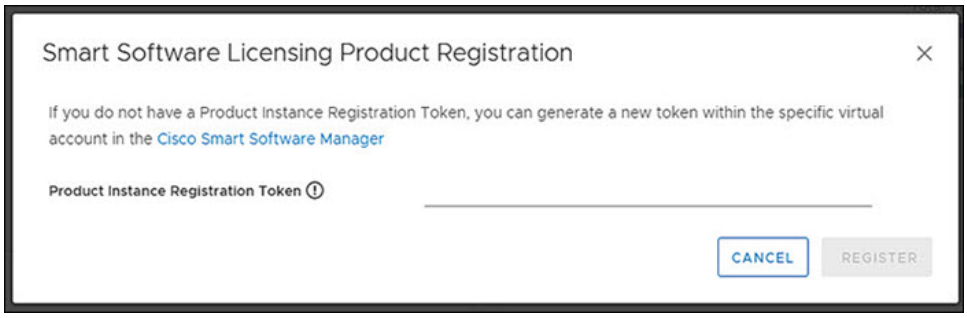
Procedure

Step 1 Starting on the vSphere web client **Summary** page, click the discovered HX cluster name to view its summary.

Step 2 In the License Type summary, click the **Register Now** link. The "Smart Software Licensing Product Registration window appears.



Step 3 Type the Product Instance Registration Token on the field provided

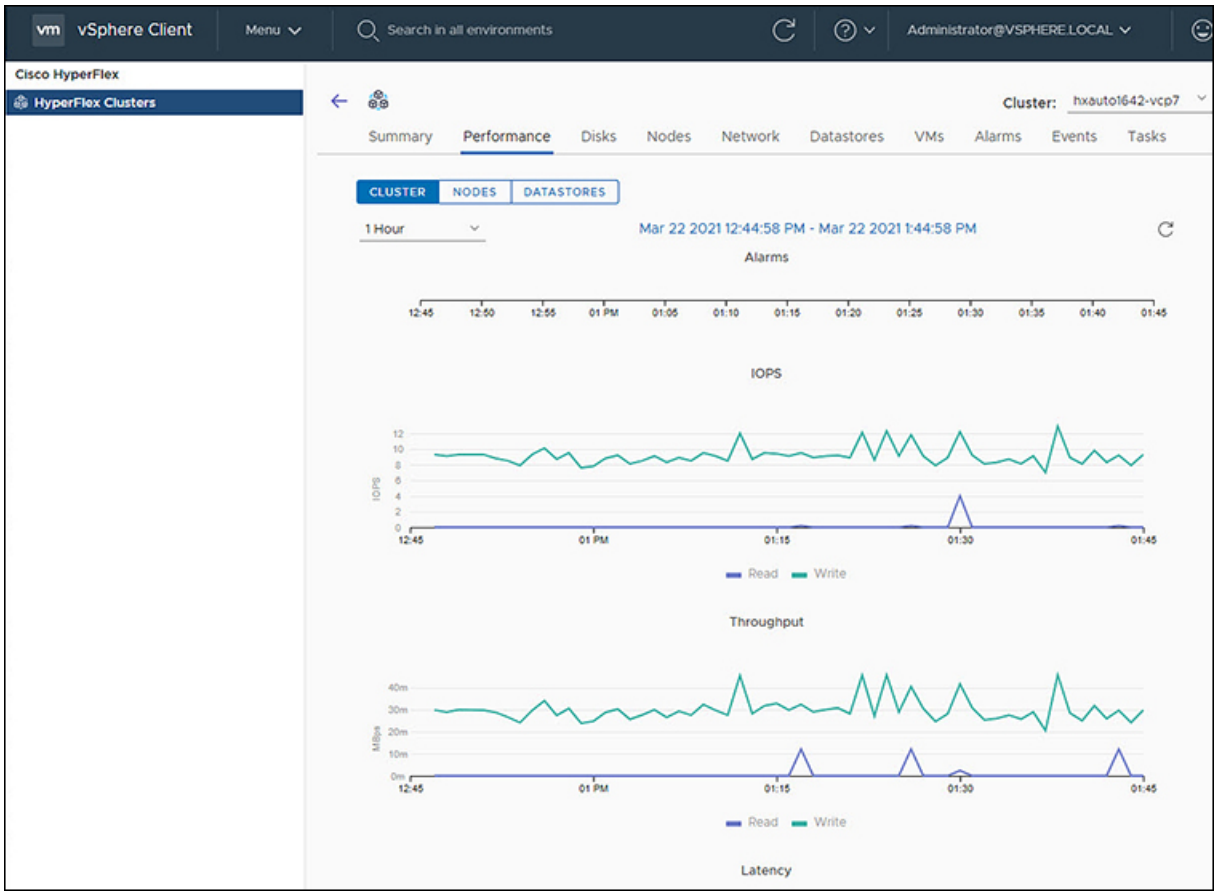


Note
If your registration token is not available, generate a new one by clicking on the **Cisco Smart Software Manager** Link and follow the prompts.

Step 4 Click **Register** to complete the action.

View Cluster and Datastore Performance Charts

The **Performance** tab displays performance details for both the cluster and for the datastore for the last hour.



General Usage:

- Click on the Time Interval list to select the length of time viewed in the performance chart.

Note The Alarms chart appears with time interval selections of 1 Month or less.

- Use the drop-down Cluster list on the top right to navigate between the clusters.
- 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.

Before you begin

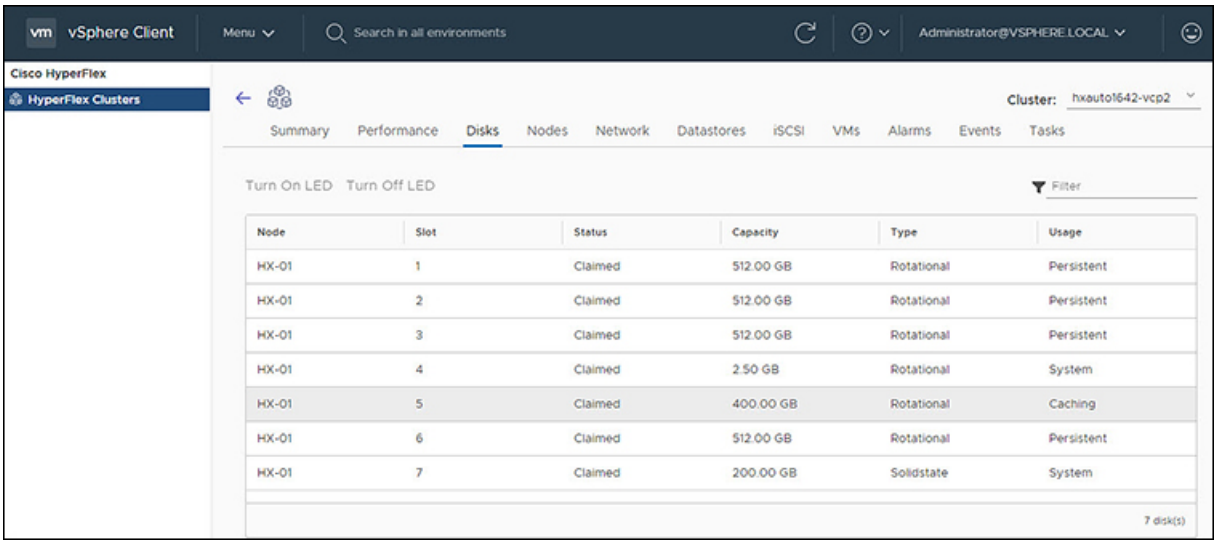
Beginning with HX Release 5.0(1a), Performance charts are only visible when the license status is In-compliance.

Procedure

Step 1 Log into the vSphere web client.
Step 2 Select **Menu** > **Cisco HyperFlex**.
Step 3 Click on HX Cluster you want to review.
Step 4 Select the **Performance** tab. The Alarms, IOPS, ThroughPut, and Latency charts appear.
Step 5 Click on the Time Interval list to select the time-span viewed in the performance chart.

Disks

To view the Disks details page, perform the following steps:



Procedure

Step 1 Log into 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 **Disks**.
The Disk Detail view appears.

Table 29: Disks Details

Field Name	Additional Information
Node	Node Name
Slot	Slot number

Field Name	Additional Information
Status	Slot Status. Valid values: Available or Claimed
Capacity	Total capacity of the slot
Type	Type of disk. Valid values include: Rotational, Solid state
Usage	How the disk is being used. Valid values include: Caching, Persistence, System

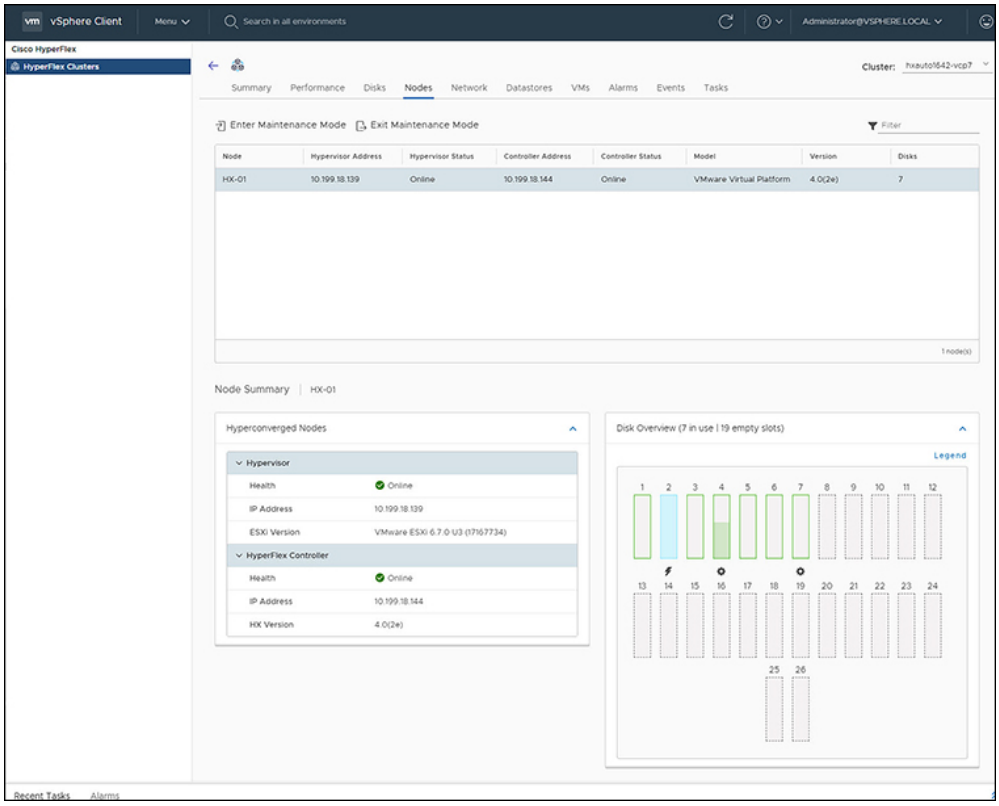
Step 5 (Optional) **Locate a physical server using the Turn On LED button**

Note
Beginning with HX Release 5.0(1a), The Turn On/Off LED button functionality requires the license status to be In-compliance.

a) Click the **Turn On LED** button to illuminate the LED light on the associated physical server.
b) When finished, click the **Turn Off LED** button to turn the LED light off.

Nodes

To view node details specific to the Cluster, Host, and VMs, perform the following steps:



Procedure

Step 1 Log into 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 **Nodes**.
The Nodes list appears.

Table 30: Nodes List Details

Field Name	Additional Information
Node	Node name.
Hypervisor Address	IP address of the Hypervisor.
Hypervisor Status	Hypervisor status. Valid values: Online and Offline.
Controller Address	IP address of the Controller.
Controller Status	Controller status. Valid values: Online and Offline.
Model	Type of node.
Version	HXDP version in use.
Disks	Number of disks associated with the node.
Site	Column is displayed only for Edge deployments.

Step 5 Click the Node name that you want to view details. The Node Summary portlet appear below the Nodes list.

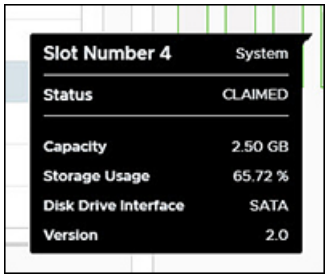
a) The Node Summary view includes two portlets with additional details about the node: Hyperconverged Nodes and Disk Overview.

Use the arrows to collapse and expand the portlet contents.

Table 31: Hyperconverged Nodes Portlet

Field Name	Additional Information
Hypervisor	**Health** - Online or Offline **IP Address** - IP address of the Hypervisor **ESXi Version** - Installed ESXi Version
HyperFlex Controller	**Health** - Online or Offline **IP Address** - IP address of the HyperFlex Controller **HX Version** - Installed HyperFlex release

Table 32: Disk Overview Portlet

Field Name	Additional Information
Disk Overview	Notes the number of slots in use and the number that are empty.
Legend	Legend for icons and colors used in the disk graphics.
Disk graphic	Hover over a disk to display details for that disk.  Details include: • Slot Number and type of usage • Disk Status: Claimed or Unclaimed • Capacity • Storage Usage as a percentage. • Disk Drive Interface • Version

Step 6 (Optional) **Enter or Exit Maintenance Mode**

a) Click the Node name that you want to put into or take out of Maintenance Mode.

b) Click **Enter Maintenance Mode** or **Exit Maintenance Mode**.

Note
Beginning with HX Release 5.0(x), the Enter and Exit Maintenance Mode button functionality is enabled when the license status is In-compliance.

Note
If you have a 3- or 4-node cluster, only one node will go into maintenance mode.

Network

Network: Create New VLAN

The Network page allows users to create a VLAN without going through UCS. To create a VLAN from the vSphere client perform the following steps:

Procedure

Step 1 Log into the vSphere client.

Step 2 Select **Menu** > **Cisco HyperFlex** > **Create VLAN**

Step 3 The Create VLAN window appears.
Complete the fields in the Create VLAN window:

Table 33: Create VLAN

Field Name	Additional Information
VLAN ID	To create one VLAN, enter a single numeric ID. A VLAN ID can: • Be between 1 and 3967 • Be between 4049 and 4093
VLAN Name	This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters, and you cannot change this name after the object has been saved
UCS Manager Host IP or FQDN	UCS Manager FQDN or IP address. For example, 10.193.211.120
UCS Username	<admin> username. For example, sample_user1
UCS Password	<root> password.

Step 4 Click **OK**.
A VLAN is created.

Note
Creating VLANs is a one-way operation. You cannot view VLANs in the HTML plugin. To see the newly created VLAN go to UCS, and then verify the newly created vLAN in the ESXi vSwitch.

Network: Configure iSCSI Network

The Network page allows users to configure an iSCSI Network. To create an iSCSI network from the vSphere client perform the following steps:

Before you begin

The iSCSI features are supported in Cisco HyperFlex Release 4.5(x) and later.

Procedure

Step 1 Log into the vSphere client.

Step 2 Select **Menu** > **Cisco HyperFlex** > **Network Configure** > **Configure**

Step 3 The Create iSCSI Network window appears.
Complete the fields in the Create iSCSI Network window:

Field Name	Additional Information
Subnet	Enter a valid Subnet
Gateway	Enter a valid Gateway
IP Range	Enter a valid IP Range Use the **Edit** button to change the IP range. All other fields are disabled.
iSCSI Storage IP	Enter a valid IP Address for iSCSI Storage
VLAN Configuration	Click in the checkbox to **Create a new VLAN (Recommended)** or **Select an existing VLAN**. 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). To Select an existing VLAN, you need to specify the VLAN ID. **Note** To configure the VLAN manually in UCS-M, use the create VLAN 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.
Set Non-Default MTU	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. To disable Jumbo Frames, click the **Set non default MTU** checkbox, and then use the pull-down to change the value to 1500. **Note** If any of the initiators are crossing a router, the router will need to allow Jumbo Frames.

Step 4 Click **OK**.
An iSCSI network is created.

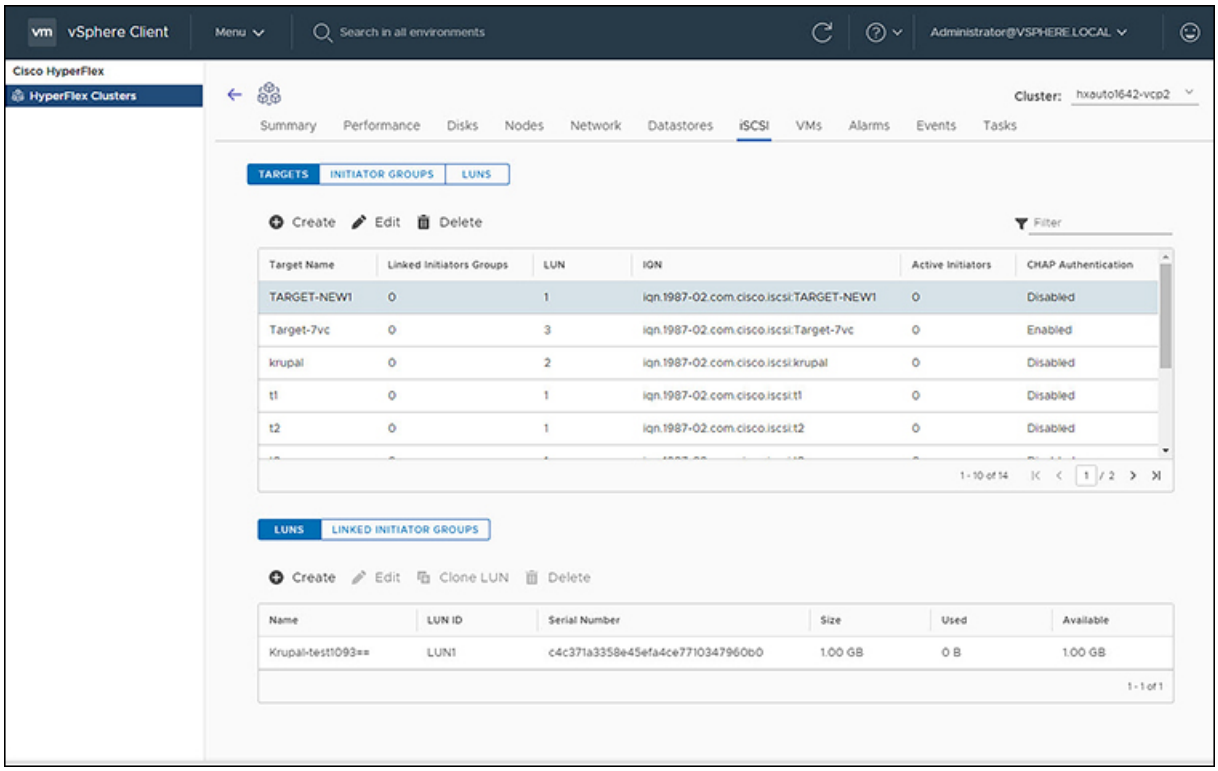
Step 5 Review the Tasks page to verify the iSCSI Network was created.

iSCSI

iSCSI: Targets

After the iSCSI network is created, the iSCSI page appears in the list of navigation tabs. The default view is **Targets**, use the **Create**, **Edit**, **Delete**, and **Clone LUN** buttons to manage iSCSI targets.

Note The iSCSI page only appears in navigaton tabs of clusters with a configured iSCSI network.



Before you begin

- The iSCSI features are supported in Cisco HyperFlex Release 4.5(x) and later.
- Beginning with HX Release 5.0(1a), the Create and Delete buttons are enabled when the license status is In-compliance.
- Create an iSCSI Network Network: Configure iSCSI Network, on page 328

Procedure

Step 1 Log into the vSphere client.

Step 2 Select **Menu** > **Cisco HyperFlex**
The list of clusters appears.

Step 3 Select a cluster with an iSCSI network configured.
The iSCSI Network page appears with the **Targets**, **Initiator Groups**, and **LUNs** buttons. Use the buttons to navigate between views.

Step 4 Click the **Targets** button to populate the table with the list of targets, along with the **Create**, **Edit**, **Clone LUN**, and **Delete** buttons.

Table 34: Target List

Field Name	Additional Information
Target Name	Name of the iSCSI storage resource on the iSCSI server.
Linked Initiators Groups	Number of Linked Initiator Groups on the cluster.
LUN	Number of LUNS in the Initiator group.
IQN	Qualified Name (IQN) for the Initiator. The IQN format takes the form `iqn.yyyy-mm.naming-authority:unique name`.
Active Initiators	Total number of active Initiators.
CHAP Authentication	Authentication scheme that uses a shared secret and a challenge message to validate the identity of remote clients.

Step 5 Select a Target from the list to display all the LUNs associated with the selected target. in the portlet below the Target list. Use the **Create**, **Edit**, **Clone LUN**, and **Delete** buttons to create, edit, clone or delete LUNs in the selected target.

Table 35: LUNs Details Portlet

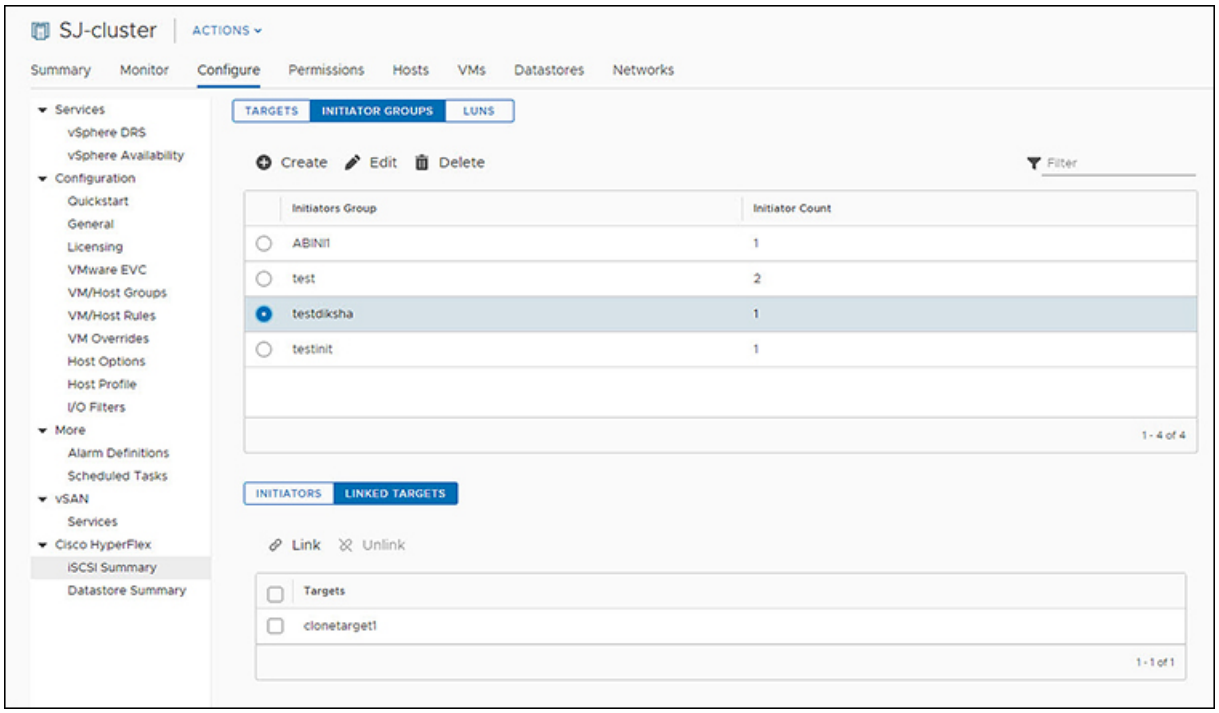
Field Name	Additional Information
Name	LUN name
LUN ID	Unique ID for the LUN
Serial Number	LUN Serial Number
Size	Total capacity size of the LUN (GB)
Used	Total capacity of the LUN used (GB)
Available	Total capacity of the LUN available (GB)

Related Topics

iSCSI LUN Page, on page 296
Creating an iSCSI LUN, on page 296
Editing an iSCSI LUN, on page 297
Cloning an iSCSI LUN, on page 299
Deleting an iSCSI LUN, on page 298
View iSCSI and Datastore Summary from the Configure Tab, on page 352

iSCSI: Initiator Groups

Use the **Initiator Groups** button on the iSCSI page to create, edit, and delete Initiator Groups.



Before you begin

- The iSCSI features are supported in Cisco HyperFlex Release 4.5(x) and later.
- Beginning with HX Release 5.0(1a), the Create and Delete buttons are enabled when the license status is In-compliance.

Procedure

Step 1 Log into the vSphere client.

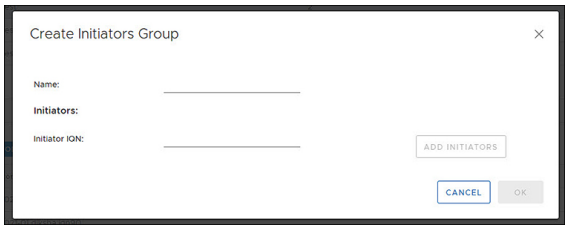
Step 2 Select **Menu** > **Cisco HyperFlex** > **iSCSI**
The iSCSI Network page appears with the **Targets**, **Initiator Groups**, and **LUNs** buttons. Use the buttons to navigate between views.

Step 3 Click the **Initiator Groups** button to populate the table with the list of Initaitor groups, along with the **Create**, **Edit**, and **Delete** buttons.

Table 36: Initiator Group List

Field Name	Additional Information
Initiators Groups	List of groups that specify which hosts can access specified LUNs on the cluster.
Initiator	Number of initiator in the group.

Table 37: Initiator Group Action Window Examples

Action Window Name	Example
Create Initiators Group	
Edit Initiators Group	

Step 4 Select an Initiators Group from the list to display the list of Initiators in the details portlet below the list.

Step 5 Click the **Initiators** button to view the individual Initiators in a group

Step 6 Click the **Linked Targets** button to view the targets associated with the selected group.

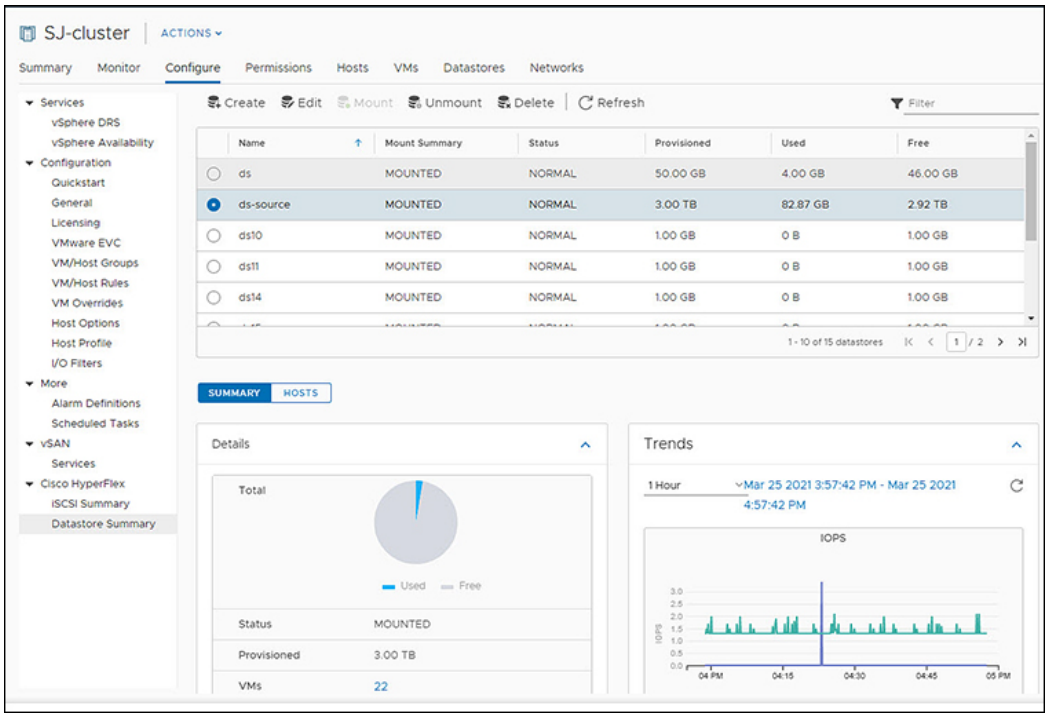
Step 7 Use the **Link**, and **Unlink** buttons to link and unlink targets to groups.

Related Topics

iSCSI Initiator Group, on page 289
Creating an iSCSI Initiator Group, on page 290
Editing an iSCSI Initiator Group, on page 291
Deleting an iSCSI Initiator Group, on page 291
Linking iSCSI Initiator Group with Targets, on page 291
Unlinking an iSCSI Initiator Group, on page 292
View iSCSI and Datastore Summary from the Configure Tab, on page 352

iSCSI: LUNs

Use the LUNS Button to Create, Edit, Delete, and Clone LUN buttons to manage LUNs.



Before you begin

- The iSCSI features are supported in Cisco HyperFlex Release 4.5(x) and later.
- Beginning with HX Release 5.0(1a), the Create, Delete, and Clone LUN buttons are enabled when the license status is In-compliance.

Procedure

Step 1 Log into the vSphere client.

Step 2 Select **Menu** > **Cisco HyperFlex** > **iSCSI**

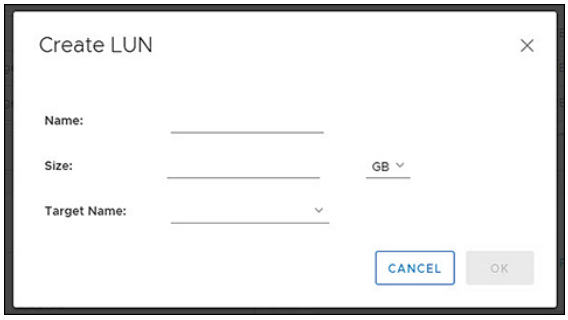
Step 3 Click the **LUNS** button to populate the table with the list of LUNs, along with the **Create**, **Edit**, **Clone LUN**, and **Delete** buttons.

Table 38: LUNs Details Portlet

Field Name	Additional Information
Name	LUN name
LUN ID	Unique ID for the LUN
Serial Number	LUN Serial Number
Size	Total capacity size of the LUN (GB)
Used	Total capacity of the LUN used (GB)

Field Name	Additional Information
Available	Total capacity of the LUN available (GB)

Table 39: iSCSI LUN Action Window Examples

Action Window Name	Example
Create LUN	
Edit LUN	
Clone LUN	

Step 4 Select a LUN on the list to display the Details Portlet and Perforrmance Charts below the LUN list.

Related Topics

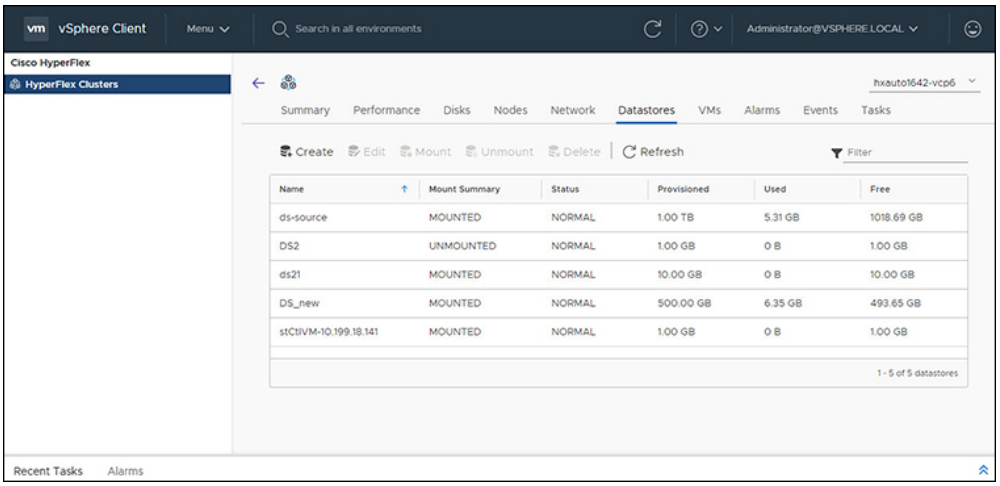
View iSCSI and Datastore Summary from the Configure Tab, on page 352

HX Datastore Management

Managing Datastores

The Datastore page allows users to view datastore details, create, edit, mount, unmount or delete datastores on a Cluster.

Note Beginning with HX Release 5.0(1a), the Create and Delete Datastore buttons are enabled when the license status is In-compliance.



Procedure

Step 1 Log into 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 40: Datastore Table Details

Field Name	Additional Information
Name	Datastore name
Mount Summary	Mounted or Unmounted
Status	Status of the Datastore: Valid values include: 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.
SUMMARY is auto-selected and the Details and Trends portlets appear below the table.

Table 41: 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 42: Trends Portlet

Field Name	Additional Information
General Usage: • 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. • The Scanning Cluster icon indicates that the cluster table is still populating. The icon disappears when the cluster list is complete. • 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

Table 43: Hosts Portlet

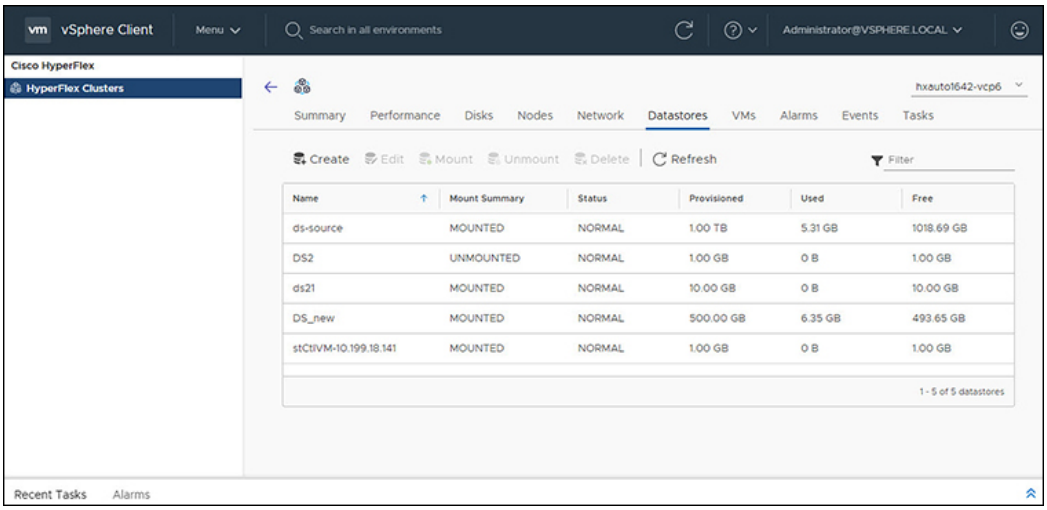
Field Name	Additional Information
Host Name	Displays the IP address of the host for the selected datastore.
Mount Status	Specifies if the Host is Mounted or Unmounted.
Accessable	Specifies if the Host is accessable or not.

Related Topics

View iSCSI and Datastore Summary from the Configure Tab, on page 352

Create New Datastore

To create a new datastore:



Before you begin

- Beginning with HX Release 5.0(1a), the Create Datastore button is enabled when the license status is In-compliance.
- Log into the vSphere client and Select **Menu** > **Cisco HyperFlex**.

Procedure

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.

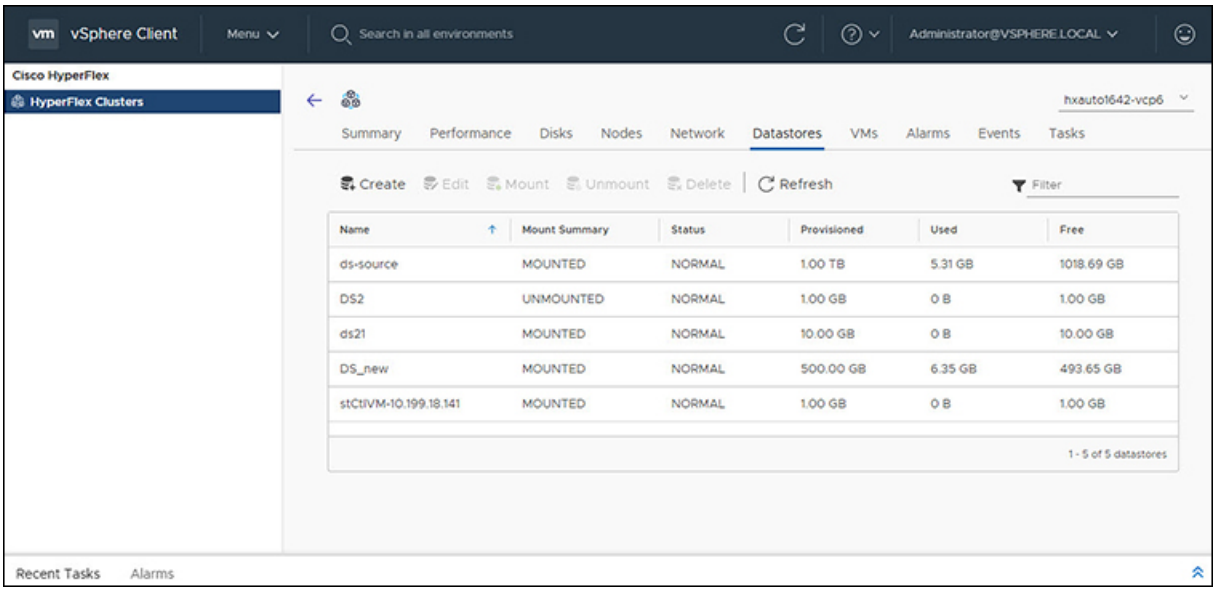
Step 4 Click **OK**. A new Datastore is created and added to the Datastore table list.

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

Log into the vSphere client and Select **Menu** > **Cisco HyperFlex**.

Procedure

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.
- Log into the vSphere client and Select **Menu** > **Cisco HyperFlex**.

Procedure

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

- Beginning with HX Release 5.0(1a), the Delete Datastore button is enabled when the license status is In-compliance.
- Remove any VMs created or registered to the datastore and unmount the datastore before starting the Delete Datastore action.
- Log into the vSphere client and Select **Menu** > **Cisco HyperFlex**.

Procedure

Step 1 Click the cluster that you want.

Step 2 Click **Datastore** and select Datastore for Delete action.

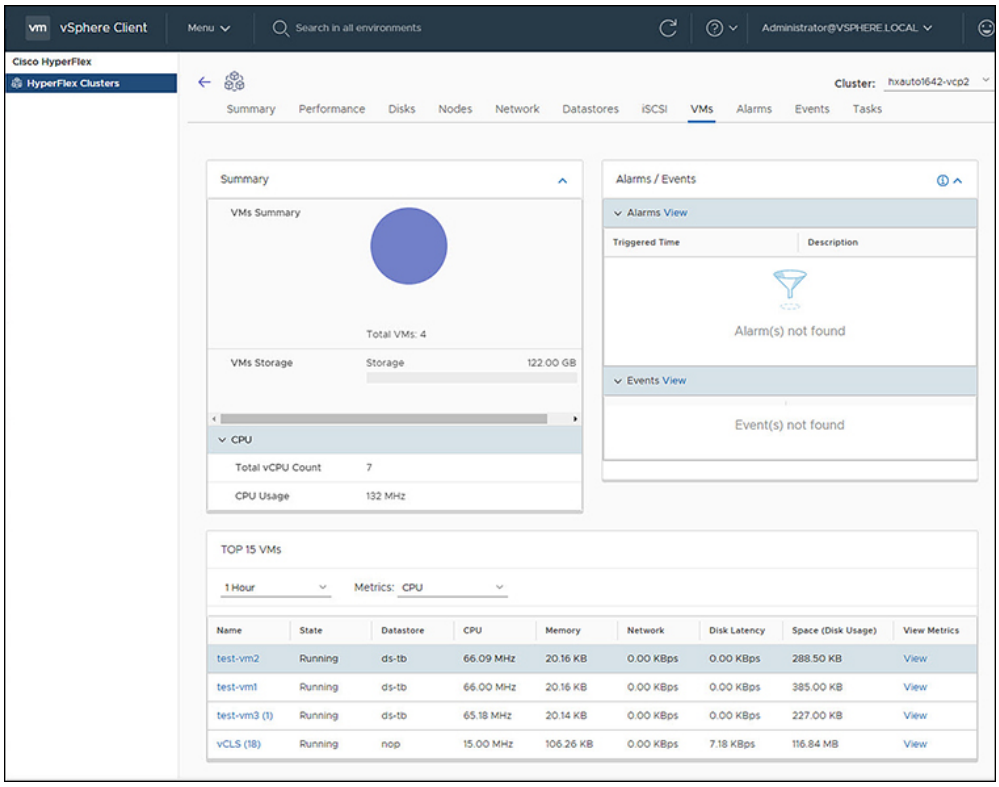
Step 3 Click the **Delete** button.

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.

VMs

To view VM details specific to the Cluster, Host, and VMs, perform the following steps:



Procedure

Step 1 Log into 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 **VMs**.
The VM Detail displays three portlets: Summary, Alarms/Events and Top 15 VMs.

Table 44: Summary Portlet

Field Name	Additional Information
VMs Summary	Usage diagram of user VMs in use. Hover over to view the number of VMs running, suspended, and off. Total VMs: The total count of all user VMs. **Note** Controller VMs are not included in the summary.
VMs Storage	The sum of all user VMs storage. Total storage capacity for all user VMs appears above image. Hover over the graphic to view the current amount of storage being consumed.

Field Name	Additional Information
VMs Memory	Amount of point-in-time memory. Total memory capacity is listed, hover over the graphic to view the current amount of memory used.
CPU	**Total vCPU Count** - Total CPU count for all VMs in the cluster. **CPU Usage** - Number of cycles per second a given CPU is using.

Table 45: Alarms/Events Portlet

Field Name	Additional Information
Alarms	Displays Alarms for the VMs during the last week (7days). Click **View** to navigate to the Alarms Details view. • **Triggered Time**-Date and time the alarm occurred. • **Description**-Alarm Description.
Events	Displays Events for theVMs during the last week (7days). Click **View** to navigate to the Events Details view. • **DateTime** - Date and time the event occurred. • **Description**- Event description.

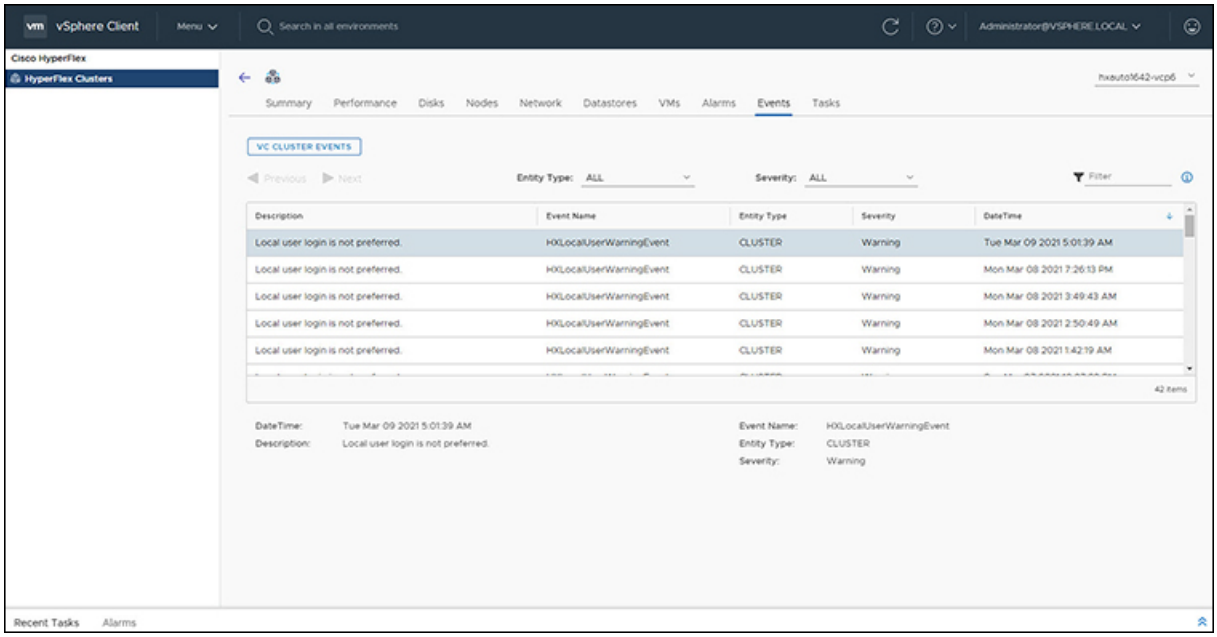
Table 46: Top 15 VMs Portlet

Field Name	Additional Information
Time list	Specify the measurement of time to show the top 15 VMs. List options include: 1 Hour, 1 Day, or 1 Week.
Metrics List	Select the metric type used to populate the table. Options include: CPU, Memory, Disk Latency, Network, and Space. • CPU, Memory, Disk Latency, Network - only report metrics for VMs in a running condition. VMs that are switched off are not included. • Space - counts all VMs regardless of their state.
Name	VM Name - Clicking on the VM name redirects users to the graph or monitoring page of VM being viewed in vCenter.
State	Current state of the VM. Valid values are Running, Off and Suspended.
Datastore	Datastore Name
CPU	CPU usage in megahertz used durineg the interval.
Memory	Amount of host physical memory consumed for backing up guest physical memory pages.

Field Name	Additional Information
Disk Latency	Highest latency value across all disks used by the host
Network Throughput	Network utilization during the interval (combined transmit and receive rates).
Space (Disk Usage)	Amount of disk space the VM is using
View Metrics	Link to view CPU, Memory, Disk Latency, and Network Throughput performance tables for that specified VM. The usage values displayed are for the five-minute average for all. Use the hover feature to display the matrices simultaneously and evaluate any visible spikes in the data.

Events

To view events specific to the Cluster, Node, Host, VM or Disk, perform the following steps:



Procedure

Step 1 Log into 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 47: 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:

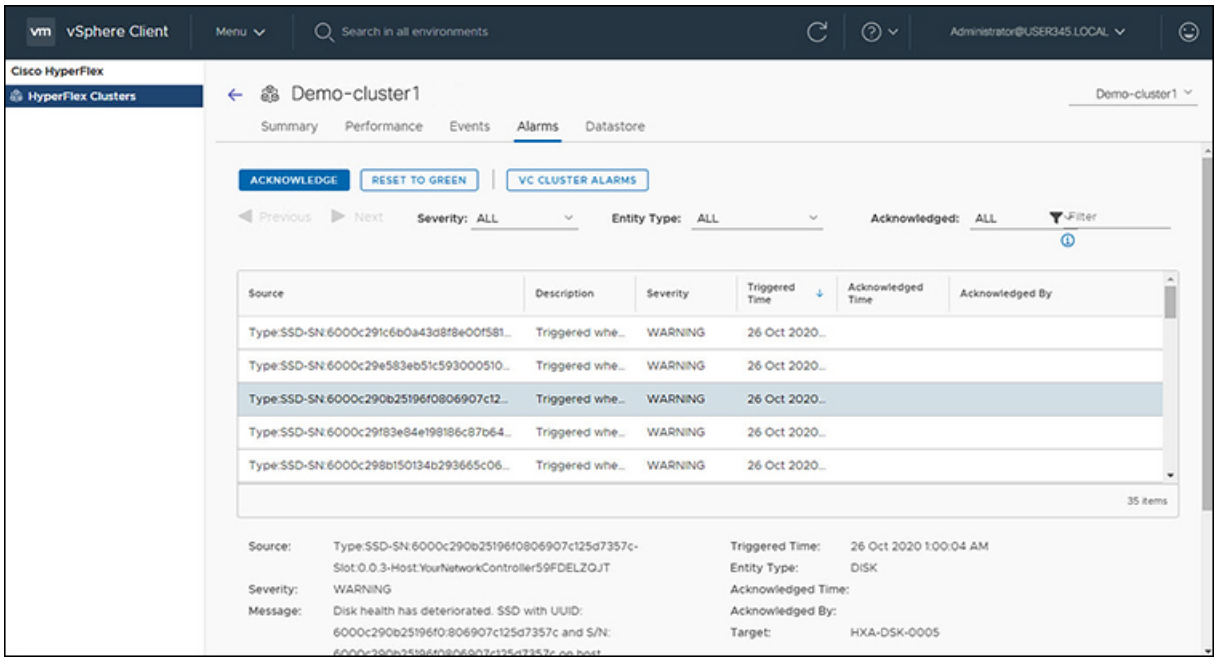
- Description
- Event Name
- Entity Type
- Severity
- DateTime

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.



Procedure

Step 1 Log into 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 48: Alarm Details

Field Name	Additional Information
Source	Text description of the alarm
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.

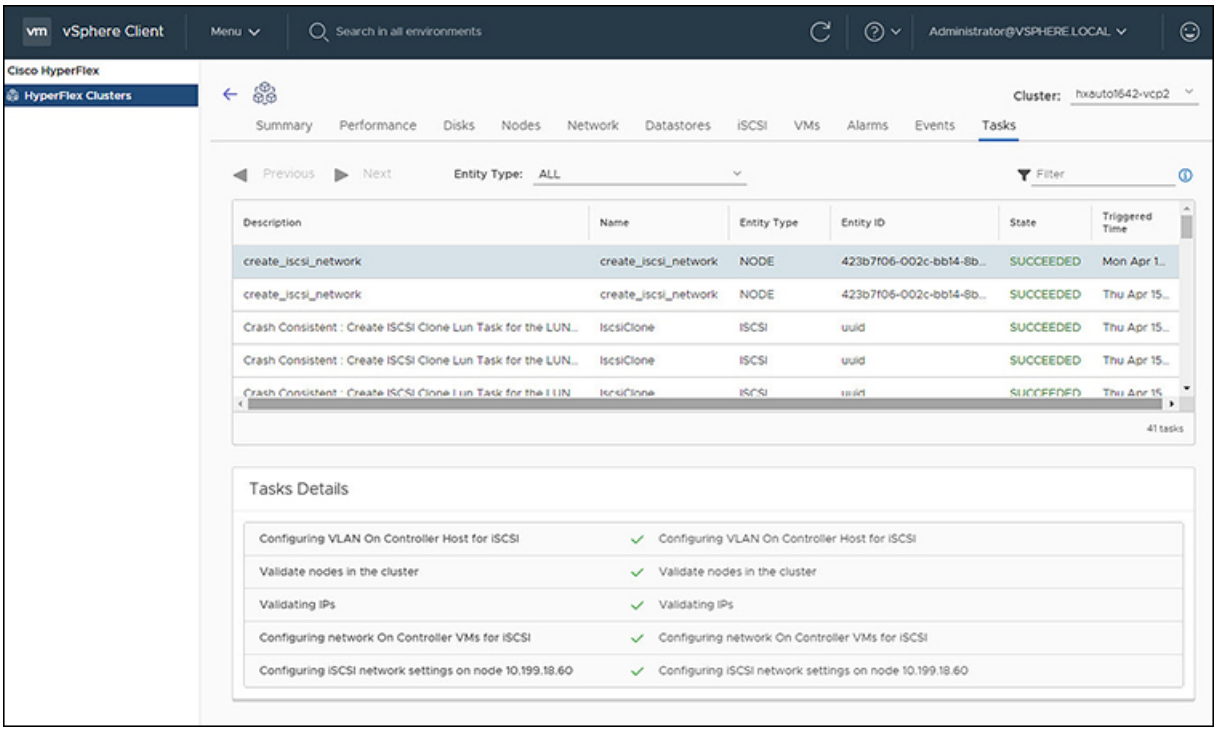
Filter	Filter Options
Severity	All, Info, Warning, and Error
Entity Type	All, Cluster, Node, Virtual Machine, Disk and Datastore
Acknowledged	All, True, and False
Filter	Type a keyword in the Filter option to filter the table contents seen in the browser

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.

Tasks

View asynchronous tasks that are happening on the platform to validate maintenance, perform the following steps:



Procedure

Step 1 Log into the vSphere client.

Step 2 Select **Menu** > **Cisco HyperFlex** > **Tasks**

Step 3 Click the Task that you want to view.
The sub-tasks appear in the table below the list of tasks.

Table 49: Task List

Field Name	Additional Information
Description	Description of the task
Name	Task Name
Entity Type	Type of task, Valid values include: NODE, DP_Summary, Virtual Machine, Disk and Datastore.
Entity ID	Device ID number
State	Specifies the success or failure of the task.
Triggered Time	Date and time the task occurred

Table 50: Task Details

Field Name	Additional Information
Sub-task name	Name of the task.
Success indicator	Description of the action and the status of the task when it completed. A check icon preceding the description identifies that the task was successful. Review this list to identify where a task failed.

Step 4 (Optional) Use the **Entity Type** list to filter the table results.

vCenter: HyperFlex Plugin Embedded Actions

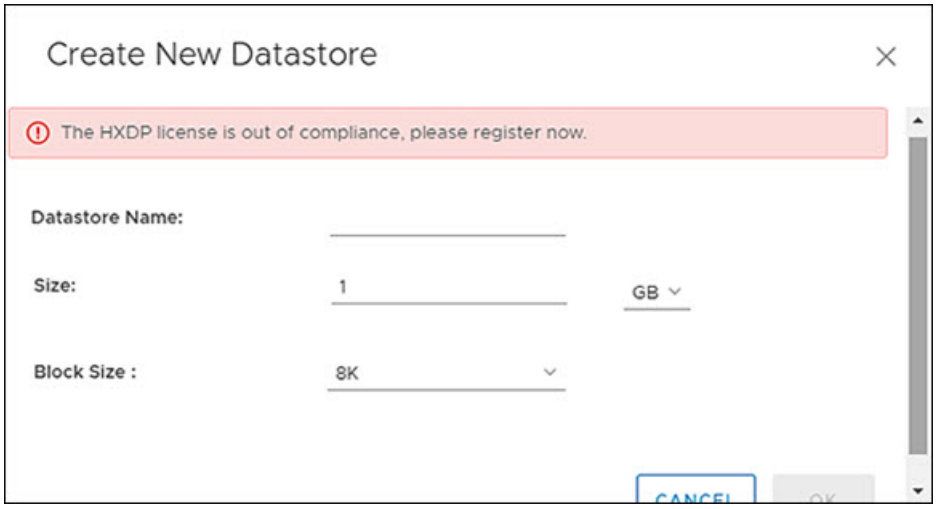
vCenter Server Actions at the Host and Cluster Level

Create New Datastore

To create a new datastore from the Hosts and Clusters level, perform the following steps:

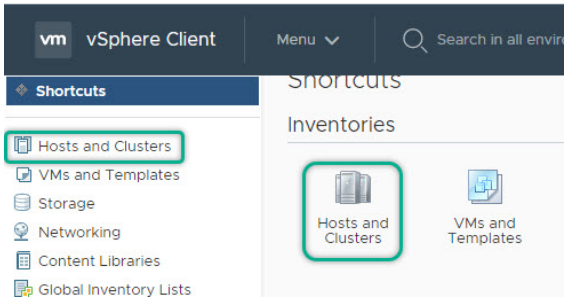
Before you begin

Beginning with HX Release 5.0(1a), this functionality is enabled when the license status is In-compliance.



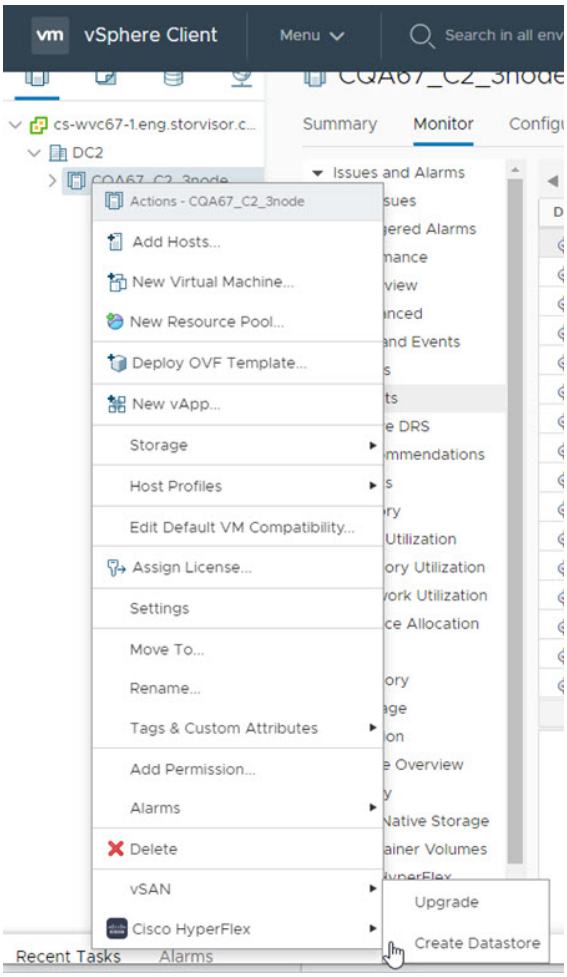
Procedure

Step 1 Access the **Hosts & Clusters** either from the vSphere Menu or the Shortcut link.



Step 2 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.

Step 3 Right click on the cluster and select **Cisco HyperFlex** > **Create Datastore**.



The Create New Datastore window appears.

Step 4 Complete the fields in the New Datastore Window.

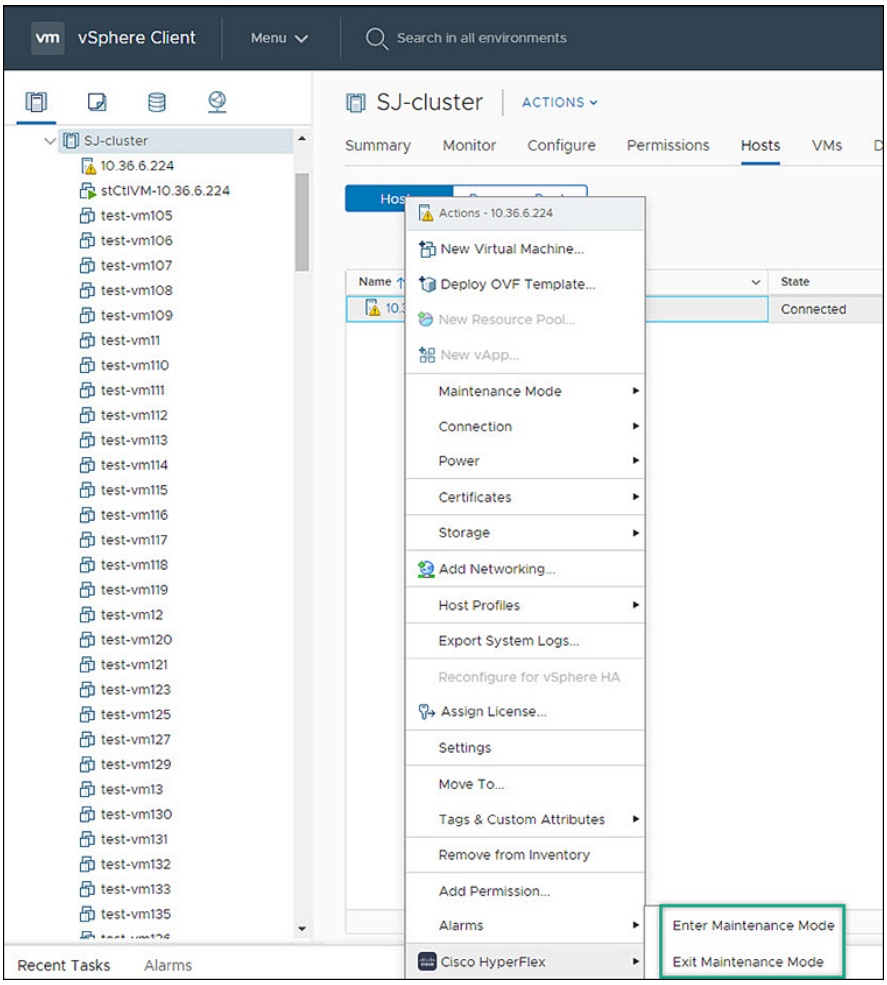
- Enter the Datastore Name
- Enter the Size, and select GB or TB.
- Select the Block Size, Select 4K or 8K.
- Click **OK**.

Related Topics

Create New Datastore, on page 338

Enter or Exit Maintenance Mode

To enter or exit Maintenance Mode from the Host level from the vSphere web UI perform the following steps:



Before you begin

Beginning with HX Release 5.0(1a), this functionality is enabled when the license status is In-compliance.

Procedure

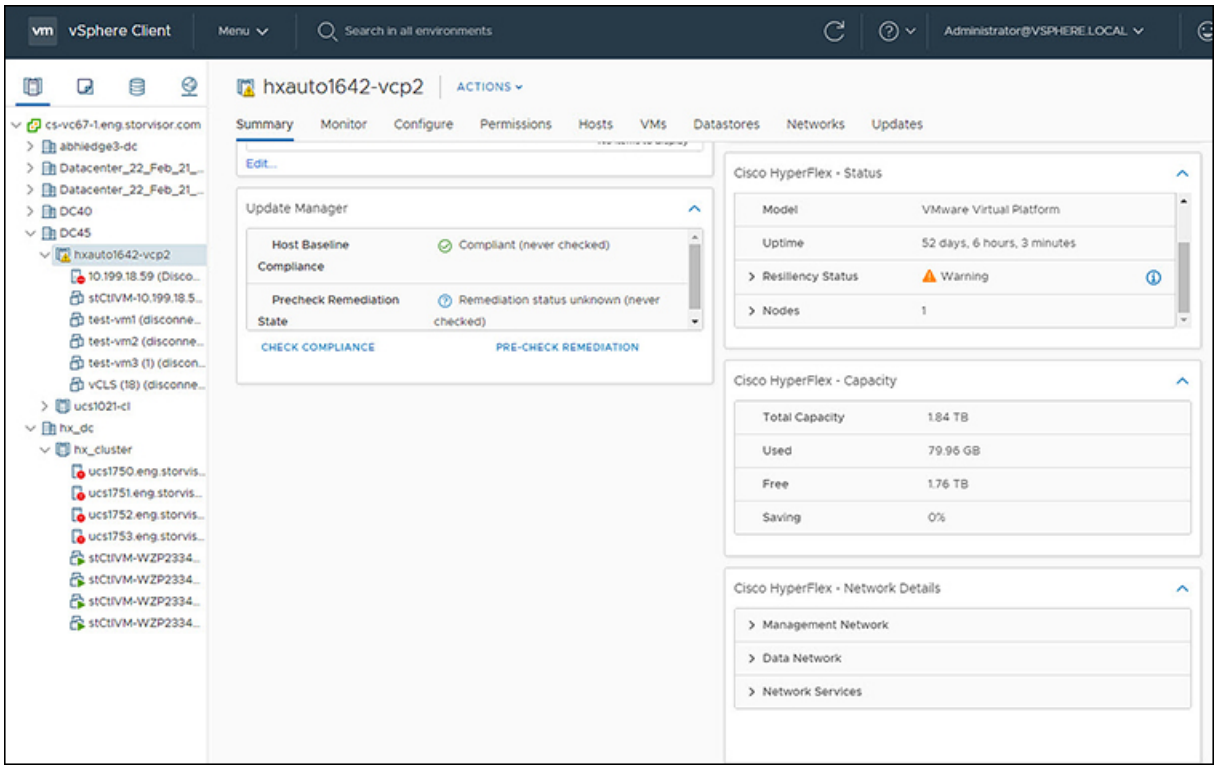
Step 1 Access **Hosts & Clusters** either from the vSphere Menu or the Shortcut link.

Step 2 Click on the cluster name and select the **Hosts** tab.
The summary page appears.

Step 3 Right click on the host and select **Cisco HyperFlex** > > **Maintenance Mode** > **Enter (or Exit) MainenanceMode** to enter or Exit HyperFlex Maintenance Mode.

View HTML5 Plugin Portlets from the Summary Tab

To view Cisco HyperFlex HTML5 Plugin Portlets from the vSphere web UI perform the following steps:



Procedure

Step 1 Access **Hosts & Clusters** either from the vSphere Menu or the Shortcut link.

Step 2 Click on a cluster name and select the **Summary** tab. The summary page appears.

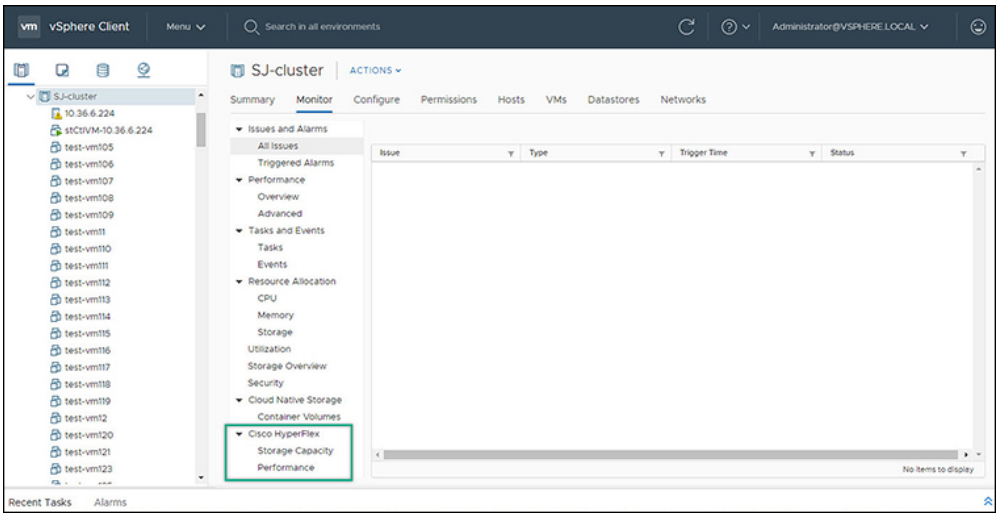
Step 3 Scroll down and use the arrow in each portlet to show or hide the portlet details.

Related Topics

View the HX Cluster Summary, on page 316

View HTML5 Plugin Portlets from the Monitor Tab

To view Cisco HyperFlex HTML5 Plugin Portlets from the vSphere web UI perform the following steps:



Procedure

Step 1 Access **Hosts & Clusters** either from the vSphere Menu or the Shortcut link.

Step 2 Click on a cluster name and select the **Monitor** tab.

Step 3 Scroll down the Monitor navigation panel and locate **Cisco HyperFlex**.

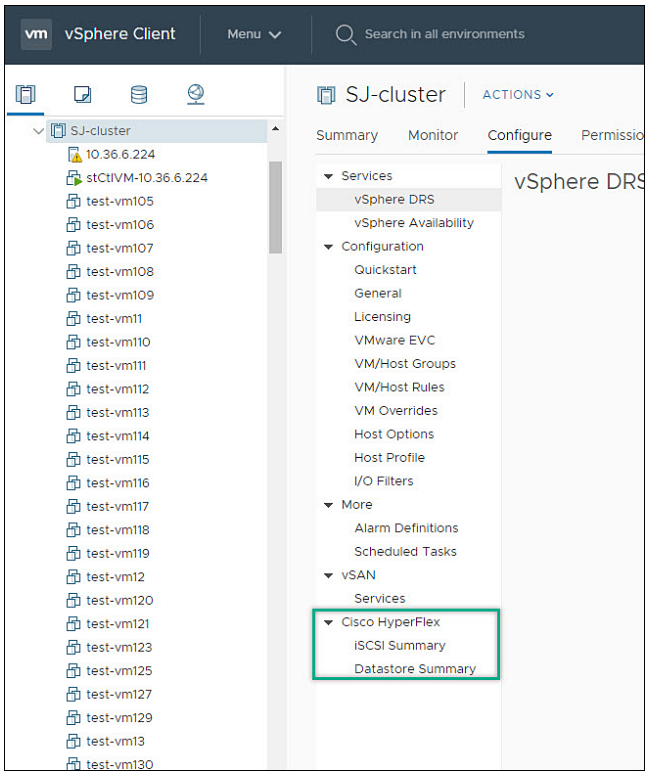
Step 4 Click on **Storage Capacity** or **Performance** to display the related Cisco HyperFlex HTML5 Plugin chart.

Related Topics

View Cluster and Datastore Performance Charts , on page 322

View iSCSI and Datastore Summary from the Configure Tab

To view the iSCSI and Datastore Summary Pages from the vSphere web UI perform the following steps:



Before you begin

The iSCSI features are supported in Cisco HyperFlex Release 4.5(x) and later.

Procedure

Step 1 Access **Hosts & Clusters** either from the vSphere Menu or the Shortcut link.

Step 2 Click on a cluster name and select the **Configure** tab.

Step 3 Scroll down the Monitor navigation panel and locate **Cisco HyperFlex**.

Step 4 Click on **iSCSI Summary** or **Datastore Summary** to display the related Cisco HyperFlex HTML5 Plugin page.

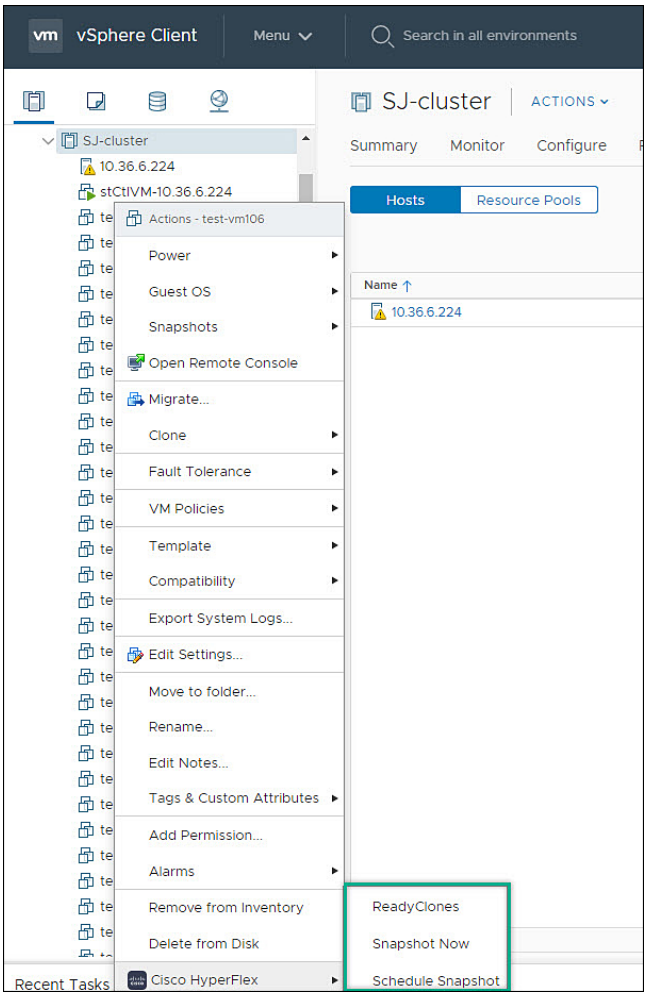
Step 5 Use the buttons to perform all maintenance tasks as defined in the related topics.

Related Topics

Managing Datastores, on page 335
Create New Datastore, on page 338
Edit Datastore, on page 338
Mount or Unmount Datastore, on page 339
Delete Datastore, on page 340
iSCSI: Targets , on page 330
iSCSI: Initiator Groups, on page 332
iSCSI: LUNs, on page 333

vCenter Server Actions at Virtual Machine Level

Snapshot Now



Before you begin

Access the **VMs and Templates** either from the vSphere Menu or the Shortcut link.

Procedure

Step 1 Right click on the virtual machine. Select **Cisco HyperFlex** > **Snapshot Now**.

Step 2 The **Take VM Native Snapshot** window appears. Complete the following fields:

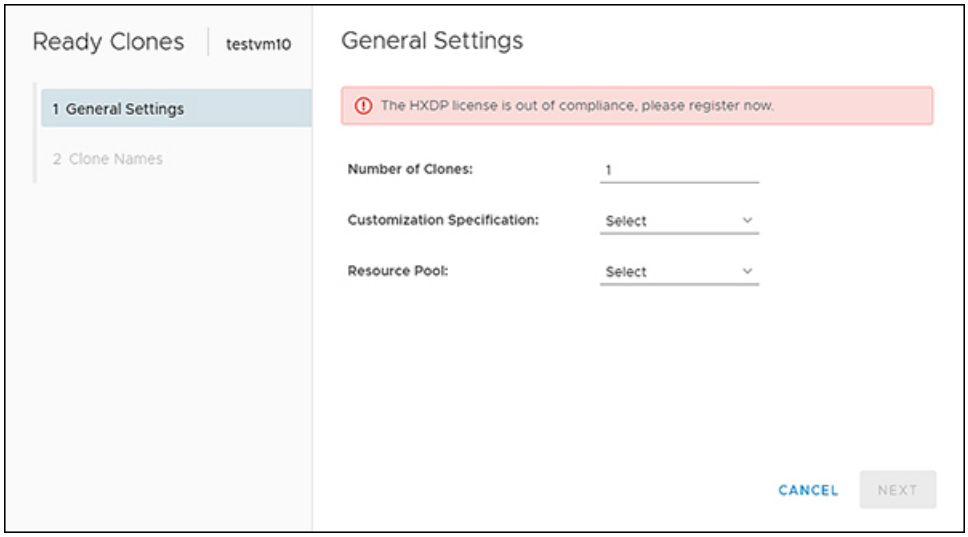
- **Name**- Snapshot name
- **Description** - Description of the snapshot
- **Quiesce guest file system** - Check box.

Step 3 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.

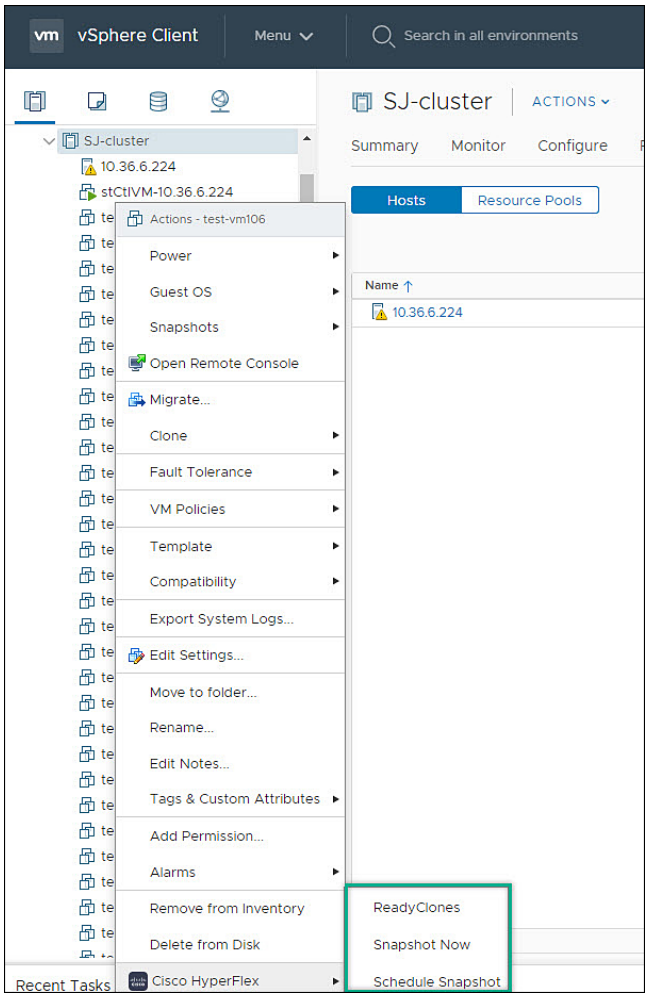
ReadyClones

Before you begin

- Beginning with HX Release 5.0(1a), the ReadyClones feature is enabled when the license status is In-compliance.



- Access the **VMs and Templates** either from the vSphere Menu or the Shortcut link.



Procedure

Step 1 Right click on the virtual machine. Select **Cisco HyperFlex** > **ReadyClones**.

Step 2 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

Step 3 Complete the Clone Name Fields:

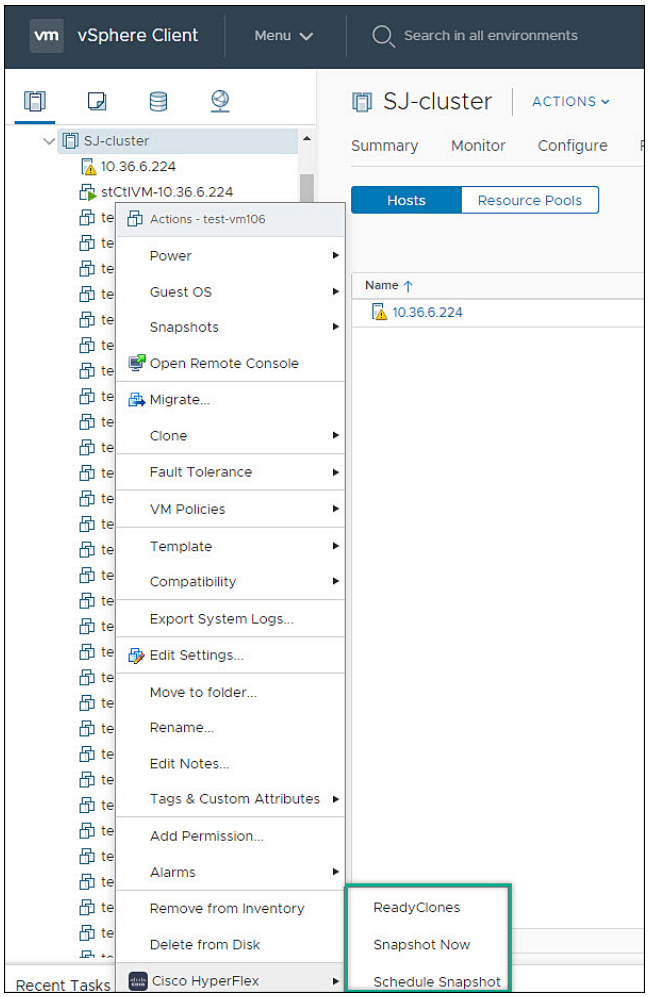
- **Power on VMs after cloning**- Check box
- **Name of VM Prefix** - Type VM prefix
- **Starting clone number**- The default is 1
- **Increment clone number by** - The default is 1

- **Use same name for Guest Name**- Uncheck to provide a guest name

Schedule Snapshot

Before you begin

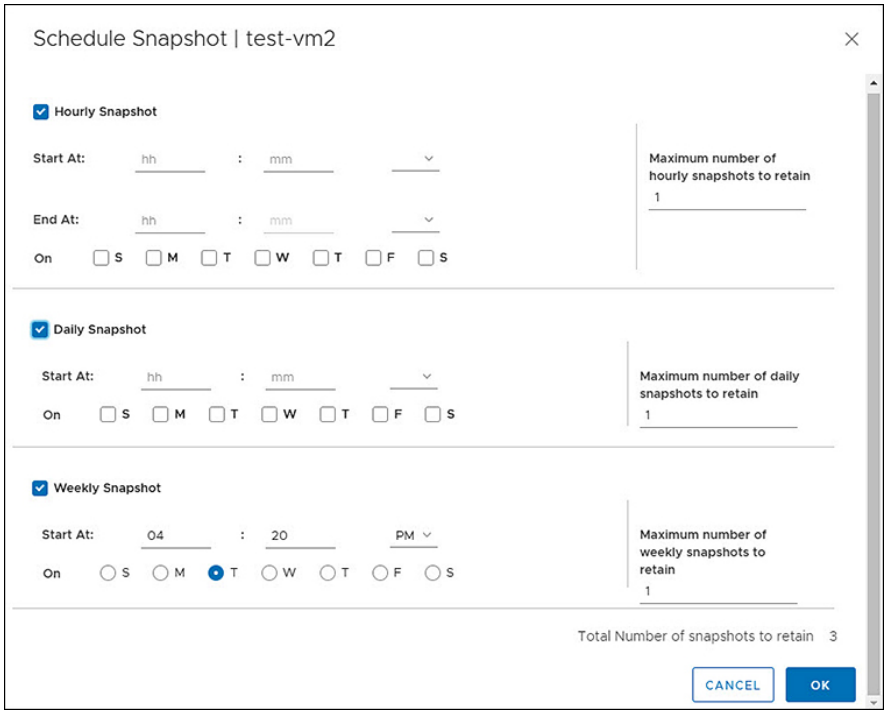
- The Schedule Snaphot feature is supported in Cisco HyperFlex Release 4.5(x) and later.
- Beginning with HX Release 5.0(1a), the Schedule Snapshot feature is enabled when the license status is In-compliance.
- Schedule Snapshot on VMs in a suspended state is not supported.
- Access the **VMs and Templates** either from the vSphere Menu or the Shortcut link.



Procedure

Step 1 Right click on the virtual machine. Select **Cisco HyperFlex** > **Schedule Snapshot** .
The **Schedule Snapshot** window appears.

Step 2 Use the checkbox to select the snapshot frequency



Step 3 Complete the fields for the Schedule Snapshot selected:

Table 51: Hourly Snapshot:

Start At:	Enter valid start time • Hour: Valid values 1-24 • Minutes: Valid values 1-60 • AM/PM: Select one from the list.
End At:	Enter valid start time • Hour: Valid values 1-24 • Minutes: Valid values 1-60 • AM/PM: Select one from the list.
On	Use the Check boxes to select the day(s) of the week for snapshots are to be taken

Maximum Number of hourly snapshots to retain	Type or select a value between 1-30.

Table 52: Daily Snapshot:

Start At:	Enter valid start time • Hour: Valid values 1-24 • Minutes: Valid values 1-60 • AM/PM: Select one from the list.
On	Use the Check boxes to select the day(s) of the week for snapshots are to be taken
Maximum Number of hourly snapshots to retain	Type or select a value between 1-30.

Table 53: Weekly Snapshot:

Start At:	Enter valid start time • Hour: Valid values 1-24 • Minutes: Valid values 1-60 • AM/PM: Select one from the list.
On	Use the Check boxes to select the starting day for the weekly snapshot.
Maximum Number of hourly snapshots to retain	Type or select a value between 1-30.

Step 4 Click **OK** to confirm your snapshot schedule.

vCenter Server Actions at the Storage Level

Edit Datastore

From the Datastore level users have the ability to edit existing datastores.

Before you begin

Access the **Datastores** either from the vSphere Menu or the Shortcut link.

Procedure

Step 1 Right click on the datastore name.

Step 2 Select **Cisco HyperFlex** > **Edit Datastore**.
The Edit Datastore window appears.

Step 3 Edit the Datastore details.

Step 4 Click **OK**to save your changes.

Related Topics

Edit Datastore, on page 338

Delete Datastore

From the Datastore level users have the ability to delete existing datastores.

Before you begin

Access the **Datastores** either from the vSphere Menu or the Shortcut link.

Procedure

Step 1 Right click on the datastore name.

Step 2 Select **Cisco HyperFlex** > **Delete Datastore**
The Delete Datastore window appears.

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.

Related Topics

Delete Datastore, on page 340

Cisco HyperFlex Remote Plugin for VMware vCenter

Starting with vSphere 8.0.0 the only architecture supported in vSphere releases is remote plugin.

The remote plugin runs on its own appliance, independent of the vCenter appliance; This autonomy improves the user experience with increased vCenter performance.

Prerequisites

The following is required for using the Remote Plugin for VMware:

- **Supported Cisco HyperFlex Release:** 5.0(2a) and later (ESXi 6.7 U3 and later).

Using the Cisco HyperFlex Remote Plugin

The following table defines feature support by plugin version:

Table 54: HTML5 Remote Plugin Feature Support

Feature	Plugin Version 3.0.0
Discover the Registered HX Cluster	✓
Rename Clusters [11]	✓
View HX Cluster Summary	✓
View Cluster and Datastore Performance Charts	✓
Disks View	✓
Nodes View	✓
HX Datastore Management	✓
VM Summary and Top VM Consumers	✓
Network Management	✓
iSCSI Management [12]	✓
Events and Alarms	✓
Manage Tasks	✓
HX Snapshots and clones at the virtual machine level	✓
Schedule Snapshot [13]	✓
Manage users and access to HX clusters	✓
Cross-launch HX Connect for upgrade	✓
Embedded vCenter server actions at the Host and Clusters level	✓
HTML 5 License Status [14]	✓
Linked Mode	✓

11 Requires HXDP Release 4.5(x) or later.
12 Requires HXDP Release 4.5(x) or later.
13 Requires HXDP Release 4.5(x) or later.
14 Requires HXDP Release 5.0(x) or later.

Note The configuration and feature functionality for the Remote and Local plugin are identical. For more information on any feature see the Cisco HyperFlex HTML5 Plugin for VMware vCenter , on page 308 section of this guide.

Install, Register and Upgrade the Remote Plugin

Remote Plugin Installation and Registration

To install and register the Remote Plugin, perform the following steps.

Before you begin

Users with an active firewall need to verify that ports 433, 9443 and 22 are open and allow traffic to pass in or out.

Default Appliance Credentials

- **Username:** vcp-admin
- **Password:** C^scohxplugin@1984

Procedure

Step 1 Download the Cisco HyperFlex HTML plugin OVA for VMware vCenter from the Cisco Software Download site.

Step 2 Login to the vCenter and select the host you want to deploy the Remote Plugin appliance.

Step 3 Right click on the host and click on **Deploy OVF Template**. Deploy the OVA in a vCenter with appropriate static/DHCP IP settings.

Step 4 Navigate to and select the local OVF downloaded from Cisco Software Download site.

Recommended configuration settings for the Remote plugin appliance:

- RAM: 4G
- Cores:2
- Datastore: 1 with 50 GB of minimum space
- Valid network adapter

Step 5 Follow the wizard to complete the deployment process.

Note
The static/DHCP IP assignment is supported through a script. Perform the following steps:

a. Open the appliance VM console and log in to the appliance with the default credentials.

b. Type **`hx-ip-address-change`** and press **Enter** to initiate the script.

c. Select whether you want to set **Static** IP configuration or **DHCP IP** settings for the MAC address of the appliance appliance.

Tip
The recommendation is to configure a static IP address to the virtual Appliance or use DHCP with MAC bound IP.

d. Type valid values for the following fields: IP address, Subnet mask, Gateway, and DNS when asked if you opt for Static IP assignment. There is no need to provide the DNS details if you select DHCP IP assignment.

e. Use the **ifconfig** command to display the Static/DHCP IP. Write these down for future use.

Step 6 Copy and paste the IP address of the appliance in a browser window.

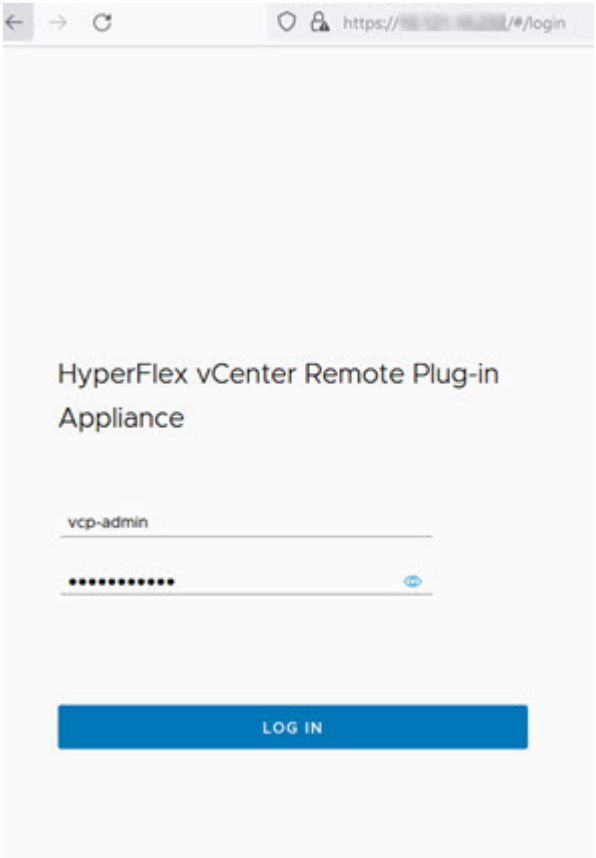
Example:

https://<appliance_ip>
The UI will open.

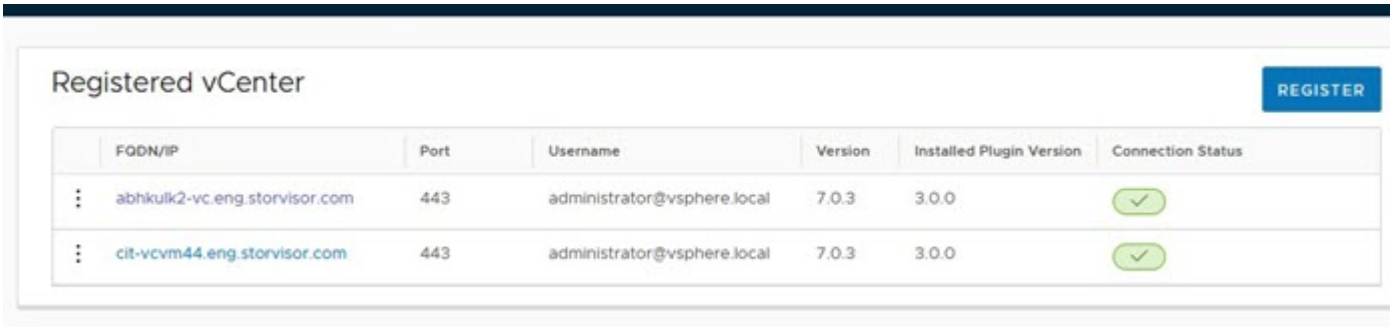
Step 7 Reset the default user (vcp-admin) password through command line with the following steps:

a) Open the appliance VM console and log in to the appliance with the default credentials.
b) Type **`passwd vcp-admin`** and press **Enter**.
c) Type the old and new password as directed in the UI and complete the password reset process.

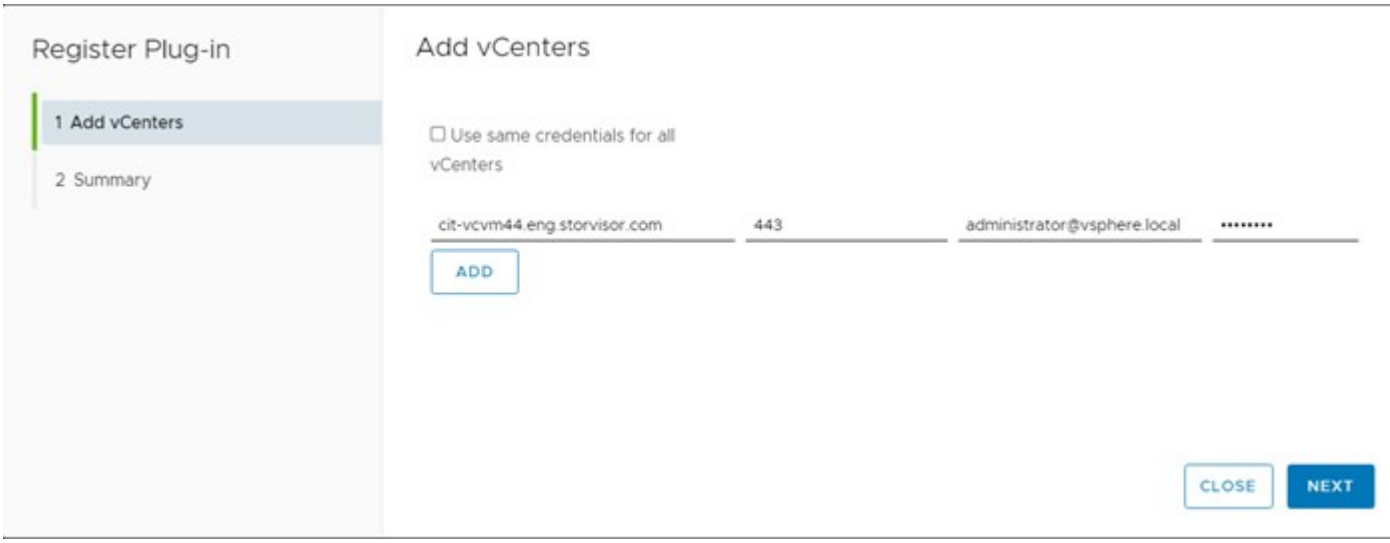
Step 8 Log in to the HX vCenter Remote Plug-in Appliance with the new credentials.



Step 9 To register the remote plugin with a vCenter, click the **Register** button. The **Add vCenters** window appears.



Step 10 In the **Add vCenters** window click the **Add** button and type the vCenter IP/FQDN, Port number, user name and password in the designated fields. Click **Next**.



If the vCenter is reachable and credentials are valid, the UI will advance to the **Register Plugin** page.

Step 11 Click **Register** and wait for the registration process to complete.

Step 12 Use the hyperlink provided in the registration UI to log on to vCenter.

Step 13 To view the Cisco HyperFlex HTML5 plugin options in the vSphere UI, log out and log in again to vCenter.

Uninstalling and Unregistering the HyperFlex Remote Plugin from vCenter

To Uninstall and unregister the HX Remote Plugin from vCenter, perform the following steps.

Procedure

Step 1 Navigate to the plugin appliance UI.

Step 2 Log in with valid credentials.

Step 3 Select the vCenter from the list that you want to unregister and click on actions buttons present before the vCenter FQDN/IP. Click **Unregister**.

Step 4 Specify the vCenter credentials and confirm that you want to unregister/uninstall the remote plugin from the vCenter.

Step 5 Log out from the vCenter, wait a few minutes and Log in again.

Upgrade the Remote Plugin Application 3.0.0 using CLI

Upgrading the remote plugin is a two-step process, perform the following steps to complete the upgrade.

Procedure

Step 1 Download the newer version of the vCenter remote plugin (hyperflex-remote-vcenter-plugin-x.x.x-xxx.tar.gz) upgrade package from the Cisco Software Download site.

Step 2 Copy it into /tmp directory on the plugin appliance VM.

Step 3 Execute the upgrade script by using the **hx-plugin-upgrade <UpgradePackagePath>** command.

Step 4 Review the messages carefully and select **Y** to continue the upgrade process.

Step 5 Once the vCenter application upgrade is complete, update the vCenter extension to the latest version on each vCenter registered with the plugin server.

a) **vCenter Extension Upgrade:** Log in to the plugin appliance UI and select the **vCenter**.
b) Click on **Update** and provide the credentials and click the **Update** button.

Encryption Support

Remote Plugin Encryption Support

To enable remote plugin encryption, perform the following steps, For more information on enabling Software Encryption on your cluster, see Enabling HyperFlex Software Encryption Workflow.

Procedure

Step 1 Select the cluster you want to encrypt.

Step 2 Click **Datastore**.

Step 3 Click the **Create** button. The Create New Datastore window appears.

a) Type the Datastore Name.
b) Type the size and select GB or TB.
c) Select the Block Size, Select 4K or 8K.
d) Check the Software Encryption check box.

Step 4 Click OK. A new Datastore is created and added to the Datastore table list

If the new Datastore does not appear in the list, click the **Refresh**arrow and recheck the list.

Generate Support Bundles

Plugin Support Bundle Generation

Currently this feature is available via command line which will help user download support bundle which includes all the required log files from appliance and vCenter.

Steps to generate Support bundle from the vCenter appliance through command line:

Before you begin

To generate support bundle for vCenter, root user credentials are required. If the root credentials are not available, download it from https://<VC_IP/FQDN>:5480 using administrator credentials.

SUMMARY STEPS

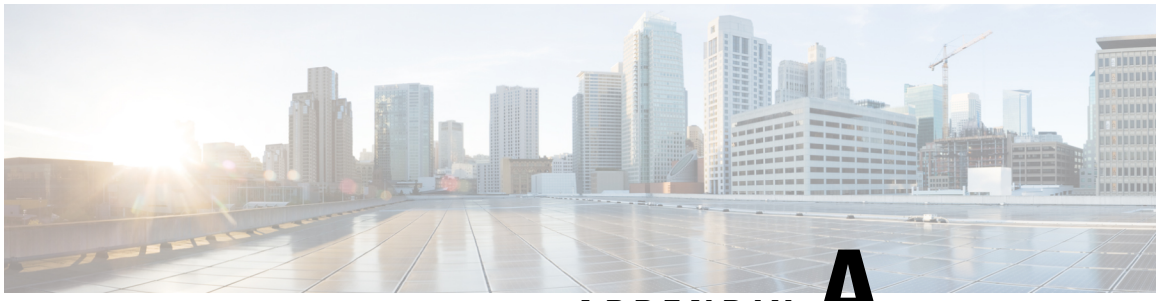
1. Log in to the vCenter appliance using pre-configured user credentials.
2. Enter the **hx-plugin-supportbundle** command. This will prompt you for the vCenter for which you want to generate the support bundle.
3. Provide the root credentials or skip the vCenter logs and continue with support bundle generation of vCenter appliance logs.
4. The default support bundle download location is: /var/log/plugin_support/.

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	Log in to the vCenter appliance using pre-configured user credentials.	
Step 2	Enter the **hx-plugin-supportbundle** command. This will prompt you for the vCenter for which you want to generate the support bundle.	
Step 3	Provide the root credentials or skip the vCenter logs and continue with support bundle generation of vCenter appliance logs.	
Step 4	The default support bundle download location is: /var/log/plugin_support/.	Specifying a custom directory to download the support bundle is supported.

Help support for the support bundle generation tool is **hx-plugin-supportbundle -help** or **hx-plugin-supportbundle -h**



APPENDIX A

Appendix

- Creating VLANs for HX Servers, on page 367
- Creating MAC Address Pools, on page 368
- Configure the vSwitches, on page 370
- Migrating vMotion Networks to Virtual Distributed Switches (VDS) or Cisco Nexus 1000v (N1Kv), on page 371

Creating VLANs for HX Servers

Procedure

Step 1 Open a web browser and enter the IP address for Cisco UCS Manager. Enter the login credentials.

Step 2 Navigate to **LAN tab** > **LAN** > **LAN Cloud** > **VLANS**.

Step 3 Right-click and select Create VLANs as shown in the table below:

VLAN Name	Description	Multicast Policy Name	VLAN ID (by default)
hx-inband-mgmt	Used for: • ESX management • SSH to storage controller VM • HX Cluster management IP - using multicast traffic. • vCenter connectivity to the HyperFlex VM for the HX Data Platform plug-in	HyperFlex	3091
hx-storage-data	Used for: • ESX NFS client(IOvisor) • HyperFlex replication/cluster • Cluster data VIP	HyperFlex	3092

VLAN Name	Description	Multicast Policy Name	VLAN ID (by default)
hx-vmotion	Used for: • VM and storage vmotion, FT, iSCSI	HyperFlex	3093
insert existing vlan name	Used for: • VM data traffic	HyperFlex	Any*

Note:

- Configuration option is Common/Global. It applies to both fabrics and uses the same configuration parameters in both cases.
- *There is no specific recommendation for VM data VLANs. You can create your own VLANs for the VM data traffic. By default, the HXDP installer will not create VLANs for the VM data traffic.
- Installer sets the VLANs as non-native by default. Ensure to configure the upstream switches to accommodate the non-native VLANs.

Creating MAC Address Pools

You can change the default MAC address blocks to avoid duplicate MAC addresses. Each block contains 100 MAC addresses by default to allow for up to 100 HX server for deployment per UCS system. We recommend that you use one MAC pool per vNIC for easier troubleshooting.



Note The 8th digit is set to A or B. "A" is set on vNICs pinned to Fabric Interconnect A. And "B" is set on vNICs pinned to Fabric Interconnect B.

Procedure

Step 1 Open a web browser and enter the IP address for Cisco UCS Manager. Enter the login credentials.

Step 2 In Cisco UCS Manager, navigate to **LAN tab** > **Pools** > **root** > **Sub-org** > **hx-cluster** > **MAC Pools**.

Step 3 Right-click **MAC Pools** and select **Create MAC Pool**.

Step 4 In the **Define Name and Description** page of the **Create MAC Pool** wizard, complete the required fields as shown below:

MAC Pool Name	Description	Assignment Order	MAC Address block
hv-mgmt-a	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:01:01-64

hv-mgmt-b	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:02:01-64
storage-data-a	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:03:01-64
storage-data-b	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:04:01-64
vm-network-a	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:05:01-64
vm-network-b	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:06:01-64
hv-vmotion-a	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:07:01-64
hv-vmotion-b	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:08:01-64

Step 5 Click **Next**.

Step 6 In the **Add MAC Addresses** page of the **Create MAC Pool** wizard, click **Add**.

Step 7 In the **Create a Block of MAC Addresses** dialog box, complete the following fields:

Name	Description
First MAC Address field	Enter the first MAC address for the block. A MAC address is 6 bytes (48 bits) long, typically displayed as 12 hexadecimal digits separated by colons. It is recommended to use the prefix 00:25:b5 to ensure unique MAC addresses in the LAN fabric. Example, 00:25:b5:xx:xx:xx.
Size field	Specify the number of MAC addresses in the block. This determines how many consecutive addresses will be generated starting from the first MAC address. The size can range from 1 to 1000 per block.

Step 8 Click **OK**.

Step 9 Click **Finish**.

After the MAC address change, ESXi will be reconfigured as it was configured earlier. If management IP was DHCP assigned, then the IP will change.

Impact of Manufacturing process on MAC address change

- The MAC address will change between the manufacturing process and the customer site, especially if the customer orders HyperFlex serves without UCS Fabric Interconnects.
- A MAC address is configured during Service Profile association. It is un-configured during Service Profile disassociation.
- At the end of manufacturing process, the service profiles are disassociated, hence the MAC addresses are un-configured.
- When a HyperFlex server is deployed, configure the MAC address pools as described above.
- VMWare supports Consistent Device Naming, but issues have been reported since 5.5.SR has been opened.

Configure the vSwitches

In both VMware ESX and ESXi host, you can configure vSwitches from either the GUI or the command line.

The CLI configurations are very helpful when you are installing multiple ESX servers and planning to script the vSwitch configuration.

After the ESX installation, configure your vSwitches on the ESX host with the following steps:

Procedure

Step 1 Log in to the command line of each ESX Server.

Step 2 Create three vSwitches on each ESX server using the listed names.

- **vswitch-hx-storage-data**

  Set the MTU to 9000 on this switch.

- **vmotion**

  Set the MTU to 9000 on this switch.

- **vswitch-hx-vm-network**

Step 3 Use the following CLI commands to create the three new vSwitches:

```
# esxcli network vswitch standard add -v vswitch-hx-storage-data
# esxcli network vswitch standard set -v vswitch-hx-storage-data -mtu= 9000
# esxcli network vswitch standard add -v vswitch-vmotion
# esxcli network vswitch standard set -v vswitch-vmotion -mtu=9000
# esxcli network vswitch standard add -v vswitch-hx-vm-network
```

Step 4 The default vSwitch **vSwitch0** created during installation of ESXi needs to be renamed to **vswitch-hx-inband-mgmt** for the Hx Data Platform node set up scripts to work properly. Use the following command to rename the switch and then reboot the host so that the vmkernel re-reads its configuration file to use the new name.

```
# sed -i 's/vSwitch0/vswitch-hx-inband-mgmt/g' /etc/vmware/esx.conf
# reboot
```

Step 5 You can verify the creation and renaming of the vSwitches after a host reboot with the following command:

```
# esxcli network vswitch standard list
```

Confirm that you see the four previously listed vSwitches in the command output. Only the switch-hx-inband-mgmt vSwitch will have Uplinks and Port groups listed. The HX Data Platform installer scripts perform the rest of the network configuration.

Migrating vMotion Networks to Virtual Distributed Switches (VDS) or Cisco Nexus 1000v (N1Kv)

Note

- The HX Data Platform can be configured with VMware DVS or Cisco Nexus 1000v for specific non-HX dependent networks:

  vMotion networks

  and virtual machine networks

- For further details, see Cisco Nexus 1000v documentation.

To migrate non-HX dependent vSwitches and associated port groups to DVS or N1Kv networks, follow the listed steps:

Procedure

Step 1 From vCenter, create DVS Switch and port groups.

a) Select **vCenter Inventory Lists** > **Datacenters** > **datacenter** > **Related Objects** > **Distributed Switches** . Click **Add Distributed Switch** icon.

b) Complete the New Distributed Switch wizard. Create each DVS switch with two uplinks.

For example: VM network and vmotion pg

- DVSwitch-VMNetwork: DVPortGroup-VMNetwork
- DVSwitch-Vmotion: DVPortGroup-Vmotion

Step 2 Migrate the vSwitch, VMNetwork. Perform the following steps to migrate VMNetwork from legacy vSwitch to DVS.

a) Select **vCenter Inventory Lists** > **Datacenters** > **datacenter** > **Related Objects** > **Distributed Switches**.

b) Select the **DVSwitch-VMNetwork** vSwitch. Click the **Add and Manage Hosts** icon. This starts the **Add and Manage Hosts** wizard.

c) On the Select task page, select **Add Hosts**. Click **Next**.

d) On the Select hosts page, click **Add New Hosts**. Select all hosts in the cluster. Click **Next**.

e) On the Select network adapter tasks page, select **Manage physical adapters** and **Migrate virtual machine** networking. Click **Next**.

f) On the Manage physical network adapters page, the physical adapters part of vswitch-hx-vm-network:VM Network are assigned to the DVSwitch-VMNetwork.

g) Under the **On other switches/unclaimed list**, select the vmnic corresponding to the **In Use by Switch**, vswitch-hx-vm-network.

h) Click **Assign** uplink.

i) Select Auto-assign.

j) Click **OK**. The page refreshes with the newly assigned vmnic listed under **On this switch**.

k) The **Analyze impact** page shows the impact of this migration. Verify the impact is all green. Click **Next**.

l) On the Migrate VM networking page, select the VMs to migrate to the new network, DVPortGroup-VMNetwork. Next

Select all the VMs, except the controller VMs, stCtlVM, from all the hosts. Select the DVPortGroup-VMNetwork. Click **Next**.

Note
The list of VMs for each host includes all the VMs, including the controller VMs. DO NOT select any controller VMs. Migrating the controller VMs will break your storage cluster.

m) On the Ready to complete page, confirm the summary of the migration. Click Finish.

Note
Post migration system generates several network related alarms. Verify and clear the alarms.

Step 3 Migrate the vSwitch to vmotion pg. Perform the following steps to migrate vmotion pg from legacy vSwitch to DVS.

a) Select **vCenter Inventory Lists** > **Datacenters > datacenter > Related Objects > Distributed Switches**.
b) Select the DVSwitch-Vmotion vSwitch. Click the **Add and Manage Hosts** icon. This starts the **Add and Manage Hosts** wizard.
c) On the Select task page, select **Add Hosts**. Click **Next**.
d) On the Select hosts page, click **Add New Hosts**. Select all hosts in the cluster. Click **Next**.
e) On the Select network adapter tasks page, select the tasks Manage physical adapters and Manage VMkernel adapters. Click **Next**.
f) On the **Manage physical network adapters** page, the physical adapters part of vmotion:vmotion pg are assigned to the DVSwitch-Vmotion.

Under the **On other switches/unclaimed** list, select the vmnic corresponding to the In Use by Switch, vmotion. Click **Assign uplink**, select Auto-assign, and click OK. The page refreshes with the newly assigned vmnic listed under **On this switch**. Click **Next**.

g) On the **Manage VMkernel network adapters** page, migrate the VMkernel adapter to the port group, DVPortGroup-Vmotion.

For each host, under the **On other switches** list, select the VMKernel adapter corresponding to the **In Use by Switch**, vmotion. Click **Assign port group**. Select the destination port group, DVPortGroup-Vmotion. Click **OK**. The page refreshes with the Reassigned VMkernel network adapters, listing the Source Port Group and Destination Port Group.

h) Select the hosts to migrate to the new network, DVPortGroup-Vmotion. Click **Next**.
i) On the Ready to complete page, confirm the summary of the migration, click **Finish**.

Step 4 Post migration step. Verify there is no impact on the VMs with respect to IO, Network connectivity and VM Migration.