



Cisco HyperFlex Data Platform Administration Guide for Hyper-V, Release 4.0

First Published: 2020-02-10

Last Modified: 2023-03-26

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 NON-INFRINGEMENT 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.

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.

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 Cisco Systems, Inc. All rights reserved.



CONTENTS

[Full Cisco Trademarks with Software License](#) ?

PREFACE

[Communications, Services, Bias-free Language, and Additional Information](#) vii

CHAPTER 1

[New and Changed Information for this Release](#) 1

[New and Changed Information](#) 1

CHAPTER 2

[Cisco HyperFlex Storage Cluster Overview](#) 3

[Cisco HX Data Platform Overview](#) 3

[Storage Cluster Physical Components Overview](#) 4

[Cisco HX Data Platform Capacity Overview](#) 5

[Understanding Capacity Savings](#) 6

[Storage Capacity Event Messages](#) 7

[Cisco HX Data Platform High Availability Overview](#) 8

[Storage Cluster Status](#) 8

[Operational Status Values](#) 9

[Resiliency Status Values](#) 9

[Cisco HX Data Platform Cluster Tolerated Failures](#) 10

[Data Replication Factor Settings](#) 11

[Cluster Access Policy](#) 12

[Responses to Storage Cluster Node Failures](#) 12

[Cisco HX Data Platform ReadyClones Overview](#) 14

[Creating ReadyClone VMs](#) 14

[Configuring Live Migration](#) 17

[Cisco HX Data Platform Hyper-V Checkpoints](#) 18

| | | |
|------------------|---|-----------|
| CHAPTER 3 | Logging in to Cisco HX Data Center Interfaces | 19 |
| | Cisco HyperFlex Cluster Interfaces Overview | 19 |
| | Guidelines for Cisco HX Data Platform Login Credentials | 20 |
| | Cisco HX Data Platform Names, Passwords, and Characters | 21 |
| | Logging into Cisco HX Connect | 23 |
| | Logging into the Controller VM (hxcli) Command Line | 24 |
| | Changing Storage Controller Password | 26 |
| | Logging Into Cisco HX Data Platform Installer | 26 |
| | Accessing the HX Data Platform REST APIs | 27 |

| | | |
|------------------|---|-----------|
| CHAPTER 4 | Monitoring Cisco HX Storage Clusters | 29 |
| | Monitoring HyperFlex Clusters | 29 |
| | Monitoring HyperFlex Clusters with HX Connect | 29 |
| | Dashboard Page | 30 |
| | Activity Page | 31 |
| | System Information Overview Page | 32 |
| | Nodes Page | 34 |
| | Disks Page | 35 |

| | | |
|------------------|---|-----------|
| CHAPTER 5 | Preparing for HX Storage Cluster Maintenance | 37 |
| | Storage Cluster Maintenance Operations Overview | 37 |
| | Serial vs. Parallel Operations | 38 |
| | Automating Updates Using Cluster-Aware Updating (CAU) | 39 |
| | Checking Cluster Status | 41 |
| | Setting a Beacon | 41 |
| | Verify Live Migration Configuration for HX Cluster | 42 |
| | Maintenance Modes for Storage Cluster Nodes | 42 |
| | Entering Cisco HyperFlex Maintenance Mode | 43 |
| | Exiting Cisco HyperFlex Maintenance Mode | 44 |
| | Creating a Backup Operation | 44 |
| | Shut Down and Power Off the Cisco HX Storage Cluster | 49 |
| | Power On and Start Up the Cisco HX Storage Cluster | 51 |
| | Restoring the Configuration for a Fabric Interconnect | 52 |

Recommendations for Verification after Storage Outages 54

Replacing a Compute Node 55

CHAPTER 6**Managing Users 57**

Managing Cisco HyperFlex Users Overview 57

User Management Terms 58

Audit Logs for AAA Accounting 59

Creating RBAC Users for Cisco HX Data Platform 59

Assigning Users Privileges 59

CHAPTER 7**Data Protection 61**

Hyper-V Checkpoints 61

Partner Solutions 62



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.



CHAPTER 1

New and Changed Information for this Release

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

New and Changed Information

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.

Table 1: New Features and Changed Features in Cisco HyperFlex Administration Guide for Hyper-V

| Feature | Description | Release or date added | Where Documented |
|--|---------------------------|-----------------------|------------------|
| Cisco HyperFlex Data Platform Administration Guide for Hyper-V | The guide was introduced. | January 15, 2020 | n/a |



CHAPTER 2

Cisco HyperFlex Storage Cluster Overview

- [Cisco HX Data Platform Overview, on page 3](#)
- [Storage Cluster Physical Components Overview, on page 4](#)
- [Cisco HX Data Platform Capacity Overview, on page 5](#)
- [Cisco HX Data Platform High Availability Overview, on page 8](#)
- [Storage Cluster Status, on page 8](#)
- [Cisco HX Data Platform Cluster Tolerated Failures, on page 10](#)
- [Responses to Storage Cluster Node Failures, on page 12](#)
- [Cisco HX Data Platform ReadyClones Overview, on page 14](#)
- [Creating ReadyClone VMs, on page 14](#)
- [Configuring Live Migration, on page 17](#)
- [Cisco HX Data Platform Hyper-V Checkpoints, on page 18](#)

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; Microsoft Hyper-V, Microsoft Windows servers with Hyper-V, Hyper-V Manager, Failover Cluster Manager, System Center Virtual Machine Manager (SCVMM) - (optional); and the Cisco HX Data Platform Installer, controller VMs, HX Connect, Powershell and hxcli commands.

Cisco HX Data Platform is installed on a virtualized platform such as Microsoft Hyper-V. 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 Cisco 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 Cisco 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 Cisco HX Data Platform Installer to add converged nodes to your storage cluster.

- **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.

Run the *Cluster Expansion* workflow from the Cisco HX Data Platform Installer to add compute nodes to your storage cluster.

- **Drives**—There are two types of drives that are a minimum 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 Cisco 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 Cisco 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.

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

Datastores are logical containers that are used by the Cisco 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.



Note Modifying permissions on HX Datastores is not supported on Hyper-V.

Cisco 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 Cisco HX Data Platform the concept of capacity is applied to both datastores and storage clusters. Values are measured in base-2 (GB/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.

Cisco HX Data Platform internal meta-data uses 0.5% to 1% space. This might cause the Cisco HX Data Platform Plug-in or Cisco 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 amount 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.

Error messages are issued if your data storage needs consume high amounts of available capacity, the performance and health of your storage cluster are affected. The error messages are displayed in, Cisco HX Connect, and ,TBD



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 affected.

Reduce the amount of storage capacity used to below the warning threshold, of 70% total HX Storage Cluster capacity.

- **SpaceAlertEvent** – Issues and 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 and 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, of 92% total HX Storage Cluster capacity.

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

Cluster storage space usage is back to normal.

Cisco HX Data Platform High Availability Overview

The Cisco 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.

Storage Cluster Status

Cisco 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, on page 9](#)
- [Resiliency Status Values, on page 9](#)

Operational Status Values

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

The Operational Status options are:

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

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

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

Resiliency Status Values

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

Resiliency Status options are:

- **Healthy**—The cluster is healthy with respect to data and availability.
- **Warning**—Either the data or the cluster availability is being adversely affected.
- **Unknown**—This is a transitional state while the cluster is coming online.

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

Cisco HX Data Platform Cluster Tolerated Failures

If nodes or disks in the Cisco 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 Cisco HX Data Platform installation and cannot be changed. The options are 2 or 3 redundant replicas of your data across the storage cluster.



Attention Data Replication Factor of 3 is recommended.

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

Cluster State with Number of Failed Nodes

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

Cluster State in 5+ Node Cluster with Number of Failed Nodes

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

Cluster State in 3- and 4-Node Clusters with Number of Failed Nodes

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

Cluster State with Number of Nodes with Failed Disks

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

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

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



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

3+ Node Cluster with Number of Nodes with Failed Disks

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

Data Replication Factor Settings



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

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

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

Choose a Data Replication Factor. The choices are:

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

Attention Data Replication Factor 3 is the recommended option.

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

Cluster Access Policy

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

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

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

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

Responses to Storage Cluster Node Failures

A storage cluster healing timeout is the length of time Cisco HX Connect or Cisco 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 Cisco 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 Cisco HX Connect and Cisco HX Data Platform Plug-in, to display reason messages that explain what is contributing to the current state.

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

| Cluster Size | Number of Simultaneous Failures | Entity Failed | Maintenance Action to Take |
|--------------|---------------------------------|---|--|
| 3 nodes | 2 | Two or more disks on two nodes are blacklisted or failed. | <ol style="list-style-type: none"> 1. If one SSD fails, the storage cluster does not automatically rebalance. Replace the faulty SSD and restore the system by rebalancing the cluster. 2. If one HDD fails or is removed, the disk is blacklisted and the storage cluster automatically begins healing with the remaining disks. 3. If more than one HDD fails, the system might not restore the storage cluster health. If the system is not restored, replace the faulty disks and rebalance the cluster. |
| 4 nodes | 1 | One node. | <p>If the node does not recover in two hours, the storage cluster automatically rebalances data on the remaining nodes.</p> <p>To recover the failed node immediately and fully restore the storage cluster health:</p> <ol style="list-style-type: none"> 1. Check that the node is powered on and restart it if necessary. If the node does not recover, you may need to replace the node. 2. Rebalance the cluster. |
| 4 nodes | 2 | Two or more disks on two nodes. | <p>If two SSDs fail, the storage cluster does not automatically rebalance. If the disk does not recover in one minute, the storage cluster automatically rebalances data on the remaining nodes.</p> |
| 5+ nodes | 2 | Up to two nodes. | <p>If the node does not recover in two hours, the storage cluster automatically rebalances data on the remaining nodes.</p> <p>To recover the failed node immediately and fully restore the storage cluster health:</p> <ol style="list-style-type: none"> 1. Check that the node is powered on and restart it if necessary. If the node does not recover, you may need to replace the node. 2. Rebalance the cluster. <p>If the storage cluster shuts down, see Troubleshooting Storage Cluster Health.</p> |
| 5+ nodes | 2 | Two nodes with two or more disk failures on each node. | The system automatically triggers a rebalance after a maintenance action to restore cluster health. |

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

Review the table above and perform the action listed.

Cisco HX Data Platform ReadyClones Overview

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

A ReadyClone, 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 ReadyClone is a separate guest VM.

Changes made to a ReadyClone do not affect the host VM. A ReadyClone'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 ReadyClone, 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.

Creating ReadyClone VMs

You can create Cisco HyperFlex Data Platform ReadyClones in a Hyper-V environment using a powershell script that is available for download from the Cisco CCO web site. The ReadyClone script automates the VM cloning process which involves exporting the original VM to a temporary folder, importing, and then registering the saved VM to a new location. After the successful creation of ReadyClone VMs, the exported temp folder is deleted automatically. The VM is later added to the cluster if that option is chosen.



Note The VM in the below example is generation 2 Windows Server 2016.

SUMMARY STEPS

1. Download the Cisco HyperFlex Data Platform Hyper-V ReadyClone powershell script from the [Cisco CCO Software Download page for HyperFlex HX Data Platform Release 4.0\(1b\)](#).
2. Run the following command:
3. A new VM created with ReadyClone will now be in the saved state. Use the Failover Cluster Manager, Hyper-V Manager or SCVMM to turn it on.

DETAILED STEPS

Step 1 Download the Cisco HyperFlex Data Platform Hyper-V ReadyClone powershell script from the [Cisco CCO Software Download page for HyperFlex HX Data Platform Release 4.0\(1b\)](#).

Step 2 Run the following command:

```
HxClone-HyperV-v4.0.1b-33133.ps1 -VmName <VM Name> -ClonePrefix <Prefix> -CloneCount <number> -AddToCluster <$false/$true>
```

```
PS C:\Users\administrator.HXHVDOM2>
PS C:\Users\administrator.HXHVDOM2> C:\HxClone-HyperV-v4.0.1b-33133.ps1 -VmName RCVMI -ClonePrefix c14 -CloneCount 1 -AddToCluster $true

Directory: \\hxbv2smb.hxbvdom2.local\hxd1\Hyper-V\Virtual Hard Disks

Mode                LastWriteTime         Length Name
----                -
d-----          9/11/2019   7:16 PM          tmp1417411279
\\hxbv2smb.hxbvdom2.local\hxd1\Hyper-V\Virtual Hard Disks\tmp1417411279\RCVMI\Virtual Machines\9b535cbb-c0a8-4b77-9142-284525fb3033.vmcx

Directory: \\hxbv2smb.hxbvdom2.local\hxd1

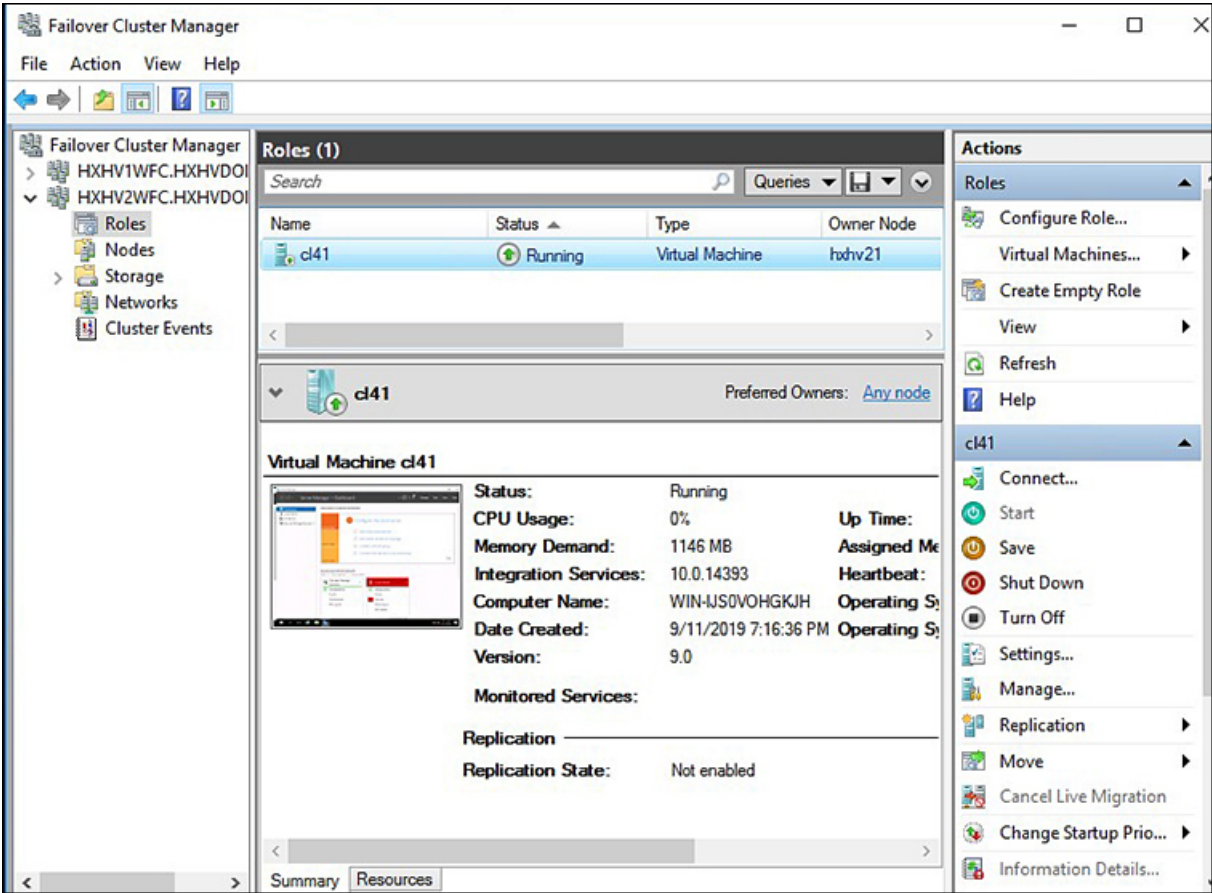
Mode                LastWriteTime         Length Name
----                -
d-----          9/11/2019   7:16 PM          c141
\\hxbv2smb.hxbvdom2.local\hxd1\c141

Name           : c141
OwnerNode      : hxbv21
State          : Offline

PS C:\Users\administrator.HXHVDOM2>
```

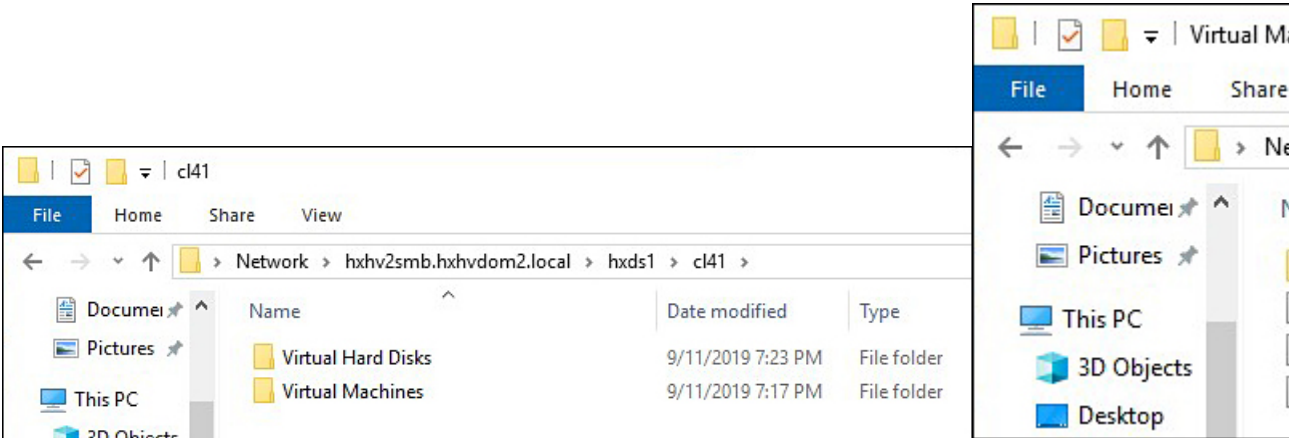
Step 3 A new VM created with ReadyClone will now be in the saved state. Use the Failover Cluster Manager, Hyper-V Manager or SCVMM to turn it on.

If the AddToCluster parameter is set to *\$true*, then the ReadyClone VMs are converted to highly available clustered roles which can be seen and managed from the Failover Cluster Manager. It will also be visible in the Hyper-V Manager.



A folder in the name of the guest VM (cl41 in this case) is created inside the HX datastore
 \\hxhv2smb.hxhvdom2.local\hxds1

This folder contains snapshots (if there are any available at the time of creating ReadyClones), Virtual Hard Disk and Virtual Machine files.



After the successful creation of ReadyClones, there is no further relationship with the original VM. During the creation of ReadyClones, the original VM is exported to a temporary folder location, and then from that

location, the VM is imported using the Copy the VM option to another location in the HX Datastore with new unique IDs for the restored VMs.

After deleting the ReadyClone VMs, the VM configuration files are deleted, but the folder structure and the virtual hard disk file remain. This may require manual clean up.

What to do next

ReadyClone powershell script parameters are documented in the following table.

Table 2: ReadyClone PowerShell Script Parameters

| Parameters | Value | Description |
|--------------|-----------------------|--|
| VmName | <Name value> | Enter the Name of the running VM used for creating ReadyClones. |
| ClonePrefix | <Prefix value> | Enter a prefix for the guest virtual machine name. This prefix is added to the name of each ReadyClone created. |
| CloneCount | <#> | Enter a value to create the number of ReadyClones. |
| AddToCluster | <\$false> <\$true> | <i>\$false</i> – creates standalone VMs (only visible in Hyper-V Manager) <i>\$true</i> – creates a highly available clustered ReadyClone VMs (visible in Failover Cluster Manager and Hyper-V Manager as well) |

Configuring Live Migration

Starting with HyperFlex 4.0(2a), the HX installer can configure Live Migration on Hyper-V cluster nodes, if the information is provided during install or expand workflows.



- Note** Additional steps may be necessary in some situations to automatically configure Live Migration during cluster expansion workflow using the HyperFlex 4.0(2a) installer. Check if the following conditions are true:
- Live Migration was not configured during a fresh cluster install workflow using the HyperFlex 4.0(2a) installer.
 - Cluster is upgraded to 4.0(2a)

In such cases, complete the following steps, then proceed to the cluster expansion workflow.

SUMMARY STEPS

1. Manually configure Live Migration IP addresses on all nodes.
2. Run `update-inventory.py` to sync your network configuration changes with HyperFlex.
3. Run the cluster expansion workflow and provide Live Migration information in the installer UI for the node(s) being expanded.

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | Manually configure Live Migration IP addresses on all nodes. | For more information, see Configuring a Static IP address for Live Migration and VM Network in the Cisco HyperFlex Systems Installation Guide for Microsoft Hyper-V, Release 4.0 . Note This is applicable only if you have not done so already using the HX Installer. |
| Step 2 | Run <code>update-inventory.py</code> to sync your network configuration changes with HyperFlex. | This file is located at <code>/usr/share/springpath/storfs-misc/update-inventory.py</code> on the cluster management IP node. This updates the HyperFlex inventory with Live Migration information on each Hyper-V node. Cluster expansion workflow will then show the corresponding Live Migration UI fields. |
| Step 3 | Run the cluster expansion workflow and provide Live Migration information in the installer UI for the node(s) being expanded. | Expansion should be aware that Live Migration is configured for the existing HX cluster and show the corresponding UI fields. |

Cisco HX Data Platform Hyper-V Checkpoints



Note Cisco HX Data Platform Native Snapshots are not supported in Hyper-V. Use Hyper-V Checkpoints.

Choose between standard or production checkpoints in Hyper-V.

Applies To: Windows Server 2016, Microsoft Hyper-V Server 2019

Starting with Windows Server 2016, you can choose between standard and production checkpoints for each virtual machine. Production checkpoints are the default for new virtual machines.

Production checkpoints are "point in time" images of a virtual machine, which can be restored later on in a way that is completely supported for all production workloads. This is achieved by using backup technology inside the guest to create the checkpoint, instead of using saved state technology.

Standard checkpoints capture the state, data, and hardware configuration of a running virtual machine and are intended for use in development and test scenarios. Standard checkpoints can be useful if you need to recreate a specific state or condition of a running virtual machine so that you can troubleshoot a problem.



CHAPTER 3

Logging in to Cisco HX Data Center Interfaces

- [Cisco HyperFlex Cluster Interfaces Overview, on page 19](#)
- [Guidelines for Cisco HX Data Platform Login Credentials, on page 20](#)
- [Cisco HX Data Platform Names, Passwords, and Characters, on page 21](#)
- [Logging into Cisco HX Connect, on page 23](#)
- [Logging into the Controller VM \(hxcli\) Command Line, on page 24](#)
- [Changing Storage Controller Password, on page 26](#)
- [Logging Into Cisco HX Data Platform Installer, on page 26](#)
- [Accessing the HX Data Platform REST APIs, on page 27](#)

Cisco HyperFlex Cluster Interfaces Overview

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

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

Additional interfaces include:

- Cisco HX Data Platform Installer—Installing HX Data Platform, deploying and expanding HX Storage Cluster, deploying stretched cluster, and deploying Hyper-V clusters.
- Cisco UCS Manager—Tasks for networking, storage and storage access, and managing resources in the HX Storage Cluster.
- Hyper-V Manager—Managing all the Hyper-V node and virtual machines
- Microsoft Failover Cluster Manager—Configure and management of failover cluster host, role and virtual machines

Guidelines for Cisco HX Data Platform Login Credentials

`hxcli` commands prompt for login credentials.

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

| Component | Permission Level | Username | Password | Notes |
|-------------------------------|--|---|---|--|
| HX Data Platform Installer VM | root | root | Cisco123 Note Systems ship with a default password of <code>Cisco123</code> that must be changed during installation. You cannot continue installation unless you specify a new user supplied password. | |
| HX Connect | administrator or read-only | Member of Active Directory "Desktop Users" group. | User defined password. | Read-only access. |
| | | Predefined <code>admin</code> or <code>root</code> users. | | Member of the Domain Admin group also have administrator access. |
| HX Storage Controller VM | admin | User defined during HX installation. Predefined <code>admin</code> or <code>root</code> users. | As specified during HX installation. Strong password required. | Must match across all nodes in storage cluster. Use the <code>hxcli</code> command when changing the password after installation. |
| Hypervisor | member of the "Domain Administrator" group | User defined password | As specified during HX installation. | - |
| UCS Manager | admin | As configured. | As configured. | - |

| Component | Permission Level | Username | Password | Notes |
|---------------------|------------------|----------------|----------------|-------|
| Fabric Interconnect | admin | As configured. | As configured. | - |

Cisco 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.

However, to simplify names and passwords, consider not using these special characters, as they are frequently assigned special purposes.

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!'

HX Storage Cluster Name—HX cluster names cannot exceed 50 characters.

HX Storage Cluster Host Names—HX cluster host names cannot exceed 15 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 15 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.
- Only special character allowed: underscore (`_`), dash (`-`), 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. Login to the management controller VM.
2. Run the `hxcli security password set [-h] [--user USER]` command.

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

UCS Manager Password Format and Character Requirements

The following is a summary of format and character requirements for UCS Manager passwords. See the Cisco UCS Manager 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.
 - **Excluded characters:**
UCS Manager passwords cannot contain the escape (\) character.

Logging into Cisco HX Connect

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

About Sessions

Each login to Cisco HX Connect is a session. Sessions are the period of activity between time when you log into Cisco 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:

- 256 concurrent sessions per user.
- 300 concurrent sessions across the Cisco 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 hypervisor and the controller VMs are in sync or near sync. If there is too large of a time skew between the hypervisor time and the cluster time, AAA authentication will fail.

-
- Step 1** Locate the Cisco 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 Cisco HX Storage Cluster management IP address in a browser.

Step 3 Enter the Cisco 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 Cisco HX Storage Cluster
 - **Read only**—Users with read only role have read (view) permissions. They cannot make any changes to the Cisco HX Storage Cluster.
- **HX pre-defined users**—To login using the Cisco 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.

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 Cisco 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 log out of Cisco HX Connect, and properly close the session, select **User** menu (top right) > **Logout**.

Logging into the Controller VM (hxcli) Command Line

All `hxcli` command are divided into commands that read Cisco HX Cluster information and commands that modify the Cisco 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 Cisco HX Data Platform `hxcli` commands, log in to the Cisco 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 in to the Cisco HX Data Platform command line interface in the Storage Controller VM in the following ways:

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

Only direct commands are supported through Cisco 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`

Step 1 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 2 From a command line terminal using `ssh`.

Note Do not include the password in an `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 root@10.198.3.24
The authenticity of host '10.198.3.24 (10.198.3.24)' can't be established.
ECDSA key fingerprint is xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx.
Are you sure you want to continue connecting (yes/no)?
```

c) Enter the password at the prompt.

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

Step 3 From HX Connect—Log in to Cisco HX Connect, select **Web CLI**.

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

Changing Storage Controller Password

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

- Step 1** Log in to a storage controller VM.
- Step 2** Change the Cisco HyperFlex storage controller password.

```
# hxcli security password set
```

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

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

To change the password on compute nodes:

- a. Live migration all the user VMs off the Hyper-V host.
- b. Launch the storage controller VM console from Hyper-V Manager and log in as the root user.
- c. Run the `passwd` command to change the password.
- d. Log out and re-login to confirm that the password changed successfully.
- e. Run the `hxcli node add -f` command to add the node back into the cluster.

- Step 3** Type in the new password.
- Step 4** Press **Enter**.

Logging Into Cisco HX Data Platform Installer

Next, you install the HX Data Platform software.



Note Before launching the Cisco HX Data Platform Installer, ensure that all the Hyper-V hosts that are in the cluster that you plan to include in the storage cluster are in maintenance mode.

- Step 1** In a browser, enter the URL for the VM where HX Data Platform Installer is installed.
- You must have this address from the earlier section on **Deploying HX Data Platform Installer**. For example `http://10.64.4.254`
- Step 2** Enter the following credentials:
- **Username:** `root`
 - **Password** (Default): `Cisco123`

Attention Systems ship with a default password of `Cisco123` that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.

Read the EULA. Click **I accept the terms and conditions**.

Verify the product version listed in the lower right corner is correct. Click **Login**.

Step 3 The HX Data Platform Installer Workflow page provides two options to navigate further.

- **Create Cluster** drop-down list—You can deploy a standard cluster, Stretched Cluster, or a Hyper-V cluster.
- **Cluster Expansion**—You can provide the data to add converged nodes and compute nodes to an existing standard storage cluster.

Accessing the HX Data Platform REST APIs

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

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

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

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

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

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

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



CHAPTER 4

Monitoring Cisco HX Storage Clusters

- [Monitoring HyperFlex Clusters, on page 29](#)
- [Monitoring HyperFlex Clusters with HX Connect, on page 29](#)

Monitoring HyperFlex Clusters

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

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

Monitoring HyperFlex Clusters with HX Connect

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

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 Guide for details.
- **Performance**—Charts for IOPS, throughput, latency, and replication network bandwidth.
- **System Information**—System overview, plus status and tasks for nodes and disks.

See the Cisco HyperFlex Systems Troubleshooting Guide for generating support bundles, [Storage Cluster Maintenance Operations Overview, on page 37](#) for entering and exiting maintenance mode, and [Setting a Beacon, on page 41](#) to set a node or disk beacon.

- **Datstores**—Status and tasks related to datastores.

The **Upgrade** page provides access to HX Data Platform 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 | <p>Provides the functional status of the HX storage cluster and application performance.</p> <p>Click Information () to access the HX storage cluster name and status data.</p> |
| Cluster License Status section | <p>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:</p> <p>Cluster License not registered link—Appears when the HX storage cluster is not registered. To register a cluster license, click this link and provide product instance registration token in the Smart Software Licensing Product Registration screen. For more information on how to get a product instance registration token, refer the Registering a Cluster with Smart Licensing section in the Cisco HyperFlex Systems Installation Guide for Microsoft Hyper-V.</p> |
| Resiliency Health section | <p>Provides the data health status and ability of the HX storage cluster to tolerate failures.</p> <p>Click Information () to access the resiliency status, and replication and failure data.</p> |
| Capacity section | <p>Displays a breakdown of the total storage versus how much storage is used or free.</p> <p>Also displays the storage optimization, compression-savings, and deduplication percentages based on the data stored in the cluster.</p> |
| Nodes section | <p>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.</p> |
| Performance section | <p>Displays an HX storage cluster performance snapshot for a configurable amount of time, showing IOPS, throughput, and latency data.</p> <p>For full details, see Performance Page.</p> |
| 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 | <p>The table automatically refreshes for dynamic updates to the HX Cluster. The timestamp indicates the last time the table was refreshed.</p> <p>Click the circular icon to refresh the content now.</p> |
| Filter field | <p>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.</p> <p>Type in the selection text in the Filter field.</p> <p>To empty the Filter field, click the x.</p> <p>To export content from other pages in the table, scroll to the bottom, click through the page numbers, and apply the filter.</p> |
| Export menu | <p>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.</p> <p>Click the down arrow to select an export file type. The file type options are: <code>cvs</code>, <code>xls</code>, and <code>doc</code>.</p> <p>To export content from other pages in the table, scroll to the bottom, click through the page numbers, and apply the export.</p> |

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, and enter/exit maintenance mode.

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

System Information Overview Page

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

HX Storage Cluster Configuration Data

Displays the basic configuration information for this HX storage cluster.

| UI Element | Essential Information |
|---------------------------------|--|
| HX storage cluster field | Name of this storage cluster. |
| HX storage cluster status field | <p>Provides functional status of the HX storage cluster.</p> <ul style="list-style-type: none"> • Online—Cluster is ready. • Offline—Cluster is not ready. • Read Only—Cluster is out of space. • Unknown—Transitional state while the cluster is coming online. |

| UI Element | Essential Information |
|---------------------------------|---|
| Hypervisor field | Hypervisor version installed on this HX storage cluster. |
| HXDP Version field | Installer package version installed on this HX storage cluster. |
| Uptime field | Length of time this HX storage cluster has been online. |
| Total Capacity field | Overall storage size of this cluster. |
| Available Capacity field | Amount of free storage in this cluster. |
| DNS Server(s) | IP address for the DNS server(s) for this HX storage cluster. |
| NTP Server(s) | IP address for the NTP server(s) for this HX storage cluster. |

Controller VM Access

You can access the controller VM using SSH as an administrator. To enable access, click **Actions** at the top of the page to enable SSH access.

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 field | Name of a node on this cluster. |
| Hypervisor Address field | IP address for the management network to this HX storage cluster. |
| Hypervisor Status field | <ul style="list-style-type: none"> • Online—Node is available. • Offline—Node is not available. • In Maintenance—The running (and powered off) node is disconnected from the host. • In Progress—a backup job is in progress. |
| Controller address | IP address for the network controller for this HX storage cluster. |
| Controller Status | <ul style="list-style-type: none"> • Online—The connection between the VM and the disk is available. • Offline—The connection between the VM and the disk is not available. • In Maintenance—the connection between the VM and the disk is powered off from the host. • In Progress—a backup job is in progress. |
| Model field | Physical hardware model number of this node. |
| HXDP Version field | Installer package version installed on this node. |

| UI Element | Essential Information |
|--------------------|--|
| Disks field | Number of persistent and caching disks in this node. |

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

Nodes Page

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

| UI Element | Essential Information |
|---|---|
| Enter HX Maintenance Mode button | <p>Select a node to access this button.</p> <p>Opens the Confirm HX Maintenance Mode dialog box.</p> |

| UI Element | Essential Information |
|--|---|
| Exit HX Maintenance Mode button | Select a node to access this button. After you complete any maintenance tasks, you must manually exit HX 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. |
| Hypervisor Status column | <ul style="list-style-type: none"> • Online—Node is available. • Offline—Node is not available. • In Maintenance—The running (and powered off) node is disconnected from the host. • In Progress—a backup job is in progress. |
| Controller Address column | IP address for the HX storage controller VM of the Node referred in the Node column. |
| Controller Status column | <ul style="list-style-type: none"> • Online—The connection between the VM and the disk is available. • Offline—The connection between the VM and the disk is not available. • In Maintenance—the connection between the VM and the disk is powered off from the host. |
| Model column | Physical hardware model number of this node. |
| Version column | HyperFlex Data Platform installer package version installed on this node. |
| Disks column | 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 | <ul style="list-style-type: none"> • 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. • Blacklisted—State when a disk is not being consumed by the cluster due to either a software error or an IO error. This could be a transitional state while the cluster attempts to repair the disk, if the disk is still available, before the state transitions to Repairing. • Claimed—State when a disk is recognized and in use. • 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. • Repairing—State when a blacklisted disk is currently being repaired. • To Be Removed—State when a disk is scheduled for RMA. | <p>The following states can be ignored:</p> <ul style="list-style-type: none"> • Invalid • Normal • Removed—State when an SED disk is removed after using the Secure Erase option. • Time out • Unknown |
| Type column | <ul style="list-style-type: none"> • Unknown • Rotational—Hybrid drive • Solid State—SSD drive | |
| Usage column | <ul style="list-style-type: none"> • Unknown • Cache • Persistent | |
| Turn On Locator LED and Turn Off Locator LED radio buttons | <p>Select a disk to access the radio buttons.</p> <p>Activates or deactivates a physical light, or beacon, on the host to help locate the disk.</p> | |



CHAPTER 5

Preparing for HX Storage Cluster Maintenance

- [Storage Cluster Maintenance Operations Overview, on page 37](#)
- [Serial vs. Parallel Operations, on page 38](#)
- [Automating Updates Using Cluster-Aware Updating \(CAU\), on page 39](#)
- [Checking Cluster Status, on page 41](#)
- [Setting a Beacon, on page 41](#)
- [Verify Live Migration Configuration for HX Cluster, on page 42](#)
- [Maintenance Modes for Storage Cluster Nodes, on page 42](#)
- [Entering Cisco HyperFlex Maintenance Mode, on page 43](#)
- [Exiting Cisco HyperFlex Maintenance Mode, on page 44](#)
- [Creating a Backup Operation, on page 44](#)
- [Shut Down and Power Off the Cisco HX Storage Cluster, on page 49](#)
- [Power On and Start Up the Cisco HX Storage Cluster, on page 51](#)
- [Restoring the Configuration for a Fabric Interconnect, on page 52](#)
- [Recommendations for Verification after Storage Outages, on page 54](#)
- [Replacing a Compute Node, on page 55](#)

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 three 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 `hxcli` command line, HX Connect, and HX Data Platform Plug-in. The `hxcli 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 38](#) for more information on these tasks.. You can perform only some tasks serially to each other.
- Put the Hyper-V host into HX Maintenance Mode prior to performing a maintenance task on the host. The HX maintenance mode performs additional storage cluster specific steps compared to the Hyper-V host provided Hyper-V maintenance mode.

Post Maintenance Tasks

After the maintenance task is completed, the nodes need to exit Cisco HX Maintenance Mode and the storage cluster needs to be restarted. In addition, some changes to the Cisco HX storage cluster require additional post maintenance tasks. For example, if you change the vNICs or vHBAs, the PCI Passthrough needs to be reconfigured.

Ensure the following:

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

Automating Updates Using Cluster-Aware Updating (CAU)

Cisco HyperFlex 4.0(2a) supports Cluster-Aware Updating (CAU), a feature on Windows systems that automates the software updating process on clustered servers. CAU enables you to update servers in a failover cluster with little or no loss in availability during the update process. During an updating run, CAU transparently performs the following tasks:

1. Puts each node of the cluster into maintenance mode.
2. Moves the clustered roles off the node.
3. Installs the updates and any dependent updates.
4. Performs a restart if necessary.
5. Brings the node out of maintenance mode.
6. Restores the clustered roles on the node.
7. Moves to update the next node.

For more information, see [Cluster-Aware Updating](#).



Note HyperFlex CAU integration does not use HyperFlex Maintenance Mode. For highly sensitive workloads, alternate patching methods which place the node into HyperFlex Maintenance Mode beforehand may need to be considered.

To use CAU, you must first configure a CAU profile.

Before you begin

Locate and run the Cluster-Aware Updating (CAU) script (called `CAU_worker.ps1`) on all nodes and verify that the cluster is online and healthy (optional).



Note If you have entered an IP address for the CIP-M field, the CAU feature is not supported. This value has to be a name, and must have a DNS entry for it.

SUMMARY STEPS

1. Create a prestaged computer account and provide full control permissions to the failover cluster object.
2. Open the Cluster-Aware Updating tool and connect to the failover cluster. From the list of cluster nodes, select the failover cluster, and then click **Connect**.
3. Configure the Cluster-Aware Updating (CAU) profile. From the **Cluster Actions** menu, select **Configure cluster self-updating options**. The Configure Self-Updating Options wizard appears.
4. Add the Clustered Role.
5. Click **Apply**. The **Add Clustered Role** indicates **Success** when done.

DETAILED STEPS

Step 1 Create a prestaged computer account and provide full control permissions to the failover cluster object.

Note When you create a failover cluster, you must specify a name for the cluster. If you have sufficient permissions when you create the cluster, the cluster creation process automatically creates a computer object in AD DS that matches the cluster name. This object is called the cluster name object or CNO. Through the CNO, virtual computer objects (VCOs) are automatically created when you configure clustered roles that use client access points. To create the CNO automatically, the user who creates the failover cluster must have the Create Computer objects permission to the organizational unit (OU) or the container where the servers that will form the cluster reside. For more information, see [Prestage cluster computer objects in Active Directory Domain Services](#).

- a) The HyperFlex installer already creates a cluster name object (CNO) in Active Directory. The CNO shares the same name as the Windows failover cluster. Note down the name of the CNO.
- b) Create a new computer object in Active Directory. This is called the virtual computer object (VCO).
- c) Right click on the VCO. Go to **Properties>Security->Add**. Provide the CNO name and give full control permissions to it.

Step 2 Open the Cluster-Aware Updating tool and connect to the failover cluster. From the list of cluster nodes, select the failover cluster, and then click **Connect**.

Step 3 Configure the Cluster-Aware Updating (CAU) profile. From the **Cluster Actions** menu, select **Configure cluster self-updating options**. The Configure Self-Updating Options wizard appears.

Step 4 Add the Clustered Role.

- a) From the **Add Clustered Role with Self-Updating Enabled** window, click on the checkbox to **Add the CAU clustered role with self-updating mode enabled to this cluster** if you want to run the updates in self-updating mode. Do not click the checkbox if you want to run the cluster updating operation in Remote updating mode.

Note If you are running Windows Core or Windows Desktop Experience on the hypervisor node, you must coordinate the cluster updating operation in Remote-updating mode. For this mode, a remote computer, which is called an Update Coordinator is configured with the CAU tools. The Update Coordinator is not a member of the cluster that is updated during the Updating Run. From the remote computer, the administrator triggers an on-demand Updating Run by using a default or custom Updating Run profile.

- b) Click on the **I have a prestaged computer object for the CAU clustered role** checkbox. Provide the VCO name in the wizard. Click **Next**.
- c) Specify the schedule by selecting the frequency of self-updating (Daily, Weekly, Monthly), Starting date, and Time of day. Click **Next**.
- d) Configure Advanced Options to set the maximum retries per node, require all nodes online, and location of the pre-update script as follows:

- MaxRetriesPerNode = 3

- RequireAllNodesOnline = True
- PreUpdateScript = c:\ProgramData\Cisco\HyperFlex\Tools\CAU\CAU_preupdate.ps1

e) From the Additional Update Options window, click on the checkbox to **Give me recommended updates the same way that I receive important updates**. Click **Next**.

Step 5 Click **Apply**. The **Add Clustered Role** indicates **Success** when done.

The Cluster-Aware Updating (CAU) process runs as configured. You can also start the update process manually by clicking **Apply Updates to this cluster** from the **Cluster Actions** menu in the CAU tool. View progress of each run in the "Log of Updates in Progress" window.

If the updating run fails, you can view the latest log file to troubleshoot the problem. The CAU log files are located in the same folder containing the CAU update scripts (i.e., c:\ProgramData\Cisco\HyperFlex\Tools\CAU).

Checking Cluster Status

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

Step 2 Verify the storage cluster is healthy.

```
# 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.

```
# hxcli cluster storage-summary
```

Example response:

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

Setting a Beacon

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

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

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

- 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 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 Live Migration 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 Live Migration. Confirm the following from your Failover Cluster Manager:

1. Verify that the Live Migration network is Up in the Networks tab.
2. Configure the Live Migration network in the Live Migration Settings located in the Actions panel.
3. Verify that you have assigned a static IP to each Live Migration NIC team, and that the static IPs for each Live Migration port group are in the same subnet.

Maintenance Modes for Storage Cluster Nodes

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

There are two types of maintenance modes.

- Cisco HX maintenance mode
- Hyper-V maintenance mode

Cisco HX Maintenance Mode

Cisco HX maintenance mode performs Cisco HX Data Platform specific functions in addition to the Hyper-V maintenance mode. Be sure to select Cisco HX maintenance mode and not Hyper-V 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 Windows updates.

Cisco HX Maintenance Mode Considerations

- When Cisco HX Maintenance Mode is entered to enable performing tasks on an Hyper-V host, be sure to exit Cisco HX Maintenance Mode after the tasks on the Hyper-V host are completed.
- Cisco HX Maintenance Mode is applied to nodes in a healthy cluster only. If the cluster is unhealthy, for example too many nodes are down, or you are shutting down the cluster, use Hyper-V Maintenance Mode.
- See [Entering Cisco HyperFlex Maintenance Mode, on page 43](#) and [Exiting Cisco HyperFlex Maintenance Mode, on page 44](#) for steps.

Hyper-V Maintenance Mode

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

To enter or exit Hyper-V maintenance mode:

- Through the System Center Virtual Machine Manager (SCVMM) select the *host*, then from the right-click menu select **Start Maintenance Mode**.

Entering Cisco HyperFlex Maintenance Mode

Using the Cisco HyperFlex (HX) Connect User Interface



Note Maintenance Mode is supported on Cisco HyperFlex Release 2.5(1a)/2.5(1b) and later.

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



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

Using the Command-Line Interface

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

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

- b. Enter the node into HX Maintenance Mode.

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

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

Exiting Cisco HyperFlex Maintenance Mode

Using the Cisco HyperFlex (HX) Connect User Interface



Note Maintenance Mode is supported on Cisco HyperFlex Release 2.5(1a)/2.5(1b) and later.

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

Using the Command-Line Interface

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

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

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

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

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

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.

- 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: <ul style="list-style-type: none"> • Enabled—Cisco UCS Manager runs the backup operation as soon as you click OK. • Disabled—Cisco UCS Manager does not run the backup operation when you click OK. If you select this option, all fields in the dialog box remain visible. However, you must manually run the backup from the Backup Configuration dialog box. |

| Name | Description |
|------------|---|
| Type field | <p>The information saved in the backup configuration file. This can be one of the following:</p> <ul style="list-style-type: none"> • Full state—A binary file that includes a snapshot of the entire system. You can use the file generated from this backup to restore the system during disaster recovery. This file can restore or rebuild the configuration on the original fabric interconnect, or recreate the configuration on a different fabric interconnect. You cannot use this file for an import. <p>Note You can only use a full state backup file to restore a system that is running the same version as the system from which the backup file was exported.</p> <ul style="list-style-type: none"> • All configuration—An XML file that includes all system and logical configuration settings. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore. This file does not include passwords for locally authenticated users. • System configuration—An XML file that includes all system configuration settings such as usernames, roles, and locales. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore. • Logical configuration—An XML file that includes all logical configuration settings such as service profiles, VLANs, VSANs, pools, and policies. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore. |

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

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

Step 7 Click **OK**.

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

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

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

Shut Down and Power Off the Cisco HX Storage Cluster

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

- A **powered-off cluster** has all the physical components of the storage cluster removed from electrical power.

Very rarely would a storage cluster need to have all the components powered off. No regular maintenance or upgrade processes require that the entire storage cluster be completely powered off.

- A **shut-down cluster** has all storage cluster processes, including the working VMs, powered down. This does not include powering down the nodes in the cluster or shutting down the Hypervisor or FI cluster.
- An **offline cluster** is one of the storage cluster operational states. A storage cluster can be offline if there is an unknown or specific error, or if the storage cluster has been shutdown.

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

Before you begin

- The storage cluster must be in a healthy state.
- Perform both the Full-State and All Configuration type backups with the Preserve Identities attribute. See [Creating a Backup Operation, on page 44](#).

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

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

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

**Step 4** On each storage cluster Hyper-V host, shutdown the controller VM (`hxctlvm`).

Using Hyper-V Manager VM Power Off

- a) From Hyper-V Manager, locate the controller VM on each Hyper-V host.
- b) Right-click the controller VM and select **Shutdown**.

This method performs a graceful guest VM shutdown.

**Step 5** Shutdown each storage cluster Hyper-V host.

- a) Log into the Hyper-V and click Power from the **Start** button.
- b) Click Shut Down from the **Power** menu.

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

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

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

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

---

The HX storage cluster is now safely powered off.

### What to do next

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

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

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




---

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

---

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

## 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 49.

#### Step 1

Plug in to power up the FIs.

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

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

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

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

#### Step 2

Connect all the Hyper-V 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 Hyper-V. If any node does not, then connect to the UCS Manager and power up the servers (nodes) from UCS Manager.

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

#### Step 3

Verify all the Hyper-V 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 (`hxctlvm`) perform the following steps:

Using Hyper-V host command line

- a) Login to a host.
- b) Identify the VMID of the `hxctlvm`.  

```
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.
- SSH to any controller VM, run the command:  

```
hxcli about
```
  - If the command returns full storage cluster information, including build number, the storage cluster is ready to be started. Proceed to restarting the storage cluster.
  - If the command does not return full storage cluster information, wait until all the services have started on the host.
- Step 8** Start the storage cluster.  
From the command line of any controller VM, run the command.  

```
hxcli cluster start
```

  
Depending upon the maintenance or upgrade task performed while the HX cluster was shutdown, the nodes might be exited from HX maintenance mode or Hyper-V maintenance mode. Ignore any error messages about an unknown host exception.
- Step 9** Wait until the storage cluster is online and returns to a healthy state.
- From any controller VM, run the command.  

```
hxcli cluster info
```
  - In the command response text, check the cluster subsection and verify the `healthstate` is `online`.  
This could take up to 30 minutes, it could take less time depending upon the last known state.
- Step 10** When the storage cluster is healthy and the datastores are remounted, power on the workload VMs.
- 

## Restoring the Configuration for a Fabric Interconnect

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

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

### Before you begin

Collect the following information to restore the system configuration:

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

---

## SUMMARY STEPS

1. Connect to the console port.
2. If the fabric interconnect is off, power on the fabric interconnect.
3. At the installation method prompt, enter **gui**.
4. If the system cannot access a DHCP server, you may be prompted to enter the following information:
5. Copy the web link from the prompt into a web browser and go to the Cisco UCS Manager GUI launch page.
6. On the launch page, select **Express Setup**.
7. On the **Express Setup** page, select **Restore From Backup** and click **Submit**.
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**
9. In the **Server Information** area, complete the following fields:
10. Click **Submit**.

## DETAILED STEPS

- 
- Step 1** Connect to the console port.
- Step 2** If the fabric interconnect is off, power on the fabric interconnect.  
You will see the power on self-test message as the fabric interconnect boots.
- Step 3** At the installation method prompt, enter **gui**.
- Step 4** If the system cannot access a DHCP server, you may be prompted to enter the following information:
- IPv4 address for the management port on the fabric interconnect
  - Subnet mask or prefix for the management port on the fabric interconnect
  - IPv4 address for the default gateway assigned to the fabric interconnect
- Step 5** Copy the web link from the prompt into a web browser and go to the Cisco UCS Manager GUI launch page.
- Step 6** On the launch page, select **Express Setup**.
- Step 7** On the **Express Setup** page, select **Restore From Backup** and click **Submit**.
- Step 8** In the **Protocol** area of the **Cisco UCS Manager Initial Setup** page, select the protocol you want to use to upload the full state backup file:

- SCP
- TFTP
- FTP
- SFTP

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

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

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

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

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

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

## Recommendations for Verification after Storage Outages

After a power (or storage) outage, you may see the same virtual machine (VM) appearing/registered on two Hyper-V nodes. To recover from this situation, proceed as follows.

### Before you begin

Confirm that the virtual machine (VM) is in running state on one host and is powered off on the other host.

### SUMMARY STEPS

1. Use Hyper-V Manager to power off the VM from the host on which it is running.

2. Copy all files of the VM to a location to serve as backup.
3. Use Hyper-V Manager to delete and un-register the VM from both hosts.
4. Refresh Hyper-V Manager, and confirm that the VM no longer appears on both hosts.
5. Restore the backup copies of the `.VMCX` and `.VMRS` files back to their original location.
6. Import the VM from Hyper-V Manager using the "Import Virtual Machine" option by specifying its original location.
7. Confirm that the VM is imported and started successfully.

## DETAILED STEPS

- 
- Step 1** Use Hyper-V Manager to power off the VM from the host on which it is running.
- Step 2** Copy all files of the VM to a location to serve as backup.
- Step 3** Use Hyper-V Manager to delete and un-register the VM from both hosts.  
This operation removes the `.VMCX` and `.VMRS` files. Other VM files will remain.
- Step 4** Refresh Hyper-V Manager, and confirm that the VM no longer appears on both hosts.
- Step 5** Restore the backup copies of the `.VMCX` and `.VMRS` files back to their original location.
- Step 6** Import the VM from Hyper-V Manager using the "Import Virtual Machine" option by specifying its original location.
- Step 7** Confirm that the VM is imported and started successfully.
- 

# Replacing a Compute Node

If a compute node boot disk or blade is corrupted and the node needs to be replaced, perform the following steps:

1. Remove the compute node from the existing Hyper-V HyperFlex Cluster.
2. Reinstall OS and re-add the compute node into the cluster.




---

**Note** Compute nodes are supported in HyperFlex release 3.5.2 and later releases.

---

This section provides the procedure for replacing a compute node that needs to be replaced due to faulty boot disk or blade.

- 
- Step 1** Use Hyper-V failover cluster manager and remove the bad compute node from the failover cluster manager.
- Step 2** Clean up the computer object of the compute node from the Active Directory.  
**Note** There is no need to clean up DNS entry of the compute node.
- Step 3** Navigate to any controller VM and run the `remcomputenode.py` script to clean up the stale entries associated with the compute node.

The remove compute node Python script can be executed by providing either the UUID or host name of the compute node as an argument.

The following sample shows how to run the script with UUID of the compute node:

```
python remcomputenode.py -u C2581942-55D2-8021-B1B1-A117F396D671
```

The following sample shows how to run the script with host name of the compute node:

```
python remcomputenode.py -n node-hv1.cloud.local
```

**Note** Ensure that the following .egg files are available in the controller VM:

- /usr/share/thrift-0.9.1.a-py2.7-linux-x86\_64.egg
- /opt/springpath/storfs-mgmt-cli/stCli-1.0-py2.7.egg

**Step 4** Replace the faulty MB, compute blade, or boot disk.

**Step 5** Run the compute node expansion workflow from the Installer VM.

- Install Windows 2016.
  - On the **HX Data Platform Installer** page, select the **I know what I'm doing...** check box.
  - Select the expansion workflow and complete the procedure.
-





# CHAPTER 6

## Managing Users

- [Managing Cisco HyperFlex Users Overview, on page 57](#)
- [Creating RBAC Users for Cisco HX Data Platform, on page 59](#)
- [Assigning Users Privileges, on page 59](#)

### 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.
- **HX service account user**—A created Cisco HX Data Platform user. This user has read and modify permissions. The password is set during user creation.
- **read-only**—Other domain admins are read only users. This user only has read permissions. The password is set during user creation.

| HX Interface               | admin                                                                                              | root      | hx_admin                                            | hx_readonly                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------|-----------|-----------------------------------------------------|------------------------------------------------------------------------------------------------|
| HX Data Platform Installer | Required                                                                                           | Optional  | Not valid                                           | Not valid                                                                                      |
| HX Connect                 | Can perform most HX tasks.<br><br>local/ prefix required for login.<br>Example:<br><br>local/admin | Not valid | Can perform most HX tasks.<br><br>A preferred user. | Can only view monitoring information.<br><br>Cannot perform HX tasks.<br><br>A preferred user. |

| HX Interface                                               | admin                                                                                                                     | root                                                                                                                     | hx_admin                   | hx_readonly                                                                                                                                                                                             |
|------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Storage Controller VM with <code>hxcli</code> command line | Can perform most HX tasks.                                                                                                | Can perform most HX tasks.                                                                                               | Can perform most HX tasks. | Can only run non-interactive <code>hxcli</code> commands to view status.<br><br>Cannot perform HX tasks.<br><br><code>vc-</code> prefix required for login. Example:<br><br><code>vc-hx_readonly</code> |
| HX REST API                                                | Can perform most HX tasks.<br><br><code>local/</code> prefix required for login. Example:<br><br><code>local/admin</code> | Can perform most HX tasks.<br><br><code>local/</code> prefix required for login. Example:<br><br><code>local/root</code> | Can perform most HX tasks. | Can only run status level REST APIs.<br><br>Cannot perform HX tasks.                                                                                                                                    |

## 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.

- **audit.log**—Contains audit records for REST API and hxcli activity.

Sample entry. Note the user name, `administrator@yourdomain.local`

```
2017-03-29-01:47:28.779 - 127.0.0.1 -> 127.0.0.1 - GET /rest/clusters 200;
administrator@yourdomain.local 454ms
```

## Creating RBAC Users for Cisco HX Data Platform

Cisco HX Data Platform supports Role-Based Access Control (RBAC) for Authentication, Authorization, and Accounting (AAA) and the AAA implementation of the Open Authorization (OAuth) protocol. Cisco HX Data Platform interfaces use the Microsoft Active Directory integration for authentication and authorization domain.

Two roles are supported. Privileges associated with these roles cannot be modified.

- **Administrator**—The role allows users to modify the HX Storage Cluster. Most tasks that can be performed on a HX Storage Cluster require administrator privileges. Administrative users create other users and assign them roles.
- **Read Only**—The role allows users to monitor status and summary information. Read only users cannot perform any task that modifies the HX Storage Cluster.

RBAC created users have access to the HX Data Platform interfaces. This includes users assigned either administrator or read only permissions. The difference between the two is what the users are allowed to do.

- Cisco HX Connect
- Storage Controller VM command line for running `hxcli` commands
- Cisco HyperFlex Systems REST APIs

## Assigning Users Privileges

### Before you begin

Create the user.

- 
- Step 1** Open Active Directory Users and Computers tool.
  - Step 2** Add user to **Administrators group** under the Builtin OU for Administrator privilege.
  - Step 3** Double click on **Administrators group** to add administrator privilege user or **Remote Desktop Users group** to add read only users.
  - Step 4** Navigate to the **Members** tab
  - Step 5** Click **Add** button
  - Step 6** Type the user in the search field and click **Check Names** button.
  - Step 7** Then click **OK** to close out of each dialog box.
-



## CHAPTER 7

# Data Protection

---

- [Hyper-V Checkpoints, on page 61](#)
- [Partner Solutions, on page 62](#)

## Hyper-V Checkpoints

Checkpoints are a built-in Microsoft Hyper-V technology that enable the ability to save the state of a virtual machine. Current versions of Hyper-V support the use of production checkpoints, which creates a data-consistent copy of a virtual machine. Data consistency is achieved using Volume Shadow Copy Service on Windows virtual machines and using File System Freeze on Linux virtual machines.

By themselves, checkpoints may be useful when making changes to a virtual machine and there is a potential need to revert to the state of the virtual machine at the point in time when the checkpoint was created. Examples include the testing of software changes, applying patches, or performing configuration changes.

Checkpoints, however, do not take the place of traditional backups. By default, a checkpoint resides on the same storage as the virtual machine it was created from. A catastrophic storage failure may render a checkpoint useless such it could not be used to successfully perform a revert operation. Additionally, checkpoints may negatively impact virtual machine performance during the time period that they are retained. This stated, checkpoints are used in conjunction with popular backup solutions to create a data-consistent point in time copy of a virtual machine. The checkpoint is deleted by the backup solution when the backup job ends. In this sense, production checkpoints form the basis for capturing a data-consistent point in time copy of a virtual machine.

Hyper-V Replica is another built-in Microsoft Hyper-V technology that enables a disaster recovery strategy where one or more virtual machines from a primary Hyper-V host are replicated to a secondary Hyper-V host. Extended replication is also possible, where virtual machines replicated to a secondary host can then be replicated to a third host.

After the initial replication of a virtual machine, where an identical replica virtual machine is created on a secondary host, change tracking is used capture changes on the virtual machine VHD. Changed data is synchronized based on the configured replication frequency. Replication frequency is typically based on the RPO of the virtual machine, and available replication network bandwidth. By default, only the most recent replica (recovery point) is saved at the replication destination, however the feature can be configured to retain additional recovery points if required.

When configured as a replica server, a Hyper-V host can receive replicas from one or more source Hyper-V hosts. Failover capabilities include “test”, “planned”, and “unplanned” operations.

## Partner Solutions

Microsoft Hyper-V can store VM files in shares using the SMB 3.0 protocol. Cisco HyperFlex clusters configured for Hyper-V present storage in the cluster as an SMB share for use by Hyper-V. Popular backup solutions support the protection of Hyper-V virtual machines residing on SMB3 shares. The partner solutions highlighted include Commvault, and Veeam.

Commvault requires the installation of a Virtual Server Agent (VSA) on all Hyper-V nodes that host VMs on SMB shares. The installation process is simple and straightforward. The Hyper-V failover cluster is then added and registered as a virtual server and each Hyper-V node acts as a proxy when performing backup and recovery operations. Commvault leverages VM checkpoints when performing backups, and promptly deletes any checkpoints it created at backup completion time.

Veeam requires the user to configure virtualization infrastructure for Hyper-V. In the case of HyperFlex, this can be configured at the cluster level. In cases where multiple HyperFlex clusters are managed using the same instance of Microsoft SCVMM, Veeam allows the user to simply configure the SCVMM instance instead of configuring multiple clusters independently. When the virtual infrastructure is configured, all required Veeam components are automatically deployed on all Hyper-V nodes in each HyperFlex cluster. Veeam leverages VM checkpoints when performing backups, and promptly deletes any checkpoints it created at backup completion time.