



Preparing for HX Storage Cluster Maintenance

- [Storage Cluster Maintenance Operations Overview, on page 1](#)
- [Serial vs. Parallel Operations, on page 3](#)
- [Checking Cluster Status, on page 3](#)
- [Setting a Beacon, on page 3](#)
- [Verify vMotion Configuration for HX Cluster, on page 4](#)
- [Maintenance Modes for Storage Cluster Nodes, on page 5](#)
- [Entering Cisco HyperFlex Maintenance Mode, on page 6](#)
- [Exiting HXDP Maintenance Mode, on page 7](#)
- [Creating a Backup Operation, on page 8](#)
- [Shut Down and Power Off the Cisco HX Storage Cluster, on page 12](#)
- [Power On and Start Up the Cisco HX Storage Cluster, on page 14](#)
- [Restoring the Configuration for a Fabric Interconnect, on page 16](#)
- [Configure PCI Passthrough After Changing vNIC or vHBAs, on page 18](#)

Storage Cluster Maintenance Operations Overview

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

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



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

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

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

Online vs Offline Maintenance

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

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

Pre-Maintenance Tasks

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

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

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

- Check the list of maintenance tasks that cannot be performed in parallel. See [Serial vs. Parallel Operations, on page 3](#) for more information on these tasks.. You can perform only some tasks serially to each other.
- Ensure that SSH is enabled on all the ESX hosts.
- Put the ESX host into HXDP Maintenance Mode prior to performing a maintenance task on the host. The HXDP Maintenance Mode performs additional storage cluster specific steps compared to the vSphere provided Host Maintenance Mode.

Post Maintenance Tasks

After the maintenance task is completed, the nodes need to exit HXDP Maintenance Mode and the storage cluster needs to be healthy. In addition, some changes to the Cisco HX storage cluster require additional post maintenance tasks. For example, if you change the vNICs or vHBAs, the PCI Passthrough needs to be reconfigured. For more information describing how to reconfigure the PCI Passthrough, see [Configure PCI Passthrough After Changing vNIC or vHBAs, on page 18](#).

Ensure the following:

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

Serial vs. Parallel Operations

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

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

Checking Cluster Status

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

Step 2 Verify the storage cluster is healthy.

```
# hxcli cluster info
```

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

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

Step 3 Verify the number of node failures.

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

Example response:

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

Setting a Beacon

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

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

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

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

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

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

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

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

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

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

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

- a) Log into HX Connect.
- b) Select **System Information > Disks**.
- c) Select a node, and then click **Turn On Locator LED** or **Turn Off Locator LED**.

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

Verify vMotion Configuration for HX Cluster

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



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

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



Note The name is case-sensitive.

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



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

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

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

Maintenance Modes for Storage Cluster Nodes

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

There are two types of maintenance modes.

- HXDP Maintenance Mode
- Host Maintenance Mode

HXDP Maintenance Mode

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

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

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

HXDP Maintenance Mode Considerations

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

resources to serve IO optimally. Each node is used to serve part of user IO as well as be responsible for doing internal bookkeeping activities.

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

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

- See [Entering Cisco HyperFlex Maintenance Mode](#) and [Exiting HXDP Maintenance Mode, on page 7](#) for steps.

Host Maintenance Mode

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

To enter or exit vSphere maintenance mode:

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

Entering Cisco HyperFlex Maintenance Mode

Using the Cisco HyperFlex (HX) Connect User Interface

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



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

Using the vSphere Web Client

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

Using the Command-Line Interface

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

- a. Identify the node ID and IP address.

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

- b. Enter the node into HXDP Maintenance Mode.

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

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

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

```
# esxcli system maintenanceMode get
```

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

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

Exiting HXDP Maintenance Mode

Using the Cisco HyperFlex (HX) Connect User Interface

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

Using the vSphere Web Client

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

Using the Command-Line Interface

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

- a. Identify the node ID and IP address.

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

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

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

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

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

```
# esxcli system maintenanceMode get
```

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

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

Creating a Backup Operation

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

Before you begin

1. Log into UCS Manager.
2. Obtain the backup server IPv4 address and authentication credentials.



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

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

Name	Description
Preserve Identities check box	<p>This checkbox remains selected for All Configuration and System Configuration type of backup operation, and provides the following functionality:</p> <ul style="list-style-type: none"> • All Configuration—The backup file preserves all identities derived from pools, including vHBAs, WWPNs, WWNN, vNICs, MACs and UUIDs. Also, the identities for Chassis, FEX, Rack Servers, and user labels for Chassis, FEX, Rack Servers, IOMs and Blade Servers are preserved. <ul style="list-style-type: none"> Note If this check box is not selected the identities will be reassigned and user labels will be lost after a restore. • System Configuration—The backup file preserves identities for Chassis, FEX, Rack Servers, and user labels for Chassis, FEX, Rack Servers, IOMs and Blade Servers. <ul style="list-style-type: none"> Note If this check box is not selected the identities will be reassigned and user labels will be lost after a restore. <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> <ul style="list-style-type: none"> Note If this check box is not selected the identities will be reassigned and user labels will be lost after a restore.
Location of the Backup File field	<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. <ul style="list-style-type: none"> Java-based Cisco UCS Manager GUI displays the Filename field with an associated Browse button that let you specify the name and location for the backup file. <ul style="list-style-type: none"> Note Once you click OK, the location cannot be changed. HTML-based Cisco UCS Manager GUI displays the Filename field. Enter a name for the backup file in <code><filename>.xml</code> format. The file is downloaded and saved to a location depending on your browser settings.

Name	Description
<p>Protocol field</p>	<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.
<p>Hostname field</p>	<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>
<p>Remote File field</p>	<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>
<p>User field</p>	<p>The username the system should use to log into the remote server. This field does not apply if the protocol is TFTP.</p>
<p>Password field</p>	<p>The password for the remote server username. This field does not apply if the protocol is TFTP.</p> <p>Cisco UCS Manager does not store this password. Therefore, you do not need to enter this password unless you intend to enable and run the backup operation immediately.</p>

Step 7 Click **OK**.

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

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

- Step 9** (Optional) To view the progress of the backup operation, do the following:
- If the operation does not display in the **Properties** area, click the operation in the **Backup Operations** table.
 - In the **Properties** area, click the down arrows on the **FSM Details** bar.

The **FSM Details** area expands and displays the operation status.

- Step 10** Click **OK** to close the **Backup Configuration** dialog box.

The backup operation continues to run until it is completed. To view the progress, re-open the **Backup Configuration** dialog box.

Shut Down and Power Off the Cisco HX Storage Cluster

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

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

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

Before you begin

- The storage cluster must be in a healthy state.
- Prior to shutdown, verify that the HyperFlex cluster has one reachable external NTP and DNS resource configured that resides outside the HyperFlex.
- Perform both the Full-State and All Configuration type backups with the Preserve Identities attribute. See [Creating a Backup Operation, on page 8](#).

- Step 1** Gracefully shut down all workload VMs on all the Cisco HX datastores.
Alternatively, use vMotion to migrate the workload VMs to another cluster.

Note Do not shut down or move the storage controller VMs (stCtlVMs).

- Step 2** Gracefully shut down the Cisco HX storage cluster.
- From any controller VM command line, run the command and wait for the shell prompt to return.

Note For clusters with a nested vCenter, performing an stcli cluster shutdown may have certain limitations. For more details, see [Known Constraints with vCenter Deployment](#).

```
# stcli cluster shutdown
```

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

```
# hxcli cluster info
```

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

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

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

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

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

Choose a method:

Using vCenter Shut Down Guest OS

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

This method performs a graceful guest VM shutdown.

Using vCenter ESX Agent Manager

- From vCenter client, open the ESX Agent Manager console.
- Locate the controller VM on each ESX host, and select **Power > Shut Down Guest OS**.

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

Using vCenter Host Maintenance Mode

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

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

Step 5 Shutdown each storage cluster ESX host.

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

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

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

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

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

The HX storage cluster is now safely powered off.

What to do next

1. Complete the task that required the storage cluster shutdown or power off. For example, an offline upgrade, physically moving the storage cluster, or performing maintenance on nodes.
 - For upgrade tasks, see the [Cisco HyperFlex Systems Upgrade Guide](#).
 - For hardware replacement tasks, see the server hardware guides.

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



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

2. To restart the Cisco HX storage cluster, proceed to [Power On and Start Up the Cisco HX Storage Cluster, on page 14](#).

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 12](#).

Step 1

Plug in to power up the FIs.

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

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

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

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

Step 2

Connect all the ESX hosts to the FIs.

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

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

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

- Step 3** Verify all the ESXi hosts are network reachable.
Ping all the management addresses.
- Step 4** Exit each node from maintenance mode.
- Note** This is automatically completed by the `hxcli cluster start` command.
- Step 5** If all the controller VMs are not automatically powered on, power on all the controller VMs (`stCtrlVM`) using one of the following methods:
- Using vSphere Client
- From the vSphere Client, view a storage controller host.
 - Right-click the `stCtrlVM` and select **Power > Power On**.
 - Repeat for each host.
- Using ESXi host command line
- Log into a host.
 - Identify the VMID of the `stCtrlVM`.

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

```
# vim-cmd vmsvc/power.on VMID
```
 - Repeat for each host.
- Step 6** Wait for all the controller VMs to boot and become network reachable. Then verify.
Ping the management addresses of each of the controller VMs.
- Step 7** Verify the storage cluster is ready to be restarted.
- SSH to any controller VM, run the command:

```
# 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 HXDP Maintenance Mode or Host Maintenance Mode. Ignore any error messages about an unknown host exception.
- Step 9** Wait until the storage cluster is online and returns to a healthy state.
- 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** Through vCenter, verify that ESX remounted the datastores.
Once the cluster is available, the datastores are automatically mounted and available.
If ESX does not recognize the datastores, from the ESX command line, run the command.

```
# esxcfg-nas -r
```

- Step 11** When the storage cluster is healthy and the datastores are remounted, power on the workload VMs.
Alternatively, use vMotion to migrate the workload VMs back to the storage cluster.

Restoring the Configuration for a Fabric Interconnect

It is recommended that you use a full state backup file to restore a system that is running the same version as the system from which the backup file was exported. You can also use a full state backup to restore a system if they have the same release train. For example, you can use a full state backup taken from a system running Release 4.5(1a) to restore a system running Release 4.5(2a).

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

Before you begin

Collect the following information to restore the system configuration:

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



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

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



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

- Step 1** Connect to the console port.
- Step 2** If the fabric interconnect is off, power on the fabric interconnect.
You will see the power on self-test message as the fabric interconnect boots.
- Step 3** At the installation method prompt, enter **gui**.

Step 4 If the system cannot access a DHCP server, enter the following information:

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

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

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

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

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

- SCP
- TFTP
- FTP
- SFTP

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

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

Step 10 Click **Submit**.

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

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

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

Configure PCI Passthrough After Changing vNIC or vHBAs

Description

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

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

Action: Update the vSphere Service Profile on the ESX Host

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

From vCenter Client: Right-click the *HX storage controller VM*, select **Edit Settings** > **Add** > **PCI Device** > **LSI Logic PCI device** > **OK**.

From vCenter Web Client: Right-click the *HX storage controller VM*, select **Edit Settings** > **PCI Device** > **Add** > **LSI Logic PCI device** > **OK**.

Step 8 Remove the ESX host from HXDP Maintenance Mode.

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