Preparing for HX Storage Cluster Maintenance

- Storage Cluster Maintenance Operations Overview, on page 1
- Serial vs. Parallel Operations, on page 2
- Automating Updates Using Cluster-Aware Updating (CAU), on page 3
- Checking Cluster Status, on page 5
- Setting a Beacon, on page 5
- Verify Live Migration Configuration for HX Cluster, on page 6
- Maintenance Modes for Storage Cluster Nodes, on page 6
- Entering Cisco HyperFlex Maintenance Mode, on page 7
- Exiting Cisco HyperFlex Maintenance Mode, on page 8
- Creating a Backup Operation, on page 8
- Shut Down and Power Off the Cisco HX Storage Cluster, on page 13
- Power On and Start Up the Cisco HX Storage Cluster, on page 15
- Restoring the Configuration for a Fabric Interconnect, on page 16
- Recommendations for Verification after Storage Outages, on page 18
- Replacing a Compute Node, on page 19

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 2 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.

Procedure

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

Procedure

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 heathy:

```
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.

Procedure

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

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

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

Step 2 Turn on and off a disk beacon using 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 7and Exiting Cisco HyperFlex Maintenance Mode, on page 8 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
```

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
```

Creating a Backup Operation

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

Before you begin

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

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

Procedure

Step 1 In the **Navigation** pane, click **Admin**.

Step 2 Click the **All** node.

Step 3 In the **Work** pane, click the **General** tab.

Step 4 In the **Actions** area, click **Backup Configuration**.

Step 5 In the **Backup Configuration** dialog box, click **Create Backup Operation**.

Step 6 In the **Create Backup Operation** dialog box, complete the following fields:

Name	Description
Admin State field	This can be one of the following: • **Enabled**—Cisco UCS Manager runs the backup operation as soon as you click **OK**. • **Disabled**—Cisco UCS Manager does not run the backup operation when you click **OK**. If you select this option, all fields in the dialog box remain visible. However, you must manually run the backup from the **Backup Configuration** dialog box.

Name	Description
Type field	The information saved in the backup configuration file. This can be one of the following: • **Full state**—A binary file that includes a snapshot of the entire system. You can use the file generated from this backup to restore the system during disaster recovery. This file can restore or rebuild the configuration on the original fabric interconnect, or recreate the configuration on a different fabric interconnect. You cannot use this file for an import. **Note** You can only use a full state backup file to restore a system that is running the same version as the system from which the backup file was exported. It is also important that the bundles from which the backup was taken remain present in the Cisco UCS Manager and should not be deleted. • **All configuration**—An XML file that includes all system and logical configuration settings. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore. This file does not include passwords for locally authenticated users. • **System configuration**—An XML file that includes all system configuration settings such as usernames, roles, and locales. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore. • **Logical configuration**—An XML file that includes all logical configuration settings such as service profiles, VLANs, VSANs, pools, and policies. You can use the file generated from this backup to import these configuration settings to the original fabric interconnect or to a different fabric interconnect. You cannot use this file for a system restore.

Name	Description
Preserve Identities check box	This checkbox remains selected for **All Configuration** and **System Configuration** type of backup operation, and provides the following functionality: • **All Configuration**—The backup file preserves all identities derived from pools, including vHBAs, WWPNs, WWNN, vNICs, MACs and UUIDs. Also, the identities for Chassis, FEX, Rack Servers, and user labels for Chassis, FEX, Rack Servers, IOMs and Blade Servers are preserved. **Note** If this check box is not selected the identities will be reassigned and user labels will be lost after a restore. • **System Configuration**—The backup file preserves identities for Chassis, FEX, Rack Servers, and user labels for Chassis, FEX, Rack Servers, IOMs and Blade Servers. **Note** If this check box is not selected the identities will be reassigned and user labels will be lost after a restore. If this checkbox is selected for **Logical Configuration** type of backup operation, the backup file preserves all identities derived from pools, including vHBAs, WWPNs, WWNN, vNICs, MACs and UUIDs. **Note** If this check box is not selected the identities will be reassigned and user labels will be lost after a restore.
Location of the Backup File field	Where the backup file should be saved. This can be one of the following: • **Remote File System**—The backup XML file is saved to a remote server. Cisco UCS Manager GUI displays the fields described below that allow you to specify the protocol, host, filename, username, and password for the remote system. • **Local File System**—The backup XML file is saved locally. Java-based Cisco UCS Manager GUI displays the **Filename** field with an associated **Browse** button that let you specify the name and location for the backup file. **Note** Once you click **OK**, the location cannot be changed. HTML-based Cisco UCS Manager GUI displays the **Filename** field. Enter a name for the backup file in **<*filename*>.xml** format. The file is downloaded and saved to a location depending on your browser settings.

Name	Description
Protocol field	The protocol to use when communicating with the remote server. This can be one of the following: • **FTP** • **TFTP** • **SCP** • **SFTP** • **USB A**—The USB drive inserted into fabric interconnect A. This option is only available for certain system configurations. • **USB B**—The USB drive inserted into fabric interconnect B. This option is only available for certain system configurations.
Hostname field	The hostname, IPv4 address of the location where the backup file is stored. This can be a server, storage array, local drive, or any read/write media that the fabric interconnect can access through the network. **Note** If you use a hostname rather than an IPv4 address, you must configure a DNS server. If the Cisco UCS domain is not registered with Cisco UCS Central or DNS management is set to **local**, configure a DNS server in Cisco UCS Manager . If the Cisco UCS domain is registered with Cisco UCS Central and DNS management is set to **global**, configure a DNS server in Cisco UCS Central. **Note** All IP addresses must be IPv4. HyperFlex does not support IPv6 addresses.
Remote File field	The full path to the backup configuration file. This field can contain the filename as well as the path. If you omit the filename, the backup procedure assigns a name to the file.
User field	The username the system should use to log into the remote server. This field does not apply if the protocol is TFTP.
Password field	The password for the remote server username. This field does not apply if the protocol is TFTP. Cisco UCS Manager does not store this password. Therefore, you do not need to enter this password unless you intend to enable and run the backup operation immediately.

Step 7 Click **OK**.

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

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

Step 9 (Optional) To view the progress of the backup operation, do the following:

a) If the operation does not display in the **Properties** area, click the operation in the **Backup Operations** table.

b) In the **Properties** area, click the down arrows on the **FSM Details** bar.

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

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

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

Shut Down and Power Off the Cisco HX Storage Cluster

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

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

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

- **A shut-down cluster** has all storage cluster processes, including the working VMs, powered down. This does not include powering down the nodes in the cluster or shutting down the 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 8.

Procedure

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.

a) 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 13.

Procedure

Step 1 Plug in to power up the FIs.

a) Power on the primary FI. Wait until you can gain access to UCS Manager.

b) Power on the secondary FI. Verify it is online in UCS Manager.

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

a. Log 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.

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

```
# hxcli about
```

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

Step 8 Start the storage cluster.

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

```
# hxcli cluster start
```

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

a) From any controller VM, run the command.

```
# hxcli cluster info
```

b) In the command response text, check the cluster subsection and verify the `healthstate` is `online`.

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

Step 10 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.

Procedure

Step 1 Connect to the console port.

Step 2 If the fabric interconnect is off, power on the fabric interconnect.

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

Step 3 At the installation method prompt, enter **`gui`**.

Step 4 If the system cannot access a DHCP server, 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
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.

Name	Description
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. It is also important that the bundles from which the backup was taken remain present in the Cisco UCS Manager and should not be deleted.
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 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.

Procedure

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.

Procedure

Step 1 Use Hyper-V failover cluster manager and remove the bad compute node from the failover cluster manager.

Step 2 Clean up the computer object of the compute node from the Active Directory.

Note
There is no need to clean up DNS entry of the compute node.

Step 3 Navigate to any controller VM and run the `remcomputenode.py` script to clean up the stale entries associated with the compute node.

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

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

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

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

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

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

- /usr/share/thrift-0.9.1.a-py2.7-linux-x86_64.egg
- /opt/springpath/storfs-mgmt-cli/stCli-1.0-py2.7.egg

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

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

a) Install Windows 2016.

b) On the **HX Data Platform Installer** page, select the **I know what I'm doing...** check box.
c) Select the expansion workflow and complete the procedure.