



Managing HX Storage Clusters

- [Changing the Cluster Access Policy Level, on page 1](#)
- [Rebalancing the Cluster, on page 1](#)
- [Handling Out of Space Errors, on page 3](#)
- [Moving the Storage Cluster from a Current vCenter Server to a New VCenter Server, on page 4](#)
- [Unregistering a Storage Cluster from a vCenter Cluster, on page 4](#)
- [Renaming Clusters, on page 9](#)
- [Replacing Self-Signed Certificate, on page 9](#)
- [Boost Mode, on page 13](#)
- [UEFI Secure Boot Mode, on page 14](#)
- [hx_edge.py Script Fails after ESXi Redeploy on Auto Secure Booted Node, on page 16](#)
- [Catalog Update, on page 17](#)

Changing the Cluster Access Policy Level

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

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

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

Rebalancing the Cluster

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



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



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

Verify rebalancing status from the storage controller VM.

a) Enter the following on the command line:

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

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

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

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

Checking Cluster Rebalance and Self-Healing Status

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



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

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

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

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

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

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

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

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

```
# stcli rebalance status
```

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

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

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

Handling Out of Space Errors

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

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



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

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

Step 2 To delete unused VMs, complete the following:

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

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

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

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

Before you begin

- Perform this task during a maintenance window.
- Ensure the cluster is healthy and upgrade state is OK and Healthy. You can view the state using the `stcli` command from the controller VM command line.

```
# stcli cluster info
```

Check response for:

```
Resiliency Health: HEALTHY
```

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

Step 1 From the current vCenter, delete the cluster.

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

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

Step 2 On the new vCenter, create a new cluster using the same cluster name.

Step 3 Add ESX hosts to new vCenter in the newly created cluster.

What to do next

Proceed to [Unregistering a Storage Cluster from a vCenter Cluster, on page 4](#).

Unregistering a Storage Cluster from a vCenter Cluster

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

Before you begin

As part of the task to move a storage cluster from one vCenter server to another vCenter server, complete the steps in [Moving the Storage Cluster from a Current vCenter Server to a New vCenter Server, on page 4](#).

**Note**

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

- Step 1** Complete the steps in [Removing HX Data Platform Files from the vSphere Client, on page 7](#).
- Step 2** Complete the steps in [Verifying HX Cluster is Unregistered from vCenter, on page 7](#).

What to do next

Proceed to [Registering a Storage Cluster with a New vCenter Cluster, on page 8](#).

Unregistering and Removing EAM Extensions

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

Before you begin

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

**Note**

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

- Step 1** Identify the HX cluster UUID.
- Every agency has a field `cluster_domain_id` which refers to the underlying vSphere extension. This extension ID uses a Managed Object ID (moid).

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

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

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

- Step 2** To unregister the storage cluster extension: Log into the vCenter server MOB extension manager
First unregister the HyperFlex cluster.
- In a browser, enter the path and command.
`https://vcenter_server/mob/?moid=ExtensionManager`
`vcenter_server` is the IP address of the vCenter where the storage cluster is currently registered.
 - Enter administrator login credentials.
- Step 3** Locate the HX storage cluster extensions with the cluster IDs. Scroll through the **Properties > extensionList** to locate the storage cluster extensions:
- `com.springpath.sysmgmt.cluster_domain_id` and `com.springpath.sysmgmt.uuid.cluster_domain_id`.
- Copy each of these strings into your clipboard. Exclude the double quotes (") on either end of string, if there are any.
- Step 4** Unregister each storage cluster extension.
- From the Methods table click `UnregisterExtension`.
 - In the **UnregisterExtension** popup, enter an extension key value, `com.springpath.sysgmt.cluster_domain_id`.
For example: `com.springpath.sysgmt.domain-26`
 - Click **Invoke Method**.
- Step 5** To remove stale EAM extensions: Log into the vCenter server MOB ESX agencies extension manager.
Second remove stale EAM extensions that were associated with the HyperFlex cluster.
- In a browser, enter the path and command.
`https://vcenter_server/eam/mob/`
`vcenter_server` is the IP address of the vCenter where the storage cluster is currently registered.
 - Enter administrator login credentials.
- Step 6** Locate the stale HX storage cluster ESX agency extensions with the cluster IDs.
- Scroll through the **Properties > agency > Value**.
 - Click an agency value.
 - In the **Agency** window, check the **Properties > solutionID > Value** extension. Verify has the correct `cluster_domain_id`.
For example: `com.springpath.sysgmt.domain-26`
- Step 7** Remove stale ESX agency extensions.
- From the **Agency** window, **Methods** table select a method.
Stale ESX agencies can be removed using either the `destroyAgency` OR `uninstall`.
 - In the *method* popup, click **Invoke Method**.
- Step 8** Refresh the **ExtensionManager** tab and verify that the **extensionList** entry does not include `com.springpath.sysgmt.cluster_domain_id` extensions.
- Step 9** Restart the vSphere Client services.
The HX Data Platform extensions are removed when the vSphere Client services are restarted. Restarting the vSphere client service temporarily disables access to vCenter through the browser. For additional information, see the *VMware*

KB, Stopping, starting, or restarting VMware vCenter Server Appliance 6.0 services (2109887) article on the VMware Customer connect site.

Removing HX Data Platform Files from the vSphere Client

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

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

Linux vCenter

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

For vCenter 6.0

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

For vCenter 5.5

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

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

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

- d) Restart the vSphere Client.

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

Windows vCenter

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

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

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

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

- d) Open the Service screen.

```
# services.msc
```

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

```
# serviceLogout
```

Verifying HX Cluster is Unregistered from vCenter

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

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

Before you begin

Complete the steps in: [Removing HX Data Platform Files from the vSphere Client, on page 7](#)

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

Registering a Storage Cluster with a New vCenter Cluster

Before you begin

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

As part of the task to move a storage cluster from one vCenter server to another vCenter server, complete the steps in [Unregistering a Storage Cluster from a vCenter Cluster, on page 4](#).

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

Example:

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

Apply additional listed options as needed.

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

Example response:


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

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

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

Contact TAC for assistance if required.

- Step 3** Re-enter your snapshot schedules.
Snapshot schedules are not moved with the storage cluster when you move the storage cluster between vCenter clusters.
- Step 4** (Optional) Once registration is successful, re-enable ESXi Lockdown mode if you disabled it prior to registering the HyperFlex cluster to vCenter.

Renaming Clusters

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



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

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

Replacing Self-Signed Certificate

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

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

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

To update the certMgmt mode:

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

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

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



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

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

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

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

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

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

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

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

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

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

Reregistering a HyperFlex cluster

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

```
hxcli license register
```



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

Recreating a Self-Signed Certificate

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

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

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

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

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

Boost Mode

Boost Mode allows the Cisco HyperFlex cluster to deliver higher IOPs by increasing the storage controller VM CPU resources by 4 vCPU. Enabling Boost Mode takes additional CPU resources from user VM for the HX data platform, and should only be enabled in deployments where support has determined that the benefit of additional CPUs, outweighs the impact to the sizing of your deployment. For more information about the CPUs supported by Boost Mode, see the [Cisco HyperFlex Spec Sheets](#).

Configuring Boost Mode

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

Before you begin

Boost Mode Support is limited to the following configurations:

- Supported Hardware:
 - All NVMe
 - All Flash C245
 - All Flash C240
 - All Flash C225
 - All Flash C220
- Hypervisor: ESX only
- Boost Mode number of controller VM vCPUs:
 - All NVMe: 16
 - All Flash C245: 12
 - All Flash C240: 12
 - All Flash C225: 12
 - All Flash C220:12
- Cluster expansion requires you to apply Boost Mode to the new nodes.
- Boost Mode is supported in Cisco HX Release 4.0(2a) and later.
- Boost Mode should be enabled after support has determined that your deployment will benefit from the additional CPUs.



Note

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

-
- Step 1** From the vCenter, right-click one controller VM and **Shut Down Guest OS**.
- Step 2** Increase the number controller VM vCPUs to 16 for all-NVMe, or 12 for all flash C220, and all flash C240. In the vSphere client, click **Edit Settings** for the VM and change the value of the CPU field in the first line.
- Note** Boost Mode number of controller VM vCPUs:
- All NVMe: 16
 - All Flash C245: 12
 - All Flash C240: 12
 - All Flash C225: 12
 - All Flash C220:12
- Step 3** Click **OK** to apply the configuration changes.
- Step 4** Power up the controller VM.
- Step 5** Log into HX Connect and wait for the cluster to become healthy.
- Step 6** Repeat the process for each host (or node) in the cluster.
-

Disabling Boost Mode

To disable Boost Mode, perform the following steps:

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

UEFI Secure Boot Mode

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

Starting with HX Data Platform release 4.5(1a), the hardening of Hypervisor (ESXi) boot security is simplified by providing an automated workflow that non-disruptively changes the boot mode of converged and compute nodes in the cluster to Unified Extensible Firmware Interface (UEFI) Secure Boot. The chain of trust is anchored by a hardware trust anchor (i.e. the Cisco Trust Anchor module) built-in to UCS rack and blade

servers. UI and API-based queries of each node's secure boot status is supported so the cluster's security posture on-demand can be audited.

The following limitations apply to the UEFI boot mode:

- For HX Edge clusters, UEFI secure boot should only be enabled on HX Edge clusters running Cisco IMC version 4.1(2a) and later. If secure boot is enabled on earlier Cisco IMC versions, secure boot will need to be temporarily disabled during firmware updates.
- Support for Secure boot is available only for HyperFlex ESXi M5/M6 Server.
- Attestation of Secure Boot of ESXi hosts by vCenter is supported. This feature requires ESXi release 6.7 or later and TPM 2.0 modules on the converged or compute nodes. The TPM and TXT parameters, which are required to enable usage of the TPM module, are automatically configured in the course of enabling secure boot. No additional steps are required to use attestation.
- All the factory prepped M.2 RAID edge nodes run HXDP server firmware version 4.1(2a) or later. If a customer downgrades in field or retrofits existing setup and tries to bring up a cluster with M.2RAID nodes with HXDP server firmware version earlier than 4.1(2a), then the install may fail with the error `UEFI boot parameters cannot be configured for Legacy boot mode`. The HXDP server firmware must be upgraded to version 4.1(2a) or later and then re-try the install.

Enabling Secure Boot Mode

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

Before you begin

- Run **Check Secure Boot Status** to confirm whether Secure Boot is already enabled and then proceed accordingly. See step 4.
- Starting with HX release 4.5(1a), UEFI Secure Boot must be enabled as a separate Day 2 operation, after a refresh of HX release 4.5(1a) install or after upgrading an existing cluster to HX 4.5(1a).
- Validate that the cluster is in a ready state to enable Secure Boot by running a pre-flight validation.
- If your cluster has legacy, UEFI, and UEFI Secure booted nodes present in same cluster, secure boot operation will still get enabled on all nodes of the cluster and any expansions after that will be secure boot enabled.
- Option to Enable Secure Boot will be available only for ESXi clusters.
- Due to the ESXi hosts going in to rolling reboot to enable secure boot, plan the activity in a maintenance window.
- Enabling Secure Boot cannot be combined with other upgrade activities.

- If the Secure Boot is already enabled, **Enable Secure Boot** option is greyed out and no further action needed.
- In the case that **Enable Secure Boot** workflow fails, then from the vCenter, confirm whether the host is still in **Maintenance Mode**. If so, then exit **Maintenance Mode** before retrying the **Enable Secure Boot** workflow.

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

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

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

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

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

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

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

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


Description

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

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

Action

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



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

Catalog Update

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

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

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

Guidelines and Limitations

- Before adding a new drive, review the [HyperFlex Release](#) notes to confirm that the current HXDP version supports the new drive model.
- Catalog Update does not guarantee a drive is supported. Hardware issues and HXDP versions may contribute to a drive being unrecognized to HXDP.
- Do not use the Catalog Upgrade for drives that require HXDP tuning for custom settings, such as a higher drive capacity point; These require a full HXDP upgrade.
- Downgrading the catalog bundle to an earlier version is not supported.

Catalog Update: HX Installer

Catalog Update: Cluster Creation using HX Installer

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

Before you begin

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

-
- Step 1** Log in to the HX Data Platform Installer.
- Step 2** Follow the **Create Cluster** workflow for a standard cluster.
- Step 3** On the Server Selection page, the installer validates the drive supportability and if unsupported drive(s) is found, the unsupported drive(s) are identified and the **Upgrade Catalog** button appears.
- Step 4** Click the **Upgrade Catalog** button. The Upgrade Catalog window appears.
- Note** The window displays the catalog version in use.
- Step 5** Upload the Catalog file saved locally. Drag and drop the file into the target or click the target to browse to the file location. The upload operation is complete.
- Step 6** Click **Upgrade** to complete the upgrade or **Close** to exit the Catalog Upgrade window.

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

Catalog Update: Cluster Expansion using HX Installer

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

Before you begin

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

-
- Step 1** Log into the HX Data Platform Installer.
- Step 2** Follow the **Expand Cluster** workflow for a standard cluster.
- Step 3** On the Server Selection page, the installer validates the drive supportability and if unsupported drive(s) is found, the unsupported drive(s) are identified and the **Upgrade Catalog** button appears.
- Step 4** Click the **Upgrade Catalog** button. The Upgrade Catalog window appears.
- Note** The window displays the catalog version in use.

- Step 5** Upload the Catalog file saved locally. Drag and drop the file into the target or click the target to browse to the file location. The upload operation is complete.
- Step 6** Click **Upgrade** to complete the upgrade or **Close** to exit the Catalog Upgrade window.

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

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

Catalog Update from the HX Installer Settings

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

-
- Step 1** Log into the HX Data Platform Installer.
- Step 2** Click the **Settings** gear icon on any page.
- Step 3** Click the **Upgrade Catalog** button. The Upgrade Catalog window appears.
- Note** The window displays the catalog version in use.
- Step 4** Upload the catalog file saved locally. Drag and drop the file into the target or click the target to browse to the file location. The upload operation is complete.
- Step 5** Click **Upgrade** to complete the upgrade or **Close** to exit the Catalog Upgrade window.

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

Catalog Update: HX Connect

Cluster Catalog Upgrade using HX Connect

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

Before you begin

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



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

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

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

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

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

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

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

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

Catalog Update: Intersight

Catalog Upgrade using Intersight

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

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