



Expand Cisco HyperFlex System Clusters

- [Cluster Expansion Guidelines, on page 1](#)
- [Prerequisites When Expanding M4/M5 Clusters, on page 3](#)
- [Mixed Cluster Expansion Guidelines, on page 3](#)
- [Steps During Mixed Cluster Expansion, on page 4](#)
- [Prerequisites for Adding a Converged \(HX220c/HX240c\) Node, on page 5](#)
- [Preparing a Converged Node, on page 6](#)
- [Adding a Converged Node to an Existing Cluster, on page 6](#)
- [Prerequisites for Adding a Compute-Only Node, on page 11](#)
- [Preparing a Compute-Only Node, on page 13](#)
- [Adding a Compute-Only Node to an Existing Cluster, on page 15](#)
- [Resolving Failure of Cluster Expansion, on page 19](#)
- [Logical Availability Zones, on page 20](#)

Cluster Expansion Guidelines

Please review these guidelines before expanding your cluster.



Note If you have LAZ configured (enabled by default for clusters of size 8 or more), please review [Logical Availability Zones, on page 20](#) prior to moving ahead with expansion.

- If you have replication configured, put replication in pause mode before performing upgrade, expansion or cluster maintenance. After the upgrade, expansion or cluster maintenance is completed, then resume replication. Perform the pause and resume on any cluster that has replication configured to or from this local cluster.
- If you are using RESTful APIs to perform cluster expansion, sometimes the task may take longer than expected.
- ESXi installation is supported on SD cards for M4 converged nodes and M.2 SATA SSD for M5 converged nodes. For compute-only nodes, ESXi installation is supported for SD Cards, SAN boot, front SSD/HDD, or single M.2 SSD (using UCS-MSTOR-M2 controller). Installing ESXi on USB Flash is not supported for compute-only nodes



Note HW RAID M.2 (UCS-M2-HWRAID and HX-M2-HWRAID) is not supported on Compute-only nodes.

- You must click on the discovered cluster to proceed with expanding a standard ESX cluster in a 3.5.x or earlier release. Not doing so results in errors.
- Use only Admin credentials for the Controller VM during expansion work flow. Using any other credentials other than Admin may cause the expansion to fail.
- In the event you see an error about unsupported drives or catalog upgrade, see the [Compatibility Catalog](#).
- Support for Edge Cluster Expansion-Starting from HyperFlex Data Platform version 4.0(2e), you can expand ESXi based 10/25 GbE HyperFlex Edge clusters with 3 nodes via Intersight. Hyperflex Edge Cluster expansion is supported only on clusters deployed using Intersight. Edge Cluster expansion using Intersight is not supported for clusters deployed through the HyperFlex OVA Installer.



Note Cluster expansion is not supported on HyperFlex 2-Node Edge clusters with this feature. Expansion of Edge Clusters below Hyperflex version 4.0(2e) is not supported.

Please refer to the Intersight documentation for all requirements: [Cluster Expansion Requirements](#).

ESXi Installation Guidelines

1. Modify boot policy for compute node.

To modify the template and boot policy for HyperFlex Stretched Cluster compute only node on M5 server:

- a. Clone the template.
- b. Uncheck the Flex flash from local boot policy, if the compute M5 node does not have flash cards.
- c. Add the SAN boot with proper WWPN to the boot order.

2. Start the DPI expansion workflow.
3. When prompted, install ESXi using an ISO image.
4. Return to the DPI expansion workflow and complete the ESXi installation workflow.



Note If the Hypervisor configuration fails with the SOL logging failure message, access the installer CLI through SSH with root and default password and configure the ESXi hypervisor. Then, run the advanced installer and check the **HX Storage Software** and **Expand Cluster** check boxes to proceed with the ESXi installation process.

Prerequisites When Expanding M4/M5 Clusters

Prior to beginning cluster expansion in M4/M5 clusters, perform the following tasks:

- **Hypercheck Health Check Utility**— Cisco recommends running this proactive health check utility on your HyperFlex cluster prior to upgrade. These checks provide early visibility into any areas that may need attention and will help ensure a seamless upgrade experience. For more information, see the [Hyperflex Health & Pre-Upgrade Check Tool](#) TechNote for full instructions on how to install and run Hypercheck.
- Upgrade the HX cluster and UCS Manager to the appropriate recommended release for your deployment. For more information, see the [Cisco HyperFlex Recommended Software Release and Requirements Guide](#).
- Download and deploy the matching HX Data Platform Installer (version should be same as cluster) version to run the expansion workflow.
- M4 Servers: Upgrade existing M4 server firmware to 3.2(1) or later firmware
- Upgrade vCenter to 6.5 or later. Without vCenter 6.5, Broadwell EVC mode cannot be enabled. Only vCenter upgrade is required. ESXi can remain on an older version subject to the VMware software interoperability matrix. Proceeding with EVC mode off is not supported and will cause operational issues in the future.

Mixed Cluster Expansion Guidelines

- Hypercheck Health Check Utility— Cisco recommends running this proactive health check utility on your HyperFlex cluster prior to upgrade. These checks provide early visibility into any areas that may need attention and help ensure a seamless **upgrade** experience. For more information on how to install and run [Hypercheck](#), see the [Hypercheck: Hyperflex Health & Pre-Upgrade Check Tool Tech Note](#).
- Expanding existing M4 cluster with M5 converged nodes is supported.
- Expanding existing M5 cluster with M4 converged nodes is not supported.
- Expanding existing mixed M4/M5 cluster with M4 or M5 converged nodes is supported.
- Adding any supported compute-only nodes is permitted with all M4, M5, and mixed M4/M5 clusters using the HX Data Platform 2.6 or later Installer. Some example combinations are listed here, many other combinations are possible.

Example combinations:

Expand mixed M4/M5 cluster with compute-only B200, C220, C240 M4/M5

Expand M4 cluster with compute-only B200 M5, C220 M5, C240M5

- Only expansion workflow is supported to create a mixed cluster. Initial cluster creation with mixed M4/M5 servers is not supported.
- All M5 servers must match the form factor (220/240), type (Hybrid/AF), security capability (Non-SED only) & disk configuration (QTY, capacity, and non-SED) of the existing M4 servers. For more information on drive compatibility, refer to the [Cisco Hyperflex Drive Compatibility](#) document.
 - HX220-M5 will use a maximum of 6 capacity disks (2 disk slots to remain empty) when mixed with HX220-M4.

- HX Edge, SED, LFF, Hyper-V, and Stretched Clusters do not support mixed M4 and M5 clusters.

Mixed Cluster Expansion Guidelines for Release 3.5

A mixed cluster is defined by having both M4 and M5 HX converged nodes within the same storage cluster. When configuring a mixed cluster, the following guidelines apply:

- Expanding existing M4 cluster with M5 converged nodes is supported.
- Expanding existing M5 cluster with M4 converged nodes is not supported.
- Expanding existing mixed M4/M5 cluster with M4 or M5 converged nodes is supported.
- Adding any supported compute-only nodes is permitted with all M4, M5, and mixed M4/M5 clusters using the HX Data Platform 2.6 or later Installer. Some example combinations are listed here, many other combinations are possible.

Example combinations:
Expand mixed M4/M5 cluster with compute-only B200, C220, C240 M4/M5
Expand M4 cluster with compute-only B200 M5, C220 M5, C240M5
- Only expansion workflow is supported to create a mixed cluster. Initial cluster creation with mixed M4/M5 servers is not supported.
- All M5 servers must match the form factor (220/240), type (Hybrid/AF), security capability (Non-SED only) & disk configuration (QTY, capacity, and non-SED) of the existing M4 servers.
 - HX220-M5 will use a maximum of 6 capacity disks (2 disk slots to remain empty) when mixed with HX220-M4.
- HyperFlex Edge does not support mixed clusters.
- SED SKUs do not support mixed clusters.

Steps During Mixed Cluster Expansion

- During the validation steps, before expansion begins, an EVC check is performed. Follow the displayed guidance to manually enable EVC mode on the existing cluster at this time.



Caution Failure to enable EVC at the time of the warning will require a complete shutdown of the storage cluster and all associated VMs at a later point in time. Do not skip this warning.

- Perform the EVC mode configuration in vCenter and then retry the validation.
- Cluster expansion will then validate a second time and then continue with the expansion.

Prerequisites for Adding a Converged (HX220c/HX240c) Node

A converged node can be added to a HyperFlex cluster after cluster creation. The storage on a converged node is automatically added to the cluster's storage capacity.

Before you start adding a converged node to an existing storage cluster, make sure that the following prerequisites are met.

- Ensure that the storage cluster state is healthy.
- Ensure that the new node meets the system requirements listed under **Installation Prerequisites**, including network and disk requirements.
- Ensure that the new node uses the same configuration as the other nodes in the storage cluster. This includes VLAN IDs and switch types (whether vSwitches), VLAN tagging with External Switch VLAN Tagging (EST), VLAN tagging with Virtual Switch Tagging (VST), or Virtual Distributed Switch.



Note If the storage cluster is in an out of space condition, when you add a new node, the system automatically rebalances the storage cluster. This is in addition to the rebalancing that is performed every 24 hours.

- Ensure that the node you add is of the same model (HX220 or HX240) type (Hybrid or All Flash), and disk configuration (SED or non-SED). In addition, ensure that the number of capacity disks matches the existing cluster nodes.
- To add a node that has a different CPU family from what is already in use in the HyperFlex cluster, enable EVC. For more details, see the *Setting up Clusters with Mixed CPUs* section in the *Cisco HyperFlex Systems Installation Guide for VMware ESXi*.
- Ensure that the software version on the node matches the Cisco HX Data Platform version, the ESXi version, and the vCenter version. To identify the software version, go to the Storage Cluster Summary tab in vCenter and check the HX Data Platform version in the top section. Upgrade if necessary.



Note If you upgraded the cluster, you must download and install a new installer VM, that matches the current version of HXDP running on the cluster.

- Ensure that the new node has at least one valid DNS and NTP server configured.
- If you are using SSO or Auto Support, ensure that the node is configured for SSO and SMTP services.
- Allow ICMP for ping between the HX Data Platform Installer and the existing cluster management IP address.

Preparing a Converged Node

Step 1 Connect the converged node to the hardware and the network of the existing storage cluster.

Step 2 Ensure that the HX node is a node prepared at factory.

Note Do not reuse a removed converged node or its disks in the original cluster.

Adding a Converged Node to an Existing Cluster



Note If you are using RESTful APIs to perform cluster expansion, the task may take longer than expected.

Step 1 Launch the Cisco HX Data Platform Installer.

- a) In your web browser, enter the IP address or the node name for the HX Data Platform Installer VM. Click **Accept** or **Continue** to bypass any SSL certificate errors. The Cisco HX Data Platform Installer login page appears. Verify the HX Data Platform Installer **Build ID** in the lower right corner of the login screen.
- b) In the login page, enter the following credentials:

Username: root

Password (Default): Cisco123

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

- c) Read the EULA, check the **I accept the terms and conditions** checkbox, and click **Login**.

Step 2 On the **Workflow** page, select **Cluster Expansion**.

Step 3 On the **Credentials** page, complete the following fields.

To perform cluster creation, you can import a *JSON configuration* file with the required configuration data. The following two steps are optional if importing a JSON file, otherwise you can input data into the required fields manually.

Note For a first-time installation, contact your Cisco representative to procure the factory preinstallation JSON file.

- a. Click **Select a file** and choose your *JSON file* to load the configuration. Select **Use Configuration**.
- b. An **Overwrite Imported Values** dialog box displays if your imported values for Cisco UCS Manager are different. Select **Use Discovered Values**.

Field	Description
UCS Manager Credentials	

Field	Description
UCS Manager Host Name	UCS Manager FQDN or IP address. For example, <i>10.193.211.120</i> .
User Name	<admin> username.
Password	<admin> password.
vCenter Credentials	
vCenter Server	vCenter server FQDN or IP address. For example, <i>10.193.211.120</i> . Note <ul style="list-style-type: none"> • A vCenter server is required before the cluster can be made operational. • The vCenter address and credentials must have root level administrator permissions to the vCenter. • vCenter server input is optional if you are building a nested vCenter. See the Nested vCenter TechNote for more details.
User Name	<admin> username. For example, <i>administrator@vsphere.local</i> .
Admin Password	<root> password.
Hypervisor Credentials	
Admin User Name	<admin> username. This is root for factory nodes.
Admin Password	<root> password. Default password is <code>Cisco123</code> for factory nodes. Note Systems ship with a default password of <code>Cisco123</code> that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.

Step 4 Click **Continue**. A **Cluster Expand Configuration** page is displayed. Select the *HX Cluster* that you want to expand. If the HX cluster to be expanded is not found, or if loading the cluster takes time, enter the IP of the Cluster Management Address in the **Management IP Address** field.

Step 5 The **Server Selection** page displays a list of unassociated HX servers under the **Unassociated** tab, and the list of discovered servers under the **Associated** tab. Select the servers under the **Unassociated** tab to include in the HyperFlex cluster.

If HX servers do not appear in this list, check Cisco UCS Manager and ensure that they have been discovered.

For each server you can use the **Actions** drop-down list to set the following:

- **Launch KVM Console**—Choose this option to launch the KVM Console directly from the HX Data Platform Installer.
- **Disassociate Server**—Choose this option to remove a service profile from that server.

Note If there are no unassociated servers, the following error message is displayed:

```
No unassociated servers found. Please login to UCS Manager and ensure server ports are enabled.
```

The **Configure Server Ports** button allows you to discover any new HX nodes. Typically, the server ports are configured in Cisco UCS Manager before you start the configuration.

Step 6 Click **Continue**. The **UCSM Configuration** page appears.

Note If you imported a JSON file at the beginning, the **Credentials** page should be populated with the required configuration data from the preexisting HX cluster. This information must match your existing cluster configuration.

Step 7 Click **Continue**. The **Hypervisor Configuration** page appears. Complete the following fields:

Attention You can skip the completion of the fields described in this step in case of a reinstall, and if ESXi networking has been completed.

Field	Description
Configure Common Hypervisor Settings	
Subnet Mask	Set the subnet mask to the appropriate level to limit and control IP addresses. For example, 255.255.0.0.
Gateway	IP address of gateway. For example, 10.193.0.1.
DNS Server(s)	IP address for the DNS Server. If you do not have a DNS server, do not enter a hostname in any of the fields on the Cluster Configuration page of the HX Data Platform installer. Use only static IP addresses and hostnames for all ESXi hosts. Note If you are providing more than one DNS server, check carefully to ensure that both DNS servers are correctly entered, separated by a comma.
Hypervisor Settings	
Ensure to select Make IP Addresses and Hostnames Sequential , to make the IP addresses sequential.	
Note You can rearrange the servers using drag and drop.	
Name	Server name.
Serial	Serial number of the server.
Static IP Address	Input static IP addresses and hostnames for all ESXi hosts.

Field	Description
Hostname	Do not leave the hostname fields empty.

Step 8

Click **Continue**. The **IP Addresses** page appears. You can add more compute or converged servers, by clicking **Add Compute Server** or **Add Converged Server**.

Ensure to select **Make IP Addresses Sequential**, to make the IP addresses sequential. For the IP addresses, specify if the network should belong to Data Network or Management Network.

For each HX node, complete the following fields for Hypervisor Management and Data IP addresses.

Field	Description
Management Hypervisor	Enter the static IP address that handles the Hypervisor management network connection between the ESXi host and the storage cluster.
Management Storage Controller	Enter the static IP address that handles the HX Data Platform storage controller VM management network connection between the storage controller VM and the storage cluster.
Data Hypervisor	Enter the static IP address that handles the Hypervisor data network connection between the ESXi host and the storage cluster.
Data Storage Controller	Enter the static IP address that handles the HX Data Platform storage controller VM data network connection between the storage controller VM and the storage cluster.
<p>When you enter IP addresses in the first row for Hypervisor (Management), Storage Controller VM (Management), Hypervisor (Data), and Storage Controller VM (Data) columns, the HX Data Platform Installer applies an incremental auto-fill to the node information for the rest of the nodes. The minimum number of nodes in the storage cluster is three. If you have more nodes, use the Add button to provide the address information.</p> <p>Note Compute-only nodes can be added only after the storage cluster is created.</p>	
Controller VM Password	<p>A default administrator username and password are applied to the controller VMs. The VMs are installed on all converged and compute-only nodes.</p> <p>Important</p> <ul style="list-style-type: none"> You cannot change the name of the controller VM or the controller VM's datastore. Use the same password for all controller VMs. The use of different passwords is not supported. Provide a complex password that includes 1 uppercase character, 1 digit, 1 special character, and a minimum of 10 characters in total. You can provide a user-defined password for the controller VMs and for the HX cluster to be created. For password character and format limitations, see the section on Guidelines for HX Data Platform Special Characters in the <i>Cisco HX Data Platform Management Guide</i>.

Field	Description
Advanced Configuration	
Jumbo frames Enable Jumbo Frames checkbox	Check to set the MTU size for the storage data network on the host vSwitches and vNICs, and each storage controller VM. The default value is 9000. Note To set your MTU size to a value other than 9000, contact Cisco TAC.
Disk Partitions Clean up Disk Partitions checkbox	Check to remove all existing data and partitions from all nodes added to the storage cluster. You must backup any data that should be retained. Important Do not select this option for factory prepared systems. The disk partitions on factory prepared systems are properly configured. For manually prepared servers, select this option to delete existing data and partitions.

Step 9 Click **Start**. A **Progress** page displays the progress of various configuration tasks.

Note If the vCenter cluster has EVC enabled, the deploy process fails with a message: The host needs to be manually added to vCenter. To successfully perform the deploy action, do the following:

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

Step 10 When cluster expansion is complete, click **Launch HyperFlex Connect** to start managing your storage cluster.

Note When you add a node to an existing storage cluster, the cluster continues to have the same HA resiliency as the original storage cluster until auto-rebalancing takes place at the scheduled time.

Rebalancing is typically scheduled during a 24-hour period, either 2 hours after a node fails or if the storage cluster is out of space.

To rebalance the storage cluster before the scheduled time, perform the following steps to manually initiate rebalance storage cluster command.

- From a storage cluster controller VM command line, run the `# stcli rebalance start --force` command:
- To monitor rebalance status, run the `# stcli rebalance status` command.

Step 11 Create the required VM Network port groups and vMotion vmkernel interfaces using HyperFlex `hx_post_install` script or manually to match the other nodes in the cluster.

- SSH to HyperFlex cluster management IP.
- Log in as the admin user.
- Run the `hx_post_install` command.

- d) Follow the on-screen instructions, starting with vMotion and VM network creation. The other configuration steps are optional.

Step 12

After the new nodes are added to the storage cluster the High Availability (HA) services are reset so that HA can recognize the added nodes.

- a) Log in to vCenter.
- b) In the vSphere Web Client, navigate to the Host: **Home > vCenter > Inventory Lists > Hosts and Clusters > vCenter > Server > Datacenter > Cluster > Host**
- c) Select the new node.
- d) Right-click and select **Reconfigure for vSphere HA**.

Prerequisites for Adding a Compute-Only Node

You can add a compute-only node to a HyperFlex cluster after cluster creation. It is added to provide extra compute resources. The Cisco UCS server does not need to have any caching or persistent drives as they do not contribute any storage capacity to the cluster.

Before you start adding a compute-only node, make sure that the following prerequisites are met.

- Ensure that the storage cluster state is healthy.
- Ensure that the new node meets the compute-only system requirements listed in *Installation Prerequisites*, including network and disk requirements.
- Install ESXi hypervisor after service profile association.
- Ensure that the new node uses the same configuration as the other nodes in the storage cluster. This includes VLAN IDs and switch types (whether vSwitches), VLAN tagging with External Switch VLAN Tagging (EST), VLAN tagging with Virtual Switch Tagging (VST), or Virtual Distributed Switch.



Note If the storage cluster is in an out of space condition, when you add a new node, the system automatically rebalances the storage cluster. This is in addition to the rebalancing that is performed every 24 hours.

- Enable EVC if the new node to be added has a different CPU family than what is already used in the HX cluster. For more details, see the *Setting up Clusters with Mixed CPUs* section in the *Cisco HyperFlex Systems Installation Guide for VMware ESXi*.
- Ensure that the software version on the node matches the Cisco HX Data Platform version, the ESXi version, and the vCenter version. To identify the software version, go to the **Storage Cluster Summary** tab in vCenter and check the *HX Data Platform version* in the top section. Upgrade if necessary.
- Ensure that the new node has at least one valid DNS and NTP server configured.
- If you are using SSO or Auto Support, ensure that the node is configured for SSO and SMTP services.
- ESXi installation is supported on SD cards for M4 converged nodes and M.2 SATA SSD for M5 converged nodes. For compute-only nodes, ESXi installation is supported for SD Cards, SAN boot, front SSD/HDD, or single M.2 SSD (using UCS-MSTOR-M2 controller). Installing ESXi on USB Flash is not supported for compute-only nodes.



Note HW RAID M.2 (UCS-M2-HWRAID and HX-M2-HWRAID) is not supported on Compute-only nodes.

- Compute-only nodes are deployed with automatic detection and configuration of disk and boot policies based on the boot hardware.

Starting with HX Data Platform version 4.5(1a) and later, compute-only nodes are deployed with automatic detection and configuration of disk and boot policies based on the inventoried boot hardware. Users cannot directly select the UCSM policies. Instead, the boot device is automatically determined based on the first acceptable boot media discovered in the server. The tables below show the priority order for M4/M5 generation servers. Reading from top to bottom, the first entry that is a match based on the inventoried hardware are selected automatically during cluster expansion. For example, when expanding with a B200 compute node with a single M.2 boot SSD, the second rule in the table below is a match and used for SPT association.

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

Table 1: Priority for M6

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

Table 2: Priority for M5

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

Table 3: Priority for M4

Priority for M4			
Priority	SPT Name	Boot Device	Number of disks
1	compute-nodes-sd	FlexFlash	1 or 2
2	compute-nodes-anyld	Any other config	Any

Preparing a Compute-Only Node

- Step 1** Ensure that the server is a supported HX server and meets the requirements. For more details, see the *Host Requirements* section in the *Cisco HyperFlex Systems Installation Guide for VMware ESXi*.
- Step 2** Log in to Cisco UCS Manager.
- Open a browser and enter the Cisco UCS Manager address for the fabric interconnect of the storage cluster network.
 - Click the **Launch UCS Manager** button.
 - If prompted, download, install, and accept Java.
 - Log in with administrator credentials.
- Username: **admin**
- Password: **<admin password>**
- Step 3** Locate the server to ensure that the server has been added to the same FI domain as the storage cluster and is an approved compute-only model. Check the latest [Release Notes for Cisco HX Data Platform](#) for a full list of compatible Compute-only nodes.

Verify the HX Data Platform Installer

- Step 1** Verify that the HX Data Platform installer is installed on a node that can communicate with all the nodes in the storage cluster and compute nodes that are being added to the storage cluster.
- Step 2** If the HX Data Platform installer is not installed, see Deploy the HX Data Platform Installer.

Apply an HX Profile on a Compute-only Node Using UCS Manager

In Cisco UCS Manager the network policies are grouped into an HX profile. The HX installer handles automatic service profile association for compute-only nodes. Manual association is not required.

Once install begins, you should monitor compute-only node service profile association in UCS Manager. Wait until the server is fully associated before continuing on to install ESXi.

Install VMware ESXi on Compute Nodes



Important

Install VMware ESXi on each compute-only node.

Install a Cisco HX Data Platform supported version of ESXi. See the [Cisco HyperFlex Data Platform Release Notes](#) for a list of supported ESXi versions.

If the compute only node already has ESXi installed, it must be re-imaged with the Cisco HX Custom image.

Before you begin

Ensure the required hardware and network settings are met. For more details, see the *Installation Prerequisites* section in the *Cisco HyperFlex Systems Installation Guide for VMware ESXi*. Ensure the service profiles in the previous step have finished associating.

Step 1 Download the *HX Custom Image for ESXi* from the Cisco.com download site for Cisco HyperFlex. See [Download Software](#).

Select a networked location that can be accessed through Cisco UCS Manager.

Step 2 Log in to Cisco UCS Manager.

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

- a) In the Navigation Pane, click **Servers > Service Profiles > Sub-Organizations > hx-cluster**.
- b) Right click the **hx-cluster** and choose **KVM Console**.

Step 4 Copy the *HX-Vmware.iso* image to the KVM path for the compute server.

Example:

HX-Vmware-ESXi-60U3-5050593-Cisco-Custom-6.0.3.1.iso

Step 5 From the KVM console session, select **Virtual Media > Map CD/DVD** and mount the *HX Custom Image for ESXi* image. If you do not see the **Map CD/DVD** option, first activate virtual devices.

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

This opens in a pop-up window.

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

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

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

There is a check mark indicating that the file is on a mapped location, once the process is complete. The mapped file's full name includes the ESXi build ID.

- Step 7** Reset the compute server.
- Click the **Reset** button on the KVM console. Click **OK** to confirm.
 - Select **Power Cycle**. Click **OK**.
- Step 8** Change the boot path to point to the *HX-Vmware.iso* file.
- Press **F6**.
 - From the **Enter boot selection** menu, use the arrow keys to highlight the *Cisco vKVM-Mapped vDVD1.22* option.
 - Press **Enter** to select.
- This launches the ESXi installer bootloader. Select one of the three compute-only node options based on desired boot type: SD Card, Local Disk, or Remote Disk. Type in **yes** (all lowercase) to confirm selection. The rest of the installation is automated. ESXi will reboot several times. It is normal to see warnings that automatically dismiss after a short wait period. Wait for the *ESXi DCUI* to fully appear, signaling the end of installation.
- Step 9** Repeat steps 3 to 8 for each Cisco HyperFlex server.
- Step 10** Once ESXi is fully installed, click **continue**. Then click **Retry Hypervisor Configuration** to complete the rest of the cluster expansion.

Adding a Compute-Only Node to an Existing Cluster

To add a HyperFlex compute-only node to an existing HyperFlex system cluster, complete the following steps.



Note If you are using RESTful APIs to perform cluster expansion, sometimes the task may take longer than expected.



Note After you add a compute-only node to an existing cluster, you must manually configure the vmk2 interface for vmotion.

- Step 1** Launch the Cisco HX Data Platform Installer.
- In your web browser, enter the IP address or the node name for the HX Data Platform Installer VM. Click **Accept** or **Continue** to bypass any SSL certificate errors. The Cisco HX Data Platform Installer login page appears. Verify the HX Data Platform Installer **Build ID** in the lower right corner of the login screen.
 - In the login page, enter the following credentials:

Username: `root`

Password (Default): `Cisco123`

Note Systems ship with a default password of `Cisco123` that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.
 - Read the EULA, check the **I accept the terms and conditions** checkbox, and click **Login**.
- Step 2** On the **Workflow** page, select **Cluster Expansion**.

Step 3 On the **Credentials** page, complete the following fields.

To perform cluster creation, you can import a *JSON configuration* file with the required configuration data. The following two steps are optional if importing a JSON file, otherwise you can input data into the required fields manually.

- Note** For a first-time installation, contact your Cisco representative to procure the factory preinstallation JSON file.
- Click **Select a file** and choose your *JSON file* to load the configuration. Select **Use Configuration**.
 - An **Overwrite Imported Values** dialog box displays if your imported values for Cisco UCS Manager are different. Select **Use Discovered Values**.

Field	Description
UCS Manager Credentials	
UCS Manager Host Name	UCS Manager FQDN or IP address. For example, <i>10.193.211.120</i> .
User Name	<admin> username.
Password	<admin> password.
vCenter Credentials	
vCenter Server	vCenter server FQDN or IP address. For example, <i>10.193.211.120</i> . Note <ul style="list-style-type: none"> A vCenter server is required before the cluster can be made operational. The vCenter address and credentials must have root level administrator permissions to the vCenter. vCenter server input is optional if you are building a nested vCenter. See the Nested vCenter TechNote for more details.
User Name	<admin> username. For example, <i>administrator@vsphere.local</i> .
Admin Password	<root> password.
Hypervisor Credentials	
Admin User Name	<admin> username. This is root for factory nodes.

Field	Description
Admin Password	<p><root> password.</p> <p>Default password is <code>Cisco123</code> for factory nodes.</p> <p>Note Systems ship with a default password of <code>Cisco123</code> that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.</p>

Step 4 Click **Continue**. A **Cluster Expand Configuration** page is displayed. Select the *HX Cluster* that you want to expand. If the HX cluster to be expanded is not found, or if loading the cluster takes time, enter the IP of the Cluster Management Address in the **Management IP Address field**.

Step 5 Click **Continue**. A **Server Selection** page is displayed. On the **Server Selection** page, the **Associated** tab lists all the HX servers that are already connected. Do not select them. On the **Unassociated** tab, select the servers you wish to add to the cluster.

Step 6 Click **Continue**. The **Hypervisor Configuration** page appears. Complete the following fields:

Attention You can skip the completion of the fields described in this step in case of a reinstall, and if ESXi networking has been completed.

Field	Description
Configure Common Hypervisor Settings	
Subnet Mask	Set the subnet mask to the appropriate level to limit and control IP addresses. For example, <code>255.255.0.0</code> .
Gateway	IP address of gateway. For example, <code>10.193.0.1</code> .
DNS Server(s)	<p>IP address for the DNS Server.</p> <p>If you do not have a DNS server, do not enter a hostname in any of the fields on the Cluster Configuration page of the HX Data Platform installer. Use only static IP addresses and hostnames for all ESXi hosts.</p> <p>Note If you are providing more than one DNS server, check carefully to ensure that both DNS servers are correctly entered, separated by a comma.</p>
Hypervisor Settings Ensure to select Make IP Addresses and Hostnames Sequential , to make the IP addresses sequential.	
Note You can rearrange the servers using drag and drop.	
Name	Server name.
Serial	Serial number of the server.
Static IP Address	Input static IP addresses and hostnames for all ESXi hosts.

Field	Description
Hostname	Do not leave the hostname fields empty.

Step 7

Click **Continue**. An **IP Addresses** page is displayed. Click **Add Compute-only Node** to add a new node.

If you are adding more than one compute-only node, select **Make IP Addresses Sequential**.

Field	Information
Management Hypervisor	Enter the static IP address that handles the Hypervisor management network connection between the ESXi host and storage cluster.
Management Storage Controller	None.
Data Hypervisor	Enter the static IP address that handles the Hypervisor data network connection between the ESXi host and the storage cluster.
Data Storage Controller	None.
Controller VM	Enter the default Admin username and password that were applied to controller VMs when they were installed on the existing HX Cluster. Note The name of the controller VM cannot be changed. Use the existing cluster password.

Step 8

Click **Start**. A **Progress** page displays the progress of various configuration tasks.

Note By default no user intervention is required if you are booting from FlexFlash (SD Card). However, if you are setting up your compute-only node to boot from a local disk, complete the following steps in Cisco UCS Manager:

- Click the service profile created by the HX Data Platform Installer.
For example, *blade-1(HX_Cluster_Name)*.
- On the **General** tab, click **Unbind from the Template**.
- In the working pane, click the **Storage** tab. Click the **Local Disk Configuration Policy** sub tab.
- In the Actions area, select **Change Local Disk Configuration Policy > Create Local Disk Configuration Policy**.
- Under **Create Local Disk Configuration Policy**, enter a name for the policy, and keep the rest as default. Click **Ok**.
- In the **Change Local Disk Configuration Policy** Actions area, select the newly created local disk configuration policy from the drop-down list. Click **Ok**.

- g. Now, go back to the HX Data Platform Installer UI and click **Continue**, and then click **Retry UCSM**

Compute Node Expansion - ESXi Installation Required

ESXi must be installed on all nodes being added at this point using the HX ESXi ISO on [cisco.com](https://www.cisco.com)

Using an existing installation of ESXi will cause installation to fail. Other ESXi ISOs other than the one posted on Cisco are not supported.

Once ESXi is installed, select Continue and then Retry to continue installation.
Full instructions can be found below.

If ESXi is already installed using the HX ESXi ISO wait for it to boot and then select Continue and Retry to continue installation.

 Instructions

 Launch UCS Manager

Continue

Configuration.

Note If the vCenter cluster has EVC enabled, the deploy process fails. The host needs to be manually added to vCenter. To successfully perform the deploy action, do the following:

- Log in to the ESXi host to be added in vSphere Client.
- Power off the controller VM.
- Add the host to the vCenter cluster in vSphere Web Client.
- In the HX installer, click **Retry Deploy**.

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

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

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

Step 11 After adding compute-only nodes to an existing cluster, you must manually configure the vmk2 interface for vmotion.

Resolving Failure of Cluster Expansion

If you receive an error dialog box and the storage cluster expansion doesn't complete, proceed with the resolution options listed below:

Step 1 **Edit Configuration** - Returns you to the Cluster Configuration page. You fix the issues listed in the validation page.

- Step 2** **Start Over** - Allows you to reverse the settings you applied by clearing progress table entries and you are returned to the Cluster Configuration page to restart a new deployment. See Technical Assistance Center (TAC).
- Step 3** **Continue** - Adds the node to the storage cluster in spite of the failure generating errors. See Technical Assistance Center (TAC).

Note Select the Continue button only if you understand the failures and are willing to accept the possibility of unpredictable behavior.

For more information about cleaning up a node for the purposes of redeploying HyperFlex, see the [HyperFlex Customer Cleanup Guides for FI and Edge](#).

Logical Availability Zones

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

Advantages of Logical Availability Zones

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

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

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

Specifying the Number of Zones and Optimizing Balance

The number of zones is set automatically by default and recommended. When you let the installer decide the number of zones, the number of zones is decided based on the number of nodes in the cluster.

To maintain the most balanced consumption of space and data distribution, it is recommended that the number of nodes in a cluster are whole multiple of number of zones, which is either 3, 4, or 5. For example, 8 nodes would evenly divide into 4 zones of 2 servers each, and 9 nodes would divide evenly into 3 zones of 3 servers each. Eleven nodes would create an unbalanced number of nodes across the zones, leading to unbalanced space consumption on the nodes. Users with a need may manually specify 3, 4 or 5 zones.

LAZ Guidelines and Considerations

- HyperFlex clusters determine which nodes participate in each zone. This configuration cannot be modified.
- When changing the number of resources, add or remove an equal number of resources from each configured zone.
- **Zone Expansion:** Perform expansions in the unit of number of zones to avoid an imbalance which could lead to non-optimal performance. For example, if your cluster has 4 zones, then it is recommended to add 4 nodes during expansion (1 node per zone is added automatically when 4 nodes are added to the cluster).
- **Imbalanced Zones:** Zones may become imbalanced due to unbalanced expansion, permanent failure of nodes from zone(s), or unbalanced nodes during install.
- **Disabling and Re-enabling LAZ:** You can disable and enable LAZ dynamically. It is not recommended to disable and re-enable LAZ in the same cluster with a different number of zones. Doing so could result in an excessive amount of movement and reorganization of data across the cluster - to comply with existing data distribution rules if LAZ is turned on in a cluster already containing data. This can result in the cluster becoming no longer zone compliant for example, if the cluster usage is already greater than 25%.

Viewing LAZ Status and Connections

- To view LAZ information from the HX Connect dashboard, log in to HX Connect and use the **System information** and **HyperFlex Connect > Dashboard** menu.
- You can also view LAZ details through CLI by running the `stcli cluster get-zone` command. The following is sample output from the `stcli cluster get-zone` command:

```
stcli cluster get-zone

zones:
-----
pNodes:
-----
state: ready
name: 10.10.18.61
-----
state: ready
name: 10.10.18.59
-----
zoneId: 0000000057eebaab:0000000000000003
numNodes: 2
-----
pNodes:
-----
state: ready
name: 10.10.18.64
-----
state: ready
name: 10.10.18.65
-----
zoneId: 0000000057eebaab:0000000000000001
numNodes: 2
-----
pNodes:
-----
state: ready
name: 10.10.18.60
```

```

-----
state: ready
name: 10.10.18.63
-----
zoneId: 0000000057eebaab:0000000000000004
numNodes: 2
-----
pNodes:
-----
state: ready
name: 10.10.18.58
-----
state: ready
name: 10.10.18.62
-----
zoneId: 0000000057eebaab:0000000000000002
numNodes: 2
-----
isClusterZoneCompliant: True
zoneType: logical
isZoneEnabled: True
numZones: 4
AboutCluster Time : 08/22/2019 2:31:39 PM PDT

```

LAZ Related Commands

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

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

Command	Description
stcli cluster get-zone	Gets the zone details. This option is used to check if the zone is enabled.
stcli cluster set-zone --zone 0	Enables or Disables zones.
stcli cluster set-zone --zone 1 stcli rebalance start	<p>(Recommended) Enables and creates zones (default number of zones)</p> <p>Important You must execute the rebalance start command after you enable and create zones.</p> <p>A cluster created without zoning enabled, will become zone compliant only after enabling zoning and successful completion of rebalance.</p> <p>Warning Rebalance is a critical background service. Disabling the service may lead to unexpected behavior including loss of cluster resiliency. Support for this command is limited to Cisco Tech support only. General use is not supported.</p> <p>Triggering rebalance activity may involve large scale data movements across several nodes in the cluster which may decrease the IO performance in the cluster.</p>

Command	Description
<pre>stcli cluster set-zone --zone 1 --numzones <integer-value> stcli rebalance start</pre>	<p>Enables zones and creates a specific number of zones.</p> <p>Important The number of zones can only be 3, 4, or 5.</p> <p>Important You must execute the rebalance start command after you enable and create zones.</p> <p>Warning Rebalance is a critical background service. Disabling the service may lead to unexpected behavior including loss of cluster resiliency. Support for this command is limited to Cisco Tech support only. General use is not supported.</p>

Expanding a Cluster with Fewer Nodes than Zones

Maintaining balance across your clusters is critical for performance. Expanding a cluster with fewer nodes than zones creates an environment that is by design out of balance. It is recommended that the following be used at the direction of TAC.

To expand the cluster, perform the following steps.

-
- Step 1** Review the guidelines and limitations.
- Step 2** Add nodes to the cluster.
- Step 3** Review the cluster configuration and determine if the cluster is zone compliant.
- ```
stcli cluster get-zone
```
- 

For example, if you have a current Logical Availability Zone (LAZ) with 3 zones and 11 nodes, you could expand it to include an additional node (Final LAZ). The Final LAZ with new node 12 will be part of the zone with a fewer number of nodes.

**Table 4: Current LAZ**

| Zone   | Nodes         |
|--------|---------------|
| Zone 1 | pnode 1,2,3,4 |
| Zone 2 | pnode 5,6,7,8 |
| Zone 3 | pnode 9,10,11 |

**Table 5: Final LAZ**

| Zone   | Nodes         |
|--------|---------------|
| Zone 1 | pnode 1,2,3,4 |

| Zone   | Nodes                    |
|--------|--------------------------|
| Zone 2 | pnode 5,6,7,8            |
| Zone 3 | pnode 9,10,11, <b>12</b> |