



## Installation Prerequisites

---

- [Recommended FI/Server Firmware - 3.5\(x\) Releases, on page 1](#)
- [Required Hardware Cables, on page 3](#)
- [Host Requirements, on page 4](#)
- [Disk Requirements, on page 4](#)
- [Browser Recommendations - 3.5\(x\) Releases, on page 6](#)
- [Port Requirements, on page 7](#)
- [HyperFlex External Connections, on page 8](#)
- [Fabric Interconnect Uplink Provisioning, on page 10](#)
- [Network Settings, on page 13](#)
- [VLAN and vSwitch Requirements, on page 14](#)
- [Cisco UCS Requirements, on page 15](#)
- [Hypervisor Requirements, on page 16](#)
- [Storage Cluster Requirements, on page 16](#)
- [vCenter Configuration Requirements, on page 18](#)
- [System Services Requirements, on page 19](#)
- [CPU Resource Reservation for Controller VMs, on page 21](#)
- [Memory Resource Reservation for Controller VMs, on page 21](#)
- [Auto Support Requirements, on page 22](#)
- [Single Sign On Requirements, on page 23](#)

## Recommended FI/Server Firmware - 3.5(x) Releases

The HX components—Cisco HX Data Platform Installer, Cisco HX Data Platform, and Cisco UCS firmware—are installed on different servers. Verify that each component on each server used with and within an HX Storage Cluster are compatible.

- HyperFlex does not support UCS Manager and UCS Server Firmware versions 4.0(4a), 4.0(4b), and 4.0(4c).



---

**Important** Do not upgrade to these versions of firmware.  
Do not upgrade to these versions of UCS Manager.

---

- Verify that the preconfigured HX servers have the same version of Cisco UCS server firmware installed. If the Cisco UCS Fabric Interconnects (FI) firmware versions are different, see the *Cisco HyperFlex Systems Upgrade Guide* for steps to align the firmware versions.
  - **M4:** For NEW hybrid or All Flash (Cisco HyperFlex HX240c M4 or HX220c M4) deployments, verify that Cisco UCS Manager 3.1(3k), 3.2(3i), or 4.0(2d) is installed.
  - **M5:** For NEW hybrid or All Flash (Cisco HyperFlex HX240c M5 or HX220c M5) deployments, verify that the recommended UCS firmware version is installed.




---

**Important** If you are upgrading Cisco UCS Manager 4.0(2a) or 4.0(2b) in the presence of more than one Nvidia GPUs, please remove the GPUs, perform the upgrade and reinstall. For more details, see [CSCvo13678](#).

---




---

**Important** For SED-based HyperFlex systems, ensure that the A (Infrastructure), B (Blade server) and C (Rack server) bundles are at Cisco UCS Manager version 4.0(2b) or later for all SED M4/M5 systems. For more details, see [CSCvh04307](#). For SED-based HyperFlex systems, also ensure that all clusters are at HyperFlex Release 3.5(2b) or later. For more information, see [Field Notice \(70234\)](#) and [CSCvk17250](#).

---

- To reinstall an HX server, download supported and compatible versions of the software. See the [Cisco HyperFlex Systems Installation Guide for VMware ESXi, Release 3.5](#) for the requirements and steps.

**Table 1: HyperFlex Software Versions for M4/M5 Servers**

HyperFlex Release	M4 Recommended FI/Server Firmware *(be sure to review important notes above)	M5 Recommended FI/Server Firmware *(be sure to review important notes above)	M4/M5 Qualified FI/Server Firmware *(be sure to review important notes above)
3.5(2i)	4.0(4k)	4.0(4k)	4.0(4k), 4.0(4l), 4.1(1d), 4.1(1e), 4.1(2a)*, 4.1(2c), 4.1(3b), 4.1(3c)
3.5(2h)	4.0(4k)	4.0(4k)	4.0(4k), 4.0(4l), 4.1(1d), 4.1(1e), 4.1(2a)*, 4.1(2c), 4.1(3b)
3.5(2g)	4.0(4k)	4.0(4k)	4.0(4h), 4.1(1d), 4.1(1e), 4.1(3b)
3.5(2f)	4.0(4e)	4.0(4e)	4.0(4d) <sup>1</sup> , 4.0(4e) <sup>2</sup> , 4.1(3b)
3.5(2e)	4.0(4e)	4.0(4e)	4.0(4g), 4.1(3b)
3.5(2d)	4.0(4e)	4.0(4e)	4.1(3b)

HyperFlex Release	M4 Recommended FI/Server Firmware <i>*(be sure to review important notes above)</i>	M5 Recommended FI/Server Firmware <i>*(be sure to review important notes above)</i>	M4/M5 Qualified FI/Server Firmware <i>*(be sure to review important notes above)</i>
3.5(2c)	Release Deferred		
3.5(2b)	4.0(2d), 3.2(3i), 3.1(3k)	4.0(2d)	4.1(3b)
3.5(2a)	4.0(1c), 3.2(3i), 3.1(3k)	4.0(1c)	4.1(3b)
3.5(1a) - Unsupported	4.0(1b), 3.2(3h), 3.1(3j)	4.0(1a)	

<sup>1</sup> 4.0(4d) qualified only for M5.

<sup>2</sup> 4.0(4e) qualified only for M5.

\*UCS Server Firmware 4.1(2a) is not supported on clusters with self-encrypting drives (SED). For more information, see [CSCvv69704](#).



#### Important

If your cluster is connected to a Fabric Interconnect 6400 series using VIC 1455/1457 with SFP-H25G-CU3M or SFP-H25G-CU5M cables, only use UCS Release 4.0(4k) and later, or 4.1(2a) and later. Do not use the any other UCS version listed in the table of qualified releases. Using a UCS Release that is not UCS Release 4.0(4k) and later, or 4.1(2a) and later may cause cluster outages.

For more information, see the [Release Notes for UCS Manager, Firmware/Drivers, and Blade BIOS](#) for any UCS issues that affect your environment and [CSCvu25233](#).

**NOTE:** If your current server firmware version is not on the recommendation list above, follow the upgrade procedure in the [Cisco HyperFlex Systems Upgrade Guide for VMware ESXi, Known Issues](#) chapter.

## Required Hardware Cables

- Use at least two 10-Gb Small Form-Factor Pluggable (SFP) cables per server when using the 6200 series FI.

Use at least two 40-GbE QSFP cables per server when using the 6300 series FI.

- Ensure that the Fabric Interconnect console cable (CAB-CONSOLE-RJ45) has an RJ-45 connector on one end and a DB9 connector on the other. This cable is used to connect into the RS-232 console connection on a laptop.
- Ensure that the standard power cords have an IEC C13 connector on the end that plugs into the power supplies. Make sure that the optional jumper power cords have an IEC C13 connector on the end that plugs into the power supplies and an IEC C14 connector on the end that plugs into an IEC C13 outlet receptacle.

For further details, see the [Cisco UCS 6300 Series Fabric Interconnect Hardware Guide](#).

- The KVM cable provides a connection for the Cisco HX-Series Servers into the system. It has a DB9 serial connector, a VGA connector for a monitor, and dual USB 2.0 ports for a keyboard and mouse.

With this cable, you can create a direct connection to the operating system and the BIOS running on the system.



---

**Note** This same KVM cable is used for both UCS rack mount and blade servers.

---

For further details on cables and ordering information for M series servers, see the respective [Cisco HyperFlex HX-Series Models](#) and [Cisco UCS B200 Blade Server Installation and Service Note](#).

## Host Requirements

A Cisco HyperFlex cluster contains a minimum of three converged HyperFlex nodes. There is an option of adding compute-only nodes to provide additional compute power if there is no need for extra storage. Each server in a HyperFlex cluster is also referred as a HyperFlex node. Make sure that each node has the following settings installed and configured before you deploy the storage cluster.

For further information, refer to the [Cisco HX240c/220c HyperFlex Node Installation Guides](#).

Ensure that the following host requirements are met.

- Use the same VLAN IDs for all the servers (node or hosts) in the cluster.
- Use the same administrator login credentials for all the ESXi servers across the storage cluster.
- Keep SSH enabled on all ESXi hosts.
- Configure DNS and NTP on all servers.
- Install and configure VMware vSphere.
- **VIC and NIC Support:** For details, see the [Cisco HyperFlex Systems—Networking Topologies](#) document.

## Disk Requirements

The disk requirements vary between converged nodes and compute-only nodes. To increase the available CPU and memory capacity, you can expand the existing cluster with compute-only nodes as needed. These compute-only nodes provide no increase to storage performance or storage capacity.

Alternatively, adding converged nodes increase storage performance and storage capacity alongside CPU and memory resources.

Servers with only Solid-State Disks (SSDs) are All-Flash servers. Servers with both SSDs and Hard Disk Drives (HDDs) are hybrid servers.

The following applies to all the disks in a HyperFlex cluster:

- All the disks in the storage cluster must have the same amount of storage capacity. All the nodes in the storage cluster must have the same number of disks.
- All **SSDs** must support TRIM and have TRIM enabled.
- All **HDDs** can be either SATA or SAS type. All SAS disks in the storage cluster must be in a pass-through mode.

- Disk partitions must be removed from SSDs and HDDs. Disks with partitions are ignored and not added to your HX storage cluster.
- Optionally, you can remove or backup existing data on disks. All existing data on a provided disk is overwritten.




---

**Note** New factory servers are shipped with appropriate disk partition settings. Do not remove disk partitions from new factory servers.

---

- Only the disks ordered directly from Cisco are supported.
- On servers with Self Encrypting Drives (SED), both the cache and persistent storage (capacity) drives must be SED capable. These servers support Data at Rest Encryption (DARE).
- In the event you see an error about unsupported drives or catalog upgrade, see the [Compatibility Catalog](#).

In addition to the disks listed in the table below, all M4 converged nodes have 2 x 64-GB SD FlexFlash cards in a mirrored configuration with ESX installed. All M5 converged nodes have M.2 SATA SSD with ESXi installed.



- 
- Note** Do not mix storage disks type or storage size on a server or across the storage cluster. Mixing storage disk types is not supported.
- When replacing cache or persistent disks, always use the same type and size as the original disk.
  - Do not mix any of the persistent drives. Use all HDD or SSD and the same size drives in a server.
  - Do not mix hybrid and All-Flash cache drive types. Use the hybrid cache device on hybrid servers and All-Flash cache devices on All-Flash servers.
  - Do not mix encrypted and non-encrypted drive types. Use SED hybrid or SED All-Flash drives. On SED servers, both the cache and persistent drives must be SED type.
  - All nodes must use same size and quantity of SSDs. Do not mix SSD types.
- 

Please refer to the corresponding server model spec sheet for details of drives capacities and number of drives supported on the different servers.

For information on compatible PIDs when performing an expansion of existing cluster, please refer to the [Cisco HyperFlex Drive Compatibility](#) document.

### Compute-Only Nodes

The following table lists the supported compute-only node configurations for compute-only functions. Storage on compute-only nodes is not included in the cache or capacity of storage clusters.



- 
- Note** When adding compute nodes to your HyperFlex cluster, the compute-only service profile template automatically configures it for booting from an SD card. If you are using another form of boot media, update the local disk configuration policy. See the *Cisco UCS Manager Server Management Guide* for server-related policies.
-

Supported Compute-Only Node Servers	Supported Methods for Booting ESXi
<ul style="list-style-type: none"> <li>• Cisco B200 M4/M5</li> <li>• B260 M4</li> <li>• B420 M4</li> <li>• B460 M4</li> <li>• C240 M4/M5</li> <li>• C220 M4/M5</li> <li>• C460 M4</li> <li>• C480 M5</li> <li>• B480 M5</li> </ul>	<p>Choose any method.</p> <p><b>Important</b> Ensure that only one form of boot media is exposed to the server for ESXi installation. Post install, you may add in additional local or remote disks.</p> <p>USB boot is not supported for HX Compute-only nodes.</p> <ul style="list-style-type: none"> <li>• SD Cards in a mirrored configuration with ESXi installed.</li> <li>• Local drive HDD or SSD.</li> <li>• SAN boot.</li> <li>• M.2 SATA SSD Drive.</li> </ul> <p><b>Note</b> HW RAID M.2 (UCS-M2-HWRAID and HX-M2-HWRAID) is not supported on Compute-only nodes.</p>

## Browser Recommendations - 3.5(x) Releases

Use one of the following browsers to run the listed HyperFlex components. These browsers have been tested and approved. Other browsers might work, but full functionality has not been tested and confirmed.

**Table 2: Supported Browsers**

Browser	Cisco UCS Manager	HX Data Platform Installer	HX Connect
Microsoft Internet Explorer	9 or higher	11 or higher	11 or higher
Google Chrome	14 or higher	56 or higher	56 or higher
Mozilla Firefox	7 or higher	52 or higher	52 or higher

### Notes

- **Cisco HyperFlex Connect:**

The minimum recommended resolution is 1024 X 768.

- **Cisco HX Data Platform Plug-In:**

The **Cisco HX Data Platform Plug-In** runs in vSphere. For VMware Host Client System browser requirements, see the [VMware documentation](#).

The **Cisco HX Data Platform Plug-In** is not displayed in the vCenter HTML client. You must use the vCenter flash client.

- **Cisco UCS Manager:**

The browser must support the following:

- Java Runtime Environment 1.6 or later.
- Adobe Flash Player 10 or higher is required for some features.

For the latest browser information about Cisco UCS Manager, refer to the most recent [Cisco UCS Manager Getting Started Guide](#).

## Port Requirements

If your network is behind a firewall, in addition to the standard port requirements, VMware recommends ports for VMware ESXi and VMware vCenter.

- CIP-M is for the cluster management IP.
- SCVM is the management IP for the controller VM.
- ESXi is the management IP for the hypervisor.

The comprehensive list of ports required for component communication for the HyperFlex solution is located in Appendix A of the [HX Data Platform Security Hardening Guide](#)



---

**Tip** If you do not have standard configurations and need different port settings, refer to [Table C-5 Port Literal Values](#) for customizing your environment.

---

## HyperFlex External Connections

External Connection	Description	IP Address/ FQDN/ Ports/Version	Essential Information
Intersight Device Connector	Supported HX systems are connected to Cisco Intersight through a device connector that is embedded in the management controller of each system.	HTTPS Port Number: 443 1.0.5-2084 or later (Auto-upgraded by Cisco Intersight)	



External Connection	Description	IP Address/ FQDN/ Ports/Version	Essential Information
			<p>All device connectors must properly resolve <code>svc.intersight.com</code> and allow outbound-initiated HTTPS connections on port 443. The current HX Installer supports the use of an HTTP proxy.</p> <p>The IP addresses of ESXi management must be reachable from Cisco UCS Manager over all the ports that are listed as being needed from installer to ESXi management, to ensure deployment of ESXi management from Cisco Intersight.</p> <p><b>Note</b> Outbound HTTPS connections on port 443 initiated by ESXi hosts can be blocked by the default ESXi firewall. The ESXi firewall can be temporarily disabled to allow this connectivity.</p> <p>To disable the ESXi firewall, use the <code>esxcli network firewall set --enabled=false</code> command and after the installation has completed use the <code>esxcli network firewall set --enabled=false</code> command to re-enable the firewall.</p>

External Connection	Description	IP Address/ FQDN/ Ports/Version	Essential Information
			For more information, see the <a href="#">Network Connectivity Requirements</a> section of the Intersight Help Center.
Auto Support	Auto Support (ASUP) is the alert notification service provided through HX Data Platform.	SMTP Port Number: 25	Enabling Auto Support is strongly recommended because it provides historical hardware counters that are valuable in diagnosing future hardware issues, such as a drive failure for a node.

## Fabric Interconnect Uplink Provisioning

Prior to setting up the HyperFlex cluster, plan the upstream bandwidth capacity for optimal network traffic management. This ensures that the flow is in steady state, even if there is a component failure or a partial network outage.

By default, the *hx-vm-network* vSwitch is configured as **active/active**. All other vSwitches are configured as **active/standby**.

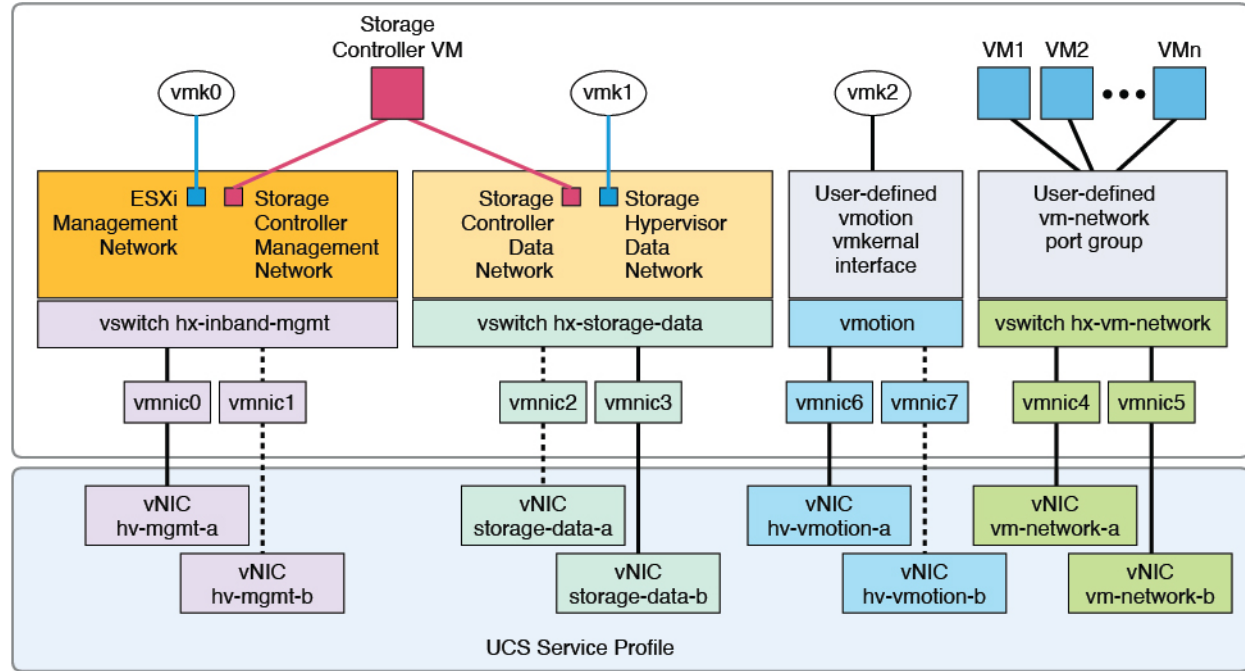



---

**Note** For clusters running Catalyst switches upstream to the FI's, set the best effort Quality of Service (QOS) MTU to 9216 (located in LAN > LAN Cloud > QoS System Class), otherwise failover will fail.

---


Figure 1: HyperFlex Data Platform Connectivity for a Single Host



- Note: 1. Dotted lines represent a “standby” link.  
 2. All “a” vNICs connect to FI-A.  
 3. All “b” vNICs connect to FI-B.  
 4. MTU of 9000 is needed for storage-data and vmotion networks.  
 5. All VLANs by default are tagged on the FI so frames are passed untagged to each vswitch.  
 6. The vm network port groups are automatically created in 1.8 installer with vlan suffix.


Set the default vSwitch NIC teaming policy and failover policy to **yes** to ensure that all management, vMotion, and storage traffic are locally forwarded to the fabric interconnects to keep the flow in steady state. When vNIC-a fails, ESXi computes the load balancing and all the virtual ports are repinned to vNIC-b. When vNIC-a comes back online, repinning does apply and virtual ports are rebalanced across vNIC-a and vNIC-b. This reduces the latency and bandwidth utilization upstream of the Cisco UCS fabric interconnects.

Figure 2: Traffic Flow in Steady State



In case one or more server links fail, for instance, if Host 1 loses connectivity to Fabric A while Host 2 loses connectivity to Fabric B, the traffic must go through the upstream switches. Therefore, the uplink network bandwidth usage increases, and you must add more uplinks.

Figure 3: Traffic Flow During Link Failure





---

**Note** When you have uplinks from a fabric interconnect to two different upstream switches, you encounter a condition called **Disjoint Layer 2** (DJL2) on the FI. This is known to happen on the FI on End Host Mode and if the DJL2 is not configured properly.

To deploy the DJL2 properly, refer to the [Cisco UCS 6300 Series Fabric Interconnect Hardware Guide—Deploy Layer 2 Disjoint Networks Upstream in End Host Mode](#) white paper.

---

## Network Settings



---

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

---

### Best Practices

- Must use different subnets and VLANs for each network.
- Directly attach each host to a Cisco UCS fabric interconnect using a 10-Gbps cable.
- Do not use VLAN 1 which is the default VLAN as it can cause networking issues, especially if Disjoint Layer 2 configuration is used.
- Installer sets the VLANs as *non-native* by default. Ensure to configure the upstream switches to accommodate the non-native VLANs.
- Uplinks from the UCS Fabric Interconnects to all top of rack switch ports must configure spanning tree in **edge trunk** or **portfast edge** mode depending on the vendor and model of the switch. This extra configuration ensures that when links flap or change state, they do not transition through unnecessary spanning tree states and incur an extra delay before traffic forwarding begins. Failure to properly configure FI uplinks in **portfast edge** mode may result in network and cluster outages during failure scenarios and during infrastructure upgrades that leverage the highly available network design native to HyperFlex.
- FI facing ports need to have Port-fast, spanning-tree port type edge trunk, or similar spanning tree configuration that immediately put ports into forwarding mode.

Each ESXi host needs the following networks.

- **Management traffic network**—From the vCenter, handles the hypervisor (ESXi server) management, and storage cluster management.
- **Data traffic network**—Handles the hypervisor and storage data traffic.
- **vMotion network**
- **VM network**

There are four vSwitches, each carrying a different network.

- **vswitch-hx-inband-mgmt**—Used for ESXi management and storage controller management.
- **vswitch-hx-storage-data**—Used for ESXi storage data and HX Data Platform replication.

These two vSwitches are further divided in two port groups with assigned static IP addresses to handle traffic between the storage cluster and the ESXi host.

- **vswitch-hx-vmotion**—Used for VM and storage vMotion.

This vSwitch, has one port group for management, defined through vSphere that connects to all the hosts in the vCenter cluster.

- **vswitch-hx-vm-network**—Used for VM data traffic.

You can add or remove VLANs on the corresponding vNIC templates in Cisco UCS Manager. See [Managing VLANs in Cisco UCS Manager](#) and [Managing vNIC templates in Cisco UCS Manager](#) for the detailed steps. To create port groups on the vSwitch, refer to [Adding Virtual Port Groups to VMware Standard vSwitch](#).



- Note**
1. The Cisco HX Data Platform Installer automatically creates the vSwitches.
  2. The following services in vSphere must be enabled after the HyperFlex storage cluster is created.
    - DRS (Optional, if licensed)
    - vMotion
    - High Availability

## VLAN and vSwitch Requirements

Provide at least three VLAN IDs. All VLANs must be configured on the fabric interconnects during the installation.

VLAN Type	Description
<b>Note</b>	Must use different subnets and VLANs for each of the following networks.
VLAN ESXi and HyperFlex Management Traffic	VLAN Name: <user-defined> (for example "hx-inband-mgmt") VLAN ID: <user-defined>
VLAN HyperFlex Storage Data	VLAN Name: <user-defined> (for example, "hx-storage-data") VLAN ID: <user-defined>
VLAN VM vMotion	VLAN Name: <user-defined> (for example, "hx-vmotion") VLAN ID: <user-defined>

VLAN Type	Description
VLAN VM Network	VLAN VM Network: <user-defined> (for example, "hx-vm-network") VLAN ID: <user-defined>

The VLAN tagging with External Switch VLAN Tagging (EST) and vSwitch settings are applied using UCS Manager profiles. The HX Data Platform Installer, simplifies this process.


**Note**

- Do not use VLAN 1 which is the default VLAN as it can cause networking issues, especially if Disjoint Layer 2 configuration is used. Use a different VLAN other than VLAN 1.  
Installer sets the VLANs as *non-native* by default. Configure the upstream switches to accommodate the non-native VLANs.
- Inband Management is not supported on VLAN 2 or VLAN 3.

## Cisco UCS Requirements

Provide the listed content for the UCS Fabric Interconnect and UCS Manager when prompted.

### Cisco UCS Fabric Interconnect Requirements

UI Element	Essential Information
Uplink Switch Model	Provide the switch type and connection type (SFP + Twin Ax or Optic).
Fabric Interconnect Cluster IP address	<IP address>.
FI-A IP Address	<IP address>.
FI-B IP Address	<IP address>.
MAC Address Pool	Check 00:00:00 MAC address pool.
IP Blocks	KVM IP pool. A minimum of 4 IP addresses.
Subnet mask	For example, 255.255.0.0.
Default Gateway	For example, 10.193.0.1.

### Cisco UCS Manager Requirements

UI Element	Essential Information
UCS Manager Host Name	Hostname or IP address.
User Name	<admin username>

UI Element	Essential Information
Password	<admin username>

## Hypervisor Requirements

Enter the IP address from the range of addresses that are available to the ESXi servers on the storage management network or storage data network through vCenter. Provide static IP addresses for all network addresses.



### Note

- Data and Management networks must be on different subnets.
- IP addresses cannot be changed after the storage cluster is created. Contact Cisco TAC for assistance.
- Though, not required by itself, if you are specifying DNS names, enable IP addresses forward and reverse DNS lookup.
- The installer IP address must be reachable from the management subnet used by the hypervisor and the storage controller VMs. The installer appliance must run on the ESXi host or on a VMware workstation that is not a part of the cluster to be installed.

Management Network IP Addresses		Data Network IP Addresses	
Hypervisor	Storage Controller	Hypervisor	Storage Controller
<IP Address >	<IP Address >	<IP Address >	<IP Address >
<IP Address >	<IP Address >	<IP Address >	<IP Address >
<IP Address >	<IP Address >	<IP Address >	<IP Address >
<IP Address >	<IP Address >	<IP Address >	<IP Address >
VLAN Tag	VLAN_ID	VLAN Tag	VLAN_ID
Subnet Mask		Subnet Mask	
Default Gateway		Default Gateway	
Installer Appliance IP Addresses			
<IP Address >		<IP Address >	

## Storage Cluster Requirements

Storage cluster is a component of the Cisco HX Data Platform which reduces storage complexity by providing a single datastore that is easily provisioned in the vSphere Web Client. Data is fully distributed across disks in all the servers that are in the storage cluster, to leverage controller resources and provide high availability.



A storage cluster is independent of the associated vCenter cluster. You can create a storage cluster using ESXi hosts that are in the vCenter cluster.

To define the storage cluster, provide the following parameters.

Field	Description
<b>Name</b>	Enter a name for the storage cluster.
<b>Management IP Address</b>	This provides the storage management network, access on each ESXi host. <ul style="list-style-type: none"><li>• The IP address must be on the same subnet as the Management IP addresses for the nodes.</li><li>• Do not allow cluster management IPs to share the last octet with another cluster on the same subnet.</li><li>• These IP addresses are in addition to the four IP addresses we assign to each node in the Hypervisor section.</li></ul>
<b>Storage Cluster Data IP Address</b>	This provides the storage data network and storage controller VM network, access on each ESXi host. The same IP address must be applied to all ESXi nodes in the cluster.

Field	Description
<b>Data Replication Factor</b>	<p>Data Replication Factor defines the number of redundant replicas of your data across the storage cluster.</p> <p>This is set during HX Data Platform installation and cannot be changed.</p> <p>Choose a <b>Data Replication Factor</b>. The choices are:</p> <ul style="list-style-type: none"> <li>• <b>Data Replication Factor 3</b>—A replication factor of three is highly recommended for all environments except HyperFlex Edge. A replication factor of two has a lower level of availability and resiliency. The risk of outage due to component or node failures should be mitigated by having active and regular backups. <ul style="list-style-type: none"> <li><b>Attention</b> This is the recommended option.</li> </ul> </li> <li>• <b>Data Replication Factor 2</b>—Keep two redundant replicas of the data. This consumes less storage resources, but reduces your data protection in the event of simultaneous node or disk failure.</li> </ul> <p>If nodes or disks in the storage cluster fail, the cluster's ability to function is affected. If more than one node fails or one node and disk(s) on a different node fail, it is called a simultaneous failure.</p>

## vCenter Configuration Requirements

Provide administrator level account and password for vCenter. Ensure that you have an existing vCenter server. Ensure that the following vSphere services are operational.

- Enable Dynamic Resource Scheduler (DRS) [Optional, enable if licensed].
- Enable vMotion.
- Enable High availability (HA) [Required to define failover capacity and for expanding the datastore heartbeat].
- User VMs must be version 9 or later [Required to use HX Data Platform, Native Snapshots, and ReadyClones].

Field	Description
<b>vCenter Server</b>	<p>Enter your current vCenter server web address.</p> <p>For example, <i>http://&lt;IP address&gt;</i>.</p>

Field	Description
<b>User Name</b>	Enter <i>&lt;admin username&gt;</i> .
<b>Password</b>	Enter <i>&lt;admin password&gt;</i> .
<b>Datacenter Name</b> <b>Note</b> An existing datacenter object can be used. If the datacenter doesn't exist in vCenter, it will be created.	Enter the required name for the vCenter datacenter.
<b>Cluster Name</b>	Enter the required name for the vCenter cluster. The cluster must contain a minimum of three ESXi servers.

## System Services Requirements

Before installing Cisco HX Data Platform, ensure that the following network connections and services are operational.

- DNS server




---

**Caution** DNS servers should reside outside of the HX storage cluster. Nested DNS servers can cause a cluster to not start after entire cluster is shutdown, such as during DC power loss.

---

- NTP server




---

**Caution** NTP servers should reside outside of the HX storage cluster. Nested NTP servers can cause a cluster to not start after entire cluster is shutdown, such as during DC power loss.

---



- Note**
- Before configuring the storage cluster, manually verify that the NTP server is working and providing a reliable source for the time.
  - Use the same NTP server for all nodes (both converged and compute) and all storage controller VMs.
  - The NTP server must be stable, continuous (for the lifetime of the cluster), and reachable through a static IP address.
  - If you are using Active Directory as an NTP server, please make sure that the NTP server is setup according to Microsoft best practices. For more information, see [Windows Time Service Tools and Settings](#). Please note that if the NTP server is not set correctly, time sync may not work, and you may need to fix the time sync on the client-side. For more information, see [Synchronizing ESXi/ESX time with a Microsoft Domain Controller](#).

- Time Zone

Field	Essential Information
DNS Server(s)	<p data-bbox="659 926 808 951">&lt;IP address&gt;</p> <p data-bbox="659 972 1479 1031">DNS server address is required if you are using hostnames while installing the HyperFlex Data Platform.</p> <p data-bbox="659 1052 1479 1142"><b>Note</b></p> <ul data-bbox="800 1052 1479 1255" style="list-style-type: none"> <li>• If you do not have a DNS server, do not enter a hostname under <b>System Services</b> in the <b>Cluster Configuration</b> page of the HX Data Platform Installer. Use only IP addresses.</li> <li>• To provide more than one <i>DNS servers address</i>, separate the address with a comma. Check carefully to ensure that DNS server addresses are entered correctly.</li> </ul>
NTP Server(s) (A reliable NTP server is required)	<p data-bbox="659 1323 808 1348">&lt;IP address&gt;</p> <p data-bbox="659 1369 1240 1394">NTP server is used for clock synchronization between:</p> <ul data-bbox="695 1415 954 1541" style="list-style-type: none"> <li>• Storage controller VM</li> <li>• ESXi hosts</li> <li>• vCenter server</li> </ul> <p data-bbox="659 1583 1479 1673"><b>Important</b> Static IP address for an NTP server is required to ensure clock synchronization between the storage controller VM, ESXi hosts, and vCenter server.</p> <p data-bbox="659 1715 1479 1803">During installation, this information is propagated to all the storage controller VMs and corresponding hosts. The servers are automatically synchronized on storage cluster startup.</p>

Field	Essential Information
Time Zone	<p>&lt;your time zone&gt;</p> <p>Select a time zone for the storage controller VMs. It is used to determine when to take scheduled snapshots.</p> <p><b>Note</b> All the VMs must be in the same time zone.</p>

## CPU Resource Reservation for Controller VMs

As the storage controller VMs provide critical functionality for the HyperFlex Data Platform, the HX Data Platform Installer configures CPU resource reservations for the controller VMs. This reservation guarantees that the controller VMs have the minimum required CPU resources. This is useful in situations where the physical CPU resources of the ESXi hypervisor host are heavily consumed by the guest VMs. The following table details the CPU resource reservation for storage controller VMs.

Product ID	Number of VM CPU	Shares	Reservation	Limit
HXAF220c-M5SN (All NVMe 220)	12	Low	10,800 MHz	Unlimited
With HX Boost Mode enabled: HXAF220c-M5SN	16	Low	10,800 MHz	Unlimited
With HX Boost Mode enabled: HXAF220c-M4/M5 HXAF240c-M4/M5SX	12	Low	10,800 MHz	Unlimited
All Other Models	8	Low	10,800 MHz	Unlimited

## Memory Resource Reservation for Controller VMs

The following table details the memory resource reservations for the storage controller VMs.

Server Model	Amount of Guest Memory	Reserve All Guest Memory
HX220c-M4/M5/M6 HX-E-220M5SX HX-E-220M6S	48 GB	Yes
HXAF220C-M4	48 GB	Yes

Server Model	Amount of Guest Memory	Reserve All Guest Memory
HXAF220C-M5/M6 HXAF-E-220M5SX HXAF-E-220M6SX	48 GB 56 GB for configurations with 7.6 TB SSDs (SED and non-SED)	Yes
HX240c-M4/M5SX/M6SX	72 GB	Yes
HXAF240c-M4/M5SX/M6SX	72 GB 84 GB for configurations with 7.6 TB SSDs (SED and non-SED)	Yes
HX240C-M5L HX240C-M6S	78 GB	Yes

- B200 compute-only blades have a lightweight storage controller VM, it is configured with only 1 vCPU and 512 MB of memory reservation.
- C240 Rack Server delivers outstanding levels of expandability and performance in a two rack-unit (2RU) form-factor.
- C220 Server delivers expandability in a one rack-unit (1RU) form-factor.

## Auto Support Requirements

Auto Support (ASUP) is the alert notification service provided through HX Data Platform. If you enable Auto Support, notifications are sent from HX Data Platform to designated email addresses or email aliases that you want to receive the notifications.

To configure Auto Support, you need the following information:

Auto Support	
<b>Enable Auto Support</b> check box	Check this box during HX storage cluster creation.
<b>Mail Server</b>	<IP address> SMTP mail server must be configured in your network to enable Auto Support. Used for handling email sent from all the storage controller VM IP addresses. <b>Note</b> Only unauthenticated SMTP is supported for ASUP.
<b>Mail Sender</b>	<username@domain.com> Email address to use for sending Auto Support notifications.
<b>ASUP Recipient</b>	List of email addresses or email aliases to receive Auto Support notifications.



**Note** Enabling Auto Support is strongly recommended because it provides historical hardware counters that are valuable in diagnosing future hardware issues, such as drive failure for a node.

## Single Sign On Requirements

The SSO URL is provided by vCenter. If it is not directly reachable from the controller VM, then configure the location explicitly using **Installer Advanced Settings**.

Single Sign On (SSO)	
SSO Server URL	SSO URL can be found in vCenter at <b>vCenter Server &gt; Manage &gt; Advanced Settings</b> , key <code>config.vpxd.sso.sts.uri</code>

