



# Installation Prerequisites

---

This chapter describes the installation and configuration requirements for the Cisco HyperFlex Systems:

- [Required Hardware Cables, on page 1](#)
- [Host Requirements, on page 2](#)
- [Disk Requirements, on page 2](#)
- [Port Requirements, on page 4](#)
- [HyperFlex External Connections, on page 12](#)
- [Fabric Interconnect Uplink Provisioning, on page 13](#)
- [Network Settings, on page 16](#)
- [VLAN and vSwitch Requirements, on page 17](#)
- [Cisco UCS Requirements, on page 18](#)
- [Hypervisor Requirements, on page 18](#)
- [Storage Cluster Requirements, on page 19](#)
- [vCenter Configuration Requirements, on page 21](#)
- [System Services Requirements, on page 21](#)
- [CPU Resource Reservation for Controller VMs, on page 23](#)
- [Memory Resource Reservation for Controller Virtual Machines, on page 23](#)
- [Auto Support Requirements, on page 24](#)
- [Single Sign On Requirements, on page 24](#)

## 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 6200 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 M4 or M5 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.
- Use single VIC only for Converged nodes or Compute-only nodes. Additional VICs or PCIe NICs are not supported.

## 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 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 M3/M4/M5</li> <li>• B260 M4</li> <li>• B420 M4</li> <li>• B460 M4</li> <li>• C240 M3/M4/M5</li> <li>• C220 M3/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>

## 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.

Verify that the following firewall ports are open:

**Time Server**

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
123	NTP/UDP	Each ESXi Node Each SCVM Node UCSM	Time Server	Bidirectional

**HX Data Platform Installer**

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
22	SSH/TCP	HX Data Platform Installer	Each ESXi Node	Management addresses
			Each SCVM Node	Management addresses
			CIP-M	Cluster management
			UCSM	UCSM management addresses
			CIMC IPs	
80	HTTP/TCP	HX Data Platform Installer	Each ESXi Node	Management addresses
			Each SCVM Node	Management addresses
			CIP-M	Cluster management
			UCSM	UCSM management addresses
443	HTTPS/TCP	HX Data Platform Installer	Each ESXi Node	Management addresses
			Each SCVM Node	Management addresses
			CIP-M	Cluster management
			UCSM	UCSM management addresses
8089	vSphere SDK/TCP	HX Data Platform Installer	Each ESXi Node	Management addresses

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
902	Heartbeat/UDP/TCP	HX Data Platform Installer	vCenter	
			Each ESXi Node	
None	Ping/ICMP	HX Data Platform Installer	ESXi IPs CVM IPs	Management addresses
9333	UDP/TCP	HX Data Platform Installer	CIP-M	Cluster management

### Mail Server

Optional for email subscription to cluster events.

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
25	SMTP/TCP	Each SCVM Node CIP-M UCSM	Mail Server	Optional

### Monitoring

Optional for monitoring UCS infrastructure.

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
161	SNMP Poll/UDP	Monitoring Server	UCSM	Optional
162	SNMP Trap/UDP	UCSM	Monitoring Server	Optional

### Name Server

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
53 (external lookups)	DNS/TCP/UDP	Each ESXi Node	Name Server	Management addresses
		Each SCVM Node	Name Server	Management addresses
		CIP-M	Name Server	Cluster management
		UCSM	Name Server	

**vCenter**

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
80	HTTP/TCP	vCenter	Each SCVM Node CIP-M	Bidirectional
443	HTTPS (Plug-in)/TCP	vCenter	Each ESXi Node Each SCVM Node CIP-M	Bidirectional
7444	HTTPS (VC SSO)/TCP	vCenter	Each ESXi Node Each SCVM Node CIP-M	Bidirectional
9443	HTTPS (Plug-in)/TCP	vCenter	Each ESXi Node Each SCVM Node CIP-M	Bidirectional
5989	CIM Server/TCP	vCenter	Each ESXi Node	
9080	CIM Server/TCP	vCenter	Each ESXi Node	Introduced in ESXi Release 6.5
902	Heartbeat/TCP/UDP	vCenter	Each ESXi Node	This port must be accessible from each host. Installation results in errors if the port is not open from the HX Installer to the ESXi hosts.

## User

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
22	SSH/TCP	User	Each ESXi Node	Management addresses
			Each SCVM Node	Management addresses
			CIP-M	Cluster management
			HX Data Platform Installer	
			UCSM	UCSM management addresses
			vCenter	
			SSO Server	
80	HTTP/TCP	User	Each SCVM Node	Management addresses
			CIP-M	Cluster management
			UCSM	
			HX Data Platform Installer	
			vCenter	
443	HTTPS/TCP	User	Each SCVM Node	
			CIP-M	
			UCSM	UCSM management addresses
			HX Data Platform Installer	
			vCenter	
7444	HTTPS (SSO)/TCP	User	vCenter SSO Server	
9443	HTTPS (Plug-in)/TCP	User	vCenter	



**SSO Server**

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
7444	HTTPS (SSO)/TCP	SSO Server	Each ESXi Node Each SCVM Node CIP-M	Bidirectional

**Stretch Witness**

Required only when deploying HyperFlex Stretched Cluster.

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
2181 2888 3888	Zookeeper/TCP	Witness	Each CVM Node	Bidirectional, management addresses
8180	Exhibitor (Zookeeper lifecycle)/TCP	Witness	Each CVM Node	Bidirectional, management addresses
80	HTTP/TCP	Witness	Each CVM Node	Potential future requirement
443	HTTPS/TCP	Witness	Each CVM Node	Potential future requirement

**Replication**

Required only when configuring native HX asynchronous cluster to cluster replication.

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
9338	Data Services Manager Peer/TCP	Each CVM Node	Each CVM Node	Bidirectional, include cluster management IP addresses
3049	Replication for CVM/TCP	Each CVM Node	Each CVM Node	Bidirectional, include cluster management IP addresses
4049	Cluster Map/TCP	Each CVM Node	Each CVM Node	Bidirectional, include cluster management IP addresses

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
4059	NR NFS/TCP	Each CVM Node	Each CVM Node	Bidirectional, include cluster management IP addresses
9098	Replication Service	Each CVM Node	Each CVM Node	Bidirectional, include cluster management IP addresses
8889	NR Master for Coordination/TCP	Each CVM Node	Each CVM Node	Bidirectional, include cluster management IP addresses
9350	Hypervisor Service/TCP	Each CVM Node	Each CVM Node	Bidirectional, include cluster management IP addresses

**SED Cluster**

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
443	HTTPS	Each SCVM Management IP (including cluster management IP)	UCSM (Fabric A, Fabric B, VIP)	Policy Configuration
5696	TLS	CIMC from each node	KVM Server	Key Exchange

**UCSM**

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
443	Encryption etc./TCP	Each CVM Node	CIMC OOB	Bidirectional for each UCS node
81	KVM/HTTP	User	UCSM	OOB KVM
743	KVM/HTTP	User	UCSM	OOB KVM encrypted

**Miscellaneous**

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
9350	Hypervisor Service/TCP	Each CVM Node	Each CVM Node	Bidirectional, include cluster management IP addresses
9097	CIP-M Failover/TCP	Each CVM Node	Each CVM Node	Bidirectional for each CVM to other CVMs
111	RPC Bind/TCP	Each SCVM node	Each SCVM node	CVM outbound to Installer
8002	Installer/TCP	Each SCVM node	Installer	Service Location Protocol
8080	Apache Tomcat/TCP	Each SCVM node	Each SCVM node	stDeploy makes connection, any request with uri /stdeploy
8082	Auth Service/TCP	Each SCVM node	Each SCVM node	Any request with uri /auth/
9335	hxRoboControl/TCP	Each SCVM node	Each SCVM node	Robo deployments
443	HTTPS/TCP	Each CVM Mgmt IP including CIP-M	UCSM A/B and VIP	Policy Configuration
5696	TLS/TCP	CIMC from each node	KMS Server	Key Exchange
8125	UDP	Each SCVM node	Each SCVM node	Graphite
427	UDP	Each SCVM node	Each SCVM node	Service Location Protocol
32768 to 65535	UDP	Each SCVM node	Each SCVM node	SCVM outbound communication



**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)	All device connectors must properly resolve <code>svc.ucsc-connect.com</code> and allow outbound-initiated HTTPS connections on port 443. The current HX Installer supports the use of an HTTP proxy.  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.
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.

External Connection	Description	IP Address/ FQDN/ Ports/Version	Essential Information
Post Installation Script	To complete the post installation tasks, you can run a post installation script on the Installer VM. The script pings across all network interfaces (management, vMotion, and storage network) to ensure full fabric availability. The script also validates the correct tagging of VLANs and jumbo frame configurations on the northbound switch.	HTTP Port Number: 80	The post install script requires name resolution to <a href="http://cs.co/hx-scripts">http://cs.co/hx-scripts</a> via port 80 (HTTP).

## Fabric Interconnect Uplink Provisioning

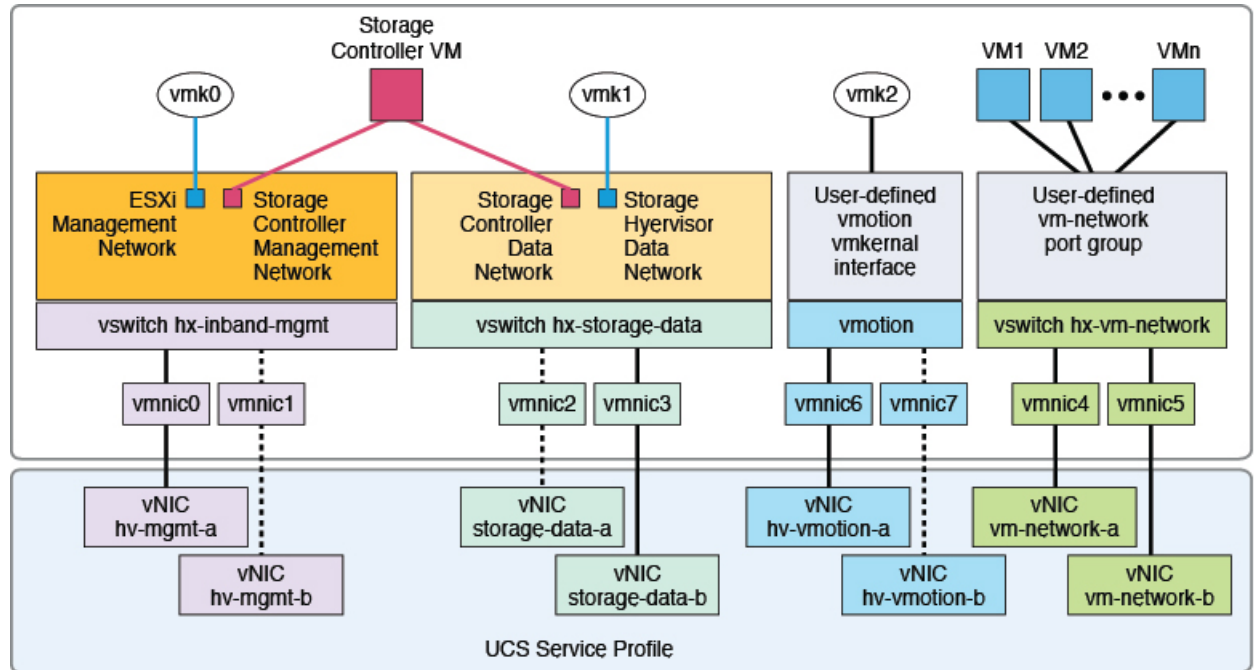
Prior to setting up the HyperFlex cluster, plan the upstream bandwidth capacity for optimal network traffic management. This ensures that the flow is in steady state, even if there is a component failure or a partial network outage.

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.

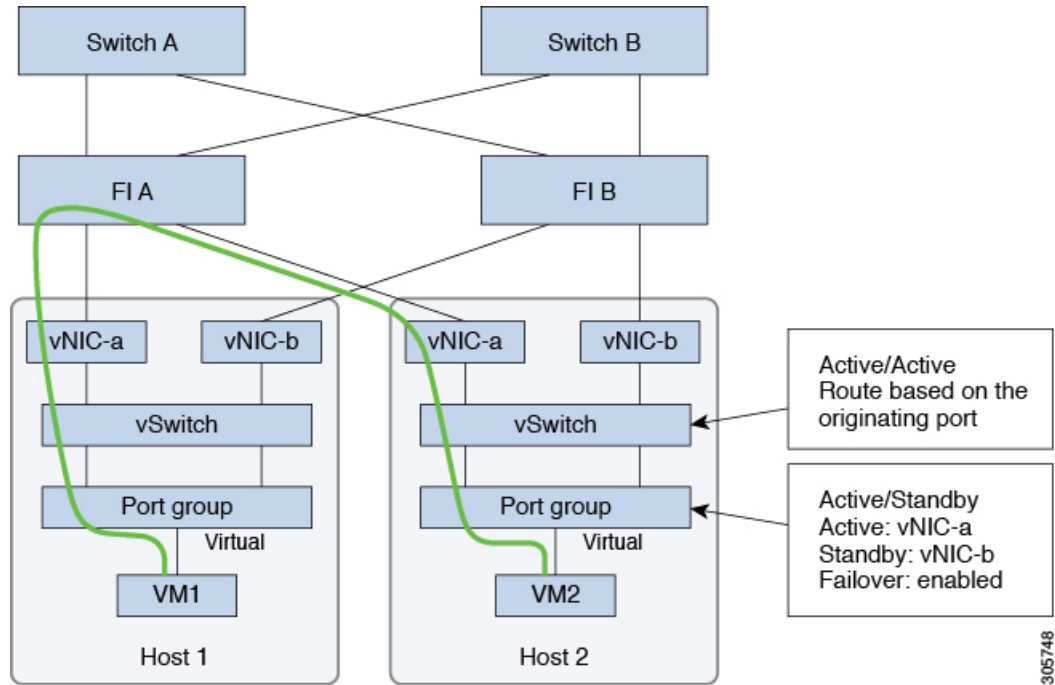
Figure 1: HyperFlex Data Platform Connectivity for a Single Host



- 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 6200 Series Fabric Interconnects—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.

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: hx-inband-mgmt VLAN ID
VLAN HyperFlex Storage Traffic	VLAN Name: hx-storage-data VLAN ID
VLAN VM vMotion	VLAN Name: hx-vmotion VLAN ID
VLAN VM data	User defined
IP Blocks	KVM IP pool One IP addresses per host.
Subnet mask	Ex: 255.255.0.0
Default Gateway	Ex: 10.193.0.1

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>
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 >
<i>VLAN Tag</i>	<i>VLAN_ID</i>	<i>VLAN Tag</i>	<i>VLAN_ID</i>
<i>Subnet Mask</i>		<i>Subnet Mask</i>	
<i>Default Gateway</i>		<i>Default Gateway</i>	
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.

Field	Description
<b>Management IP Address</b>	<p>This provides the storage management network, access on each ESXi host.</p> <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>	<p>This provides the storage data network and storage controller VM network, access on each ESXi host.</p> <p>The same IP address must be applied to all ESXi nodes in the cluster.</p>
<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. <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> </li> </ul>

## 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>	Enter your current vCenter server web address. For example, <i>http://&lt;IP address&gt;</i> .
<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="662 1119 808 1146">&lt;IP address&gt;</p> <p data-bbox="662 1163 1484 1224">DNS server address is required if you are using hostnames while installing the HyperFlex Data Platform.</p> <p data-bbox="662 1245 708 1272"><b>Note</b></p> <ul style="list-style-type: none"> <li data-bbox="802 1245 1479 1339">• 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 data-bbox="802 1360 1463 1455">• 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>

Field	Essential Information
NTP Server(s) (A reliable NTP server is required)	<p>&lt;IP address&gt;</p> <p>NTP server is used for clock synchronization between:</p> <ul style="list-style-type: none"> <li>• Storage controller VM</li> <li>• ESXi hosts</li> <li>• vCenter server</li> </ul> <p><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>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>
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.

Number of VM CPU	Shares	Reservation	Limit
8	Low	10,800 MHz	Unlimited

## Memory Resource Reservation for Controller Virtual Machines

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/M5s HXAF220C-M4/M5s	48 GB	Yes
HX240c-M4/M5SX HXAF240C-M4/M5SX	72 GB	Yes

Server Model	Amount of Guest Memory	Reserve All Guest Memory
HX240C-M5L	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>

