



Installation Prerequisites

- [Supported Versions and System Requirements for Cisco HXDP](#) , on page 1
- [Required Hardware Cables](#), on page 2
- [Non Pre-configured Cisco HyperFlex Systems](#), on page 2
- [Host Requirements](#), on page 3
- [Disk Requirements](#), on page 3
- [Port Requirements](#), on page 5
- [HyperFlex External Connections](#), on page 6
- [Fabric Interconnect Uplink Provisioning](#), on page 8
- [Network Settings](#), on page 11
- [VLAN and vSwitch Requirements](#), on page 12
- [Cisco UCS Requirements](#), on page 13
- [Hypervisor Requirements](#), on page 14
- [Storage Cluster Requirements](#), on page 15
- [vCenter Configuration Requirements](#), on page 16
- [System Services Requirements](#), on page 16
- [CPU Resource Reservation for Controller VMs](#), on page 18
- [Memory Resource Reservation for Controller VMs](#), on page 19
- [Memory Usage for Controller VMs - NVMe Drives](#), on page 21
- [Changing Controller Memory on a Cluster](#), on page 21
- [Auto Support Requirements](#), on page 23
- [Single Sign On Requirements](#), on page 24

Supported Versions and System Requirements for Cisco HXDP

Cisco HX Data Platform requires specific software and hardware versions, and networking settings for successful installation.

Table 1: Supported Versions and System Requirements for Cisco HXDP Release 5.0(x)

Requirement	Link to Details
Confirm the component firmware on the server meets the minimum versions supported.	For more information, see the " FI/Server Firmware - 5.0(x) Releases " topic in the Cisco HyperFlex Software Requirements and Recommendations document.
List of recommended browsers.	For more information, see the " Browser Recommendations " topic in the Cisco HyperFlex Software Requirements and Recommendations document.

Required Hardware Cables

- **6200/6400/6500 series FI:** Use at least two 10-Gb Small Form-Factor Pluggable (SFP) cables per server.
- **6300 series FI:** Use at least two 40-GbE QSFP cables per server.
- 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.

[Non Pre-configured Cisco HyperFlex Systems, on page 2](#)

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](#).

Non Pre-configured Cisco HyperFlex Systems

The Cisco HyperFlex System must have VMware ESXi installed before starting the actual Cisco HyperFlex Installation. In the event your system does not have VMware ESXi preinstalled, perform the tasks in the [Cisco HyperFlex Systems Customized Installation Method](#) chapter of this guide.

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.



Note The root user is created with the same password as the admin user during cluster creation. It is important to track the root user password because future changes to the admin password do not automatically update the root password .

- 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.

- Moving operational disks between servers within same cluster or moving them into expansion nodes within the same active cluster is not supported.
- 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/M6 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"> • Cisco B200 M4/M5 • B260 M4 • B420 M4 • B460 M4 • C240 M4/M5/M6 • C220 M4/M5/M6 • C460 M4 • C480 M5 • B480 M5 	<p>Choose any method.</p> <p>Important 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"> • SD Cards in a mirrored configuration with ESXi installed. • Local drive HDD or SSD. • SAN boot. • M.2 SATA SSD Drive. <p>Note HW RAID M.2 (UCS-M2-HWRAID and HX-M2-HWRAID) is a supported boot configuration starting with HX Data Platform version 4.5(1a) and later.</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.

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 Appendix A of the [HX Data Platform Security Hardening Guide](#) 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>Note 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 Network Connectivity Requirements 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 All VLANs (including storage-data and vmotion) must be configured upstream. It is imperative that storage-data is configured upstream to prevent fail-back timing issues which require temporary upstream connectivity.



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.



Note You cannot use the storage-data VLAN for User VM traffic.

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 need to have Port-fast, spanning-tree port type edge trunk, or similar spanning tree configuration that immediately put ports into forwarding mode. 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.

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 and is required to be a unique VLAN per Hyperflex Cluster.
- **vMotion network**
- **VM network**

There are four vSwitches, each carrying a different network.

- **vswitch-hx-inband-mgmt**—Used for ESXi management, storage controller management, and 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-storage-data**—Used for ESXi storage data and HX Data Platform ISCSI. 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. 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.
- All VLANs (including storage-data and vmotion) must be configured upstream. It is imperative that storage-data is configured upstream to prevent fail-back timing issues which require temporary upstream connectivity.

VLAN Type	Description
Note	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 Type	Description
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 VM Network	VLAN VM Network: <user-defined> (for example, "hx-vm-network"). This is required to be a unique VLAN HyperFlex Cluster. 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	<i>Hostname or IP address.</i>
User Name	<i><admin username></i>
Password	<i><admin username></i>

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.
- Data network IP addresses cannot be changed after the storage cluster is created. Contact Cisco TAC for assistance changing Management Network IPs.
- 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
<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>
<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>
<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>
<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>
<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			
<i><IP Address ></i>		<i><IP Address ></i>	

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
Name	Enter a name for the storage cluster.
Management IP Address	<p>This provides the storage management network, access on each ESXi host.</p> <ul style="list-style-type: none"> • The IP address must be on the same subnet as the Management IP addresses for the nodes. • Do not allow cluster management IPs to share the last octet with another cluster on the same subnet. • These IP addresses are in addition to the four IP addresses we assign to each node in the Hypervisor section.
Storage Cluster Data IP Address	<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>
Data Replication Factor	<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 Data Replication Factor. The choices are:</p> <ul style="list-style-type: none"> • Data Replication Factor 3—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. <p>Attention This is the recommended option.</p> <ul style="list-style-type: none"> • Data Replication Factor 2—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>

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
vCenter Server	Enter your current vCenter server web address. For example, <i>http://<IP address></i> .
User Name	Enter <i><admin username></i> .
Password	Enter <i><admin password></i> .
Datacenter Name Note 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.
Cluster Name	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="695 1119 846 1146"><IP address></p> <p data-bbox="695 1163 1523 1226">DNS server address is required if you are using hostnames while installing the HyperFlex Data Platform.</p> <p data-bbox="695 1245 748 1272">Note</p> <ul style="list-style-type: none"> <li data-bbox="862 1245 1515 1339">• If you do not have a DNS server, do not enter a hostname under System Services in the Cluster Configuration page of the HX Data Platform Installer. Use only IP addresses. <li data-bbox="862 1358 1515 1453">• 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.

Field	Essential Information
NTP Server(s) (A reliable NTP server is required)	<p><IP address></p> <p>NTP server is used for clock synchronization between:</p> <ul style="list-style-type: none"> • Storage controller VM • ESXi hosts • vCenter server <p>Important 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><your time zone></p> <p>Select a time zone for the storage controller VMs. It is used to determine when to take scheduled snapshots.</p> <p>Note 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 / M6SN (All NVMe 220) HXAF240C-M6SN (All NVMe 240)	12	Low	10,800 MHz	Unlimited
With HX Boost Mode enabled: HXAF225C-M6 HXAF220C-M5SN / M6SN (All NVMe 220) HX245C-M6 HXAF240C-M6SN (All NVMe 240)	16	Low	10,800 MHz	Unlimited

Product ID	Number of VM CPU	Shares	Reservation	Limit
With HX Boost Mode enabled: HXAF220C-M5/M6 HXAF240C-M5SX/M6	12	Low	10,800 MHz	Unlimited
All Other Models	8	Low	10,800 MHz	Unlimited



Note Enabling HX Boost Mode does not change the controller VM CPU reservation. It merely changes the upper limit on how much CPU the controller VM can consume.

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
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 HX-E-240M5SX HX-E-240M6SX	72 GB	Yes
HXAF240C-M4/M5SX/M6SX HXAF-E-240M5SX HXAF-E-240M6SX	72 GB 88 GB for configurations with 7.6 TB SSDs (SED and non-SED)	Yes
HX240C-M5L HX240C-M6L	78 GB	Yes

Server Model	Amount of Guest Memory	Reserve All Guest Memory
HXAF220C-M5SN (All NVMe 220)	72 GB for configurations using < 8 TB NVMe capacity drives 84 GB for configurations with 8 TB NVMe capacity drives Note 15 TB capacity drives are supported on M6 servers only.	Yes
HXAF220C-M6SN (All NVMe 220)	70 GB for configurations using 8 TB or lower NVMe capacity drives 74 GB for configurations with 15 TB NVMe capacity drives Note 15 TB capacity drives are supported on M6 servers only.	Yes
HXAF240C-M6SN (All NVMe 240)	72 GB for configurations using < 8 TB NVMe capacity drives 84 GB for configurations with 8 TB NVMe capacity drives ¹ 94 GB for configurations with 15 TB NVMe capacity drives ²	Yes
HXAF240C-M5SD HX240C-M5SD (Short depth 240)	48 GB Note 7.6 TB drives are supported on HXAF240C-M5SD but they do not require a higher memory configuration.	Yes

¹ Requires HX 5.0(2b) or higher version

² Requires HX 5.0(2b) or higher version

- 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.

- Support for 16 TB LFF drives is not enabled.



Note Starting with HX 5.0(2b), new controller VM memory allocations are made for any clusters deployed using HX 5.0(2b) and later as a fresh install or SW redeploy, or for nodes which are expanded after 5.0(2b). If you are upgrading clusters to HX 5.0(2b) or later, note that the memory assigned to controller VMs is not changed automatically. If you wish to reduce the CVM memory, you will need to do so manually. For more information on changing the CVM memory manually, see [Changing Controller Memory on a Cluster, on page 21](#).

Memory Usage for Controller VMs - NVMe Drives

The following table details the memory usage for storage controller VMs for NVMe drives.

Table 2: Memory Usage for Controller VMs - NVMe Drives

Server Model	Drive Size		
	1 / 4 TB	8 TB	15 TB (HX Release 5.0(2a) and later)
M5 C220 All NVMe	72 GB	84 GB	No Support
M6 C220 All NVMe	72 GB	84 GB	96 GB
M6 240C All NVMe	72 GB	84 GB	110 GB

- 15 TB persistent drives on All Flash configurations is not supported.
- Support for 16 TB LFF drives is not enabled.

Changing Controller Memory on a Cluster

Starting with HX 5.0(2b), new controller VM memory allocations are made for any clusters deployed using HX 5.0(2b) and later as a fresh install or software redeployments, or for nodes which are expanded after 5.0(2b). If you are upgrading clusters to HX 5.0(2b) or later, note that the memory assigned to controller VMs is not changed automatically. If you wish to reduce the CVM memory, you will need to do so manually.



Note Changes you make will not generally apply to compute nodes, only converged nodes.

To change the controller VM memory on a node, proceed as follows:

- Step 1** Put the node in Maintenance Mode. For more information, see [Entering Cisco HyperFlex Maintenance Mode](#) in the [HX Administration Guide, 5.0](#).
- Step 2** From the vSphere web client, select the SCVM and proceed as follows:

- a) Edit Settings.
- b) Adjust the memory. For more information on adjusting your memory appropriately, see [Memory Resource Reservation for Controller VMs, on page 19](#).

Step 3 Exit Maintenance Mode. For more information, see [Exiting Cisco HyperFlex Maintenance Mode](#) in the [HX Administration Guide, 5.0](#).

Step 4 Type `top -n 1 -b | grep Mem` or `free -m` to confirm the memory change.

The example below is for SCVM memory being 96GB

Example:

```
root@SpringpathControllerA01BC2DEFG:~# top -n 1 -b | grep Mem
Mem: 99010800k total, 7274456k used, 91736344k free, 19620k buffers
```

```
root@SpringpathController55I37QHCXA:~# free -m
total used free shared buff/cache available
Mem: 98304 37692 58095 6 2517 34263
Swap: 0 0 0
```

Step 5 Login to the cluster, and verify that cluster is online and healthy by running the `hxcli cluster --detail` command.

Example:

```
root@SpringpathControllerA01BC2DEFG:~# hxcli cluster --detail
address: 192.168.255.165
name: hx-3
state: online
uptime: 181 days 3 hours 9 minutes 2 seconds
activeNodes: 4 of 4
compressionSavings: 38.5514792631
deduplicationSavings: 96.9983763142
freeCapacity: 7.9T
healingInfo:
  inProgress: False
resiliencyDetails:
  current ensemble size:4
  # of caching failures before cluster shuts down:3
  minimum cache copies remaining:3
  minimum data copies available for some user data:3
  minimum metadata copies available for cluster metadata:3
  # of unavailable nodes:0
  # of nodes failure tolerable for cluster to be available:1      health state reason:storage cluster is healthy.
  # of node failures before cluster shuts down:3
  # of node failures before cluster goes into readonly:3
  # of persistent devices failures tolerable for cluster to be available:2
  # of node failures before cluster goes to enospace warn trying to move the existing data:na
  # of persistent devices failures before cluster shuts down:3
  # of persistent devices failures before cluster goes into readonly:3
  # of caching failures before cluster goes into readonly:na
  # of caching devices failures tolerable for cluster to be available:2
resiliencyInfo:
  messages:
    Storage cluster is healthy.
  state: 1
  nodeFailuresTolerable: 1
  cachingDeviceFailuresTolerable: 2
  persistentDeviceFailuresTolerable: 2
```

```
spaceStatus: normal
totalCapacity: 8.0T
totalSavings: 98.155546647
usedCapacity: 127.3G
clusterAccessPolicy: lenient
dataReplicationCompliance: compliant
dataReplicationFactor: 3
```

What to do next

Once the cluster becomes healthy, repeat these steps on each SCVM (one at a time) until all SCVMs have their memory changed.

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	
Enable Auto Support check box	Check this box during HX storage cluster creation.
Mail Server	<p><IP address></p> <p>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.</p> <p>Note Only unauthenticated SMTP is supported for ASUP.</p>
Mail Sender	<p><username@domain.com></p> <p>Email address to use for sending Auto Support notifications.</p>
ASUP Recipient	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 vCenter Server > Manage > Advanced Settings , key <code>config.vpxd.sso.sts.uri</code>