



Managing Disks

- [Managing Disks in the Cluster, on page 1](#)
- [Disk Requirements, on page 1](#)
- [Replacing Self Encrypted Drives \(SEDs\), on page 6](#)
- [Replacing SSDs, on page 8](#)
- [Replacing NVMe SSDs, on page 9](#)
- [Replacing Housekeeping SSDs, on page 10](#)
- [Replacing or Adding Hard Disk Drives, on page 12](#)

Managing Disks in the Cluster

Disks, SSDs or HDDs, might fail. If this occurs, you need to remove the failed disk and replace it. Follow the server hardware instructions for removing and replacing the disks in the host. The HX Data Platform identifies the SSD or HDD and incorporates it into the storage cluster.

To increase the datastore capacity of a storage cluster add the same size and type SSDs or HDDs to each converged node in the storage cluster. For hybrid servers, add hard disk drives (HDDs). For all flash servers, add SSDs.



Note When performing a hot-plug pull and replace on multiple drives from different vendors or of different types, pause for a few moments (30 seconds) between each action. Pull, pause for about 30 seconds and replace a drive, pause for 30 seconds. Then, pull, pause for 30 seconds and replace the next drive.

Sometimes, when a disk is removed it continues to be listed in cluster summary information. To refresh this, restart the HX cluster.

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).

Converged Nodes

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.
-

The following tables list the compatible drives for each HX server type. Drives are located in the front slots of the server, unless otherwise indicated. Multiple drives listed are options. Use one drive size for capacity per server. Minimum and maximum number of drives are listed for each component.

HX240 M5 Servers

Component	Qty	Hybrid	All Flash	Hybrid SED	All Flash SED
System SSD for logs	1	240 GB SSD	240 GB SSD	240 GB SSD	240 GB SSD
Cache SSD	1 (back)	1.6 TB SSD	1.6 TB NVMe 400 GB SSD	1.6 TB SSD	800 GB SSD
Persistent	6-23	1.2 TB HDD 1.8 TB HDD	960 GB SSD 3.8 TB SSD	1.2 TB HDD	800 GB SSD 960 GB SSD 3.8 TB SSD



Note For information on disk requirements for HX240 M5 LFF servers, see [Disk Requirements for LFF Converged Nodes Hardware and Software Requirements, on page 4](#).

HX240 M4 Servers

Component	Qty	Hybrid	All Flash	Hybrid SED	All Flash SED
System SSD for logs	1	120 GB SSD 240 GB SSD	120 GB SSD 240 GB SSD	120 GB SSD 240 GB SSD	120 GB SSD 240 GB SSD
Cache SSD	1	1.6 TB SSD	1.6 TB NVMe 400 GB SSD	1.6 TB SSD	1.6 TB NVMe 800 GB SSD
Persistent	6-23	1.2 TB HDD 1.8 TB HDD	960 GB SSD 3.8 TB SSD	1.2 TB HDD	800 GB SSD 960 GB SSD 3.8 TB SSD

HX220 M5 Servers

Component	Qty	Hybrid	All Flash	Hybrid SED	All Flash SED
System SSD for logs	1	240 GB SSD	240 GB SSD	240 GB SSD	240 GB SSD
Cache SSD	1	480 GB SSD 800 GB SSD	1.6 TB NVMe 400 GB SSD	800 GB SSD	800 GB SSD
Persistent	6-8	1.2 TB HDD 1.8 TB HDD	960 GB SSD 3.8 TB SSD	1.2 TB HDD	800 GB SSD 960 GB SSD 3.8 TB SSD

HX 220 M4 Servers

Component	Qty	Hybrid	All Flash	Hybrid SED	All Flash SED
System SSD for logs	1	120 GB SSD 240 GB SSD	120 GB SSD 240 GB SSD	120 GB SSD 240 GB SSD	120 GB SSD 240 GB SSD
Cache SSD	1	480 GB SSD	400 GB SSD	800 GB SSD	800 GB SSD
Persistent	6	1.2 TB HDD 1.8 TB HDD	960 GB SSD 3.8 TB SSD	1.2 TB HDD	800 GB SSD 960 GB SSD 3.8 TB SSD

HX220 M5 Servers for Edge Clusters

Component	Qty	Hybrid	All Flash	Hybrid SED	All Flash SED
System SSD for logs	1	240 GB SSD	240 GB SSD	240 GB SSD	240 GB SSD
Cache SSD	1	480 GB SSD 800 GB SSD	1.6 TB NVMe 400 GB SSD	800 GB SSD	800 GB SSD
Persistent	3-8	1.2 TB HDD	960 GB SSD 3.8 TB SSD	1.2 TB HDD	800 GB SSD 960 GB SSD 3.8 TB SSD

HX 220 M4 Servers for Edge Clusters

Component	Qty	Hybrid	All Flash	Hybrid SED	All Flash SED
System SSD for logs	1	120 GB SSD 240 GB SSD	120 GB SSD 240 GB SSD	120 GB SSD 240 GB SSD	120 GB SSD 240 GB SSD
Cache SSD	1	480 GB SSD	400 GB SSD	800 GB SSD	800 GB SSD
Persistent	3-6	1.2 TB HDD	960 GB SSD 3.8 TB SSD	1.2 TB HDD	800 GB SSD 960 GB SSD 3.8 TB SSD

Disk Requirements for LFF Converged Nodes

The following table lists the supported HX240 M5 Server Large-Form-Factor (LFF) converged node configurations:

Table 1: HX240 M5 Server Large-Form-Factor (LFF) Configuration

	Description	Part Number	Quantity
Memory	16GB or 32GB or 64GB or 128GB DDR4-2666-MHz	HX-MR-X16G1RS-H HX-MR-X32G2RS-H HX-MR-X64G4RS-H HX-MR-128G8RS-H	Min. 128 MB
Processor	Processor Choices: Supported Skylake parts on HX 240 M5	Varies	2
Drive Controller	Cisco 12Gbps Modular SAS HBA	HX-SAS-M5	1
SSD1 (Boot SSD)	240GB 2.5 inch Enterprise Value 6G SATA SSD	HX-SD240G61X-EV	1
SSD2 (Cache/WL)	3.2TB 2.5 inch Enterprise Performance 12G SAS SSD(3X)	HX-SD32T123X-EP	
HDD (Capacity/Data)	6TB 12G SAS 7.2K RPM LFF HDD (4K) OR 8TB 12G SAS 7.2K RPM LFF HDD (4K)	HX-HD6T7KL4KN OR HX-HD8T7KL4KN	6 - 12
Network	Cisco VIC 1387 Dual Port 40GB QSFP CNA MLOM	HX-MLOM-C40Q-03	1
Boot Device	240GB SATA M.2	HX-M2-240GB	1
Software	Cisco HX Data Platform 1, 2, 3, or 4 or 5yr SW subscription	HXDP-001-xYR	1
Optional VMware License	Factory Installed – VMware vSphere6 Enterprise Plus/Standard SW License & Subscription		2
FI Support	2G FI and 3G FI		

Hardware and Software Requirements

Hardware

- Memory Configurable
- CPU Configurable
- HDD Storage Quantity

Software

- Storage Controller
 - Reserves 72GB RAM
 - Reserves 8 vCPU, 10.800 GHz CPU

- VAAI VIB
- IO Visor VIB

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 M3/M4/M5 • B260 M4 • B420 M4 • B460 M4 • C240 M3/M4/M5 • C220 M3/M4/M5 • 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 not supported on Compute-only nodes.</p>

Replacing Self Encrypted Drives (SEDs)

Cisco HyperFlex Systems offers Data-At-Rest protection through Self-Encrypting Drives (SEDs) and Enterprise Key Management Support.

- Servers that are data at rest capable refer to servers with self encrypting drives.
- All servers in an encrypted HX Cluster must be data at rest capable.
- Encryption is configured on a HX Cluster, after the cluster is created, using HX Connect.
- Servers with self encrypting drives can be either solid state drive (SSD) or hybrid.



Important To ensure the encrypted data remains secure, the data on the drive must be **securely erased** prior to removing the SED.

Before you begin

Determine if the encryption is applied to the HX Cluster.

- **Encryption not configured**—No encryption related prerequisite steps are required to remove or replace the SED. See [Replacing SSDs, on page 8](#) or [Replacing or Adding Hard Disk Drives, on page 12](#) and the hardware guide for your server.
- **Encryption is configured**—Ensure the following:
 1. If you are replacing the SED, obtain a Return to Manufacturer Authorization (RMA). Contact TAC.
 2. If you are using a local key for encryption, locate the key. You will be prompted to provide it.
 3. To prevent data loss, ensure the data on the disk is not the last primary copy of the data.
If needed, add disks to the servers on the cluster. Initiate or wait until a rebalance completes.
 4. Complete the steps below before removing any SED.

Step 1 Ensure the HX Cluster is healthy.

Step 2 Login to HX Connect.

Step 3 Select **System Information** > **Disks** page.

Step 4 Identify and verify the disk to remove.

- a. Use the Turn On Locator LED button.
- b. Physically view the disks on the server.
- c. Use the Turn Off Locator LED button.

Step 5 Select the corresponding **Slot** row for the disk to be removed.

Step 6 Click **Secure erase**. This button is available only after a disk is selected.

Step 7 If you are using a local encryption key, enter the **Encryption Key** in the field and click **Secure erase**.

If you are using a remote encryption server, no action is needed.

Step 8 Confirm deleting the data on this disk, click **Yes, erase this disk**.

Warning This deletes all your data from the disk.

Step 9 Wait until the **Status** for the selected **Disk Slot** changes to **Ok To Remove**, then physically remove the disk as directed.

What to do next



Note Do not reuse a removed drive in a different server in this, or any other, HX Cluster. If you need to reuse the removed drive, contact TAC.

1. After securely erasing the data on the SED, proceed to the disk replacing tasks appropriate to the disk type: SSD or hybrid.

Check the **Type** column for the disk type.

- **Solid State** (SSDs)—See [Replacing SSDs, on page 8](#) and the hardware guide for your server.
- **Rotational** (hybrid drives)—See [Replacing or Adding Hard Disk Drives, on page 12](#) and the hardware guide for your server.

2. Check the status of removed and replaced SEDs.

When the SED is removed:

- **Status**—Remains **Ok To Remove**.
- **Encryption**—Changes from **Enabled** to **Unknown**.

When the SED is replaced, the new SED is automatically consumed by the HX Cluster. If encryption is not applied, the disk is listed the same as any other consumable disk. If encryption is applied, the security key is applied to the new disk.

- **Status**—Transitions from **Ignored** > **Claimed** > **Available**.
- **Encryption**—Transitions from **Disabled** > **Enabled** after the encryption key is applied.

Replacing SSDs

The procedures for replacing an SSD vary depending upon the type of SSD. Identify the failed SSD and perform the associated steps.



Note Mixing storage disks type or size on a server or across the storage cluster is not supported.

- Use all HDD, or all 3.8 TB SSD, or all 960 GB SSD
- Use the hybrid cache device on hybrid servers and all flash cache devices on all flash servers.
- When replacing cache or persistent disks, always use the same type and size as the original disk.

Step 1 Identify the failed SSD.

- For cache or persistent SSDs, perform a disk beacon check. See [Setting a Beacon](#).

Only cache and persistent SSDs respond to the beacon request. NVMe cache SSDs and housekeeping SSDs do not respond to beacon requests.

- For cache NVMe SSDs, perform a physical check. These drives are in Drive Bay 1 of the HX servers.
- For housekeeping SSDs on HXAF240c or HX240c servers, perform a physical check at the back of the server.
- For housekeeping SSDs on HXAF220c or HX220c servers, perform a physical check at Drive Bay 2 of the server.

Step 2 If the failed SSD is a housekeeping SSD, proceed based on the type of server.

- For HXAF220c or HX220c servers, proceed to Step 3.
- For HXAF240c or HX240c servers, contact Technical Assistance Center (TAC).

Step 3 If a failed SSD is a cache or persistent SSD, proceed based on the type of disk.

- For NVMe SSDs, see [Replacing NVMe SSDs, on page 9](#).
- For all other SSDs, follow the instructions for removing and replacing a failed SSD in the host, per the server hardware guide.

After the cache or persistent drive is replaced, the HX Data Platform identifies the SDD and updates the storage cluster. When disks are added to a node, the disks are immediately available for HX consumption.

Step 4 To enable the Cisco UCS Manager to include new disks in the **UCS Manager > Equipment > Server > Inventory > Storage** tab, re-acknowledge the server node. This applies to cache and persistent disks.

Note Re-acknowledging a server is disruptive. Place the server into HX Maintenance Mode before doing so.

Step 5 If you replaced an SSD, and see a message *Disk successfully scheduled for repair*, it means that the disk is present, but is still not functioning properly. Check that the disk has been added correctly per the server hardware guide procedures.

Replacing NVMe SSDs

The procedures for replacing an SSD vary depending upon the type of SSD. This topic describes the steps for replacing NVMe cache SSDs.



Note Mixing storage disks type or size on a server or across the storage cluster is not supported. When replacing NVMe disks, always use the same type and size as the original disk.

Before you begin

Ensure the following conditions are met when using NVMe SSDs in HX Cluster servers.

- NVMe SSDs are supported in HX240 and HX220 All-Flash servers.
- Replacing NVMe SSDs with an HGST SN200 disk requires HX Data Platform version 2.5.1a or later.

- NVMe SSDs are only allowed in slot 1 of the server. Other server slots do not detect NVMe SSDs.
- NVMe SSDs are only used for cache.
 - Using them for persistent storage is not supported.
 - Using them as the housekeeping drive is not supported.
 - Using them for hybrid servers is not supported.

-
- Step 1** Confirm the failed disk is an NVMe cache SSD.
- Perform a physical check. These drives are in Drive Bay 1 of the HX servers. NVMe cache SSDs and housekeeping SSDs do not respond to beacon requests.
- If the failed SSD is not an NVMe SSD, see [Replacing SSDs, on page 8](#).
- Step 2** Put ESXi host into HX Maintenance Mode.
- Login to HX Connect.
 - Select **System Information > Nodes > node > Enter HX Maintenance Mode**.
- Step 3** Follow the instructions for removing and replacing a failed SSD in the host, per the server hardware guide.
- Note** When you remove an HGST NVMe disk, the controller VM will fail until you reinsert a disk of the same type into the same slot or reboot the host.
- After the cache or persistent drive is replaced, the HX Data Platform identifies the SDD and updates the storage cluster. When disks are added to a node, the disks are immediately available for HX consumption.
- Step 4** Reboot the ESXi host. This enables ESXi to discover the NVMe SSD.
- Step 5** Exit ESXi host from HX Maintenance Mode.
- Step 6** To enable the Cisco UCS Manager to include new disks in the **UCS Manager > Equipment > Server > Inventory > Storage** tab, re-acknowledge the server node. This applies to cache and persistent disks.
- Note** Re-acknowledging a server is disruptive. Place the server into HX Maintenance Mode before doing so.
- Step 7** If you replaced an SSD, and see a message *Disk successfully scheduled for repair*, it means that the disk is present, but is still not functioning properly. Check that the disk has been added correctly per the server hardware guide procedures.
-

Replacing Housekeeping SSDs

Identify the failed housekeeping SSD and perform the associated steps.

- Step 1** Identify the failed housekeeping SSD.
- Physically check the SSD drives, as housekeeping drives are not listed through a beacon check.
- Step 2** Remove the SSD and replace with a new SSD of the same kind and size. Follow the steps in the server hardware guide. The server hardware guide describes the physical steps required to replace the SSD.

Note Before performing the hardware steps, enter the node into Cisco HX Maintenance Mode. After performing the hardware steps, exit the node from Cisco HX Maintenance Mode.

Step 3 Using `SSH`, login into the storage controller VM of the affected node and run the following command.

```
# /usr/share/springpath/storfs-appliance/config-bootdev.sh -r -y
```

This command consumes the new disk, adding it into the storage cluster.

Sample response

```
Creating partition of size 65536 MB for /var/stv ...
Creating ext4 filesystem on /dev/sdg1 ...
Creating partition of size 24576 MB for /var/zookeeper ...
Creating ext4 filesystem on /dev/sdg2 ...
Model: ATA INTEL SSDSC2BB12 (scsi)
Disk /dev/sdg: 120034MB
Sector size (logical/physical): 512B/4096B
Partition Table: gpt ....
discovered. Rebooting in 60 seconds
```

Step 4 Wait for the storage controller VM to automatically reboot.

Step 5 When the storage controller VM completes its reboot, verify that partitions are created on the newly added SSD. Run the command.

```
# df -ah
```

Sample response

```
.....
/dev/sdb1 63G 324M 60G 1%
/var/stv /dev/sdb2 24G 173M 23G 1% /var/zookeeper
```

Step 6 Identify the HX Data Platform installer package version installed on the existing storage cluster.

```
# stcli cluster version
```

The same version must be installed on all the storage cluster nodes. Run this command on the controller VM of any node in the storage cluster, but not the node with the new SSD.

Step 7 Copy the HX Data Platform installer packages into the storage controller VM in `/tmp` folder.

```
# scp <hxdp_installer_vm_ip>:/opt/springpath/packages/storfs-packages-<hxdp_installer>.tgz /tmp
# cd /tmp
# tar zxvf storfs-packages-<hxdp_installer>.tgz
```

Step 8 Run the HX Data Platform installer deployment script.

```
# ./inst-packages.sh
```

For additional information on installing the HX Data Platform, see the appropriate [Cisco HX Data Platform Install Guide](#).

Step 9 After the package installation, HX Data Platform starts automatically. Check the status.

```
# status storfs
```

Sample response

```
storfs running
```

The node with the new SSD re-joins the existing cluster and the cluster returns to a healthy state.

Replacing or Adding Hard Disk Drives



- Note** Mixing storage disks type or size on a server or across the storage cluster is not supported.
- Use all HDD, or all 3.8 TB SSD, or all 960 GB SSD
 - Use the hybrid cache device on hybrid servers and all flash cache devices on all flash servers.
 - When replacing cache or persistent disks, always use the same type and size as the original disk.
-

Step 1 Refer to the hardware guide for your server and follow the directions for adding or replacing disks.

Step 2 Add HDDs of the same size to each node in the storage cluster.

Step 3 Add the HDDs to each node within a reasonable amount of time.

The storage starts being consumed by storage cluster immediately.

The vCenter Event log displays messages reflecting the changes to the nodes.

Note When disks are added to a node, the disks are immediately available for HX consumption although they will not be seen in the UCSM server node inventory. This includes cache and persistent disks. To include the disks in the **UCS Manager > Equipment > Server > Inventory > Storage** tab, re-acknowledge the server node.

Note Re-acknowledging a server is disruptive. Place the server into HX Maintenance Mode before doing so.
