



HX Storage Cluster Overview

- [Cisco HX Data Platform Overview, on page 1](#)
- [Storage Cluster Physical Components Overview, on page 1](#)
- [HX Data Platform Capacity Overview, on page 3](#)
- [HX Data Platform High Availability Overview, on page 6](#)
- [Storage Cluster Status, on page 7](#)
- [HX Data Platform Cluster Tolerated Failures, on page 8](#)
- [Responses to Storage Cluster Node Failures, on page 10](#)
- [HX Data Platform Ready Clones Overview, on page 12](#)
- [HX Native Snapshots Overview, on page 12](#)

Cisco HX Data Platform Overview

Cisco HyperFlex Data Platform (HX Data Platform) is a hyperconverged software appliance that transforms Cisco servers into a single pool of compute and storage resources. It eliminates the need for network storage and enables seamless interoperability between computing and storage in virtual environments. The Cisco HX Data Platform provides a highly fault-tolerant distributed storage system that preserves data integrity and optimizes performance for virtual machine (VM) storage workloads. In addition, native compression and deduplication reduce storage space occupied by the VMs and VM workloads.

Cisco HX Data Platform has many integrated components. These include: Cisco Fabric Interconnects (FIs), Cisco UCS Manager, Cisco HX specific servers, and Cisco compute only servers; VMware vSphere, ESXi servers, and vCenter; and the Cisco HX Data Platform Installer, controller VMs, HX Connect, vSphere HX Data Platform Plug-in, and `hxcli` commands.

Cisco HX Data Platform is installed on a virtualized platform such as VMware vSphere. During installation, after specifying the Cisco HyperFlex HX Cluster name, and the HX Data Platform creates a hyperconverged storage cluster on each of the nodes. As your storage needs to increase and you add nodes in the HX cluster, the HX Data Platform balances the storage across the additional resources. Compute only nodes can be added to increase compute only resources to the storage cluster.

Storage Cluster Physical Components Overview

Cisco HyperFlex storage clusters contain the following objects. These objects are monitored by the HX Data Platform for the storage cluster. They can be added and removed from the HX storage cluster.

- **Converged nodes**—Converged nodes are the physical hardware on which the VM runs. They provide computing and storage resources such as disk space, memory, processing, power, and network I/O.

When a converged node is added to the storage cluster, a storage controller VM is installed. The HX Data Platform services are handled through the storage controller VM. Converged nodes add storage resources to your storage cluster through their associated drives.

Run the *Cluster Expansion* workflow from the HX Data Platform Installer to add converged nodes to your storage cluster. You can remove converged nodes using *hxcli* commands.

- **Compute nodes**—Compute nodes add compute resource but not storage capacity to the storage cluster. They are used as a means to add compute resources, including CPU and memory. They do not need to have any caching (SSD) or storage (HDD) drives. Compute nodes are optional in a HX storage cluster.

When a compute node is added to the storage cluster, an agent controller VM is installed. The HX Data Platform services are handled through the agent controller VM.

Run the *Cluster Expansion* workflow from the HX Data Platform Installer to add compute nodes to your storage cluster. You can remove compute nodes using *hxcli* commands.

- **Drives**—There are two types of drives that are required for any node in the storage cluster: Solid State Drive (SSD) and Hard Disk Drive (HDD). HDD typically provides the physical storage units associated with converged nodes. SSD typically supports management.

Adding HDD to existing converged nodes, also adds storage capacity to the storage cluster. When storage is added to a HX node in the storage cluster, an equal amount of storage must be added to every node in the storage cluster.

When disks are added or removed, the HX Data Platform rebalances the storage cluster to adjust for the change in storage resources.

Adding or removing disks on your converged nodes is not performed through the HX Data Platform. Before adding or removing disks, review the best practices. See the server hardware guides for specific instructions to add or remove disks in nodes.

NVMe Caching SSD's slot information is unavailable from HX-Connect for all AF server PIDs except for the All-NVMe server PIDs. Please refer to UCSM management console for NVMe SSD slot information.

- **Datastores**—Storage capacity and datastore capacity. This is the combined consumable physical storage available to the storage cluster through datastores, and managed by the HX Data Platform.

Datastores are logical containers that are used by the HX Data Platform to manage your storage use and storage resources.

Datastores are where the host places virtual disk files and other VM files. Datastores hide the specifics of physical storage devices and provide a uniform model for storing VM files.

HX Data Platform Capacity Overview



Note Capacity addition in a cluster through the addition of disks or nodes can result in a rebalance. This background activity can cause interference with regular User IO on the cluster and increase the latency. You must note the time duration for the storage capacity at the time where performance impact can be tolerated. Also, this operation may be performed in urgent situations that may warrant capacity addition.

In the HX Data Platform the concept of capacity is applied to both datastores and storage clusters. Values are measured in base-2 (GiB/TiB), but for simplicity and consistency are labeled as GB or TB.

- **Cleaner**—A process run on all the storage cluster datastores. After it completes, all the storage cluster datastores total capacity should be in a similar range to the total storage cluster capacity, excluding the metadata. Datastore capacity listed typically will not match the HX storage cluster capacity. See the [Cisco HX Data Platform Command Line Interface Reference Guide](#) for information on the `cleaner` command.

- **Cluster capacity**—All the storage from all the disks on all the nodes in the storage cluster. This includes uncleaned data and the metadata overhead for each disk.

The total/used/free capacity of cluster is based on overall storage capacity and how much storage is used.

- **Condition**—When the HX Storage Cluster enters a space event state, the **Free Space Status** fields are displayed. The **Condition** field lists the space event state. The options are: **Warning**, **Critical**, and **Alert**.

- **Available Datastore capacity**—The amount of storage available for provisioning to datastores without over-provisioning. Generally, this is similar to the cleaned storage cluster capacity, but it is not an exact match. It does not include metadata or uncleaned data.

The provisioned/used/free capacity of each datastore is based on datastore (thin) provisioned capacity. Because the datastore is thin provisioned, the provisioned capacity (specified by the administrator when creating the datastore) can be well above the actual storage.

- **Free Capacity, storage cluster**—Same as available capacity. For the storage cluster, this is the difference between the amount available to the storage cluster and the amount used in the storage cluster.

- **Free capacity, datastore**—Same as available capacity. For all the storage cluster datastores, this is the difference between the amount provisioned to all the storage cluster datastores and the amount used on all the storage cluster datastores.

The amount used on the whole storage cluster is not included in this datastore calculation. Because datastores are frequently over provisioned, the free capacity can indicate a large availability on all the storage cluster datastores, while the storage cluster capacity can indicate a much lower availability.

- **Multiple users**—Can have different datastores with different provisioned capacities. At any point in time, users do not fully utilize their allocated datastore capacity. When allocating datastore capacity to multiple users, it is up to the administrator to ensure that each user's provisioned capacity is honored at all time.

- **Over-provisioning**—Occurs when the amount of storage capacity allocated to all the datastores exceeds the amount available to the storage cluster.

It is a common practice to initially over-provision. It allows administrators to allocate the capacity now and backfill the actual storage later.

The value is the difference between the usable capacity and provisioned capacity.

It displays zero (0) value, unless more space has been allocated than the maximum physical amount possible.

Review the over provisioned capacity and ensure that your system does not reach an out-of-space condition.

- **Provisioned**—Amount of capacity allowed to be used by and allocated to the storage cluster datastores.

The provisioned amount is not set aside for the sole use of the storage cluster datastores. Multiple datastores can be provisioned storage from the same storage capacity.

- **Space Needed**—When the HX Storage Cluster enters a space event state, the **Free Space Status** fields are displayed. **Space Needed** indicates the amount of storage that needs to be made available to clear the listed **Condition**.

- **Used**—Amount of storage capacity consumed by the listed storage cluster or datastore.

HX Data Platform internal meta-data uses 0.5% to 1% space. This might cause the HX Data Platform Plug-in or HX Connect to display a Used Storage value even if you have no data in your datastore.

Storage Used shows how much datastore space is occupied by virtual machine files, including configuration and log files, snapshots, and clones. When the virtual machine is running, the used storage space also includes swap files.

- **Usable Capacity**—Amount of storage in the storage cluster available for use to store data.

Understanding Capacity Savings

The Capacity portlet on the Summary tab displays the deduplication and compression savings provided by the storage cluster. For example, with 50% overall savings, a 6TB capacity storage cluster can actually store 9 TB of data.

The total storage capacity saved by the HX Data Platform system is a calculation of two elements:

- **Compression**—How much of the data is compressed.
- **Deduplication**—How much data is deduplicated. Deduplication is a method of reducing storage space by eliminating redundant data. It stores only one unique instance of the data.

Deduplication savings and compression savings are not simply added together. They are not independent operations. They are correlated using the following elements where essentially the number of unique bytes used for storage is reduced through deduplication. Then the deduplicated storage consumption is compressed to make even more storage available to the storage cluster.

Deduplication and compression savings are useful when working with VM clones.

If the savings is showing 0%, this indicates the storage cluster is new. The total ingested data to the storage cluster is insufficient to determine meaningful storage savings. Wait until sufficient data is written to the storage cluster.

For example:

1. Initial values

Given a VM of 100 GB that is cloned 2 times.

Total Unique Used Space (TUUS) = 100GB

Total Addressable Space (TAS) = 100x2 = 200 GB

Given, for this example:

Total Unique Bytes (TUB) = 25 GB

2. Deduplication savings

$$= (1 - \text{TUUS}/\text{TAS}) * 100$$

$$= (1 - 100\text{GB} / 200\text{GB}) * 100$$

$$= 50\%$$

3. Compression Savings

$$= (1 - \text{TUB}/\text{TUUS}) * 100$$

$$= (1 - 25\text{GB} / 100\text{GB}) * 100$$

$$= 75\%$$

4. Total savings calculated

$$= (1 - \text{TUB}/\text{TAS}) * 100$$

$$= (1 - 25\text{GB} / 200\text{GB}) * 100$$

$$= 87.5\%$$

Storage Capacity Event Messages

Cluster storage capacity includes all the storage from all the disks on all the nodes in the storage cluster. This available capacity is used to manage your data.

Calculating Cluster Capacity

A HyperFlex HX Data Platform cluster capacity is calculated as follows:

$$(((\text{capacity disk size in GB} * 10^9) / 1024^3) * \text{number of capacity disks per node} * \text{number of HyperFlex nodes} * 0.92) / \text{replication factor}$$

Divide the result by 1024 to get a value in TiB. The replication factor value is 3 if the HX cluster is set to RF=3, and the value is 2 if the HX cluster is set to RF=2. The 0.92 multiplier accounts for an 8% reservation set aside on each disk by the HX Data Platform software for various internal filesystem functions.

Calculation Example: <capacity disk size in GB> = 1200 for 1.2 TB disks <number of capacity disks per node> = 15 for an HX240c-M6SX model server <number of HyperFlex nodes> = 8
replication factor = 3

Result: $((1200 * 10^9) / 1024^3) * 15 * 8 * 0.92 / 3 = 41127.2049$ 41127.2049 / 1024 = 40.16 TiB



Note This formula for calculating cluster capacity does not apply for Large Form Factor (LFF) clusters.

Error Messages

Error messages are issued if your data storage needs to consume high amounts of available capacity, the performance and health of your storage cluster are affected. The error messages are displayed in vCenter Alarms panels, HX Connect, and HX Data Platform Plug-in Alarms and Events pages.



Note The event and alarm details provided on vCenter and HX Connect are not always a 1:1 relationship. When reviewing messages in HX Connect, it is a best practice to also review the events and tasks in vCenter.



Note **When the warning or critical errors appear:**

Add additional drives or nodes to expand capacity. Additionally, consider deleting unused virtual machines and snapshots. Performance is impacted until storage capacity is reduced.

- **SpaceWarningEvent** – Issues an error. This is a first level warning.

Cluster performance is impacted due to increased cleaner activity to reclaim the space as fast as possible. The effect on throughput and latency depend on the workload and how much read and writes are being performed.

Reduce the amount of storage capacity used to below the warning threshold, of 76% total HX Storage Cluster capacity.

- **SpaceAlertEvent** – Issues an error. Space capacity usage remains at error level.

This alert is issued after storage capacity has been reduced, but is still above the warning threshold.

Cluster performance is affected.

Continue to reduce the amount of storage capacity used, until it is below the warning threshold, of 80% total HX Storage Cluster capacity.

- **SpaceCriticalEvent** – Issues an error. This is a critical level warning.

Cluster is in a read only state.

Do not continue the storage cluster operations until you reduce the amount of storage capacity used to below this warning threshold, that is, 100% of the available disk space.

- **SpaceRecoveredEvent** - This is informational. The cluster capacity has returned to normal range.

Cluster storage space usage is back to normal.

HX Data Platform High Availability Overview

The HX Data Platform High Availability (HA) feature ensures that the storage cluster maintains at least two copies of all your data during normal operation with three or more fully functional nodes.

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*.

The number of nodes in the storage cluster, combined with the Data Replication Factor and Access Policy settings, determine the state of the storage cluster that results from node failures.



Note Before using the HX Data Platform HA feature, enable DRS and vMotion on the vSphere Web Client.

Storage Cluster Status

HX Data Platform storage cluster status information is available through HX Connect, the HX Data Platform Plug-in, and the storage controller VM `hxcli` commands. Storage cluster status is described through resiliency and operational status values.

Storage cluster status is described through the following reported status elements:

- **Operational Status**—Describes the ability of the storage cluster to perform the functions storage management and storage cluster management of the cluster. Describes how well the storage cluster can perform operations.
- **Resiliency Status**—Describes the ability of the storage clusters to tolerate node failures within the storage cluster. Describes how well the storage cluster can handle disruptions.

The following settings take effect when the storage cluster transitions into particular operational and resiliency status states.

- **Data Replication Factor** —Sets the number of redundant data replicas.
- **Cluster Access Policy**—Sets the level of data protection and data loss.

Operational Status Values

Cluster Operational Status indicates the operational status of the storage cluster and the ability for the applications to perform I/O.

The Operational Status options are:

- **Online**—Cluster is ready for IO.
- **Offline**—Cluster is not ready for IO.
- **Out of space**—Either the entire cluster is out of space or one or more disks are out of space. In both cases, the cluster cannot accept write transactions, but can continue to display static cluster information.
- **Readonly**—Cluster cannot accept write transactions, but can continue to display static cluster information.
- **Unknown**—This is a transitional state while the cluster is coming online.

Other transitional states might be displayed during cluster upgrades and cluster creation.

Color coding and icons are used to indicate various status states. Click icons to display additional information such as reason messages that explain what is contributing to the current state.

Resiliency Status Values

Resiliency status is the data resiliency health status and ability of the storage cluster to tolerate failures.

Resiliency Status options are:

- **Healthy**—The cluster is healthy with respect to data and availability.
- **Warning**—Either the data or the cluster availability is being adversely affected.
- **Unknown**—This is a transitional state while the cluster is coming online.

Color coding and icons are used to indicate various status states. Click an icon to display additional information, such as reason messages that explain what is contributing to the current state.

HX Data Platform Cluster Tolerated Failures

If nodes or disks in the HX 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*.

How the number of node failures affect the storage cluster is dependent upon:

- **Number of nodes in the cluster**—The response by the storage cluster is different for clusters with 3 to 4 nodes and 5 or greater nodes.
- **Data Replication Factor**—Set during HX Data Platform installation and cannot be changed. The options are 2 or 3 redundant replicas of your data across the storage cluster. Production clusters should always use RF3. RF2 should be reserved for use in labs and demos.



Important Production clusters should be set to Data Replication Factor 3.

- **Access Policy**—Can be changed from the default setting after the storage cluster is created. The options are strict for protecting against data loss, or lenient, to support longer storage cluster availability.

Cluster State with Number of Failed Nodes

The tables below list how the storage cluster functionality changes with the listed number of simultaneous node failures.

Table 1: Cluster State in 5+ Node Cluster with Number of Failed Nodes, HX Release 4.5(x) and later

Replication Factor	Access Policy	Number of Failed Nodes	
		Read/Write	Read-Only
3	Lenient	2	--
3	Strict	1	2
2	Lenient	1	--
2	Strict	--	1

Table 2: Cluster State in 3 - 4 Node Clusters with Number of Failed Nodes HX Release 4.5(x) and later.

Replication Factor	Access Policy	Number of Failed Nodes	
		Read/Write	Read-Only
3	Lenient or Strict	1	--
2	Lenient	1	--
2	Strict	--	1

Cluster State with Number of Nodes with Failed Disks

The table below lists how the storage cluster functionality changes with the number of nodes that have one or more failed disks. Note that the node itself has not failed but disk(s) within the node have failed. **For example:** 2 indicates that there are 2 nodes that each have at least one failed disk.

There are two possible types of disks on the servers: SSDs and HDDs. When we talk about multiple disk failures in the table below, it's referring to the disks used for storage capacity. **For example:** If a cache SSD fails on one node and a capacity SSD or HDD fails on another node the storage cluster remains highly available, even with an Access Policy strict setting.

The table below lists the worst case scenario with the listed number of failed disks. This applies to any storage cluster 3 or more nodes. **For example:** A 3 node cluster with Replication Factor 3, while self-healing is in progress, only shuts down if there is a total of 3 simultaneous disk failures on 3 separate nodes.



Note HX storage clusters are capable of sustaining serial disk failures, (separate disk failures over time). The only requirement is that there is sufficient storage capacity available for support self-healing. The worst-case scenarios listed in this table only apply during the small window while HX is completing the automatic self-healing and rebalancing.

3+ Node Cluster with Number of Nodes with Failed Disks

Replication Factor	Access Policy	Failed Disks on Number of Different Nodes	
		Read/Write	Read Only
3	Lenient	2	--
3	Strict	1	2
2	Lenient	1	--
2	Strict	--	1

Data Replication Factor Settings



Important Data Replication Factor cannot be changed after the storage cluster is configured.

Data Replication Factor is set when you configure the storage cluster. Data Replication Factor defines the number of redundant replicas of your data across the storage cluster. The options are 2 or 3 redundant replicas of your data.

- If you have hybrid servers (servers that contain both SSD and HDDs), then the default is 3.
- If you have all flash servers (servers that contain only SSDs), then you must explicitly select either 2 or 3 during HX Data Platform installation.

Choose a Data Replication Factor. The choices are:

- Data Replication Factor 3 — (**Recommended Usage:**All production environments) Keep three redundant replicas of the data. This consumes more storage resources, and ensures the maximum protection for your data in the event of node or disk failure.
 - Data Replication Factor 2 — (**Recommended Usage:**Non production labs and demos) Keep two redundant replicas of the data. This consumes fewer storage resources, but reduces your data protection in the event of node or disk failure.
-

Cluster Access Policy

The Cluster Access Policy works with the Data Replication Factor to set levels of data protection and data loss prevention. There are two Cluster Access Policy options. The default is `lenient`. It is not configurable during installation, but can be changed after installation and initial storage cluster configuration.

- **Strict** - Applies policies to protect against data loss.

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. The strict setting helps protect the data in event of simultaneous failures.

- **Lenient** - Applies policies to support longer storage cluster availability. This is the default.

Responses to Storage Cluster Node Failures

A storage cluster healing timeout is the length of time HX Connect or HX Data Platform Plug-in waits before automatically healing the storage cluster. If a disk fails, the healing timeout is 1 minute. If a node fails, the healing timeout is 2 hours. A node failure timeout takes priority if a disk and a node fail at same time or if a disk fails after node failure, but before the healing is finished.

When the cluster resiliency status is Warning, the HX Data Platform system supports the following storage cluster failures and responses.

Optionally, click the associated Cluster Status/Operational Status or Resiliency Status/Resiliency Health in HX Connect and HX Data Platform Plug-in, to display reason messages that explain what is contributing to the current state.

Review the table and perform the appropriate action.

Cluster Size	Number of Simultaneous Failures	Entity Failed	Maintenance Action to Take
3 nodes	1	One node.	The storage cluster does not automatically heal. Replace the failed node to restore storage cluster health.
3 nodes	2	Two or more disks on two nodes are blocklisted or failed.	<ul style="list-style-type: none"> a. If one cache SSD fails, the storage cluster does not automatically heal. b. If one HDD fails or is removed, the disk is blocklisted immediately. The storage cluster automatically begins healing within a minute. c. If more than one HDD fails, the system might not automatically restore storage cluster health. If the system is not restored, replace the faulty disk and restore the system by rebalancing the cluster.
4 nodes	1	One node.	<p>If the node does not recover in two hours, the storage cluster starts healing by rebalancing data on the remaining nodes.</p> <p>To recover the failed node immediately and fully restore the storage cluster:</p> <ul style="list-style-type: none"> a. Check that the node is powered on and restart it if possible. You might need to replace the node. b. Rebalance the cluster.
4 nodes	2	Two or more disks on two nodes.	<p>If two SSDs fail, the storage cluster does not automatically heal.</p> <p>If the disk does not recover in one minute, the storage cluster starts healing by rebalancing data on the remaining nodes.</p>
5+ nodes	2	Up to two nodes.	<p>If the node does not recover in two hours, the storage cluster starts healing by rebalancing data on the remaining nodes.</p> <p>To recover the failed node immediately and fully restore the storage cluster:</p> <ul style="list-style-type: none"> a. Check that the node is powered on and restart it if possible. You might need to replace the node. b. Rebalance the cluster. <p>If the storage cluster shuts down, see Troubleshooting, Two Nodes Fail Simultaneously Causes the Storage Cluster to Shutdown section.</p>
5+ nodes	2	Two nodes with two or more disk failures on each node.	The system automatically triggers a rebalance after a minute to restore storage cluster health.

Cluster Size	Number of Simultaneous Failures	Entity Failed	Maintenance Action to Take
5+ nodes	2	One node and One or more disks on a different node.	<p>If the disk does not recover in one minute, the storage cluster starts healing by rebalancing data on the remaining nodes.</p> <p>If the node does not recover in two hours, the storage cluster starts healing by rebalancing data on the remaining nodes.</p> <p>If a node in the storage cluster fails and a disk on a different node also fails, the storage cluster starts healing the failed disk (without touching the data on the failed node) in one minute. If the failed node does not come back up after two hours, the storage cluster starts healing the failed node as well.</p> <p>To recover the failed node immediately and fully restore the storage cluster:</p> <ol style="list-style-type: none"> a. Check that the node is powered on and restart it if possible. You might need to replace the node. b. Rebalance the cluster.

HX Data Platform Ready Clones Overview

HX Data Platform Ready Clones is a pioneer storage technology that enables you to rapidly create and customize multiple cloned VMs from a host VM. It enables you to create multiple copies of VMs that can then be used as standalone VMs.

A Ready Clone, similar to a standard clone, is a copy of an existing VM. The existing VM is called the host VM. When the cloning operation is complete, the Ready Clone is a separate guest VM.

Changes made to a Ready Clone do not affect the host VM. A Ready Clone's MAC address and UUID are different from that of the host VM.

Installing a guest operating system and applications can be time consuming. With Ready Clone, you can make many copies of a VM from a single installation and configuration process.

Clones are useful when you deploy many identical VMs to a group.

HX Native Snapshots Overview

HX native snapshots are a backup feature that saves versions (states) of VMs. VMs can be reverted back to a prior saved version using an HX native snapshot. A native snapshot is a reproduction of a VM that includes the state of the data on all VM disks and the VM powerstate (on, off, or suspended) at the time the native snapshot is taken. Taking a native snapshot to save the current state of a VM gives you the ability to revert back to the saved state.

The following methodologies are used in the administration of HX native Snapshots:

- Support for HX native Snapshot in the vSphere client plug-in for HTML 5 was introduced in plugin version 2.0.0. For more information, see [Snapshot Now](#).
- Support for Schedule Snapshot the vSphere client plug-in for HTML 5 was introduced in plugin version 2.1.0. For more information, see [Schedule Snapshot](#)
- The vSphere “Manage Snapshots” function can revert to a specific HX native snapshot, or delete all snapshots.
- Cisco HyperFlex Connect can create on-demand and schedule HX native snapshots.
- The HyperFlex command line user interface can create HX native snapshots.
- HX REST APIs can create and delete HX native snapshots.
- Significant changes in Cisco HXDP Release 5.5(x) and later:
 - ESXi versions 6.5, 6.7 and 7.0 U1 are not supported.
 - VMware VAAI snapshot workflow is used instead of the Sentinel Snapshot Create workflow.

For additional information about VMware snapshots, see the "Overview of virtual machine snapshots in vSphere (KB 1015180)" on the VMware Customer Connect site.

