



Managing Nodes

- [Managing Nodes, on page 1](#)
- [Identify Node Maintenance Methods, on page 3](#)
- [Searching by DNS Address or Host Name, on page 5](#)
- [Changing ESXi Host Root Password, on page 6](#)
- [Reinstalling Node Software, on page 6](#)
- [Changing Node Identification Form in vCenter Cluster from IP to FQDN, on page 7](#)
- [Replacing Node Components, on page 8](#)
- [Removing a Node, on page 10](#)

Managing Nodes

Nodes are initially added to a storage cluster using the Create Cluster feature of the HX Data Platform Installer. Nodes are added to an existing storage cluster using the Expand Cluster feature of the HX Data Platform Installer. When nodes are added or removed from the storage cluster, the HX Data Platform adjusts the storage cluster status accordingly.

- Tasks for node maintenance with a failed node.
 - The ESXi or HX software needs to be reinstalled.
 - A node component needs to be replaced.
 - The node needs to be replaced.
 - The node needs to be removed.
- Tasks for node maintenance with a non-failed node.
 - Putting the node into maintenance mode.
 - Changing the ESX password.



Note Though there are subtle differences, the terms **server**, **host**, and **node** are used interchangeably throughout the HyperFlex documentation. Generally a server is a physical unit that runs software dedicated to a specific purpose. A node is a server within a larger group, typically a software cluster or a rack of servers. Cisco hardware documentation tends to use the term node. A host is a server that is running the virtualization and/or HyperFlex storage software, as it is 'host' to virtual machines. VMware documentation tends to use the term host.

Step 1 Monitor the nodes in the cluster.

HX storage cluster, node, and node component status is monitored and reported to HX Connect, HX Data Platform Plug-in, vCenter UI, and assorted logs as Operational status (online, offline) and Resiliency (healthy, warning) status values.

Note Functional state distinctions contribute to, but are separate from, the storage cluster operational and resiliency status reported in the HX Connect and HX Data Platform Plug-in views. For each Data Replication Factor (2 or 3), Cluster Access Policy (lenient or strict), and given number of nodes in the storage cluster, the storage cluster shifts between Read and Write, Read Only, or Shutdown state, depending on the number of failed nodes or failed disks in nodes.

Note 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 and should not be used in a production environment. The risk of outage due to component or node failures should be mitigated by having active and regular backups.

Step 2 Analyze the node failure and determine the action to take.

This frequently requires monitoring the node state through HX Connect, HX Data Platform Plug-in, vCenter, or ESXi; checking the server beacons; and collecting and analyzing logs.

Step 3 Complete the identified tasks.

- Reinstall or upgrade software.

For steps to reinstall ESXi or the HX Data Platform see [Cisco HyperFlex Systems Installation Guide for VMware ESXi](#). For steps to upgrade software, see the [Cisco HyperFlex Systems Upgrade Guide](#).

- Repair a component in the node.

Node components, such as solid state drives (SSD), hard disk drives (HDD), power supply units (PSU), and network interface cards (NIC) components are not configurable through HX Connect or HX Data Platform Plug-in, but the HX Data Platform monitors them and adjusts the storage cluster status when any of these items are disrupted, added, removed, or replaced.

The steps to add or remove disks, depends upon the type of disk. Field replaceable units (FRUs), such as PSUs and NICs are replaced following steps described in the server hardware guides.

- To replace a node in the cluster, review the [Replacing Node Components, on page 8](#).

Replacing a node in a storage cluster typically requires TAC assistance. Provided the requirements are met, nodes can be replaced without TAC assistance while the storage cluster is online (5+ node clusters only) or offline (4+ node clusters).

- To remove a node from the cluster, review the [Removing a Node from an Online Storage Cluster, on page 13](#) or [Removing a Node from an Offline Storage Cluster, on page 15](#)

Note Removing the node must not reduce the number of available nodes below the minimum 3 nodes, as this makes the storage cluster unhealthy. To remove a node in a 3 node cluster always requires TAC assistance. You can remove a maximum of 2 nodes from an offline cluster.

Identify Node Maintenance Methods

When performing maintenance tasks on nodes, some of these tasks are performed while the storage cluster is offline, others can be performed while the cluster is online and only require that the node is in HXDP Maintenance Mode.

- **Online tasks** - require that the storage cluster is healthy before the task begins.

- **Offline tasks** - require that the storage cluster will be shutdown.

If 2 or more nodes are down, then the storage cluster is automatically offline.

- **TAC assisted tasks** - typically require steps that are performed by the TAC representative.

The following tables lists the methods available to perform the associated node maintenance task.

Repair Node Software

ESX and HX Data Platform software is installed on every node in the storage cluster. If it is determined after node failure analysis that either software item needs to be re-installed, see the [Cisco HyperFlex Systems Installation Guide for VMware ESXi](#). For steps to upgrade software, see the [Cisco HyperFlex Systems Upgrade Guide](#).

Repair Node Hardware

A reparable item on node fails. This includes FRUs and disks. Some node components require TAC assistance. Replacing a node's mother board, for example, requires TAC assistance.

No. Nodes in Cluster	No. Failed Nodes in Cluster	Method	Notes
3	1 or more	TAC assisted only node repair.	Node does not need to be removed to perform repair. Includes replacing disks on node.
4-8	1	Online or Offline node repair.	Node does not need to be removed to perform repair. Includes replacing disks on node.

Remove Node

A non-reparable item on node fails. Disks on the removed node are not reused in the storage cluster.

No. Nodes in Cluster	No. Failed Nodes in Cluster	Method	Notes
4	1	Offline node remove.	A 4 node cluster with 2 nodes down, requires TAC assistance.
5 or more	1	Online or Offline node remove.	
5 or more	2	Offline 2 node remove.	A 5 node cluster with 3 nodes down, requires TAC assistance.

Replace Node and Discard Storage

A non-reparable item on node fails. Disks on the removed node are not reused in the storage cluster.

No. Nodes in Cluster	No. Failed Nodes in Cluster	Method	Notes
3	1	TAC assisted only node replace.	TAC assisted node replacement required to return cluster to minimum 3 nodes. A 3 node cluster with 1 node down, requires TAC assistance.
4	1	Offline replace node. Not reusing the disks.	Use Expand cluster to add new nodes. All other nodes must be up and running. A 4 node cluster with 2 nodes down, requires TAC assistance.
5 or more	1	Online or offline replace node. Not reusing the disks.	Use Expand cluster to add new nodes. All other nodes must be up and running.
5 or more	2	Offline replace 1 or 2 nodes. Not reusing the disks.	Use Expand cluster to add new nodes. All other nodes must be up and running. Replacing up to 2 nodes is supported. Replacing 3 or more nodes requires TAC assistance.

Replace Node and Reuse Storage

A non-reparable item on node fails. Disks on the removed node are reused in the storage cluster.

No. Nodes in Cluster	No. Failed Nodes in Cluster	Method	Notes
3 or more	1 or more	TAC assisted only.	<p>TAC assisted node replacement required to return cluster to minimum 3 nodes.</p> <p>Note Reusing disks requires assigning old node UUID to new node. Disks UUIDs to node UUID relationship is fixed and cannot be reassigned. This is a TAC assisted task.</p>

Searching by DNS Address or Host Name

Sometimes for troubleshooting purposes it is useful to be able to search by the DNS server address or DNS server host name. This is an optional task.

Step 1

Assign DNS search addresses

- Log into the HX Data Platform Installer virtual machine. Use either `ssh` or the vSphere console interface.
- Edit `resolv.conf.d` file.

```
# vi /etc/resolvconf/resolv.conf.d/base
```

- Confirm the change.

```
# resolvconf -u
# cat /etc/resolv.conf
```

- Confirm the DNS server can be queried from either the IP address or the host name.

```
# nslookup ip_address
# nslookup newhostname
```

Step 2

Assign a DNS host name.

- Log into the HX Data Platform Installer virtual machine. Use either `ssh` or the vSphere console interface.
- Open the hosts file for editing.

```
# vi /etc/hosts
```

- Add the following line and save the file.

```
ip_address ubuntu newhostname
```

For each host `ip_address`, enter the host `newhostname`.

- Add the `newhostname` to `hostname`.

```
# hostname newhostname
```

Changing ESXi Host Root Password

You can change the default ESXi password for the following scenarios:

- During creation of a standard and stretch cluster (supports only converged nodes)
- During expansion of a standard cluster (supports both converged or compute node expansion)
- During Edge cluster creation



Note In the above cases, the ESXi root password is secured as soon as installation is complete. In the event a subsequent password change is required, the procedure outlined below may be used after installation to manually change the root password.

As the ESXi comes up with the factory default password, you should change the password for security reasons. To change the default ESXi root password post-installation, do the following.



Note If you have forgotten the ESXi root password, for password recovery please contact Cisco TAC.

Step 1 Log into the ESXi host service control using SSH.

Step 2 Acquire root privileges.

```
su -
```

Step 3 Enter the current root password.

Step 4 Change the root password.

```
passwd root
```

Step 5 Enter the new password, and press **Enter**. Enter the password a second time for confirmation.

Note If the password entered the second time does not match, you must start over.

Reinstalling Node Software

To re-install software on a node that is a member of an existing storage cluster, contact TAC. This task must be performed with TAC assistance.

Step 1 Reinstall ESX following the directions from TAC.

Ensure the server meets the required hardware and configuration listed in Host ESX Server Setting Requirements. HX configuration settings are applied during the HX Data Platform process.

Step 2 Reinstall HX Data Platform, following the directions from TAC.

The HX Data Platform must always be re-installed after ESX is re-installed.

Changing Node Identification Form in vCenter Cluster from IP to FQDN

This task describes how to change how vCenter identifies the nodes in the cluster, from IP address to Fully Qualified Domain Name (FQDN).

Step 1 Schedule a maintenance window to perform this task.

Step 2 Ensure the storage cluster is healthy.

Check the storage cluster status through either HX Connect, HX Data Platform Plug-in, or from the `stcli cluster info` command on the storage controller VM.

Step 3 Lookup the FQDN for each ESXi host in the storage cluster.

a) From the ESXi host command line.

```
# cat /etc/hosts
```

In this example, the FQDN is `sjs-hx-3-esxi-01.sjs.local`.

```
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1        localhost.localdomain localhost
::1             localhost.localdomain localhost
172.16.67.157   sjs-hx-3-esxi-01.sjs.local sjs-hx-3-esxi-01
```

b) Repeat for each ESXi host in the storage cluster.

Step 4 Verify the FQDNs for each ESXi host are resolvable from vCenter, each other ESXi host, and the controller VMs.

a) From the vCenter command line.

```
# nslookup <fqdn_esx_host1>
# nslookup <fqdn_esx_host2>
# nslookup <fqdn_esx_host3>
...
```

b) Repeat for each ESXi host from an ESXi host.

c) Repeat for each ESXi host from each controller VM.

Step 5 If the FQDN name is not resolvable, then verify the DNS configuration on each ESXi host and each controller VM.

a) Check that the controller VMs have the correct IP address for the DNS server.

From a controller VM command line.

```
# stcli services dns show
10.192.0.31
```

a) Check the ESXi hosts have the same DNS configuration as the controller VMs.

From vCenter, select each ESXi host then **Configuration > DNS Servers**.

Step 6 Locate and note the Datacenter Name and the Cluster Name.

From vCenter client or web client, scroll through to see the Datacenter Name and Cluster Name. Write them down. They will be used in a later step.

Step 7 Delete the **cluster** from vCenter.

From vCenter, select **datacenter** > **cluster**. Right-click the **cluster** and select **Delete**.

Note Do not delete the **datacenter**.

Step 8 Recreate the **cluster** in vCenter.

- From vCenter, right-click the **datacenter**. Select **New Cluster**.
- Enter the exact same name for the **Cluster Name** as the cluster you deleted. This is the name you wrote down from Step 6.

Step 9 Add ESXi hosts (nodes) to the **cluster** using the FQDN name. Perform these steps for all ESXi hosts.

- From vCenter, right-click the **datacenter** > **cluster**. Select **Add Host**.
- Select an ESXi host using their FQDN.
- Repeat for each ESXi host in the cluster.

Step 10 Reregister the cluster with vCenter.

```
# # stcli cluster reregister
--vcenter-datacenter <datacenter_name>
--vcenter-cluster <hx_cluster_name>
--vcenter-url <FQDN_name>
--vcenter-user <vCenter_username>
--vcenter-password <vCenter_Password>
```

The SSO URL is not required for HX version 1.8.1c or later. See [Registering a Storage Cluster with a New vCenter Cluster](#) for additional information on reregistering a cluster.

Step 11 Enable VMware cluster HA and DRS using the post install script:

- Log into the HX cluster IP as admin and run the command **# hx_post_install**.
- Select Option 1 - "New/Existing Cluster" and input all login credentials
- Type "y" if you want to enter a new license key
- Type "y" to enable HA and DRS in the cluster
- Select 'n' for all other options and exit the script.

Replacing Node Components

Selected components on a node can be replaced. Some components can be replaced while the node is up and running. Replacing some components requires that the node be placed into a maintenance mode and shutdown. Refer to the hardware installation guide for your specific server for a complete list of field replaceable units (FRUs). Some components cannot be replaced or can only be replaced with TAC assistance. The following is a general list of components that can be replaced in a node.



Note When disks are removed, the disk UUIDs continue to be listed, even when not physically present. To reuse disks on another node in the same cluster see TAC for assistance.

- Components that do not require the node be shutdown. These are hot-swappable.
 - HDD data drives. Front bays
See [Managing Disks](#) for the storage cluster tasks and the hardware installation guides for the hardware focused tasks. Both sets of tasks are required to replace this component.
 - SSD cache drive. Front bay 1
See [Managing Disks](#) for the storage cluster tasks and the hardware installation guides for the hardware focused tasks. Both sets of tasks are required to replace this component.
 - Fan Modules
See the hardware installation guides to replace this component.
 - Power Supplies
See the hardware installation guides to replace this component.
- Components that do required the node be put into maintenance mode and shutdown.
For all of the following components, see the hardware installation guides.
 - Housekeeping SSD
Both the storage cluster tasks, and hardware focused tasks are required to replace this component.
 - RTC Battery on motherboard



Note The motherboard itself is not a replaceable component. You must purchase a battery from your local hardware store and replace it.

- DIMMS
- CPUs and Heatsinks
- Internal SD Card
- Internal USB Port
- Modular HBA Riser (HX 220c servers)
- Modular HBA Card
- PCIe Riser Assembly
- PCIe Card
- Trusted Platform Module
- mLOM Card

- RAID Controller
- Virtual Interface Card (VIC)
- Graphic Processing Unit (GPU)

Removing a Node

Removing a node is supported on the following cluster types:

Table 1: Cluster Types that Support Node Removal

Cluster Type	Converged	Compute
Standard	Yes	Yes
Stretch	No	Yes
Edge	Yes (see Note) Note Removing a node (compute or converged) is supported only on a Edge clusters with more than 3 nodes. For Edge clusters with 4 nodes, you will need to follow the offline node removal process. For Edge clusters with 5 or more nodes, the online and offline node removal process is supported. The online node removal method is however recommended .	

Depending upon the number of nodes in a cluster, you can remove a node when the cluster is either online or you need to make the cluster offline. Before you do so, you must first ensure that you have completed the required preparation steps.

The affecting context is based on the number of converged nodes. The number of compute nodes does not affect the process to remove a node.

You can only remove 1 converged node at any time.

For clusters with 4 converged nodes, follow the offline node removal process. For clusters with 5 converged nodes or more, follow the online node removal process.



Note Removing a converged node from a 3-node cluster is not supported



Note If you remove a node when the cluster is offline, you cannot add the node back to the cluster.

Prior to removing a node or nodes for HyperFlex clusters with Logical Availability Zones (LAZ) configured, LAZ must be disabled.

If LAZ is utilized in the HyperFlex cluster, then the number of remaining nodes must be in a balanced configuration that supports LAZ per the [LAZ Guidelines and Considerations](#) prior to reenabling LAZ.

Preparing to Remove a Node

Before you remove a node from a storage cluster, complete the following steps.

Step 1 Ensure the cluster is healthy.

```
# stcli cluster info
```

Example response that indicates the storage cluster is online and healthy:

```
locale: English (United States)
state: online
upgradeState: ok
healthState: healthy
state: online
state: online
```

Step 2 Ensure that SSH is enabled in ESX on all the nodes in the storage cluster.

Step 3 Ensure that the Distributed Resource Scheduler (DRS) is enabled.

DRS migrates only powered-on VMs. If your network has powered-off VMs, you must manually migrate them to a node in the storage cluster that will not be removed.

Note If DRS is not available then manually move the Virtual Machines from the node.

Step 4 Make a note of zkEnsemble. It contains the data IP of controller VMs (CVM).

Example:

```
admin:~$ cat /etc/springpath/storfs.cfg | grep -i ense
crmZKEnsemble=10.104.18.37:2181,10.104.18.38:2181,10.104.18.39:2181, 10.104.18.40:2181,
10.104.18.41:2181
```

If the node which was removed was ucs-308 whose CVM data IP is 10.104.18.40, then that CVMs data IP should no longer appear after you run the above command after the node removal.

Step 5 Put the node to be removed into Maintenance mode. Choose a method: vSphere GUI, controller VM command line (CLI) or HyperFlex Connect System Information panel:

GUI

a) Right-click each host, scroll down the list, and select **Maintenance Mode > Enter Maintenance Mode**.

The vSphere Maintenance Mode option is at the top of the host right-click menu. Be sure to scroll to the bottom of the list to select Maintenance Mode.

b) In HX Connect, from the **MANAGE > System Information panel's Node tab**, select the node, and then click on the button to **Enter HXDP Maintenance Mode**.

CLI

a) Log in to a controller VM as an admin user.

b) Run **stcli cluster info** and look for `stNodes:` section

```
# stcli cluster info

stNodes:
-----
type: node
id: 689324b2-b30c-c440-a08e-5b37c7e0eeef
name: ucs-305
```

```

-----
type: node
id: 9314ac70-77aa-4345-8f35-7854f71a0d0c
name: ucs-306
-----
type: node
id: 9e6ba2e3-4bb6-214c-8723-019fa483a308
name: ucs-307
-----
type: node
id: 575ace48-1513-0b4f-bfe1-e6abd5ff6895
name: ucs-308
-----
type: node
id: 875ebe8-1617-0d4c-afe1-c6aed4ff6732
name: ucs-309

```

Under the `stNodes` section, the `id` is listed for each node in the cluster. Find the node `id` or `name` you need to remove.

- c) Move the ESX host into Maintenance mode.

```
# stcli node maintenanceMode (--id ID | --ip NAME) --mode enter
```

(see also `stcli node maintenanceMode --help`)

For example, to remove node `ucs-308`:

Example:

```
stcli node maintenanceMode -id 575ace48-1513-0b4f-bfe1-e6abd5ff6895
or
stcli node maintenanceMode -ip 10.104.18.40
```

Step 6

Wait for 2 hours, monitor healing info in `stcli cluster storage-summary`. You should wait until you see "Storage cluster is healthy." as shows in the following example:

Example:

```
admin:$ stcli cluster storage-summary | grep -i heali -A 8
healingInfo:
  inProgress: False
resiliencyInfo:
  messages:
    -----
    Storage node 10.104.18.40 is unavailable.
    -----
    Storage cluster is healthy.
    -----
```

Before the healing starts you will see following:

```
admin:$ date; stcli cluster storage-summary | grep -i heali -A 8
Thu Sep 30 12:33:57 PDT 2021
healingInfo:
  inProgress: False
resiliencyInfo:
  messages:
    -----
    Storage cluster is unhealthy .
    -----
    Storage node 10.104.18.40 is unavailable .
    -----
```

After 2 hours + you will see following:

```
admin:$ stcli cluster storage-summary | grep -i heali -A 8
healingInfo:
```

```

messages:
Space rebalancing in progress, 83 % completed.
InProgress: True
percentComplete: 83
estimatedCompletionTimeInSeconds: 211
resiliencyInfo:
messages:

```

What to do next

Proceed to Removing a Node. Choose the Online or Offline method based on the number of nodes in your storage cluster.

Removing a Node from an Online Storage Cluster

Use the `stcli node remove` to clean up a deployment or remove a node from a storage cluster.



Note You can remove multiple nodes in a series, as long as it is done one at a time and when the cluster is healthy between each successive node removal. You must also have followed the steps required to prepare to remove a node. For more information, see [Preparing to Remove a Node, on page 11](#).



Note Prior to removing a node or nodes for HyperFlex clusters with Logical Availability Zones (LAZ) configured, LAZ must be disabled.

If LAZ is utilized in the HyperFlex cluster, then the number of remaining nodes must be in a balanced configuration that supports LAZ per the [LAZ Guidelines and Considerations](#) prior to reenabling LAZ.



Note Do not remove the controller VM or other HX Data Platform components before you complete the steps in this task.

Step 1

Run the `stcli cluster info` command and look for `stNodes:` section to find the node which needs to be removed.. This information is also available when you put the node in maintenance mode.

Example:

```

-----
stNodes:
type: node
id: 689324b2-b30c-c440-a08e-5b37c7e0eefe
name: ucs305

type: node
id: 9314ac70-77aa-4345-8f35-7854f71a0d0c
name: ucs306

type: node
id: 9e6ba2e3-4bb6-214c-8723-019fa483a308

```

```
name: ucs307

type: node
id: 575ace48-1513-0b4f-bfe1-e6abd5ff6895
name: ucs308

type: node
id: 875ebe8-1617-0d4c-af
name: ucs 309
```

The **stcli node remove** command to remove nodes from the 5-node cluster are:

- **stcli node remove --ip-1 ucs 308** or
- **stcli node remove --id-1 575ace48-1513-0b4f-bfe1-e6abd5ff6895**

After the **stcli node remove** command completes successfully, the system rebalances the storage cluster until the storage cluster state is Healthy. Do not perform any failure tests during this time. The storage cluster remains healthy.

Because the node is no longer in the storage cluster, you do not need to exit HXDP Maintenance Mode.

Note It is highly recommended that you work with TAC when removing a converged node in a storage cluster. Do not reuse a removed converged node or its disks in the original cluster.

Note If you want to reuse a removed node in another storage cluster, contact Technical Assistance Center (TAC). Additional steps are required to prepare the node for another storage cluster.

Step 2 Confirm that the node is removed from the storage cluster.

a) Check the storage cluster information.

```
# stcli cluster storage-summary
```

b) Check the **ActiveNodes** entry in the response to verify the cluster has one less node.

c) Check that the node which was removed is not part of Ensemble. For example:

Example:

```
admin:~$ cat /etc/springpath/storfs.cfg | grep -i ense
crmZKEnsemble=10.104.18.37:2181,10.104.18.38:2181,10.104.18.39:2181, 10.104.18.40:2181,
10.104.18.41:2181
```

For example, if the node which was removed was ucs-308 whose CVM data IP is 10.104.18.40, then that CVMs data IP should no longer appear after you run the above command after the node removal as seen above.

If there are more than 5 nodes and the removed node was part of ensemble, then the new node ip appears in the **crmZKEnsemble**. For example, if the cluster initially has 7 nodes (10.104.18.37 to 10.104.18.43), and **crmZKEnsemble** has 10.104.18.37:2181,10.104.18.38:2181,10.104.18.39:2181, 10.104.18.40:2181, 10.104.18.41:2181, then after removal of 10.104.18.40, **crmZKEnsemble** has either:

10.104.18.37:2181,10.104.18.38:2181,10.104.18.39:2181, 10.104.18.42:2181, 10.104.18.41:2181, or:

10.104.18.37:2181,10.104.18.38:2181,10.104.18.39:2181, 10.104.18.43:2181, 10.104.18.41:2181

Step 3 Verify that disks from the removed node no longer appear by running the **hxcli disk list** command:

```
admin:~$ hxcli disk list --no-loader
+-----+-----+-----+-----+-----+-----+-----+
| NODE NAME | HYPERVSIOR STATUS | SLOT | CAPACITY | STATUS | TYPE | USAGE |
+-----+-----+-----+-----+-----+-----+-----+
| ucs305 | ONLINE | 1 | 111.8 GB | Claimed | Solidstate | System |
```

```

| ucs305 | ONLINE | 2 | 894.3 GB | Claimed | Solidstate | Caching |
| ucs305 | ONLINE | 3 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 4 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 5 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 6 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 7 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 8|0B| Unknown | | |
| ucs306 | ONLINE | 1 | 111.8 GB | Claimed | Solidstate | System |
| ucs306 | ONLINE | 2 | 894.3 GB | Claimed | Solidstate | Caching |
| ucs306 | ONLINE | 3 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 4 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 5 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 6 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 7 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 8|0B| Unknown | | |
| ucs307 | ONLINE | 1 | 111.8 GB | Claimed | Solidstate | System |
| ucs307 | ONLINE | 2 | 894.3 GB | Claimed | Solidstate | Caching |
| ucs307 | ONLINE | 3 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 4 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 5 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 6 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 7 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 8|0B| Unknown | | |
| ucs309 | ONLINE | 1 | 111.8 GB | Claimed | Solidstate | System |
| ucs309 | ONLINE | 2 | 894.3 GB | Claimed | Solidstate | Caching |
| ucs309 | ONLINE | 3 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs309 | ONLINE | 4 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs309 | ONLINE | 5 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs309 | ONLINE | 6 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs309 | ONLINE | 7 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs309 | ONLINE | 8|0B| Unknown | | |
+-----+-----+-----+-----+-----+-----+-----+

```

For example, if you removed ucs-308, then its disks no longer appear.

Step 4 Remove the host from the vCenter **Hosts and Cluster** view.

- Log in to vSphere Web Client Navigator. Navigate to **Host** in the vSphere Inventory.
- Right-click on the host and select **All vCenter Actions > Remove from Inventory**. Click **Yes**.

Step 5 Confirm that all the node-associated datastores are removed. For example, run the following command in ESXi:

```

[root@ucs308:~] esxcfg-nas -l
ds4 is 169.254.226.1:ds4 from 6894152532647392862-8789011882433394379 mounted available
ds3 is 169.254.226.1:ds3 from 6894152532647392862-8789011882433394379 mounted available
ds2 is 169.254.226.1:ds2 from 6894152532647392862-8789011882433394379 mounted available
ds5 is 169.254.226.1:ds5 from 6894152532647392862-8789011882433394379 mounted available
ds1 is 169.254.226.1:ds1 from 6894152532647392862-8789011882433394379 mounted available

```

Note If any node-associated datastores are listed, then unmount and delete those datastores manually.

Removing a Node from an Offline Storage Cluster

Use the `stcli node remove` to clean up a deployment or remove a node from a storage cluster.



Note It is highly recommended that you work with TAC when removing a converged node in a storage cluster.

Number of nodes in cluster	Number of failed nodes in cluster	Method	Notes
3	1	See TAC to remove and replace the node.	-
4	1	Offline node remove. Note If " <i>storage cluster manager is not configured</i> " error is seen in <code>hxconnect</code> or as part of <code>stcli cluster storage-summary grep -i heali -A 8</code> command, Contact TAC for assistance.	A 4 node cluster with 2 nodes down, requires TAC assistance.
5 or more	1	Cluster must be offline.	Online mode is recommended.
5 or more	2	Cluster can be online.	A 5 node cluster with 3 nodes down, requires TAC assistance.



Note Do not remove the controller VM or other HX Data Platform components before you complete the steps in this task.

Step 1 Follow the process to prepare for removing a node. For more information, see [Preparing to Remove a Node, on page 11](#).

Step 2 (For 4-node clusters only) Prepare to shutdown, then shutdown the storage cluster.

a) Gracefully shutdown all resident VMs on all the HX datastores.

Optionally, vMotion the VMs.

b) Gracefully shutdown all VMs on non-HX datastores on HX storage cluster nodes, and unmount.

c) From any controller VM command line, issue the `stcli cluster shutdown` command.

```
# stcli cluster shutdown
```

Step 3 Run the `stcli cluster info` command and look for `stNodes:` section to find the node which needs to be removed.. This information is also available when you put the node in maintenance mode.

Example:

```
-----
type: node
id: 569c03dc-9af3-c646-8ac5-34b1f7e04b5c
name: example1
-----
type: node
id: 0e0701a2-2452-8242-b6d4-bce8d29f8f17
```



```

name: example2
-----
type: node
id: a2b43640-cf94-b042-a091-341358fdd3f4
name: example3
-----
type: node
id: d2d43691-daf5-50c4-d096-941358fede374
name: example5

```

Step 4 Remove the desired node using the `stcli node remove` command.

For example:

To remove 1 node

- **stcli node remove -ip-1 example5** or
- **stcli node remove -id-1 d2d43691-daf5-50c4-d096-941358fede374**

Response:

```
Successfully removed node: EntityRef(type=3, id='', name='10.10.2.4')
```

This command unmounts all datastores, removes from the cluster ensemble, resets the EAM for this node, stops all services (stores, cluster management IP), and removes all firewall rules.

This command does not remove the node from vCenter. The node remains in vCenter. This command also does not remove the installed HX Data Platform elements, such as the controller VM.

Due to the node no longer being in the storage cluster, you do not need to exit HXDP Maintenance Mode.

Note If you want to reuse a removed node in another storage cluster, contact Technical Assistance Center (TAC). Additional steps are required to prepare the node for another storage cluster.

Step 5 Restart the cluster.

```
# hxcli cluster start
```

Step 6 Confirm that the node is removed from the storage cluster once the cluster is up.

a) Check the storage cluster information.

```
# stcli cluster storage-summary
```

- b) Check the `ActiveNodes` entry in the response to verify the cluster has one less node.
- c) Check that the node which was removed is not part of Ensemble.

Example:

```
admin:~$ cat /etc/springpath/storfs.cfg | grep -i ense
crmZKEnsemble=10.104.18.37:2181,10.104.18.38:2181,10.104.18.39:2181, 10.104.18.40:2181,
10.104.18.41:2181
```

For example, if you removed 10.104.18.40, note that it no longer appears.

- d) If the node remove action is successful for a cluster with 5 or fewer nodes, the ensemble may contain 4 participant nodes, and 1 observer node. The only participant nodes are updated in `stMgr.cfg`; the observer node details are not updated.

Example:

```
Ensemble after node remove server.0=10.107.16.111:2888:3888:participant;10.107.16.111:2181
```

```
server.1=10.107.16.107:2888:3888:participant;10.107.16.107:2181
```

```
server.2=10.107.16.108:2888:3888:participant;10.107.16.108:2181
server.4=10.107.16.110:2888:3888:participant;10.107.16.110:2181
server.5=10.107.16.109:2888:3888:observer;10.107.16.109:2181
version=10000558b
```

Step 7 Verify that disks from the removed node no longer appear by running the `hxcli disk list` command:

```
admin:~$ hxcli disk list --no-loader
+-----+-----+-----+-----+-----+-----+-----+-----+
| NODE NAME | HYPERVSIOR | STATUS | SLOT | CAPACITY | STATUS | TYPE | USAGE |
+-----+-----+-----+-----+-----+-----+-----+-----+
| ucs305 | ONLINE | 1 | 111.8 GB | Claimed | Solidstate | System |
| ucs305 | ONLINE | 2 | 894.3 GB | Claimed | Solidstate | Caching |
| ucs305 | ONLINE | 3 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 4 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 5 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 6 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 7 | 1.1 TB | Ignored | Rotational | Persistence |
| ucs305 | ONLINE | 8 | 0 B | Unknown | | |
| ucs306 | ONLINE | 1 | 111.8 GB | Claimed | Solidstate | System |
| ucs306 | ONLINE | 2 | 894.3 GB | Claimed | Solidstate | Caching |
| ucs306 | ONLINE | 3 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 4 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 5 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 6 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 7 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs306 | ONLINE | 8 | 0 B | Unknown | | |
| ucs307 | ONLINE | 1 | 111.8 GB | Claimed | Solidstate | System |
| ucs307 | ONLINE | 2 | 894.3 GB | Claimed | Solidstate | Caching |
| ucs307 | ONLINE | 3 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 4 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 5 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 6 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 7 | 1.1 TB | Claimed | Rotational | Persistence |
| ucs307 | ONLINE | 8 | 0 B | Unknown | | |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

For example, if you removed ucs-308, then its disks no longer appear.

Step 8 Remove the host from the vCenter **Hosts and Cluster** view.

- Log in to vSphere Web Client Navigator. Navigate to **Host** in the vSphere Inventory.
- Right-click on the host and select **All vCenter Actions > Remove from Inventory**. Click **Yes**.

Step 9 Confirm that all the node-associated datastores are removed. For example, run the following command in ESXi:

```
[root@ucs308:~] esxcfg-nas -l
ds4 is 169.254.226.1:ds4 from 6894152532647392862-8789011882433394379 mounted available
ds3 is 169.254.226.1:ds3 from 6894152532647392862-8789011882433394379 mounted available
ds2 is 169.254.226.1:ds2 from 6894152532647392862-8789011882433394379 mounted available
ds5 is 169.254.226.1:ds5 from 6894152532647392862-8789011882433394379 mounted available
ds1 is 169.254.226.1:ds1 from 6894152532647392862-8789011882433394379 mounted available
```

Note If any node-associated datastores are listed, then unmount and delete those datastores manually.

Removing a Compute Node

-
- Step 1** Migrate all the VMs from a compute node that needs to be removed.
- Step 2** Unmount the datastore from the compute node.
- Step 3** Check if the cluster is in the healthy state, by running the following command:
- ```
stcli cluster info --summary
```
- Step 4** Put ESXi host in the HXDP Maintenance Mode.
- Step 5** Remove the compute node using the `stcli node remove` command, from CMIP (use the Cisco HX connect IP address as it is the cluster IP address).
- ```
stcli node remove --id-1
```
- Or
- ```
stcli node remove --ip-1
```
- Where, IP is the IP address of the node to be removed.
- Step 6** Remove any DVS from the ESXi host in vCenter, if there is a DVS.
- Step 7** Remove the ESXi host from vCenter.
- Step 8** Check if the cluster is in the healthy state, by running the following command:
- ```
stcli cluster info --summary
```
- Step 9** If compute node virtnode entry still exists in `stcli cluster info` output, perform the following:
Restart the stMgr management service using `priv service stMgr restart` on the SCVMs.
- Step 10** Clear stale entries in the compute node by logging out of Cisco HX Connect and then logging into Cisco HX Connect.
- Step 11** Disable and re-enable the High Availability (HA) and Distributed Resource Scheduler (DRS) services to reconfigure the services after node removal.
-

Reuse a Previously Removed Node Within the Same Cluster

To reuse a node within the same cluster that was previously removed, perform the following steps:

Before you begin

- Nodes must be removed on a cluster running HXDP 4.5(2b) or later in order to be reused in the same cluster.
- The HX Node must be removed only while the cluster is online. If the node is removed while the cluster is offline, then the node may not be re-used within the same cluster.
- The HX Node removal must be performed with the `stcli node` command listed within the administration guide. If the node is not properly removed from the cluster then it may not be reused within the same cluster.
- For a 4 node cluster, use [Removing a Node from an Offline Storage Cluster, on page 15](#) only. Node reuse is not supported for 4 node clusters.

-
- Step 1** Remove the ESXi host from the vCenter Cluster inventory.
- Step 2** Re-install ESXi with the same version that matches the rest of the HX Cluster.
- Step 3** Delete only the UCS Service Profile from UCS Manager that was previously associated with this removed node.
- Step 4** Use the HX Installer of the same version and run an Expansion Workflow.

Note Be sure to check the **Clear Disk Partitions** box during the expansion workflow.
