



Managing HX Storage Clusters

- [Changing the Cluster Access Policy Level, on page 1](#)
- [Rebalancing the Cluster, on page 1](#)
- [Handling Out of Space Errors, on page 3](#)
- [Planning to Move a Storage Cluster Between vCenters, on page 4](#)
- [Renaming Clusters, on page 8](#)
- [Replacing Self-Signed Certificate, on page 8](#)

Changing the Cluster Access Policy Level

Step 1 The storage cluster must be in a healthy state prior to changing the Cluster Access Policy to strict.

Step 2 From the command line of a storage controller VM in the storage cluster, type:

```
# stcli cluster get-cluster-access-policy  
  
# stcli cluster set-cluster-access-policy --name {strict,lenient}
```

Rebalancing the Cluster

The storage cluster is rebalanced on a regular schedule. It is used to realign the distribution of stored data across changes in available storage and to restore storage cluster health. If you add or remove a node in the storage cluster, you can manually initiate a storage cluster rebalance using the `stcli rebalance` command.



Note Rebalancing might take some time depending on the disk capacity used on the failed node or disk.

Step 1 Start rebalancing the storage cluster.

- a) Login to a controller VM in the storage cluster.
- b) From the controller VM command line, run the command:

```
# stcli rebalance start --force
```

Step 2 Verify rebalancing status from the storage controller VM.

a) Enter the following on the command line:

```
# stcli rebalance status
rebalanceStatus:
rebalanceState:
cluster_rebalance_ongoing
percentComplete: 10
rebalanceEnabled: True
```

b) Reenter the command line to confirm the process completes:

```
# stcli rebalance status
rebalanceStatus:
rebalanceState: cluster_rebalance_not_running
rebalanceEnabled: True
```

This sample indicates that `rebalance` is enabled, and ready to perform a rebalance, but is not currently rebalancing the storage cluster.

Checking Cluster Rebalance and Self Healing Status

The storage cluster is rebalanced on a regular schedule and when the amount of available storage in the cluster changes. A rebalance is also triggered when there is a change in the amount of available storage. This is an automatic self healing function.



Important Rebalance typically occurs only when a single disk usage exceeds 50% or cluster aggregate disk usage is greater than 50%.

You can check rebalance status through the HX Data Platform plug-in or through the storage controller VM command line.

Step 1 Check the rebalance status through HX Data Platform plug-in.

a) From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Summary**.

The **Status** portlet lists the **Self healing status**.

b) Expand the 'Resiliency Status' to see the 'Self Healing status' section. The Self healing status field lists the rebalance activity or N/A, when rebalance is not currently active.

Step 2 Check the rebalance status through the storage controller VM command line.

a) Login to a controller VM using `ssh`.
b) From the controller VM command line, run the command.

```
# stcli rebalance status
```

The following output indicates that rebalance is not currently running on the storage cluster.

```
rebalanceStatus:  
percentComplete: 0  
rebalanceState: cluster_rebalance_not_running  
rebalanceEnabled: True
```

The Recent Tasks tab in the HX Data Platform plug-in displays a status message.

Handling Out of Space Errors

If your system displays an Out of Space error, you can either add a node to increase free capacity or delete existing unused VMs to release space.

When there is an Out of Space condition, the VMs are unresponsive.



Note Do not delete storage controller VMs. Storage controller VM names have the prefix `stCtlVM`.

Step 1 To add a node, use the Expand Cluster feature of the HX Data Platform Installer.

Step 2 To delete unused VMs, complete the following:

- a) Determine which guest VMs you can delete. You can consider factors such as disk space used by the VM or naming conventions.
- b) Go to **vCenter > Virtual Machines** to display the virtual machines in the inventory.
- c) Double-click a VM that you want to delete.
- d) Select the **Summary > Answer Questions** to display a dialog box.
- e) Click the **Cancel** radio button and click **OK**.
- f) Power off the VM.
- g) Delete the VM.

Step 3 After the Out of Space condition is cleared, complete the following:

- a) Go to **vCenter > Virtual Machines** to display the VM in the inventory.
 - b) Double-click a VM that you want to use.
 - c) Select the **Summary > Answer Questions** to display a dialog box.
 - d) Click the **Retry** radio button and click **OK**.
-

Checking Cleaner Schedule

The `stcli cleaner` command typically runs in the background continuously. `cleaner` goes into sleep mode when it is not needed and wakes when policy defined conditions are met. For example, if your storage cluster is experiencing ENOSPC condition, the cleaner automatically runs at High Priority.

Do not expand the cluster while the `cleaner` is running. Check the cleaner schedule or adjust the schedule, as needed.

Step 1 Login to any controller VM in the storage cluster. Run the listed commands from the controller VM command line.

Step 2 View the cleaner schedule.

```
# stcli cleaner get-schedule --id ID | --ip NAME
```

Parameter	Description
--id ID	ID of storage cluster node
--ip NAME	IP address of storage cluster node

Planning to Move a Storage Cluster Between vCenters

When you rename the vCenter datacenter or vCenter cluster, you must re-register the HX storage cluster.

Moving a storage cluster from one vCenter cluster to another requires the listed steps. See the following topics for detailed information.

1. Meet the prerequisites to this task. See [Moving the Storage Cluster from a Current vCenter Server to a New VCenter Server, on page 4](#).
2. Delete the cluster from the old vCenter, create a new cluster on the new vCenter. Use the same cluster name. See [Moving the Storage Cluster from a Current vCenter Server to a New VCenter Server, on page 4](#).
3. Unregister HX Data Platform using the vCenter Extension Manager. See [Unregistering a Storage Cluster from a vCenter Cluster, on page 5](#).
4. Use the `stcli cluster reregister` command to associate the HX Storage Cluster with a new vCenter. See [Registering a Storage Cluster with a New vCenter Cluster, on page 7](#).

Moving the Storage Cluster from a Current vCenter Server to a New VCenter Server

Before you begin

- If your HX Cluster is running HX Data Platform version older than 1.8(1c), upgrade before attempting to reregister to a new vCenter.
- Perform this task during a maintenance window.
- Ensure the cluster is healthy and upgrade state is OK and Healthy. You can view the state using the `stcli` command from the controller VM command line.

```
# stcli cluster info
```

Check response for:

```
upgradeState: ok
healthState: healthy
```

- Ensure vCenter must be up and running.
- Snapshot schedules are not moved with the storage cluster when you move the storage cluster between vCenter clusters.

-
- Step 1** From the current vCenter, delete the cluster.
This is the vCenter cluster specified when the HX storage cluster was created.
- Step 2** On the new vCenter, create a new cluster using the same cluster name.
- Step 3** Add ESX hosts to new vCenter in the newly created cluster.
-

What to do next

Proceed to [Unregistering a Storage Cluster from a vCenter Cluster, on page 5](#).

Unregistering a Storage Cluster from a vCenter Cluster

This step is optional and not required. It is recommended to leave the HX Data Platform Plug-in registration alone in the old vCenter.

Before you begin

As part of the task to move a storage cluster from one vCenter server to another vCenter server, complete the steps in [Moving the Storage Cluster from a Current vCenter Server to a New vCenter Server, on page 4](#).



Note

- If multiple HX clusters are registered to the same vCenter, do not attempt this procedure until all HX clusters have been fully migrated to different vCenter. Running this procedure is disruptive to any existing HX clusters registered to the vCenter.
-

-
- Step 1** Complete the steps in [Removing HX Data Platform Files from the vSphere Client, on page 5](#).
- Step 2** Complete the steps in [Verifying HX Cluster is Unregistered from vCenter, on page 6](#).
-

What to do next

Proceed to [Registering a Storage Cluster with a New vCenter Cluster, on page 7](#).

Removing HX Data Platform Files from the vSphere Client

This task is a step in unregistering a HX Storage Cluster from vCenter.

Remove the HX Data Platform files from the vSphere Client. Select a method.

Linux vCenter

- a) Login to the Linux vCenter server using `ssh` as a root user.
- b) Change to the folder containing the HX Data Platform Plug-in folder.

For vCenter 6.0

```
# cd /etc/vmware/vsphere-client/vc-packages/vsphere-client-serenity/
```

For vCenter 5.5

```
# cd /var/lib/just/vmware/vsphere-client/vc-packages/vsphere-client-serenity/
```

- c) Remove the HX Data Platform Plug-in folder and files.

```
# rm -rf com.springpath*
```

- d) Restart the vSphere Client.

```
# service vsphere-client restart
```

Windows vCenter

- a) Login to the Windows vCenter system command line using Remote Desktop Protocol (RDP).
- b) Change to the folder containing the HX Data Platform Plug-in folder.

```
# cd "%PROGRAMDATA%\VMware\VSphere Web Client\vc-packages\vsphere-client-serenity
```

- c) Remove the HX Data Platform Plug-in folder and files.

```
# rmdir /com.springpath*
```

- d) Open the Service screen.

```
# services.msc
```

- e) Restart the vSphere Web Client to logout of vCenter.

```
# serviceLogout
```

Verifying HX Cluster is Unregistered from vCenter

This task is a step in unregistering a HX Storage Cluster from vCenter.

Verify that the HX cluster is no longer on the old vCenter.

Before you begin

Complete the steps in:

- [Removing HX Data Platform Files from the vSphere Client, on page 5](#)
-

Step 1 Clear your cache before logging back into vCenter.

Step 2 Log out of the old vCenter.

Step 3 Log in again to the old vCenter and verify the HX Data Platform Plug-in has been removed.

Registering a Storage Cluster with a New vCenter Cluster

Before you begin

As part of the task to move a storage cluster from one vCenter server to another vCenter server, complete the steps in [Unregistering a Storage Cluster from a vCenter Cluster, on page 5](#).

Step 1 Login to a controller VM.

Step 2 Run the `stcli cluster reregister` command.

```
stcli cluster reregister [-h] --vcenter-datacenter NEWDATACENTER --vcenter-cluster NEWVCENTERCLUSTER
--vcenter-url NEWVCENTERURL [--vcenter-sso-url NEWVCENTERSSOURL] --vcenter-user
NEWVCENTERUSER
```

Apply additional listed options as needed.

Syntax	Description	Option	Required or Optional	Description
		--vcenter-cluster NEWVCENTERCLUSTER	Required	Name of the new vCenter cluster.
		--vcenter-datacenter NEWDATACENTER	Required	Name of the new vCenter datacenter.
		--vcenter-sso-url NEWVCENTERSSOURL	Optional	URL of the new vCenter SSO server. This is inferred from --vcenter-url, if not specified.
		--vcenter-url NEWVCENTERURL	Required	URL of the new vCenter, <vcentername>. Where <vcentername> can be FQDN or IP.
		--vcenter-user NEWVCENTERUSER	Required	User name of the new vCenter administrator. Enter vCenter administrator password when prompted.

Example response:

```
Reregister StorFS cluster with a new vCenter ...
Enter NEW vCenter Administrator password:
Waiting for Cluster creation to finish ...
```

If, after your storage cluster is re-registered, your compute only nodes fail to register with EAM, or are not present in the EAM client, and not under the resource pool in vCenter, then run the command below to re-add the compute only nodes:

```
# stcli node add --node-ips <computeNodeIP> --controller-root-password <ctlvm-pwd> --esx-username
<esx-user> --esx-password <esx-pwd>
```

Contact TAC for assistance if required.

Step 3 Re-enter your snapshot schedules.

Snapshot schedules are not moved with the storage cluster when you move the storage cluster between vCenter clusters.

Renaming Clusters

After you create a HX Data Platform storage cluster, you can rename it without disrupting any processes.



Note These steps apply to renaming the HX Cluster, not the vCenter cluster.

-
- Step 1** From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster** to rename.
- Step 2** Open the **Rename Cluster** dialog box. Either right-click on the storage cluster or click the **Actions** drop-down list at the top of the tab.
- Step 3** Select **Rename Cluster**.
- Step 4** Enter a new name for the storage cluster in the text field.
HX cluster names cannot exceed 50 characters.
- Step 5** Click **OK** to apply the new name.
-

Replacing Self-Signed Certificate

Replacing Self-Signed Certificate with External CA Certificate on a vCenter Server

Set the certMgmt mode in vCenter to **Custom** to add the ESXi hosts with third party certificate to vCenter.

Note By default, the certMgmt mode is **vmsa**. In the default **vmsa** mode, you can add only the ESX host with self signed certificates. If you try to add an ESX with CA certificate to a vCenter, it will not allow you to add the ESX host unless CA certificate is replaced with self-signed certificate.

To update the certMgmt mode:

- Select the vCenter server that manages the hosts and click **Settings**.
- Click **Advanced Settings**, and click **Edit**.
- In the **Filter** box, enter **certmgmt** to display only certificate management keys.
- Change the value of **vpxd.certmgmt.mode** to **custom** and click **OK**.
- Restart the vCenter server service.

To restart services, enter the following link in a browser and then click **Enter**:

`https://<VC URL>:5480/ui/services`



- Note** The behavior of host addition in vCenter varies according to the certificate and certMgmt mode.
- When the host has self-signed certificate with the certMgmt mode set to the default value of **vmrsa** in vCenter:
 - Only ESX host with self-signed certificate can be added.
 - The addition of ESX with third party CA certificate is not allowed.
 - If you try to add an ESX to a vCenter after replacing the self-signed certificate with a third party CA certificate, the system will prompt you to replace third party CA certificate with self-signed certificate. You can add the ESX host after replacing CA certificate with self-signed certificate.
 - When the host has self-signed certificate with the certMgmt mode set to **custom** in vCenter:
 - If you try to add an ESX to a vCenter after replacing the self-signed certificate with a third party CA certificate, the system throws an error: `ssl thumbprint mismatch and add host fails`. In this case, do the following to replace the third party CA certificate with the self-signed certificate:
 1. Place the host in the maintenance mode (MM mode).
 2. Replace the certified `ruicert.crt` and `ruicert.key` files with the backed up previous key and certificate.
 3. Restart the `hostd` and `vxpxd` service. The CA certificate comes up in the new node.
 4. Right-click and connect to vCenter. The host loses the CA certificate and gets replaced with self-signed certification in VMware.
 - When the host has third party CA certificate with the certMgmt mode set to the default value of **vmrsa** in vCenter:
 - ESX host with self-signed certificate can be added.
 - The addition of ESX with third party CA certificate is not allowed.
 - When the host has third party CA certificate with the certMgmt mode set to **custom** in vCenter:
 - ESX host with self-signed certificate cannot be added.
 - The self-signed certificate in ESX host needs to be replaced with a CA certificate of vCenter.

Replacing Self-Signed Certificate with External CA Certificate on a ESXi Host

Step 1 Generate the host certificate (`ruicert.crt`) and key (`ruicert.key`) files and send the files to the certificate authority.

Note Ensure that a proper hostname or FQDN of the ESX host is provided while generating the `ruicert.key` and `ruicert.crt` files.

Step 2 Replace the certified host certificate (`ruicert.crt`) and key (`ruicert.key`) files in the `/etc/vmware/ssl` directory on each ESXi host after taking backup of the original host certificate (`ruicert.crt`) and key (`ruicert.key`) files.

Note Replace host certificate (rui.crt) and key (rui.key) files in a rolling fashion by placing only one host in maintenance mode and then wait for the cluster to be healthy and then replace the certificates for the other nodes.

- a) Log in to the ESXi host from an SSH client with administrator privileges.
- b) Place the host in the maintenance mode (MM mode).
- c) Take a backup of the previous key and certificate to the rui.bak file in the /etc/vmware/ssl/ directory.
- d) Upload the new certified rui.crt and rui.key files to the /etc/vmware/ssl/ directory.
- e) Restart the hostd and vpxa service, and check the running status using the following commands:

```
/etc/init.d/hostd restart
/etc/init.d/vpxa restart
/etc/init.d/hostd status
/etc/init.d/vpxa status
```

- f) Reconnect the host to vCenter and exit the maintenance mode.

Note Repeat the same procedure on all the nodes. You can verify the certificate of each node by accessing it through web.

Reregistering a HyperFlex cluster

After adding all the hosts to the vCenter after replacing the certified files, reregister the HX cluster to the vCenter using the following command:

```
stcli cluster reregister
```

Recreating a Self-Signed Certificate

If you face any issue with the host certificate after replacing external CA certificate, you can recreate the self-signed certificate by executing the following procedure:

1. Log in to the ESXi host from an SSH client.
2. Delete the rui.key and rui.crt files from the /etc/vmware/ssl/ directory.
3. Recreate the self-signed certificate for the host using the following command:

```
/sbin/generate-certificates
```

4. Restart the hostd and vpxa service using the following commands:

```
/etc/init.d/hostd restart
/etc/init.d/vpxa restart
```