



Pre-Upgrade Validation Checks

- [HyperFlex Node Upgrade Validations, on page 1](#)
- [Viewing HyperFlex Cluster Health, on page 2](#)
- [Checking Cluster Storage Capacity, on page 3](#)
- [Verifying If DRS Is Enabled, on page 3](#)
- [Verifying and Configuring the **Net.TeamPolicyUpDelay** Default Value, on page 3](#)
- [Viewing ESX Agent Manager, on page 4](#)
- [Verify the Health of a HyperFlex Cluster In Cisco UCS Manager, on page 4](#)
- [Verify UCS Server Firmware \(C-Bundle\) Version, on page 5](#)
- [Configuring vMotion Interfaces, on page 5](#)
- [Configure Lenient Mode, on page 6](#)

HyperFlex Node Upgrade Validations

Perform the following validations on each HyperFlex node before moving on to upgrade the next node in the cluster.

- Verify that the HyperFlex cluster is healthy and online. Verify all HyperFlex cluster nodes are connected to the vCenter and are online.
- Verify that no major **Alarms** are reported for the HyperFlex cluster in HyperFlex Connect.
- Verify that DRS is enabled and set to fully automated.
- Verify that vSphere services are running and ESXi Agent Manager (EAM) health is normal.
- Verify the health of the cluster in Cisco UCS Manager.
- SSH into each controller VM in the HX cluster and run the command `df -h` to ensure that there is at least 50% free space in `/var/stv`.

Output example

```
root@SpringpathControllerG9ES5WGWDG:~# df -h
Filesystem      Size  Used Avail Use% Mounted on
udev            24G   4.0K   24G   1% /dev
tmpfs           4.8G   7.0M   4.8G   1% /run
/dev/sdal       2.4G   1.5G   779M  67% /
none            5.0M    0    5.0M   0% /run/lock
none           24G    0    24G   0% /run/shm
none           100M    0    100M   0% /run/user
```

/dev/sdb1	158G	2.7G	147G	2%	/var/stv
/dev/sdb2	32G	50M	30G	1%	/var/zookeeper
none	4.0K	0	4.0K	0%	/sys/fs/cgroup



Important If there is not enough space (usage is over 50%), contact Cisco TAC for assistance.

Viewing HyperFlex Cluster Health

Using GUI

- From HyperFlex Connect, select **System Information** > **Nodes** page. Verify if the HyperFlex cluster is healthy and online.
- From the vSphere Web Client Navigator, select **vCenter Global Inventory Lists** > **Cisco HyperFlex Systems** > **Cisco HX Data Platform** > *cluster* > **Summary**. View the *cluster widget* to verify if the HyperFlex cluster is healthy and online.
- From the vSphere Web Client Navigator, select **vCenter Global Inventory Lists** > **Clusters** > *cluster* > **Summary**. Verify if all HX Cluster nodes are connected to the vCenter and they are online.

Using CLI

Log in to any controller VM in the storage cluster. Run the command `stcli cluster storage-summary --detail`.

```
address: 192.168.100.82
name: HX-Cluster01
state: online
uptime: 0 days 12 hours 16 minutes 44 seconds
activeNodes: 5 of 5
compressionSavings: 78.1228617455
deduplicationSavings: 0.0
freeCapacity: 38.1T
healingInfo:
  inProgress: False
resiliencyDetails:
  current ensemble size:5
  # of ssd failures before cluster shuts down:3
  minimum cache copies remaining:3
  minimum data copies available for some user data:3
  minimum metadata copies available for cluster metadata:3
  # of unavailable nodes:0
  # of nodes failure tolerable for cluster to be available:2
health state reason:storage cluster is healthy.
  # of node failures before cluster shuts down:3
  # of node failures before cluster goes into readonly:3
  # of hdd failures tolerable for cluster to be available:2
  # of node failures before cluster goes to enospace warn trying to move the existing
data:na
  # of hdd failures before cluster shuts down:3
  # of hdd failures before cluster goes into readonly:3
  # of ssd failures before cluster goes into readonly:na
  # of ssd failures tolerable for cluster to be available:2
```

```
resiliencyInfo:
  messages:
    Storage cluster is healthy.
  state: healthy
  hddFailuresTolerable: 2
  nodeFailuresTolerable: 1
  ssdFailuresTolerable: 2
spaceStatus: normal
totalCapacity: 38.5T
totalSavings: 78.1228617455
usedCapacity: 373.3G
clusterAccessPolicy: lenient
dataReplicationCompliance: compliant
dataReplicationFactor: 3
```

Sample response that indicates the HyperFlex storage cluster is online and healthy.

Checking Cluster Storage Capacity

We recommend that you check the cluster storage capacity before starting the upgrade of an existing installation of Cisco HX Data Platform. If the storage cluster capacity is above 70%, it is highly recommended to either reduce the amount of storage capacity used or increase the storage capacity by adding new nodes or disks. This confirmation of cluster storage capacity is important because if a node goes down in such a situation, the cluster will not be able to rebalance and will stay unhealthy (online).

Refer to the *HX Storage Cluster Overview* chapter in the [Cisco HyperFlex Data Platform Administration Guide](#) for background details about checking cluster storage capacity.

Verifying If DRS Is Enabled

Procedure

- Step 1** From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Clusters > cluster > Configure** tab.
- Verify that DRS is **Enabled**.
- Step 2** Click the **vSphere DRS** tab.
- Check if **Migration Automation Level** is set to **Fully Automated**.
-

Verifying and Configuring the Net.TeamPolicyUpDelay Default Value

To avoid loss of storage access during Fabric Interconnect reboot for firmware update, perform this check prior to UCSM Infrastructure upgrade.

Upgrades to 3.0(1), 3.5(1), and 3.5(2) require that the default value of the ESXi host `Net.TeamPolicyUpDelay` be set to 30000. Complete the following steps to verify, and if needed, modify the default value of the ESXi host `Net.TeamPolicyUpDelay` to 30000.



Note This change does not require a reboot of the ESXi host, and is persistent across a reboot.

Procedure

- Step 1** From the vSphere Web Client Navigator, click on each ESXi Host > **Configure** > **System** > **Advanced System Settings**.
 - Step 2** In **Advanced System** Settings, scroll down to **Net.TeamPolicyUpDelay**.
 - Step 3** If needed, change the value to *30000*. The default value is 100.
-

Viewing ESX Agent Manager

Procedure

From the vSphere Web Client Navigator, select **Administration** > **vCenter Server Extensions** > **vSphere ESX Agent Manager** > **Summary**.

Verify that ESX Agent Manager (EAM) health is normal.

Verify the Health of a HyperFlex Cluster In Cisco UCS Manager

Procedure

- Step 1** Verify if the high availability status of the fabric interconnects shows that both the fabric interconnects are up and running. See the [Cisco UCS Manager System Monitoring Guide](#) for more information.
 - Step 2** Verify that the data path is up and running. See the [Cisco UCS Manager Firmware Management Guide](#) for more information.
 - Step 3** Verify that the HyperFlex servers have no faults.
 - Step 4** Verify that vNIC faults are cleared to ensure VMware ESXi vSwitch uplinks are up and operational.
 - Step 5** Verify if all servers have been discovered.
-

Verify UCS Server Firmware (C-Bundle) Version

Using UCS Manager

1. Log in to UCS Manager.
2. Select the **Server** tab.
3. Select the Host Firmware Package policy by navigating to, **Policies > Root > Sub-Organizations > <hx-cluster> > Host Firmware Packages > HyperFlex.**



Note Ensure that you select the desired cluster under the sub-org list.

4. Under properties, note the current Rack Package version. It is listed as X.Y(Z)C. For example, 3.1(2g)C.

Using HX Connect

1. Log in to HX Connect.
2. In the Navigation pane, select **Upgrade**.
3. Select the **UCS Firmware** check box and click **Discover**.
4. Note the current C-bundle version displayed.

Configuring vMotion Interfaces

Complete the following steps to add the VMkernel interface necessary for vMotion to work:

Before you begin

Only default TCP/IP stack is supported for vMotion vmkernel adapters.

You must pre-define vMotion networking by creating a vSwitch and defining the vNICs and VLANs in UCS Manager.

Procedure

- Step 1** In the vSphere Web Client Navigator, click on **Host > Inventory > Manage > Networking > VMkernel adapters**.
- Step 2** Click **Add Host Networking**.
- Step 3** Select **VMkernel Network Adapter**.
- Step 4** Select the existing **vmotion vSwitch** by selecting browse.
- Step 5** Provide a name, and refer to table below to enter the appropriate **VLAN ID**.

Cluster Installation Version	VLAN ID
1.7.x	0 (default)
1.8.x and later	same as vMotion network

Step 6 Provide a **Static IP Address** and complete the wizard.

Step 7 (Optional) To use jumbo frames, edit the **vmk2** and set the **MTU** to **9000**. Your upstream switch must be configured to pass jumbo frames on the **vMotion VLAN**.

Step 8 Repeat steps 1 to 6 for all hosts in the cluster.

Configure Lenient Mode

Cluster access policy is set by default to `lenient mode`. To manually set the cluster access policy to `lenient`, use the following procedure.

Procedure

	Command or Action	Purpose
Step 1	SSH to any one of the controller VMs and login as root.	
Step 2	Check if lenient mode is already configured.	<code>#stcli cluster get-cluster-access-policy</code>
Step 3	If set to strict, change to lenient. If already set to lenient, no further action is required.	<code>~/#stcli cluster set-cluster-access-policy --name lenient</code>
Step 4	Confirm the change.	<code>stcli cluster info grep -i policy</code>

Example

```
~/#stcli cluster get-cluster-access-policy strict
~/#stcli cluster set-cluster-access-policy --name lenient
stcli cluster info | grep -i policy
```