



## Appendix

---

- [Configuring HX Server QoS System Classes, on page 1](#)
- [Creating VLANs for HX Servers, on page 2](#)
- [Creating MAC Address Pools, on page 3](#)
- [Creating a QoS Policy, on page 5](#)
- [Creating vNIC Templates for HX Servers, on page 6](#)
- [Configure the vSwitches, on page 10](#)
- [Migrating vMotion Networks to Virtual Distributed Switches \(VDS\) or Cisco Nexus 1000v \(N1Kv\), on page 11](#)
- [Reset Stats Daemon, on page 13](#)
- [HyperFlex HX-Series Server Disk Details, on page 13](#)
- [Validation Checks for Installing HyperFlex, on page 14](#)
- [Non-Default BIOS Settings, on page 17](#)

## Configuring HX Server QoS System Classes

Cisco UCS uses Data Center Ethernet (DCE) to handle all traffic inside a Cisco UCS domain . This industry standard enhancement to Ethernet divides the bandwidth of the Ethernet pipe into eight virtual lanes. Two virtual lanes are reserved for internal system and management traffic. You can configure quality of service (QoS) for the other six virtual lanes. System classes determine how the DCE bandwidth in these six virtual lanes is allocated across the entire Cisco UCS domain.

Each system class reserves a specific segment of the bandwidth for a specific type of traffic, which provides a level of traffic management, even in an oversubscribed system. For example, you can configure the Fibre Channel Priority system class to determine the percentage of DCE bandwidth allocated to FCoE traffic.

To configure QoS System Classes in UCS Manager, follow the steps below:

### Procedure

---

- Step 1** Open a web browser and enter the IP address for Cisco UCS Manager . Enter the login credentials.
- Step 2** In Cisco UCS Manager, navigate to the **LAN** tab.
- Step 3** In the **LAN** tab, expand **LAN > LAN Cloud**.
- Step 4** Select the **QoS System Class** node.
- Step 5** In the **Work** pane, click the **General** tab

**Step 6** Update the properties for the system class that you want to configure to meet the traffic management needs of the system.

Priority	Enabled	QoS	Packet Drop	Weight	MTU	Multicast Optimized
<b>Platinum</b>	<i>Check the box for the system that you want to enable</i>	5	No	4	9216	no
<b>Gold</b>		4	Yes	4	normal	no
<b>Silver</b>		2	Yes	Best Effort	normal	yes
<b>Bronze</b>		1	Yes	Best Effort	9216	no
<b>Best Effort</b>		Any	Yes	Best Effort	normal	no
<b>fibre Channel</b>		3	No	5	fc	N/A

**Note** When you change the MTU values of the platinum and bronze system classes in a brand new Cisco UCS domain setup, you might receive the following warnings:

- Validating MTU change for QoS system class: 'platinum'. Failed: Cannot change MTU from '1500' to '9216' since it is currently in use by '0' service profile(s).
- Validating MTU change for QoS system class: 'bronze'. Failed: Cannot change MTU from '1500' to '9216' since it is currently in use by '0' service profile(s).

**Step 7** Click **Save Changes**.

---

## Creating VLANs for HX Servers

### Procedure

---

- Step 1** Open a web browser and enter the IP address for Cisco UCS Manager . Enter the login credentials.
- Step 2** Navigate to **LAN tab > LAN > LAN Cloud > VLANS**.
- Step 3** Right-click and select Create VLANs as shown in the table below:

VLAN Name	Description	Multicast Policy Name	VLAN ID (by default)
<b>hx-inband-mgmt</b>	Used for: <ul style="list-style-type: none"> <li>• ESX management</li> <li>• SSH to storage controller VM</li> <li>• HX Cluster management IP - using multicast traffic.</li> <li>• vCenter connectivity to the HyperFlex VM for the HX Data Platform plug-in</li> </ul>	HyperFlex	3091
<b>hx-storage-data</b>	Used for: <ul style="list-style-type: none"> <li>• ESX NFS client(IOvisor)</li> <li>• HyperFlex replication/cluster</li> <li>• Cluster data VIP</li> </ul>	HyperFlex	3092
<b>hx-vmotion</b>	Used for: <ul style="list-style-type: none"> <li>• VM and storage vmotion, FT, iSCSI</li> </ul>	HyperFlex	3093
<b>insert existing vlan name</b>	Used for: <ul style="list-style-type: none"> <li>• VM data traffic</li> </ul>	HyperFlex	Any*

Note:

- Configuration option is Common/Global. It applies to both fabrics and uses the same configuration parameters in both cases.
- \*There is no specific recommendation for VM data VLANs. You can create your own VLANs for the VM data traffic. By default, the HXDP installer will not create VLANs for the VM data traffic.
- Installer sets the VLANs as non-native by default. Ensure to configure the upstream switches to accommodate the non-native VLANs.

## Creating MAC Address Pools

You can change the default MAC address blocks to avoid duplicate MAC addresses that may already exist. Each block contains 100 MAC addresses by default to allow for up to 100 HX server for deployment per UCS system. We recommend that you use one MAC pool per vNIC for easier troubleshooting.



**Note** The 8th digit is set to A or B. "A" is set on vNICs pinned to Fabric Interconnect A. And "B" is set on vNICs pinned to Fabric Interconnect B.

**Procedure**

- Step 1** Open a web browser and enter the IP address for Cisco UCS Manager . Enter the login credentials.
- Step 2** In Cisco UCS Manager, navigate to **LAN tab > Pools > root > Sub-org > hx-cluster > MAC Pools**.
- Step 3** Right-click **MAC Pools** and select **Create MAC Pool**.
- Step 4** In the **Define Name and Description** page of the **Create MAC Pool** wizard, complete the required fields as shown below:

MAC Pool Name	Description	Assignment Order	MAC Address block
<b>hv-mgmt-a</b>	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:01:01-64
<b>hv-mgmt-b</b>	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:02:01-64
<b>storage-data-a</b>	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:03:01-64
<b>storage-data-b</b>	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:04:01-64
<b>vm-network-a</b>	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:05:01-64
<b>vm-network-b</b>	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:06:01-64
<b>hv-vmotion-a</b>	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:07:01-64
<b>hv-vmotion-b</b>	MAC pool for HyperFlex System	Sequential	00:25:B5:XX:08:01-64

- Step 5** Click **Next**.
- Step 6** In the **Add MAC Addresses** page of the **Create MAC Pool** wizard, click **Add**.
- Step 7** In the **Create a Block of MAC Addresses** dialog box, complete the following fields:

Name	Description
<b>First MAC Address</b> field	The first MAC address in the block.
<b>Size</b> field	The number of MAC addresses in the block.

- Step 8** Click **OK**.
- Step 9** Click **Finish**.

After the MAC address change, ESXi will be reconfigured as it was configured earlier. But if management IP was DHCP assigned then the IP would change.

**Impact of Manufacturing process on MAC address change**

- The MAC address will change between the manufacturing process and the customer site, especially if the customer orders HyperFlex servers without UCS Fabric Interconnects.
- A MAC address is configured during Service Profile association. It is un-configured during Service Profile disassociation.
- At the end of manufacturing process, the service profiles are disassociated, hence the MAC addresses are un-configured.
- When a HyperFlex server is deployed, configure the MAC address pools as described above.
- VMWare supports Consistent Device Naming, but issues have been reported since 5.5.SR has been opened.

## Creating a QoS Policy

A quality of service (QoS) policy assigns a system class to the outgoing traffic for a vNIC or vHBA. This system class determines the quality of service for that traffic.

You must include a QoS policy in a vNIC policy or vHBA policy and then include that policy in a service profile to configure the vNIC or vHBA.

The following table describes the system classes that you can configure.

**Table 1: System Classes**

System Class	Description
Platinum Gold Silver Bronze	A configurable set of system classes that you can include in the QoS policy for a service profile. Each system class manages one lane of traffic.  All properties of these system classes are available for you to assign custom settings and policies.
Best Effort	A system class that sets the quality of service for the lane reserved for basic Ethernet traffic.  Some properties of this system class are preset and cannot be modified. For example, this class has a drop policy that allows it to drop data packets if required. You cannot disable this system class.
Fibre Channel	A system class that sets the quality of service for the lane reserved for Fibre Channel over Ethernet traffic.  Some properties of this system class are preset and cannot be modified. For example, this class has a no-drop policy that ensures it never drops data packets. You cannot disable this system class.  <b>Note</b> FCoE traffic has a reserved QoS system class that should not be used by any other type of traffic. If any other type of traffic has a CoS value that is used by FCoE, the value is remarked to 0.

To create a QoS Policy in UCS Manager, follow the steps below:

### Procedure

- Step 1** Open a web browser and enter the IP address for Cisco UCS Manager . Enter the login credentials.
- Step 2** In the **Navigation** pane, click **LAN**.
- Step 3** In the **LAN** tab, expand **LAN > Policies**.
- Step 4** Expand the **root** node > **Sub-org** > **hx-cluster**
- Step 5** Right-click **QoS Policy** and select **Create QoS Policy**.
- Step 6** In the **Create QoS Policy** dialog box, complete the required fields as shown in the table below:

QoS Policy Name	QoS Class	Burst Size	Rate	Host Control
<b>Platinum</b>	Platinum	10240	Line-rate	none
<b>Gold</b>	Gold	10240	Line-rate	none
<b>Silver</b>	Silver	10240	Line-rate	none
<b>Bronze</b>	Bronze	10240	Line-rate	none
<b>Best Effort</b>	Best Effort	10240	Line-rate	none

- Step 7** Click **OK**.

### What to do next

Include the QoS policy in a vNIC or vHBA template.

## Creating vNIC Templates for HX Servers

### Before you begin

This policy requires that one or more of the following resources already exist in the system:

- Named VLAN
- MAC pool
- QoS policy
- LAN pin group
- Statistics threshold policy

In this procedure you create a total of eight vNIC templates: one each for the traffic management, storage management, network management, and VMotion for FI(A), and then the same set for FI(B).

## Procedure

- Step 1** In Cisco UCS Manager, navigate to **LAN tab > Policies > root > Sub-Organization > Hyperflex > vNIC Templates**.
- Step 2** Right-click the **vNIC Templates** node and select **Create vNIC Template**.
- Step 3** In the **Create Network Policy** dialog box, complete the required fields as follows:

vNIC Template Name	Fabric ID	VLAN	Use VLAN	MAC address Pool	MTU	QoS Policy	Network Control Policy	Description
hv-mgmt-a	A	hx-inband-mgmt	No	hv-mgmt-a	1500	Silver	Network Control Policy: <b>hyperflex-infra</b>	Used for: <ul style="list-style-type: none"> <li>• ESX management</li> <li>• SSH to storage controller VM</li> <li>• Cluster management IP</li> <li>• vCenter connectivity to the HX Controller VM for the HXDP plug-in.</li> <li>• <b>hv-mgmt-a</b> and <b>hv-mgmt-b</b> are used as uplinks for virtual switch <b>vswitch-hx-inband-mgmt</b> in vCenter.</li> </ul>
hv-mgmt-b	B	hx-inband-mgmt	No	hv-mgmt-b				
storage-data-a	A	storage-data	No	storage-data-a	9000	Platinum	Network Control Policy: <b>hyperflex-infra</b>	Used for: <ul style="list-style-type: none"> <li>• ESXi NFS client (IOSvisor)</li> <li>• HXDP replication/cluster</li> <li>• Cluster data VIP</li> <li>• <b>storage-data-a</b> and <b>storage-data-b</b> are used as uplinks for virtual switch <b>vswitch-hx-storage-data</b> in vCenter.</li> <li>• NFS traffic should be on a dedicated vNIC and VLAN due to security and QoS considerations.</li> </ul>
storage-data-b	B	storage-data	No	storage-data-b				

vNIC Template Name	Fabric ID	VLAN	Name VLAN	MAC address Pool	MTU	QoS Policy	Network Control Policy	Description
<b>vm-network-a</b>	A	(customer vlan name)	No	<b>vm-network-a</b>	1500	Gold	Network Control Policy: <b>hyperflex-vm</b>	Used for: <ul style="list-style-type: none"> <li>• VM data traffic (VDI, database...)</li> <li>• <b>vm-network-a</b> and <b>vm-network-b</b> are used as uplinks for virtual switch <b>vswitch-hx-vm-network</b> in vCenter.</li> </ul>
<b>vm-network-b</b>	B	(customer VLAN name)	No	<b>vm-network-b</b>				
<b>hv-vmotion-a</b>	A	hv-motion-a	No	<b>hv-vmotion-a</b>	9000	Bronze	Network Control Policy: <b>hyperflex-infra</b>	Used for: <ul style="list-style-type: none"> <li>• VM and storage vMotion, FT</li> <li>• <b>hv-vmotion-a</b> and <b>hv-vmotion-b</b> are used as uplinks for virtual switch <b>vmotion</b> in vCenter.</li> </ul>
<b>hv-vmotion-b</b>	B	hv-motion-b	No	<b>hv-vmotion-a</b>				

In the **General** area, set all the properties according to the following reference table across all the eight vNICs:

<b>Failover</b>	Disabled
<b>Target</b>	Adapter
<b>Template Type</b>	Updating
<b>Pin Group</b>	not set
<b>Stats Threshold Policy</b>	default
<b>Dynamic vNIC Connection Policy</b>	not set
<b>VLANs</b>	Configured as shown in the following table below for each of the vNIC templates



Table 2: Configured VLANs on the vNIC templates

vNIC Name	VLANs	Comments
hv-mgmt-a	hx-inband-mgmt	<p>The HXDP Installer configures a single VLAN on the UCSM LCP vNIC as follows:</p> <ul style="list-style-type: none"> <li>Set the VLAN name to <b>hx-inband-mgmt</b></li> <li>Set the VLAN ID to 3091 by default.</li> </ul> <p><b>Note</b> You can change the VLAN ID in the HXDP Installer.</p> <ul style="list-style-type: none"> <li>Post HXDP install, you can open UCSM and create more VLANs to add to the <b>hv-mgmt-a</b> and <b>hv-mgmt-b</b> vNIC templates.</li> </ul> <p><b>Note</b> You can use these additional VLANs to access external systems such as NetApp NFS/ISCSI filer.</p> <ul style="list-style-type: none"> <li>Port Group name is <b>Storage Controller Management network</b> backed by VLAN <b>hx-inband-mgmt</b></li> </ul>
hv-mgmt-b		
storage-data-a	hx-storage-data	<p>The HXDP Installer configures a single VLAN as follows:</p> <ul style="list-style-type: none"> <li>Set the VLAN name to <b>hx-storage-data</b></li> <li>Set as the native VLAN</li> <li>Set the VLAN ID to 3092 by default</li> </ul> <p><b>Note</b> You can change the VLAN id in the HXDP Installer, but it cannot be same as <b>hx-inband-mgmt</b>, or ESXi routing will get confused.</p> <ul style="list-style-type: none"> <li>Port Group names are: <ul style="list-style-type: none"> <li><b>Storage Controller Data Network</b> backed by VLAN hx-storage-data</li> <li><b>VMK Storage Hypervisor Data Network</b> backed by VLAN hx-storage-data</li> </ul> </li> <li>Subnet 10</li> </ul>
storage-data-b		
vm-network-a	user created VLANs	<ul style="list-style-type: none"> <li>Manually create one or more VLANs in UCSM</li> <li>Manually create port groups backed by user-created VLANs</li> <li>You can create additional VLANs in UCSM and assign them to the <b>vm-network-a</b> and <b>vm-network-b</b> vNIC templates for VM traffic</li> </ul> <p><b>Note</b> The HXDP Installer does not configure any VLAN or Port group.</p>
vm-network-b		

vNIC Name	VLANs	Comments
hv-vmotion-a	hx-vmotion	The HXDP Installer configures a single VLAN as follows: <ul style="list-style-type: none"> <li>• VMotion: VLAN <b>hx-vmotion</b></li> <li>• Set the VLAN ID</li> <li>• Sets as the native VLAN</li> <li>• VLAN ID is 3093 by default</li> <li>• Subnet 10</li> </ul>
hv-vmotion-b		

**Step 4** Click **OK** when finished.

## Configure the vSwitches

In both VMware ESX and ESXi host, you can configure vSwitches from either the GUI or the command line. The CLI configurations are very helpful when you are installing multiple ESX servers and planning to script the vSwitch configuration.

After the ESX installation, configure your vSwitches on the ESX host with the following steps:

### Procedure

**Step 1** Log in to the command line of each ESX Server.

**Step 2** Create three vSwitches on each ESX server using the listed names.

- **vswitch-hx-storage-data**  
Set the MTU to 9000 on this switch.
- **vmotion**  
Set the MTU to 9000 on this switch.
- **vswitch-hx-vm-network**

**Step 3** Use the following CLI commands to create the three new vSwitches:

```
# esxcli network vswitch standard add -v vswitch-hx-storage-data
# esxcli network vswitch standard set -v vswitch-hx-storage-data -mtu= 9000
# esxcli network vswitch standard add -v vswitch-vmotion
# esxcli network vswitch standard set -v vswitch-vmotion -mtu=9000
# esxcli network vswitch standard add -v vswitch-hx-vm-network
```

**Step 4** The default vSwitch **vSwitch0** created during installation of ESXi needs to be renamed to **vswitch-hx-inband-mgmt** for the Hx Data Platform node set up scripts to work properly. Use the following

command to rename the switch and then reboot the host so that the vmkernel re-reads its configuration file to use the new name.

```
# sed -i 's/vSwitch0/vswitch-hx-inband-mgmt/g' /etc/vmware/esx.conf
# reboot
```

**Step 5** You can verify the creation and renaming of the vSwitches after a host reboot with the following command:

```
# esxcli network vswitch standard list
```

Confirm that you see the four previously listed vSwitches in the command output. Only the switch-hx-inband-mgmt vSwitch will have Uplinks and Port groups listed. The HX Data Platform installer scripts perform the rest of the network configuration.

---

## Migrating vMotion Networks to Virtual Distributed Switches (VDS) or Cisco Nexus 1000v (N1Kv)



### Note

- The HX Data Platform can be configured with VMware DVS or Cisco Nexus 1000v for specific non-HX dependent networks:
  - vMotion networks
  - and virtual machine networks
- For further details, see [Cisco Nexus 1000v documentation](#).

---

To migrate non-HX dependent vSwitches and associated port groups to DVS or N1Kv networks, follow the listed steps:

### Procedure

---

**Step 1** From vCenter, create DVS Switch and port groups.

- a) Select **vCenter Inventory Lists > Datacenters > datacenter > Related Objects > Distributed Switches**. Click **Add Distributed Switch** icon.
- b) Complete the New Distributed Switch wizard. Create each DVS switch with two uplinks.

For example: VM network and vmotion pg

- DVSwitch-VMNetwork: DVPortGroup-VMNetwork
- DVSwitch-Vmotion: DVPortGroup-Vmotion

**Step 2** Migrate the vSwitch, VMNetwork. Perform the following steps to migrate VMNetwork from legacy vSwitch to DVS.

- a) Select **vCenter Inventory Lists > Datacenters > datacenter > Related Objects > Distributed Switches**.

- b) Select the **DVSwitch-VMNetwork** vSwitch. Click the **Add and Manage Hosts** icon. This starts the **Add and Manage Hosts** wizard.
- c) On the Select task page, select **Add Hosts**. Click **Next**.
- d) On the Select hosts page, click **Add New Hosts**. Select all hosts in the cluster. Click **Next**.
- e) On the Select network adapter tasks page, select **Manage physical adapters** and **Migrate virtual machine networking**. Click **Next**.
- f) On the Manage physical network adapters page, the physical adapters part of vswitch-hx-vm-network:VM Network are assigned to the DVSwitch-VMNetwork.
- g) Under the **On other switches/unclaimed list**, select the vmnic corresponding to the **In Use by Switch**, vswitch-hx-vm-network.
- h) Click **Assign uplink**.
- i) Select Auto-assign.
- j) Click **OK**. The page refreshes with the newly assigned vmnic listed under **On this switch**.
- k) The **Analyze impact** page shows the impact of this migration. Verify the impact is all green. Click **Next**.
- l) On the Migrate VM networking page, select the VMs to migrate to the new network, DVPortGroup-VMNetwork.

Next

Select all the VMs, except the controller VMs, stCtlVM, from all the hosts. Select the DVPortGroup-VMNetwork. Click **Next**.

**Note** The list of VMs for each host includes all the VMs, including the controller VMs. DO NOT select any controller VMs. Migrating the controller VMs will break your storage cluster.

- m) On the Ready to complete page, confirm the summary of the migration. Click Finish.

**Note** Post migration system generates several network related alarms. Verify and clear the alarms.

### Step 3

Migrate the vSwitch to vmotion pg. Perform the following steps to migrate vmotion pg from legacy vSwitch to DVS.

- a) Select **vCenter Inventory Lists > Datacenters > datacenter > Related Objects > Distributed Switches**.
- b) Select the DVSwitch-Vmotion vSwitch. Click the **Add and Manage Hosts** icon. This starts the **Add and Manage Hosts** wizard.
- c) On the Select task page, select **Add Hosts**. Click **Next**.
- d) On the Select hosts page, click **Add New Hosts**. Select all hosts in the cluster. Click **Next**.
- e) On the Select network adapter tasks page, select the tasks **Manage physical adapters** and **Manage VMkernel adapters**. Click **Next**.
- f) On the **Manage physical network adapters** page, the physical adapters part of vmotion:vmotion pg are assigned to the DVSwitch-Vmotion.

Under the **On other switches/unclaimed list**, select the vmnic corresponding to the In Use by Switch, vmotion. Click **Assign uplink**, select Auto-assign, and click OK. The page refreshes with the newly assigned vmnic listed under **On this switch**. Click **Next**.

- g) On the **Manage VMkernel network adapters** page, migrate the VMkernel adapter to the port group, DVPortGroup-Vmotion.

For each host, under the **On other switches** list, select the VMKernel adapter corresponding to the **In Use by Switch**, vmotion. Click **Assign port group**. Select the destination port group, DVPortGroup-Vmotion. Click **OK**. The page refreshes with the Reassigned VMkernel network adapters, listing the Source Port Group and Destination Port Group.

- h) Select the hosts to migrate to the new network, DVPortGroup-Vmotion. Click **Next**.
- i) On the Ready to complete page, confirm the summary of the migration, click **Finish**.

**Step 4** Post migration step. Verify there is no impact on the VMs with respect to IO, Network connectivity and VM Migration.

---

## Reset Stats Daemon

### Description

A network daemon listens for statistics, like counters and timers, sent over UDP or TCP and sends aggregates to one or more pluggable backend services.

After manually re-installing ESX on your HX Data Platform servers, reset the stats daemon to ensure performance statistics display correctly.

**Action: restart stats daemon**

### Procedure

---

**Step 1** Login to the command line of the controller VM of the ESX host.

**Step 2** Run the restart command.

```
# /etc/init.d/statsd restart
```

**Step 3** Repeat Step 1 and Step 2 on the controller VM of every ESX host in the storage cluster.

---

## HyperFlex HX-Series Server Disk Details



---

**Note** Only the disks (and servers) that are available in the order tool with specific PIDs, are supported.

---

All the HX-series servers have to be the same type in a cluster and must have the same number of disks.

All the disks in the storage cluster must have the same amount of storage capacity.

All SSDs must support TRIM and have TRIM enabled.

All HDDs can be either SATA or SAS type. All SAS disks in the storage cluster must be in a pass-through mode.

HX220c	HX240c	Hybrid
<b>Cisco 220c M4 with:</b> 2 x 2.5GHz E5-2680 v3 24 x 16GB DDR4 RAM 1 x Cisco 12GBps SAS RAID Controller 2 x 64 GB SD FlexFlash cards 1 x 120 GB SATA SSD 1 x 480 GB SATA SSD 6 x 1.2 TB, 12 GBps SAS 10k RPM HDDs VIC 1227 MLOM (2 x 10Gb ports)	<b>Cisco 240c M4 with:</b> 2 x 2.5 GHz E5-2680 v3 24 x 16GB DDR4 RAM 1 x 120GB back SSD 1 x Cisco 12GBps SAS RAID Controller 1 x 120 GB SATA SSD 1 x 1.6TB front SATA SSD Up to 23 x 1.2 TB SAS 10k RPM HDDs (for data) VIC 1227 MLOM	<b>Cisco B200 M4 (+ HX240c ):</b> 2 x 2.5GHz E5-2680 v3 24 x 16GB DDR4 RAMs 2 x 64GB SD FlexFlash Cards VIC 1340

## Validation Checks for Installing HyperFlex

### New Hardware Features in Release 1.7.1-14835

The HyperFlex installer will perform the following validation checks before configuring the HX system.

1. Firmware
2. Hardware
3. QoS
4. VLAN

Validation results in one of the following severity levels:

- OK – Check succeeded
- Warning – Check failed. Ideally, a fix should be made, but installation can proceed.
- Error – Check failed. Installation will not proceed until the error is fixed.

### 1. Firmware Validations

Configuration file: firmware\_support.json

- **Version Compatibility:** The HyperFlex version should be compatible with the UCS version. Please refer to the Cisco HyperFlex HCL documentation to obtain up-to-date information about version compatibility.

Severity: Error

Config: Advanced -> firmware\_support -> ucs\_version

- **Disk Space:** There should be enough disk space on the installer VM to download the UCS A, B, C bundles and ESXi images. Currently the requirement is around 3 GB of free space.

Severity: Error

## 2. Hardware Validations

These validations, check the necessary hardware configuration for a HX system. The server specification sheet can provide the detail hardware configuration supported by a particular model of server.

Configuration file: hardware\_support.json

- **FI Connection:** The HX server should have connection to both Fabric Interconnects.

Severity: Error

- **FI Connection Symmetric:** The HX server should be connected to same server ports on both Fabric Interconnects.

Severity: Warning

- **FlexFlash Controller:** The HX server should have SD Card FlexFlash controller.

Severity: Error

- **FlexFlash Controller Error:** The SD Card FlexFlash controller should have no errors.

Severity: Error

- **SD Card:** The HX server should have 2 SD Cards.

Severity: Error

- **SAS Controller:** The HX server should have supported SAS controller.

Severity: Error

Config: Advanced -> hardware\_support-> <ServerModel> -> sas\_ctrl\_supp\_pid

- **Disk:** The HX server should have supported disks.

Severity: Error

Since the disks are used for 3 different functionalities, there are 3 configs for each:

- SDS Logs SSD

Config: Advanced -> hardware\_support-> <ServerModel> -> ssd\_sds\_logs\_supp\_pid

- Caching SSD

Config: Advanced -> hardware\_support-> <ServerModel> -> ssd\_caching\_supp\_pid

- Data HDD

Config: Advanced -> hardware\_support-> <ServerModel> -> hdd\_supp\_pid

- **SDS Logs SSD:**

The HX server should have supported SAS controller.

The HX240C\_M4SX server should have the SDS Logs SSD attached to PCH controller (back SSD).

Severity: Error

- **Caching SSD:** The HX server should have the Caching SSD attached to SAS controller.

Severity: Error

- **Same data disks:** The HX server should have all the data disks of same PID.

Severity: Error

- **Data disk sector size:** The HX server should have all the data disks of sector size: 512 bytes.

Severity: Error

Config: Advanced -> hardware\_support-> <ServerModel> -> hdd\_sector\_size

- **Minimum data disks:** The HX server should have minimum 3 data disks.

Severity: Warning

Config: Advanced -> hardware\_support-> <ServerModel> -> hdd\_min\_req

### 3. QoS Validations

The QoS validations will check whether the desired QoS configuration can be applied. The following table shows the desired MTU of the system QoS classes for HyperFlex Traffic.

QoS class	HyperFlex Traffic	Desired MTU
Platinum	Storage Data	9216
Gold	VM data	1500
Silver	Management Traffic	1500
Bronze	VMotion Traffic	9216
Best Effort	None	1500

- **MTU Conflict:** The installer will try to configure the MTU of the system QoS classes as shown above. However, if any of the above QoS class is being used by any associated Service Profiles, the installer will try to set Best Effort QoS class for all HyperFlex traffic.

Severity: Warning

- **Best Effort MTU:** In the event Best Effort QoS class, needs to be used for all HyperFlex traffic, the Best Effort MTU should be normal (1500) or jumbo (9000).

Severity: Error

- **QoS Parameter Change:** If any of the following QoS parameters needs to be changed.

- admin\_state
- weight
- multicast\_optimize
- cos
- drop

Severity: Warning



#### 4. VLAN Validations

- **Storage/Mgmt Conflict:** The ‘storage-data’ and ‘hv-mgmt’ vNICs should not have same VLAN.  
Severity: Warning
- **Storage/vMotion:** The ‘storage-data’ and ‘hv-vmotion’ vNICs should not have same VLAN.  
Severity: Warning
- **Management/vMotion:** The ‘hv-mgmt’ and ‘hv-vmotion’ vNICs should not have same VLAN.  
Severity: Warning
- **Storage VLAN:** The ‘storage-data’ vNIC should not specify VLAN which is in use already.  
Severity: Warning

## Non-Default BIOS Settings

The following list are non-platform default BIOS settings that are common for both compute and converged nodes:



**Note** Specific settings are needed to obtain maximum performance, availability, and failover characteristics. These BIOS settings should never be manually changed and should be left as configured by the HyperFlex installer.

```
cdn_control = "enabled" for HyperV deployments
cdn_control = "platformdefault" for Esxi deployments
vp_console_redirection = "serialporta"
vp_serial_port_a_enable = "enabled"
vp_processor_c_state = "disabled"
vp_processor_c1_e = "disabled"
vp_processor_c3_report = "disabled"
vp_processor_c6_report = "disabled"
vp_processor_c7_report = "disabled"
vp_power_technology = "performance"
vp_energy_performance = "performance"
vp_altitude = "auto"
vp_cpu_performance = "hpc"
vp_cdn_control = cdn_control
vp_com_spcr_enable = "enabled"
vp_direct_cache_access = "enabled"
vp_intel_vtd_coherency_support = "disabled"
vp_intel_vtd_interrupt_remapping = "enabled"
vp_intel_vt_for_directed_io = "enabled"
vp_intel_virtualization_technology = "enabled"
vp_memory_mapped_io_above4_gb = "enabled"
```

