Cisco HyperFlex Systems Network and External Storage Management Guide

First Published: 2018-04-13

Last Modified: 2023-10-13

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

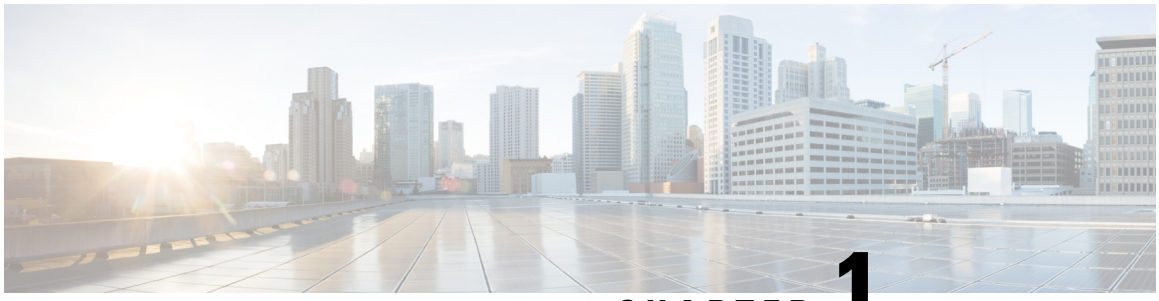

CHAPTER 1

New and Changed

New and Changed Information

The following table provides an overview of the significant changes to this guide. The table does not provide an exhaustive list of all changes made to this guide.

Table 1: New Features and Changed Behavior in Cisco HX Data Platform

Description	Date Added	Where Documented
HX Release 5.0(1a)	November 4, 2021	This guide.
Added note describing VSAN creation in the SAN Cloud tab.	July 12, 2019	Direct Attached Storage, on page 45

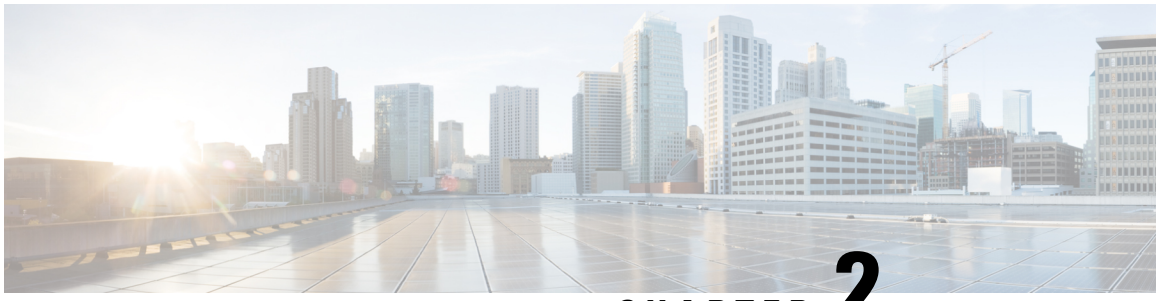

CHAPTER 2

Overview

About this Guide

This guide presents an overview of the network and external storage architecture for Cisco HyperFlex Systems. It also describes network and external management procedures that are typically performed after configuration of the initial HyperFlex cluster.

This guide is for use with all supported HX versions.

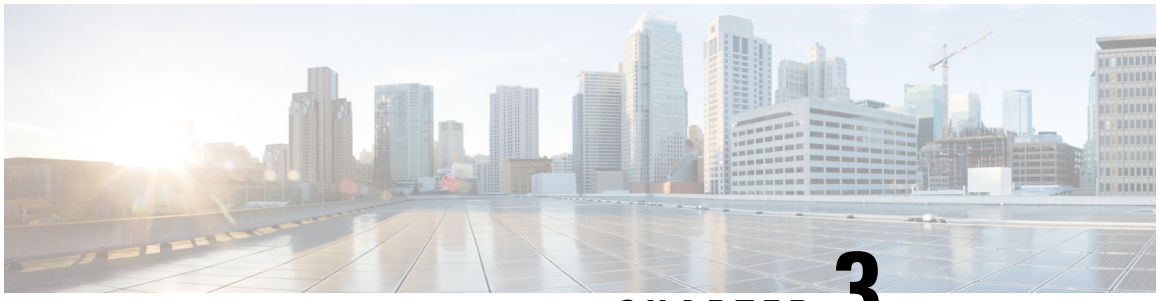

CHAPTER 3

Network Management

Physical Network

Cisco UCS Uplink Connectivity

Cisco UCS network uplinks connect *northbound* from the pair of UCS Fabric Interconnects (FIs) to the LAN in the customer datacenter. All UCS uplinks operate as trunks, carrying multiple 802.1Q VLAN IDs across the uplinks. By default, the UCS software assumes that all VLAN IDs defined in the UCS configuration are eligible to trunk across all available uplinks.

Figure 1: Logical Network Design

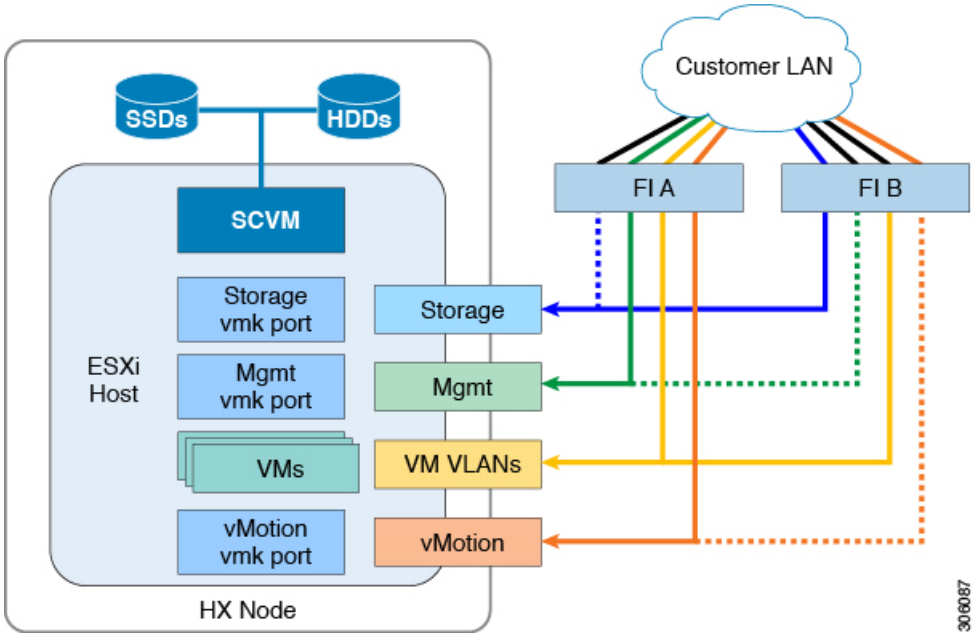

Cisco FIs appear on the network as a collection of endpoints versus another network switch. Internally, the FIs do not participate in spanning-tree protocol (STP) domains, and the FIs cannot form a network loop, as they are not connected to each other with a layer 2 Ethernet link. The upstream root bridges make all link up/down decisions through STP.

Uplinks need to be connected and active from both FIs. For redundancy, you can use multiple uplinks on each FI, either as 802.3ad Link Aggregation Control Protocol (LACP) port-channels or using individual links. For the best level of performance and redundancy, make uplinks as LACP port-channels to multiple upstream Cisco switches using the virtual port channel (vPC) feature. Using vPC uplinks allows all uplinks to be active passing data, plus protects against any individual link failure and the failure of an upstream switch. Other uplink configurations can be redundant, but spanning-tree protocol loop avoidance may disable links if vPC is unavailable.

All uplink connectivity methods must allow for traffic to pass from one FI to the other, or from fabric A to fabric B. Scenarios can occur where cable, port, or link failures require traffic that normally does not leave the UCS domain to now be forced over the UCS uplinks. In addition, you can briefly see this traffic flow pattern maintenance procedures, such as during firmware updates on the FI, which requires them to be rebooted.

VLANs and Subnets

For a Cisco HyperFlex system configuration, you must carry multiple VLANs to the UCS domain from the upstream LAN. You define these VLANs in the UCS configuration.

Table 2: HyperFlex Installer-Created VLANs

VLAN Name	VLAN ID	Purpose
hx-inband-mgmt	Customer supplied	ESXi host management interfaces HX Storage Controller VM management interfaces HX Storage Cluster roaming management interface
hx-storage-data	Customer supplied	ESXi host storage vmkernel interfaces HX Storage Controller storage network interfaces HX Storage Cluster roaming storage interface
hx-vm-data	Customer supplied	Guest VM network interfaces
hx-vmotion	Customer supplied	VMWare ESXi host vMotion vmkernel interfaces

Note Data centers often use a dedicated network or subnet for physical device management. In this scenario, the mgmt0 interfaces of the two FIs must connect to that dedicated network or subnet. HyperFlex installations consider this a valid configuration with the following caveat: you must deploy the HyperFlex installer in a location where it has IP connectivity to the following subnets:

- Subnet of the mgmt0 interfaces of the FIs
- Subnets used by the hx-inband-mgmt VLANs previously listed

Jumbo Frames

Configure all Cisco HyperFlex storage traffic that traverses the hx-storage-data VLAN and subnet to use jumbo frames; that means you configure all communication to send IP packets with a Maximum Transmission Unit (MTU) size of 9000 bytes. When you use a larger MTU value, each IP packet sent carries a larger payload, so it transmits more data per packet, and consequently sends and receives data faster. This requirement also means that you must configure the Cisco UCS uplinks to pass jumbo frames. Failure to configure the Cisco UCS uplink switches to allow jumbo frames can lead to service interruptions during some failure scenarios, particularly when cable or port failures cause storage traffic to traverse the northbound Cisco UCS uplink switches.

Logical Network

The Cisco HyperFlex system has communication pathways that fall into the following defined zones:

Table 3: Defined Communication Pathway Zones

Zone	Description
Management Zone	Comprises the connections needed to manage the physical hardware, the hypervisor hosts, and the storage platform controller virtual machines (SCVM). Make these interfaces and IP addresses available to all staff who will administer the HX system, throughout the LAN/WAN. This zone must provide access to Domain Name System (DNS) and Network Time Protocol (NTP) services, and allow Secure Shell (SSH) communication.
VM Zone	Comprises the connections needed to service network IO to the guest VMs that run inside the HyperFlex hyperconverged system. This zone typically contains multiple VLANs that are trunked to the Cisco UCS Fabric Interconnects (FIs) through the network uplinks and tagged with 802.1Q VLAN IDs. Make these interfaces and IP addresses available to all staff and other computer endpoints that need to communicate with the guest VMs in the HX system, throughout the LAN/WAN.
Storage Zone	Comprises the connections used by the Cisco HX Data Platform software, ESXi hosts, and the storage controller VMs to service the HX Distributed Data Filesystem. These interfaces and IP addresses must be able to communicate with each other at all times for proper operation. During normal operation, this traffic all occurs within the UCS domain; however, there are hardware failure scenarios where this traffic needs to traverse the network northbound of the Cisco UCS domain. For that reason, the VLAN used for HX storage traffic must be able to traverse the network uplinks from the UCS domain, reaching FI A from FI B, and vice-versa. This zone contains primarily jumbo frame traffic, so jumbo frames must be enabled on the UCS uplinks.
vMotion Zone	Comprises the connections used by the ESXi hosts to enable vMotion of the guest VMs from host to host. During normal operation, this traffic all occurs within the Cisco UCS domain; however, there are hardware failure scenarios where this traffic needs to traverse the network northbound of the Cisco UCS domain. For that reason, the VLAN used for HX storage traffic must be able to traverse the network uplinks from the Cisco UCS domain, reaching FI A from FI B, and vice-versa.

Virtual Network

The Cisco HyperFlex system has a pre-defined virtual network design at the hypervisor level. The HyperFlex installer creates four different virtual switches (vSwitches). Each switch uses two uplinks, which are each serviced by a vNIC defined in the Cisco UCS service profile.

Figure 2: ESXi Network Design

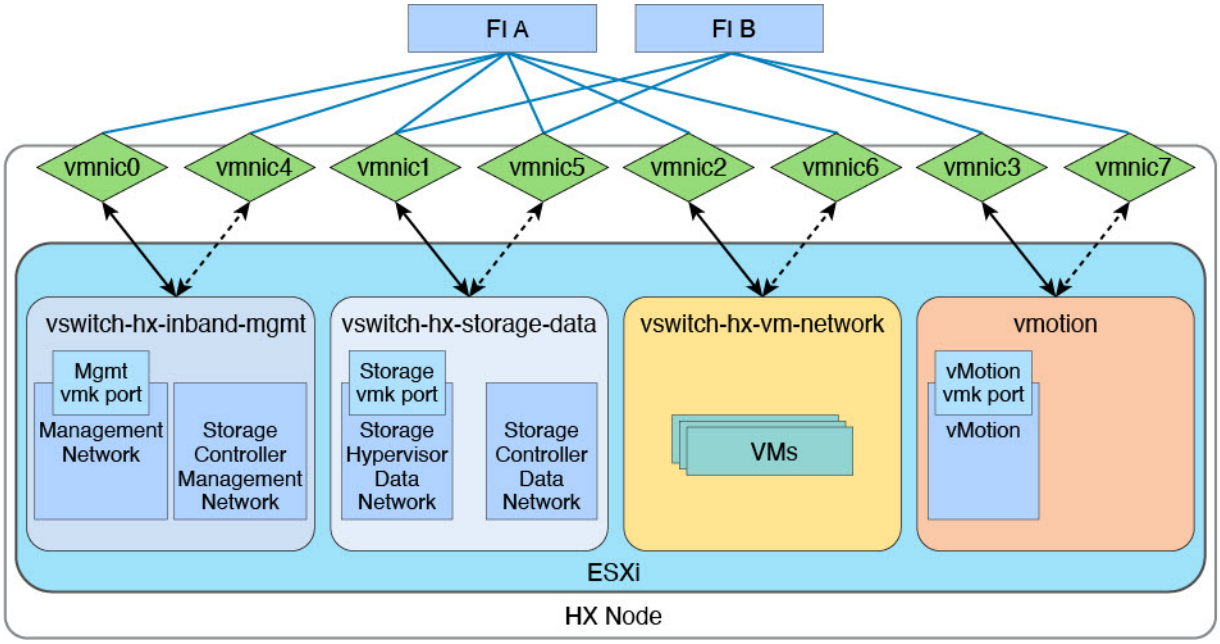

Table 4: Installer-Created vSwitches

vSwitches	Description
vswitch-hx-inband-mgmt	Default vSwitch0. Renamed by the ESXi kickstart file as part of the automated installation. The installer configures the default vmkernel port, vmk0, in the standard Management Network port group. The switch has two uplinks, active on fabric A and standby on fabric B, without jumbo frames. The installer also creates a second port group for the Storage Platform Controller VMs to connect to with their individual management interfaces. The VLAN is not a Native VLAN as assigned to the vNIC template, and therefore assigned in ESXi/vSphere.
vswitch-hx-storage-data	Created as part of the automated installation. The installer configures a vmkernel port, vmk1, in the Storage Hypervisor Data Network port group. The system uses the interface for connectivity to the HX Datastores through NFS. The switch has two uplinks, active on fabric B and standby on fabric A, with jumbo frames required. The installer also creates a second port for the Storage Platform Controller VMs to connect to with their individual storage interfaces. The VLAN is not a Native VLAN as assigned to the vNIC template, and therefore assigned in ESXi/vSphere.
vswitch-hx-vm-network	Created as part of the automated installation. The switch has two uplinks, active on both fabrics A and B, and without jumbo frames. The VLAN is not a Native VLAN as assigned to the vNIC template, and therefore assigned in ESXi/vSphere.
vMotion	Created as part of the automated installation. The switch has two uplinks, active on fabric A and standby on fabric B, with jumbo frames required. The VLAN is not a Native VLAN as assigned to the vNIC template, and therefore assigned in ESXi/vSphere.

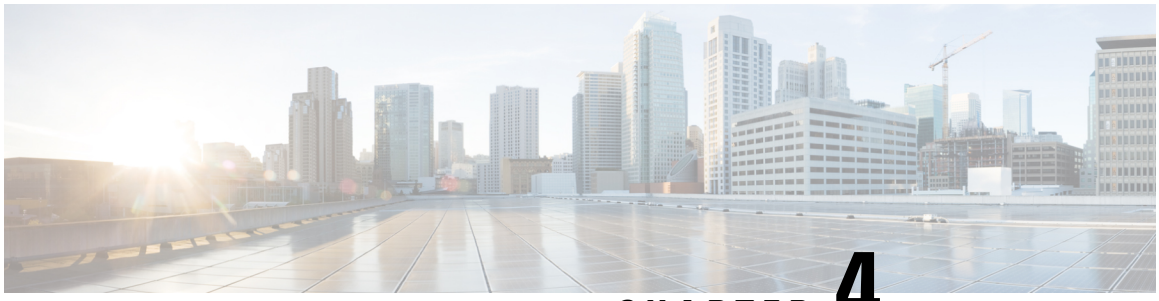

CHAPTER 4

Network Configuration after Cluster Setup

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.

You can configure the system classes shown in the following table:

Table 5: System Classes

System Class	Description
Platinum Gold Silver Bronze	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.

System Class	Description
Best Effort	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	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. **Note** 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, perform the following steps:

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, complete the fields for your systems class as shown in the following table:

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

Step 7 Click **OK**.

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 servers 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 either A or B. The *A* is set on vNICs pinned to Fabric Interconnect (FI) A. The *B* is set on vNICs pinned to FI B.

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 in the following table:

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

Entry	**Description**
First MAC Address field	The first MAC address in the block.
Size field	The number of MAC addresses in the block.

Step 8 Click **OK**.

Step 9 Click **Finish**.

After the MAC address change, the software reconfigures ESXi to how it was configured earlier. But, if management IP was DHCP assigned, then the IP changes.

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 serves 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 earlier.
- VMware supports consistent device naming, but issues have been reported since 5.5.SR has been opened.

Creating VLANs for HX Servers

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	
hx-inband-mgmt	Used for: • ESX management • SSH to storage controller VM • HX Cluster management IP - using multicast traffic. • vCenter connectivity to the HyperFlex VM for the HX Data Platform plug-in	HyperFlex	3091
hx-storage-data	Used for: • ESX NFS client (IOvisor) • HyperFlex replication/cluster • Cluster data VIP	HyperFlex	3092
hx-vmotion	Used for: • VM and storage vmotion, FT, iSCSI	HyperFlex	3093

VLAN Name	Description	Multicast Policy Name	
insert existing vlan name	Used for: • VM data traffic	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.

About Private VLANs

A private VLAN partitions the Layer 2 broadcast domain of a VLAN into subdomains, allowing you to isolate the ports on the switch from each other. A subdomain consists of a primary VLAN and one or more secondary VLANs. A private VLAN domain has only one primary VLAN. Each port in a private VLAN domain is a member of the primary VLAN, and the primary VLAN is the entire private VLAN domain.

Understanding Private VLAN Ports

Table 6: Types of Private VLAN Ports

VLAN Port	Description
Promiscuous Primary VLAN	Belongs to the primary VLAN. Can communicate with all interfaces that belong to those secondary VLANs that are associated to the promiscuous port and associated with the primary VLAN. Those interfaces include the community and isolated host ports. All packets from the secondary VLANs go through this VLAN.
Isolated Secondary VLAN	Host port that belongs to an isolated secondary VLAN. This port has complete isolation from other ports within the same private VLAN domain, except that it can communicate with associated promiscuous ports.
Community Secondary VLAN	Host port that belongs to a community secondary VLAN. Community ports communicate with other ports in the same community VLAN and with associated promiscuous ports.

Following HX deployment, a VM network uses a regular VLAN by default. To use a Private VLAN for the VM network, see the following sections:

- Configuring a Private VLAN on a VM Network without Existing VMs, on page 17.
- Configuring a Private VLAN on a VM Network with Existing VMs, on page 17.

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:

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.

Configuring a Private VLAN on a VM Network without Existing VMs

Step 1 To configure a private VLAN on Cisco UCS Manager, see the Cisco UCS Manager Network Management Guide.

Step 2 To configure a private VLAN on the upstream switch, see the Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide.

Step 3 To configure a private VLAN on ESX hosts, see Configuring Private VLAN on ESX Hosts, on page 17.

Configuring a Private VLAN on a VM Network with Existing VMs

Step 1 To configure a private VLAN on Cisco UCS Manager, see the Cisco UCS Manager Network Management Guide.

Step 2 To configure a private VLAN on the upstream switch, see the Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide.

Step 3 To configure a private VLAN on ESX hosts, see Configuring Private VLAN on ESX Hosts, on page 17

Step 4 Migrate VMs from vSphere standard switch to the newly created vSphere distributed switch.

a) Right-click the vCenter Virtual Machine and click **Migrate Virtual Machine Networking**.
b) Choose **source network** and **destination network** from the drop-down list.
c) Click **Next**.
d) Select the **Virtual Machines** that you want to migrate.
e) Click **Finish**.

Step 5 Change network connection of the network adapter on the VMs to private VLAN.

a) Right-click the vCenter Virtual Machine and click **Edit Settings**.
b) Under the **Hardware** tab, select the network adapter you want to modify.
c) Select the **Network Connection** you want to use from the **Network Label** drop-down list.
d) Click **OK**.

Configuring Private VLAN on ESX Hosts

To configure private VLANs on the ESX hosts do the following:

Step 1 Delete VMNICs on the vSphere Standard Switches from the VMware vSphere Client.

Step 2 Create new vSphere Distributed Switch with the VMNICs deleted from the previous step.

Step 3 Create promiscuous, isolated, and community VLAN.

Deleting VMNICs on the vSphere Standard Switch

Step 1 Log on to VMware vSphere Client.
Step 2 Select **Home** > **Hosts and Clusters**.
Step 3 Select the ESX host from which you want to delete the VMNIC.
Step 4 Open the **Configuration** tab.
Step 5 Click **Networking**.
Step 6 Select the **switch** you wish to remove a VMNIC from.
Step 7 Click the **Manage the physical adapters connected to the selected switch** button.
Step 8 Select the **`vminc`** you want to delete and click **Remove**.
Step 9 Confirm your selection by clicking **Yes**.
Step 10 Click **Close**.

Creating vSphere Distributed Switch

Step 1 Log on to the VMware vSphere Client.
Step 2 Select **Home** > **Networking**.
Step 3 Right click on the cluster **Distributed Switch** > **New Distributed Switch**.
Step 4 In the **Name and Location** dialog box, enter a name for the distributed switch.
Step 5 In the **Select Version** dialog box, select the distributed switch version that correlates to your version and configuration requirements.
Step 6 Click **Next**.
Step 7 In the **Edit Settings** dialog box, specify the following:

- Number of uplink ports
- **Enable** Network I/O Control.
- **Create a default port group** should be checked.
- Enter the default port group **name** in the **Port Group Name** box.

Step 8 Click **Next**.
Step 9 Review the settings in the **Ready to complete** dialog box.
Step 10 Click **Finish**.

Creating Private VLANs on vSphere Distributed Switch

Step 1 From the VMware vSphere Client, select **Inventory** > **Networking**.

Step 2 Right-click on the dvSwitch.

Step 3 Click **Edit Settings**.

Step 4 Select the **Private VLAN** tab.

Step 5 On the **Primary private VLAN ID** tab, enter a `private VLAN ID`.

Step 6 On the **Secondary private VLAN ID** tab, enter a `private VLAN ID`.

Step 7 Select the type of VLAN from the **Type** drop-down list. Valid values include:

- **Isolated**
- **Community**
- **Promiscuous** (Default)

Step 8 Click **OK**.

Set Private VLAN in Distributed Port Group

Before you begin

Create Private VLAN on the vSphere Distribute Switch.

Step 1 Right click **dvPortGroup** under **dvSwitch**, and click **Edit Settings**.

Step 2 Click **Policies** > **VLAN**.

Step 3 Select **Private VLAN**, from the **VLAN type** drop-down list.

Step 4 From the **Private VLAN Entry** drop-down list, select the type of private VLAN. It can be one of the following:

- **Isolated**
- **Community**

Note Community private VLAN is recommended.

Promiscuous ports are not supported

Step 5 Click **OK**.

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:

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

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.

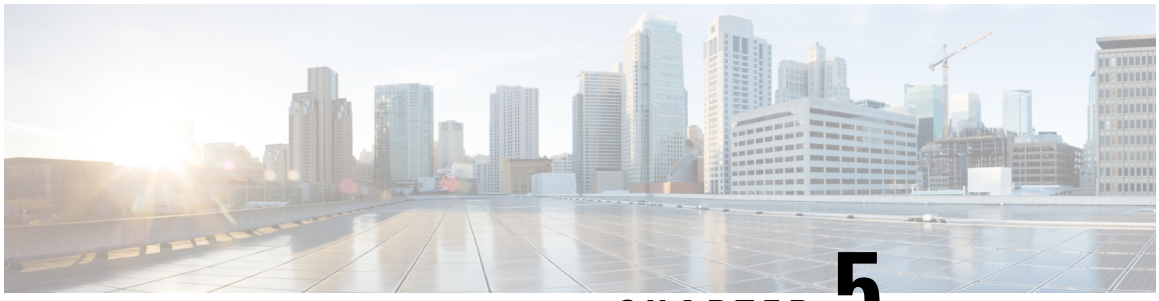

CHAPTER 5

External Storage Management

External Storage Management Overview

A Cisco HyperFlex System provides consolidated access to both SAN storage and Network Attached Storage (NAS) over the unified fabric. By unifying storage access, a Cisco Unified Computing System can access storage over Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE), and iSCSI.

The following image depicts a Cisco HyperFlex System integrated with external storage.

Figure 3: Integrating External Storage with Cisco HyperFlex Systems

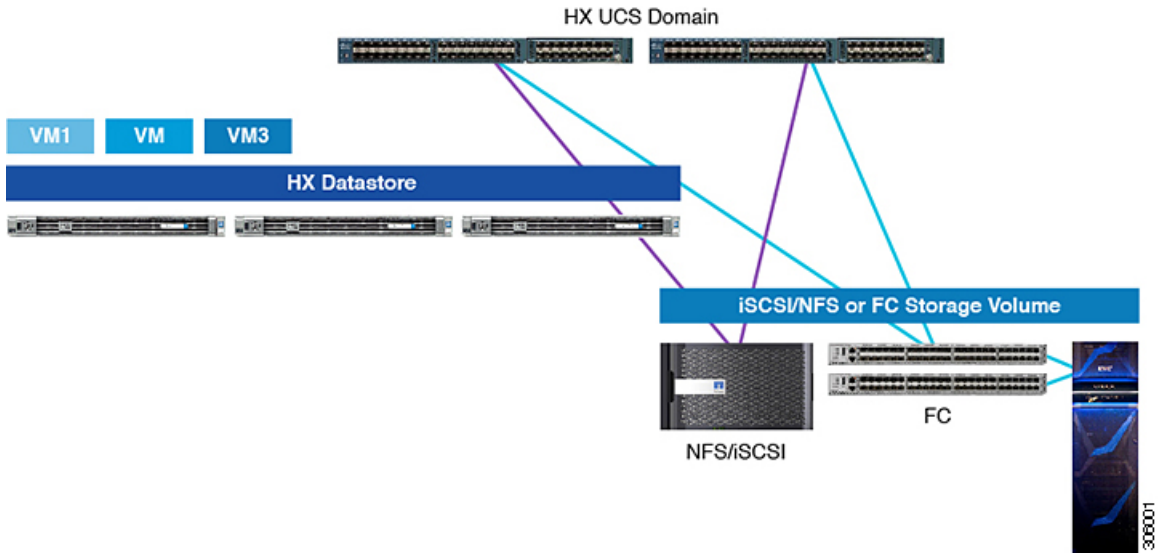

External Fiber Channel Storage

Connecting HyperFlex Nodes to External Fibre Channel Storage

This document provides detailed instructions about how to connect a Tier 1 external Fibre Channel (FC) storage array to the HyperFlex nodes. You can connect external FC storage to the HX nodes in FC end-host mode and Ethernet end-host mode as follows:

- Fabric attach FC storage
- Fabric attach FCoE storage

Storage Design Considerations

Consider the following design characteristics for HX SAN connectivity:

- Northbound storage physical connectivity does not support Virtual Port-Channels (vPCs) like LAN connectivity.
- Port channels or trunking is supported to combine multiple storage uplink ports that provide physical link redundancy.
- Storage handles the redundancy of storage resources and it varies by vendor.
- Connecting the external storage directly to the HX domain increases the fabric interconnect physical port's consumption due to additional processing.
- Software configuration including VSANs and zoning is required for providing access to storage resources.
- When utilizing external storage connectivity, it is imperative to have each cluster connecting to the storage in its own domain because the LAN connectivity policy likely differs between the two clusters.

- When integrating HyperFlex with an existing UCS domain that has NetApp IP Storage, modify the default QoS Gold class to 9216 bytes, so COS 4 allows Jumbo Frames. For more information, see NetApp KB Article number 000003500.

Storage Configuration Sequence

To connect HX servers to external storage, perform the following steps:

Step 1 Log in to the HX Data Platform Installer using admin level credentials.

Step 2 At the initial workflow selection, click **I know what I'm doing, let me customize my workflow**.

Step 3 Select only **Run UCSM Configuration** and **Run ESX Configuration**.

Step 4 Follow the wizard to complete the configuration.

The wizard creates the HX policies, Service Profile Templates, and the Service Profiles to be associated with the HX Cluster.

For detailed steps, see the Configuring HX Data Platform chapter in the HyperFlex Getting Started Guide.

Note When the cluster comes online, there will be no vHBA.

Step 5 Attach one or *both* of the following types of storage to the HX FI Domain:

- For Fibre Channel storage, see Attaching Fibre Channel Storage to HX .
- For iSCSI, see Attaching iSCSI Storage to HX.

Step 6 Re-log in to HX Installer, and then click **Start Over**.

Step 7 At the initial workflow selection, click **I know what I'm doing, let me customize my workflow**.

Step 8 Select **Deploy HX Software**, and then follow the wizard to create the HX Cluster.

Attaching Fibre Channel Storage to HX

This procedure describes the high-level steps for attaching Fibre Channel (FC) storage to the HX FI Domain:

Step 1 Log in to Cisco UCS Manager GUI.

Step 2 Configure Unified Ports as Fibre Channel. For details, see the LAN Ports and Port Channels chapter of the Cisco UCS Manager Network Management Guide.

Step 3 Create a VSAN for Fibre Channel communication. For details, see Creating VSAN for Fibre Channel, on page 26.

Step 4 Create WWNN pool and WWNN block for HyperFlex. For details, see Creating WWNN Pools.

Step 5 Create fabric-specific (hx-a and hx-b) WWPN pool and WWPN block. For details, see Creating a WWPN Pool, on page 27.

Step 6 Create a pair of vHBA templates with the previously created WWPN Pool associated with Fabric-A and Fabric-B, respectively.

Step 7 Create HyperFlex SAN connectivity Policy. For details, see Creating SAN Connectivity Policy, on page 29.

Step 8 Assign the HX SAN connectivity policy to the HX Service Profile Template (SPT) that is used for the cluster. This step triggers a pending-ack to be raised on all nodes within the cluster created from the modified SPT. Acknowledge all pending-acks for all Service Profiles in the cluster to trigger re-configuration of Service Profiles with vHBAs.

Creating VSAN for Fibre Channel

FCoE VLANs in the SAN cloud and VLANs in the LAN cloud must have different IDs. Using the same ID results in a critical fault and traffic disruption for all vNICs and uplink ports that use that FCoE VLAN. Ethernet traffic is dropped on any VLAN with an ID that overlaps with an FCoE VLAN ID.

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Click the **SAN Cloud** > **VSAN** node.

Step 3 Right-click the **VSAN** node, and select **Create Storage VSAN**.

Step 4 In the **Create VSAN** dialog box, complete the following fields:

Name	Description
Name field	The name assigned to the network. This name can be between 1 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
FC Zoning* field	Make sure the **Disable** radio button for the Fabric Interconnects (FIs) in FC end-host mode is selected. **Note** Make sure that the FI is not connected to an upstream switch.
Configuration	Select a configuration for your environment. • Click the **Common/Global** radio button so that VSAN maps to the same VSAN ID in all available fabrics. • Click **Both Fabrics Configured Differently** radio button to create two VSANS, with different IDs for Fabric A and Fabric B.
VSAN ID field	The unique identifier assigned to the network. For FC end-host mode, the range between 3840 to 4079 is also a reserved VSAN ID range.
FCoE VLAN field	The unique identifier assigned to the VLAN used for Fibre Channel connections. VLAN 4048 is user configurable. However, Cisco UCS Manager uses VLAN 4048 for the following default values. If you want to assign 4048 to a VLAN, you must reconfigure these values.

Creating WWNN Pools

A World Wide Node Name (WWNN) pool is a World Wide Name (WWN) pool that contains only World Wide Node Names. If you include a pool of WWNNs in a service profile, the software assigns the associated server a WWNN from that pool.

Important A WWN pool can include only WWNNs or WWPNs in the ranges from 20:00:00:00:00:00:00:00 to 20:FF:FF:FF:FF:FF:FF:FF or from 50:00:00:00:00:00:00:00 to 5F:FF:FF:FF:FF:FF:FF:FF. All other WWN ranges are reserved. To ensure the uniqueness of the Cisco UCS WWNNs and WWPNs in the SAN fabric, it is recommended to use the following WWN prefix for all blocks in a pool: 20:00:00:25:B5:XX:XX:XX

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Pools** > **root** > **Sub-Organizations** > **hx-cluster**.

Step 3 Expand the **hx-cluster** sub-organization to create the pool.

Step 4 Right-click WWNN Pools and select **Create WWNN Pool**.

Step 5 In the **Define Name and Description** dialog box of the **Create WWNN Pool** wizard, enter **`HyperFlex`**.

Step 6 Click **Next**.

Step 7 In the **Add WWN Blocks** page of the **Create WWNN Pool** wizard, click **Add**.

Step 8 In the **Create WWN Block** dialog box, complete the following fields:

Form field: The first WWN in the block.

Size field: The number of WWNs in the block.

For WWN pools, the pool size must be a multiple of ports-per-node + 1. For example, if there are seven ports per node, the pool size must be a multiple of eight. If there are 63 ports per node, the pool size must be a multiple of 64.

Step 9 Click **OK**.

Step 10 Click **Finish**.

What to do next

Create WWPN Pool.

Creating a WWPN Pool

To create a WWWPN Pool, perform the following steps.

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Pools** > **root** > **Sub-Organizations** > **hx-cluster**.

Step 3 Right-click WWPN Pools and select **Create WWPN Pool**.

Step 4 In the **Define Name and Description** dialog box of the **Create WWPN Pool** wizard, enter **`hx-a`**.

Step 5 Click **Next**.

Step 6 In the **Add WWN Blocks** page of the **Create WWNN Pool** wizard, click **Add**.

Step 7 In the **Create WWN Block** dialog box, complete the following fields:

Form field: The first WWN in the block.

Size field: The number of WWNs in the block.

For WWN pools, the pool size must be a multiple of ports-per-node + 1. For example, if there are seven ports per node, the pool size must be a multiple of eight. If there are 63 ports per node, the pool size must be a multiple of 64.

Step 8 Click **OK**.

Step 9 Click **Finish**.

What to do next

Create WWPN Pool **hx-b**. Follow the steps above.

Creating a vHBA Template

This template is a policy that defines how a vHBA on a server connects to the SAN. It is also referred to as a vHBA SAN connectivity template. Include this policy in a service profile for it to take effect.

Before you begin

Before creating the vHBA template policy, make sure that one or more of the following resources exist in the system:

- Named VSAN
- WWNN pool or WWPN pool
- SAN pin group
- Statistics threshold policy

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Policies** > **root** > **Sub-Organizations** > **hx-cluster**.

Step 3 Right-click the **vHBA Templates** node and choose **Create vHBA Template**.

Step 4 In the **Create vHBA Template** dialog box, complete the following fields:

Name	Description
Name field	Enter **vhba-a**. The name of the virtual HBA template. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than hyphen (-), underscore(_), period(.), and colon (:). You cannot change this name after the object is saved.
Description field	Enter up to 256 characters. A user-defined description of the template.
Fabric ID field	Select **A**.

Name	Description
Select VSAN drop-down list	Select the VSAN created earlier for Fabric A, to associate with this vHBA.
Template Type field	Select **Updating Template**. vHBAs created from this template are updated if the template changes.
Max Data Field Size field	Default: 2048 This the maximum size of the Fibre Channel frame payload bytes that the vHBA supports.
WWPN Pool drop-down list	Assign **hx-a**.
QoS Policy drop-down list	*<Not set>*
Pin Group drop-down list	*<Not set>*
Stats Threshold Policy drop-down list	*<Not set>*

Step 5 Click **Ok**.

SAN Connectivity Policy

Connectivity policies determine the connections and the network communication resources between the server and the SAN in the network. These policies use pools to assign MAC addresses, WWNs, and WWPNs to servers and to identify the vNICs and vHBAs that the servers use to communicate with the network.

Note We do not recommend that you use static IDs in connectivity policies, because these policies are included in service profiles and service profile templates. Also, connectivity policies can be used to configure multiple servers.

Creating SAN Connectivity Policy

Step 1 In the **Navigation** pane, click **SAN**.

Step 2 Expand **SAN** > **Policies** > **root** > **Sub-Organizations** > **hx-cluster**.

Step 3 Right-click **SAN Connectivity Policies** and choose **Create SAN Connectivity Policy**.

Step 4 In the **Create SAN Connectivity Policy** dialog box, enter a **Name** **`Hyperflex`** and optional **Description**.

Step 5 From the **WWNN Assignment** drop-down, select the **`hyperflex`** pool.

Each pool name is followed by two numbers in parentheses that show the number of WWNs still available in the pool and the total number of WWNs in the pool.

Step 6 Click **Add**.

Step 7 In the **Create vHBAs** dialog box, enter the name **`vhba-a`**.

Step 8 Check **Use vHBA template**.

Step 9 Select **vHBA template `vhba-a`** from the drop-down list.

Step 10 Select **Adapter Policy `VMware`** from the drop-down list.

Step 11 Click **Ok**.

Step 12 Repeat steps 7- 12 to create vHBA **`vhba-b`** and assign **`vhba-b`** template to it.

What to do next

Include the SAN connectivity policy to the HX node service profile template.

Including SAN Connectivity Policy to the HX Node Service Profile Template

This procedure causes the Service Profiles associated with this SPT to require user acknowledgement and the HX node to reboot.

Step 1 In the **Navigation** pane, click **Server**.

Step 2 Expand **Servers** > **Service Profile Template** > **root** > **Sub-Organizations** > **hx-cluster**.

Step 3 Select **`Service Template hx-nodes`**, select **`vHBA`**.

Step 4 In the work pane, on **Storage** tab, select **`HyperFlex`** from the drop-down list under the **SAN Connectivity Policy** section.

Step 5 Click **Save**.

Adding Additional vNICs to an Existing Cluster

Before you begin

In order to connect to other storage systems such as FlexPod via iSCSI or NFS, or a FC SAN, it is recommended that the additional vNICs be added during the creation of the HX cluster. The HyperFlex installer prompts for the optional creation of additional iSCSI vNICs or FC HBAs at the time of installation and should be configured if external storage is required now or at some point in the future.

HyperFlex supports adding additional vNICs after cluster creation. To add additional vNICs to an existing cluster, perform the following actions:

Note Do not reboot multiple nodes at once while making these hardware changes, as it could lead to the storage cluster going offline. Validate the health state of each host, and the HX cluster before moving onto subsequent nodes.

Note In some rare cases, vmnics in ESXi may re-order and require manual reconfiguration to restore network services. Before beginning this procedure, run and save the output of the following commands on **every** ESXi host in the cluster via SSH:

```
esxcli network nic list
esxcli network vswitch standard list
esxcli network vswitch standard policy failover get -v vswitch-hx-inband-mgmt
esxcli network vswitch standard policy failover get -v vswitch-hx-storage-data
esxcli network vswitch standard policy failover get -v vmotion
esxcli network vswitch standard policy failover get -v vswitch-hx-vm-network
```

Step 1 Login to UCSM and click on **LAN** tab, navigate to **Policies** > **root** > **Sub-Organizations** > ***Name of the Suborg for this Cluster*** > **vNIC Templates**. Right-click and select **Create vNIC Template**.

Step 2 From the **LAN** tab, navigate to **Policies** > **root** > **Sub-Organizations** > ***Name of the Suborg for this Cluster*** > **LAN Connectivity Policies** > **HyperFlex**. Click **Add** at the bottom of the table. Specify a name, check the box for **Use vNIC Template**, and then select the template created in **Step 1**. Finally click **Save Changes** and review any warnings that may be triggered.

After adding the vNIC template to the LAN connectivity policy, the servers go in to **Pending Reboot** state and require a reboot to add the new interface.

Note Do not reboot HX servers at this time.

Step 3 Log in to your vCenter Server as a user with administrative privileges on the cluster.

Step 4 Place one of the existing HX ESXi Hosts in to **Maintenance Mode**.

Step 5 After the host has entered **Maintenance Mode**, reboot the associated node to complete the addition of the new hardware.

Step 6 In some configurations, after the node has rebooted, the HXDP software detects that the DirectPath I/O configuration has changed, and must be reconfigured. This results in one additional automatic reboot of the node.

Note After the second reboot, exit the ESXi host from maintenance mode, the SCVM should start automatically without errors.

Step 7 Check the health status of the cluster, validating that the cluster is healthy before proceeding to reboot the next node. The cluster health status can be viewed from HyperFlex Connect.

Step 8 Repeat steps 3 through 7 for each node in the cluster as necessary, until all of the nodes have been rebooted and the new vNICs are visible to ESXi as new vmnic interfaces.

Step 9 Create a new vSwitch and assign the new vmnics as uplinks. Do not alter the existing HyperFlex vSwitches or vmnics.

Note Be sure to create new vSwitches for any additional vNICs added to the cluster.

Adding vHBAs to an Existing HyperFlex Cluster

In order to connect to an external block storage system using FC SAN, it is recommended that the FC vHBAs be added during the creation of the HX cluster. The HyperFlex installer prompts for the optional creation of FC vHBAs at the time of installation and should be configured if external storage is required now or at some point in the future.

HyperFlex supports adding new FC HBAs for a new SAN storage connectivity or adding additional vHBAs to an existing SAN Connectivity policy for additional SAN connection after cluster creation with a SAN connectivity policy.

Note Do not reboot multiple nodes at once while making these hardware changes, as it could lead to the storage cluster going offline. Validate the health state of each host, and the HX cluster before moving onto subsequent nodes.

Creating a SAN Connection Policy without SAN Storage Connected

Use these steps in order to create a new SAN connectivity policy and attach it to existing HyperFlex cluster nodes for connecting new external SAN storage. If your cluster already has a external SAN storage connected, see Creating a FC HBA to an Existing SAN Connection Policy with External Storage Connected, on page 34.

Step 1 Login to UCSM and perform the following steps:

a) Click on the **SAN** tab and navigate to **Pools** > **root** > **Sub-Organizations** > ***Name of the Suborg for this Cluster*** > **WWNN Pools**.
b) Right-click on **WWNN Pools** and select **Create WWNN Pool**.
c) Enter the WWNN Pool name and click **Next** and the **Add** at the bottom of the table.
d) You have the option to edit the last 6 characters of WWNN and select the size.

Note The size should be equal or more than the number of nodes in the HyperFlex cluster.

e) Click **Finish**.

Step 2 Create 2 WWPN policies, one for **SAN A** and one for **SAN B** by performing the following steps:

a) Login to your UCSM.
b) Click on the **SAN** tab and navigate to **Pools** > **root** > **Sub-Organizations** > ***Name of the Suborg for this Cluster*** > **WWPN Pools**.
c) Right-click on **WWPN Pools** and select **Create WWPN Pool**.
d) Enter the WWPN Pool name and click **Next** and the **Add** at the bottom of the table.
e) You have the option to edit the last 6 characters of WWPN and select the size.

Note It is recommended to change one or more characters of the last 6 characters of WWPN so that its identified easily for each SAN fabric. The size should be equal or more than the number of nodes in the HyperFlex cluster.

f) Click **Finish** and repeat this process for FC SAN B.

Step 3 From the **SAN** tab, navigate to **SAN Cloud** > **Fabric A** > **VSANs** and perform the following steps:

a) Right-click and select **Create VSAN**.
b) Right-click on **WWNN Pools** and select **Create WWNN Pool**.
c) Enter the **VSAN Name** and select **Fabric A** from the radio button options.
d) Enter the **VSAN ID** and the corresponding **FCoE VSAN ID**.

Note FCoE VSAN ID can be the same as the VSAN ID.

Step 4 From the **SAN** tab, navigate to **SAN Cloud** > **Fabric B** > **VSANs** and perform the following steps:

a) Right-click and select **Create VSAN**.
b) Right-click on **WWNN Pools** and select **Create WWNN Pool**.
c) Enter the **VSAN Name** and select **Fabric B** from the radio button options.
d) Enter the **VSAN ID** and the corresponding **FCoE VSAN ID**.

Note FCoE VSAN ID can be the same as the VSAN ID.

Note Make sure to use different VSAN IDs in Fabric A and Fabric B.

Step 5 From the **SAN** tab, navigate to **Policies** > **root** > **Sub-Organizations** > ***Name of the Suborg for this Cluster*** > **vHBA Templates** and perform the following steps:

a) Right-click and select **Create vHBA Template**.
b) Enter the **vHBA Name** and select **Fabric ID A**.
c) From the **Select VSAN** drop-down, select the VSAN previously created for SAN A in Step 3.
d) From the **Template Type** field, select **Updating Template**. In the dropdown for **WWPN Pool**, select the WWPN Pool created for SAN A in Step 2.

Step 6 From the **SAN** tab, navigate to **Policies** > **root** > **Sub-Organizations** > ***Name of the Suborg for this Cluster*** > **vHBA Templates** and perform the following steps:

a) Right-click and select **Create vHBA Template**.
b) Enter the **vHBA Name** and select **Fabric ID B**.
c) From the **Select VSAN** drop-down, select the VSAN previously created for SAN B in Step 4.
d) From the **Template Type** field, select **Updating Template**. In the dropdown for **WWPN Pool**, select the WWPN Pool created for SAN B in Step 2.

Step 7 From the **SAN** tab, navigate to **Policies** > **root** > **Sub-Organizations** > ***Name of the Suborg for this Cluster*** > **SAN Connectivity Policies** and perform the following steps:

a) Right-click and select **Create SAN Connectivity Policy**.
b) Enter the **SAN Connectivity Policy** name.
c) From the **WWNN Assignment** drop-down, select the WWNN Pool previously created in Step 1.
d) Click **Add** at the bottom of the table and enter the WWNN Pool name.
e) Select **Use vHBA Template** and in the vHBA Template drop-down, select the vHBA Template for SAN A previously created in Step 5 and click **OK**.
f) Click **Add** at the bottom of the table again and enter the WWNN Pool name.
g) Select **Use vHBA Template** and in the vHBA Template drop-down, select the vHBA Template for SAN B previously created in Step 6 and click **OK**.

Step 8 Navigate to **Servers** > **Service Profiles** > **root** > **Sub-Organizations** > ***Name of the Suborg for this Cluster***.

a) Click on one of the service profiles, and from the **General** tab, click on **Template Instance**.
b) From the **Service Template** pop-up window under **Properties**, navigate to the **Storage** > **vHBA** tab.
c) In the **SAN Connectivity Policy** section, select the SAN Connectivity policy created in Step 7 and click **Apply**. Click **Yes** in the pop-up window.

Note If you have a cluster with Mixed Node types such as M4/M5/Compute, make sure to identify the Service Profile Template for different node types and update the Service Profile Template to add SAN Connectivity Policy.

d) After adding the SAN Connectivity Policy to the Service Profile Template, the servers go in to Pending Reboot state and require a reboot to add the new FC HBA interface.

Note Do not reboot HX servers at this time.

Step 9 Log in to your vCenter Server as a user with administrative privileges on the cluster.

Step 10 Place one of the existing HX ESXi Hosts in to **Maintenance Mode**.

Step 11 After the host has entered **Maintenance Mode**, reboot the associated node to complete the addition of the new hardware.

Step 12 After the reboot, exit the ESXi host from **Maintenance Mode**, the SCVM should start automatically without errors.

Step 13 Check the health status of the cluster, validating that the cluster is healthy before proceeding to reboot the next node. The cluster health status can be viewed from HyperFlex Connect.

Step 14 Repeat steps 10 through 13 for each node in the cluster as necessary, until all of the nodes have been rebooted and the new FC HBAs are visible to ESXi as new vHBA interfaces.

Step 15 Check and confirm that there are no more pending acknowledgements after completing the reboot of all hosts in the cluster.

Creating a FC HBA to an Existing SAN Connection Policy with External Storage Connected

Use these steps in order to create a new FC HBA and add to existing SAN connectivity policy. If your cluster doesn't have a SAN connection policy, see Creating a SAN Connection Policy without SAN Storage Connected, on page 32.

Step 1 Login to UCSM and click on the **SAN** tab, navigate to **Policies** > **root** > **Sub-Organizations** > ***Name of the Suborg for this Cluster*** > **vHBA Templates** and perform the following steps:

a) Right-click and select **Create vHBA Template**.
b) Enter the **vHBA Name** and select **Fabric ID A**.
c) From the **Select VSAN** drop-down, select the VSAN for SAN A.
d) From the **Template Type** field, select **Updating Template**. In the dropdown for **WWPN Pool**, select the WWPN Pool created for SAN A.

Step 2 Login to UCSM and click on the **SAN** tab, navigate to **Policies** > **root** > **Sub-Organizations** > ***Name of the Suborg for this Cluster*** > **vHBA Templates** and perform the following steps:

a) Right-click and select **Create vHBA Template**.
b) Enter the **vHBA Name** and select **Fabric ID B**.
c) From the **Select VSAN** drop-down, select the VSAN for SAN B.
d) From the **Template Type** field, select **Updating Template**. In the dropdown for **WWPN Pool**, select the WWPN Pool created for SAN B.

Note If you want to use a new VSAN for the additional FC HBAs, you can create new VSANs under **SAN** > **SAN Cloud** > **Fabric A/Fabric B**.

Step 3 From the **SAN** tab, navigate to **Policies** > **root** > **Sub-Organizations** > ***Name of the Suborg for this Cluster*** > **SAN Connectivity Policies** > **Hyperflex** and perform the following steps:

a) Click **Add** at the bottom of the table and enter the WWNN Pool name.
b) Select **Use vHBA Template** and in the vHBA Template drop-down, select the vHBA Template for SAN A previously created in Step 1 and click **OK**.
c) Repeat this step for vHBA Template for SAN B.

Step 4 Click **Save Changes** and review any warnings that may be triggered.

Step 5 After adding the SAN Connectivity Policy to the Service Profile Template, the servers go in to Pending Reboot state and require a reboot to add the new FC HBA interface.

Note Do not reboot HX servers at this time.

Step 6 Log in to your vCenter Server as a user with administrative privileges on the cluster.

Step 7 Place one of the existing HX ESXi Hosts in to **Maintenance Mode**.

Step 8 After the host has entered **Maintenance Mode**, reboot the associated node to complete the addition of the new hardware.

Step 9 After the reboot, exit the ESXi host from **Maintenance Mode**, the SCVM should start automatically without errors.

Step 10 Check the health status of the cluster, validating that the cluster is healthy before proceeding to reboot the next node. The cluster health status can be viewed from HyperFlex Connect.

Step 11 Repeat steps 7 through 10 for each node in the cluster as necessary, until all of the nodes have been rebooted and the new FC HBAs are visible to ESXi as new vHBA interfaces.

Step 12 Check and confirm that there are no more pending acknowledgements after completing the reboot of all hosts in the cluster.

Attaching iSCSI Storage to HX

This procedure describes the high-level steps for attaching iSCSI storage to the HX FI Domain:

Step 1 Log in to the Cisco UCS Manager GUI.

Step 2 Create a VLAN.

Step 3 Create MAC Pool Addresses for iSCSI storage. For details, see Creating MAC Address Pools for External Storage, on page 37.

Step 4 Create a pair of vNIC templates associated with Fabric-A and Fabric-B respectively. See Creating a vNIC Template for iSCSI Storage, on page 38 for detailed steps.

Step 5 Create HyperFlex LAN connectivity Policy. See Creating a LAN Connectivity Policy, on page 40 for detailed steps.

Step 6 Assign the HX LAN connectivity policy to the HX Service Profile Template (SPT) used for the cluster. It triggers a pending-ack to be raised on all nodes within the cluster created from the modified SPT. Acknowledge all pending-ack's for all Service Profiles in the cluster to trigger re-configuration of Service Profiles with vNICs. For detailed steps, refer to Creating a LAN Connectivity Policy, on page 40.

Step 7 Add Network and Storage Adaptors. See Adding Network Adapters, on page 40 for detailed steps.

iSCSI SAN Concepts

The iSCSI SANs use Ethernet connections between computer systems, or host servers, and high-performance storage subsystems. SAN components include iSCSI Host Bus Adapters (HBAs) or Network Interface Cards (NICs) in the host servers, switches, and routers that transport the storage traffic, cables, storage processors, and storage disk systems.

The iSCSI SANs use a client-server architecture. The client, called an iSCSI initiator, operates on the host. The client initiates iSCSI sessions by issuing iSCSI commands and transmitting them, encapsulated using the iSCSI protocol, to a server. The server, called an iSCSI target, represents a physical storage system on the network. The target can also be provided by a virtual iSCSI SAN, for example, an iSCSI target emulator

running in a virtual machine. The iSCSI target responds to the initiator's commands by transmitting the required iSCSI data.

Discovery, Authentication, and Access Control

You can use several mechanisms to discover your storage and to limit access to it. Configure the host and the internet SCSI (iSCSI) storage system to support your storage access control policy.

How Virtual Machines Access Data on an iSCSI SAN

ESXi stores the disk files from a virtual machine within a VMFS datastore that resides on a SAN storage device. When virtual machine guest operating systems issue iSCSI commands to their virtual disks, the SCSI virtualization layer translates these commands to VMFS file operations. Depending on which port the iSCSI initiator uses to connect to the network, Ethernet switches and routers carry the request to the storage device that the host wants to access.

Using ESXi with iSCSI SAN

Using ESXi together with a SAN provides storage consolidation, improves reliability, and helps with disaster recovery. When you set up ESXi hosts to use iSCSI SAN storage systems, you must be aware of certain special considerations that exist.

Best Practices for iSCSI Storage

When using ESXi with the iSCSI SAN, follow VMware best practices to avoid problems.

Check with your storage representative if your storage system supports Storage API - Array Integration hardware acceleration features. If it does, refer to your vendor documentation for information on how to enable hardware acceleration support on the storage system side.

Preventing iSCSI SAN Problems

When using ESXi with the SAN, follow these specific guidelines to avoid problems with the SAN configuration:

- Place only one VMFS datastore on each LUN. Multiple VMFS datastores on a LUN is not recommended.
- Do not change the path policy the system sets, unless you understand the implications of making such a change.
- Document everything, include information about configuration, access control, storage, switch, server, iSCSI HBA configuration, software and firmware versions, and storage cable plan.
- Plan for failure. Make several copies of the topology maps. For each element, consider what happens to the SAN if the element fails.
- Cross off different links, switches, HBAs, and other elements to ensure that you did not miss a critical failure point in your design.
- Ensure that the iSCSI HBAs are installed in the correct slots in the ESXi host, based on the slot and bus speed. Balance PCI bus load among the available buses in the server.
- Become familiar with the various monitor points in your storage network, at all visibility points, including ESXi performance charts, Ethernet switch statistics, and storage performance statistics.
- Be cautious when changing IDs of the LUNs that have VMFS datastores being used by the host. If the ID is changed, the virtual machines running on the VMFS datastore fail. If there are no running virtual

machines on the VMFS datastore, after the ID of the LUN is changed, use rescan to reset the ID on your host. For information on using rescan, see `Storage Refresh and Rescan Operations`.

- If you need to change the default iSCSI name of your iSCSI adapter, make sure the name used is worldwide unique and properly formatted. To avoid storage access problems, never assign the same iSCSI name to different adapters, even on different hosts.
- Ensure that the iSCSI traffic and uplinks are segregated on their own dedicated vSwitch.

Creating VLAN to add iSCSI Storage to HX FI Domain

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 following table:

VLAN Name	Description	Multicast Policy Name	VLAN ID (by default)
hx-extstorage-iscsi	Used for adding external storage connectivity	HyperFlex	4201

Note
- Configuration option is Common/Global. It applies to both fabrics and uses the same configuration parameters in both cases.
- Sharing type is set to None.

Step 4 Click **Ok**.

What to do next

Create MAC Pool for the external storage.

Creating MAC Address Pools for External Storage

Change the default MAC address blocks to avoid duplicate MAC addresses that already exist. Each block contains **100 MAC addresses** by default to allow for up to 100 HX servers 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. The *A* is set on vNICs pinned to Fabric Interconnect (FI) A. The *B* is set on vNICs pinned to FI B.

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 in the following table:

MAC Pool Name	**Description**	**Assignment Order**	**MAC Address block**
hx-extstorage-a	MAC pool for adding external storage to HyperFlex System	Sequential	00:25:B5:XX:01:01-63

Note Make sure to check the last block of MAC addresses and use next order of block to create the new MAC pools for both fabrics.

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**
First MAC Address field	The first MAC address in the block.
Size field	The number of MAC addresses in the block.

Step 8 Click **OK**.

Step 9 Click **Finish**.

What to do next

Repeat steps to create MAC Pool - **h-extstorage-b** for FI b.

Creating a vNIC Template for iSCSI Storage

This template is a policy that defines how the vNIC on a server connects to the LAN. It is also called a vNIC LAN connectivity template. You must include this policy in a service profile for it to take effect.

Before you begin

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

- Named VLAN
- MAC pool
- Jumbo MTU
- QoS policy

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 choose **Create vNIC Template**.

Step 3 In the **Create vNIC Template** dialog box, complete the following fields:

Name	Description
Name field	Enter **extstorage_iscsi-a** This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A user-defined description of the template. Enter up to 256 characters.
Fabric ID field	Select **A**
Redundancy drop-down list	Primary
Target	Adapter
Template Type field	Select **Updating Template**. vNICs created from this template are updated if the template changes.
VLAN field	hx-extstorage-iscsi (what you created above)
CDN Source	vNIC Name
MTU drop-down list	**9000**
MAC Pool	**hx-extstorage-a** (created earlier)
QoS Policy drop-down list	**Bronze**
Connection	**Dynamic**

Step 4 Click **Ok**.

Step 5 Repeat the workflow to create a vNIC template with Fabric ID B as primary.

LAN Connectivity Policy

Connectivity policies determine the connections and the network communication resources between the server and the LAN in the network. These policies use pools to assign MAC addresses to servers and to identify the vNICs that the servers use to communicate with the network.

Note We recommend that you *do not* use static IDs in connectivity policies, because these policies are included in service profiles and service profile templates. Also, connectivity policies can be used to configure multiple servers.

Creating a LAN Connectivity Policy

Step 1 In the Navigation pane, click the **LAN** tab.

Step 2 On the **LAN** tab, expand **LAN > Sub-Org > hx-cluster > LAN Connectivity Policies> HyperFlex**.

Step 3 Click **Add vNICs**.

Step 4 In the **Create vNIC** dialog box, enter a name. **Check Use vNIC Template** and **Redundancy Pair**.

Example: iscsi-A

Step 5 Enter **Peer Name**.

Example: iscsi-B

Step 6 Select **vNIC Template** name *iscsi-A* from the drop-down list. Click **Ok**.

Step 7 Repeat steps 3 - 6 to create vNIC *iscsi-B* and assign *vNIC-b* template to it.

Step 8 Click **Save Changes**. In the Save Changes box that displays, click **Yes** to accept the changes.
Includes the LAN connectivity policy to the HX node service profile template.

Including LAN Connectivity Policy to the HX Node Service Profile Template

Step 1 Navigate to the **Server** tab and expand **root > Sub-Org > hx-cluster > Service Template hx-nodes**.

Step 2 In the work pane, on the **Network** tab, select **`HyperFlex`** from the drop-down list under the **LAN Connectivity Policy** section.

Step 3 Click **Modify vNIC/HBA Placement**. Check the iscsi vNIC for proper order. Make sure they are last in the order. Re-arrange as necessary.

Note If you are adding both FC and iSCSI storage, then the order of vHBAs precedes the order of the vNICs.

Step 4 Click **Save**.

This procedure causes the Service Profiles associated with this SPT to require user acknowledgment and the HX node reboot.

Adding Network Adapters

Step 1 Go to VMware vCenter.

Step 2 Select the **HX node**.

Step 3 Navigate to **Configuration** > **Hardware**.

Step 4 Click **Network Adapters**.

Step 5 Make sure the iSCSI vmnics that were created are visible in **Network Adapters**. Click **Networking**.

Step 6 Select **Create vSphere Standard Switch** to create a new vSwitch and select the two vNICs that you added earlier.

Step 7 Click **Next**.

Step 8 Create a port group with connection type set as **`VMkernel`**. Under **Port Group Properties**, enter the **`Network Label`**, **`VLAN ID`**, and the **`Network Type`**.

Step 9 Click **Next** and enter the **`IP Address`**, which is the address of the iSCSI initiator.

Step 10 Click **Finish**. The vswitch gets created with one vmkernel port group.

Attention If you choose to have multipath, add another vmkernel port by clicking **Properties** of the vswitch you just created and clicking **Add**. Follow Steps 8 and 9. Click **Finish**.

If you have multipathing with two vmkernel ports, set the NIC teaming policy.

Click the first vmkernel port **> Edit > NIC teaming** check box to override the switch failover order. Push one vnic adapter to active and the other to unused. For the second vmkernel port, do the opposite.

Adding Storage Adapters

Step 1 Go to VMware vCenter.

Step 2 Select the HX node.

Step 3 Navigate to **Configuration** > **Hardware**.

Step 4 Click **Storage Adapters**.

Step 5 Select the **USB Storage Controller**.

Step 6 Click **Add**.

Step 7 Click **Add iSCSI Software Adapter**.

Step 8 Click **Ok**.

Step 9 A new **Software iSCSI adapter** is added to the **Storage Adapter** list. After it has been added, select the **Software iSCSI Adapter** from the list.

Step 10 Click **Properties** > **Network Configuration tab**.

Step 11 Click **Add**.

Step 12 You will now see *vNIC 10* and *vNIC 11*. Choose both and click **Ok**.

Step 13 Navigate to the **Dynamic Discovery tab**.

It will be blank.

Step 14 Click **Add the iSCSI target IP against the iSCSI adapter**. Use as default port.

Step 15 Click **Ok**.

The target is populated. Repeat the above steps for second IP multipathing.

Note The software asks for rescan. Click **Yes**. If you already have an iSCSI LUN assigned, it displays under the devices in **Software- Configuration** (thick client).

Step 16 Navigate to **Storage**.

Step 17 On the **Configuration tab**, click **Add Storage**.

Step 18 Select **Storage Type**, **`Disk or LUN`**. Click **Next**.

Step 19 The LUN shows up now. Select it and click **Next**.

Step 20 Enter the **Datastore Name** and click **Finish**.

Connecting Cisco HX Servers to External NFS Storage

Network File System

An NFS client built into ESXi uses the Network File System (NFS) protocol over TCP/IP to access a designated NFS volume that is located on a NAS server. The ESXi host can mount the volume and use it for its storage needs.

ESXi supports the following storage capabilities on most NFS volumes:

- VMotion and Storage vMotion
- High Availability (HA)
- Distributed Resource Scheduler (DRS)

NFS Storage Guidelines and Requirements

When using NFS storage, use the following configuration, networking, and NFS datastore guidelines.

NFS Server Configuration Guidelines

- Make sure that NFS servers you use are listed in the VMware HCL. Use the correct version for the server firmware.
- When configuring NFS storage, follow the recommendations of your storage vendor.
- Ensure that the NFS volume is exported using NFS over TCP.
- Ensure that each host has root access to the volume. If the NAS server does not grant root access, you might still be able to mount the NFS datastore on the host. However, you will not be able to create any virtual machines on the datastore.
- Make sure that the NFS server does not provide both protocol versions for the same share.
- If the underlying NFS volume, on which files are stored, is read-only, make sure that the volume is exported as a read-only share by the NFS server, or configure it as a read-only datastore on the ESXi host. Otherwise, the host considers the datastore to be read-write and might not be able to open the files.

NFS Networking Guidelines

- For network connectivity, the host requires a standard network adapter.
- ESXi supports Layer 2 and Layer 3 Network switches. If you use Layer 3 switches, ESXi hosts and NFS storage arrays must be on different subnets and the network switch must handle the routing information.
- Ensure that the NFS traffic and uplinks are segregated on their own dedicated vSwitch.

- A VMkernel port group is required for NFS storage. Add a VMkernel port group for IP storage on a new virtual switch (vSwitch). The vSwitch can be a vSphere Standard Switch (VSS) or a vSphere Distributed Switch (VDS).
- If you use multiple ports for NFS traffic, make sure that you correctly configure your virtual switches and physical switches. For information, see the vSphere Networking documentation.

Note For details on configuring NFS storage, consult your storage vendor documentation.

Creating the vSwitch, Adapter, and Port Group for NFS Storage

Step 1 In vCenter web client, go to **Inventory > Hosts and Clusters > DC > host** and perform the following steps:

a) Under **Configuration tab**, click **Networking > Add Networking**.

b) In the wizard box, select **VMkernel**, and then click **Next**.

Step 2 Select **Create vSphere Standard switch** and select the available vmnic. Click **Next**.

Step 3 Enter a name for the Port Group. For example: NFS

For VLAN ID, leave it as default in case your environment has native VLANs.

Step 4 Enter the **IP Setting** information, click **Next**, and then click **Finish**.

Setup NFS Storage Environment

Step 1 On the NFS server, configure an NFS volume and export it to be mounted on the ESXi hosts. Note the IP address or the DNS name of the NFS server and the full path, or folder name, for the NFS share.

Note Make sure that each host that mounts this datastore is a part of an Active Directory domain and its NFS authentication credentials are set.

Step 2 In vCenter thick client, go to **Storage**. Under **Configuration tab**, click **Add Storage**. In the wizard box, select **Network File System (NFS)**. And click **Next**.

Step 3 You are at the Local NFS wizard. Enter the **`target server IP address`** and the **`path for the address`**. Give your datastore a **`name`**. Click **Next** and **Finish**.

Note Note: No change is required in Cisco UCS Manager for service profile templates, service profiles, and policies. For more information, see the **vSphere Networking documentation**.

Fibre Channel Zoning

Fibre Channel (FC) zoning allows you to partition the FC fabric into one or more zones. Each zone defines the set of FC initiators and FC targets that can communicate with each other in a VSAN. Zoning also enables you to set up access control between hosts and storage devices or user groups.

Information About Zones

A zone consists of multiple zone members and has the following characteristics:

- Members in a zone can access each other; members in different zones cannot access each other.
- Zones can vary in size.
- Devices can belong to more than one zone.

A physical fabric can have a maximum of 8,000 zones.

Fibre Channel Zoning in Cisco UCS Manager

Cisco UCS Manager supports switch-based Fibre Channel (FC) zoning and Cisco UCS Manager-based FC zoning. You cannot configure a combination of zoning types in the same Cisco UCS domain. You can configure a Cisco UCS domain with one of the following types of zoning:

- Cisco UCS Manager-based FC zoning — Combines direct attached storage with local zoning. FC or FCoE storage connects directly to the fabric interconnects (FIs). You perform zoning in Cisco UCS Manager, using Cisco UCS local zoning. Disable any existing FC or FCoE uplink connections. Cisco UCS does not currently support active FC or FCoE uplink connections coexisting with the utilization of the Cisco UCS Local Zoning feature.
- Switch-based FC zoning — Combines direct attached storage with uplink zoning. The FC or FCoE storage connects directly to the FIs. You perform zoning externally to the Cisco UCS domain through an MDS or Nexus 5000 switch. This configuration does not support local zoning in the Cisco UCS domain. With switch-based zoning, a Cisco UCS domain inherits the zoning configuration from the upstream switch.

Note Zoning is configured on a per-VSAN basis. You cannot enable zoning at the fabric level.

Recommendations

- If you want Cisco UCS Manager to handle FC zoning, the FIs must be in Fibre Channel Switch mode. You cannot configure FC zoning in End-Host mode.
- If a Cisco UCS domain is configured for high availability with two FIs, we recommend that you configure both FIs with the same set of VSANs.

Configuring Fibre Channel Zoning

SUMMARY STEPS

1. If you have not already done so, disconnect the fabric interconnects in the Cisco UCS domain from any external Fibre Channel switches, such as an MDS.
2. If the Cisco UCS domain still includes zones that were managed by the external Fibre Channel switch, run the `clear-unmanaged-fc-zone-all` command on every affected VSAN to remove those zones.
3. Configure the Fibre Channel switching mode for both fabric interconnects in Fibre Channel Switch mode.
4. Configure the Fiber Channel and FCoE storage ports that you require to carry traffic for the Fibre Channel zones.

DETAILED STEPS

	Command or Action	Purpose
Step 1	If you have not already done so, disconnect the fabric interconnects in the Cisco UCS domain from any external Fibre Channel switches, such as an MDS.	
Step 2	If the Cisco UCS domain still includes zones that were managed by the external Fibre Channel switch, run the `clear-unmanaged-fc-zone-all` command on every affected VSAN to remove those zones.	This functionality is not currently available in the Cisco UCS Manager GUI. You must perform this step in the Cisco UCS Manager CLI.
Step 3	Configure the Fibre Channel switching mode for both fabric interconnects in Fibre Channel Switch mode.	You cannot configure Fibre Channel zoning in End-Host mode. See http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Storage-Mgmt/3-1/b_UCSM_GUI_Storage_Management_Guide_3_1/b_UCSM_GUI_Storage_Management_Guide_3_1_chapter_01110.html#task_B6E0C2A15FE84D498503ADC19CDB160B
Step 4	Configure the Fiber Channel and FCoE storage ports that you require to carry traffic for the Fibre Channel zones.	See Configuring an Ethernet Port as an FCoE Storage Port and Configuring a Fibre Channel Storage Port. Refer the following link: http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Storage-Mgmt/3-1/b_UCSM_GUI_Storage_Management_Guide_3_1/b_UCSM_GUI_Storage_Management_Guide_3_1_chapter_01100.html#task_A33D13CA58924EB1AD35EBA473B92625

Direct Attached Storage

A typical Direct Attached Storage (DAS) system is made of a data storage device connected directly to a computer through a host bus adapter (HBA). Between those two points there is no network device (like a

switch or router). The main protocols used for DAS connections are ATA, SATA, eSATA, SCSI, SAS, USB, USB 3.0, IEEE 1394 and Fibre Channel.

Cisco UCS Manager allows you to have DAS without the need for a SAN switch to push the zoning configuration. The DAS configuration described assumes that the physical cables are already connected between the storage array ports and the Fabric Interconnects.

Note VSAN is created in the SAN Cloud tab, even when the storage is directly attached.

Fibre Channel Switching Mode

The Fibre Channel switching mode determines how the fabric interconnect behaves as a switching device between the servers and storage devices. The fabric interconnect operates in either of the following Fibre Channel switching modes:

End-Host Mode

End-host mode allows the fabric interconnect to act as an end host to the connected fibre channel networks, representing all servers (hosts) connected to it through virtual host bus adapters (vHBAs). This behavior is achieved by pinning (either dynamically pinned or hard pinned) vHBAs to Fibre Channel uplink ports, which makes the Fibre Channel ports appear as server ports (N-ports) to the rest of the fabric. When in end-host mode, the fabric interconnect avoids loops by denying uplink ports from receiving traffic from one another.

End-host mode is synonymous with N Port Virtualization (NPV) mode. This mode is the default Fibre Channel Switching mode.

Note When you enable end-host mode, if a vHBA is hard pinned to an uplink Fibre Channel port and this uplink port goes down, the system cannot re-pin the vHBA, and the vHBA remains down.

Switch Mode

Switch mode is the traditional Fibre Channel switching mode. Switch mode allows the fabric interconnect to connect directly to a storage device. Enabling Fibre Channel switch mode is useful in Pod models where there is no SAN (for example, a single Cisco UCS domain that is connected directly to storage), or where a SAN exists (with an upstream MDS). Switch mode is not the default Fibre Channel switching mode.

Note In Fibre Channel switch mode, SAN pin groups are irrelevant. Any existing SAN pin groups are ignored.

Configuring Fibre Channel Switching Mode

Important When you change the Fibre Channel switching mode, Cisco UCS Manager's behaviour depends on its version.

In UCS Manager version 3.1(1) and earlier releases, Cisco UCS Manager restarts both fabric interconnects simultaneously.

When the Fibre Channel switching mode is changed, both Cisco UCS fabric interconnects will reload simultaneously. Reloading of fabric interconnects will cause a system-wide downtime for approximately 10-15 minutes.

In UCS Manager version 3.1(2), when the Fibre Channel switching mode is changed, the UCS fabric interconnects reload sequentially, the second fabric interconnect can take several minutes to complete the change in Fibre Channel switching mode and become system ready.

In UCS Manager Release 3.1(3) and later releases, the subordinate fabric interconnect reboots first as a result of the change in switching mode. The primary fabric interconnect reboots only after you acknowledge it in Pending Activities. The primary fabric interconnect can take several minutes to complete the change in Fibre Channel switching mode and become system ready.

For more information, see the Cisco UCS Manager Storage Management Guide.

Step 1 In the **Navigation** pane, click **Equipment**.

Step 2 Expand **Equipment** > **Fabric Interconnects** > **Fabric_Interconnect_Name**.

Step 3 In the Work pane, click the **General** tab.

Step 4 In the **Actions** area of the **General** tab, click one of the following links:

- **Set FC Switching Mode**
- **Set FC End-Host Mode**

  The link for the current mode is dimmed.

Step 5 In the dialog box, click **Yes**.

Cisco UCS Manager restarts the fabric interconnect, logs you out, and disconnects Cisco UCS Manager GUI.