



Network Management

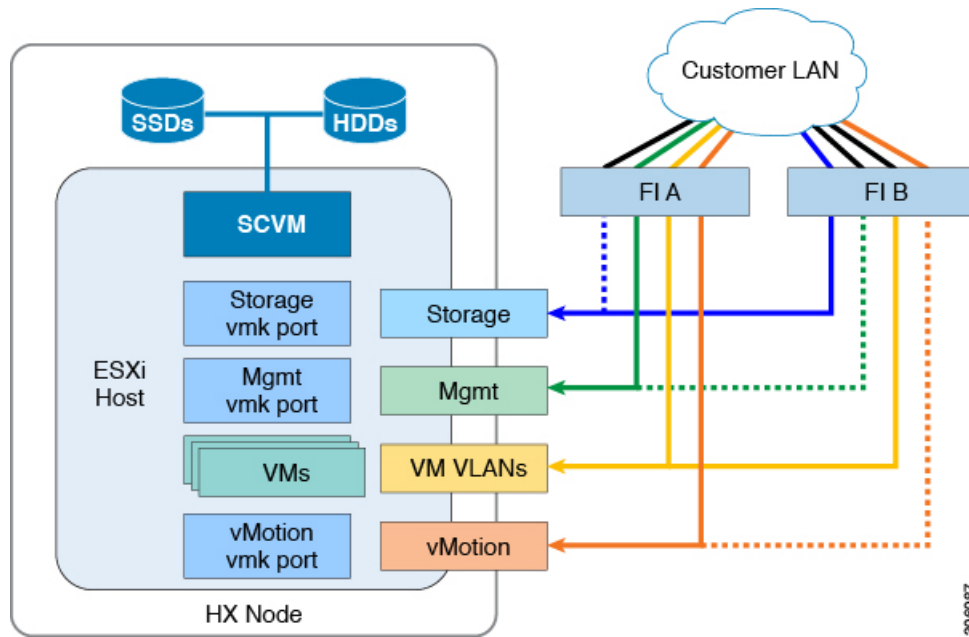
- [Physical Network, on page 1](#)
- [VLANs and Subnets, on page 2](#)
- [Jumbo Frames, on page 3](#)
- [Logical Network, on page 3](#)
- [Virtual Network, on page 4](#)

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 1: 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:

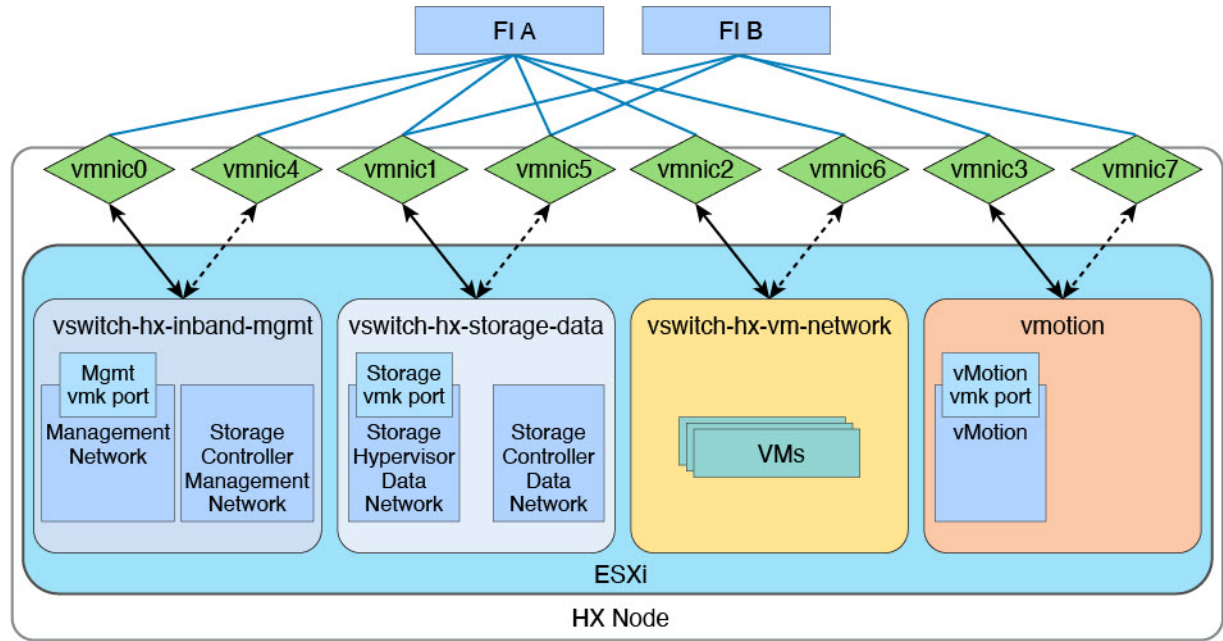
Table 2: Defined Communication Pathway Zones

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



306088

Table 3: 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.

