



Cisco HyperFlex Edge Deployment

- [Cisco HyperFlex Edge, on page 1](#)
- [HyperFlex Edge Deployment Options, on page 4](#)
- [Select your 2-Node Network Topology, on page 5](#)
- [Selecting your 3- or 4-Node Network Topology, on page 25](#)
- [Selecting your 2-Node 2-Room Network Topology, on page 34](#)
- [Installation Overview, on page 54](#)
- [\(1GE Only\) Run Post Installation Script , on page 69](#)
- [\(10/25GE Only\) Run Post Installation Script , on page 71](#)
- [Configuring vMotion Automatically, on page 73](#)
- [Configuring vMotion Manually, on page 74](#)
- [Configuring Traffic Shaping Manually, on page 74](#)
- [\(10/25GE\) Using Additional VIC Ports \(Optional\), on page 75](#)
- [Sample Network Configurations, on page 75](#)

Cisco HyperFlex Edge

Introduction

Cisco HyperFlex Edge brings the simplicity of hyperconvergence to remote and branch office (ROBO) and edge environments. This document describes the deployment for HyperFlex Edge.



Important HX Edge clusters deployment with M6 nodes require Intersight for initial deployment and ongoing management.

Limitations and Supportability Summary

Limitation for	Support
Cluster size and type	

Limitation for	Support
	<p>2-Node clusters:</p> <ul style="list-style-type: none"> • HX220c M6 Hybrid/HXAF220c M6 All-Flash • HX225c M6 Hybrid/HXAF225c M6 All-Flash • HX220c M5 Hybrid/HXAF220c M5 All-Flash • HX240c M6 Hybrid/HXAF240c M6 All-Flash • HX245c M6 Hybrid/HXAF245c M6 All-Flash • HX240c M5 Hybrid/HXAF240c M5 All-Flash • HX240c M5SD Hybrid/HX240c M5SD All-Flash <p>Note 2-Node clusters require Intersight for initial deployment and ongoing management.</p> <p>3-Node clusters:</p> <ul style="list-style-type: none"> • HX220c M6 Hybrid/HXAF220c M6 All-Flash • HX225c M6 Hybrid/HXAF225c M6 All-Flash • HX220c M5 Hybrid/HXAF220c M5 All-Flash • HX220c M4 Hybrid/HXAF220c M4 All-Flash • HX240c M6 Hybrid/HXAF240c M6 All-Flash • HX245c M6 Hybrid/HXAF245c M6 All-Flash • HX240c M5 Hybrid/HXAF240c M5 All-Flash • HX240c M5SD Hybrid/HX240c M5SD All-Flash <p>Note HX Edge clusters deployment with M6 or HX 240 Edge (short depth and full depth) nodes require Intersight for initial deployment and ongoing management.</p> <p>4-Node clusters:</p> <ul style="list-style-type: none"> • HX220c M6 Hybrid/HXAF220c M6 All-Flash • HX225c M6 Hybrid/HXAF225c M6 All-Flash • HX220c M5 Hybrid/HXAF220c M5 All-Flash • HX240c M6 Hybrid/HXAF240c M6 All-Flash • HX245c M6 Hybrid/HXAF245c M6 All-Flash • HX240c M5 Hybrid/HXAF240c M5 All-Flash • HX240c M5SD Hybrid/HX240c M5SD All-Flash <p>Note HX Edge clusters deployment with M6 or HX 240 (short depth and</p>

Limitation for	Support
	full depth) nodes require Intersight for initial deployment and ongoing management.
Replication Factor	Replication Factor recommendations: <ul style="list-style-type: none"> • 3- or 4-Node edge clusters: 3 • 2-Node edge clusters: 2 <p>Note If RF2 is selected, a reliable backup strategy is strongly recommended to ensure that production data is adequately protected.</p>
Networking	1GE or 10/25GE networking without Cisco UCS Fabric Interconnects. HX Edge Systems do not implement QoS.
HX clusters per vCenter	Up to 100.

HyperFlex Edge Deployment Options

HyperFlex Edge can be deployed using Cisco Intersight from the cloud or by using the on-premises installer appliance. You can choose between the following two options depending on your requirements:

- **HyperFlex On-Premises OVA Installer**—Use this option for on-premises Edge deployments for three and four node clusters. This type of deployment supports all three network topologies, and requires download and installation of the appliance along with local network access.



Note Use of the on-premises installer is not supported for two node HyperFlex Edge clusters.

- **Intersight Installer**—Use this option for Edge to deploy HyperFlex Edge from the cloud. This deployment option supports all Edge cluster sizes and network topologies.



Note Deployment of 2-node edge clusters and NIC-Based clusters is supported from Intersight only.

This guide covers deployment using the on-premises OVA installer only.

To deploy an HyperFlex Edge cluster using Cisco Intersight, see [Cisco HyperFlex Systems Installation Guide for Cisco Intersight](#) for detailed deployment instructions. The Cisco Intersight HX installer rapidly deploys HyperFlex Edge clusters. The installer constructs a pre-configuration definition of your cluster, called an HX Cluster Profile. This definition is a logical representation of the HX nodes in your HX Edge cluster. Each HX node provisioned in Cisco Intersight is specified in a HX Cluster profile.

Additional guest VM VLANs are optional. You may use the same management VLAN above for guest VM traffic in environments that wish to keep a simplified flat network design.



Note Each cluster should use a unique storage data VLAN to keep all storage traffic isolated. Reuse of this VLAN across multiple clusters is highly discouraged.



Note Due to the nature of the Cisco VIC carving up multiple vNICs from the same physical port, it is not possible for guest VM traffic configured on vswitch-hx-vm-network to communicate L2 to interfaces or services running on the same host. It is recommended to either a) use a separate VLAN and perform L3 routing or b) ensure any guest VMs that need access to management interfaces be placed on the vswitch-hx-inband-mgmt vSwitch. In general, guest VMs should not be put on any of the HyperFlex configured vSwitches except for the vm-network vSwitch. An example use case would be if you need to run vCenter on one of the nodes and it requires connectivity to manage the ESXi host it is running on. In this case, use one of the recommendations above to ensure uninterrupted connectivity.

Select your 2-Node Network Topology

When selecting your 2-Node topology, keep in mind that the network topology chosen during initial deployment cannot be changed or upgraded without full reinstallation. Choose your network topology carefully with future needs in mind and take into account the following Cisco HyperFlex offerings:

- 10/25Gigabit (GE) topology with Cisco VIC-based hardware or Intel NIC-Based adapters.
- 1GE topology, for clusters that will not need node expansion and where the top-of-rack ToR switch does not have 10GE ports available.

For more specific information on Cisco IMC Connectivity, physical cabling, network design, and configuration guidelines, select from the following list of available topologies:

- [10 or 25GE VIC-Based Topology, on page 5](#)
- [1 Gigabit Ethernet Topology, on page 19](#)

After completing the 10/25GE or 1GE ToR physical network and cabling section, continue with the [Common Network Requirement Checklist, on page 47](#).

10 or 25 Gigabit Ethernet Topology

10 or 25GE VIC-Based Topology

The 10 or 25 Gigabit Ethernet (GE) switch topology provides a fully redundant design that protects against switch (if using dual or stacked switches), link and port failures. The 10/25GE switch may be one or two standalone switches or may be formed as a switch stack.

Cisco IMC Connectivity for 10/25GE VIC-Based Topology

Choose one of the following Cisco IMC Connectivity options for the 2-node 10/25 Gigabit Ethernet (GE) topology:

- Use of a dedicated 1GE Cisco IMC management port is recommended. This option requires additional switch ports and cables, however it avoids network contention and ensures always on, out of band access to each physical server.
- Use of shared LOM extended mode (EXT). In this mode, single wire management is used and Cisco IMC traffic is multiplexed onto the 10/25GE VIC connections. When operating in this mode, multiple streams of traffic are shared on the same physical link and uninterrupted reachability is not guaranteed. This deployment option is not recommended.
 - In fabric interconnect-based environments, built in QoS ensures uninterrupted access to Cisco IMC and server management when using single wire management. In HyperFlex Edge environments, QoS is not enforced and hence the use of a dedicated management port is recommended.
- Assign an IPv4 management address to the Cisco IMC following the procedures in the [Server Installation and Service Guide](#) for the equivalent Cisco UCS C-series server. HyperFlex does not support IPv6 addresses.

Physical Network and Cabling for 10/25GE VIC-Based Topology

A managed switch (1 or 2) with VLAN capability is required. Cisco fully tests and provides reference configurations for Catalyst and Nexus switching platforms. Choosing one of these switches provides the highest level of compatibility and ensures a smooth deployment and seamless ongoing operations.

Dual switch configuration provides a slightly more complex topology with full redundancy that protects against: switch failure, link failure, and port failure. It requires two switches that may be standalone or stacked, and two 10/25GE ports, one 1GE port for CIMC management, and one Cisco VIC 1457 per server. Trunk ports are the only supported network port configuration.

Single switch configuration provides a simple topology requiring only a single switch, and two 10/25GE ports, one 1GE port for CIMC management, and one Cisco VIC 1457 per server. Switch level redundancy is not provided, however all links/ports and associated network services are fully redundant and can tolerate failures.

Requirements for both 10 and 25GE Topologies

The following requirements are common to both 10/25GE topologies and must be met before starting deployment:

- Dedicated 1 Gigabit Ethernet (GE) Cisco IMC management port per server (recommended)
 - 2 x 1GE ToR switch ports and two (2) Category 6 ethernet cables for dedicated Cisco IMC management port (customer supplied)
- Cisco VIC 1457 (installed in the MLOM slot in each server)
 - Prior generation Cisco VIC hardware is not supported for 2 node or 4 node HX Edge clusters.
 - 4 x 10/25GE ToR switch ports and 4 x 10/25GE SFP+ or SFP28 cables (customer supplied. Ensure the cables you select are compatible with your switch model).
 - Cisco VIC 1457 supports 10GE interface speeds in Cisco HyperFlex Release 4.0(1a) and later.
 - Cisco VIC 1457 supports 25GE interface speeds in Cisco HyperFlex Release 4.0(2a) and late

- Cisco VIC 1457 does not support 40GE internet speeds.

Requirements for HX Edge clusters using 25GE



Note Using 25GE mode typically requires the use of forward error correction (FEC) depending on the transceiver or the type & length of cabling selected. The VIC 1400 series by default is configured in CL91 FEC mode (FEC mode “auto” if available in the Cisco IMC UI is the same as CL91) and does not support auto FEC negotiation. Certain switches will need to be manually set to match this FEC mode to bring the link state up. The FEC mode must match on both the switch and VIC port for the link to come up. If the switch in use does not support CL91, you may configure the VIC ports to use CL74 to match the FEC mode available on the switch. This will require a manual FEC mode change in the CIMC UI under the VIC configuration tab. Do not start a HyperFlex Edge deployment until the link state is up as reported by the switch and the VIC ports. CL74 is also known as FC-FEC (Firecode) and CL91 is also known as RS-FEC (Reed Solomon). See the [Cisco UCS C-Series Integrated Management Controller GUI Configuration Guide, Release 4.1](#) for further details on how to change the FEC mode configured on the VIC using the Cisco IMC GUI.

Next Step:

Select either a single switch or dual switch configuration to continue with physical cabling:

10/25GE VIC-Based Dual Switch Physical Cabling



Warning Proper cabling is important to ensure full network redundancy.

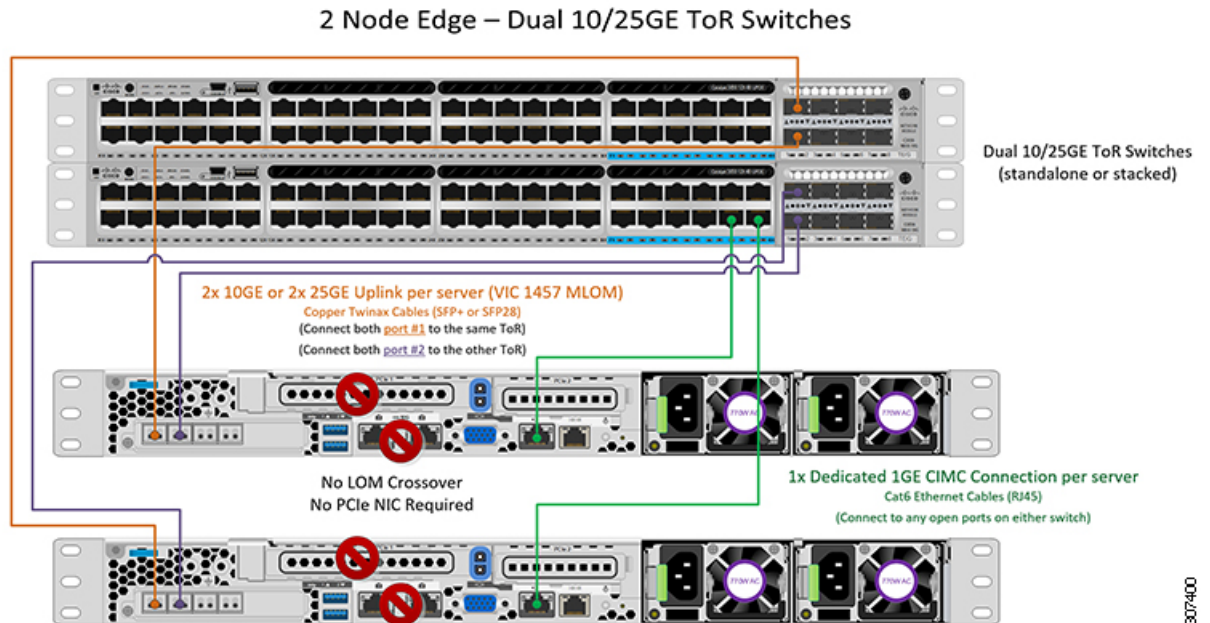
To deploy with dual ToR switches for extra redundancy (see diagram below for a visual layout):

- If using dedicated Cisco IMC, connect the 1GE management port on each server (Labeled M on the back of the server) to one of the two switches.
- Connect one out of the four 10/25GE ports on the Cisco VIC from each server to the same ToR switch.
 - Use the same port number on each server to connect to the same switch.



Note Failure to use the same VIC port numbers will result in an extra hop for traffic between servers and will unnecessarily consume bandwidth between the two switches.

- Connect a second 10/25GE port on the Cisco VIC from each server to the other ToR switch. Use the same port number on each server to connect to the same switch.
- Do not connect additional 10/25GE ports prior to cluster installation. After cluster deployment, you may optionally use the additional two 10/25GE ports for guest VM traffic.



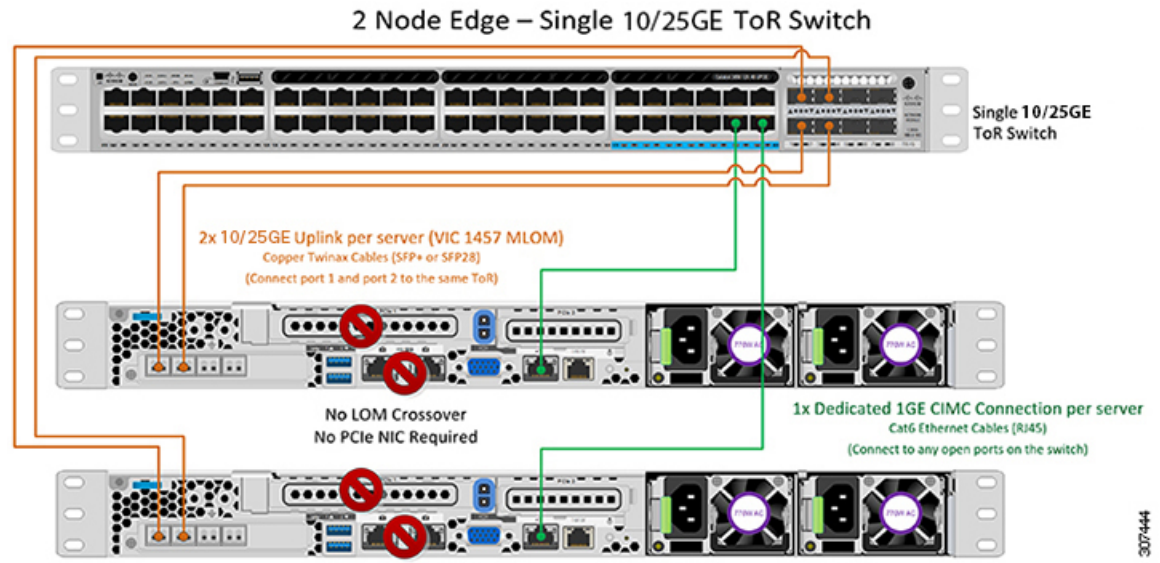
10/25GE VIC-Based Single Switch Physical Cabling



Warning Proper cabling is important to ensure full network redundancy.

To deploy with a single ToR (see diagram below for a visual layout):

- If using dedicated Cisco IMC, connect the 1GE management port on each server (Labeled M on the back of the server) to the switch.
- Connect any two out of the four 10/25GE ports on the Cisco VIC from each server to the same ToR switch.
- Do not connect additional 10/25GE ports prior to cluster installation. After cluster deployment, you may optionally use the additional two 10/25GE ports for guest VM traffic.



Virtual Networking Design for 2-Node 10/25GE VIC-Based Topology


This section details the virtual network setup. No action is required as all of the virtual networking is set up automatically by the HyperFlex deployment process. These extra details are included below for informational and troubleshooting purposes.

Virtual Switches:

Four vSwitches are required:

- **vswitch-hx-inband-mgmt**—ESXi management (vmk0), storage controller management network
- **vswitch-hx-storage-data**—ESXi storage interface (vmk1), HX storage controller data network
- **vmotion**—vMotion interface (vmk2)
- **vswitch-hx-vm-network**—VM guest portgroups

Network Topology



Failover Order:

- **vswitch-hx-inband-mgmt**—entire vSwitch is set for active/standby. All services by default consume a single uplink port and failover when needed.
- **vswitch-hx-storage-data**—HyperFlex storage data network and vmk1 are with the opposite failover order as inband-mgmt and vmotion vSwitches to ensure traffic is load balanced.
- **vmotion**—The vMotion VMkernel port (vmk2) is configured when using the post_install script. Failover order is set for active/standby.
- **vswitch-hx-vm-network**—vSwitch is set for active/active. Individual portgroups can be overridden as needed.

10/25GE VIC-based Switch Configuration Guidelines

3 VLANs are required at a minimum.

- 1 VLAN for the following connections: VMware ESXi management, Storage Controller VM management and Cisco IMC management.
 - VMware ESXi management and Storage Controller VM management must be on the same subnet and VLAN.

- A dedicated Cisco IMC management port may share the same VLAN with the management interfaces above or may optionally use a dedicated subnet and VLAN. If using a separate VLAN, it must have L3 connectivity to the management VLAN above and must meet Intersight connectivity requirements.
- If using shared LOM extended mode for Cisco IMC management, a dedicated VLAN is recommended.
- 1 VLAN for Cisco HyperFlex storage traffic. This can and should be an isolated and non-routed VLAN. It must be unique and cannot overlap with the management VLAN.
- 1 VLAN for vMotion traffic. This can be an isolated and non-routed VLAN.



Note It is not possible to collapse or eliminate the need for these VLANs. The installation will fail if attempted.

- Additional VLANs as needed for guest VM traffic. These VLANs will be configured as additional portgroups in ESXi and should be trunked and allowed on all server facing ports on the ToR switch.
 - These additional guest VM VLANs are optional. You may use the same management VLAN above for guest VM traffic in environments that wish to keep a simplified flat network design.



Note Due to the nature of the Cisco VIC carving up multiple vNICs from the same physical port, it is not possible for guest VM traffic configured on vswitch-hx-vm-network to communicate L2 to interfaces or services running on the same host. It is recommended to either a) use a separate VLAN and perform L3 routing or b) ensure any guest VMs that need access to management interfaces be placed on the vswitch-hx-inband-mgmt vSwitch. In general, guest VMs should not be put on any of the HyperFlex configured vSwitches except for the vm-network vSwitch. An example use case would be if you need to run vCenter on one of the nodes and it requires connectivity to manage the ESXi host it is running on. In this case, use one of the recommendations above to ensure uninterrupted connectivity.

- Switchports connected to the Cisco VIC should be configured in trunk mode with the appropriate VLANs allowed to pass.
- Switchports connected to the dedicated Cisco IMC management port should be configured in 'Access Mode' on the appropriate VLAN.
- All cluster traffic will traverse the ToR switches in the 10/25GE topology
- Spanning tree portfast trunk (trunk ports) should be enabled for all network ports



Note Failure to configure portfast may cause intermittent disconnects during ESXi bootup and longer than necessary network re-convergence during physical link failure

Additional Considerations:

- Additional 3rd party NIC cards may be installed in the HX Edge nodes as needed. See the section in chapter 1 with the link to the networking guide.
- All non-VIC interfaces must be shut down or left un-cabled until installation is completed
- Only a single VIC is supported per HX Edge node in the MLOM slot. PCIe based VIC adapters are not supported with HX Edge nodes.

Jumbo Frames for 10/25 GE VIC-Based

Jumbo frames are typically used to reduce the number of packets transmitted on your network and increase efficiency. The following describes the guidelines to using jumbo frames on your 10/25GE topology.

- The option to enable jumbo frames is only provided during initial install and cannot be changed later.
- Jumbo Frames are not required. If opting out of jumbo frames, leave the MTU set to 1500 bytes on all network switches.
- For highest performance, jumbo frames may be optionally enabled. Ensure full path MTU is 9000 bytes or greater. Keep the following considerations in mind when enabling jumbo frames:
 - When running a dual switch setup, it is imperative that all switch interconnects and switch uplinks have jumbo frames enabled. Failure to ensure full path MTU could result in a cluster outage if traffic is not allowed to pass after link or switch failure.
 - The HyperFlex installer will perform a one-time test on initial deployment that will force the failover order to use the standby link on one of the nodes. If the switches are cabled correctly, this will test the end to end path MTU. Do not bypass this warning if a failure is detected. Correct the issue and retry the installer to ensure the validation check passes.
 - For these reasons and to reduce complexity, it is recommended to disable jumbo frames when using a dual switch setup.
- The option to enable jumbo frames is found in the HyperFlex Cluster profile, under the Network Configuration policy. Checking the box will enable jumbo frames. Leaving the box unchecked will keep jumbo frames disabled.

Next Steps:

Complete the [Common Network Requirement Checklist](#), on page 47.

Common Network Requirement Checklist

Before you begin installation, confirm that your environment meets the following specific software and hardware requirements.

VLAN Requirements



Important **Reserved VLAN IDs** - The VLAN IDs you specify must be supported in the Top of Rack (ToR) switch where the HyperFlex nodes are connected. For example, VLAN IDs 3968 to 4095 are reserved by Nexus switches and VLAN IDs 1002 to 1005 are reserved by Catalyst switches. Before you decide the VLAN IDs for HyperFlex use, make sure that the same VLAN IDs are available on your switch.

Network	VLAN ID	Description
Use a separate subnet and VLANs for each of the following networks:		
VLAN for VMware ESXi, and Cisco HyperFlex management		Used for management traffic among ESXi, HyperFlex, and VMware vCenter, and must be routable. Note This VLAN must have access to Intersight (if deploying with Intersight).
CIMC VLAN		Can be same or different from the Management VLAN. Note This VLAN must have access to Intersight (if deploying with Intersight).
VLAN for HX storage traffic		Used for storage traffic and requires only L2 connectivity.
VLAN for VMware vMotion		Used for vMotion VLAN, if applicable. Note Can be the same as the management VLAN but not recommended.
VLAN(s) for VM network(s)		Used for VM/application network. Note Can be multiple VLANs separated by a VM portgroup in ESXi.

Supported vCenter Topologies

Use the following table to determine the topology supported for vCenter.

Topology	Description	Recommendation
Single vCenter	Virtual or physical vCenter that runs on an external server and is local to the site. A management rack mount server can be used for this purpose.	Highly recommended

Topology	Description	Recommendation
Centralized vCenter	vCenter that manages multiple sites across a WAN.	Highly recommended
Nested vCenter	vCenter that runs within the cluster you plan to deploy.	Installation for a HyperFlex Edge cluster may be initially performed without a vCenter. Alternatively, you may deploy with an external vCenter and migrate it into the cluster. In either case, the cluster must be registered to a vCenter server before running production workloads. For the latest information, see the How to Deploy vCenter on the HX Data Platform tech note.

3-Node Customer Deployment Information

A typical three-node HyperFlex Edge deployment requires 13 IP addresses – 10 IP addresses for the management network and 3 IP addresses for the vMotion network.



Important All IP addresses must be IPv4. HyperFlex does not support IPv6 addresses.

4-Node Customer Deployment Information

A typical four-node HyperFlex Edge deployment requires 17 IP addresses – 13 IP addresses for the management network and 4 IP addresses for the vMotion network.



Important All IP addresses must be IPv4. HyperFlex does not support IPv6 addresses.

CIMC Management IP Addresses

Server	CIMC Management IP Addresses
Server 1:	
Server 2:	
Server 3:	
Server 4:	
Subnet mask	
Gateway	

Server	CIMC Management IP Addresses
DNS Server	
NTP Server	
Note NTP configuration on CIMC is required for proper Intersight connectivity.	

Network IP Addresses



Note By default, the HX Installer automatically assigns IP addresses in the 169.254.1.X range, to the Hypervisor Data Network and the Storage Controller Data Network. This IP subnet is not user configurable.



Note Spanning Tree portfast trunk (trunk ports) should be enabled for all network ports.
Failure to configure portfast may cause intermittent disconnects during ESXi bootup and longer than necessary network re-convergence during physical link failure.

Management Network IP Addresses (must be routable)	
Hypervisor Management Network	Storage Controller Management Network
Server 1:	Server 1:
Server 2:	Server 2:
Server 3:	Server 3:
Server 4:	Server 4:
Storage Cluster Management IP address	Cluster IP:
Subnet mask	
Default gateway	

VMware vMotion Network IP Addresses

For vMotion services, you may configure a unique VMkernel port or, if necessary, reuse the vmk0 if you are using the management VLAN for vMotion (not recommended).

Server	vMotion Network IP Addresses (configured using the post_install script)
Server 1:	
Server 2:	

Server	vMotion Network IP Addresses (configured using the post_install script)
Server 3:	
Server 4:	
Subnet mask	
Gateway	

VMware vCenter Configuration



Note HyperFlex communicates with vCenter through standard ports. Port 80 is used for reverse HTTP proxy and may be changed with TAC assistance. Port 443 is used for secure communication to the vCenter SDK and may not be changed.

vCenter admin username <i>username@domain</i>	
vCenter admin password	
vCenter data center name Note An existing datacenter object can be used. If the datacenter doesn't exist in vCenter, it will be created.	
VMware vSphere compute cluster and storage cluster name Note Cluster name you will see in vCenter.	

Port Requirements



Important Ensure that the following port requirements are met in addition to the prerequisites listed for [Intersight Connectivity, on page 18](#).

If your network is behind a firewall, in addition to the standard port requirements, VMware recommends ports for VMware ESXi and VMware vCenter.

- CIP-M is for the cluster management IP.
- SCVM is the management IP for the controller VM.
- ESXi is the management IP for the hypervisor.

The comprehensive list of ports required for component communication for the HyperFlex solution is located in Appendix A of the [HX Data Platform Security Hardening Guide](#)



Tip If you do not have standard configurations and need different port settings, refer to [Table C-5 Port Literal Values](#) for customizing your environment.

Network Services



Note

- DNS and NTP servers should reside outside of the HX storage cluster.
- To ensure your cluster works properly and to avoid any issues when your cluster is deployed through Intersight, create the A and PTR DNS records for the SCVMs hostnames.
- Use an internally-hosted NTP server to provide a reliable source for the time.
- All DNS servers should be pre-configured with forward (A) and reverse (PTR) DNS records for each ESXi host before starting deployment. When DNS is configured correctly in advance, the ESXi hosts are added to vCenter via FQDN rather than IP address.

Skipping this step will result in the hosts being added to the vCenter inventory via IP address and require users to change to FQDN using the following procedure: [Changing Node Identification Form in vCenter Cluster from IP to FQDN](#).

DNS Servers <i><Primary DNS Server IP address, Secondary DNS Server IP address, ...></i>	
NTP servers <i><Primary NTP Server IP address, Secondary NTP Server IP address, ...></i>	
Time zone <i>Example: US/Eastern, US/Pacific</i>	

Connected Services

Enable Connected Services (Recommended) <i>Yes or No required</i>	
Email for service request notifications <i>Example: name@company.com</i>	

Proxy Server

- Use of a proxy server is optional if direct connectivity to Intersight is not available.
- When using a proxy, the device connectors in each server must be configured to use the proxy in order to claim the servers into an Intersight account. In addition, the proxy information must be provided in the HX Cluster Profile to ensure the HyperFlex Data Platform can be successfully downloaded.
- Use of username/password is optional

Proxy required: Yes or No	
Proxy Host	
Proxy Port	
Username	
Password	

Guest VM Traffic

Considerations for guest VM traffic are given above based on the topology selection. In general, guest port groups may be created as needed so long as they are applied to the correct vSwitch:

- 10/25GE Topology: use **vswitch-hx-vm-network** to create new VM port groups.

Cisco recommends you run the `post_install` script to add more VLANs automatically to the correct vSwitches on all hosts in the cluster. Execute **hx_post_install --vlan** (space and two dashes) to add new guest VLANs to the cluster at any point in the future.

Additional vSwitches may be created that use leftover vmnics or third party network adapters. Care should be taken to ensure no changes are made to the vSwitches defined by HyperFlex.



Note Additional user created vSwitches are the sole responsibility of the administrator, and are not managed by HyperFlex.

Intersight Connectivity

Consider the following prerequisites pertaining to Intersight connectivity:

- Before installing the HX cluster on a set of HX servers, make sure that the device connector on the corresponding Cisco IMC instance is properly configured to connect to Cisco Intersight and claimed.
- Communication between CIMC and vCenter via ports 80, 443 and 8089 during installation phase.
- All device connectors must properly resolve `svc.intersight.com` and allow outbound initiated HTTPS connections on port 443. The current version of the HX Installer supports the use of an HTTP proxy.
- All controller VM management interfaces must properly resolve `svc.intersight.com` and allow outbound initiated HTTPS connections on port 443. The current version of HX Installer supports the use of an HTTP proxy if direct Internet connectivity is unavailable.

- IP connectivity (L2 or L3) is required from the CIMC management IP on each server to all of the following: ESXi management interfaces, HyperFlex controller VM management interfaces, and vCenter server. Any firewalls in this path should be configured to allow the necessary ports as outlined in the [Hyperflex Hardening Guide](#).
- Starting with HXDP release 3.5(2a), the Intersight installer does not require a factory installed controller VM to be present on the HyperFlex servers.

When redeploying HyperFlex on the same servers, new controller VMs must be downloaded from Intersight into all ESXi hosts. This requires each ESXi host to be able to resolve `svc.intersight.com` and allow outbound initiated HTTPS connections on port 443. Use of a proxy server for controller VM downloads is supported and can be configured in the HyperFlex Cluster Profile if desired.

- Post-cluster deployment, the new HX cluster is automatically claimed in Intersight for ongoing management.

Cisco HyperFlex Edge Invisible Cloud Witness

The Cisco HyperFlex Edge Invisible Cloud Witness is an innovative technology for Cisco HyperFlex Edge Deployments that eliminates the need for witness VMs or arbitration software.

The Cisco HyperFlex Edge invisible cloud witness is only required for 2-node HX Edge deployments. The witness does not require any additional infrastructure, setup, configuration, backup, patching, or management of any kind. This feature is automatically configured as part of a 2-node HyperFlex Edge installation. Outbound access at the remote site must be present for connectivity to Intersight (either Intersight.com or to the Intersight Virtual Appliance). HyperFlex Edge 2-node clusters cannot operate without this connectivity in place.

For additional information about the benefits, operations, and failure scenarios of the Invisible Cloud Witness feature, see <https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/whitepaper-c11-741999.pdf>

Ordering Cisco HyperFlex Edge Servers

When ordering Cisco HyperFlex Edge servers, be sure to choose the correct components as outlined in the [HyperFlex Edge spec sheets](#). Pay attention to the network topology selection to ensure it matches your desired configuration. Further details on network topology PID selection can be found in the supplemental material section of the spec sheet.

1 Gigabit Ethernet Topology

1 Gigabit Ethernet Topology

The 1 Gigabit Ethernet (GE) switch topology provides a fully redundant design that protects against switch (if using dual or stacked switches), link and port failures. The 1GE switch may be one or two standalone switches or may be formed as a switch stack.



Note Single or dual 1Gb switch connectivity limits the maximum performance that virtual machines can get and is not recommended for applications requiring high performance.

Cisco IMC Connectivity for 1 Gigabit Ethernet Topology

Cisco IMC Connectivity for your 2-node 1 Gigabit Ethernet (GE) topology requires the use of the dedicated 1GE Cisco IMC management port. Other operating modes, including shared LOM mode, are not available due to the use of direct connect cables in this topology.

Assign an IPv4 management address to the Cisco IMC following the procedures in the [Server Installation and Service Guide](#) for the equivalent Cisco UCS C-series server. HyperFlex does not support IPv6 addresses.

Physical Network and Cabling for 1 GE Topology

A managed switch (1 or 2) with VLAN capability is required. Cisco fully tests and provides reference configurations for Catalyst and Nexus switching platforms. Choosing one of these switches provides the highest level of compatibility and ensures a smooth deployment and seamless ongoing operations.

Dual switch cabling provides a slightly more complex topology with full redundancy that protects against: switch failure, link failure, and switch port failure. It requires two switches that may be standalone or stacked, and three 1 Gigabit Ethernet (GE) switch ports per server. Single switch cabling provides a simple topology requiring only single switch and three 1GE switch ports per server. Switch level redundancy is not provided, however all links/ports and associated network services are fully redundant and can tolerate failures.

The 1GE topology uses direct-connect cables for high speed, redundant, 10GE connectivity between the two nodes without the need for a 10GE capable switch.



Note This topology does not support future node expansion capability and should be avoided where requirements may dictate adding more HX Edge nodes in the future.

The following requirements are common to both 1GE topologies and must be met before starting deployment:

- Dedicated 1 Gigabit Ethernet (GE) Cisco IMC management port per server (required)
- Intel i350 Quad Port PCIe NIC Card (installed in a PCIe slot in each server) (required)
 - Cisco VIC is not used in this topology
- 2 x 10GE DirectConnect LAN-on-Motherboard (LOM) connections (do not consume switchports)
 - 2 x Category 6 straight through ethernet cables for direct connect links (customer supplied)
- 6 x 1GE Top of Rack (ToR) switchports and 6x Category 6 ethernet cables (customer supplied)

Select either a single switch or dual switch configuration to continue with physical cabling:

1 Gigabit Ethernet Dual Switch Cabling



Warning Proper cabling is important to ensure full network redundancy.

To deploy with dual ToR switches for extra redundancy (see diagram below for a visual layout):

- Connect the 1GE dedicated Cisco IMC management port on each server (Labeled M on the back of the server) to one of the two switches.


- Connect the Lan-on-motherboard (LOM) port 1 on one server to the LOM port 1 on the other server using a regular ethernet cable.
- Connect LOM port 2 on one server to LOM port 2 on the second server.
- Connect one out of the four 1GE ports on the i350 NIC from each server to the same ToR switch. Use the same port number on each server to connect to the same switch.



Note Failure to use the same port numbers will result in an extra hop for traffic between servers and will unnecessarily consume bandwidth between the two switches.

- Connect a second 1GE port on the i350 NIC from each server to the other ToR switch. Use the same port number on each server to connect to the same switch.
- Do not connect additional 1GE ports prior to cluster installation. After cluster deployment, you may optionally use the additional two 1GE ports for guest VM traffic.

2 Node Edge – Dual 1GE ToR Switches



1 Gigabit Ethernet Single Switch Cabling

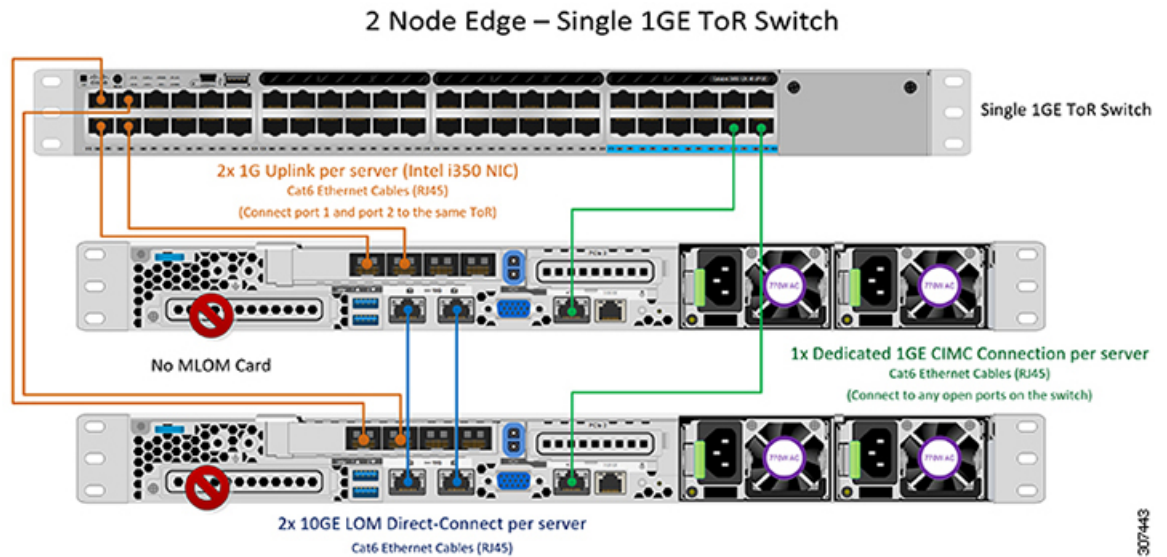


Warning Proper cabling is important to ensure full network redundancy.

To deploy with a single ToR (see diagram below for a visual layout):

- Connect the 1GE dedicated Cisco IMC management port on each server (Labeled M on the back of the server) to the switch.

- Connect the Lan-on-motherboard (LOM) port 1 on one server to the LOM port 1 on the other server using a regular ethernet cable.
- Connect LOM port 2 on one server to LOM port 2 on the second server.
- Connect any two out of the four 1GE ports on the i350 NIC from each server to the same ToR switch.
- Do not connect additional 1GE ports prior to cluster installation. After cluster deployment, you may optionally use the additional two 1GE ports for guest VM traffic.



Virtual Networking Design for 2-Node 1 Gigabit Ethernet Topology

This section details the virtual network setup. No action is required as all of the virtual networking is set up automatically by the HyperFlex deployment process. These extra details are included below for informational and troubleshooting purposes.

Virtual Switches:

The recommended configuration for each ESXi calls for the following networks to be separated:

- Management traffic network
- Data traffic network
- vMotion network
- VM network

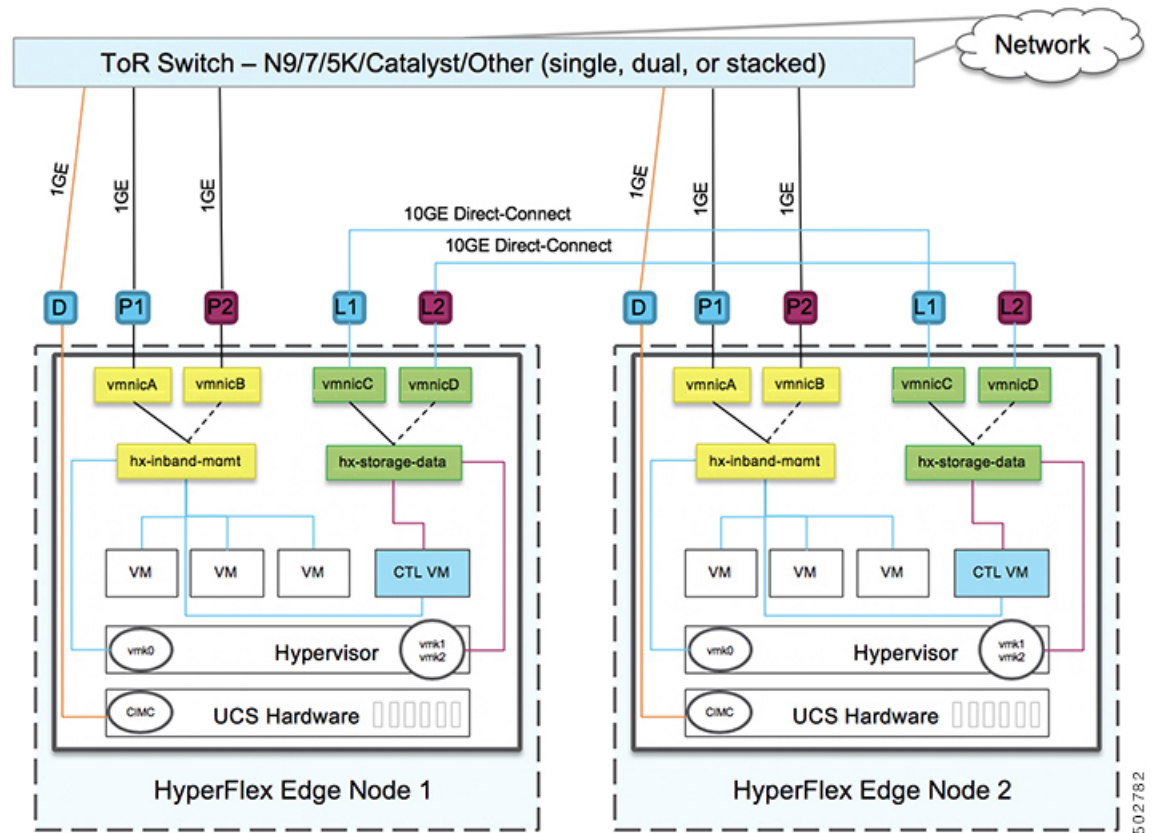
The minimum network configuration requires at least two separate networks:

- Management network (includes vMotion and VM network).
- Data network (for storage traffic)

Two vSwitches each carrying different networks are required:

- **vswitch-hx-inband-mgmt**—ESXi management (vmk0), HyperFlex storage controller management network, VM guest portgroups.
- **vswitch-hx-storage-data**—ESXi storage interface (vmk1), HyperFlex interface (vmk2), HyperFlex storage controller data network.

Network Topology



Failover Order:

vswitch-hx-inband-mgmt—entire vSwitch is set for active/standby across the two uplinks. All services by default consume a single uplink port and failover when needed. Failover order for guest VM portgroups may be overridden as needed and to achieve better load balancing.

vswitch-hx-storage-data—HyperFlex storage data network and vmk1 are set to the same active/standby order. The vMotion Vmkernel port is set to use the opposite order when configured using the post_install script. This ensures full utilization of the direct connect links.

1 Gigabit Ethernet Switch Configuration Guidelines

- 1 VLAN minimum for the following connections: VMware ESXi management, Storage Controller VM Management and Cisco IMC Management.
 - VMware ESXi management and Storage Controller VM management must be on the same subnet & VLAN

- The dedicated Cisco IMC management port may share the same VLAN with the management interfaces above or may optionally use a dedicated subnet & VLAN. If using a separate VLAN, it must have L3 connectivity to the management VLAN above and must meet Intersight connectivity requirements.
- Additional VLANs as needed for guest VM traffic. These VLANs will be configured as additional portgroups in ESXi and should be trunked on all connections to the ToR switch.
 - These additional guest VM VLANs are optional. You may use the same management VLAN above for guest VM traffic in environments that wish to keep a simplified flat network design.
- Switchports connected to the Intel i350 should be configured in trunk mode with the appropriate VLANs allowed to pass.
- Switchports connected to the dedicated Cisco IMC management port should be configured in 'Access Mode' on the appropriate VLAN.
- VMware vMotion traffic and Cisco HyperFlex data traffic will traverse the direct connect LOMs and will therefore not utilize the top of rack switch. Hence additional VLANs are not required for these services.
 - Configuration of Jumbo Frames on the ToR switch is not required in this topology due to all traffic remaining local without need to traverse upstream switches. This topology therefore defaults vMotion traffic to use jumbo frames for high performance.
- Spanning tree portfast trunk (trunk ports) should be enabled for all network ports



Note Failure to configure portfast may cause intermittent disconnects during ESXi bootup and longer than necessary network re-convergence during physical link failure

Jumbo Frames for 1 Gigabit Ethernet

Jumbo frames are typically used to reduce the number of packets transferred on your network. The following describes the guidelines to using jumbo frames on your 1GE topology.

- Jumbo Frames are automatically configured on the vMotion network as there is no additional setup required.
- The option to enable jumbo frames is found in the HyperFlex Cluster profile, under the Network Configuration policy. When using the 1GE topology, you may choose to enable jumbo frames by ensuring the check box is enabled before starting deployment.

Next Steps:

Complete the [Common Network Requirement Checklist](#), on page 47.

Selecting your 3- or 4-Node Network Topology

When selecting your 3- or 4-Node topology, keep in mind that the network topology chosen during initial deployment cannot be changed or upgraded without full reinstallation. Choose your network topology carefully with future needs in mind and take into account the following Cisco HyperFlex offerings:

- 10/25Gigabit (GE) topology with Cisco VIC-based hardware or Intel NIC-Based adapters.
- 1GE topology, for clusters that will not need node expansion and where the top-of-rack (ToR) switch does not have 10GE ports available.

For more specific information on Cisco IMC Connectivity, physical cabling, network design, and configuration guidelines, select from the following list of available topologies:

- [10 or 25GE VIC-Based Topology, on page 25](#)
- [1 Gigabit Ethernet Topology, on page 29](#)

After completing the 10/25GE or 1GE ToR physical network and cabling section below, continue with the [Common Network Requirement Checklist, on page 12](#).

10 or 25GE VIC-Based Topology

The 10 or 25 Gigabit Ethernet (GE) switch topology provides a fully redundant design that protects against switch (if using dual or stacked switches), link and port failures. The 10/25GE switch may be one or two standalone switches or may be formed as a switch stack.

Cisco IMC Connectivity for 10/25GE VIC-Based Topology

Choose one of the following Cisco IMC Connectivity options for the 3-Node and 4-Node 10/25 Gigabit Ethernet (GE) topology:

- Use of a dedicated 1GE Cisco IMC management port is recommended. This option requires additional switch ports and cables, however it avoids network contention and ensures always on, out of band access to each physical server.
- Use of shared LOM extended mode (EXT). In this mode, single wire management is used and Cisco IMC traffic is multiplexed onto the 10/25GE VIC connections. When operating in this mode, multiple streams of traffic are shared on the same physical link and uninterrupted reachability is not guaranteed. This deployment option is not recommended.
 - In fabric interconnect-based environments, built in QoS ensures uninterrupted access to Cisco IMC and server management when using single wire management. In HyperFlex Edge environments, QoS is not enforced and hence the use of a dedicated management port is recommended.

Regardless of the Cisco IMC connectivity choice above, you must assign an IPv4 management address to the Cisco IMC following the procedures in the [Server Installation and Service Guide](#) for the equivalent Cisco UCS C-series server. HyperFlex does not support IPv6 addresses.

Physical Network and Cabling for 10/25GE VIC-Based Topology

A managed switch (1 or 2) with VLAN capability is required. Cisco fully tests and provides reference configurations for Catalyst and Nexus switching platforms. Choosing one of these switches provides the highest level of compatibility and ensures a smooth deployment and seamless ongoing operations.

Dual switch configuration provides a slightly more complex topology with full redundancy that protects against: switch failure, link failure, and port failure. It requires two switches that may be standalone or stacked, and two 10/25GE ports, one 1GE port for CIMC management, and one Cisco VIC 1457 per server. Trunk ports are the only supported network port configuration.

Single switch configuration provides a simple topology requiring only a single switch, and two 10/25GE ports, one 1GE port for CIMC management, and one Cisco VIC 1457 per server. Switch level redundancy is not provided, however all links/ports and associated network services are fully redundant and can tolerate failures.

Requirements for both 10 and 25GE Topologies

The following requirements are common to both 10/25GE topologies and must be met before starting deployment:

- Dedicated 1 Gigabit Ethernet (GE) Cisco IMC management port per server (recommended)
 - 1 x 1GE ToR switch ports and one (1) Category 6 ethernet cable for dedicated Cisco IMC management port per HyperFlex node (customer supplied)
- Cisco VIC 1457 (installed in the MLOM slot in each server)
 - Prior generation Cisco VIC hardware is not supported for 2 node or 4 node HX Edge clusters.
 - 2 x 10/25GE ToR switch ports and 2 x 10/25GE SFP+ or SFP28 cables per HyperFlex node (customer supplied. Ensure the cables you select are compatible with your switch model).
 - Cisco VIC 1457 supports 10GE interface speed in Cisco HyperFlex Release 4.0(1a) and later.
 - Cisco VIC 1457 supports 25GE interface speed in Cisco HyperFlex Release 4.0(2a) and later.
 - 40GE interfaces speed is not supported is not supported on the Cisco VIC 1457

Requirements for HX Edge clusters using 25GE



Note Using 25GE mode typically requires the use of forward error correction (FEC) depending on the transceiver or the type & length of cabling selected. The VIC 1400 series by default is configured in CL91 FEC mode (FEC mode “auto” if available in the Cisco IMC UI is the same as CL91) and does not support auto FEC negotiation. Certain switches will need to be manually set to match this FEC mode to bring the link state up. The FEC mode must match on both the switch and VIC port for the link to come up. If the switch in use does not support CL91, you may configure the VIC ports to use CL74 to match the FEC mode available on the switch. This will require a manual FEC mode change in the CIMC UI under the VIC configuration tab. Do not start a HyperFlex Edge deployment until the link state is up as reported by the switch and the VIC ports. CL74 is also known as FC-FEC (Firecode) and CL91 is also known as RS-FEC (Reed Solomon). See the [Cisco UCS C-Series Integrated Management Controller GUI Configuration Guide, Release 4.1](#) for further details on how to change the FEC mode configured on the VIC using the Cisco IMC GUI.

Select either a single switch or dual switch configuration to continue with physical cabling:

10/25GE VIC-Based Switch Configuration Guidelines

3 VLANs are required at a minimum.

- 1 VLAN for the following connections: VMware ESXi management, Storage Controller VM management and Cisco IMC management.
 - VMware ESXi management and Storage Controller VM management must be on the same subnet and VLAN.
 - A dedicated Cisco IMC management port may share the same VLAN with the management interfaces above or may optionally use a dedicated subnet and VLAN. If using a separate VLAN, it must have L3 connectivity to the management VLAN above and must meet Intersight connectivity requirements (if managed by Cisco Intersight).
 - If using shared LOM extended mode for Cisco IMC management, a dedicated VLAN is recommended.
- 1 VLAN for Cisco HyperFlex storage traffic. This can and should be an isolated and non-routed VLAN. It must be unique and cannot overlap with the management VLAN.
- 1 VLAN for vMotion traffic. This can be an isolated and non-routed VLAN.



Note It is not possible to collapse or eliminate the need for these VLANs. The installation will fail if attempted.

- Additional VLANs as needed for guest VM traffic. These VLANs will be configured as additional portgroups in ESXi and should be trunked and allowed on all server facing ports on the ToR switch.
 - These additional guest VM VLANs are optional. You may use the same management VLAN above for guest VM traffic in environments that wish to keep a simplified flat network design.



Note Due to the nature of the Cisco VIC carving up multiple vNICs from the same physical port, it is not possible for guest VM traffic configured on vswitch-hx-vm-network to communicate L2 to interfaces or services running on the same host. It is recommended to either a) use a separate VLAN and perform L3 routing or b) ensure any guest VMs that need access to management interfaces be placed on the vswitch-hx-inband-mgmt vSwitch. In general, guest VMs should not be put on any of the HyperFlex configured vSwitches except for the vm-network vSwitch. An example use case would be if you need to run vCenter on one of the nodes and it requires connectivity to manage the ESXi host it is running on. In this case, use one of the recommendations above to ensure uninterrupted connectivity.

- Switchports connected to the Cisco VIC should be configured in trunk mode with the appropriate VLANs allowed to pass.

- Switchports connected to the dedicated Cisco IMC management port should be configured in ‘Access Mode’ on the appropriate VLAN.
- All cluster traffic will traverse the ToR switches in the 10/25GE topology
- Spanning tree portfast trunk (trunk ports) should be enabled for all network ports



Note Failure to configure portfast may cause intermittent disconnects during ESXi bootup and longer than necessary network re-convergence during physical link failure

Additional Considerations:

- Additional 3rd party NIC cards may be installed in the HX Edge nodes as needed. See the section in chapter 1 with the link to the networking guide.
- All non-VIC interfaces must be shutdown or left un-cabled until install is completed
- Only a single VIC is supported per HX Edge node in the MLOM slot. PCIe based VIC adapters are not supported with HX Edge nodes.


Virtual Networking Design for 3- and 4-Node 10/25GE VIC-Based Topology

This section details the virtual network setup. No action is required as all of the virtual networking is set up automatically by the HyperFlex deployment process. These extra details are included below for informational and troubleshooting purposes.

Virtual Switches

Four vSwitches are required:

- **vswitch-hx-inband-mgmt**—ESXi management (vmk0), storage controller management network
- **vswitch-hx-storage-data**—ESXi storage interface (vmk1), HX storage controller data network
- **vmotion**—vMotion interface (vmk2)
- **vswitch-hx-vm-network**—VM guest portgroups

Network Topology:

5.02781

Failover Order:

- **vswitch-hx-inband-mgmt**—entire vSwitch is set for active/standby. All services by default consume a single uplink port and failover when needed.
- **vswitch-hx-storage-data**—HyperFlex storage data network and vmk1 are with the opposite failover order as inband-mgmt and vmotion vSwitches to ensure traffic is load balanced.
- **vmotion**—The vMotion VMkernel port (vmk2) is configured when using the post_install script. Failover order is set for active/standby.
- **vswitch-hx-vm-network**—vSwitch is set for active/active. Individual portgroups can be overridden as needed.

1 Gigabit Ethernet Topology

The 1 Gigabit Ethernet (GE) switch topology provides two designs depending on requirements. The dual switch design is fully redundant and protects against switch (using dual or stacked switches), link and port failures. The other single switch topology does not provide network redundancy, and is not recommended for production clusters.

Cisco IMC Connectivity for 1 Gigabit Ethernet Topology

Choose one of the following Cisco IMC Connectivity options for the 3-Node and 4-Node 10 Gigabit Ethernet (GE) topology:

- Use of a dedicated 1GE Cisco IMC management port is recommended. This option requires additional switch ports and cables, however it avoids network contention and ensures always on, out of band access to each physical server.
- Use of shared LOM extended mode (EXT). In this mode, single wire management is used and Cisco IMC traffic is multiplexed onto the 1GE LOM connections. When operating in this mode, multiple streams of traffic are shared on the same physical link and uninterrupted reachability is not guaranteed. This deployment option is not recommended.
 - In fabric interconnect-based environments, built in QoS ensures uninterrupted access to Cisco IMC and server management when using single wire management. In HyperFlex Edge environments, QoS is not enforced and hence the use of a dedicated management port is recommended.

Regardless of the Cisco IMC connectivity choice above, you must assign an IPv4 management address to the Cisco IMC following the procedures in the [Server Installation and Service Guide](#) for the equivalent Cisco UCS C-series server. HyperFlex does not support IPv6 addresses.

Physical Network and Cabling for 1GE Topology

A managed switch (1 or 2) with VLAN capability is required. Cisco fully tests and provides reference configurations for Cisco Catalyst and Cisco Nexus switching platforms. Choosing one of these switches provides the highest level of compatibility and ensures a smooth deployment and seamless ongoing operations.

Dual switch cabling provides a slightly more complex topology with full redundancy that protects against: switch failure, link failure, switch port failure, and LOM/PCIe NIC HW failures. It requires two switches that may be standalone or stacked, and four 1GE ports for cluster and VM traffic, one 1GE port for CIMC management, and one Intel i350 PCIe NIC per server. Trunk ports are the only supported network port configuration.

Single switch configuration provides a simple topology requiring only a single switch, two 1GE ports for cluster and VM traffic, one 1GE port for CIMC management, and no additional PCIe NICs. Link or switch redundancy is not provided. Access ports and trunk ports are the two supported network port configurations.



Note The lack of redundancy makes the single switch 1GE configuration only recommended for non-production environments.

Select either a single switch or dual switch configuration to continue with physical cabling:

1 Gigabit Ethernet Switch Configuration Guidelines

- 1 VLAN minimum for the following connections: VMware ESXi management, Storage Controller VM Management and Cisco IMC Management.
 - VMware ESXi management and Storage Controller VM management must be on the same subnet & VLAN

- The dedicated Cisco IMC management port may share the same VLAN with the management interfaces above or may optionally use a dedicated subnet & VLAN. If using a separate VLAN, it must have L3 connectivity to the management VLAN above and must meet Intersight connectivity requirements (if managed by Cisco Intersight).
- 1 VLAN for Cisco HyperFlex storage traffic. This can and should be an isolated and non-routed VLAN. It must be unique and cannot overlap with the management VLAN.



Note It is not possible to collapse or eliminate the need for both a management VLAN and a second data VLAN. The installation will fail if attempted.

- Additional VLANs as needed for guest VM traffic. These VLANs will be configured as additional portgroups in ESXi and should be trunked on all connections to the ToR switch.
 - These additional guest VM VLANs are optional. You may use the same management VLAN above for guest VM traffic in environments that wish to keep a simplified flat network design.
- Switchports connected to the Intel i350 should be configured in trunk mode with the appropriate VLANs allowed to pass.
- Switchports connected to the dedicated Cisco IMC management port should be configured in 'Access Mode' on the appropriate VLAN.
- VMware vMotion traffic will follow one of these two paths:
 - **Dual Switch Topologies** - vMotion will use the opposite failover order as the storage data network and will have a dedicated 1GE path when there are no network failures. Using the post_install script will set up the VMkernel interface on the correct vSwitch with the correct failover settings. A dedicated VLAN is required since a new interface in ESXi is created (vmk2).
 - **Single Switch Topologies** - vMotion will be shared with the management network. Using the post_install script will a new ESX interface (vmk2) with a default traffic shaper to ensure vMotion doesn't fully saturate the link. A dedicated VLAN is required since a new interface is created.

For more information VMware vMotion traffic, see the Post Installation Tasks section of the [Cisco HyperFlex Edge Deployment Guide](#).

- Spanning tree portfast trunk (trunk ports) should be enabled for all network ports



Note Failure to configure portfast may cause intermittent disconnects during ESXi bootup and longer than necessary network re-convergence during physical link failure.

Virtual Networking Design for 3- and 4-Node 1 Gigabit Ethernet Topology

This section details the virtual network setup. No action is required as all of the virtual networking is set up automatically by the HyperFlex deployment process. These extra details are included below for informational and troubleshooting purposes.

Virtual Switches

The recommended configuration for each ESXi host calls for the following networks to be separated:

- Management traffic network
- Data traffic network
- vMotion network
- VM network

The minimum network configuration requires at least two separate networks:

- Management network (includes vMotion and VM network)
- Data network (for storage traffic)

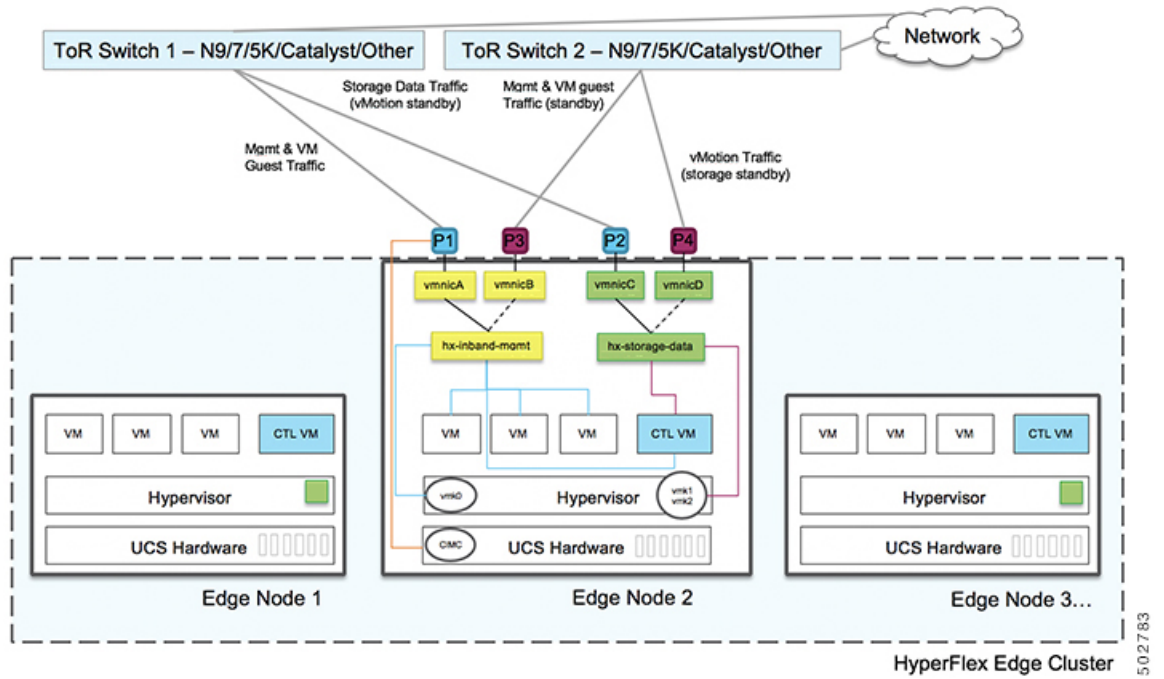
Two vSwitches each carrying different networks are required:

- **vswitch-hx-inband-mgmt**—ESXi management (vmk0), HyperFlex storage controller management network, VM guest portgroups.
- **vswitch-hx-storage-data**—HyperFlex ESXi storage interface (vmk1), HyperFlex storage data network, vMotion (vmk2).

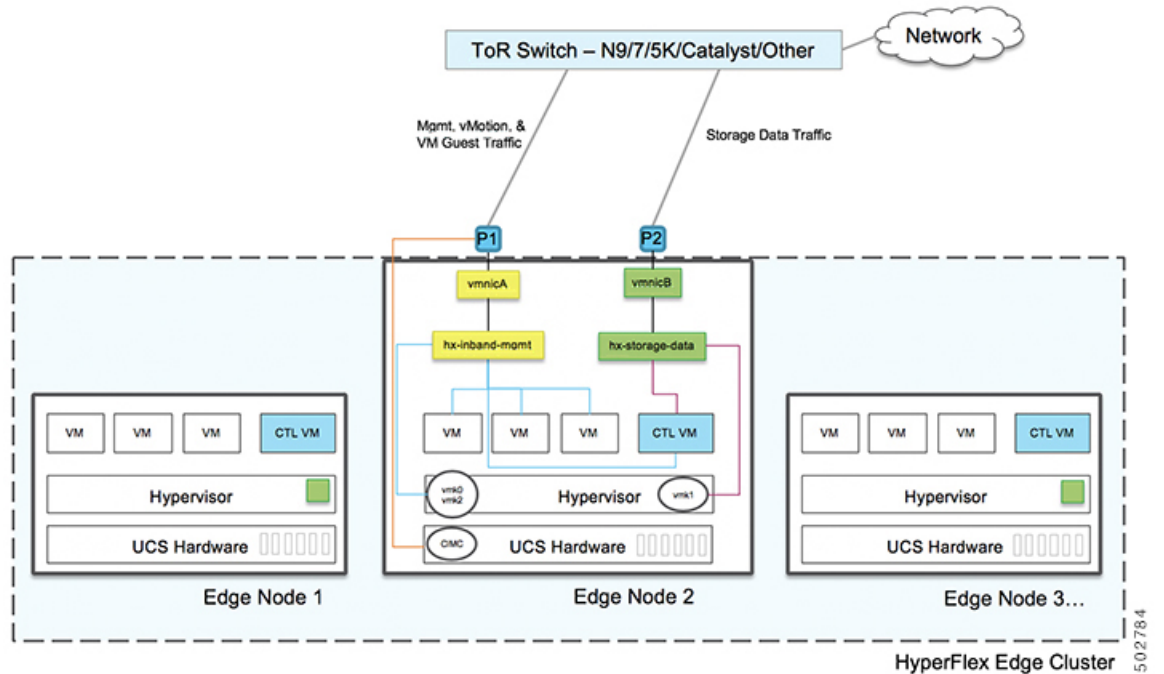


Note After some HyperFlex Edge deployments using the single switch configuration, it is normal to see the storage data vSwitch and associated portgroup failover order with only a standby adapter populated. The missing active adapter does not cause any functional issue with the cluster and we recommend leaving the failover order as configured by the installation process.

Network Topology: Dual Switch Configuration



Network Topology: Single Switch Configuration



Failover Order - Dual switch configuration only:

vswitch-hx-inband-mgmt— entire vSwitch is set for active/standby across the two uplinks. All services by default consume a single uplink port and failover when needed. Failover order for guest VM portgroups may be overridden as needed and to achieve better load balancing.

vswitch-hx-storage-data— HyperFlex storage data network and vmk1 are set to the same active/standby order. The vMotion Vmkernel port is set to use the opposite order when configured using the post_install script. This ensures full utilization of the direct connect links.

Selecting your 2-Node 2-Room Network Topology

To get started, select from one of the available network topologies below. Topologies are listed in priority order based on Cisco's recommendations.

- [10 or 25 Gigabit Ethernet Cross Connect Topology, on page 35](#)
- [10 or 25 Gigabit Ethernet Stacked Switches Per Room Topology, on page 41](#)
- [10 or 25 Gigabit Ethernet Single Switch Per Room Topology, on page 44](#)

After completing the physical network and cabling section, continue with the [Common Network Requirement Checklist, on page 47](#).

2-Node 2-Room Use Case

HyperFlex Edge offers many flexible deployment options depending on workload requirements. Standard topologies are covered in [Select your 2-Node Network Topology, on page 5](#) and [Selecting your 3- or 4-Node Network Topology, on page 25](#) that include single switch, dual switch, 1GE, 10GE, and 25GE options. Some designs call for placing a two-node cluster “stretched” across two rooms within a building or a campus. This type of network topology will further be referred to as a 2-node 2-room design to distinguish this type of topology from a full HyperFlex Stretched Cluster deployment.

This design is sometimes chosen as an attempt to boost the cluster availability and its ability to tolerate certain failure scenarios. Cisco does not currently recommend deploying this type of topology and recommends a properly designed 2-node cluster within the same rack. The following are some reasons why this topology is not considered a Cisco recommended best practice:

- The ability to mitigate power failures can be handled with reliable power and use of an interruptible power supply (UPS)
- Introduces more single points of failure – extra switching infrastructure with inter-switch links that can become oversubscribed and require proper QoS implementation
- Complicates upgrade procedures, requiring careful planning to upgrade all components end to end.
- Does not provide the same level of availability for mission critical applications as a HyperFlex Stretched Cluster (for more information, see the [Cisco HyperFlex Systems Stretched Cluster Guide, Release 4.5](#). HyperFlex Edge is designed to run Edge workloads and does not provide the same performance, data resiliency, and availability guarantees. Deploy a proper stretched cluster when running mission critical applications.
- Requirements for 10GE end to end, maximum 1.5ms RTT, and independent network paths to Intersight or local witness, described in further detail below

- Increases overall complexity to an otherwise simple design

It is possible that a 2-node 2-room topology could unintentionally reduce availability by adding unnecessary complexity to the environment that could be otherwise mitigated through simpler means (e.g., dual redundant switches, redundant power/UPS, etc.).

Despite these best practice recommendations, it is possible and fully supported to deploy HyperFlex Edge using this topology choice. The remainder of this chapter will cover the various requirements and details to deploy such a topology.



Note 2-node 2-room topologies will never be permitted to expand beyond two converged nodes. Expansion to larger clusters is possible for other 10GE+ topologies as outlined in earlier chapters. Do not deploy this topology if cluster expansion may be required in the future.

2-Node 2-Room Requirements

The following requirements must be met when planning a 2-node 2-room deployment.

- Networking speeds must be a minimum of 10/25GE end-to-end. This means all servers must connect to top of rack (ToR) switches using native 10/25GE and all switches must be interconnected by at least one 10GE interface, preferably more.
- Round-Trip-Time (RTT) = the time it takes traffic to go both ways, must not exceed 1.5ms between each server room. Exceeding this threshold will result in substantial reduction in storage cluster performance. Unlike a HyperFlex Stretched Cluster with site affinity for optimized local reads, all reads and writes in a 2-node 2-room design will traverse the inter switch link (ISL) and performance is directly proportional to the network latency. For these reasons, this topology must never be used beyond campus distances (e.g., <1 km).
- Quality of service (QoS) should be implemented at a minimum for the storage data network to prevent other background traffic from saturating the ISL and impacting storage performance. The appendix includes a sample QoS configuration for Catalyst 9300 switches.
- Both rooms must have independent network paths to Intersight (SaaS or Appliance), which serves as the cluster witness. Without independent paths, there is no ability to tolerate the loss of either room. For example, if the Internet connection for room #1 and room #2 is serviced out of room #1, it would be impossible for room #1 to fail and for the Internet in room #2 to remain operational. This strict requirement may disqualify some environments from using a 2-node 2-room design.
- A local witness can also be used with the design. In this case, the same principle applies; both rooms must have independent paths with no dependency on each other to be able to reach the local witness server.
- The HyperFlex Edge 2-node, 2-room topology was introduced and is supported in HyperFlex Data Platform (HXDP) Release 4.5(1a) and later.

10 or 25 Gigabit Ethernet Cross Connect Topology

The cross connect 10 or 25 Gigabit Ethernet (GE) switch topology provides a fully redundant design that protects against room, switch, link and port failures. A single 10/25GE switch is required in each room.

In this topology, each server is cross connected directly to both rooms. This provides dedicated links and prevents oversubscription to the Inter-Switch Link (ISL). This topology still requires a minimum 10GE ISL between each room to handle high bandwidth during server link failure cases.

Physical Network and Cabling for 10/25GE Cross Connect Topology

Each room requires a managed 10GE switch with VLAN capability. Cisco fully tests and provides reference configurations for Catalyst and Nexus switching platforms. Choosing one of these switches provides the highest level of compatibility and ensures a smooth deployment and seamless ongoing operations.

Each room requires: a single switch, and two 10/25GE ports, one 1GE port for CIMC management, and one Cisco VIC 1457 per server. Redundancy is provided at the room level and can tolerate the loss of either room as well as any smaller failure (e.g., switch failure, link failure, port failure).

Requirements for 10/25GE Cross Connect Topology

The following requirements must be met across both rooms before starting deployment:

- Dedicated 1 Gigabit Ethernet (GE) Cisco IMC management port per server (recommended)
- 2 x 1GE ToR switch ports and two (2) Category 6 ethernet cables for dedicated Cisco IMC management port (customer supplied)
- Cisco VIC 1457 (installed in the MLOM slot in each server)
- Prior generation Cisco VIC hardware is not supported for 2 node HX Edge clusters.
- 4 x 10/25GE ToR switch ports and 4 x 10/25GE SFP+ or SFP28 cables (customer supplied. Ensure the cables you select are compatible with your switch model).
- Cisco VIC 1457 supports 10GE or 25GE interface speeds.
- Cisco VIC 1457 does not support 40GE interface speeds.

Requirements for HX Edge Clusters using 25GE



Note Using 25GE mode typically requires the use of forward error correction (FEC) depending on the transceiver or the type & length of cabling selected. The VIC 1400 series by default is configured in CL91 FEC mode (FEC mode “auto” if available in the Cisco IMC UI is the same as CL91) and does not support auto FEC negotiation. Certain switches will need to be manually set to match this FEC mode to bring the link state up. The FEC mode must match on both the switch and VIC port for the link to come up. If the switch in use does not support CL91, you may configure the VIC ports to use CL74 to match the FEC mode available on the switch. This will require a manual FEC mode change in the CIMC UI under the VIC configuration tab. Do not start a HyperFlex Edge deployment until the link state is up as reported by the switch and the VIC ports. CL74 is also known as FC-FEC (Firecode) and CL91 is also known as RS-FEC (Reed Solomon). For more information on how to change the FEC mode configured on the VIC using the Cisco IMC GUI, see the [Cisco UCS C-Series Integrated Management Controller GUI Configuration Guide, Release 4.1](#).

10/25 Gigabit Ethernet Cross Connect Physical Cabling



Warning Proper cabling is important to ensure full network redundancy.

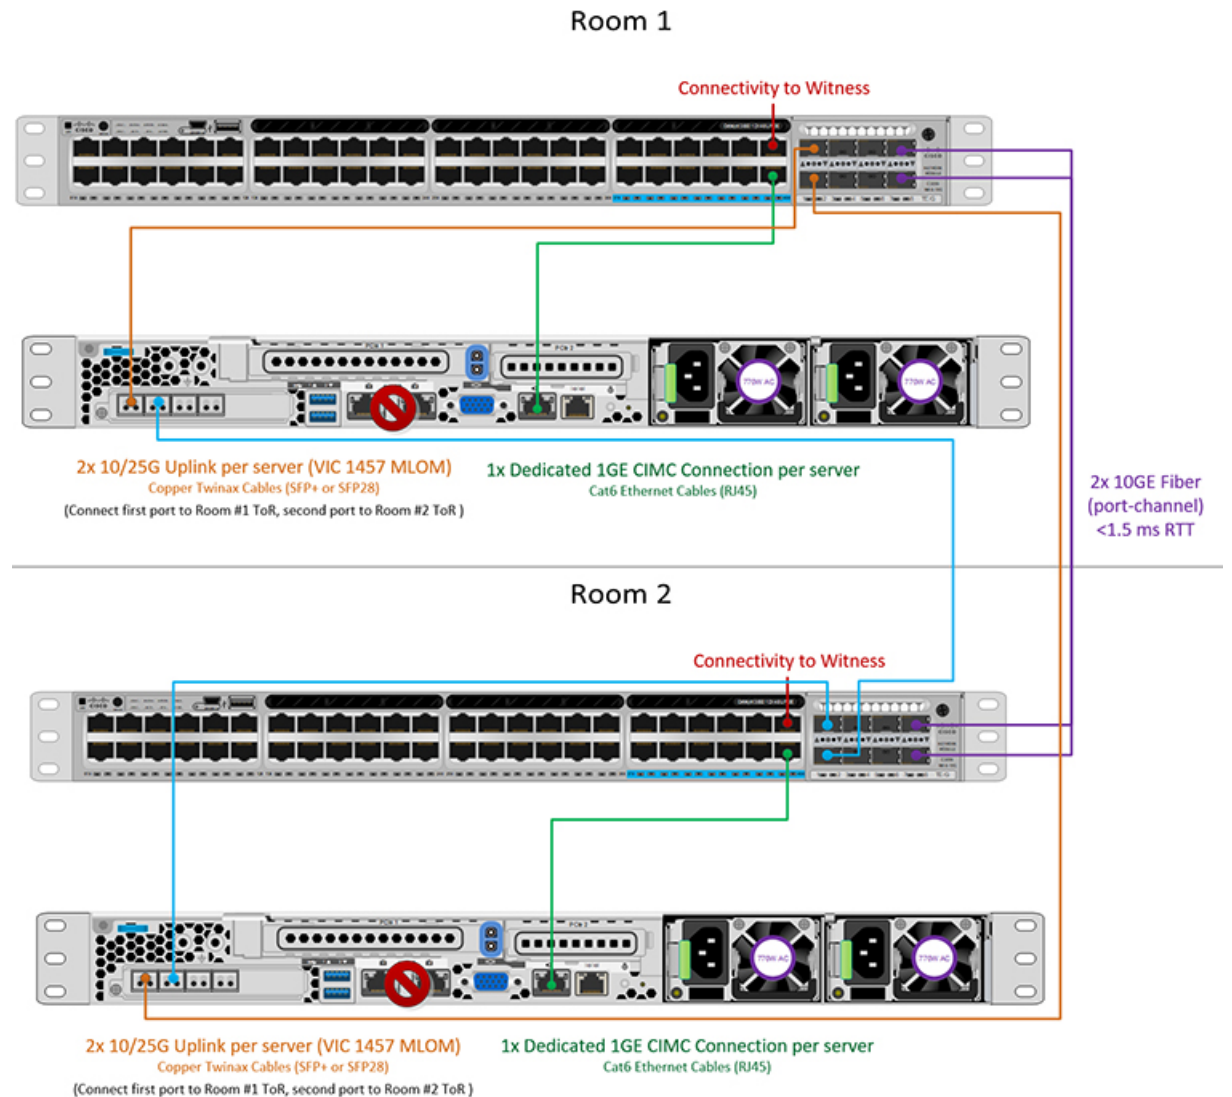
- If using dedicated Cisco IMC, connect the 1GE management port on each server (Labeled M on the back of the server) to the local switch.
- Connect one out of the four 10/25GE ports on the Cisco VIC from each server to the same ToR switch in room 1.
 - Use the same port number on each server to connect to the same switch.



Note Failure to use the same VIC port numbers will result in an extra hop for traffic between servers and will unnecessarily consume bandwidth between the two switches.

- Connect a second 10/25GE port on the Cisco VIC from each server to the ToR switch in room 2.
- Do not connect additional 10/25GE ports prior to cluster installation. After cluster deployment, you may optionally use the additional two 10/25GE ports for guest VM traffic.
- Ensure each switch has an independent network path to Intersight or a local witness server.

2-Node 2-Room Cross Connect



Cisco IMC Connectivity for All 2-Node 2-Room Topologies

Choose one of the following Cisco IMC Connectivity options for the 2-node 10/25 Gigabit Ethernet (GE) topology:

- Use of a dedicated 1GE Cisco IMC management port is recommended. This option requires additional switch ports and cables, however it avoids network contention and ensures always on, out of band access to each physical server.
- Use of shared LOM extended mode (EXT). In this mode, single wire management is used and Cisco IMC traffic is multiplexed onto the 10/25GE VIC connections. When operating in this mode, multiple streams of traffic are shared on the same physical link and uninterrupted reachability is not guaranteed. This deployment option is not recommended.

- In fabric interconnect-based environments, built in QoS ensures uninterrupted access to Cisco IMC and server management when using single wire management. In HyperFlex Edge environments, QoS is not enforced and hence the use of a dedicated management port is recommended.
- Assign an IPv4 management address to the Cisco IMC. For more information, see the procedures in the [Server Installation and Service Guide for the equivalent Cisco UCS C-series server](#). HyperFlex does not support IPv6 addresses.

10/25GE VIC-based Switch Configuration Guidelines

3 VLANs are required at a minimum.

- 1 VLAN for the following connections: VMware ESXi management, Storage Controller VM management and Cisco IMC management.
 - VMware ESXi management and Storage Controller VM management must be on the same subnet and VLAN.
 - A dedicated Cisco IMC management port may share the same VLAN with the management interfaces above or may optionally use a dedicated subnet and VLAN. If using a separate VLAN, it must have L3 connectivity to the management VLAN above and must meet Intersight connectivity requirements.
 - If using shared LOM extended mode for Cisco IMC management, a dedicated VLAN is recommended.
- 1 VLAN for Cisco HyperFlex storage traffic. This can and should be an isolated and non-routed VLAN. It must be unique and cannot overlap with the management VLAN.
- 1 VLAN for vMotion traffic. This can be an isolated and non-routed VLAN.



Note It is not possible to collapse or eliminate the need for these VLANs. The installation will fail if attempted.

- Additional VLANs as needed for guest VM traffic. These VLANs will be configured as additional portgroups in ESXi and should be trunked and allowed on all server facing ports on the ToR switch.
 - These additional guest VM VLANs are optional. You may use the same management VLAN above for guest VM traffic in environments that wish to keep a simplified flat network design.



Note Due to the nature of the Cisco VIC carving up multiple vNICs from the same physical port, it is not possible for guest VM traffic configured on vswitch-hx-vm-network to communicate L2 to interfaces or services running on the same host. It is recommended to either a) use a separate VLAN and perform L3 routing or b) ensure any guest VMs that need access to management interfaces be placed on the vswitch-hx-inband-mgmt vSwitch. In general, guest VMs should not be put on any of the HyperFlex configured vSwitches except for the vm-network vSwitch. An example use case would be if you need to run vCenter on one of the nodes and it requires connectivity to manage the ESXi host it is running on. In this case, use one of the recommendations above to ensure uninterrupted connectivity.

- Switchports connected to the Cisco VIC should be configured in trunk mode with the appropriate VLANs allowed to pass.
- Switchports connected to the dedicated Cisco IMC management port should be configured in ‘Access Mode’ on the appropriate VLAN.
- All cluster traffic will traverse the ToR switches in the 10/25GE topology
- Spanning tree portfast trunk (trunk ports) should be enabled for all network ports



Note Failure to configure portfast may cause intermittent disconnects during ESXi bootup and longer than necessary network re-convergence during physical link failure

Additional Considerations:

- Additional 3rd party NIC cards may be installed in the HX Edge nodes as needed. See the section in chapter 1 with the link to the networking guide.
- All non-VIC interfaces must be shut down or left un-cabled until installation is completed
- Only a single VIC is supported per HX Edge node in the MLOM slot. PCIe based VIC adapters are not supported with HX Edge nodes.

Virtual Networking Design for 2-Node 10/25GE VIC-Based Topology

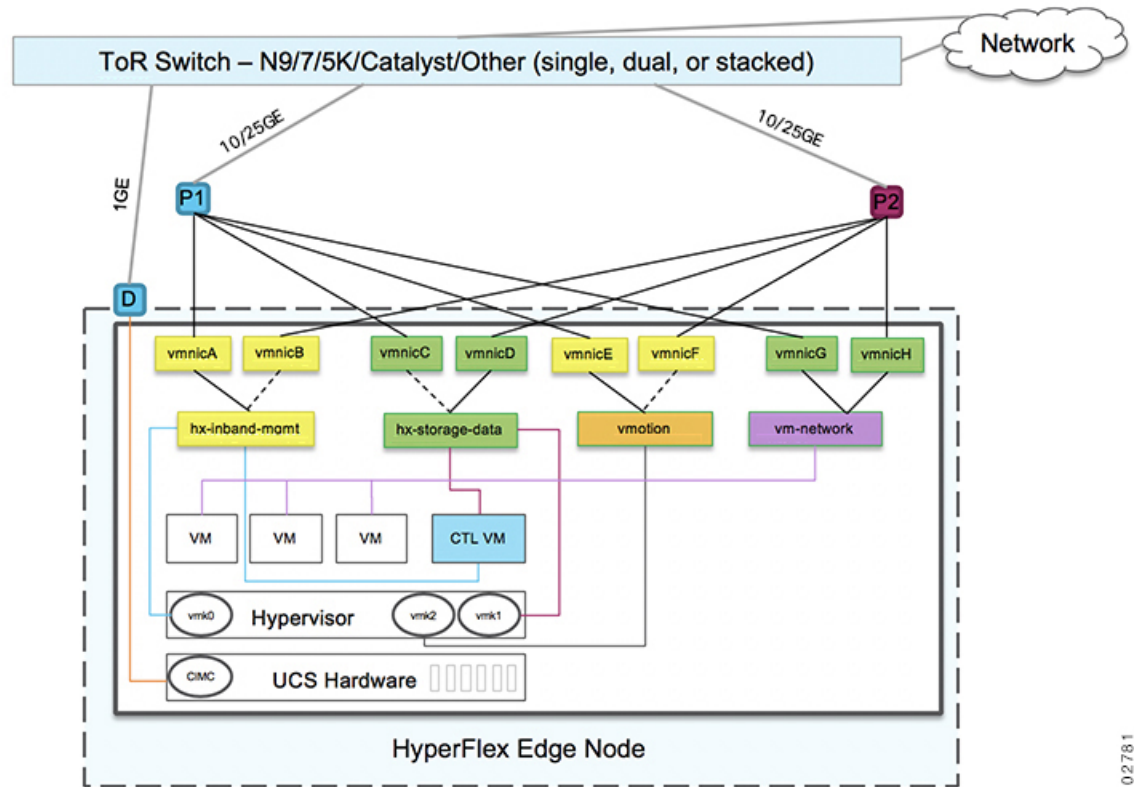
This section details the virtual network setup. No action is required as all of the virtual networking is set up automatically by the HyperFlex deployment process. These extra details are included below for informational and troubleshooting purposes.

Virtual Switches:

Four vSwitches are required:

- **vswitch-hx-inband-mgmt**—ESXi management (vmk0), storage controller management network
- **vswitch-hx-storage-data**—ESXi storage interface (vmk1), HX storage controller data network
- **vmotion**—vMotion interface (vmk2)
- **vswitch-hx-vm-network**—VM guest portgroups

Network Topology



5.02781

Failover Order:

- **vswitch-hx-inband-mgmt**—entire vSwitch is set for active/standby. All services by default consume a single uplink port and failover when needed.
- **vswitch-hx-storage-data**—HyperFlex storage data network and vmk1 are with the opposite failover order as inband-mgmt and vmotion vSwitches to ensure traffic is load balanced.
- **vmotion**—The vMotion VMkernel port (vmk2) is configured when using the post_install script. Failover order is set for active/standby.
- **vswitch-hx-vm-network**—vSwitch is set for active/active. Individual portgroups can be overridden as needed.

10 or 25 Gigabit Ethernet Stacked Switches Per Room Topology

This 10 or 25 Gigabit Ethernet (GE) switch topology provides a fully redundant design that protects against room, switch, link and port failures. A switch stack of at least two 10/25GE switches is required in each room. If a switch stack is not available, dual standalone switches can be combined to achieve similar results. Ensure there is ample bandwidth between the two switches in each room and between both switch stacks across rooms.

In this topology, each server is directly connected to just the local switches in each room. Unlike the cross connect topology, the inter-switch link (ISL) is a vital component used to carry all cluster storage and

management traffic between each room. The ISL must run at a minimum of 10GE with a maximum RTT latency of 1.5ms and should consist of multiple links in a port channel to ensure the links do not become saturated. With this topology, implementing quality of service (QoS) for storage data traffic is imperative as storage traffic is mixed alongside all other background traffic between the two rooms. To ensure HyperFlex storage remains reliable and performance, implement some form of priority queuing for the storage traffic.

Physical Network and Cabling for 10/25GE Stacked Switches Per Room Topology



Warning Proper cabling is important to ensure full network redundancy.

To deploy with dual or stacked switches per room (see diagram below for a visual layout):

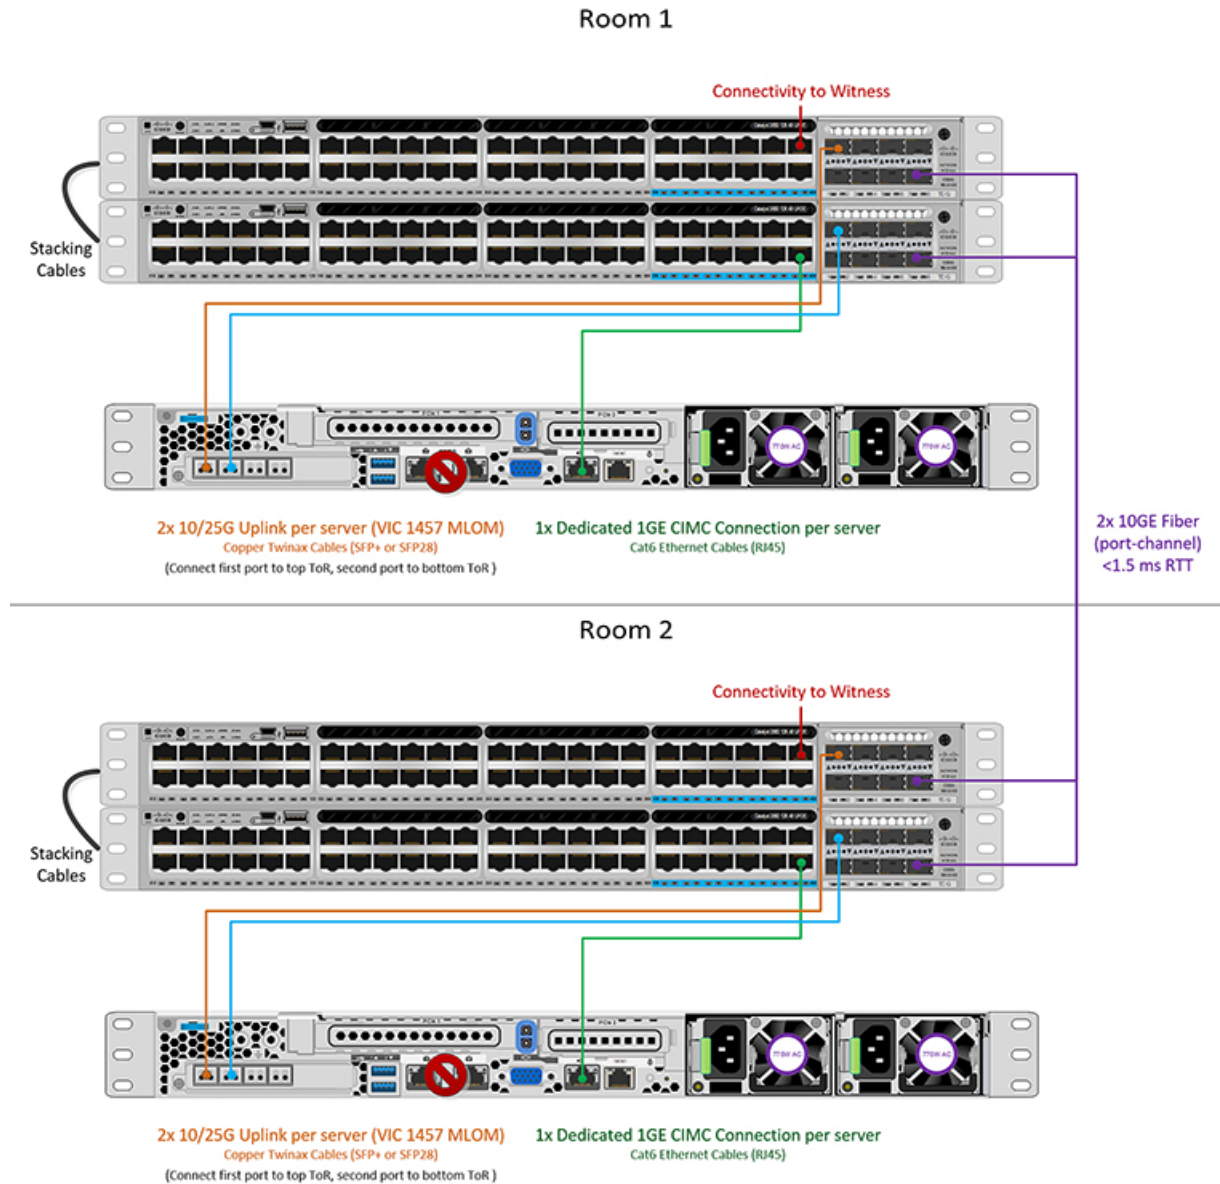
- If using dedicated Cisco IMC, connect the 1GE management port on each server (Labeled M on the back of the server) to one of the two switches.
- Connect one out of the four 10/25GE ports on the Cisco VIC from each server to the first ToR switch **in the same room**.
 - Use the same port number on each server to connect to the same switch.



Note Failure to use the same VIC port numbers will result in an extra hop for traffic between servers and will unnecessarily consume bandwidth between the two switches.

- Connect a second 10/25GE port on the Cisco VIC from each server to the second ToR switch **in the same room**. Use the same port number on each server to connect to the same switch.
- Do not connect additional 10/25GE ports prior to cluster installation. After cluster deployment, you may optionally use the additional two 10/25GE ports for guest VM traffic.
- Ensure each switch has an independent network path to Intersight or a local witness server.

2-Node 2-Room Dual/Stacked Switches



10/25 Gigabit Ethernet Stacked Switches Physical Cabling

Each room requires a pair of managed 10GE switches with VLAN capability. Cisco fully tests and provides reference configurations for Catalyst and Nexus switching platforms. Choosing one of these switches provides the highest level of compatibility and ensures a smooth deployment and seamless ongoing operations.

Each room requires the following: dual or stacked switches, and two 10/25GE ports, one 1GE port for CIMC management, and one Cisco VIC 1457 per server. Redundancy is provided at the room level and can tolerate the loss of either room as well as any smaller failure (e.g., switch failure, link failure, port failure).

Requirements for 10/25GE Stacked Switches Topology

The following requirements must be met across both rooms before starting deployment:

- Dedicated 1 Gigabit Ethernet (GE) Cisco IMC management port per server (recommended)
- 2 x 1GE ToR switch ports and two (2) Category 6 ethernet cables for dedicated Cisco IMC management port (customer supplied)
- Cisco VIC 1457 (installed in the MLOM slot in each server)
- Prior generation Cisco VIC hardware is not supported for 2 node HX Edge clusters.
- 4 x 10/25GE ToR switch ports and 4 x 10/25GE SFP+ or SFP28 cables (customer supplied. Ensure the cables you select are compatible with your switch model).
- Cisco VIC 1457 supports 10GE or 25GE interface speeds.
- Cisco VIC 1457 does not support 40GE interface speeds.

Requirements for HX Edge Clusters using 25GE



Note Using 25GE mode typically requires the use of forward error correction (FEC) depending on the transceiver or the type & length of cabling selected. The VIC 1400 series by default is configured in CL91 FEC mode (FEC mode “auto” if available in the Cisco IMC UI is the same as CL91) and does not support auto FEC negotiation. Certain switches will need to be manually set to match this FEC mode to bring the link state up. The FEC mode must match on both the switch and VIC port for the link to come up. If the switch in use does not support CL91, you may configure the VIC ports to use CL74 to match the FEC mode available on the switch. This will require a manual FEC mode change in the CIMC UI under the VIC configuration tab. Do not start a HyperFlex Edge deployment until the link state is up as reported by the switch and the VIC ports. CL74 is also known as FC-FEC (Firecode) and CL91 is also known as RS-FEC (Reed Solomon). For more information on how to change the FEC mode configured on the VIC using the Cisco IMC GUI, see the [Cisco UCS C-Series Integrated Management Controller GUI Configuration Guide, Release 4.1](#).

10 or 25 Gigabit Ethernet Single Switch Per Room Topology

This 10 or 25 Gigabit Ethernet (GE) switch topology provides a fully redundant design that protects against room, switch, link and port failures. A single 10/25GE switch is required in each room. Ensure there is ample bandwidth between the two switches in each room.

In this topology, each server is directly connected to just the local switch in each room. Unlike the cross connect topology, the inter-switch link (ISL) is a vital component used to carry all cluster storage and management traffic between each room. The ISL must run at a minimum of 10GE with a maximum RTT latency of 1.5ms and should consist of multiple links in a port channel to ensure the links do not become saturated. With this topology, implementing quality of service (QoS) for storage data traffic is imperative as storage traffic is mixed alongside all other background traffic between the two rooms. To ensure HyperFlex storage remains reliable and performance, implement some form of priority queueing for the storage traffic.

Physical Network and Cabling for 10/25GE Single Switch Per Room Topology

Each room requires a managed 10GE switch with VLAN capability. Cisco fully tests and provides reference configurations for Catalyst and Nexus switching platforms. Choosing one of these switches provides the highest level of compatibility and ensures a smooth deployment and seamless ongoing operations.

Each room requires the following: a single 10/25GE switch, and two 10/25GE ports, one 1GE port for CIMC management, and one Cisco VIC 1457 per server. Redundancy is provided at the room level and can tolerate the loss of either room as well as any smaller failure (e.g., switch failure, link failure, port failure).

Requirements for 10/25GE Single Switch Topology

The following requirements must be met across both rooms before starting deployment:

- Dedicated 1 Gigabit Ethernet (GE) Cisco IMC management port per server (recommended)
- 2 x 1GE ToR switch ports and two (2) Category 6 ethernet cables for dedicated Cisco IMC management port (customer supplied)
- Cisco VIC 1457 (installed in the MLOM slot in each server)
- Prior generation Cisco VIC hardware is not supported for 2 node HX Edge clusters.
- 4 x 10/25GE ToR switch ports and 4 x 10/25GE SFP+ or SFP28 cables (customer supplied. Ensure the cables you select are compatible with your switch model).
- Cisco VIC 1457 supports 10GE or 25GE interface speeds.
- Cisco VIC 1457 does not support 40GE interface speeds.

Requirements for HX Edge Clusters using 25GE



Note Using 25GE mode typically requires the use of forward error correction (FEC) depending on the transceiver or the type & length of cabling selected. The VIC 1400 series by default is configured in CL91 FEC mode (FEC mode “auto” if available in the Cisco IMC UI is the same as CL91) and does not support auto FEC negotiation. Certain switches will need to be manually set to match this FEC mode to bring the link state up. The FEC mode must match on both the switch and VIC port for the link to come up. If the switch in use does not support CL91, you may configure the VIC ports to use CL74 to match the FEC mode available on the switch. This will require a manual FEC mode change in the CIMC UI under the VIC configuration tab. Do not start a HyperFlex Edge deployment until the link state is up as reported by the switch and the VIC ports. CL74 is also known as FC-FEC (Firecode) and CL91 is also known as RS-FEC (Reed Solomon). For more information on how to change the FEC mode configured on the VIC using the Cisco IMC GUI, see the [Cisco UCS C-Series Integrated Management Controller GUI Configuration Guide, Release 4.1](#).

10/25 Gigabit Ethernet Single Switch Physical Cabling




Warning Proper cabling is important to ensure full network redundancy.

To deploy with a single switch per room (see diagram below for a visual layout):

- If using dedicated Cisco IMC, connect the 1GE management port on each server (Labeled M on the back of the server) to the local switch.
- Connect one out of the four 10/25GE ports on the Cisco VIC from each server to the ToR switch **in the same room**.
- Connect a second 10/25GE port on the Cisco VIC from each server to the ToR switch **in the same room**.
- Do not connect additional 10/25GE ports prior to cluster installation. After cluster deployment, you may optionally use the additional two 10/25GE ports for guest VM traffic.
- Ensure each switch has an independent network path to Intersight or a local witness server.

2-Node 2-Room Single Switch



Quality of Service (QoS)

In all the topologies listed in this chapter, it is highly recommended to implement QoS on the HyperFlex storage data traffic at a minimum. These 2-node 2-room configurations rely heavily on the inter-site link (ISL) for carrying storage traffic between the two HyperFlex nodes and the link could become saturated by other background traffic. Cisco recommends the following:

- Ensure ample bandwidth and link redundancy for the ISL. Using multiple high bandwidth links in a port channel helps to reduce the need for QoS by ensuring ample capacity for all types of traffic between rooms. Avoid link speed mismatches along the end-to-end storage path as speed mismatches can create network bottlenecks.
- Classify incoming traffic to the switch based on IP address. HyperFlex Edge does not pre-mark any traffic and it is up to the switch to classify traffic. Use the HyperFlex Data Platform storage network IP addresses for this classification. Typically, these IP addresses exist in the 169.254.x.x range as a /24 network. You can find the proper range by investigating the controller VM configuration in vCenter or running ifconfig command on Controller VM and noting the subnet in use for the eth1 interface.
- It is recommended to match the entire /24 subnet so that as clusters are expanded with more nodes, all storage traffic continues to be properly classified.
- Mark storage traffic according to environmental needs. In the example configurations with Catalyst 9000, DSCP EF is used. End-to-end QoS is achieved using DSCP header values only.
- Queue based on your switch platform's capabilities. For the Catalyst 9000 example, one of the priority queues is used to prioritize the HX storage traffic (marked EF) across the inter-site link. HyperFlex storage traffic performs best on a high priority queue with low latency and high bandwidth. Increasing the assigned buffer of the queue will also help reduce packet loss when there is link transmission delay.
- Apply the QoS configuration to the ingress interfaces (for marking) and egress interfaces (for queueing).
- Apply additional QoS configurations as needed for management traffic, vMotion, and application traffic. It is recommended to prioritize traffic in the following order:
 1. Management - DSCP CS6
 2. VM or application traffic – DSCP CS4
 3. vMotion – DSCP CS0

The above DSCP values are recommended. You can however, use any values as necessary to meet environmental needs. For each type of traffic, create an ACL for marking based on IP range. Then create a class-map to match the ACL. Add to the existing marking policy class and specify a set action. Finally, update the egress queueing policy with a dedicated class per traffic type that matches the DSCP marking and specifies the desired bandwidth.

Common Network Requirement Checklist

Before you begin installation, confirm that your environment meets the following specific software and hardware requirements.

VLAN Requirements



Important **Reserved VLAN IDs** - The VLAN IDs you specify must be supported in the Top of Rack (ToR) switch where the HyperFlex nodes are connected. For example, VLAN IDs 3968 to 4095 are reserved by Nexus switches and VLAN IDs 1002 to 1005 are reserved by Catalyst switches. Before you decide the VLAN IDs for HyperFlex use, make sure that the same VLAN IDs are available on your switch.

Network	VLAN ID	Description
Use a separate subnet and VLANs for each of the following networks:		
VLAN for VMware ESXi, and Cisco HyperFlex management		Used for management traffic among ESXi, HyperFlex, and VMware vCenter, and must be routable. Note This VLAN must have access to Intersight (Intersight is required for 2-Node deployment).
CIMC VLAN		Can be same or different from the Management VLAN. Note This VLAN must have access to Intersight (Intersight is required for 2-Node deployment).
VLAN for HX storage traffic		Used for raw storage traffic and requires only L2 connectivity.
VLAN for VMware vMotion		Used for vMotion VLAN.
VLAN(s) for VM network(s)		Used for VM/application network. Note Can be multiple VLANs, each backed by a different VM portgroup in ESXi.

Supported vCenter Topologies

Use the following table to determine the topology supported for vCenter.

Topology	Description	Recommendation
Single vCenter	Virtual or physical vCenter that runs on an external server and is local to the site. A management rack mount server can be used for this purpose.	Highly recommended

Topology	Description	Recommendation
Centralized vCenter	vCenter that manages multiple sites across a WAN.	Highly recommended
Nested vCenter	vCenter that runs within the cluster you plan to deploy.	Installation for a HyperFlex Edge cluster may be initially performed without a vCenter. Alternatively, you may deploy with an external vCenter and migrate it into the cluster. In either case, the cluster must be registered to a vCenter server before running production workloads. For the latest information, see the How to Deploy vCenter on the HX Data Platform tech note.

Customer Deployment Information

A typical two-node HyperFlex Edge deployment requires 9 IP addresses – 7 IP addresses for the management network and 2 IP addresses for the vMotion network.



Important All IP addresses must be IPv4. HyperFlex does not support IPv6 addresses.

CIMC Management IP Addresses

Server	CIMC Management IP Addresses
Server 1:	
Server 2:	
Subnet mask	
Gateway	
DNS Server	
NTP Server	
Note NTP configuration on CIMC is required for proper Intersight connectivity.	

Network IP Addresses



Note By default, the HX Installer automatically assigns IP addresses in the 169.254.X.X range as a /24 network, to the Hypervisor Data Network and the Storage Controller Data Network. This IP subnet is not user configurable.



Note Spanning Tree portfast trunk (trunk ports) should be enabled for all network ports. Failure to configure portfast may cause intermittent disconnects during ESXi bootup and longer than necessary network re-convergence during physical link failure.

Management Network IP Addresses (must be routable)	
Hypervisor Management Network	Storage Controller Management Network
Server 1:	Server 1:
Server 2:	Server 2:
Storage Cluster Management IP address	Cluster IP:
Subnet mask	
Default gateway	

VMware vMotion Network IP Addresses

For vMotion services, you may configure a unique VMkernel port or, if necessary, reuse the vmk0 if you are using the management VLAN for vMotion (not recommended).

Server	vMotion Network IP Addresses (configured using the post_install script)
Server 1:	
Server 2:	
Subnet mask	
Gateway	

VMware vCenter Configuration



Note HyperFlex communicates with vCenter through standard ports. Port 80 is used for reverse HTTP proxy and may be changed with TAC assistance. Port 443 is used for secure communication to the vCenter SDK and may not be changed.

vCenter admin username <i>username@domain</i>	
vCenter admin password	
vCenter data center name Note An existing datacenter object can be used. If the datacenter doesn't exist in vCenter, it will be created.	
VMware vSphere compute cluster and storage cluster name Note Cluster name you will see in vCenter.	

Port Requirements



Important Ensure that the following port requirements are met in addition to the prerequisites listed for [Intersight Connectivity](#), on page 18.

If your network is behind a firewall, in addition to the standard port requirements, VMware recommends ports for VMware ESXi and VMware vCenter.

- CIP-M is for the cluster management IP.
- SCVM is the management IP for the controller VM.
- ESXi is the management IP for the hypervisor.

The comprehensive list of ports required for component communication for the HyperFlex solution is located in Appendix A of the [HX Data Platform Security Hardening Guide](#)



Tip If you do not have standard configurations and need different port settings, refer to [Table C-5 Port Literal Values](#) for customizing your environment.

Network Services



- Note**
- DNS and NTP servers should reside outside of the HX storage cluster.
 - Use an internally-hosted NTP server to provide a reliable source for the time.
 - All DNS servers should be pre-configured with forward (A) and reverse (PTR) DNS records for each ESXi host before starting deployment. When DNS is configured correctly in advance, the ESXi hosts are added to vCenter via FQDN rather than IP address.
- Skipping this step will result in the hosts being added to the vCenter inventory via IP address and require users to change to FQDN using the following procedure: [Changing Node Identification Form in vCenter Cluster from IP to FQDN](#).

DNS Servers <i><Primary DNS Server IP address, Secondary DNS Server IP address, ... ></i>	
NTP servers <i><Primary NTP Server IP address, Secondary NTP Server IP address, ... ></i>	
Time zone <i>Example: US/Eastern, US/Pacific</i>	

Connected Services

Enable Connected Services (Recommended) <i>Yes or No required</i>	
Email for service request notifications <i>Example: name@company.com</i>	

Proxy Server

- Use of a proxy server is optional if direct connectivity to Intersight is not available.
- When using a proxy, the device connectors in each server must be configured to use the proxy in order to claim the servers into an Intersight account. In addition, the proxy information must be provided in the HX Cluster Profile to ensure the HyperFlex Data Platform can be successfully downloaded.
- Use of username/password is optional

Proxy required: Yes or No	
----------------------------------	--

Proxy Host	
Proxy Port	
Username	
Password	

Guest VM Traffic

Considerations for guest VM traffic are given above based on the topology selection. In general, guest port groups may be created as needed so long as they are applied to the correct vSwitch:

- 10/25GE Topology: use **vswitch-hx-vm-network** to create new VM port groups.

Cisco recommends you run the `post_install` script to add more VLANs automatically to the correct vSwitches on all hosts in the cluster. Execute **hx_post_install --vlan** (space and two dashes) to add new guest VLANs to the cluster at any point in the future.

Additional vSwitches may be created that use leftover vmnics or third party network adapters. Care should be taken to ensure no changes are made to the vSwitches defined by HyperFlex.



Note Additional user created vSwitches are the sole responsibility of the administrator, and are not managed by HyperFlex.

Intersight Connectivity

Consider the following prerequisites pertaining to Intersight connectivity:

- Before installing the HX cluster on a set of HX servers, make sure that the device connector on the corresponding Cisco IMC instance is properly configured to connect to Cisco Intersight and claimed.
- Communication between CIMC and vCenter via ports 80, 443 and 8089 during installation phase.
- All device connectors must properly resolve *svc.intersight.com* and allow outbound initiated HTTPS connections on port 443. The current version of the HX Installer supports the use of an HTTP proxy.
- All controller VM management interfaces must properly resolve *svc.intersight.com* and allow outbound initiated HTTPS connections on port 443. The current version of HX Installer supports the use of an HTTP proxy if direct Internet connectivity is unavailable.
- IP connectivity (L2 or L3) is required from the CIMC management IP on each server to all of the following: ESXi management interfaces, HyperFlex controller VM management interfaces, and vCenter server. Any firewalls in this path should be configured to allow the necessary ports as outlined in the [Hyperflex Hardening Guide](#).
- Starting with HXDP release 3.5(2a), the Intersight installer does not require a factory installed controller VM to be present on the HyperFlex servers.

When redeploying HyperFlex on the same servers, new controller VMs must be downloaded from Intersight into all ESXi hosts. This requires each ESXi host to be able to resolve *svc.intersight.com* and

allow outbound initiated HTTPS connections on port 443. Use of a proxy server for controller VM downloads is supported and can be configured in the HyperFlex Cluster Profile if desired.

- Post-cluster deployment, the new HX cluster is automatically claimed in Intersight for ongoing management.

Cisco HyperFlex Edge Invisible Cloud Witness

The Cisco HyperFlex Edge Invisible Cloud Witness is an innovative technology for Cisco HyperFlex Edge Deployments that eliminates the need for witness VMs or arbitration software.

The Cisco HyperFlex Edge invisible cloud witness is only required for 2-node HX Edge deployments. The witness does not require any additional infrastructure, setup, configuration, backup, patching, or management of any kind. This feature is automatically configured as part of a 2-node HyperFlex Edge installation. Outbound access at the remote site must be present for connectivity to Intersight (either Intersight.com or to the Intersight Virtual Appliance). HyperFlex Edge 2-node clusters cannot operate without this connectivity in place.

For additional information about the benefits, operations, and failure scenarios of the Invisible Cloud Witness feature, see <https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/whitepaper-c11-741999.pdf>

Ordering Cisco HyperFlex Edge Servers

When ordering Cisco HyperFlex Edge servers, be sure to choose the correct components as outlined in the [HyperFlex Edge spec sheets](#). Pay attention to the network topology selection to ensure it matches your desired configuration. Further details on network topology PID selection can be found in the supplemental material section of the spec sheet.

Installation Overview



Note If the HyperFlex cluster nodes were part of any other HyperFlex cluster before (or not factory shipped), follow the node cleanup procedure before starting the cluster deployment. For more information, see [HyperFlex Customer Cleanup Guides for FI and Edge](#).

Refer to the following table that summarizes the installation workflow for Edge deployments. The Steps 1-3 are common between 1GE and 10/25GE deployments. However, Step 4 is applicable for 1GE deployments, while the remaining Steps 5-6 are for 10/25GE deployments.

Step#	Description	Reference	Applicable for 1GE & 10/25GE, 1GE, or 10/25GE
1	Complete preinstallation checklist.	<p>Make a selection from below based on your switch configuration:</p> <p>2-Node Edge Deployments:</p> <ul style="list-style-type: none"> • 10 or 25GE VIC-Based Topology, on page 5 or, • 1 Gigabit Ethernet Topology, on page 19 <p>3- and 4-Node Edge Deployments:</p> <ul style="list-style-type: none"> • 10 or 25GE VIC-Based Topology, on page 25 or, • 1 Gigabit Ethernet Topology, on page 29 	1GE & 10/25GE
2	Complete installation prerequisites.	<ul style="list-style-type: none"> • Rack Cisco HyperFlex Nodes, on page 55 • Cisco Integrated Management Controller Configuration, on page 56 • Verifying Firmware Versions, on page 57 	1GE & 10/25GE
3	Download and deploy Cisco HX Data Platform Installer.	Deploying Cisco HX Data Platform Installer, on page 58	1GE & 10/25GE
4	Deploy HyperFlex Edge cluster.	<p>Complete the following steps to configure your Edge cluster and verify successful installation.</p> <ul style="list-style-type: none"> • Configuring Your HyperFlex Cluster, on page 59 • Verifying Cisco HX Data Platform Software Installation, on page 65 	1GE only
6	Deploy HyperFlex Edge cluster.	(10/25GE only) Configuring Your HyperFlex Cluster, on page 65	10/25GE only

Rack Cisco HyperFlex Nodes

For details on installation of Cisco HX220c M5 HyperFlex Nodes or Cisco HX220c M6 HyperFlex Nodes, review the [Cisco Hardware Install Guides](#).



Important You can use a console dongle to connect the VGA monitor and keyboard for CIMC configuration. You can also directly connect to the VGA and USB ports on the rear of the server. Alternatively, you can perform a lights-out configuration of CIMC if a DHCP server is available in the network.

Cisco Integrated Management Controller Configuration

Choose one method for CIMC network configuration: static assignment or DHCP assignment.

Configuring CIMC: Static Assignment

To configure Cisco Integrated Management Controller (CIMC), you must enable CIMC standalone mode, configure the CIMC password and settings, and configure a static IP address manually using a KVM. This requires physical access to each server with a monitor and keyboard. Each server must be configured one at a time.

Customers may opt to use the dedicated CIMC management port for out-of-band use. Users should account for this third 1GE port when planning their upstream switch configuration. Additionally, the user should set the CIMC to dedicated mode during CIMC configuration. Follow [Cisco UCS C-series documentation](#) to configure the CIMC in dedicated NIC mode. Under **NIC properties**, set the NIC mode to **dedicated** before saving the configuration.

Before you begin

- Ensure that all Ethernet cables are connected as described in the Physical Cabling section of this guide that applies to your deployment.
- Attach the VGA dongle to the server and connect a monitor and USB keyboard.

-
- Step 1** Power on the server, and wait for the screen with the Cisco logo to display.
- Step 2** When prompted for boot options, press **F8** to enter the **Cisco IMC Configuration** utility.
- Step 3** In **CIMC User Details**, enter **password** for the **current CIMC password**, enter your **new CIMC password** twice, and press **Enter** to save your new password.
- Important** Systems ship with a default password of *password* that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.
- Step 4** For **IP (Basic)**, check **IPv4**, uncheck **DHCP enabled**, and enter values for **CIMC IP**, **Prefix/Subnet** mask, and **Gateway**.
- Step 5** For **VLAN (Advanced)**, check **VLAN enabled**, and:
- If you are using trunk ports, set the appropriate **VLAN ID**.
 - If you are using access ports, leave this field blank.
- Step 6** Leave the rest of the settings as default, press **F10** to save your configuration, and press **ESC** to exit the utility.
- Step 7** In a web browser, navigate directly to the CIMC page at **https://CIMC IP address**.
- Step 8** Enter the username **admin** and your new CIMC password, and click **Log In**.
- Step 9** Manually set the power policy to match the desired operation from **Server > Power Policies**.

Servers default to the Power Off power-restore policy set at the factory.

What to do next

You can use the virtual KVM console or continue to use the physical KVM. The SD cards have ESXi preinstalled from the factory and boot automatically during installation.

Configuring CIMC: DHCP Assignment

To configure Cisco Integrated Management Controller (CIMC), you must enable CIMC standalone mode, configure the CIMC password and settings, and configure a dynamic IP address obtained through DHCP. This requires more network setup but eases configuration by enabling a lights-out setup of HyperFlex Edge nodes. All servers lease addresses automatically and in parallel, reducing deployment time.

Before you begin

- Ensure that all Ethernet cables are connected as described in the Physical Cabling section of this guide that applies to your deployment.
- Ensure the DHCP server is properly configured and running with a valid scope.
- Ensure the DHCP server is directly listening on the management VLAN or you have an IP helper configured on your switch(es).
- Decide on inband versus out-of-band CIMC:
 - If using inband CIMC, configure the native VLAN for all HyperFlex Edge switch ports to match the correct DHCP VLAN. This is the only way to ensure that the CIMC can lease an address automatically.
 - If using out-of-band CIMC, configure the dedicated switch port for access mode on the DHCP VLAN.

Step 1 Connect power cables.

Step 2 Access the DHCP logs or lease table to determine the CIMC addresses obtained

Step 3 Search the hostnames for **C220-*<S/N>*** to find your HyperFlex servers, and make note of the addresses for required inputs into the HX Data Platform Installer.

What to do next

When using DHCP, you must manually set a user defined CIMC password before beginning HyperFlex Data Platform installation. Use either the web UI or a CLI session to set a new password. The default password of **password** must be changed or installation fails.

Verifying Firmware Versions

You need to view current BIOS, CIMC, SAS HBA, and drive firmware versions, and verify that those versions match data in the Release Notes.

-
- Step 1** In your browser, log into the CIMC web UI by navigating to **https://<CIMC IP>**.
- Step 2** In the Navigation pane, click **Server**.
- Step 3** On the **Server** page, click **Summary**.
- Step 4** In the **Cisco Integrated Management Controller (CIMC) Information** section of the **Server Summary** page, locate and make a note of the **BIOS Version** and **CIMC Firmware Version**.
- Step 5** In CIMC, navigate to **Inventory > PCIe Adapters**, and locate and make a note of the **SAS HBA Version**.
- Step 6** In CIMC, navigate to **Storage** depending on which server type you are using, navigate to one of the following:
- For M4, **Cisco 12G Modular SAS > Physical Drive Info**, and make a note of the drive type, manufacturer, and firmware version.
 - For M5 and M6, **Cisco 12G SAS HBA > Physical Drive Info**, and make a note of the drive type, manufacturer, and firmware version.
- Step 7** Compare the current BIOS, CIMC, SAS HBA, and drive firmware versions with the versions listed in the Release Notes.
- Step 8** If the minimum versions are not met, use the Host Update Utility (HUU) Download Links in the compatibility matrix to upgrade the firmware versions running on the system, including Cisco Virtual Interface Cards (VIC), PCI Adapter, RAID controllers, and drive (HDD/SSD) firmware. You can find current and previous releases of the Cisco HUU User Guide at this location: <http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-c-series-rack-servers/products-user-guide-list.html>.
-

Deploying Cisco HX Data Platform Installer

HX Data Platform Installer can be deployed on an ESXi server, as well as a VMware Workstation, VMware Fusion, or Virtual Box. The HyperFlex software is distributed as a deployable virtual machine, contained in an Open Virtual Appliance (OVA) file format. Use the following procedure to deploy HX Data Platform Installer using a VMware vSphere (thick) Client.

-
- Step 1** Download the HX Data Platform Installer OVA from [Cisco.com](http://www.cisco.com), and save the package locally. Verify the downloaded version matches the recommended version for your deployment.
- Step 2** Log into vCenter using the vSphere client.
- Step 3** Select **File > Deploy OVF Template**.
- Step 4** In the **Deploy OVF Template** wizard, on the **Source** page, specify the source location, and click **Next**.
- Step 5** On the **OVF Template Details** page, view the information, and click **Next**.
- Step 6** (Optional) On the **Name and Location** page, edit the name and location for the virtual appliance, and click **Next**.
- Step 7** On the **Host/Cluster** page, select the host or cluster on which you want to deploy, and click **Next**.
- Step 8** On the **Resource Pool** page, select the resource pool where you want to run the OVF template, and click **Next**.
- Step 9** On the **Storage** page, select a datastore to store the deployed OVF template, and click **Next**.
- Step 10** On the **Disk Format** page, select the disk format to store the virtual machine virtual disks, and click **Next**.
- Step 11** On the **Network Mapping** page, for each network specified in the OVF template, right-click the **Destination Network** column to select a network in your infrastructure, and click **Next**.
- Step 12** Provide the OVF properties for the installer VM, namely: hostname, default gateway, DNS server, IP address, and subnet mask.

Alternatively, leave all of the OVF properties blank for a DHCP assigned address.

Step 13 On the **Ready to Complete** page, select **Power On After Deployment**, and click **Finish**.

Configuring Your HyperFlex Cluster

Step 1 In your web browser, enter the IP address of the installer VM, and click **Accept** or **Continue** to bypass any SSL certificate errors.

Step 2 Verify the HyperFlex installer **Build ID** in the lower right corner of the login screen.

Step 3 Log into Cisco HX Data Platform Installer using username **root** and password **Cisco123**.

Important Systems ship with a default password of *Cisco123* that must be changed during installation. The HyperFlex on-premises installer requires changing the root password as part of deployment. You cannot continue installation unless you specify a new password. Use the new password at this point in the configuration procedure.

Step 4 Read the End User Licensing Agreement, check **I accept the terms and conditions**, and click **Login**.

Step 5 On the **Workflow** page, click **Cluster Creation with HyperFlex Edge**.

Step 6 To perform cluster creation, you can import a *JSON configuration* file with the required configuration data. The following two steps are optional if importing a JSON file, otherwise you can input data into the required fields manually.

Note For a first-time installation, contact your Cisco representative to procure the factory preinstallation JSON file.

- a. Click **Select a file** and choose your *JSON file* to load the configuration. Select **Use Configuration**.
- b. An **Overwrite Imported Values** dialog box displays if your imported values for Cisco UCS Manager are different. Select **Use Discovered Values**.

Step 7 On the **Credentials** page, complete the following fields, and click **Continue**.

Name	Description
Cisco IMC Credentials	
Cisco IMC User Name	Cisco IMC username. By default, the username is admin .
Password	CIMC password. By default, the password is password .
vCenter Credentials	
Configuring Your HyperFlex Cluster Server	FQDN or IP address of the vCenter server. You must use an account with vCenter root-level admin permissions.
User Name	Administrator username.
Admin Password	Administrator password.
Hypervisor Credentials	

Name	Description
Admin User Name	Administrator username. By default, the username is root .
Hypervisor Password	Default password is Cisco123. Important Systems ship with a default password of <i>Cisco123</i> that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.

Use the following screenshot as a reference to complete the fields in this page.

The screenshot shows the 'HyperFlex Installer' interface with the 'Credentials' tab selected. It contains the following fields:

- Cisco IMC Credentials:** Cisco IMC User Name (eng-f55.eng.com), Password (masked).
- vCenter Credentials:** vCenter Server (10.64.53.54), User Name (administrator@vsphere.local), Admin Password (masked).
- Hypervisor Credentials:** Admin User name (root), a checkbox for 'The hypervisor on this node uses the factory default password' (unchecked), and Hypervisor Password (masked).
- Configuration Panel:** A dashed box with the text 'Drag and drop configuration files here or' and a 'Select a File' button.
- Navigation:** '< Back' and 'Continue' buttons at the bottom.

Step 8 On the **IP Addresses** page, enter the assigned addresses for each server.

Name	Description
Cisco IMC	IP Address of Cisco IMC
Hypervisor	Management IP for Hypervisor
Storage Controller	Management IP for Storage Controller
Cluster IP Address	Cluster management IP address
Subnet mask	Subnet mask for cluster management
Gateway	Gateway IP address for cluster management IP

Use the following screenshot as a reference to complete the fields in this page.

The screenshot shows the 'HyperFlex Installer' interface. The main section is titled 'IP Addresses' and includes a table for configuring three servers. The table has columns for Name, Cisco IMC, Hypervisor, and Storage Controller. Below the table are fields for Cluster IP Address, Subnet Mask, and Gateway. A 'Configuration' sidebar on the right shows 'Credentials' with fields for vCenter Server, User Name, Cisco IMC User Name, and Admin User name. A 'Continue' button is visible at the bottom right.

Step 9 On the **Cluster Configuration** page, complete the following fields, and click **Continue**.

Note Complete all the fields using your pre-install worksheet.

Name	Description
Cisco HX Cluster	
Cluster Name	User-supplied name for the HyperFlex storage cluster.
Replication Factor	Support for Replication Factor 3 for 3- and 4-Node edge clusters was introduced in HXDP Release 4.5. The default Replication Factor for a 2-Node edge cluster is 2 .
Controller VM	
Create Admin Password	There is no default password for the Controller VM. User needs to set this field
Confirm Admin Password	Confirm the Administrator password.
vCenter Configuration	
vCenter Datacenter Name	The name of the vCenter datacenter where the HyperFlex hosts were added.
vCenter Cluster Name	The name of the vCenter cluster where the HyperFlex hosts were added.
System Services	

Name	Description
DNS Server(s)	A comma-separated list of IP addresses for each DNS Server.
NTP Server(s)	A comma-separated list of IP addresses for each NTP Server. Important A highly reliable NTP server is required.
Time Zone	The local time zone for the controller VM.
Connected Services	
Enable Connected Services (Recommended)	Check to Enable Connected Services. Note We highly recommend enabling Connected Services to enable sending email alerts to Cisco TAC.
Send service ticket notifications to: <i>Example: admin@cisco.com</i>	Email address to receive service request notifications.
Advanced Networking	
Management VLAN Tag Data VLAN Tag	Enter the correct VLAN tags if you are using trunk ports. The VLAN tags must be different when using trunk mode. Enter 0 for both VLAN tags if you are using access ports. Note Do not enter 0 if you are using trunk ports.
Management vSwitch Data vSwitch	Do not change the auto-populated vSwitch name.
Advanced Configuration	
Enable Jumbo Frames on Data Network	Do not check to ensure HyperFlex Edge deployments use regularly-sized packets. You may optionally enable jumbo frames for 10/25GE deployments depending on your network configuration. For ease of deployment, it is recommended to uncheck this option.
Clean up disk partitions	Check to remove all existing data and partitions from every node in the HX storage cluster. For example, if this is not the first time installing the software on the cluster.
Optimize for VDi only deployment	Check to optimize VDI deployments. By default HyperFlex is performance optimized for Virtual Server Infrastructure (VSI). Check this box to tune the performance parameters for VDI deployments. This option has no affect on all-flash HX models and only needs to be enabled for hybrid HX clusters. If you are running mixed VDI and VSI workloads, do not select this option.
vCenter Single-Sign-On Server	Fill in this field only if instructed by Cisco TAC.

Use the following screenshot as a reference to complete the fields in this page.

HyperFlex Installer

Credentials | IP Addresses | Cluster Configuration

Cisco HX Cluster

Cluster Name: Replication Factor:

Controller VM

Create Admin Password: Confirm Admin Password:

vCenter Configuration

vCenter Datacenter Name: vCenter Cluster Name:

System Services

DNS Server(s): NTP Server(s): DNS Domain Name:

Time Zone:

Connected Services

Connected Services: Enable Connected Services (Recommended)

Send service ticket notifications to:

Advanced Networking

Management VLAN Tag: Management vSwitch:

Data VLAN Tag: Data vSwitch:

Advanced Configuration

Jumbo Frames: Enable Jumbo Frames on Data Network

Disk Partitions: Clean up disk partitions

Virtual Desktop (VDI): Optimize for VDI only deployment

vCenter Single-Sign-On Server:

Configuration

Credentials

vCenter Server	10.64.53.54
User Name	administrator@vsphere.local
Cisco IMC User Name	eng-f55.eng.com
Admin User name	root

IP Addresses

Cluster Name	EDGE_CLUSTER
Management Cluster	10.10.10.40
Management Subnet Mask	255.0.0.0
Management Gateway	10.10.10.1

Server 1

Management Cisco IMC	10.10.10.10
Management Hypervisor	10.10.10.20
Management Storage Controller	10.10.10.30

Server 2

Management Cisco IMC	10.10.10.11
----------------------	-------------

Step 10 After deployment finishes, the **Summary Deployment** page displays a summary of the deployment details.

What to do next

Confirm HX Data Platform Plug-in installation. See [Verifying Cisco HX Data Platform Software Installation, on page 65](#)

Verifying Cisco HX Data Platform Software Installation

Step 1 Launch **vSphere**, and log into the vCenter Server as an administrator.

Step 2 Under **vCenter Inventory Lists**, verify that **Cisco HX Data Platform** displays.

If the entry for Cisco HX Data Platform does not appear, log out of vCenter, close the browser, and log back in. In most cases, the issue is resolved by this action.

If logging out of vCenter does not fix the issue, you may have to restart the vSphere Web Client . SSH to the VCSA and run `service vsphere-client restart`. For a Windows vCenter, restart VMware vSphere Web Client in the services page in mmc.

Step 3 Ensure that your new cluster is online and registered.

(10/25GE only) Configuring Your HyperFlex Cluster

Step 1 In your web browser, enter the IP address of the installer VM, and click **Accept** or **Continue** to bypass any SSL certificate errors.

Step 2 Verify the HyperFlex installer **Build ID** in the lower right corner of the login screen.

Step 3 Log in to Cisco HX Data Platform Installer using username **root** and password **Cisco123**.

Important Systems ship with a default password of *Cisco123* that must be changed during installation. The HyperFlex on-premises installer requires changing the root password as part of deployment. You cannot continue installation unless you specify a new password. Use the new password at this point in the configuration procedure.

Step 4 Read the End User Licensing Agreement, check **I accept the terms and conditions**, and click **Login**.

Step 5 On the **Workflow** page, click **Cluster Creation with HyperFlex Edge**.

Step 6 On the **Credentials** page, complete the following fields, and click **Continue**.

Name	Description
vCenter Credentials	
vCenter Server	FQDN or IP address of the vCenter server. You must use an account with vCenter root-level admin permissions.
User Name	Administrator username.

Name	Description
Admin Password	Administrator password.
CIMC Credentials	
CIMC User Name	CIMC username. By default, the username is admin .
Password	CIMC password. By default, the password is password .
Hypervisor Credentials	
Admin User Name	Administrator username. By default, the username is root .
Admin Password	Default password is Cisco123. Important Systems ship with a default password of <i>Cisco123</i> that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.

Step 7 On the **IP Addresses** page, enter the assigned addresses for each server:

Name	Description
Cisco IMC	IP Address of Cisco IMC
Hypervisor	Management IP for Hypervisor
Storage Controller	Management IP for Storage Controller
Cluster IP Address	Cluster management IP address
Subnet mask	Subnet mask for cluster management
Gateway	Gateway IP address for cluster management IP

Step 8 On the **Cluster Configuration** page, complete the following fields, and click **Continue**.

Note Complete all the fields using your pre-install worksheet.

Name	Description
Cisco HX Cluster	
Cluster Name	User-supplied name for the HyperFlex storage cluster.
Replication Factor	
Controller VM	

Name	Description
Create Admin Password	Default password is Cisco123. Important Systems ship with a default password of <i>Cisco123</i> that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.
Confirm Admin Password	Confirm the Administrator password.
vCenter Configuration	
vCenter Datacenter Name	The name of the vCenter datacenter where the HyperFlex hosts were added.
vCenter Cluster Name	The name of the vCenter cluster where the HyperFlex hosts were added.
System Services	
DNS Server(s)	A comma-separated list of IP addresses for each DNS Server.
NTP Server(s)	A comma-separated list of IP addresses for each NTP Server. Important A highly reliable NTP server is required.
Time Zone	The local time zone for the controller VM.
Auto Support	
Enable Auto Support (Recommended)	Check to enable Auto Support. Note We highly recommend enabling Auto Support to enable sending email alerts to Cisco TAC.
Send service tickets notifications to, for example: name@company.com	Email address to receive service request notifications.

Step 9

On the **Advanced Cluster Configuration** page, complete the following fields, and click **Start**.

Name	Description
Advanced Networking	
Uplink Switch Speed	Check the radio button for 10/25GE . The MAC Address Prefix field appears. Provide the MAC Address Prefix. Note The MAC Address Prefix is used to assign unique MAC addresses to the virtual interfaces of the Cisco VIC. Ensure you select a unique range to avoid any overlap with existing network equipment.

Name	Description
Management VLAN Tag Data VLAN Tag	Enter the correct VLAN tags if you are using trunk ports. The VLAN tags must be different when using trunk mode. Enter 0 for both VLAN tags if you are using access ports. Note Do not enter 0 if you are using trunk ports.
Management vSwitch Data vSwitch	Do not change the auto-populated vSwitch name.
Advanced Configuration	
Enable Jumbo Frames on Data Network	Check to enable Jumbo Frames for 10/25G deployments.
Clean up disk partitions	Check to remove all existing data and partitions from every node in the HX storage cluster. For example, if this is not the first time installing the software on the cluster.
Optimize for VDI only deployment	Check to optimize VDI deployments. By default HyperFlex is performance optimized for Virtual Server Infrastructure (VSI). Check this box to tune the performance parameters for VDI deployments. This option has no affect on all-flash HX models and only needs to be enabled for hybrid HX clusters. If you are running mixed VDI and VSI workloads, do not select this option.
vCenter Single-Sign-On Server	Fill in this field only if instructed by Cisco TAC.

Review the progress for any of the cluster configuration tasks on the **Progress** page. The deployment can take anywhere between 20-45 minutes to complete.

Step 10

After deployment finishes, the **Summary Deployment** page displays a summary of the deployment details.

What to do next

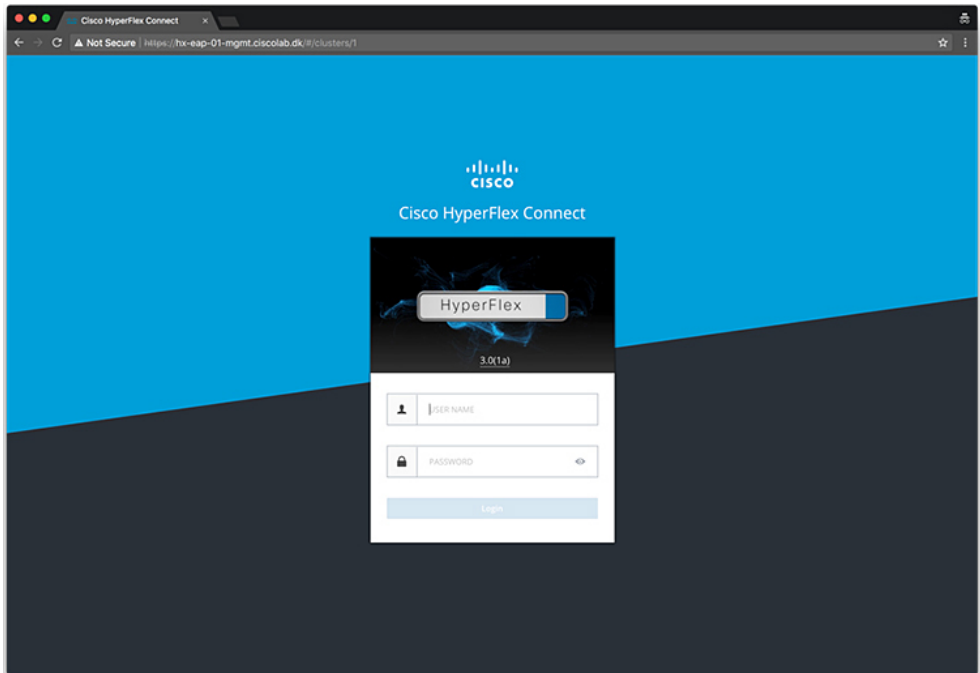
Confirm HX Data Platform Plug-in installation.

Logging into HX Connect

Cisco HyperFlex Connect provides an HTML5 based access to HX Storage monitoring, and replication, encryption, datastore, and virtual machine tasks. This procedure provides a summary on launching and logging into HX Connect. For the detailed procedure on logging in to HX Connect, see the [Cisco HyperFlex Data Platform Administration Guide](#).

Step 1

Launch HX Connect UI from a browser of your choice from `https://Cluster_IP/` or `https://FQDN`.



Step 2 Log in with the following credentials:

- **Username**—**admin**
- **Password**—Use the password set during cluster installation.

What to do next

Run the post install script before the HyperFlex cluster is ready for production use. Depending on whether you are running in a 1GE or 10/25GE switch configuration, see:

(1GE Only) Run Post Installation Script

Step 1 In your web browser, navigate to `http://<installer VM IP address>/mssh`, and log in using username **admin** and your password, and run `hx_post_install`.

Step 2 Press **Enter** to start post installation tasks in the web-based SSH window.

On execution of the post installation script, choose one of the following options as per the requirement:

- **1**—To run the post installation script on a newly created cluster or on an existing cluster. On selection of this option, the script runs the post installation operations on all the nodes in the cluster.
- **2**—To run the post installation script on expanded nodes or on newly added nodes after executing the expansion workflow. On selection of this option, the script fetches the list of the expanded nodes and runs the post installation operations on the expanded nodes.

- **3**—To have unique SSL certificate in the cluster. On selection of this option, the current certificate is replaced with the newly created SSL certificate. This option is not required for the cluster expansion.

Step 3 Follow the on-screen prompts to complete the installation.
The `hx_post_install` script completes the following:

- License the vCenter host.
- Enable HA/DRS on the cluster per best practices.
- Suppress SSH/Shell warnings in vCenter.
- Configure vMotion per best practices.
- Add additional guest VLANs/portgroups.
- Perform HyperFlex Edge configuration check.

On successful completion of the `post_install` workflow, the summary of the configuration executed based on the chosen option is displayed under **Cluster Summary**.

Sample Post-Install Script is as follows:

```
Select post_install workflow-
```

1. New/Existing Cluster
2. Expanded Cluster (for non-edge clusters)
3. Generate Certificate

Note: Workflow No.3 is mandatory to have unique SSL certificate in the cluster. By Generating this certificate, it will replace your current certificate. If you're performing cluster expansion, then this option is not required.

```
Selection: 1
Logging in to controller localhost
HX CVM admin password:
Getting ESX hosts from HX cluster...
vCenter URL: 10.121.48.111
Enter vCenter username (user@domain): administrator@vsphere.local
vCenter Password:
Found datacenter ucs659_dc
Found cluster ucs659-hx-cluster
```

```
post_install to be run for the following hosts:
ucs659.eng.storvisor.com
ucs660.eng.storvisor.com
```

```
Enter ESX root password:
HX Edge configuration detected
Uplink speed is detected as: 1G
Uplink count is detected as: 2
```

```
Enter vSphere license key? (y/n) n
```

```
Enable HA/DRS on cluster? (y/n) y
Successfully completed configuring cluster HA.
Successfully completed configuring cluster DRS.
```

```
Disable SSH warning? (y/n) y
```

```
Add vmotion interfaces? (y/n) y
```

```
Netmask for vMotion: 255.255.240.0
vMotion IP for ucs659.eng.storvisor.com: 10.64.73.131
Adding vmotion to ucs659.eng.storvisor.com
Adding vmkernel to ucs659.eng.storvisor.com
Updating portgroup vmotion on ucs659.eng.storvisor.com
Successfully updated portgroup vmotion on host ucs659.eng.storvisor.com.  activeNic: vmnic0 standbyNic:
vmnic1
vMotion IP for ucs660.eng.storvisor.com: 10.64.73.132
Adding vmotion to ucs660.eng.storvisor.com
Adding vmkernel to ucs660.eng.storvisor.com
Updating portgroup vmotion on ucs660.eng.storvisor.com
Successfully updated portgroup vmotion on host ucs660.eng.storvisor.com.  activeNic: vmnic0 standbyNic:
vmnic

Add VM network VLANs? (y/n) y

Run health check? (y/n) y

Validating cluster health and configuration...

Cluster Summary:
  Version - 4.0.2f-35930
  Model - HX220C-M5SX
  Health - HEALTHY
  ASUP enabled - False
```

(10/25GE Only) Run Post Installation Script

Step 1 In your web browser, navigate to `http://<installer VM IP address>/mssh`, and log in using username **admin** and your password, and run `hx_post_install`.

Step 2 Press **Enter** to start post installation tasks in the web-based SSH window.

On execution of the post installation script, choose one of the following options as per the requirement:

- **1**—To run the post installation script on a newly created cluster or on an existing cluster. On selection of this option, the script runs the post installation operations on all the nodes in the cluster.
- **2**—To run the post installation script on expanded nodes or on newly added nodes after executing the expansion workflow. On selection of this option, the script fetches the list of the expanded nodes and runs the post installation operations on the expanded nodes.
- **3**—To have unique SSL certificate in the cluster. On selection of this option, the current certificate is replaced with the newly created SSL certificate. This option is not required for the cluster expansion.

Step 3 Follow the on-screen prompts to complete the installation. The `hx_post_install` script completes the following:

- License the vCenter host.
- Enable HA/DRS on the cluster per best practices.
- Remove SSH/Shell warnings in vCenter.
- Configure vMotion per best practices.

- Add new VM portgroups
- Perform HyperFlex Edge health check.

On successful completion of the post_install workflow, the summary of the configuration executed based on the chosen option is displayed under **Cluster Summary**.

Sample Post-Install Script is as follows:

```
Select post_install workflow-
```

1. New/Existing Cluster
2. Expanded Cluster (for non-edge clusters)
3. Generate Certificate

Note: Workflow No.3 is mandatory to have unique SSL certificate in the cluster. By Generating this certificate, it will replace your current certificate. If you're performing cluster expansion, then this option is not required.

```
Selection: 1
Cluster IP/FQDN : 10.1.22.13
HX CVM admin password:
Getting ESX hosts from HX cluster...
vCenter URL: 10.1.22.150
Enter vCenter username (user@domain): administrator@vsphere.local
vCenter Password:
Found datacenter spiderman
Found cluster spiderman
```

```
post_install to be run for the following hosts:
hx-node-1.spiderman.hx.local
hx-node-2.spiderman.hx.local
```

```
Enter ESX root password:
HX Edge configuration detected
  Uplink speed is detected as: 10G
  Uplink count is detected as: 2
```

```
Enter vSphere license key? (y/n) n
```

```
Enable HA/DRS on cluster? (y/n) y
Successfully completed configuring cluster HA.
```

```
Disable SSH warning? (y/n) y
```

```
Add vmotion interfaces? (y/n) y
Netmask for vMotion: 255.255.255.0
VLAN ID: (0-4096) 2032
vMotion MTU is set to use jumbo frames (9000 bytes). Do you want to change to 1500 bytes? (y/n) n
Do you wish to enter the range of vMotion IPs (y/n) y
Please enter vMotion Ip range (format: IP_start-IP_end) 10.20.32.16-10.20.32.17
Vmotion ip 10.20.32.16 used for hx-node-1.spiderman.hx.local
Adding vmkernel to hx-node-1.spiderman.hx.local
Vmotion ip 10.20.32.17 used for hx-node-2.spiderman.hx.local
Adding vmkernel to hx-node-2.spiderman.hx.local
```

```
Add VM network VLANs? (y/n) y
  Port Group Name to add (VLAN ID will be appended to the name in ESXi host): infra
  VLAN ID: (0-4096) 199
  Adding infra-199 to hx-node-1.spiderman.hx.local
  Adding infra-199 to hx-node-2.spiderman.hx.local
Add additional VM network VLANs? (y/n) n
```



```
Run health check? (y/n) y

Validating cluster health and configuration...

Cluster Summary:
Version - 5.0.2a-41212
Model - HXAF220C-M5SX
Health - HEALTHY
ASUP enabled - False
hxshell:~$
```

Configuring vMotion Automatically

The `hx_post_install` script automatically configures vMotion based on network topology.

1GE Single Switch Considerations

- Automated configuration supports only trunk ports and only configurations using a dedicated vMotion VLAN.
- If using access ports or using a shared vMotion VLAN, you have to manually configure vMotion on the existing management VMkernel port (vmk0).
- vMotion is shared on the 1GE management and VM network uplink.
- A new VMKernel port (vmk2) is created with a default 500Mbps traffic shaper to ensure vMotion doesn't fully saturate the link. This default value may be changed after running `hx_post_install`. See [Configuring Traffic Shaping Manually, on page 74](#).

1GE Dual Switch Considerations

- vMotion is configured on a dedicated 1GE uplink.
- A new VMKernel port (vmk2) is created. Failover order is auto-configured such that storage data and vMotion are separated under normal network conditions.
- No traffic shaper is required in this configuration.

10/25GE Switch Considerations

- vMotion is configured on dedicated vMotion vSwitch with dedicated active/standby vNICs.
- A new VMKernel port (vmk2) is created. Failover order is auto-configured such that storage data and vMotion are separated under normal network conditions.
- No traffic shaper is required in this configuration, although bandwidth is shared among management, vMotion, and guest VM port groups. You may apply an optional traffic shaper depending on your networking requirements.

Configuring vMotion Manually

vMotion can be configured in a number of different ways depending on environmental needs. This task covers one possible configuration and variations to this procedure are expected and permitted.

This configuration leverages a unique VLAN for vMotion that is trunked across port 1.

-
- Step 1** Launch vSphere, and log into the vCenter Server as an administrator.
 - Step 2** From the **vCenter Inventory Lists**, click the HyperFlex host, and navigate to **Manage > Networking > Virtual Switches**.
 - Step 3** Click **Add Host Networking**.
 - Step 4** On the **Add Network Wizard: Connection Type** page, click **VMkernel**, and click **Next**.
 - Step 5** Click **Use vswitch-hx-inband-mgmt**, and click **Next**.
 - Step 6** Enter a distinctive **Network Label**, such as **vMotion**, enter the correct **VLAN ID**, check **Use this port group for vMotion**, and click **Next**.
 - Step 7** Click **Use the following IP settings**, enter a static IPv4 address and **Subnet Mask**, and click **Next**.
 - Step 8** Review the settings, and click **Finish**.
 - Step 9** Repeat this procedure for all HyperFlex hosts and compute-only hosts in the HyperFlex storage cluster.
-

Configuring Traffic Shaping Manually

For 1GE single switch deployments, it is a best practice to enable traffic shaping on the vMotion interface to prevent network congestion on shared uplinks. Failure to configure a traffic shaper could result in vMotion traffic starving management and VM guest traffic sharing the same physical 1GE port. vMotion standard switches only allow for egress traffic shaping.

-
- Step 1** Launch vSphere, and log into the vCenter Server as an administrator.
 - Step 2** From the **vCenter Inventory Lists**, click the HyperFlex host, and navigate to **Manage > Networking > Virtual Switches**.
 - Step 3** Select the vSwitch that contains the vMotion portgroup.
 - Step 4** Click the vMotion portgroup name, and click **Edit Settings** (pencil icon).
 - Step 5** On the left menu, select **Traffic shaping**.
 - Step 6** Check the override box to enable the traffic shaper.
 - Step 7** Set the average and peak bandwidth to meet environmental needs. One possible value to use is 500,000 Kbits/sec for both, representing 50% of total bandwidth available on a 1GE uplink.
 - Step 8** Select **OK** to save settings.

Note Be careful to set average bandwidth to the desired setting. Peak bandwidth works only for bursting traffic and is quickly exhausted for vMotion operations.

(10/25GE) Using Additional VIC Ports (Optional)

In cases where uplinks to different switches (for example, different VLANs), or additional dedicated bandwidth is needed for guest VMs, you can connect ports 3 and 4 from the VIC after HX is installed. This section describes the configuration created by default and how you can create additional vNICs on the additional ports.

Default 10GE VIC Configuration:

During installation, HyperFlex configures the VIC 1457 as follows:

- Disables the port channel
- Configures the 8 vNICs that HyperFlex needs to operate (note the Uplink Port # set to 0 or 1 accordingly which correspond to the first two ports of the VIC)

The screenshot shows the vNIC configuration interface with a table of Host Ethernet Interfaces. The table has columns for Name, CON, MAC Address, MTU, vsNIC, Uplink Port, CoS, VLAN, VLAN Mode, iSCSI Boot, PXE Boot, Channel, Port Profile, and Uplink Fallover. The Uplink Port column is highlighted in yellow, showing values 0 and 1 for the vNICs.

Name	CON	MAC Address	MTU	vsNIC	Uplink Port	CoS	VLAN	VLAN Mode	iSCSI Boot	PXE Boot	Channel	Port Profile	Uplink Fallover
<input type="checkbox"/> eth0	hv-mgmt-a	00:25:B5:50:A1:01	9000	0	0	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/> eth1	hv-mgmt-b	00:25:B5:50:B2:01	9000	0	1	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/> eth2	storage-d...	00:25:B5:50:A3:01	9000	0	0	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/> eth3	storage-d...	00:25:B5:50:B4:01	9000	0	1	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/> eth4	vm-neteo...	00:25:B5:50:A5:01	9000	0	0	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/> eth5	vm-neteo...	00:25:B5:50:B6:01	9000	0	1	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/> eth6	hv-vmotio...	00:25:B5:50:A7:01	9000	0	0	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/> eth7	hv-vmotio...	00:25:B5:50:B8:01	9000	0	1	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A

Sample Network Configurations

1GE Single Switch

Nexus 5548 using trunk ports

```

vlan 101
  name HX-MGMT
vlan 102
  name HX-STORAGE
vlan 103
  name HX-VMOTION
vlan 104
  name HX-GUESTVM
...
interface Ethernet2/11
  description HX-01-Port1
  switchport mode trunk
  switchport trunk allowed vlan 101-104
  spanning-tree port type edge trunk
  speed 1000
interface Ethernet2/12

```

```

    description HX-01-Port2
    switchport mode trunk
    switchport trunk allowed vlan 101-104
    spanning-tree port type edge trunk
    speed 1000
interface Ethernet2/13
    description HX-02-Port1
    switchport mode trunk
    switchport trunk allowed vlan 101-104
    spanning-tree port type edge trunk
    speed 1000
interface Ethernet2/14
    description HX-02-Port2
    switchport mode trunk
    switchport trunk allowed vlan 101-104
    spanning-tree port type edge trunk
    speed 1000
interface Ethernet2/15
    description HX-03-Port1
    switchport mode trunk
    switchport trunk allowed vlan 101-104
    spanning-tree port type edge trunk
    speed 1000
interface Ethernet2/16
    description HX-03-Port2
    switchport mode trunk
    switchport trunk allowed vlan 101-104
    spanning-tree port type edge trunk
    speed 1000

```

Catalyst 3850-48T using trunk ports

```

vlan 101
    name HX-MGMT
vlan 102
    name HX-STORAGE
vlan 103
    name HX-vMOTION
vlan 104
    name HX-GUESTVM
...
interface GigabitEthernet1/0/1
    description HX-01-Port1
    switchport trunk allowed vlan 101-104
    switchport mode trunk
    speed 1000
    spanning-tree portfast trunk
interface GigabitEthernet1/0/2
    description HX-01-Port2
    switchport trunk allowed vlan 101-104
    switchport mode trunk
    speed 1000
    spanning-tree portfast trunk
interface GigabitEthernet1/0/3
    description HX-02-Port1
    switchport trunk allowed vlan 101-104
    switchport mode trunk
    speed 1000
    spanning-tree portfast trunk
interface GigabitEthernet1/0/4
    description HX-02-Port2
    switchport trunk allowed vlan 101-104
    switchport mode trunk
    speed 1000

```

```

spanning-tree portfast trunk
interface GigabitEthernet1/0/5
  description HX-03-Port1
  switchport trunk allowed vlan 101-104
  switchport mode trunk
  speed 1000
  spanning-tree portfast trunk
interface GigabitEthernet1/0/6
  description HX-03-Port2
  switchport trunk allowed vlan 101-104
  switchport mode trunk
  speed 1000
  spanning-tree portfast trunk

```

1GE Dual Switch

Nexus 5548 using trunk ports

This configuration uses DHCP with in-band management using native vlan 105. This switch connects to both 1GE LOMs and uses dhcp relay.

```

ip dhcp relay
...
interface Vlan105
  ip address 10.1.2.1/24
  ip dhcp relay address 10.1.1.2
  no shutdown
vlan 101
  name HX-MGMT
vlan 102
  name HX-STORAGE
vlan 103
  name HX-vmOTION
vlan 104
  name HX-GUESTVM
vlan 105
  name HX-DHCP-CIMC
...
interface Ethernet2/11
  description HX-01-Port1
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk
  speed 1000
interface Ethernet2/12
  description HX-01-Port2
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk
  speed 1000
interface Ethernet2/13
  description HX-02-Port1
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk
  speed 1000
interface Ethernet2/14
  description HX-02-Port2
  switchport mode trunk
  switchport trunk native vlan 105

```

```

        switchport trunk allowed vlan 101-105
        spanning-tree port type edge trunk
        speed 1000
interface Ethernet2/15
    description HX-03-Port1
    switchport mode trunk
    switchport trunk native vlan 105
    switchport trunk allowed vlan 101-105
    spanning-tree port type edge trunk
    speed 1000
interface Ethernet2/16
    description HX-03-Port2
    switchport mode trunk
    switchport trunk native vlan 105
    switchport trunk allowed vlan 101-105
    spanning-tree port type edge trunk
    speed 1000

```

Repeat the same configuration on switch #2. Eliminate the `dhcp relay` and `interface Vlan 105` commands.

Catalyst 3850-48T using trunk ports

This configuration uses statically-assigned CIMC IPs on `vlan 105`. All `vlan`s are allowed on all trunk interfaces. For security purposes, we recommend restricting the VLANs to those required for a HyperFlex deployment by adding the `switchport trunk allowed vlan` statement into all your port configurations.

```

vlan 101
    name HX-MGMT
vlan 102
    name HX-STORAGE
vlan 103
    name HX-VMOTION
vlan 104
    name HX-GUESTVM
vlan 105
    name HX-CIMC
...
interface GigabitEthernet1/0/1
    description HX-01-Port1
    switchport mode trunk
    speed 1000
    spanning-tree portfast trunk
interface GigabitEthernet1/0/2
    description HX-01-Port2
    switchport mode trunk
    speed 1000
    spanning-tree portfast trunk
interface GigabitEthernet1/0/3
    description HX-02-Port1
    switchport mode trunk
    speed 1000
    spanning-tree portfast trunk
interface GigabitEthernet1/0/4
    description HX-02-Port2
    switchport mode trunk
    speed 1000
    spanning-tree portfast trunk
interface GigabitEthernet1/0/5
    description HX-03-Port1
    switchport mode trunk
    speed 1000
    spanning-tree portfast trunk
interface GigabitEthernet1/0/6
    description HX-03-Port2

```

```
switchport mode trunk
speed 1000
spanning-tree portfast trunk
```

Repeat the same configuration on switch #2.

10GE Dual Switch

Nexus 9000 using trunk ports

```
vlan 101
  name HX-MGMT
vlan 102
  name HX-STORAGE
vlan 103
  name HX-VMOTION
vlan 104
  name HX-GUESTVM
vlan 105
  name HX-DHCP-CIMC
...
interface Ethernet1/35
  description M5-Edge-Node1-VIC1
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk

interface Ethernet1/36
  description M5-Edge-Node1-VIC2
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk

interface Ethernet1/37
  description M5-Edge-Node2-VIC1
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk

interface Ethernet1/38
  description M5-Edge-Node2-VIC2
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk

interface Ethernet1/39
  description M5-Edge-Node3-VIC1
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk

interface Ethernet1/40
  description M5-Edge-Node3-VIC2
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk
```

Catalyst 9300 using trunk ports

```

vlan 101
  name HX-MGMT
vlan 102
  name HX-STORAGE
vlan 103
  name HX-vMOTION
vlan 104
  name HX-GUESTVM
vlan 105
  name HX-CIMC
...
interface GigabitEthernet1/0/1
  description M5-Edge-16W9-LOM1
  switchport trunk allowed vlan 101-105
  switchport mode trunk
  spanning-tree portfast trunk

interface GigabitEthernet1/0/2
  description M5-Edge-16W9-LOM2
  switchport trunk allowed vlan 101-105
  switchport mode trunk
  spanning-tree portfast trunk

interface GigabitEthernet1/0/3
  description M5-Edge-16UQ-LOM1
  switchport trunk allowed vlan 101-105
  switchport mode trunk
  spanning-tree portfast trunk

interface GigabitEthernet1/0/4
  description M5-Edge-16UQ-LOM2
  switchport trunk allowed vlan 101-105
  switchport mode trunk
  spanning-tree portfast trunk

interface GigabitEthernet1/0/5
  description M5-Edge-05G9-LOM1
  switchport trunk allowed vlan 101-105
  switchport mode trunk
  spanning-tree portfast trunk

interface GigabitEthernet1/0/6
  description M5-Edge-05G9-LOM2
  switchport trunk allowed vlan 101-105
  switchport mode trunk
  spanning-tree portfast trunk

```

10/25GE 2-Node 2-Room**Catalyst 9300 with QoS**

This configuration uses quality of service to mark and prioritize HyperFlex storage traffic using the 10 or 25 Gigabit Ethernet Stacked Switches Per Room Topology

```

class-map match-all PQ_Storage
  match dscp ef
class-map match-all Storage
  match access-group name Storage
...
policy-map Storage_Mark

```



```
class Storage
  set dscp ef
class class-default
policy-map Storage_Queue
class PQ_Storage
  priority level 1
  queue-buffers ratio 80
class class-default
  bandwidth remaining percent 100
  queue-buffers ratio 20
...
interface Port-channel98
  switchport trunk allowed vlan 101,102,103,104,105
  switchport mode trunk
!
interface GigabitEthernet1/0/3
  description SERVER1-Dedicated-CIMC
  switchport access vlan 145
  switchport mode access
  spanning-tree portfast
!
interface TenGigabitEthernet1/1/1
  description SERVER1-VIC-1
  switchport trunk allowed vlan 101,102,103,104,105
  switchport mode trunk
  spanning-tree portfast trunk
  service-policy input Storage_Mark
!
interface TenGigabitEthernet2/1/1
  description SERVER1-VIC-2
  switchport trunk allowed vlan 101,102,103,104,105
  switchport mode trunk
  spanning-tree portfast trunk
  service-policy input Storage_Mark
!

interface TenGigabitEthernet1/1/8
  description cross-connect-01
  switchport trunk allowed vlan 101,102,103,104,105
  switchport mode trunk
  channel-group 98 mode on
  service-policy output Storage_Queue
!
interface TenGigabitEthernet2/1/8
  description cross-connect-02
  switchport trunk allowed vlan 101,102,103,104,105
  switchport mode trunk
  shutdown
  channel-group 98 mode on
  service-policy output Storage_Queue
!
...
ip access-list extended Storage
 10 permit ip 169.254.1.0 0.0.0.255 169.254.1.0 0.0.0.255
```

Repeat the same configuration on switch stack #2.

