



Cisco HyperFlex Edge Deployment Guide, Release 4.0

First Published: 2019-04-29

Last Modified: 2021-07-26

Americas Headquarters

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2019–2021 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE	Communications, Services, Bias-free Language, and Additional Information	v
----------------	---	----------

CHAPTER 1	Technology Use Case	1
	Cisco HyperFlex Edge	1

CHAPTER 2	Deployment Options	3
	HyperFlex Edge Deployment Options	3

CHAPTER 3	Preinstallation Checklist	5
	2-Node Network Topology	5
	Selecting your 2-Node Network Topology	5
	10 or 25 Gigabit Ethernet Topology	5
	1 Gigabit Ethernet Topology	12
	Common Network Requirement Checklist	18
	Guest VM Traffic	23
	Intersight Connectivity	23
	Cisco HyperFlex Edge Invisible Cloud Witness	24
	Ordering Cisco Hyperflex Edge Servers	24
	3- and 4-Node Network Topology	24
	Selecting your 3- or 4-Node Network Topology	24
	10 or 25 Gigabit Ethernet Topology	25
	1 Gigabit Ethernet Topology	32
	HyperFlex Edge Compatibility and Software Requirements: HyperFlex Release 4.0(x)	40
	Common Network Requirement Checklist	40
	Guest VM Traffic	46
	Intersight Connectivity	46

Cisco HyperFlex Edge Invisible Cloud Witness 47
 Ordering Cisco Hyperflex Edge Servers 47

CHAPTER 4

Installation 49

Installation Overview 49
 Applicable for 1GE and 10/25GE Switch Configurations 51
 Rack Cisco HyperFlex Nodes 51
 Cisco Integrated Management Controller Configuration 51
 Configuring CIMC: Static Assignment 51
 Configuring CIMC: DHCP Assignment 52
 Verifying Firmware Versions 53
 Deploying Cisco HX Data Platform Installer 53
 Applicable for 1GE Switch Configurations 54
 Configuring Your HyperFlex Cluster 54
 Verifying Cisco HX Data Platform Software Installation 60
 Applicable for 10/25GE Switch Configurations 60
 (10/25GE only) Configuring Your HyperFlex Cluster 60
 Logging into HX Connect 63

CHAPTER 5

Post Installation Tasks 65

(1GE Only) Run Post Installation Script 65
 (10/25GE Only) Run Post Installation Script 67
 Configuring vMotion Automatically 69
 Configuring vMotion Manually 70
 Configuring Traffic Shaping Manually 70
 (10/25GE) Using Additional VIC Ports (Optional) 71
 Creating Additional VIC Ports 71

CHAPTER 6

Appendix 73

Sample Network Configuration 10GE Dual Switch 73
 Sample Network Configuration 1GE Single Switch 75
 Sample Network Configuration 1GE Dual Switch 76
 Sample Network Configuration 10/25GE 2-Node 2-Room 78



Communications, Services, Bias-free Language, and Additional Information

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you're looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

Documentation Feedback

To provide feedback about Cisco technical documentation, use the feedback form available in the right pane of every online document.

Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

Bias-Free Language

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.



CHAPTER 1

Technology Use Case

- [Cisco HyperFlex Edge, on page 1](#)

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.

Limitations and Supportability Summary

Limitation for	Support
Cluster size and type	<p>2-Node clusters:</p> <ul style="list-style-type: none"> • HX220c M5 Hybrid/HXAF220c M5 All-Flash • HX240c M5SD Hybrid/HX240c M5SD All-Flash¹ <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 M5 Hybrid/HXAF220c M5 All-Flash • HX220c M4 Hybrid/HXAF220c M4 All-Flash • HX240c M5SD Hybrid/HX240c M5SD All-Flash² <p>Note HX 240 Edge (short depth) requires Intersight for initial deployment and ongoing management.</p> <p>4-Node clusters:</p> <ul style="list-style-type: none"> • HX220c M5 Hybrid/HXAF220c M5 All-Flash • HX240c M5SD Hybrid/HX240c M5SD All-Flash³ <p>Note HX 240 Edge (short depth) requires Intersight for initial deployment and ongoing management.</p>
Replication Factor	<p>2</p> <p>Note A reliable backup strategy is strongly recommended to ensure that production data is adequately protected.</p>
Networking	<p>1GE or 10/25GE networking without Cisco UCS Fabric Interconnects.</p> <p>HX Edge Systems do not implement QoS.</p>
HX clusters per vCenter	Up to 100

¹ Supported in Cisco HyperFlex Release 4.0(2c) and later.

² Supported in Cisco HyperFlex Release 4.0(2c) and later.

³ Supported in Cisco HyperFlex Release 4.0(2c) and later.



CHAPTER 2

Deployment Options

- [HyperFlex Edge Deployment Options, on page 3](#)

HyperFlex Edge Deployment Options

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

- **HyperFlex On-Premise OVA Installer**—Use this option for on-premise 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 support 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.

This guide covers deployment using the on-premise 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.



CHAPTER 3

Preinstallation Checklist

- [2-Node Network Topology, on page 5](#)
- [3- and 4-Node Network Topology, on page 24](#)

2-Node Network Topology

Selecting your 2-Node Network Topology

Cisco HyperFlex Edge offers both a 1 Gigabit Ethernet (GE) and a 10/25GE installation option. Both topologies support single top-of-rack (ToR) and dual ToR switch options for ultimate network flexibility and redundancy. Consider the following when determining the best topology for your cluster:

- Cisco highly recommends the 10/25GE topology for higher performance and future node expansion capabilities.
- The 1GE topology is reserved for clusters that will never require node expansion, and instances where the ToR switch does not have 10GE ports available.



Note A network topology is chosen during initial deployment and cannot be changed or upgraded without a full reinstallation. Choose your network topology carefully and with future needs in mind.

To get started, select your network topology:

- [10 or 25 Gigabit Ethernet Topology, on page 5](#)
- [1 Gigabit Ethernet Topology, on page 12](#)

After completing the 10/25GE or 1GE ToR physical network and cabling section, continue with the [Common Network Requirement Checklist, on page 18](#).

10 or 25 Gigabit Ethernet 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 10/25 Gigabit Ethernet 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 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.

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

10/25 Gigabit Ethernet 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):

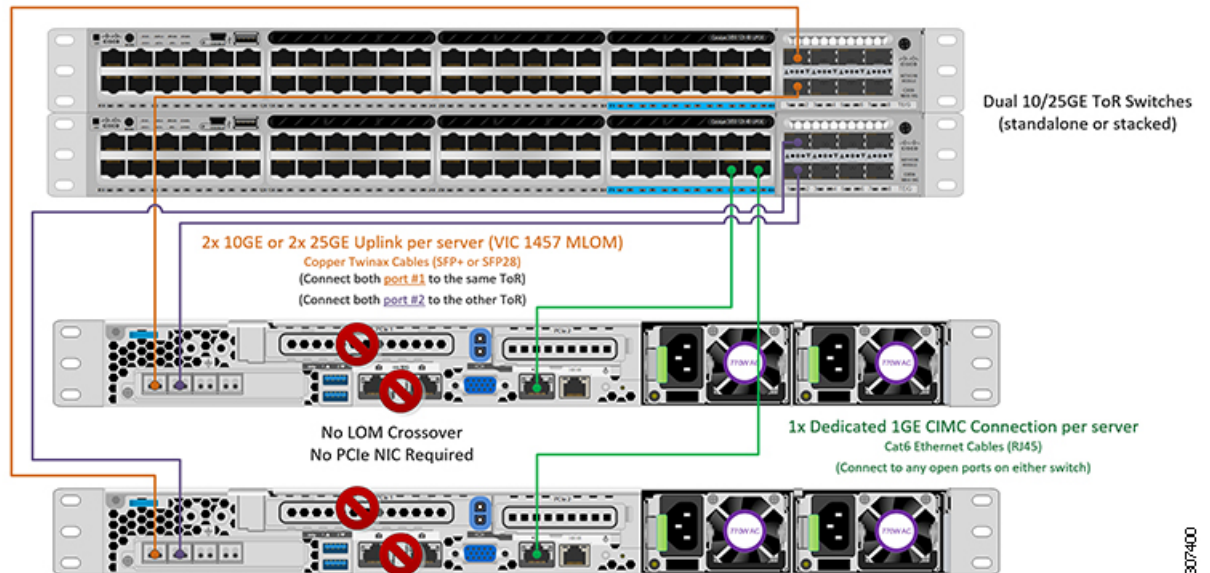
- If using dedicated Cisco IMC, connect the 1GE management port on each server (Labeled M on the back of the server) to one of the two switches.
- Connect one out of the four 10/25GE ports on the Cisco VIC from each server to the same ToR switch.
 - Use the same port number on each server to connect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

2 Node Edge – Dual 10/25GE ToR Switches



307-400

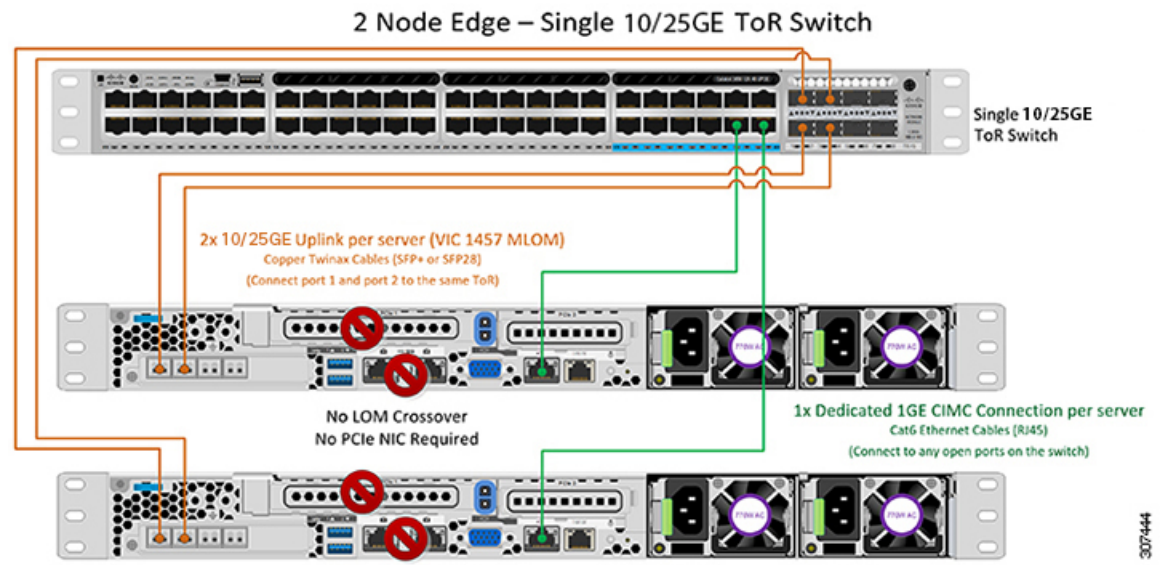
10/25 Gigabit Ethernet Single Switch Physical Cabling



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 10/25 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:

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.

10/25 Gigabit Ethernet 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 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.

Jumbo Frames for 10/25 Gigabit Ethernet

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 18.

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



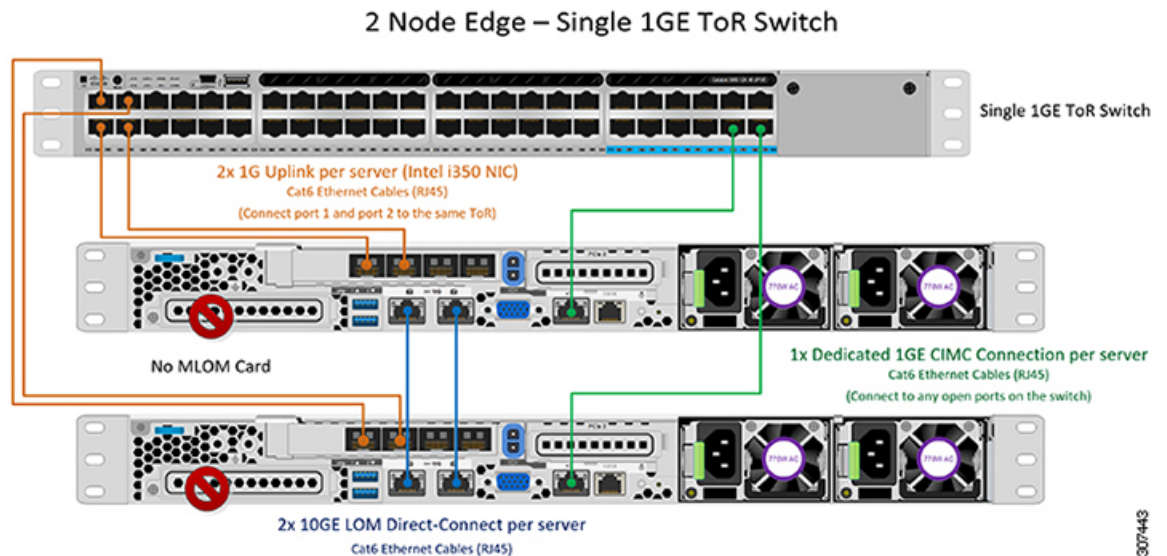
1 Gigabit Ethernet Single Switch Cabling



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

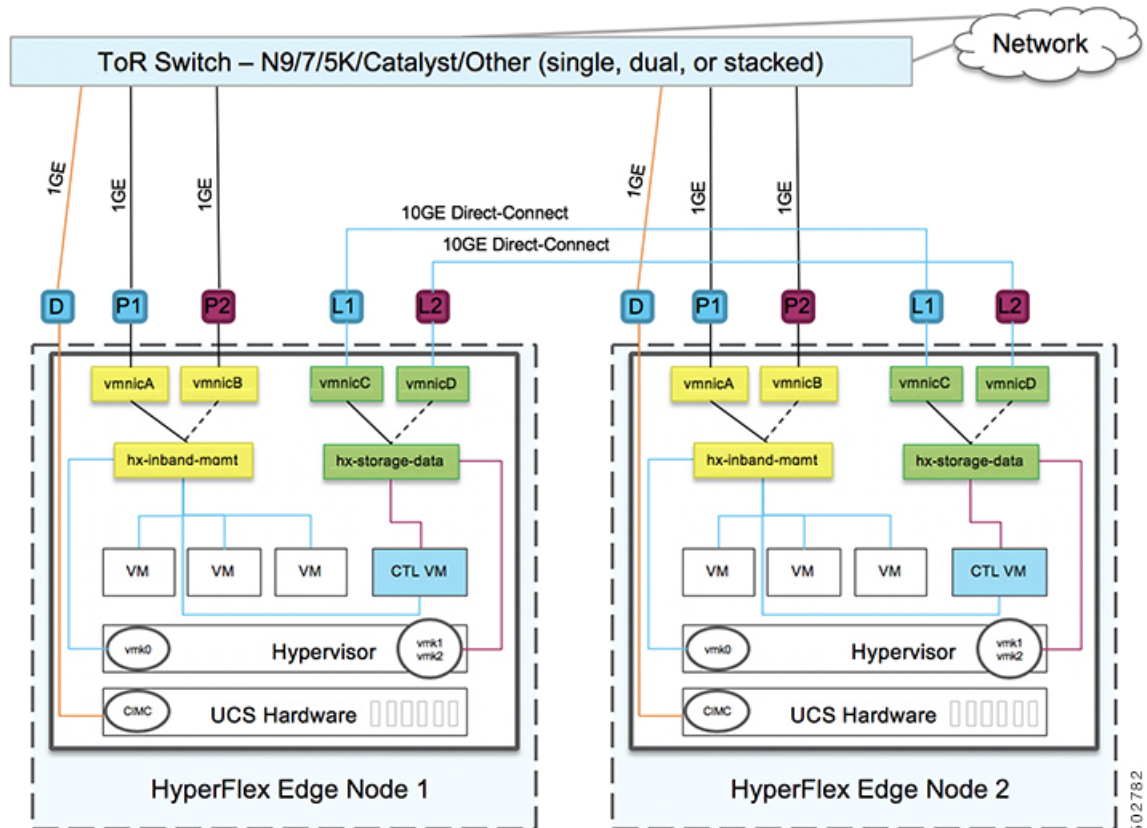
The minimum network configuration requires at least two separate networks:

- 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 18](#).

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
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 23.

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](#).

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

Connected Services

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

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.

3- and 4-Node Network Topology

Selecting your 3- or 4-Node Network Topology

Cisco HyperFlex Edge 3-Node and 4-Node clusters may be deployed either through Cisco Intersight, or an on-premises installer VM. Cisco Intersight provides advanced multi-cluster monitoring and management capabilities; the Cisco Intersight HyperFlex Installer is the recommended installer, and should be used whenever possible.

Cisco HyperFlex Edge offers both a 1 Gigabit Ethernet (GE) and a 10/25GE installation option. Both topologies support single top-of-rack (ToR) and dual ToR switch options for ultimate network flexibility and redundancy. Consider the following when determining the best topology for your cluster:

- Cisco highly recommends the 10/25GE topology for higher performance and future node expansion capabilities.
- The 1GE topology is reserved for clusters that will never require node expansion, and instances where the ToR switch does not have 10GE ports available.



Note A network topology is chosen during initial deployment and cannot be changed or upgraded without a full reinstallation. Choose your network topology carefully and with future needs in mind.

To get started, select your network topology:

- [10 or 25 Gigabit Ethernet Topology, on page 25](#)
- [1 Gigabit Ethernet Topology, on page 32](#)

After completing the 10/25GE or 1GE ToR physical network and cabling section below, continue with the [Common Network Requirement Checklist, on page 18](#).

10 or 25 Gigabit Ethernet 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 10/25 Gigabit Ethernet 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 ensure 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 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:

10/25 Gigabit Ethernet Dual Switch Physical Cabling



Warning

Proper cabling is important to ensure full network redundancy.

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 2 x 10/25GE ports, 1 x 1GE port (dedicated CIMC), and 1 x Cisco VIC 1457 MLOM card for each HyperFlex node. Trunk ports are the only supported network port configuration.

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

Upstream Network Requirements


- Two managed switches with VLAN capability (standalone or stacked)
- 2 x 10/25GE ports and 1 x 1GE port for each HyperFlex node.
All 10/25GE ports must trunk and allow all applicable VLANs. All 1GE ports may be trunked or in access mode when connected to the dedicated CIMC port.
- Jumbo frames are not required to be configured
- Portfast trunk should be configured on all ports to ensure uninterrupted access to Cisco Integrated Management Controller (CIMC)
- 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, or to an out-of-band management switch.
- 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.

3 & 4 Node Edge – Dual 10/25GE ToR Switches



10/25 Gigabit Ethernet Single Switch Physical Cabling

**Warning**


Proper cabling is important to ensure full network redundancy.

Single switch configuration provides a simple topology requiring only a single switch, and 2 x 10/25GE and 1 x 1GE port per server. Link level redundancy is provided for all HyperFlex network services. However, switch level redundancy is not provided when operating with a single ToR switch. Trunk ports are the only supported network port configuration.

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 or to an out-of-band management 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.

3 & 4 Node Edge – Single 10/25GE ToR Switch



502780

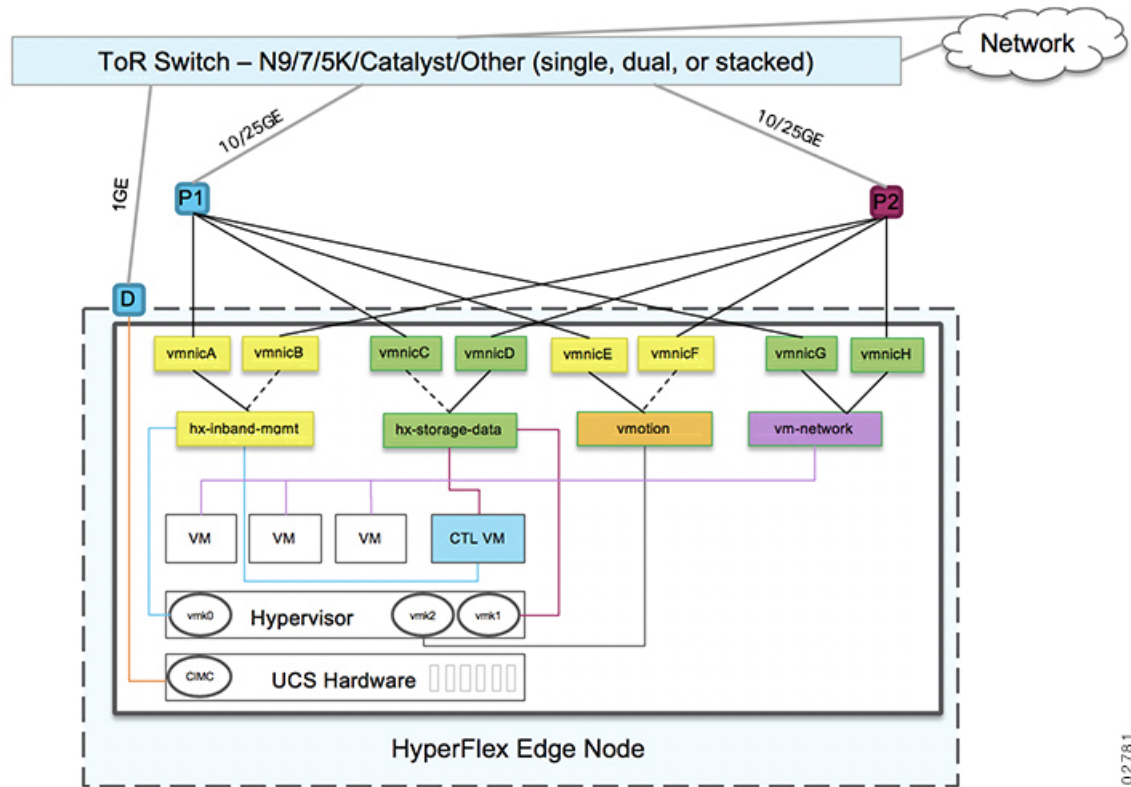
Virtual Networking Design for 3- and 4-Node 10/25 Gigabit Ethernet Topology

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

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:

502781

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.

10/25 Gigabit Ethernet 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.

Jumbo Frames for 10/25 Gigabit Ethernet

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 40](#).

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.

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)

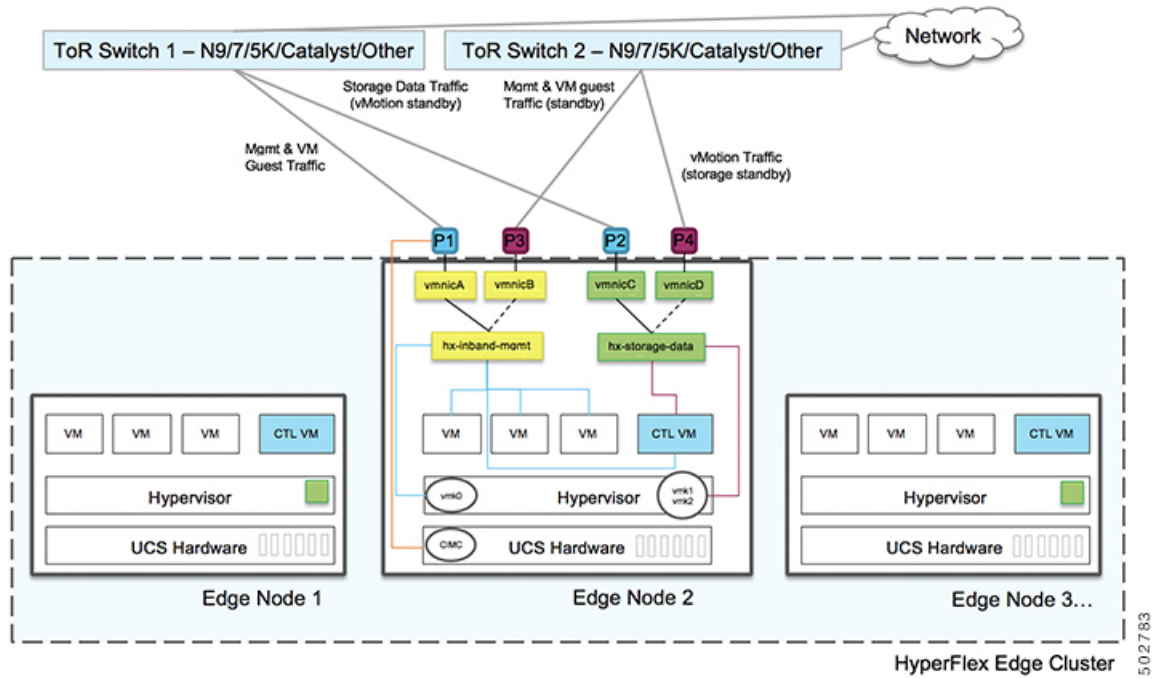
Two vSwitches each carrying different networks are required:

- **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.

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 Dual Switch Cabling

Warning Proper cabling is important to ensure full network redundancy.

The following requirements must be met before starting deployment:

- Dedicated 1 Gigabit Ethernet (GE) Cisco IMC management port per server (recommended).
 - 1 x 1GE ToR switch port and 1 x Category 6 ethernet cable for dedicated Cisco IMC management port per HyperFlex server (customer supplied)
- Intel i350 PCIe NIC [HX-PCIE-IRJ45] (installed in a PCIe slot in each server).
 - This NIC may be selected at ordering time and shipped preinstalled from the factory. The NIC may also be field-installed if ordered separately. Either riser #1 or #2 may be used, although riser #1 is recommended as it supports single socket CPU configurations.
 - 2 x 1GE ToR switch ports and 2 x Category 6 Ethernet Cables per HyperFlex server (customer supplied).
 - Cisco VIC is not used in this topology.

- Intel i350 in MLOM form factor is not supported.
- Intel x550 Lan-on-motherboard LOM (built into Cisco UCS motherboard)
 - 2 x 1GE ToR switch ports and 2 x Category 6 Ethernet Cables per HyperFlex server (customer supplied)



Note Only 1GE speed is supported for this topology. Use of 10GE LOM ports with 10GbaseT switches is not supported. Instead, set the speed manually for 1GE or use one of the supported 10GE topologies described in this guide.

To deploy with dual ToR switches for extra redundancy:

- 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 or to an out-of-band management switch.
- Connect both integrated Lan-on-motherboard (LOM) ports on all servers to the same ToR switch.



Note Redundancy occurs at the vSwitch level and includes one uplink port from the integrated LOM and one uplink port from the PCIe NIC for each vSwitch. Do not connect LOM ports to different switches.

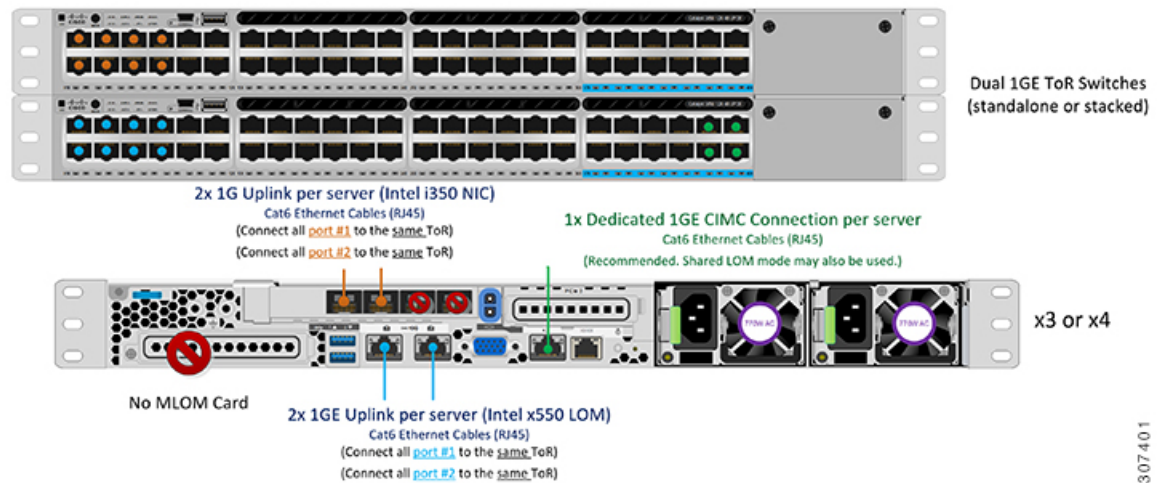
- Connect any two out of the four 1GE ports on the i350 NIC from each server to the same ToR switch. Use the same port number on each server to connect to the same switch.



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.

- Do not connect more than two 1GE ports from the i350 NIC prior to cluster installation. After cluster deployment, you may optionally use the additional two 1GE ports for guest VM traffic. see [Cisco HyperFlex Systems—Networking Topologies](#) for guidelines on using extra available NIC ports.

3 & 4 Node Edge – Dual 1GE ToR Switches



307401

1 Gigabit Ethernet Single Switch Cabling

**Warning**

Proper cabling is important to ensure full network redundancy.

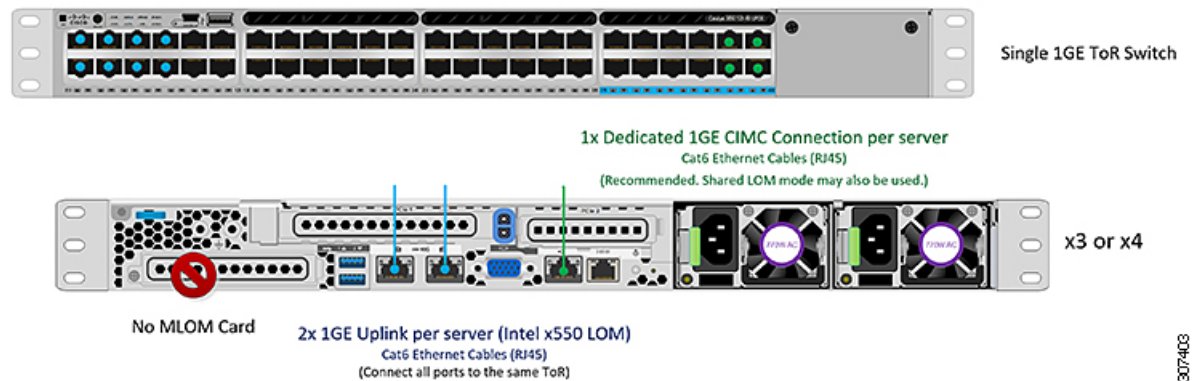
To deploy with a single ToR (see diagram below for a visual layout):

- If using dedicated Cisco IMC, connect the 1GE management port on each server (Labeled M on the back of the server) to the ToR switch or to an out-of-band management switch.
- Connect both integrated Lan-on-motherboard (LOM) ports on all servers to the same ToR switch.

**Note**

Only 1GE speed is supported for this topology. Use of 10GE LOM ports with 10GbaseT switches is not supported. Instead, set the speed manually for 1GE or use one of the supported 10GE topologies described in this guide.

3 & 4 Node Edge – Single 1GE ToR Switch

**About Access and Trunk Ports**

Ethernet interfaces can be configured either as access ports or trunk ports, as follows:

- An access port can have only one VLAN configured on the interface; it can carry traffic for only one VLAN.
- A trunk port can have one or more VLANs configured on the interface; it can carry traffic for several VLANs simultaneously.

The following table summarizes the differences between access and trunk ports. You can use the details described in this table to determine which ports to use for your deployment.

**Important**

Trunk ports are assumed in this guide, and is highly recommended for your deployment.

Trunk Ports	Access Ports
Requires more setup and definition of VLAN tags within CIMC, ESXi, and HX Data Platform Installer.	Provides a simpler deployment process than trunk ports.
Provides the ability to logically separate management, vMotion, and VM guest traffic on separate subnets.	Requires that management, vMotion, and VM guest traffic must share a single subnet.
Provides flexibility to bring in additional L2 networks to ESXi.	Requires a managed switch to configure ports 1 and 2 on discrete VLANs; storage traffic must use a dedicated VLAN, no exceptions.

**Note**

Both trunk and access ports require a managed switch to configure ports 1 and 2 on discrete VLANs.

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.

Jumbo Frames for 1 Gigabit Ethernet

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 10GE 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 40.

HyperFlex Edge Compatibility and Software Requirements: HyperFlex Release 4.0(x)

For details about compatibility and software requirements for Cisco HX Release 4.0(x), review the [Cisco HX Release 4.0\(x\) - Software Requirements](#) chapter of the Cisco HyperFlex Recommended Software Release and Requirements Guide.

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 23.

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.



CHAPTER 4

Installation

- [Installation Overview, on page 49](#)
- [Applicable for 1GE and 10/25GE Switch Configurations, on page 51](#)
- [Applicable for 1GE Switch Configurations, on page 54](#)
- [Applicable for 10/25GE Switch Configurations, on page 60](#)
- [Logging into HX Connect, on page 63](#)

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 25 Gigabit Ethernet Topology, on page 5 or, • 1 Gigabit Ethernet Topology, on page 12 <p>3- and 4-Node Edge Deployments:</p> <ul style="list-style-type: none"> • 10 or 25 Gigabit Ethernet Topology, on page 25or, • 1 Gigabit Ethernet Topology, on page 32 	1GE & 10/25GE
2	Complete installation prerequisites.	<ul style="list-style-type: none"> • Rack Cisco HyperFlex Nodes, on page 51 • Cisco Integrated Management Controller Configuration • Verifying Firmware Versions, on page 53 	1GE & 10/25GE
3	Download and deploy Cisco HX Data Platform Installer.	Deploying Cisco HX Data Platform Installer, on page 53	1GE & 10/25GE
4	Deploy Cisco 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 54 • Verifying Cisco HX Data Platform Software Installation, on page 60 	1GE only
6	Deploy Cisco HyperFlex Edge cluster.	(10/25GE only) Configuring Your HyperFlex Cluster, on page 60	10/25GE only

Applicable for 1GE and 10/25GE Switch Configurations

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 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 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 in to 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, and save the package locally.
Verify the downloaded version matches the recommended version for your deployment.
- Step 2** Log in to 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**.

Applicable for 1GE Switch Configurations

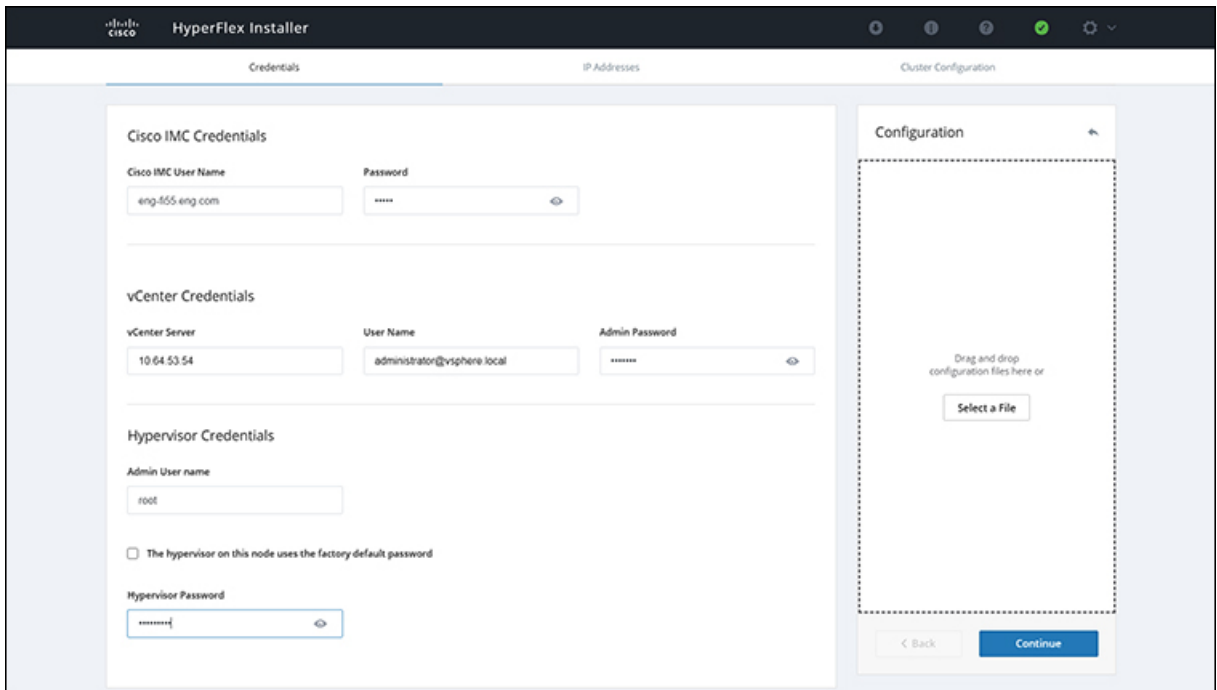
Configuring Your HyperFlex Cluster

- 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-premise 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.
- Click **Select a file** and choose your *JSON file* to load the configuration. Select **Use Configuration**.
 - 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

Name	Description
Cisco IMC User Name	Cisco IMC username. By default, the user name 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	
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.




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

Name	Description
Cisco IMC	IP Address of Cisco IMC

Name	Description
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.



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	For HXDP 3.0(x) and 4.0(x), the Replication Factor must be set to 2 for HX Edge deployments. Replication factor 3 is not supported for HX Edge deployments until HXDP release 4.5.
Controller VM	

Name	Description
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	
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	

Name	Description
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.

Cisco HX Cluster

Cluster Name: EDGE_CLUSTER
 Replication Factor: 2

Controller VM

Create Admin Password: [masked]
 Confirm Admin Password: [masked]

vCenter Configuration

vCenter Datacenter Name: datacenter
 vCenter Cluster Name: cluster

System Services

DNS Server(s): 10.64.48.21
 NTP Server(s): 10.64.1.48
 DNS Domain Name: [empty]
 Time Zone: (UTC+05:30) Chennai, Sri Jayawardenepura

Connected Services

Enable Connected Services (Recommended):
 Send service ticket notifications to: admin@cisco.com

Advanced Networking

Management VLAN Tag: 10
 Management vSwitch: vswitch-hx-inband-mgmt
 Data VLAN Tag: 11
 Data vSwitch: vswitch-hx-storage-data

Advanced Configuration

Enable Jumbo Frames on Data Network:
 Clean up disk partitions:
 Optimize for VDI only deployment:
 vCenter Single-Sign-On Server: ex. https://<address>:7444/sts/STSService

Configuration Summary (Right Sidebar):

- Credentials:** vCenter Server: 10.64.53.54, User Name: administrator@vsphere.local, Cisco IMC User Name: eng-fis5.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

Buttons: < Back, Start

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 60](#)

Verifying Cisco HX Data Platform Software Installation

Step 1 Launch **vSphere**, and log in to 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 VMware vSphere Web Client in the services page in mmc.

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

Applicable for 10/25GE Switch Configurations

(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-premise 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	

Name	Description
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.
Admin Password	Administrator password.
CIMC Credentials	
CIMC User Name	CIMC username. By default, the user name 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.

Name	Description
Replication Factor	For HXDP 3.0(x) and 4.0(x), the Replication Factor must be set to 2 for HX Edge deployments. Replication factor 3 is not supported for HX Edge deployments until HXDP release 4.5.
Controller VM	
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	

Name	Description
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.
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

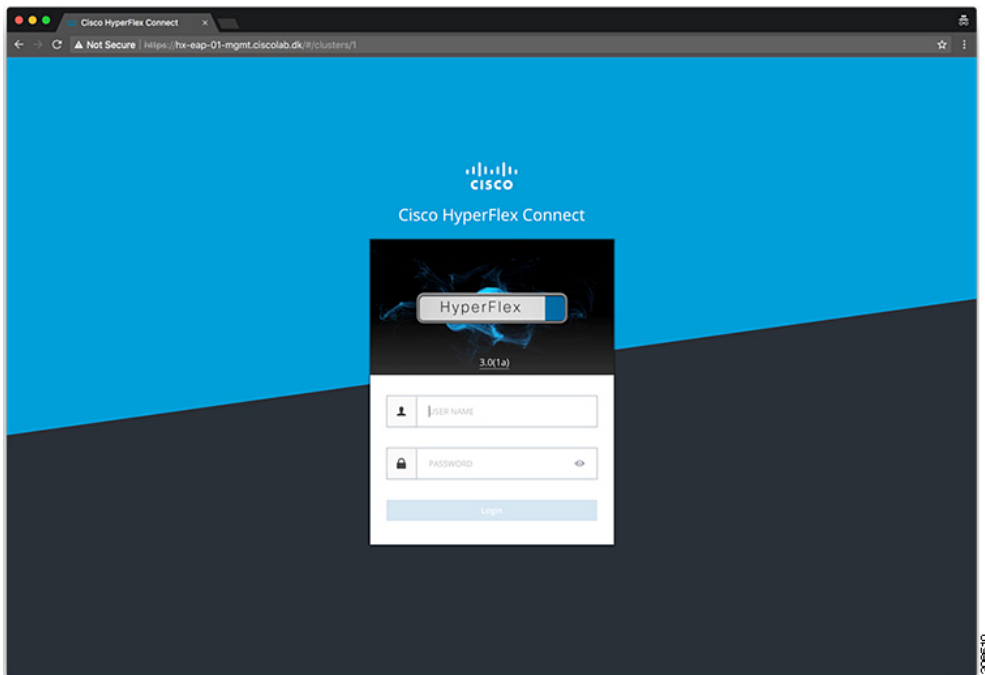
Confirm HX Data Platform Plug-in installation.

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:



CHAPTER 5

Post Installation Tasks

- [\(1GE Only\) Run Post Installation Script](#) , on page 65
- [\(10/25GE Only\) Run Post Installation Script](#) , on page 67
- [Configuring vMotion Automatically](#), on page 69
- [Configuring vMotion Manually](#), on page 70
- [Configuring Traffic Shaping Manually](#), on page 70
- [\(10/25GE\) Using Additional VIC Ports \(Optional\)](#), on page 71

(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:

```
admin@SpringpathController:~$ hx_post_install
```

```
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 HX-01-cmip.example.com
```

```
HX CVM admin password:
```

```
Getting ESX hosts from HX cluster...
```

```
vCenter URL: 192.168.202.35
```

```
Enter vCenter username (user@domain): administrator@vsphere.local
```

```
vCenter Password:
```

```
Found datacenter HX-Clusters
```

```
Found cluster HX-01
```

```
post_install to be run for the following hosts:
```

```
HX-01-esxi-01.example.com
```

```
HX-01-esxi-02.example.com
```

```
HX-01-esxi-03.example.com
```

```
Enter ESX root password:
```

```
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: 255.255.254.0
```

```
VLAN ID: (0-4096) 208
```

```
vMotion MTU is set to use jumbo frames (9000 bytes). Do you want to change to 1500 bytes? (y/n) y
```

```
vMotion IP for HX-01-esxi-01.example.com: 192.168.208.17
```

```
Adding vmotion-208 to HX-01-esxi-01.example.com
```

```
Adding vmkernel to HX-01-esxi-01.example.com
```

```
vMotion IP for HX-01-esxi-02.example.com: 192.168.208.18
```

```
Adding vmotion-208 to HX-01-esxi-02.example.com
```

```
Adding vmkernel to HX-01-esxi-02.example.com
```

```
vMotion IP for HX-01-esxi-03.example.com: 192.168.208.19
```

```
Adding vmotion-208 to HX-01-esxi-03.example.com
```

```
Adding vmkernel to HX-01-esxi-03.example.com
```

```
Add VM network VLANs? (y/n) y
```

```
Attempting to find UCSM IP
```

```
Found UCSM 10.75.61.254, logging with username admin. Org is HX-Cluster
```

```
UCSM Password:
```

```
Port Group Name to add (VLAN ID will be appended to the name): USERS
```

```
VLAN ID: (0-4096) 1219
```

```
Adding VLAN 1219 to FI
```



```
Adding VLAN 1219 to vm-network-a VNIC template
Adding USERS-1219 to HX-01-esxi-01.example.com
Adding USERS-1219 to HX-01-esxi-02.example.com
Adding USERS-1219 to HX-01-esxi-03.example.com
Add additional VM network VLANs? (y/n) n
```

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

```
Validating cluster health and configuration...
```

```
Cluster Summary:
Version - 4.0(2c)
Model - HXAF220C-M5SX
Health - HEALTHY
ASUP enabled - False
admin@SpringpathController:~$
```

(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:

(10/25GE Only) Run Post Installation Script

```

admin@SpringpathController:~$ hx_post_install

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 HX-01-cmip.example.com
HX CVM admin password:
Getting ESX hosts from HX cluster...
vCenter URL: 192.168.202.35
Enter vCenter username (user@domain): administrator@vsphere.local
vCenter Password:
Found datacenter HX-Clusters
Found cluster HX-01

post_install to be run for the following hosts:
HX-01-esxi-01.example.com
HX-01-esxi-02.example.com
HX-01-esxi-03.example.com

Enter ESX root password:

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: 255.255.254.0
VLAN ID: (0-4096) 208
vMotion MTU is set to use jumbo frames (9000 bytes). Do you want to change to 1500 bytes? (y/n) y
vMotion IP for HX-01-esxi-01.example.com: 192.168.208.17
Adding vmotion-208 to HX-01-esxi-01.example.com
Adding vmkernel to HX-01-esxi-01.example.com
vMotion IP for HX-01-esxi-02.example.com: 192.168.208.18
Adding vmotion-208 to HX-01-esxi-02.example.com
Adding vmkernel to HX-01-esxi-02.example.com
vMotion IP for HX-01-esxi-03.example.com: 192.168.208.19
Adding vmotion-208 to HX-01-esxi-03.example.com
Adding vmkernel to HX-01-esxi-03.example.com

Add VM network VLANs? (y/n) y
Attempting to find UCSM IP
Found UCSM 10.75.61.254, logging with username admin. Org is HX-Cluster
UCSM Password:
Port Group Name to add (VLAN ID will be appended to the name): USERS
VLAN ID: (0-4096) 1219
Adding VLAN 1219 to FI
Adding VLAN 1219 to vm-network-a VNIC template
Adding USERS-1219 to HX-01-esxi-01.example.com
Adding USERS-1219 to HX-01-esxi-02.example.com
Adding USERS-1219 to HX-01-esxi-03.example.com
Add additional VM network VLANs? (y/n) n

Run health check? (y/n) y

```

```
Validating cluster health and configuration...
```

```
Cluster Summary:  
Version - 4.0(2c)  
Model - HXAF220C-M5SX  
Health - HEALTHY  
ASUP enabled - False  
admin@SpringpathController:~$
```

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 70](#).

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 vMotion is configured on dedicated vMotion vSwitch with dedicate 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 in to 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 in to 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. On the left, a list of vNICs (eth0 through eth7) is visible. The main area displays a table of Host Ethernet Interfaces with the following columns: Name, CON, MAC Address, MTU, usNIC, Uplink Port, CoS, VLAN, VLAN Mode, iSCSI Boot, PXE Boot, Channel, Port Profile, and Uplink Failover. The 'Uplink Port' column is highlighted in yellow, showing values of 0 or 1 for each interface.

Name	CON	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode	iSCSI Boot	PXE Boot	Channel	Port Profile	Uplink Failover
<input type="checkbox"/> eth0	lv-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	lv-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-netwo...	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-netwo...	00:25:B5:50:B6:01	9000	0	1	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/> eth6	lv-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	lv-vmotio...	00:25:B5:50:B8:01	9000	0	1	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A

Creating Additional VIC Ports

Guidelines for creating additional vNICs:

- You can use the remaining ports to create additional vNICs as required. They will appear in VMware as unused vmnics. You can then create a new vSwitch and use these network ports (for guest VM use).
- Do not modify any of the existing vNICs, vSwitches, or portgroups that are created automatically by the installer. Create new vNICs, vSwitches and portgroups as needed only to support your requirement.
- Do not reuse existing HX vSwitches with extra uplinks. Create new vSwitches with these uplinks.
- When creating new vNICs, do not place them on the same physical 2 ports used for the HX services (for example, mgmt., storage traffic, vmotion, and vm-network). Be sure to create them on the unused physical ports.

Step 1 Confirm that the HX cluster is healthy.

Step 2 Place the ESXi host in HX Maintenance Mode.

Step 3 Create the new vNIC(s) and then place them in the corresponding Uplink Port # 2 or 3 which correspond to the last 2 ports of the VIC.

Add vNIC

▼ General

Name:

CDN:

MTU: (1500 - 9000)

Uplink Port: ▼

MAC Address:

Class of Service: (0 - 6)

Trust Host CoS:

PCI Order: (0 - 12)

Default VLAN: None (1 - 4094)

Step 4 Reboot ESXi and then confirm that you see the additional vmnics which you can then add to new vSwitches.

Step 5 Exit HX Maintenance Mode.

Step 6 Repeat this process accordingly for remaining hosts.



CHAPTER 6

Appendix

- [Sample Network Configuration 10GE Dual Switch, on page 73](#)
- [Sample Network Configuration 1GE Single Switch, on page 75](#)
- [Sample Network Configuration 1GE Dual Switch, on page 76](#)
- [Sample Network Configuration 10/25GE 2-Node 2-Room, on page 78](#)

Sample Network Configuration 10GE Dual Switch

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
```

Sample Network Configuration 1GE Single Switch

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

```

Sample Network Configuration 1GE Dual Switch

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.

Sample Network Configuration 10/25GE 2-Node 2-Room

10/25GE 2-Node 2-Room

Catalyst 9000 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

```

qos queue-softmax-multiplier 1200
qos queue-stats-frame-count
...
class-map match-all Storage_PQ
match dscp ef
class-map match-all Storage_Mark

```

```
match access-group name Storage
...
policy-map Storage_Mark
class Storage_Mark
    set dscp ef
class class-default
policy-map Storage_Queue
class Storage_PQ
    priority level 1 percent 80 //Adjust this value based on traffic mix. This guarantees
80% bandwidth for storage when needed.
    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
service-policy output Storage_Queue
!
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
service-policy output Storage_Queue
!
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 input Storage_Mark
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
channel-group 98 mode on
service-policy input Storage_Mark
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.

