



Cisco HyperFlex Systems Installation Guide for Cisco Intersight

First Published: 2018-01-10

Last Modified: 2024-03-21

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)

© 2018–2024 Cisco Systems, Inc. All rights reserved.



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

Overview

- [Overview, on page 1](#)
- [HyperFlex Systems Supported Releases, on page 1](#)
- [HyperFlex Systems Unsupported Releases, on page 3](#)
- [Limitations, on page 11](#)
- [HyperFlex Cluster Policies in Intersight, on page 12](#)
- [HyperFlex Server Personality, on page 13](#)

Overview

Cisco Intersight provides an installation wizard to install, configure, and deploy Cisco HyperFlex clusters — *HyperFlex Edge*, *FI-attached*, and *HyperFlex Datacenter without Fabric Interconnect*. The wizard constructs a pre-configuration definition of your cluster called an HyperFlex Cluster Profile. This definition is a logical representation of the HyperFlex nodes in your HyperFlex cluster and includes:

- **Security**—credentials for HyperFlex cluster such as controller VM password, Hypervisor username, and password.
- **Configuration**—server requirements, firmware, etc.
- **Connectivity**—upstream network, virtual network, etc.

HyperFlex Cluster Profiles are built on policies which administrator defined sets of rules and operating characteristics such as the node identity, interfaces, and network connectivity. Every active node in your HyperFlex cluster must be associated with an HyperFlex Cluster Profile.

After gathering the node configuration settings to build the HyperFlex Cluster Profile, the installation wizard will validate and deploy the HyperFlex Cluster Profile in your HyperFlex cluster. You can clone a successfully deployed HyperFlex Cluster Profile, and then use that copy as the basis to create a new cluster. For instructions on cloning HyperFlex cluster profiles, see [Cloning HyperFlex Cluster Profiles](#).

HyperFlex Systems Supported Releases

Intersight supports the following HyperFlex Data Platform versions for HyperFlex installation:

- 5.5(1a)
- 5.0(2e), 5.0(2g)



- Note**
- HXDP versions 5.0(2a), 5.0(2b), 5.0(2c), 5.0(2d), 4.5(2a), 4.5(2b), 4.5(2c), 4.5(2d), and 4.5(2e) are still supported for cluster expansion only.
 - Upgrades from HXDP 4.0.2x are supported provided the ESXi version is compatible with 4.5(2x).

HyperFlex Data Platform Release	Features Added in the Release	Description	Reference
5.5(1a)	<ul style="list-style-type: none"> • VMware ESXi 8.0 U1 support • UCS FI 6536 support 	Feature release with ESXi 8.0 U1 support	Supported Models/Versions for HyperFlex Fabric Interconnect Cluster Deployment Supported Models/Versions for HyperFlex Edge Cluster Deployment HyperFlex Datacenter Without Fabric Interconnect Cluster Deployment
5.0(2g)	N/A	Maintenance release with bug fixes.	Supported Models/Versions for HyperFlex Fabric Interconnect Cluster Deployment Supported Models/Versions for HyperFlex Edge Cluster Deployment HyperFlex Datacenter Without Fabric Interconnect Cluster Deployment

HyperFlex Data Platform Release	Features Added in the Release	Description	Reference
5.0(2e)	N/A	Maintenance release with bug fixes.	Supported Models/Versions for HyperFlex Fabric Interconnect Cluster Deployment Supported Models/Versions for HyperFlex Edge Cluster Deployment HyperFlex Datacenter Without Fabric Interconnect Cluster Deployment

HyperFlex Systems Unsupported Releases

HyperFlex cluster deployments and upgrade are no longer supported in Intersight for the following HyperFlex Data Platform releases:

- 5.0(2a), 5.0(2b), 5.0(2c), 5.0(2d)
- 4.5(2a), 4.5(2b), 4.5(2c), 4.5(2d), 4.5(2e)
- 4.0(2a), 4.0(2b), 4.0(2c), 4.0(2d), 4.0(2e), 4.0(2f)

HyperFlex Data Platform releases that have reached End of Support:

- 5.0(1a), 5.0(1b), 5.0(1c)
- 4.5(1a)
- 4.0(1a), 4.0(1b)



Note

- Upgrades from HXDP 4.0.2x are supported provided the ESXi version is compatible with 4.5(2x). You can upgrade both at the same time. For example, upgrade HXDP 4.0.2x + ESXi 6.0 to HXDP 4.5 + ESXi 6.5.
- HXDP versions 5.0(2a), 5.0(2b), 5.0(2c), 5.0(2d) 4.5(2a) and 4.5(2b), 4.5(2c), 4.5(2d) and 4.5(2e) are still supported for cluster expansion only.

HyperFlex Data Platform Release Feature Matrix [Unsupported Releases]

Unsupported HyperFlex Data Platform Release	Features Added in the Release	Description	Reference
5.0(2d)	N/A	Maintenance release with bug fixes.	Supported Models/Versions for HyperFlex Fabric Interconnect Cluster Deployment Supported Models/Versions for HyperFlex Edge Cluster Deployment HyperFlex Datacenter Without Fabric Interconnect Cluster Deployment
5.0(2c)	N/A	Maintenance release with bug fixes.	Supported Models/Versions for HyperFlex Fabric Interconnect Cluster Deployment Supported Models/Versions for HyperFlex Edge Cluster Deployment HyperFlex Datacenter Without Fabric Interconnect Cluster Deployment
5.0(2b)	B200 M6 compute node support	Support for cluster expansion with B200 M6 compute nodes for Datacenter and DC-no-FI clusters	Supported Models/Versions for HyperFlex Fabric Interconnect Cluster Deployment Supported Models/Versions for HyperFlex Edge Cluster Deployment HyperFlex Datacenter Without Fabric Interconnect Cluster Deployment

Unsupported HyperFlex Data Platform Release	Features Added in the Release	Description	Reference
5.0(2a)	HyperFlex HX245C/225C M6 All Flash/Hybrid Server nodes	Added support for HyperFlex HX245C/225C M6 All Flash/Hybrid Server nodes.	Supported Models/Versions for HyperFlex Fabric Interconnect Cluster Deployment
5.0(1c)	N/A	Maintenance release with bug fixes.	
5.0(1b)	HyperFlex Software Encryption	The HyperFlex software encryption feature offers file level end-to-end data encryption to provide confidentiality of data at-rest from theft of storage media. Intersight manages the keys natively with Intersight Key Manager which increases both security and simplicity by eliminating the overhead of key management.	HyperFlex Software Encryption

Unsupported HyperFlex Data Platform Release	Features Added in the Release	Description	Reference
5.0(1a)	Cisco UCS/HyperFlex M6 Node Support	Supported models are Cisco UCS 220c M6 and Cisco UCS 240c M6 HyperFlex nodes. You can use Intersight to install and upgrade a HyperFlex converged node on HyperFlex-Series M6 servers and BOM compliant UCS C-Series M6 servers.	
	HyperFlex Cluster Deployment	You can use HyperFlex-Series M6 servers, as well as BOM compliant UCS C-Series M6 servers, to deploy HyperFlex Clusters. HyperFlex installation is no longer restricted to Cisco UCS C-series nodes manufactured as HyperFlex nodes.	Cisco HyperFlex Cluster Deployment
	HyperFlex Server Personality	With HyperFlex server personalities, you can use a BOM compliant C-Series M6 server to deploy a HyperFlex cluster. The HyperFlex Series M6-servers have the HyperFlex Server Personality configured. A BOM compliant UCS C-Series M6 servers in the deployed HyperFlex cluster will also have a HyperFlex Server Personality configured.	HyperFlex Server Personality
	Backup and Restore for HyperFlex Clusters using Intersight		N:1 Replication for Cisco HyperFlex Clusters

Unsupported HyperFlex Data Platform Release	Features Added in the Release	Description	Reference
		<ul style="list-style-type: none"> • Ability to configure different retention counts on the source and target clusters. • Ability to use HyperFlex Fabric Interconnected clusters as source clusters. • Backup dashboard enhancements to include error reporting for failed snapshots and replication, drill down options, and consolidated view of failed backups and restores in the last 24 hours. 	
	Secure Boot	Secure boot is enabled by default.	
4.5(2e)	HyperFlex Data Platform 4.5(2e)	Added HyperFlex Data Platform 4.5(2e) support for HyperFlex cluster deployment and upgrade.	Supported Models/Versions for HyperFlex Fabric Interconnect Cluster Deployment Supported Models/Versions for HyperFlex Edge Cluster Deployment HyperFlex Datacenter Without Fabric Interconnect Cluster Deployment

Unsupported HyperFlex Data Platform Release	Features Added in the Release	Description	Reference
4.5(2d)	HyperFlex Data Platform 4.5(2d)	Added HyperFlex Data Platform 4.5(2d) support for HyperFlex cluster deployment and upgrade. Added support for NVMe cache in All Flash models, All NVMe hardware models, and 40/100 GE Networking with DC-No-FI HyperFlex clusters.	Supported Models/Versions for HyperFlex Fabric Interconnect Cluster Deployment HyperFlex Datacenter Without Fabric Interconnect Cluster Deployment
4.5(2c)	Support for VMware ESXi 7.0 U3	Added support for VMware ESXi 7.0 U3 version.	Supported Systems
4.5(2b)	End of Life Advisories for HyperFlex Data Platform Software Releases	Cisco Intersight alerts users about End of Life and End of Support dates for Cisco HyperFlex Data Platform software releases that are no longer supported with a list of devices that are affected.	Cisco End of Life Advisories
4.5(2a)	Health Check for HyperFlex Clusters	Provides that ability to run pre-defined health checks on HyperFlex clusters and view granular details about the health of HyperFlex clusters	Health Check for HyperFlex Clusters
	Backup and Restore for HyperFlex Edge Clusters using Intersight	Provides the ability for HyperFlex Edge clusters to take snapshots of Virtual Machines and restore using Intersight.	N:1 Replication for Cisco HyperFlex Clusters

Unsupported HyperFlex Data Platform Release	Features Added in the Release	Description	Reference
4.5(1a)	Backup and Restore for HyperFlex Edge Clusters using Intersight	Provides the ability for HyperFlex Edge clusters to take snapshots of Virtual Machines and restore using Intersight.	N:1 Replication for Cisco HyperFlex Clusters
	Health Check for HyperFlex Clusters	Provides that ability to run pre-defined health checks on HyperFlex clusters and view granular details about the health of HyperFlex clusters	Health Check for HyperFlex Clusters
	Upgrade of HyperFlex Clusters	Provides the ability to upgrade HyperFlex clusters	Upgrade Cisco HyperFlex Systems in Cisco Intersight
4.0(2a), 4.0(2b), 4.0(2c), 4.0(2d), 4.0(2e), 4.0(2f)	Upgrade of HyperFlex Edge Clusters	Combined upgrade of VMware ESXi and HyperFlex Data Platform.	Upgrade Cisco HyperFlex Systems in Cisco Intersight
	HyperFlex Cluster Deployment	Support for 25GE networking topology for HyperFlex Edge.	Cisco HyperFlex Cluster Deployment

Unsupported HyperFlex Data Platform Release	Features Added in the Release	Description	Reference
4.0(1b)	HyperFlex Cluster Deployment	<p>View the progress and history of the HyperFlex Cluster Profile deployment both from the Requests page and the HyperFlex Cluster Profile Results page.</p> <p>Note Ability to view the progress in the Requests page is not available in HyperFlex Data Platform 4.0(2a) and 3.5(x).</p>	Cisco HyperFlex Cluster Deployment
	Upgrade of HyperFlex Edge Clusters	<p>Multi-site remote HyperFlex Data Platform upgrade of HyperFlex Edge clusters installed using Intersight.</p> <p>Note Use HyperFlex Connect to upgrade Edge clusters at versions before HyperFlex Data Platform 4.0(1a)</p>	Upgrade Cisco HyperFlex Systems in Cisco Intersight

Unsupported HyperFlex Data Platform Release	Features Added in the Release	Description	Reference
4.0(1a)	Alarms and Health Status	Supports alarms and health status for HyperFlex with Hyper-V	
	HyperFlex Cluster Deployment	Comprehensive lifecycle management and includes remote cloud-based installation, and invisible witnessing of HyperFlex Edge clusters in 1GE and 10GE networking topology options. This release extends support to 2-Node Edge clusters to run HyperFlex in environments requiring a small footprint and 4-Node Edge clusters to enable scaling-up HyperFlex Edge clusters.	Cisco HyperFlex Cluster Deployment
	Invisible Cloud Witness for Cisco HyperFlex Edge	Invisible Cloud Witness architecture uses a new deployment methodology that eliminates the need for witness VMs, ongoing patching and maintenance, or additional infrastructure at a third site to ensure data consistency during a node loss in a clustered file system. For a 2-Node Edge cluster, Intersight acts as the arbitrator to form a quorum to maintain data consistency in case of a node failure. To use this feature, you can use the basic Intersight requirements for Port and firewall rules.	Invisible Cloud Witness for Cisco HyperFlex Edge

Limitations

The following list outlines the limitations in performing cluster installation through Intersight:

- Intersight cluster install is not supported on Stretched Clusters.
- Intersight cluster install is not supported on Hyper-V Clusters.
- 10G+ NIC-based cluster deployment is supported on HXDP version 5.0(2a) and later for HX Edge and DC-No-FI clusters.
- HXDP release 5.5(1a) and above doesn't support VMware ESXi 6.5, 6.7 and 7.0 U1.
- HXDP release 5.5(1a) and above doesn't support HyperFlex M4 platforms. If your cluster has M4 nodes, please plan and replace them with supported platform before upgrading to HXDP 5.5(x).
- NIC-based cluster deployment is supported on VMware ESXi version 7.0 U3 and later.
- NIC-based cluster deployment is supported on M6 platform only for HX Edge and DC-No-FI clusters.
- Mixing of NIC-based and VIC-based nodes within the same cluster is not supported.

HyperFlex Cluster Policies in Intersight

Navigate to **CONFIGURE > Policies > Create Policy > HyperFlex Cluster**. HyperFlex Policies in Cisco Intersight provide different configurations including Auto Support, External Storage (such as FC and iSCSI), security, network configuration and more. A policy that is once configured can be assigned to any number of servers in order to provide a configuration baseline.

HyperFlex Policies can also be cloned by using the Policy Clone wizard with properties that are similar to the existing policies. The clone policy action is available on both the policies list and detailed views.

All HyperFlex policies can be shared by HyperFlex Edge and HyperFlex with Fabric Interconnect, unless otherwise mentioned below:

- **Auto Support Policy**—Auto Support is the alert notification service provided through HyperFlex Data Platform in Intersight. If enabled, notifications are sent to designated email addresses or email aliases that you want to receive the notifications. Typically, Auto Support is configured during HyperFlex storage cluster creation by configuring the SMTP mail server and adding email recipients.
- **Backup Configuration Policy**—The backup policy consists of the Edge cluster datastores being protected, the backup interval, snapshot retention value, and the backup target cluster. All virtual machines residing in the protected Edge cluster datastore will be automatically protected by the backup policy. This includes VMs created in the protected datastore, as well as VMs migrated into the protected datastore.
- **DNS, NTP, and Timezone Policy**—Configures DNS, NTP, and Timezone on all servers. DNS and NTP servers should reside outside of the HyperFlex storage cluster. Use an internally-hosted NTP server to provide a reliable source for the time.
- **External FC Storage Policy**—Enables the use of external FC Storage and configures the WWxN name and the associated VSAN details. This policy is not applicable to HyperFlex Edge clusters.
- **External iSCSI Storage Policy**—Enables the use of external iSCSI Storage and configures the associated VLAN details. This policy is not applicable to HyperFlex Edge clusters.
- **HTTP Proxy**—Specifies the HTTP proxy settings to be used by the HyperFlex installation process and the HyperFlex Storage Controller VMs. This policy is required when the internet access of your servers including CIMC and HyperFlex Storage Controller VMs is secured by an HTTP proxy.

- **Network Configuration Policy**—Configures the VLAN and KVM for the management network in Fabric Interconnects; and Uplink Speed, VLAN, and Jumbo Frames for the management network in Edge clusters. The VLAN must have access to Intersight. This policy cannot be shared by HyperFlex Edge and HyperFlex with Fabric Interconnect clusters.
- **Node IP Ranges Policy**—Configures the management IP ranges for hypervisors and controller VMs. The data IPs are automatically assigned in a /24 subnet in the range *169.254.x.2* to *169.254.239.254*.
- **Replication Network Configuration Policy**—The replication network policy consists of replication VLAN, gateway, subnet mask, bandwidth, MTU, and replication IP address range parameters. The replication network policy is unique for each Edge cluster configured to use N:1 Replication.
- **Security Policy**—Configures ESXi and Controller VM password for the HyperFlex cluster. This policy presents an option to update the Hypervisor password in Intersight, if you have not already changed it on the Hypervisor.
- **Storage Configuration Policy**—Configures the options for VDI Optimization (for hybrid HyperFlex systems). For HyperFlex with Fabric Interconnect and DC-No-FI, this policy provides the option to enable Logical Availability Zones as part of the Cluster Deployment workflow. Logical Availability Zones are not supported on HyperFlex Edge clusters.



Note Logical Availability Zones are automatic partitions of the physical cluster into multiple logical zones. They are created to avoid multiple node and component failures on large clusters, and to increase cluster resiliency. HyperFlex Data Platform intelligently places a copy of data in every zone. When there is a node failure in a single zone, it does not cause the entire cluster to fail because the other zones contain data replicas. Logical Availability Zones can only be enabled with clusters that are 8 converged nodes or larger.

The LAZ option in the Storage Configuration policy is recommended for clusters greater than 8 nodes.

- **vCenter Policy**—An optional policy during installation of the HyperFlex cluster. However, post-installation, you must register the cluster to vCenter to ensure that the cluster functions smoothly. This policy is optional to the installation of the HyperFlex cluster.
- **DC-No-FI Policy**—Cisco HyperFlex Datacenter without Fabric Interconnect (DC-No-FI) brings the simplicity of hyperconvergence to data center deployments without the requirement of connecting the converged nodes to Cisco Fabric Interconnect.

HyperFlex Server Personality

The HyperFlex Series M6-servers have the HyperFlex Server Personality configured. A BOM compliant UCS C-Series M6 servers in the deployed HyperFlex cluster will also have a HyperFlex Server Personality configured. Navigate to **OPERATE > Servers** and select the server(s) in the servers table view. You can view the server personality details under the **General** tab. Server personality is set by during cluster deployment to include information like the disk type, node role, and SED capability. Server Personality for a HyperFlex server can have the following values:

- Compute Node

- Converged Node

During cluster deployment, you can view the node role based on Server Personality in the **Node Type** column. If you choose a node that has a *HyperFlex Compute Server* or no personality, you must ensure that the required hardware is available in the server for successful cluster deployment. For information about the Product Identification Standards (PIDs) that are supported by Cisco Intersight, see [Cisco HyperFlex HX-Series Data Sheet](#)

Select a node and click the Ellipsis (...) icon to change the personality of a node from **Compute** to **Converged** or from **Converged** to **Compute**.

A server configured as HyperFlex node with a personality of either Compute or Converged node cannot be reset to factory default no-personality mode from Intersight. Use Cisco IMC or UCS Manager API to reset the personality to factory default *No Personality* state.

- To reset the server personality using Cisco IMC, see [Cisco UCS C-Series Servers Integrated Management Controller CLI Configuration Guide](#).
- To reset the server personality using the *GetPersonality* UCSM API, see [Cisco UCS Manager XML API Programmer's Guide](#).



CHAPTER 2

Deploy HyperFlex Edge Clusters

- [Installation Overview, on page 15](#)
- [Preinstallation Checklist for HyperFlex Edge, on page 16](#)
- [Supported Models/Versions for HyperFlex Edge Cluster Deployments, on page 16](#)
- [Installation, on page 17](#)
- [Post Installation, on page 26](#)

Installation Overview

The following table summarizes the installation workflow for Cisco HyperFlex Edge:

Step	Description	Reference
1.	Complete the preinstallation checklist.	Preinstallation Checklist for Cisco HyperFlex Edge
2.	Ensure that the network is set up.	
3.	Log in to Cisco Intersight.	Log In to Cisco Intersight, on page 17
4.	Claim Targets. Note Skip if you have already claimed HyperFlex Nodes.	Claim Edge Targets, on page 18
5.	Verify Cisco UCS Firmware versions.	Verify Firmware Version for HyperFlex Edge, on page 19
6.	Run the HyperFlex Cluster Profile Wizard.	Configure HyperFlex Edge Clusters, on page 19
7.	Run the post installation script through the controller VM.	Post Installation Tasks, on page 26

Preinstallation Checklist for HyperFlex Edge

Ensure that your system meets the following installation and configuration requirements before you begin to install a Cisco HyperFlex Edge system. Refer to the [Preinstallation Checklist for Cisco HyperFlex Edge](#) for detailed preinstallation requirements.

Supported Models/Versions for HyperFlex Edge Cluster Deployments

The following table lists the supported hardware platforms and software versions for HyperFlex Edge cluster deployments. For information about the Product Identification Standards (PIDs) that are supported by Cisco Intersight, see [Cisco HyperFlex HX-Series Data Sheet](#).

Component	Models/Versions
M6 Servers	<ul style="list-style-type: none"> • HXAF-E-225M6S • HX-E-225M6S • HXAF-E-245-M6SX • HX-E-245-M6SX • HX-E-240-M6SX • HXAF-E-240-M6SX • HX-E-220-M6S • HXAF-E-220-M6S
M5 Servers	<ul style="list-style-type: none"> • HX240C-M5SD • HX240C-M5SD • HX220C-M5SX • HXAF220C-M5SX

Component	Models/Versions
Cisco HX Data Platform (HXDP)	<ul style="list-style-type: none"> • 5.5(1a) • 5.0(2e), 5.0(2g) <p>Note</p> <ul style="list-style-type: none"> • HXDP versions 5.0(2a), 5.0(2b), 5.0(2c), 5.0(2d), 4.5(2a), 4.5(2b), 4.5(2c), 4.5(2d), and 4.5(2e) are still supported for cluster expansion only. • Upgrades from HXDP 4.0.2x are supported provided the ESXi version is compatible with 4.5(2x). • M6 servers require HXDP 5.0(1a) or later.
NIC Mode	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • Dedicated Management Port • Shared LOM
Device Connector	Auto-upgraded by Cisco Intersight
Network Topologies	1GE and 10G+
Connectivity Type	<p>Types:</p> <ul style="list-style-type: none"> • VIC based • NIC-based (10G+ NIC-based clusters require HXDP version 5.0(2a) or later)

Installation

Log In to Cisco Intersight

Log In using Cisco ID

To login to Cisco Intersight, you must have a valid **Cisco ID** to create a Cisco Intersight account. If you do not have a Cisco ID, create one [here](#).



Important

The device connector does not mandate the format of the login credentials, they are passed as is to the configured HTTP proxy server. Whether or not the username must be qualified with a domain name will depend on the configuration of the HTTP proxy server.

Log In using Single Sign-On

Single Sign-On (SSO) authentication enables you to use a single set of credentials to log in to multiple applications. With SSO authentication, you can log in to Intersight with your corporate credentials instead of your Cisco ID. Intersight supports SSO through SAML 2.0, and acts as a service provider (SP), and enables integration with Identity Providers (IdPs) for SSO authentication. You can configure your account to sign in to Intersight with your Cisco ID and SSO. Learn more about SSO with Intersight [here](#).

Claim Edge Targets

Complete the following steps to claim one or more Targets to be managed by Cisco Intersight:

Before you begin

This procedure assumes that you are an existing user with a Cisco account. If not, see [Log In to Cisco Intersight, on page 17](#).

-
- Step 1** In the Cisco Intersight, left navigation pane, select **ADMIN > Targets**.
- Step 2** In the **Targets** details page, click **Claim a New Target**.
- Step 3** In the **Claim a New Target** wizard, select **Hyperconverged > Cisco HyperFlex Cluster** and complete the following fields:

- Note** You can locate the **Device ID** and the **Claim Code** information in:
- Cisco IMC by navigating to **Admin > Device Connector**.
 - Cisco HyperFlex by navigating to **HyperFlex Connect UI > Settings > Device Connector**.

UI Element	Essential Information
Device ID	Enter the applicable Device ID. <ul style="list-style-type: none"> For a Cisco UCS C-Series Standalone server, use serial number. Example: NGTR12345 For HyperFlex, use Cluster UUID. Example: xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx
Claim Code	Enter target claim code. You can find this code in the Device Connector for the target type. <p>Note Before you gather the Claim Code, ensure that the Device Connector has outbound network access to Cisco Intersight, and is in the “Not Claimed” state.</p>

- Step 4** Click **Claim**.

Note Refresh the Targets page to view the newly claimed target.

Verify Firmware Version for HyperFlex Edge

View current BIOS, CIMC, SAS HBA, and drive firmware versions, and verify that those versions match the Cisco HyperFlex Edge and Firmware Compatibility Matrix in the *Common Network Requirements*. Refer to the [Preinstallation Checklist for Cisco HyperFlex Edge](#) for 2-Node, 3-Node, and 4-Node Edge clusters for more details.

- Step 1** In your browser, log in to the CIMC web UI by navigating to **https://<CIMC IP>**. You can also cross launch CIMC from Cisco Intersight from the **Servers** tableview.
 - 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 > Storage**. Double-click on **Cisco 12G Modular SAS HBA (max 16 drives) (MRAID)** and navigate to **Details > Physical Drive Info**.
 - Step 6** Compare the current BIOS, CIMC, SAS HBA, and drive firmware versions with the versions listed in the Cisco HyperFlex Edge and Firmware Compatibility Matrix in the *Common Network Requirements*. Refer to the [Preinstallation Checklist for Cisco HyperFlex Edge](#) for 2-Node, 3-Node, and 4-Node Edge clusters for more details.
 - Step 7** 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>.
-

Configure HyperFlex Edge Clusters

To configure a HyperFlex Edge Cluster in Intersight, do the following:

- Step 1** Log in to Intersight with HyperFlex Cluster administrator or Account Administrator privileges.
- Step 2** Navigate to **CONFIGURE > Profiles**.
- Step 3** In the **Profiles** page, make sure that the **HyperFlex Cluster Profiles** tab is selected, and click **Create HyperFlex Cluster Profile** to launch the Create HyperFlex Cluster Profile installation wizard.
- Step 4** Select **Edge** as the deployment type. Click **Start**.
- Step 5** In the **General** page, complete the following fields:

Field	Description
Organization drop-down list	<p>You can make the HyperFlex Cluster Profile belong to the default organization or a specific organization. Choose:</p> <ul style="list-style-type: none"> • default—To make the Cluster Profile belong to the default organization. All the policies that belong to the default organization will be listed on the Create HyperFlex Cluster Profile wizard. • Specific Organization—To make the HyperFlex Cluster Profile belong to the specified organization only. Only the policies that belong to the selected organization will be listed on the Create HyperFlex Cluster Profile wizard. <p>For example, if HyperFlex nodes are shared across two organizations and are associated to a Cluster Profile in one organization, you cannot associate the same node to a Cluster Profile in another organization. The Cluster Profile will be available only to users who belong to the specified Organization.</p>
Name field	<p>Enter a name for the HyperFlex cluster.</p> <p>The cluster name will be used as the vCenter cluster name, HyperFlex storage controller name, and HyperFlex storage cluster name.</p> <p>Note The name of the HyperFlex Cluster Profile belonging to an organization must be unique. You may create a HyperFlex Cluster Profile with the same name in a different organization.</p>
HyperFlex Data Platform Version drop-down list	<p>Select the version of the Cisco HyperFlex Data Platform to be installed. This can be one of the following:</p> <ul style="list-style-type: none"> • 5.5(1a) • 5.0(2e), 5.0(2g) <p>Note The version that you select impacts the types of HyperFlex policies that you can choose later in the configuration wizard.</p>
(Optional) Description field	Add a description for the HyperFlex cluster profile.
(Optional) Set Tags field	Enter a tag key.

Click **Next**.

Step 6

In the **Nodes Assignment** page, you can assign nodes now or optionally, you can choose to assign the nodes later. To assign nodes, click the **Assign nodes** check box and select the node you want to assign.

You can view the node role based on Server Personality in the **Node Type** column. If you choose a node that has a *HyperFlex Compute Server* or no personality, you must ensure that the required hardware is available in the server for successful cluster deployment. For information about the Product Identification Standards (PIDs) that are supported by Cisco Intersight, see [Cisco HyperFlex HX-Series Data Sheet](#)

Important Cisco HyperFlex Edge cluster only allows a minimum of 2 to a maximum of 4 nodes.

Note The expansion of edge clusters beyond 4 nodes, changes the deployment type from Edge type to DC-No-FI type.

Click **Next**.

Step 7

In the **Cluster Configuration** page, complete the following fields:

Note For the various cluster configuration tasks, you can enter the configuration details or import the required configuration data from policies. To use pre-configured policies, click **Select Policy**, next to the configuration task and choose the appropriate policy from the list.

Field	Description
<i>Security</i>	
Hypervisor Admin field	Enter the Hypervisor administrator username. Note Use root account for ESXi deployments.
Hypervisor Password field	Enter the Hypervisor password, this can be one of the following: Remember The default ESXi password of Cisco123 must be changed as part of installation. For fresh ESXi installation, ensure the checkbox for The Hypervisor on this node uses the factory default password is checked . Provide a new ESXi root password that will be set on all nodes during installation. If the ESXi installation has a non-default root password, ensure the checkbox The Hypervisor on this node uses the factory default password is unchecked. Provide the ESXi root password that you configured. This password will not be changed during installation.
Hypervisor Password Confirmation field	Retype the Hypervisor password.
Controller VM Admin Password field	Enter a user-supplied HyperFlex storage controller VM password. Important Make a note of this password as it will be used for the administrator account.
Controller VM Admin Password Confirmation field	Retype Controller VM administrator password.

Field	Description
<i>DNS, NTP, and Timezone</i>	
Timezone field	Select the local timezone.
DNS Suffix field	Enter the suffix for the DNS. This is applicable only for HX Data Platform 3.0 and later.
DNS Servers field	Enter one or more DNS servers. A DNS server that can resolve public domains is required for Intersight.
NTP Servers field	Enter one or more NTP servers (IP address or FQDN). A local NTP server is highly recommended.
<i>vCenter (Optional Policy)</i>	
vCenter Server FQDN or IP field	Enter the vCenter server FQDN or IP address.
vCenter Username field	Enter the vCenter username. For example, <i>administrator@vsphere.local</i>
vCenter Password field	Enter the vCenter password.
vCenter Datacenter Name field	Enter the vCenter datacenter name.
<i>Storage Configuration (Optional Policy)</i>	
VDI Optimization check box	Check this check box to enable VDI optimization (hybrid HyperFlex systems only).
<i>Auto Support (Optional Policy)</i>	
Auto Support check box	Check this check box to enable Auto Support.
Send Service Ticket Notifications To field	Enter the email address recipient for support tickets.
<i>Node IP Ranges</i>	
Note	This section configures the management IP pool. You must complete the management network fields to define a range of IPs for deployment. On the node configuration screen, these IPs will be automatically assigned to the selected nodes. If you wish to assign a secondary range of IPs for the controller VM Management network, you may optionally fill out the additional fields below. Both IP ranges must be part of the same subnet.
Management Network Starting IP field	The starting IP address for the management IP pool.
Management Network Ending IP field	The ending IP address for the management IP pool.
Management Network Subnet Mask field	The subnet mask for the management VLAN.
Management Network Gateway field	The default gateway for the management VLAN.
Controller VM Management Network Starting IP field (Optional)	The starting IP address for the controller VM management network.

Field	Description
Controller VM Management Network Ending IP field (Optional)	The ending IP address for the controller VM management network.
Controller VM Management Network Subnet Mask field (Optional)	The subnet mask for the controller VM management network.
Controller VM Management Network Gateway field (Optional)	The default gateway for the controller VM management network.
<i>Cluster Network</i>	

Field	Description
Uplink Speed drop-down list	<p>Select the link speed of the server adapter port to the upstream switch. The Uplink speed can be:</p> <ul style="list-style-type: none"> • 1G (HyperFlex Edge) • 10G+ (HyperFlex Edge) <p>When the policy is attached to a cluster profile with Edge management platform, the uplink speed can be '1G' or '10G+'. When the policy is attached to a cluster profile with Fabric Interconnect management platform, the uplink speed can be 'default' only.</p> <p>Refer to the Preinstallation Checklist for Cisco HyperFlex Edge for more details of the supported Network Topologies.</p> <p>Attention Using 10G+ 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 for further details on how to change the FEC mode configured on the VIC using the Cisco IMC GUI.</p>
Management Network VLAN ID field	<p>Enter the VLAN ID for the management network. VLAN must have access to Intersight.</p> <p>An ID of 0 means the traffic is untagged. The VLAN ID can be any number between 0 and 4095, inclusive.</p>

Field	Description
Jumbo Frames check box	Check this check box to enable Jumbo Frames. Jumbo Frames are optional and can remain disabled for HyperFlex Edge deployments.
<i>Proxy Setting</i> (Optional Policy)	
Hostname field	Enter the HTTP proxy server FQDN or IP address.
Port field	Enter the proxy port number.
Username field	Enter the HTTP Proxy username.
Password field	Enter the HTTP Proxy password.
<i>HyperFlex Storage Network</i>	
Storage Network VLAN ID field	Enter the VLAN ID for the storage VLAN traffic. The VLAN must be unique per HyperFlex cluster. Note The storage VLAN must be unique per HyperFlex cluster. This VLAN does not need to be routable and can remain layer 2 only. IP addresses from the link local range <i>169.254.0.0/16</i> are automatically assigned to storage interfaces. A storage VLAN is not required for two node HyperFlex Edge 1GE configurations, and you should enter 1 for this field.

Click **Next**.

Step 8

In the **Nodes Configuration** page, you can view the IP and Hostname settings that were automatically assigned. Intersight will make an attempt to auto-allocate IP addresses. Complete the following fields:

Field	Description
Cluster Management IP Address field	The cluster management IP should belong to the same subnet as the Management IPs.
MAC Prefix Address field	The MAC Prefix Address is auto-allocated for NIC-based and 1G HyperFlex Edge clusters. For 10G+ HyperFlex Edge clusters you can overwrite the MAC Prefix address, using a MAC Prefix address from the range 00:25:B5:00 to 00:25:B5:EF. Attention Ensure that the MAC prefix is unique across all clusters for successful HyperFlex cluster deployment. Intersight does a validation for duplicate MAC prefix and shows appropriate warning if any duplicate MAC prefix is found.

Field	Description
Replication Factor radio button	The number of copies of each data block written. The options are 2 or 3 redundant replicas of your data across the storage cluster. Important Replication factor 3 is the recommended option.
Hostname Prefix field	The specified Hostname Prefix will be applied to all nodes.

Step 9 In the **Summary** page, you can view the cluster configuration and node configuration details. Review and confirm that all information entered is correct. Ensure that there are no errors triggered under the **Errors/Warnings** tab.

Note When deploying 2 Node Edge clusters a warning will be displayed reminding of the importance of connectivity to Intersight. Ensure that your cluster remains connected to Intersight at all times. A second warning will be shown reminding users to implement a backup strategy to ensure that all user data is protected.

Step 10 Click **Validate and Deploy** to begin the deployment. Optionally, click **Validate**, and click **Save & Close** to complete deployment later. The **Results** page displays the progress of the various configuration tasks. You can also view the progress of the HyperFlex Cluster Profile deployment from the **Requests** page.

What to do next

Monitoring cluster deployment

Check your cluster deployment progress in the following ways:

- You can remain on the **Results** page to watch the cluster deployment progress in real time.
- You can also close the current view and allow the installation to continue in the background. To return to the results screen, navigate to **CONFIGURE > Profiles > HyperFlex Cluster Profiles**, and click on the name of your cluster.
- You can see the current state of your deployment in the status column in the HyperFlex Cluster Profile Table view.

Post Installation

Post Installation Tasks

- Step 1** Confirm that the HyperFlex Cluster is claimed in Intersight.
- Step 2** Confirm that the cluster is registered to vCenter.
- Step 3** Navigate to **HyperFlex Clusters**, select your cluster and click ... to launch HyperFlex Connect.
- Step 4** SSH to the cluster management IP address and login using **admin** username and the controller VM password provided during installation. Verify the cluster is online and healthy.

Step 5 Paste the following command in the Shell, and hit enter:

```
hx_post_install
```

Step 6 Follow the on-screen prompts to complete the installation. The **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 configuration check.
-



CHAPTER 3

Deploy HyperFlex Fabric Interconnect Clusters

- [Installation Overview, on page 29](#)
- [Preinstallation Checklist for HyperFlex Fabric Interconnect-attached Clusters, on page 30](#)
- [Supported Models/Versions for HyperFlex with Fabric Interconnect Cluster Deployments, on page 30](#)
- [Installation, on page 31](#)
- [Post Installation, on page 41](#)

Installation Overview

The following table summarizes the installation workflow for a configuring a Fabric Interconnect-attached HyperFlex cluster:

Step	Description	Reference
1.	Complete the preinstallation checklist.	Preinstallation Checklist for Cisco HX Data platform
2.	Ensure that the network is set up.	
3.	Log in to Cisco Intersight.	Log In to Cisco Intersight, on page 17
4.	Claim Targets. Note Skip if you have already claimed HyperFlex Nodes.	Claim Targets
5.	(Optional) Verify Cisco UCS Firmware versions.	Verify Firmware Version for Fabric Interconnect, on page 33
6.	Run the Create HyperFlex Cluster Profile Wizard.	Configure HyperFlex Fabric Interconnect Clusters, on page 33
7.	Run the post installation script through the controller VM.	Post Installation Tasks, on page 26

Preinstallation Checklist for HyperFlex Fabric Interconnect-attached Clusters

Ensure that your system meets the installation and configuration requirements before you begin to install Cisco HyperFlex Fabric Interconnects-attached clusters. Refer to the [Preinstallation Checklist for Cisco HX Data Platform](#) for detailed preinstallation requirements.

Supported Models/Versions for HyperFlex with Fabric Interconnect Cluster Deployments

The following table lists the supported hardware platforms and software versions for HyperFlex with Fabric Interconnect cluster deployments. For information about the Product Identification Standards (PIDs) that are supported by Cisco Intersight, see [Cisco HyperFlex HX-Series Data Sheet](#).

Component	Models/Versions
M6 Servers	<ul style="list-style-type: none"> • HXAF245C-M6SX • HX245C-M6SX • HXAF225C-M6SX • HX225C-M6SX • HX220C-M6S • HX240C-M6SX • HX240C-M6L • HXAF220C-M6S • HXAF220C-M6SN • HXAF240C-M6SX • HXAF240C-M6SN
M5 Servers	<ul style="list-style-type: none"> • HX220C-M5SX • HXAF220C-M5SX • HX240C-M5L • HX240C-M5SX • HXAF240C-M5SX

Component	Models/Versions
Cisco HyperFlex HX Data Platform (HXDP)	<ul style="list-style-type: none"> • 5.5(1a) • 5.0(2e), 5.0(2g) <p>Note</p> <ul style="list-style-type: none"> • HXDP versions 5.0(2a), 5.0(2b), 5.0(2c), 5.0(2d), 4.5(2a), 4.5(2b), 4.5(2c), 4.5(2d), and 4.5(2e) are still supported for cluster expansion only. • Upgrades from HXDP 4.0.2x are supported provided the ESXi version is compatible with 4.5(2x). • M6 servers require HXDP 5.0(1a) or higher.
Device Connector	Auto-upgraded by Cisco Intersight

Installation

Log In to Cisco Intersight

Log In using Cisco ID

To login to Cisco Intersight, you must have a valid **Cisco ID** to create a Cisco Intersight account. If you do not have a Cisco ID, create one [here](#).



Important The device connector does not mandate the format of the login credentials, they are passed as is to the configured HTTP proxy server. Whether or not the username must be qualified with a domain name will depend on the configuration of the HTTP proxy server.

Log In using Single Sign-On

Single Sign-On (SSO) authentication enables you to use a single set of credentials to log in to multiple applications. With SSO authentication, you can log in to Intersight with your corporate credentials instead of your Cisco ID. Intersight supports SSO through SAML 2.0, and acts as a service provider (SP), and enables integration with Identity Providers (IdPs) for SSO authentication. You can configure your account to sign in to Intersight with your Cisco ID and SSO. Learn more about SSO with Intersight [here](#).

Claim Fabric Interconnect Targets

Complete the following steps to claim one or more Targets to be managed by Cisco Intersight:

Before you begin

This procedure assumes that you are an existing user with a Cisco account. If not, see [Log In to Cisco Intersight, on page 17](#).

- Step 1** In the Cisco Intersight, left navigation pane, select **ADMIN > Targets**.
- Step 2** In the **Targets** details page, click **Claim a New Target**.
- Step 3** In the **Claim a New Target** wizard, select **Hyperconverged > Cisco HyperFlex Cluster** and complete the following fields:

- Note** You can locate the **Device ID** and the **Claim Code** information in:
- Cisco UCS Manager and Cisco IMC by navigating to **Admin > Device Connector**.
 - Cisco HyperFlex by navigating to **HyperFlex Connect UI > Settings > Device Connector**.
 - Cisco UCS Director by navigating to **Administration > Device Connector**.

UI Element	Essential Information
Device ID	Enter the applicable Device ID. <ul style="list-style-type: none"> For a UCS Domain, use the serial number of the primary and subordinate FIs, in this format: (Serial number of FI-A & Serial number of FI-B). Example: [SAL1924GKV6&SAL1913CJ7V] For a standalone server, use serial number. Example: NGTR12345 For HyperFlex, use Cluster UUID. Example: xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx. For Cisco UCS Director, use Device ID. Example Example: xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx.
Claim Code	Enter device claim code. You can find this code in the Device Connector for the target type. <p>Note Before you gather the Claim Code, ensure that the Device Connector has outbound network access to Cisco Intersight, and is in the “Not Claimed” state.</p>

- Step 4** Click **Claim**.

- Note** Refresh the Targets page to view the newly claimed target.

Verify Firmware Version for Fabric Interconnect

In Cisco UCS Manager, from **Equipment > Firmware Management > Installed Firmware** tab, verify for the correct firmware version.

For a complete list of hardware and software inter-dependencies, refer to respective UCSM release version using the [UCS Hardware and Software Compatibility tool](#).

Configure HyperFlex Fabric Interconnect Clusters

To configure a HyperFlex Fabric Interconnect Cluster in Intersight, do the following:

- Step 1** Log in to Intersight with HyperFlex Cluster administrator or Account Administrator privileges.
- Step 2** Navigate to **CONFIGURE > Profiles**.
- Step 3** In the **Profiles** page, make sure that the **HyperFlex Cluster Profiles** tab is selected, and click **Create HyperFlex Cluster Profile** to launch the Create HyperFlex Cluster Profile installation wizard.
- Step 4** Select **Standard** as the deployment type. Click **Start**.
- Step 5** In the **General** page, complete the following fields:

Field	Description
Organization drop-down list	<p>You can make the HyperFlex Cluster Profile belong to the default organization or a specific organization. Choose:</p> <ul style="list-style-type: none"> • default—To make the Cluster Profile belong to the default organization. All the policies that belong to the default organization will be listed on the Create HyperFlex Cluster Profile wizard. • Specific Organization—To make the HyperFlex Cluster Profile belong to the specified organization only. Only the policies that belong to the selected organization will be listed on the Create HyperFlex Cluster Profile wizard. <p>For example, if HyperFlex nodes are shared across two organizations and are associated to a Cluster Profile in one organization, you cannot associate the same node to a Cluster Profile in another organization. The Cluster Profile will be available only to users who belong to the specified Organization.</p>
Name field	<p>Enter a name for the HyperFlex cluster.</p> <p>The cluster name will be used as the vCenter cluster name, HyperFlex storage controller name, and HyperFlex storage cluster name.</p>

Field	Description
HyperFlex Data Platform Version drop-down list	Select the version of the Cisco HyperFlex Data Platform to be installed. This can be one of the following: <ul style="list-style-type: none"> • 5.5(1a) • 5.0(2e), 5.0(2g) <p>Note The version that you select impacts the types of HyperFlex policies that you can choose later in the configuration wizard.</p>
Server Firmware Version drop-down list	Manually enter the server firmware bundle version used for server components such as CIMC, adapters, BIOS, etc. Refer to the Release Notes for Cisco HX Data Platform for the recommended firmware for each HX Data Platform release.
(Optional) Description field	Add a description for the HyperFlex cluster profile.
(Optional) Set Tags field	A tag key.

Click **Next**.

Step 6

In the **Nodes Assignment** page, you can assign nodes now or optionally, you can choose to assign the nodes later. To Assign nodes, click the **Assign nodes** check box and select the node you want to assign. You can view the node role based on Server Personality.

You can view the node role based on Server Personality in the **Node Type** column. If you choose a node that has a *HyperFlex Compute Server* or no personality, you must ensure that the required hardware is available in the server for successful cluster deployment. For information about the Product Identification Standards (PIDs) that are supported by Cisco Intersight, see [Cisco HyperFlex HX-Series Data Sheet](#)

Important Cisco HyperFlex Fabric Interconnect cluster allows a minimum of 3 to a maximum of 32 nodes. All selected nodes should belong to the same UCS domain.

Click **Next**.

Step 7

In the **Cluster Configuration** page, complete the following fields:

Note For the various cluster configuration tasks, you can enter the configuration details or import the required configuration data from policies. To use pre-configured policies, click **Select Policy**, next to the configuration task and choose the appropriate policy from the list.

Field	Description
<i>Security</i>	
Hypervisor Admin field	Hypervisor administrator username must contain only alphanumeric characters. <p>Note Use root account for ESXi deployments.</p>

Field	Description
Hypervisor Password field	<p>Enter the Hypervisor password, this can be one of the following:</p> <p>Remember The default ESXi password of Cisco123 must be changed as part of installation. For fresh ESXi installation, ensure the checkbox for The Hypervisor on this node uses the factory default password is checked. Provide a new ESXi root password that will be set on all nodes during installation.</p> <p>If the ESXi installation has a non-default root password, ensure the checkbox The Hypervisor on this node uses the factory default password is unchecked. Provide the ESXi root password that you configured. This password will not be changed during installation.</p>
Hypervisor Password Confirmation field	Retype the Hypervisor password.
Controller VM Admin Password field	<p>Enter a user-supplied HyperFlex storage controller VM password. The password must contain a minimum of 10 characters, with at least 1 lowercase, 1 uppercase, 1 numeric, and 1 of these - _@#%&^&!* special characters.</p> <p>Note Make a note of this password as it will be used for the administrator account.</p>
Controller VM Admin Password Confirmation field	Retype the Controller VM administrator password.
<i>DNS, NTP, and Timezone</i>	
Timezone field	Select the local timezone.
DNS Servers field	Enter one or more DNS servers. A DNS server that can resolve public domains is required for Intersight.
NTP Servers field	Enter one or more NTP servers (IP address or FQDN). A local NTP server is highly recommended.
(Optional) DNS Suffix field	<p>Enter the DNS search suffix.</p> <p>This field is applicable only for Cisco HX Data Platform release 3.0 and later.</p>
<i>vCenter (Optional Policy)</i>	
vCenter Server FQDN or IP field	Enter the vCenter server FQDN or IP address.
vCenter Username field	Enter the vCenter username. For example, <i>administrator@vsphere.local</i>

Field	Description
vCenter Password field	Enter the vCenter password.
vCenter Datacenter Name field	Enter the vCenter datacenter name.
<i>Storage Configuration</i> (Optional Policy)	
VDI Optimization check box	Check this check box to enable VDI optimization (hybrid HyperFlex systems only).
Logical Availability Zones check box	<p>Logical Availability Zones configuration is recommended. This option is available for HyperFlex cluster sizes with 8 or more nodes and clusters deployed on Cisco HX Data Platform, Release 3.0 or higher.</p> <p>Note Logical Availability Zones configuration is recommended for HyperFlex Clusters with 8 or more nodes connected to Fabric Interconnect.</p>
<i>Auto Support</i> (Optional Policy)	
Auto Support check box	Check this check box to enable Auto Support.
Send Service Ticket Notifications To field	Enter the email address recipient for support tickets.
<i>Node IP Ranges</i>	
Note	This section configures the management IP pool. You must complete the management network fields to define a range of IP addresses for deployment. On the node configuration screen, these IP addresses will be automatically assigned to the selected nodes. If you wish to assign a secondary range of IP addresses for the controller VM Management network, you may optionally fill out the additional fields below. Both IP ranges must be part of the same subnet.
Management Network Starting IP field	Enter the starting IP address for the management IP pool.
Management Network Ending IP field	Enter the ending IP address for the management IP pool.
Management Network Subnet Mask field	Enter the subnet mask for the management VLAN.
Management Network Gateway field	Enter the default gateway for the management VLAN.
Controller VM Management Network Starting IP (Optional)	Enter the starting IP address for the controller VM management network.
Controller VM Management Network Ending IP (Optional)	Enter the ending IP address for the controller VM management network.
Controller VM Management Network Subnet Mask (Optional)	Enter the subnet mask for the controller VM management network.
Controller VM Management Network Gateway (Optional)	Enter the default gateway for the controller VM management network.

Field	Description
<i>Cluster Network</i>	
VM Migration VLAN Name field	Enter the VLAN name for the VM Migration. The name can be from 1 to 32 characters long and can contain a combination of alphanumeric characters, underscores, and hyphens.
VM Migration VLAN ID field	Enter the VLAN ID for the VM Migration. An ID of 0 means the traffic is untagged. The ID can be any number between 0 and 4095, inclusive.
VM Network VLAN Name field	Enter the VLAN name for the VM Network. The name can be from 1 to 32 characters long and can contain a combination of alphanumeric characters, underscores, and hyphens. You can add additional VM Network VLAN Names if necessary.
VM Network VLAN ID field	Enter the VLAN ID for the VM Network. The ID can be any number between 0 and 4095, inclusive.
KVM Starting IP field	Enter the Start IP address for out-of-band KVM access. One IP address per node is required. The range must fall in the same address range as UCS Manager management interfaces.
KVM Ending IP field	Enter the End IP address for out-of-band KVM access.
KVM Subnet Mask field	Enter the Subnet mask for the KVM.
KVM Gateway field	Enter the Gateway for the KVM.
Management Network VLAN Name field	Enter the VLAN name for the management network. The name can be from 1 to 32 characters long and can contain a combination of alphanumeric characters, underscores, and hyphens.
Management Network VLAN ID field	Enter the VLAN ID for the management network. VLAN must have access to Intersight. The name can be from 1 to 32 characters long and can contain a combination of alphanumeric characters, underscores, and hyphens.
Jumbo Frames check box	For best performance, enable Jumbo Frames.
<i>External FC Storage</i> (Optional Policy) If you want to add external storage, configure FC Storage by completing the following fields:	
Enable FC Storage check box	Check this check box to enable FC Storage.

Field	Description
VSAN A Name field	<p>The name of the VSAN for the primary Fabric Interconnect (FI-A).</p> <p>Enter the name of the first Virtual Storage Area Network. The name can be from 1 to 32 characters long and can contain a combination of alphanumeric characters, underscores, and hyphens.</p>
VSAN A ID field	<p>The unique identifier assigned to the network for the primary Fabric Interconnect (FI-A).</p> <p>Attention Do not enter VSAN IDs that are currently used on the UCS or HyperFlex system. If you enter an existing VSAN ID in the installer which utilizes UCS zoning, zoning will be disabled in your existing environment for that VSAN ID.</p> <p>Enter the ID of the first Virtual Storage Area Network. The ID can be any number between 1 and 4093, inclusive.</p>
VSAN B Name field	<p>The unique identifier assigned to the network for the subordinate Fabric Interconnect (FI-B).</p> <p>Note Do not enter VSAN IDs that are currently used on the UCS or HyperFlex system. If you enter an existing VSAN ID in the installer which utilizes UCS zoning, zoning will be disabled in your existing environment for that VSAN ID.</p> <p>Enter the name of the second Virtual Storage Area Network. The name can be from 1 to 32 characters long and can contain a combination of alphanumeric characters, underscores, and hyphens.</p>
VSAN B ID field	<p>Enter the ID of the second Virtual Storage Area Network. The ID can be any number between 1 and 4093, inclusive.</p>
WWxN Range Starting Address field	<p>Enter the start address of the WWxN range in the form of 20:00:00:25:B5:XX.</p>
WWxN Range Ending Address field	<p>Enter the end address of the WWxN range in the form of 20:00:00:25:B5:XX.</p>
<p><i>External iSCSI Storage</i> (Optional Policy)</p> <p>If you want to add external storage, configure iSCSI Storage by completing the following fields:</p>	
Enable iSCSI Storage check box	<p>Check this check box to enable iSCSI Storage.</p>

Field	Description
VLAN A Name field	Name of the VLAN associated with the iSCSI vNIC, on the primary Fabric Interconnect (FI-A). Enter the name of the first Virtual Local Area Network. The name can be from 1 to 32 characters long and can contain a combination of alphanumeric characters, underscores, and hyphens.
VLAN A ID field	ID of the VLAN associated with the iSCSI vNIC, on the primary Fabric Interconnect (FI-A). Enter the ID of the first Virtual Local Area Network. The ID can be any number between 0 and 4095, inclusive.
VLAN B Name field	Name of the VLAN associated with the iSCSI vNIC, on the subordinate Fabric Interconnect (FI-B). Enter the name of the second Virtual Local Area Network. The name can be from 1 to 32 characters long and can contain a combination of alphanumeric characters, underscores, and hyphens.
VLAN B ID field	ID of the VLAN associated with the iSCSI vNIC, on the subordinate Fabric Interconnect (FI-A). Enter the ID of the second Virtual Local Area Network. The ID can be any number between 0 and 4095, inclusive.
Proxy Setting (Optional Policy)	
Hostname field	Enter the HTTP proxy server FQDN or IP address.
Port field	Enter the proxy port number.
Username field	Enter the HTTP Proxy username.
Password field	Enter the HTTP Proxy password.
HyperFlex Storage Network	
Storage Network VLAN Name field	Enter the VLAN name for the storage network. The name can be from 1 to 32 characters long and can contain a combination of alphanumeric characters, underscores, and hyphens.

Field	Description
Storage Network VLAN ID field	<p>Enter the VLAN ID for the storage VLAN traffic. The VLAN must be unique per HyperFlex cluster. The ID can be any number between 0 and 4095, inclusive.</p> <p>Note The storage VLAN must be unique per HyperFlex cluster. This VLAN does not need to be routable and can remain layer 2 only. IP addresses from the link local range <i>169.254.x.x/24</i> are automatically assigned to storage interfaces.</p>

Click **Next**.

Step 8

In the **Nodes Configuration** page, you can view the IP and Hostname settings that were automatically assigned. Intersight will make an attempt to auto-allocate IP addresses. Complete the following fields:

Field	Description
Cluster Management IP Address field	The cluster management IP should belong to the same subnet as the Management IPs.
MAC Prefix Address field	<p>The MAC Prefix Address is auto-allocated. You can overwrite the MAC Prefix address, using a MAC Prefix address from the range 00:25:B5:00 to 00:25:B5:EF.</p> <p>Attention Ensure that the MAC prefix is unique across all clusters for successful HyperFlex cluster deployment. Intersight does a validation for duplicate MAC prefix and shows appropriate warning if any duplicate MAC prefix is found.</p>
Replication Factor radio button	<p>The number of copies of each data block written.</p> <p>The options are 2 or 3 redundant replicas of your data across the storage cluster. Replication factor 3 is the recommended option.</p>
Hostname Prefix field	The specified Hostname Prefix will be applied to all nodes.

Step 9

In the **Summary** page, you can view the cluster configuration and node configuration details. Review and confirm that all information entered is correct. Ensure that there are no errors triggered under the **Errors/Warnings** tab.

Step 10

Click **Validate and Deploy** to begin the deployment. Optionally, click **Validate**, and then click **Save & Close** to complete deployment later. The **Results** page displays the progress of the various configuration tasks.

What to do next

Monitoring cluster deployment

Check your cluster deployment progress in the following ways:

- You can remain on the **Results** page to watch the cluster deployment progress in real time.
- You can also close the current view and allow the installation to continue in the background. To return to the results screen, navigate to **CONFIGURE > Profiles > HyperFlex Cluster Profiles**, and click on the name of your cluster.
- You can see the current state of your deployment in the status column in the HyperFlex Cluster Profile Table view.
- You can also view the progress of the HyperFlex Cluster Profile deployment from the **Requests** page.

Post Installation

Post Installation Tasks

- Step 1** Confirm that the HyperFlex Cluster is claimed in Intersight.
- Step 2** Confirm that the cluster is registered to vCenter.
- Step 3** Navigate to **HyperFlex Clusters**, select your cluster and click ... to launch HyperFlex Connect.
- Step 4** SSH to the cluster management IP address and login using **admin** username and the controller VM password provided during installation. Verify the cluster is online and healthy.
- Step 5** Paste the following command in the Shell, and hit enter:
- ```
hx_post_install
```
- Step 6** Follow the on-screen prompts to complete the installation. The **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 configuration check.
-





## CHAPTER 4

# Deploy HyperFlex Datacenter Without Fabric Interconnect Clusters

---

- [DC-No-FI Overview, on page 43](#)
- [Supported Models/Versions for HyperFlex Datacenter without Fabric Interconnect Deployments, on page 45](#)
- [Preinstallation Checklist for Datacenter Without Fabric Interconnect, on page 47](#)
- [Installation, on page 59](#)
- [Post Installation, on page 67](#)

## DC-No-FI Overview

Cisco HyperFlex Datacenter without Fabric Interconnect (DC-No-FI) brings the simplicity of hyperconvergence to data center deployments without the requirement of connecting the converged nodes to Cisco Fabric Interconnect.

Starting with HyperFlex Data Platform Release 4.5(2b) and later:

- Support for DC-No-FI deployment from 3 to 12 converged nodes.
- For clusters larger than 8 nodes, it is recommended to enable Logical Availability Zones (LAZ) as part of your cluster deployment.
- Support for cluster expansion with converged and compute nodes on HyperFlex DC-No-FI clusters. For more information see [Expand Cisco HyperFlex Clusters in Cisco Intersight](#).
- The expansion of HyperFlex Edge clusters beyond 4 nodes changes the deployment type from Edge type to DC-No-FI type.
- Support for DC-No-FI as a target cluster for N:1 Replication.
- VIC-based and NIC-based clusters are supported. See the Preinstall Checklist for more details in the following section.



**Note** Starting with HXDP release 5.0(2a), DC-no-FI clusters support the HX nodes connected to different pairs of leaf switches for better redundancy and rack distribution, allowing you to scale the cluster as needed. This is supported with spine-leaf network architecture where all the HX nodes in a cluster belong to single network fabric in same datacenter.

HyperFlex Data Platform Datacenter Advantage license or higher is required. For more information on HyperFlex licensing, see *Cisco HyperFlex Software Licensing* in the [Cisco HyperFlex Systems Ordering and Licensing Guide](#).



**Note** 1:1 converged:compute ratio requires HXDP DC Advantage license or higher and 1:2 converged:compute ratio requires HXDP DC Premier license.

The Cisco Intersight HX installer rapidly deploys HyperFlex 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 HyperFlex DC-No-FI 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.

The following table summarizes the installation workflow for DC-No-FI clusters:

| Step | Description                             | Reference                                                                            |
|------|-----------------------------------------|--------------------------------------------------------------------------------------|
| 1.   | Complete the preinstallation checklist. | <a href="#">Preinstallation Checklist for Datacenter without Fabric Interconnect</a> |
| 2.   | Ensure that the network is set up.      |                                                                                      |
| 3.   | Log in to Cisco Intersight.             | <a href="#">Log In to Cisco Intersight, on page 17</a>                               |



| Step | Description                                                                     | Reference                                                                      |
|------|---------------------------------------------------------------------------------|--------------------------------------------------------------------------------|
| 4.   | Claim Targets.<br><b>Note</b> Skip if you have already claimed HyperFlex Nodes. | Claim Targets for DC-no-FI Clusters                                            |
| 5.   | Verify Cisco UCS Firmware versions.                                             | Verify Firmware Version for HyperFlex DC-No-FI Clusters                        |
| 6.   | Run the HyperFlex Cluster Profile Wizard.                                       | Configure and Deploy HyperFlex Datacenter without Fabric Interconnect Clusters |
| 7.   | Run the post installation script through the controller VM.                     | Post Installation Tasks for DC-no-FI Clusters                                  |

## Supported Models/Versions for HyperFlex Datacenter without Fabric Interconnect Deployments

The following table lists the supported hardware platforms and software versions for HyperFlex DC-No-FI clusters. For information about the Product Identification Standards (PIDs) that are supported by Cisco Intersight, see [Cisco HyperFlex HX-Series Data Sheet](#).

| Component  | Models/Versions                                                                                                                                                                                                                                                                          |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| M6 Servers | <ul style="list-style-type: none"> <li>• HXAF245C-M6SX</li> <li>• HX245C-M6SX</li> <li>• HXAF225C-M6SX</li> <li>• HX225C-M6SX</li> <li>• HXAF220C-M6SN</li> <li>• HXAF240C-M6SN</li> <li>• HX240C-M6SX</li> <li>• HXAF240C-M6SX</li> <li>• HX220C-M6S</li> <li>• HXAF220C-M6S</li> </ul> |

| Component                     | Models/Versions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| M5 Servers                    | <ul style="list-style-type: none"> <li>• HXAF220C-M5SN</li> <li>• HX220C-M5SX</li> <li>• HXAF220C-M5SX</li> <li>• HX240C-M5SX</li> <li>• HXAF240C-M5SX</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Cisco HX Data Platform (HXDP) | <ul style="list-style-type: none"> <li>• 5.5(1a)</li> <li>• 5.0(2e), 5.0(2g)</li> </ul> <p><b>Note</b></p> <ul style="list-style-type: none"> <li>• HXDP versions 5.0(2a), 5.0(2b), 5.0(2c), 5.0(2d), 4.5(2a), 4.5(2b), 4.5(2c), 4.5(2d), and 4.5(2e) and are still supported for cluster expansion only.</li> <li>• Upgrades from HXDP 4.0.2x are supported provided the ESXi version is compatible with 4.5(2x).</li> <li>• M6 servers require HXDP 5.0(1a) or later.</li> <li>• M5SN servers require HXDP 4.5(2c) or later.</li> <li>• Compute only nodes are supported on M5/M6 rack servers.</li> </ul> |
| NIC Mode                      | <p>This can be one of the following:</p> <ul style="list-style-type: none"> <li>• Dedicated Management Port</li> <li>• Shared LOM</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Device Connector              | Auto-upgraded by Cisco Intersight                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Network Topologies            | <ul style="list-style-type: none"> <li>• M5 Servers—10/25/40 GE</li> <li>• M6 Servers—10/25/40/100 GE</li> </ul> <p><b>Note</b> Greater than 10 GE is recommended for All NVMe clusters.</p>                                                                                                                                                                                                                                                                                                                                                                                                                 |

| Component         | Models/Versions                                                                                                                                           |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Connectivity Type | Types: <ul style="list-style-type: none"> <li>• VIC based</li> <li>• NIC-based (10G+ NIC-based clusters require HXDP version 5.0(2a) or later)</li> </ul> |

## Preinstallation Checklist for Datacenter Without Fabric Interconnect

Ensure that your system meets the following installation and configuration requirements before you begin to install a Cisco HyperFlex Datacenter Without Fabric Interconnect (DC-No-FI) system.

### 10/25/40/100 Gigabit Ethernet Topology and IMC Connectivity (VIC-based)

Cisco HyperFlex Data Center 3-Node to 12-Node DC-no-FI clusters are deployed through Cisco Intersight. Cisco Intersight provides advanced multi-cluster monitoring and management capabilities and the topology supports 10/25/40/100GE installation, and dual ToR switch options for ultimate network flexibility and redundancy.

Cisco recommends the 10/25/40/100 GE topology for best performance and future node expansion capabilities.

The 10/25/40/100 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 10/25/40/100 GE switch may be two standalone switches or may be formed as a switch stack.

Use the following Cisco IMC Connectivity option for the 3-Node to 12-Node 10/25/40/100 Gigabit Ethernet (GE) topology:

- Use of shared LOM extended mode (EXT). In this mode, single wire management is used and Cisco IMC traffic is multiplexed onto the 10/25/40/100GE 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 DC-no-FI 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.

### VIC-based Physical Network and Cabling for 10/25/40/100 GE Topology

A managed switch 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 10/25/40/100 GE ports, one 1GE port for CIMC management, and one Cisco VIC 1457 per server. Trunk ports are the only supported network port configuration.

Select dual switch configuration to continue with physical cabling:

### 10/25/40/100 Gigabit Ethernet Dual Switch Physical Cabling (VIC-based)



**Warning** Proper cabling is important to ensure full network redundancy.

Dual switch configuration provides 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 10/25/40/100GE 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 10/25/40/100GE ports and 1 x 1GE port for each HyperFlex node.
  - All 10/25/40/100GE 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 but is recommended.
- 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 10/25/40/100GE 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 10/25/40/100GE 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 10/25/40/100GE ports prior to cluster installation. After cluster deployment, you may optionally use the additional two 10/25/40/100GE ports for guest VM traffic.

## 3 to 12 Node - Dual 10GE ToR Switches


**Virtual Networking Design for 3- to 12-Node 10/25/40/100 Gigabit Ethernet Topology (VIC-based)**

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.

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



## 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.

## 10/25/40/100 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

---

### Jumbo Frames for 10/25/40/100 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 10/25/40/100 GE 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.

### 10GBASE-T Copper Support

HX supports the use of Cisco copper 10G transceivers (SFP-10G-T-X) for use with switches that have 10G copper (RJ45) ports. In all of the 10GE topologies listed in this chapter, supported twinax, fiber, or 10G copper transceivers may be used. For more information on supported optics and cables, see the [Cisco UCS Virtual Interface Card 1400/14000 Series Data Sheet](#).

When using SFP-10G-T-X transceivers with HyperFlex, the following limitations apply:

- Minimum Cisco IMC firmware version 4.1(3d) and HyperFlex Data Platform version 4.5(2b).
- Maximum of two SFP-10G-T-X may be used per VIC. Do not use the additional two ports.
- The server must not use Cisco Card or Shared LOM Extended NIC modes. Use the Dedicated or Shared LOM NIC modes only.

## 10 or 25GE NIC-Based Topology and IMC Connectivity

The 10 or 25 Gigabit Ethernet (GE) switch NIC-based 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.

The 10 or 25 Gigabit Ethernet (GE) network interface card (NIC)-based topology (two standalone switches or can be a switch stack) is an option in place of a VIC-based topology. Both NIC- and VIC-based topologies provide a fully redundant design that protects against switch (if using dual or stacked switches) link and port



failures. The 10/25GE switches may be two standalone switches or may be formed as a switch stack. Before you consider deploying a NIC-based topology, consider the following requirements and supported hardware.

The following requirements and hardware must be considered before starting deployment:

- NIC-based deployment is supported on HXDP release 5.0(2a) and later
- VMware ESXi 7.0 U3 or later
- NIC-Based cluster is supported for Intersight deployment only and requires an Intersight Essentials License
- NIC-Based HX deployments are supported with HX 220/225/240/245 M6 nodes only.
- Support for Edge and DC-no-FI clusters only • 10/25GE Dual Top of Rack (ToR) Switches.
- One Intel 710/810 quad port NIC or two Intel 710/810 series dual port NICs installed on Cisco HX hardware. Supported NIC options are;
  - Intel X710-DA2 Dual Port 10Gb SFP+ NIC (HX-PCIE-ID10GF)
  - Intel X710 Quad-port 10G SFP+ NIC (HX-PCIE-IQ10GF)
  - Cisco-Intel E810XXVDA2 2x25/10 GbE SFP28 PCIe NIC (HX-P-I8D25GF)
  - Cisco-Intel E810XXVDA4L 4x25/10 GbE SFP28 PCIe NIC (HX-P-I8Q25GF)
  - Cisco-Intel X710T2LG 2x10 GbE RJ45 PCIe NIC (HX-P-ID10GC)

Cisco HyperFlex Data Center 3-Node to 12-Node DC-no-FI clusters are deployed through Cisco Intersight. Cisco Intersight provides advanced multi-cluster monitoring and management capabilities and the topology supports 10/25GE installation, and dual ToR switch options for ultimate network flexibility and redundancy.



---

**Note** Mixing VIC-based and NIC-based topologies in the same cluster is not supported.

---

Cisco recommends the 10/25 GE topology for best performance and future node expansion capabilities.

The 10/25 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 10/25 GE switch may be two standalone switches or may be formed as a switch stack.



---

**Note** NIC-Based HyperFlex DC-no-FI clusters support only 10/25GE uplink connectivity.

---

Use the following Cisco IMC Connectivity option for the 3-Node to 12-Node 10/25 Gigabit Ethernet (GE) topology:

- 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 DC-no-FI 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.

## NIC-based Physical Network and Cabling for 10/25 GE Topology

Two managed switches with VLAN capability are 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 either one Intel 710/810 series quad port NIC or two Intel 710/810 series dual port NICs two 10/25/40/100 GE ports. Trunk ports are the only supported network port configuration.

Select dual switch configuration to continue with physical cabling:

### 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)
- One Intel Quad port NIC or two Intel dual port NICs installed in the PCIE slots as below:
  - HX 220/225 Nodes: Use PCIE slot 1 for quad port NIC or use PCIE slots 1 and 2 for dual port
  - NICs HX 240/245 Nodes: Use PCIE slot 4 for quad port NIC or use PCIE slot 4 & 6 for dual port NICs

### Upstream Network Requirements

- Two managed switches with VLAN capability (standalone or stacked).
- 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 but is recommended.
- 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.

### NIC-based 10/25 Gigabit Ethernet Dual Switch Physical Cabling



---

**Warning** Proper cabling is important to ensure full network redundancy.

---

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



---

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

---

- Connect the first NIC port (going from left) from each node to first ToR switch (switchA).
- Connect the second NIC port (going from left) from each node to the second ToR switch (switchB).
- Connect the third NIC port (going from left) from each node to first ToR switch (switchA).
- Connect the fourth NIC port (going from left) from each node to the second ToR switch (switchB).



---

**Note** Use the same port number on each server to connect to the same switch. Refer to topology diagram below for connectivity details.

---

- Do not connect LOM ports or any additional ports prior to cluster installation. After cluster deployment, you may optionally use the additional ports for guest VM traffic




---

**Note** Please follow the above guidelines for cabling. Deviating from the above recommendation may fail the cluster deployment.

---


### Network Cabling Diagram for 1 x Quad Port NIC

### 3 or greater Node Edge/DC-no-FI - Dual 10/25GE



#### Network Cabling Diagram for 2 x Dual Port NICs

### 3 or greater Node Edge/DC-no-FI - Dual 10/25GE



#### NIC-based Virtual Networking Design for 3- to 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.

Four vSwitches are required:

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

- **vswitch-hx-storage-data**—ESXi storage interface (vmk1), HX storage controller data network

### 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.

## 10/25 GE NIC-based 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.
  - This VLAN should be configured as the trunk VLAN on all the switch ports connected to port 1 and port 2 from left on each node.
  - 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. This VLAN should be configured as a trunk VLAN on all the switch ports connected to port 3 and port 4 from the left on each node.
- 1 VLAN for vMotion traffic. This can be an isolated and non-routed VLAN. In a NIC-Based HX cluster, the vSwitch vswitch-hx-inband-mgmt is used for vMotion and guest VM networking. So, the VLANs used for vMotion and guest VM networking should be trunked on all switch ports connected to port 1 and port 2 from the left on each node.




---

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

---

- Switch ports connected to the NICs in a NIC-based cluster should be operating at dedicated 10/25GE speed.
- 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 NIC cards may be installed in the HX Edge nodes as needed.
- All other VIC or NIC cards in slots other than 1 and 2 in HX 220/225 nodes, or slots 4 and 6 in HX 240/245 nodes, must be shut down or left un-cabled until installation is complete.

# Installation

## Log In to Cisco Intersight

### Log In using Cisco ID

To login to Cisco Intersight, you must have a valid **Cisco ID** to create a Cisco Intersight account. If you do not have a Cisco ID, create one [here](#).



**Important** The device connector does not mandate the format of the login credentials, they are passed as is to the configured HTTP proxy server. Whether or not the username must be qualified with a domain name will depend on the configuration of the HTTP proxy server.

### Log In using Single Sign-On

Single Sign-On (SSO) authentication enables you to use a single set of credentials to log in to multiple applications. With SSO authentication, you can log in to Intersight with your corporate credentials instead of your Cisco ID. Intersight supports SSO through SAML 2.0, and acts as a service provider (SP), and enables integration with Identity Providers (IdPs) for SSO authentication. You can configure your account to sign in to Intersight with your Cisco ID and SSO. Learn more about SSO with Intersight [here](#).

## Claim DC-No-FI Targets

Complete the following steps to claim one or more Targets to be managed by Cisco Intersight:

### Before you begin

This procedure assumes that you are an existing user with a Cisco account. If not, see [Log In to Cisco Intersight, on page 17](#).

**Step 1** In the Cisco Intersight, left navigation pane, select **ADMIN > Targets**.

**Step 2** In the **Targets** details page, click **Claim a New Target**.

**Step 3** In the **Claim a New Target** wizard, select **All > Cisco UCS Server (Standalone)** and complete the following fields:

**Note** You can locate the **Device ID** and the **Claim Code** information in Cisco IMC by navigating to **Admin > Device Connector**.

| UI Element | Essential Information                                                                                                |
|------------|----------------------------------------------------------------------------------------------------------------------|
| Device ID  | Enter the applicable Device ID. For a Cisco UCS C-Series Standalone server, use serial number.<br>Example: NGTR12345 |

| UI Element | Essential Information                                                                                                                                                                                                                                                         |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Claim Code | <p>Enter target claim code. You can find this code in the Device Connector for the target type.</p> <p><b>Note</b> Before you gather the Claim Code, ensure that the Device Connector has outbound network access to Cisco Intersight, and is in the “Not Claimed” state.</p> |

**Step 4** Click **Claim**.

**Note** Refresh the Targets page to view the newly claimed target.

## Verify Firmware Version for HyperFlex DC-No-FI Clusters

View current BIOS, CIMC, SAS HBA, and drive firmware versions, and verify that those versions match the Cisco HyperFlex Edge and Firmware Compatibility Matrix in the *Common Network Requirements*. Refer to the [Preinstallation Checklist for Datacenter Without Fabric Interconnect, on page 47](#) for 3-Node and 12-Node DC-No-FI clusters for more details.

- Step 1** In your browser, log in to the CIMC web UI by navigating to **https://<CIMC IP>**. You can also cross launch CIMC from Cisco Intersight from the **Servers** tableview.
- 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 > Storage**. Double-click on **Cisco 12G Modular SAS HBA (MRAID)** and navigate to **Details > Physical Drive Info**.
- Step 6** Compare the current BIOS, CIMC, SAS HBA, and drive firmware versions with the versions listed in the Cisco HyperFlex Edge and Firmware Compatibility Matrix in the *Common Network Requirements*. Refer to the [Preinstallation Checklist for Datacenter Without Fabric Interconnect, on page 47](#) for 3-Node and 12-Node DC-No-FI clusters for more details.
- Step 7** 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>.

## Configure HyperFlex Datacenter Without Fabric Interconnect Clusters

To configure a HyperFlex Datacenter without Fabric Interconnect (DC-No-FI) cluster in Intersight, do the following:



- Step 1** Log in to Intersight with HyperFlex Cluster administrator or Account Administrator privileges.
- Step 2** Navigate to **CONFIGURE > Profiles**.
- Step 3** In the **Profiles** page, make sure that the **HyperFlex Cluster Profiles** tab is selected, and click **Create HyperFlex Cluster Profile** to launch the Create HyperFlex Cluster Profile installation wizard.
- Step 4** Select **Data Center** as the deployment type and uncheck the **Use Fabric Interconnect** box. Click **Start**.
- Step 5** In the **General** page, complete the following fields:

| Field                                                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Organization</b> drop-down list                    | <p>You can make the HyperFlex Cluster Profile belong to the default organization or a specific organization. Choose:</p> <ul style="list-style-type: none"> <li>• <b>default</b>—To make the Cluster Profile belong to the default organization. All the policies that belong to the default organization will be listed on the Create HyperFlex Cluster Profile wizard.</li> <li>• <b>Specific Organization</b>—To make the HyperFlex Cluster Profile belong to the specified organization only. Only the policies that belong to the selected organization will be listed on the Create HyperFlex Cluster Profile wizard.</li> </ul> <p>For example, if HyperFlex nodes are shared across two organizations and are associated to a Cluster Profile in one organization, you cannot associate the same node to a Cluster Profile in another organization. The Cluster Profile will be available only to users who belong to the specified Organization.</p> |
| <b>Name</b> field                                     | <p>Enter a name for the HyperFlex cluster.</p> <p>The cluster name will be used as the vCenter cluster name, HyperFlex storage controller name, and HyperFlex storage cluster name.</p> <p><b>Note</b> The name of the HyperFlex Cluster Profile belonging to an organization must be unique. You may create a HyperFlex Cluster Profile with the same name in a different organization.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>HyperFlex Data Platform Version</b> drop-down list | <p>Select the version of the Cisco HyperFlex Data Platform to be installed. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• 5.5(1a)</li> <li>• 5.0(2e), 5.0(2g)</li> </ul> <p><b>Note</b> The version that you select impacts the types of HyperFlex policies that you can choose later in the configuration wizard.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

| Field                               | Description                                          |
|-------------------------------------|------------------------------------------------------|
| (Optional) <b>Description</b> field | Add a description for the HyperFlex cluster profile. |
| (Optional) <b>Set Tags</b> field    | Enter a tag key.                                     |

Click **Next**.

### Step 6

In the **Nodes Assignment** page, you can assign nodes now or optionally, you can choose to assign the nodes later. To assign nodes, click the **Assign nodes** check box and select the node you want to assign.

You can view the node role based on Server Personality in the **Node Type** column. If you choose a node that has a *HyperFlex Compute Server* or no personality, you must ensure that the required hardware is available in the server for successful cluster deployment. For information about the Product Identification Standards (PIDs) that are supported by Cisco Intersight, see [Cisco HyperFlex HX-Series Spec Sheet](#)

**Important** Cisco HyperFlex DC-No-FI cluster only allows a minimum of 3 to a maximum of 12 nodes.

For clusters larger than 8 nodes, it is recommended to enable Logical Availability Zones (LAZ) as part of your cluster deployment.

Click **Next**.

### Step 7

In the **Cluster Configuration** page, complete the following fields:

**Note** For the various cluster configuration tasks, you can enter the configuration details or import the required configuration data from policies. To use pre-configured policies, click **Select Policy**, next to the configuration task and choose the appropriate policy from the list.

| Field                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Security</i>                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hypervisor Admin</b> field    | Enter the Hypervisor administrator username.<br><b>Note</b> Use root account for ESXi deployments.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hypervisor Password</b> field | Enter the Hypervisor password, this can be one of the following:<br><b>Remember</b> The default ESXi password of Cisco123 must be changed as part of installation. For fresh ESXi installation, ensure the checkbox for <b>The Hypervisor on this node uses the factory default password is checked</b> . Provide a new ESXi root password that will be set on all nodes during installation.<br><br>If the ESXi installation has a non-default root password, ensure the checkbox <b>The Hypervisor on this node uses the factory default password</b> is unchecked. Provide the ESXi root password that you configured. This password will not be changed during installation. |

| Field                                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hypervisor Password Confirmation</b> field          | Retype the Hypervisor password.                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Controller VM Admin Password</b> field              | Enter a user-supplied HyperFlex storage controller VM password.<br><br><b>Important</b> Make a note of this password as it will be used for the administrator account.                                                                                                                                                                                                                                                                       |
| <b>Controller VM Admin Password Confirmation</b> field | Retype Controller VM administrator password.                                                                                                                                                                                                                                                                                                                                                                                                 |
| <i>DNS, NTP, and Timezone</i>                          |                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Timezone</b> field                                  | Select the local timezone.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>DNS Suffix</b> field                                | Enter the suffix for the DNS. This is applicable only for HX Data Platform 3.0 and later.                                                                                                                                                                                                                                                                                                                                                    |
| <b>DNS Servers</b> field                               | Enter one or more DNS servers. A DNS server that can resolve public domains is required for Intersight.                                                                                                                                                                                                                                                                                                                                      |
| <b>NTP Servers</b> field                               | Enter one or more NTP servers (IP address or FQDN). A local NTP server is highly recommended.                                                                                                                                                                                                                                                                                                                                                |
| <i>vCenter (Optional Policy)</i>                       |                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>vCenter Server FQDN or IP</b> field                 | Enter the vCenter server FQDN or IP address.                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>vCenter Username</b> field                          | Enter the vCenter username. For example, <i>administrator@vsphere.local</i>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>vCenter Password</b> field                          | Enter the vCenter password.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>vCenter Datacenter Name</b> field                   | Enter the vCenter datacenter name.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <i>Storage Configuration (Optional Policy)</i>         |                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>VDI Optimization</b> check box                      | Check this check box to enable VDI optimization (hybrid HyperFlex systems only).                                                                                                                                                                                                                                                                                                                                                             |
| <i>Auto Support (Optional Policy)</i>                  |                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Auto Support</b> check box                          | Check this check box to enable Auto Support.                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Send Service Ticket Notifications To</b> field      | Enter the email address recipient for support tickets.                                                                                                                                                                                                                                                                                                                                                                                       |
| <i>Node IP Ranges</i>                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Note</b>                                            | This section configures the management IP pool. You must complete the management network fields to define a range of IPs for deployment. On the node configuration screen, these IPs will be automatically assigned to the selected nodes. If you wish to assign a secondary range of IPs for the controller VM Management network, you may optionally fill out the additional fields below. Both IP ranges must be part of the same subnet. |

| <b>Field</b>                                                         | <b>Description</b>                                                |
|----------------------------------------------------------------------|-------------------------------------------------------------------|
| <b>Management Network Starting IP</b> field                          | The starting IP address for the management IP pool.               |
| <b>Management Network Ending IP</b> field                            | The ending IP address for the management IP pool.                 |
| <b>Management Network Subnet Mask</b> field                          | The subnet mask for the management VLAN.                          |
| <b>Management Network Gateway</b> field                              | The default gateway for the management VLAN.                      |
| <b>Controller VM Management Network Starting IP</b> field (Optional) | The starting IP address for the controller VM management network. |
| <b>Controller VM Management Network Ending IP</b> field (Optional)   | The ending IP address for the controller VM management network.   |
| <b>Controller VM Management Network Subnet Mask</b> field (Optional) | The subnet mask for the controller VM management network.         |
| <b>Controller VM Management Network Gateway</b> field (Optional)     | The default gateway for the controller VM management network.     |
| <i>Cluster Network</i>                                               |                                                                   |

| Field                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Uplink Speed field                     | <p>The Uplink speed is 10G+. Refer to "Preinstallation Checklist for Datacenter Without Fabric Interconnect" for more details of the supported Network Topologies.</p> <p><b>Attention</b> Using 10G+ mode typically requires the use of forward error correction (FEC) depending on the transceiver or the type &amp; 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 the 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 <a href="#">Cisco UCS C-Series Integrated Management Controller GUI Configuration Guide</a> for further details on how to change the FEC mode configured on the VIC using the Cisco IMC GUI.</p> |
| Management Network VLAN ID field       | <p>Enter the VLAN ID for the management network. VLAN must have access to Intersight.</p> <p>An ID of 0 means the traffic is untagged. The VLAN ID can be any number between 0 and 4095, inclusive.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Jumbo Frames check box                 | <p>Check this check box to enable Jumbo Frames.</p> <p>Jumbo Frames are optional and can remain disabled for HyperFlex DC-NO-FI deployments.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <i>Proxy Setting</i> (Optional Policy) |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Hostname field                         | Enter the HTTP proxy server FQDN or IP address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Port field                             | Enter the proxy port number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Username field                         | Enter the HTTP Proxy username.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Password field                         | Enter the HTTP Proxy password.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

| Field                                | Description                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>HyperFlex Storage Network</i>     |                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Storage Network VLAN ID</b> field | <p>Enter the VLAN ID for the storage VLAN traffic. The VLAN must be unique per HyperFlex cluster.</p> <p><b>Note</b> The storage VLAN must be unique per HyperFlex cluster. This VLAN does not need to be routable and can remain layer 2 only. IP addresses from the link local range <i>169.254.0.0/16</i> are automatically assigned to storage interfaces.</p> |

Click **Next**.

### Step 8

In the **Nodes Configuration** page, you can view the IP and Hostname settings that were automatically assigned. Intersight will make an attempt to auto-allocate IP addresses. Complete the following fields:

| Field                                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Cluster Management IP Address</b> field | The cluster management IP should belong to the same subnet as the Management IPs.                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>MAC Prefix Address</b> field            | <p>The MAC Prefix Address is auto-allocated for NIC-based HyperFlex Edge clusters. For 10G+ HyperFlex Edge clusters you can overwrite the MAC Prefix address, using a MAC Prefix address from the range 00:25:B5:00 to 00:25:B5:EF.</p> <p><b>Attention</b> Ensure that the MAC prefix is unique across all clusters for successful HyperFlex cluster deployment. Intersight does a validation for duplicate MAC prefix and shows appropriate warning if any duplicate MAC prefix is found.</p> |
| <b>Replication Factor</b> radio button     | <p>The number of copies of each data block written. The options are 2 or 3 redundant replicas of your data across the storage cluster.</p> <p><b>Important</b> Replication factor 3 is the recommended option.</p>                                                                                                                                                                                                                                                                              |
| <b>Hostname Prefix</b> field               | The specified Hostname Prefix will be applied to all nodes.                                                                                                                                                                                                                                                                                                                                                                                                                                     |

### Step 9

In the **Summary** page, you can view the cluster configuration and node configuration details. Review and confirm that all information entered is correct. Ensure that there are no errors triggered under the **Errors/Warnings** tab.

### Step 10

Click **Validate and Deploy** to begin the deployment. Optionally, click **Validate**, and click **Save & Close** to complete deployment later. The **Results** page displays the progress of the various configuration tasks. You can also view the progress of the HyperFlex Cluster Profile deployment from the **Requests** page.

## What to do next

### Monitoring cluster deployment

Check your cluster deployment progress in the following ways:

- You can remain on the **Results** page to watch the cluster deployment progress in real time.
- You can also close the current view and allow the installation to continue in the background. To return to the results screen, navigate to **CONFIGURE > Profiles > HyperFlex Cluster Profiles**, and click on the name of your cluster.
- You can see the current state of your deployment in the status column in the HyperFlex Cluster Profile Table view.
- Once deployed, the cluster deployment type is displayed as DC-No-FI.

# Post Installation

## Post Installation Tasks

- 
- Step 1** Confirm that the HyperFlex Cluster is claimed in Intersight.
- Step 2** Confirm that the cluster is registered to vCenter.
- Step 3** Navigate to **HyperFlex Clusters**, select your cluster and click **...** to launch HyperFlex Connect.
- Step 4** SSH to the cluster management IP address and login using **admin** username and the controller VM password provided during installation. Verify the cluster is online and healthy.
- Step 5** Paste the following command in the Shell, and hit enter:
- ```
hx_post_install
```
- Step 6** Follow the on-screen prompts to complete the installation. The **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 configuration check.
-



CHAPTER 5

HyperFlex Cluster Profile Operations

- [Cloning HyperFlex Cluster Profiles, on page 69](#)
- [Unassigning HyperFlex Cluster Profiles, on page 69](#)

Cloning HyperFlex Cluster Profiles

You can clone a successfully deployed HyperFlex Cluster Profile, and then use that copy as the basis to create new clusters. To clone a HyperFlex Cluster Profile, do the following:

-
- Step 1** Navigate to **CONFIGURE > Profiles > HyperFlex Cluster Profiles**. Click on the Ellipsis icon (...) and select **Clone**.
- Step 2** In the **Clone HyperFlex Cluster Profile** dialog box, enter the name for the cloned cluster, select the **Organization**, and choose the number of clones. Click **Clone**.
- Note** For multiple clones of a profile, a consecutive index is appended to the given clone name.
- Step 3** The **HyperFlex Cluster Profiles** page now contains a copy of the cloned cluster. Before deploying the cloned cluster you can review the configuration, and make any necessary changes to the policies.
-

Unassigning HyperFlex Cluster Profiles

Unassigning removes all assigned servers from the cluster profile. This action is specific to Intersight and will not result in a factory reset or affect the operation of the running cluster. To fully clean up this cluster from Intersight, please unclaim the HyperFlex cluster and associated servers.

-
- Step 1** Navigate to **CONFIGURE > Profiles > HyperFlex Cluster Profiles** page and identify the cluster to unassign.
- Step 2** Click on the Ellipsis icon (...) and select **Unassign**.
- Step 3** Confirm the Unassign operation.
-



CHAPTER 6

Appendix

- [Troubleshooting](#), on page 71
- [Sample Network Configurations](#), on page 72

Troubleshooting

Symptom	Possible Cause	Resolution
HyperFlex Device Connector is unable to connect to Intersight. Device Connector UI displays the type of error and corrective actions to be taken under Details and Recommendations	No outbound network access, NTP or DNS misconfiguration.	Use the following commands to check connectivity from the CIMC. <pre>\$ scope CIMC \$ scope network \$ ping cisco.com \$ # Ensure that you receive a response \$ ping svc.ucs-connect.com \$ # Ensure that an IP address is displayed. \$</pre> <p>If an IP address is not shown, DNS is not resolving. This IP address will not respond to pings and should only be used to test DNS name resolution. Use the Cisco.com ping test to confirm outbound connectivity.</p> <p>Ensure time is synced via NTP.</p>
Failed to validate the SSL certificate for intersight.com.	Improper network access or DNS may be misconfigured.	SSH to one of the controller VMs and ensure network access via ping to Cisco.com.

Sample Network Configurations

1GE Single Switch

Nexus 5548 using trunk ports

```

vlan 101
  name HX-MGMT
vlan 102
  name HX-STORAGE
vlan 103
  name HX-vMOTION
vlan 104
  name HX-GUESTVM
...
interface Ethernet2/11
  description HX-01-Port1
  switchport mode trunk
  switchport trunk allowed vlan 101-104
  spanning-tree port type edge trunk
  speed 1000
interface Ethernet2/12
  description HX-01-Port2
  switchport mode trunk
  switchport trunk allowed vlan 101-104
  spanning-tree port type edge trunk
  speed 1000
interface Ethernet2/13
  description HX-02-Port1
  switchport mode trunk
  switchport trunk allowed vlan 101-104
  spanning-tree port type edge trunk
  speed 1000
interface Ethernet2/14
  description HX-02-Port2
  switchport mode trunk
  switchport trunk allowed vlan 101-104
  spanning-tree port type edge trunk
  speed 1000
interface Ethernet2/15
  description HX-03-Port1
  switchport mode trunk
  switchport trunk allowed vlan 101-104
  spanning-tree port type edge trunk
  speed 1000
interface Ethernet2/16
  description HX-03-Port2
  switchport mode trunk
  switchport trunk allowed vlan 101-104
  spanning-tree port type edge trunk
  speed 1000

```

Catalyst 3850-48T using trunk ports

```

vlan 101
  name HX-MGMT
vlan 102
  name HX-STORAGE
vlan 103
  name HX-vMOTION

```

```

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

```

1GE Dual Switch

Nexus 5548 using trunk ports

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

```

ip dhcp relay
...
interface Vlan105
  ip address 10.1.2.1/24
  ip dhcp relay address 10.1.1.2
  no shutdown
vlan 101
  name HX-MGMT
vlan 102
  name HX-STORAGE
vlan 103
  name HX-VMOTION
vlan 104
  name HX-GUESTVM
vlan 105
  name HX-DHCP-CIMC

```

```

...
interface Ethernet2/11
  description HX-01-Port1
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk
  speed 1000
interface Ethernet2/12
  description HX-01-Port2
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk
  speed 1000
interface Ethernet2/13
  description HX-02-Port1
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk
  speed 1000
interface Ethernet2/14
  description HX-02-Port2
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk
  speed 1000
interface Ethernet2/15
  description HX-03-Port1
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk
  speed 1000
interface Ethernet2/16
  description HX-03-Port2
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 101-105
  spanning-tree port type edge trunk
  speed 1000

```

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

Catalyst 3850-48T using trunk ports

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

```

vlan 101
  name HX-MGMT
vlan 102
  name HX-STORAGE
vlan 103
  name HX-vmOTION
vlan 104
  name HX-GUESTVM
vlan 105
  name HX-CIMC
...

```

```

interface GigabitEthernet1/0/1
  description HX-01-Port1
  switchport mode trunk
  speed 1000
  spanning-tree portfast trunk
interface GigabitEthernet1/0/2
  description HX-01-Port2
  switchport mode trunk
  speed 1000
  spanning-tree portfast trunk
interface GigabitEthernet1/0/3
  description HX-02-Port1
  switchport mode trunk
  speed 1000
  spanning-tree portfast trunk
interface GigabitEthernet1/0/4
  description HX-02-Port2
  switchport mode trunk
  speed 1000
  spanning-tree portfast trunk
interface GigabitEthernet1/0/5
  description HX-03-Port1
  switchport mode trunk
  speed 1000
  spanning-tree portfast trunk
interface GigabitEthernet1/0/6
  description HX-03-Port2
  switchport mode trunk
  speed 1000
  spanning-tree portfast trunk

```

Repeat the same configuration on switch #2.

10GE Dual Switch

Nexus 9000 using trunk ports

```

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

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

interface Ethernet1/37

```

```

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

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

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

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

```

Catalyst 9300 using trunk ports

```

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

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

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

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

```



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

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

