



Deploy HyperFlex Datacenter Without Fabric Interconnect Clusters

- [DC-No-FI Overview, on page 1](#)
- [Supported Models/Versions for HyperFlex Datacenter without Fabric Interconnect Deployments, on page 3](#)
- [Preinstallation Checklist for Datacenter Without Fabric Interconnect, on page 5](#)
- [Installation, on page 17](#)
- [Post Installation, on page 25](#)

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.	Preinstallation Checklist for Datacenter without Fabric Interconnect
2.	Ensure that the network is set up.	
3.	Log in to Cisco Intersight.	Log In to Cisco Intersight

Step	Description	Reference
4.	Claim Targets. Note 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"> • HXAF245C-M6SX • HX245C-M6SX • HXAF225C-M6SX • HX225C-M6SX • HXAF220C-M6SN • HXAF240C-M6SN • HX240C-M6SX • HXAF240C-M6SX • HX220C-M6S • HXAF220C-M6S

Component	Models/Versions
M5 Servers	<ul style="list-style-type: none"> • HXAF220C-M5SN • HX220C-M5SX • HXAF220C-M5SX • HX240C-M5SX • HXAF240C-M5SX
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) and 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. • M5SN servers require HXDP 4.5(2c) or later. • Compute only nodes are supported on M5/M6 rack servers.
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	<ul style="list-style-type: none"> • M5 Servers—10/25/40 GE • M6 Servers—10/25/40/100 GE <p>Note Greater than 10 GE is recommended for All NVMe clusters.</p>

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

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

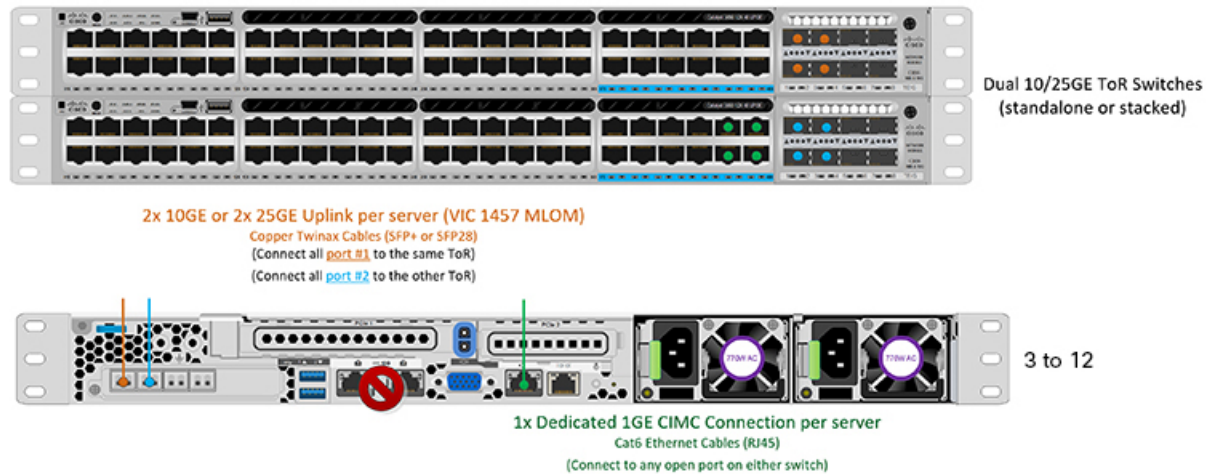
- Two managed switches with VLAN capability (standalone or stacked).
- 2 x 10/25/40/100GE ports and 1 x 1GE port for each HyperFlex node.
 - All 10/25/40/100GE ports must trunk and allow all applicable VLANs. All 1GE ports may be trunked or in access mode when connected to the dedicated CIMC port.
- 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

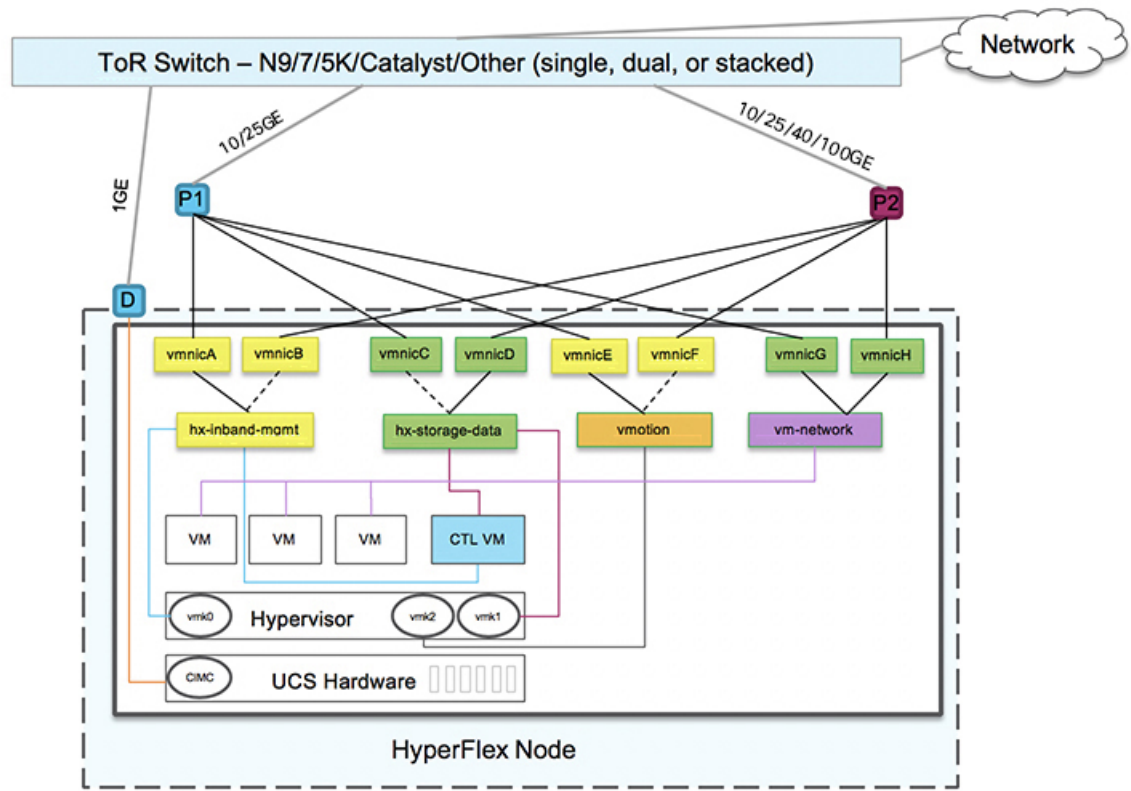
**Virtual Networking Design for 3- to 12-Node 10/25/40/100 Gigabit Ethernet Topology (VIC-based)**

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

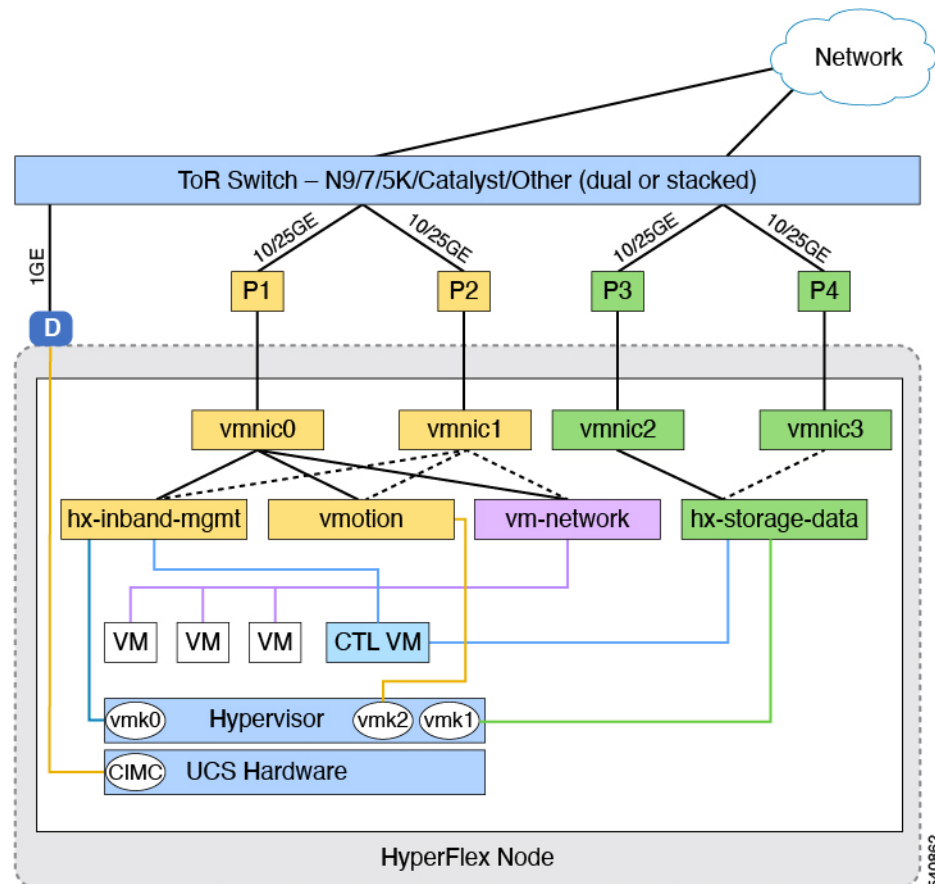
This section details the virtual network setup. No action is required as all of the virtual networking is set up automatically by the HyperFlex deployment process. These extra details are included below for informational and troubleshooting purposes.

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](#).

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. 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>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 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 5](#) 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 5](#) 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
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>

Field	Description
(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 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>	
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.

Field	Description
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.
<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.

Field	Description
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.
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 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>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 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 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>
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>	
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.

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 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.
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.

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.
-

