



Preinstallation Checklist for 3- and 4-Node Edge Deployments

Cisco recommends the use of Cisco Intersight for all HyperFlex Edge deployments to ensure a seamless global management experience. Cisco HyperFlex Edge 3-Node and 4-Node clusters may be deployed either through Cisco Intersight, or an on-premises installer VM. Cisco HyperFlex Edge 2-Node clusters require the use of Cisco Intersight for initial deployment and ongoing operations using the Invisible Cloud Witness. HyperFlex Edge 2-Node clusters cannot be deployed using the on-premises HyperFlex installer VM.

Cisco HyperFlex Edge offers both a 1 Gigabit Ethernet (GE) and a 10/25GE installation option. Both topologies support single top-of-rack (ToR) and dual ToR switch options for ultimate network flexibility and redundancy. A network topology is chosen during initial deployment and cannot be changed or upgraded without a full reinstallation. Choose your network topology carefully and with future needs in mind. Consider the following when determining the best topology for your cluster:

- Higher performance and future node expansion capabilities: Select the 10/25GE topology. You can choose Cisco VIC-based hardware or Intel NIC-Based adapters.
- Clusters that will never require node expansion, and instances where the ToR switch does not have 10GE ports available: Select the 1GE topology

[Selecting your 3- or 4-Node Network Topology, on page 1](#)

- [3- and 4-Node Network Topology, on page 1](#)

3- and 4-Node Network Topology

Selecting your 3- or 4-Node Network Topology

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

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

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

- [10 or 25GE VIC-Based Topology, on page 2](#)
- [10 or 25GE NIC-Based Topology, on page 8](#)
- [1 Gigabit Ethernet Topology, on page 13](#)

After completing the 10/25GE or 1GE ToR physical network and cabling section below, continue with the [Common Network Requirement Checklist, on page 21](#).

10 or 25GE VIC-Based Topology

10 or 25GE VIC-Based Topology

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

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

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

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

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

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

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

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

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

Requirements for both 10 and 25GE Topologies

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

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

Requirements for HX Edge clusters using 25GE



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

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

10/25GE VIC-Based Dual Switch Physical Cabling



Warning Proper cabling is important to ensure full network redundancy.

Dual switch configuration provides a slightly more complex topology with full redundancy that protects against: switch failure, link failure, and port failure. It requires two switches, that may be standalone or stacked, and 2 x 10/25GE ports, 1 x 1GE port (dedicated CIMC), and 1 x Cisco VIC 1457 MLOM card for each HyperFlex node. Trunk ports are the only supported network port configuration.

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

Upstream Network Requirements


- Two managed switches with VLAN capability (standalone or stacked)
- 2 x 10/25GE ports and 1 x 1GE port for each HyperFlex node.
All 10/25GE ports must trunk and allow all applicable VLANs. All 1GE ports may be trunked or in access mode when connected to the dedicated CIMC port.
- Jumbo frames are not required to be configured
- Portfast trunk should be configured on all ports to ensure uninterrupted access to Cisco Integrated Management Controller (CIMC)
- If using dedicated Cisco IMC, connect the 1GE management port on each server (Labeled M on the back of the server) to one of the two switches, or to an out-of-band management switch.
- Connect one out of the four 10/25GE ports on the Cisco VIC from each server to the same ToR switch.
 - Use the same port number on each server to connect to the same switch.



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

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

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



10/25GE VIC-Based Single Switch Physical Cabling




Warning Proper cabling is important to ensure full network redundancy.

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

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

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

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




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

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

Virtual Switches

Four vSwitches are required:

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

Network Topology:

502781

Failover Order:

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

10/25GE VIC-Based Switch Configuration Guidelines

3 VLANs are required at a minimum.

- 1 VLAN for the following connections: VMware ESXi management, Storage Controller VM management and Cisco IMC management.
 - VMware ESXi management and Storage Controller VM management must be on the same subnet and VLAN.
 - A dedicated Cisco IMC management port may share the same VLAN with the management interfaces above or may optionally use a dedicated subnet and VLAN. If using a separate VLAN, it must have

L3 connectivity to the management VLAN above and must meet Intersight connectivity requirements (if managed by Cisco Intersight).

- If using shared LOM extended mode for Cisco IMC management, a dedicated VLAN is recommended.
- 1 VLAN for Cisco HyperFlex storage traffic. This can and should be an isolated and non-routed VLAN. It must be unique and cannot overlap with the management VLAN.
- 1 VLAN for vMotion traffic. This can be an isolated and non-routed VLAN.



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

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



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

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



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

Additional Considerations:

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

Jumbo Frames for VIC-based 10/25GE

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

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

Next Steps:

Complete the [Common Network Requirement Checklist](#), on page 21.

10 or 25GE NIC-Based Topology

10 or 25GE NIC-Based Topology

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 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.

Requirements for NIC-Based Topology

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 IMC Connectivity for 10/25GE NIC-Based Topology

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

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

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

Physical Network and Cabling for 10/25GE NIC-Based 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 four 10/25GE ports, one 1GE port for CIMC management, and one quad port or two dual port NICs per server. Trunk ports are the only supported network port configuration.

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 & 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

Next Step:

After completing the 10/25GE or 1GE ToR physical network and cabling section, continue with the [Common Network Requirement Checklist, on page 21](#).

10/25GE NIC-Based Dual Switch Physical Cabling



Warning Proper cabling is important to ensure full network redundancy.

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

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



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 the 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 Follow the guidelines above for cabling. Deviating from the recommendations above may result in cluster deployment failure.



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

Network Cabling Diagram for 1 x Quad Port NIC

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



Dual 10/25GE ToR Switches (standalone or stacked)

Intel X710 quad-port 10G / Cisco-Intel E810XXVDA4L 4x25/10G



x3 or x4 for edge
x3 and up to 12 for DC-no-FI

- Legend
- HX/ESXi Management
 - HX storage data
 - CIMC port

Network Cabling Diagram for 2 x Dual Port NICs

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



Dual 10/25GE ToR Switches (standalone or stacked)

Intel X710-DA2 Dual Port 10G / Cisco-Intel E810XXVDA2 2x25/10 G



x3 or x4 for edge
x3 and up to 12 for DC-no-FI

- Legend
- HX/ESXi Management
 - HX storage data
 - CIMC port

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

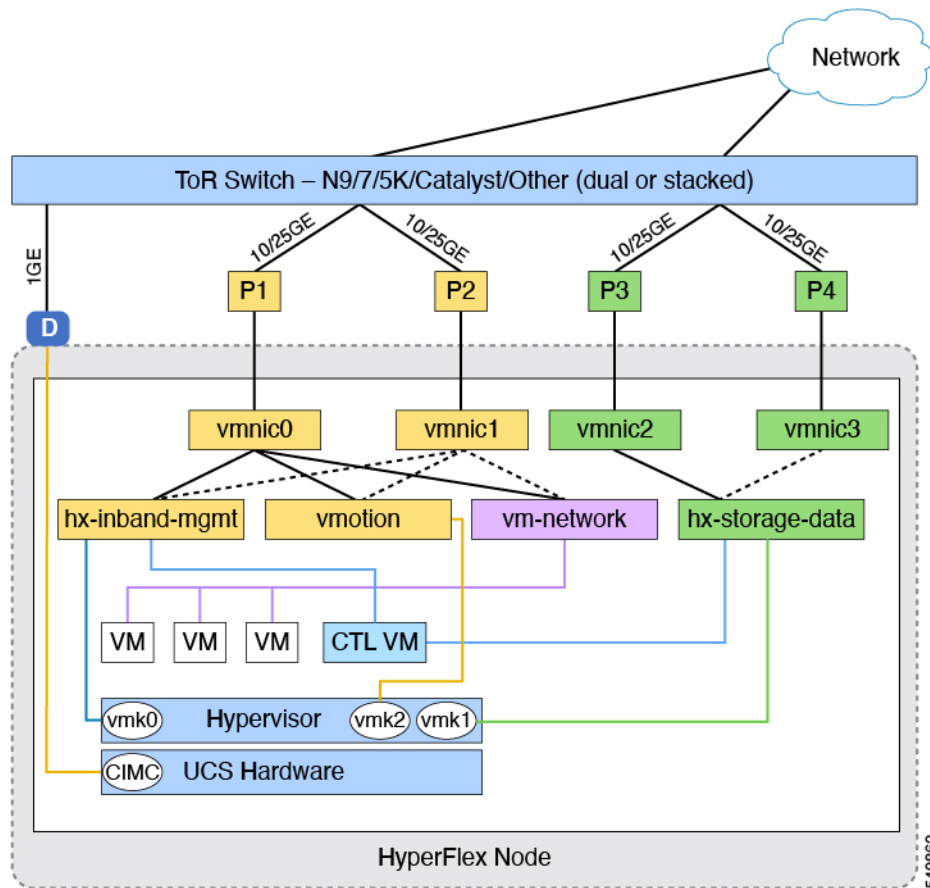
This section details the virtual network setup. No action is required as all of the virtual networking is set up automatically by the HyperFlex deployment process. These extra details are included below for informational and troubleshooting purposes.

Virtual Switches:

Four vSwitches are required:

- **vswitch-hx-inband-mgmt**—ESXi management (vmk0), storage controller management network, 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.

Jumbo Frames for NIC-Based 10/25GE

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

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

Next Steps:

Complete the [Common Network Requirement Checklist, on page 21](#).

1 Gigabit Ethernet Topology

1 Gigabit Ethernet Topology

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

Cisco IMC Connectivity for 1 Gigabit Ethernet Topology

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

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

- In fabric interconnect-based environments, built-in QoS ensures uninterrupted access to Cisco IMC and server management when using single wire management. In HyperFlex Edge environments, QoS is not enforced and hence the use of a dedicated management port is recommended.

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

Physical Network and Cabling for 1GE Topology

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

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

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



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

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

1 Gigabit Ethernet Dual Switch Cabling



Warning Proper cabling is important to ensure full network redundancy.

The following requirements must be met before starting deployment:

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

- Intel x550 Lan-on-motherboard LOM (built into Cisco UCS motherboard)
 - 2 x 1GE ToR switch ports and 2 x Category 6 Ethernet Cables per HyperFlex server (customer supplied)



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

To deploy with dual ToR switches for extra redundancy:

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



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


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



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

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

3 & 4 Node Edge – Dual 1GE ToR Switches



1 Gigabit Ethernet Single Switch Cabling



Warning Proper cabling is important to ensure full network redundancy.

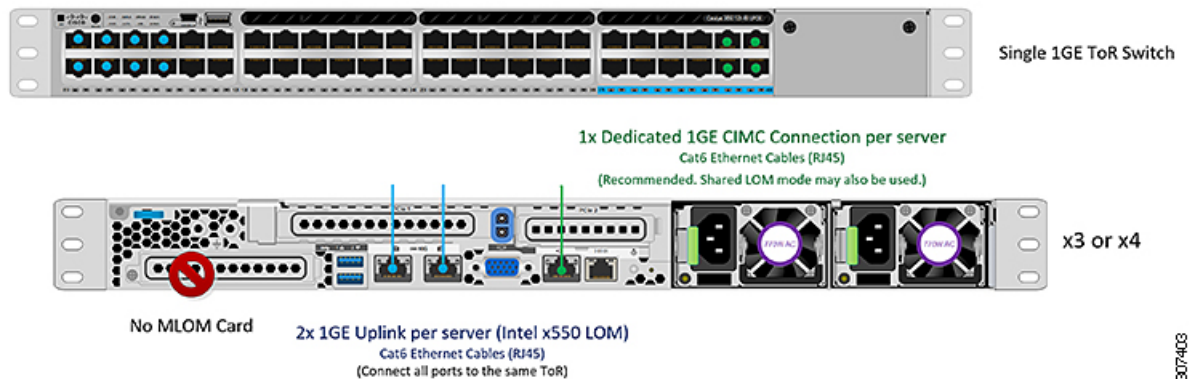
To deploy with a single ToR (see diagram below for a visual layout):

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



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

3 & 4 Node Edge – Single 1GE ToR Switch



About Access and Trunk Ports

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

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

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



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

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



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

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

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

Virtual Switches

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

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

The minimum network configuration requires at least two separate networks:

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


Two vSwitches each carrying different networks are required:

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




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

Network Topology: Dual Switch Configuration



Network Topology: Single Switch Configuration



Failover Order - Dual switch configuration only:

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

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

1 Gigabit Ethernet Switch Configuration Guidelines

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



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

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

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

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



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

Jumbo Frames for 1 Gigabit Ethernet

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

- The option to enable jumbo frames is only provided during initial install and cannot be changed later.
- Jumbo Frames are a best practice, but are not required. If opting out of jumbo frames, leave the MTU set to 1500 bytes on all network switches.
- For highest performance, jumbo frames may be optionally enabled. Ensure full path MTU is 9000 bytes or greater. Keep the following considerations in mind when enabling jumbo frames:
 - When running a dual switch setup, it is imperative that all switch interconnects and switch uplinks have jumbo frames enabled. Failure to ensure full path MTU could result in a cluster outage if traffic is not allowed to pass after link or switch failure.

- The HyperFlex installer will perform a one-time test on initial deployment that will force the failover order to use the standby link on one of the nodes. If the switches are cabled correctly, this will test the end to end path MTU. Do not bypass this warning if a failure is detected. Correct the issue and retry the installer to ensure the validation check passes.
- For these reasons and to reduce complexity, it is recommended to disable jumbo frames when using a dual switch setup.
- The option to enable jumbo frames is found in the HyperFlex Cluster profile, under the Network Configuration policy. Checking the box will enable jumbo frames. Leaving the box unchecked will keep jumbo frames disabled.

Next Steps:

Complete the [Common Network Requirement Checklist, on page 21](#).

10GBASE-T Copper Support

HX Edge 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](#).

Limitations

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

- Minimum Cisco IMC firmware version 4.1(3d) and HyperFlex Data Platform version 4.5(2a).
- 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.

Common Network Requirement Checklist

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

VLAN Requirements



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

Network	VLAN ID	Description
Use a separate subnet and VLANs for each of the following networks:		

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

Supported vCenter Topologies

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

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

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

3-Node Customer Deployment Information

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



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

4-Node Customer Deployment Information

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



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

CIMC Management IP Addresses

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

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

Network IP Addresses



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



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

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

VMware vMotion Network IP Addresses

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

Server	vMotion Network IP Addresses (configured using the post_install script)
Server 1:	
Server 2:	
Server 3:	

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

VMware vCenter Configuration



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

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

Port Requirements



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

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

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

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



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

Network Services



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

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

Connected Services

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

Proxy Server

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

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

Guest VM Traffic

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

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

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

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



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

Intersight Connectivity

Consider the following prerequisites pertaining to Intersight connectivity:

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

- IP connectivity (L2 or L3) is required from the CIMC management IP on each server to all of the following: ESXi management interfaces, HyperFlex controller VM management interfaces, and vCenter server. Any firewalls in this path should be configured to allow the necessary ports as outlined in the [Hyperflex Hardening Guide](#).
- When redeploying HyperFlex on the same servers, new controller VMs must be downloaded from Intersight into all ESXi hosts. This requires each ESXi host to be able to resolve svc.intersight.com and allow outbound initiated HTTPS connections on port 443. Use of a proxy server for controller VM downloads is supported and can be configured in the HyperFlex Cluster Profile if desired.
- Post-cluster deployment, the new HX cluster is automatically claimed in Intersight for ongoing management.

Cisco HyperFlex Edge Invisible Cloud Witness

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

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

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

Ordering Cisco HyperFlex Edge Servers

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