Revised: September 1, 2026

Cisco Nexus Hyperfabric — Switches

Configuring the switches and ports

The Blueprint Designer provides the basic design and port configuration for the fabric. During the fabric design phase or after fabric deployment, you can perform general configuration tasks. For example, you can enable host ports for your network devices, create routed ports for external connections, or add sub-interfaces. To configure and apply your changes, ensure that you are in edit mode and that you commit your changes.

Modify a switch configuration

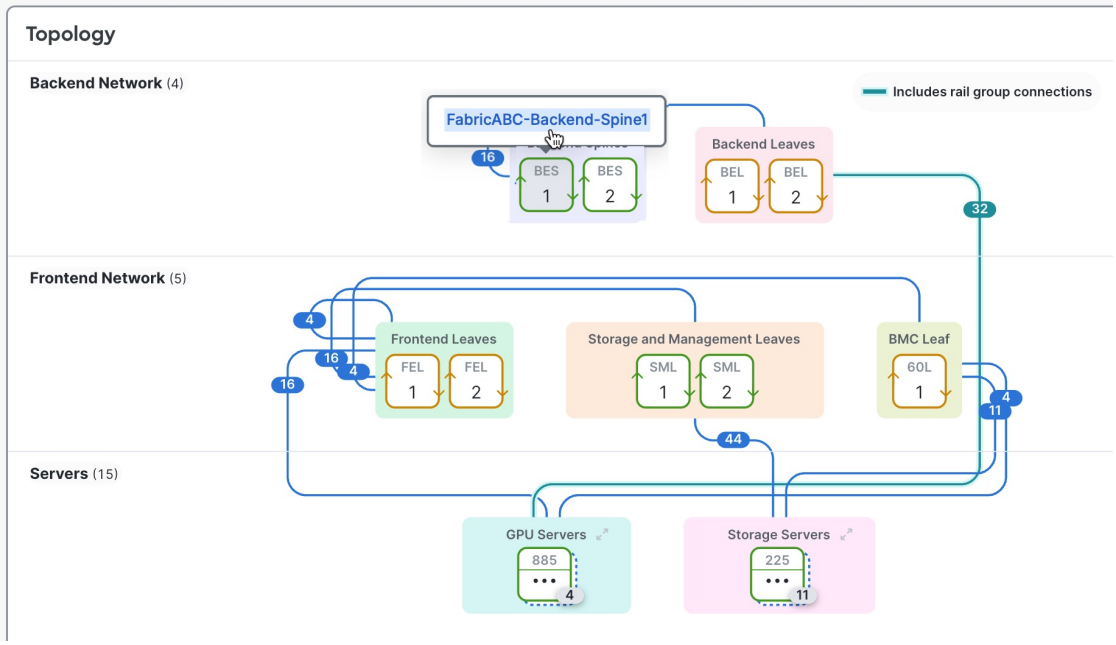
You can edit the switch connections, ports, VLANs, and route tables (VRFs) for a specific switch.

Step 1 Navigate to the switch you are interested in.

a) Choose **Fabrics**, then select the fabric that has the switch you want to configure.

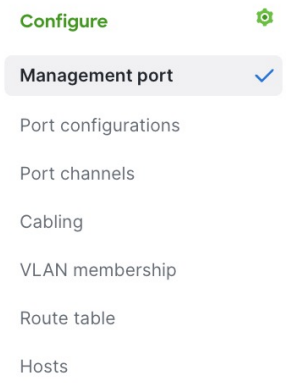
b) In the **Topology** area, select the switch you want to view, then select the switch name. You may have to expand the switch group in order to see the switch.

Figure 1: Switch in Topology area



Step 2 In the **Configure** area, select a switch property to view or modify.

Figure 2: Switch properties



Step 3 Configure the selected switch property.

The configuration procedures are described in these links.

- Management port: Allows you to configure management port IP address, DNS, proxy, and the switch console password.
- Port configurations: Allows you to configure port role, speed, and port breakouts.
- Port channels: Allows you to configure port channels to combine multiple physical Ethernet links into a single logical link.
- Cabling: Allows you to configure port switch connections within the fabric.
- VLAN membership: Allows you to configure VLAN members to be mapped to a logical network.
- Route table: Allows you to add notes and labels to route tables.
- **Hosts**: Lists hosts discovered from this switch.

Configure the management interface of a switch

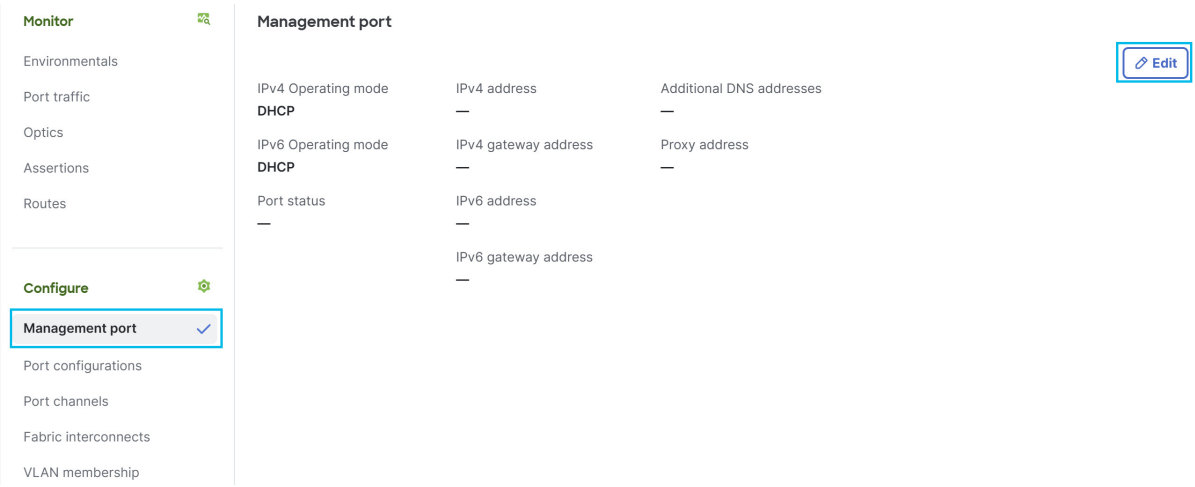
You typically configure the management interface and Craft console password prior to onboarding, but you may configure or modify these parameters after onboarding. You may want to modify the management interface configuration to update the IP addressing scheme, improve security settings, or adjust operational parameters.

Step 1 Choose **Fabrics**, then select the fabric that has the switch you want to modify.

Step 2 In the **Topology** area, select the switch you want to configure, then select the switch name.

Step 3 In the **Configure** area, select **Management port**.

Figure 3: Management port



Step 4 Select **Edit**.

Step 5 In the **Edit management ports** dialog box, select how the management port obtains IP addresses.

At least one IPv4 or IPv6 address is required for cloud connectivity.

a) To obtain an IPv4 address by DHCP, toggle on **Automatically obtain IP addresses** under **IPv4**. Otherwise, enter an IPv4 address and IPv4 gateway.

b) To obtain an IPv6 address by DHCP, toggle on **Automatically obtain IP addresses** under **IPv6**. Otherwise, enter an IPv6 address and IPv6 gateway.

c) If you do not set **Automatically obtain IP addresses** , you must provide one or more **Additional DNS addresses**. You can enter multiple addresses as a comma-separated list.

Step 6 (Optional) Enter **Proxy address** , **Proxy username** , and **Proxy password**.

Step 7 (Optional) In the **Forgot your console password?** area, enter the new password in both fields, then select **Update**. For more information, see Change the Craft console password of a switch, on page 13.

Step 8 Select **Update**.

Managing switch ports

Dynamic port breakout

Dynamic port breakout enables you to divide a high-speed, channelized port on a network element into multiple low-speed ports. Each low-speed port can connect to a different network element, maximizing the use of the high-speed port's bandwidth. Breakout is suitable for very short links and offers a cost effective way to connect within racks and across adjacent racks. For example, you can split a 400 Gigabit (Gb) port into four independent and logical 100Gb ports.

You configure breakout on the down links (also known as the access-facing ports or downlink ports) and fabric links of the switches. Fabric links form the connections between switches.

For how to configure a port breakout, see Configure a port at the switch level, on page 5.

View the ports of a switch

After you configure the ports of a switch, you can look at the status of all ports. At a glance, you can see which ports are up, not connected, disabled, or configured for breakout.

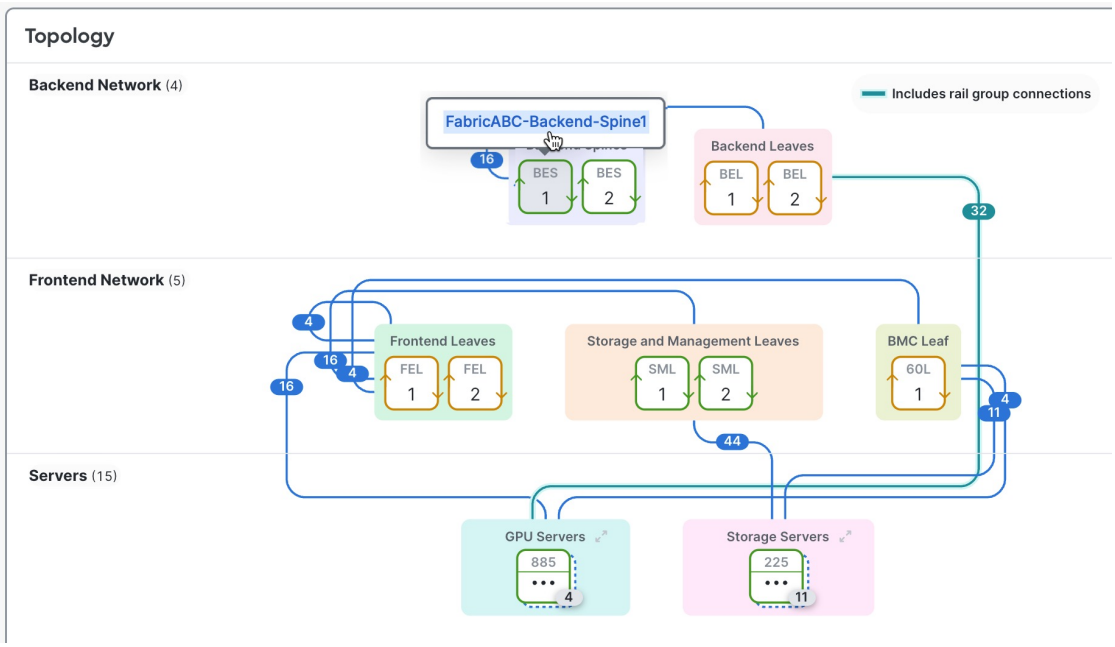
Follow these steps to view the ports of a switch.

Step 1 Navigate to the switch you are interested in.

a) Choose **Fabrics**, then select the fabric that has the switch you want to configure.

b) In the **Topology** area, select the switch you want to view, then select the switch name. You may have to expand the switch group in order to see the switch.

Figure 4: Switch in Topology area



Step 2 In the **Ports** area, select **Port status** or **Face plate**.

Figure 5: Ports area



- **Port status**: This view shows all ports of the device, as well as the status of each port. The icon next to the port number indicates whether the port is up or configured for breakout. The port image itself indicates whether the port is enabled or disabled. To understand the icons, see the legend in the upper right of the **Ports** area.

  A green color indicates that a port has no issues. A yellow color indicates that a port has an unexpected condition and there are issues detected, but the issues should not cause operation disruption. A red color indicates that a port has an unexpected condition and there are issues detected that can cause operation disruption; you should investigate the issues. See Assertions.

  If a port is configured for breakout, the port shows the broken out ports as well as the color status for each breakout port. You can select the broken out port to see an enlarge view of the ports, which enables you to see the icon of each port.

- **Face plate**: This view shows a more realistic representation of the ports. This view is not interactive.

Step 3 To see a list of all ports on the switch, in the **Configure** area, select **Port configurations**. See Configure a port at the switch level.

Step 4 To see details about a specific port:

a) In the **Port status** view, select that port. This can be a port in the enlarged view of a broken out port.

The **Details** drawer opens, which contains more information about the port. You can perform various actions, such as disabling the port, changing the admin state, or editing the description.

b) Select the ***port-interface-name*** **details** link to view more information about the port.

c) In the **Details** area, select the port property you want to view.

Guidelines for configuring ports

With the Cisco HF6100-60L4D switch, when you configure the speed of one of the SFP56 ports in a port group, the speed of the other port gets configured with the same speed.

Configure a port at the switch level

You can configure various properties of a switch's port, such as the role and speed, and breakout status. The port role specifies the type of connection provided by a port.

Follow these steps to configure a port at the switch level.

Step 1 Choose **Fabrics**, then select the fabric that contains the switch.

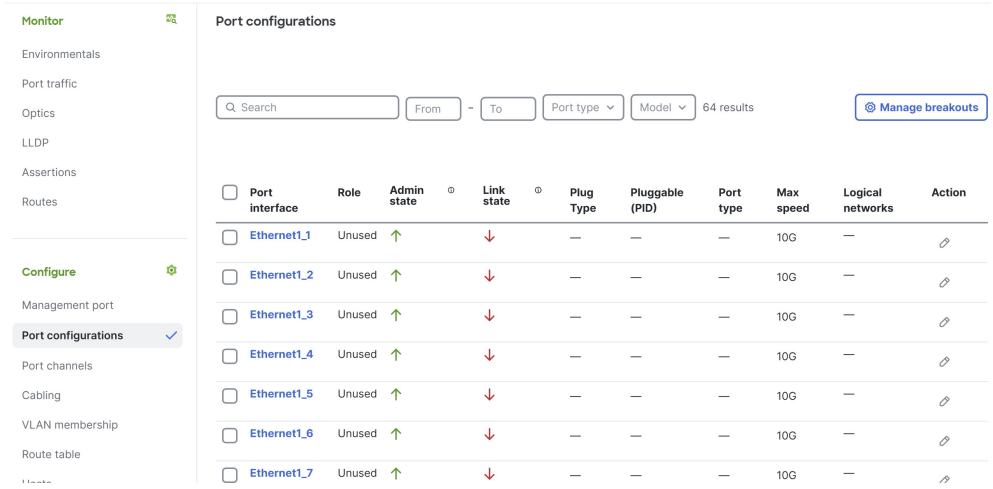
a) If the fabric is not in the edit mode, select **Switch to edit mode** ( ).

Step 2 In the **Topology** area, select the switch you want to configure, then select the switch name.

Step 3 In the **Configure** area, select **Port configurations**.

The **Port configurations** table lists all ports of the switch.

Figure 6: Port configurations



Step 4 (Optional) If you want to enable or disable breakout for any of the ports, select **Manage breakouts**.

a) For **Breakout option**, select **Disable** to disable breakout. Otherwise, select the desired option to enable breakout.

The value in parentheses indicates the speed of each lane and the number of lanes per port. For example, the value "2x 200G(50G x 4 lanes)" indicates 2 ports of 200G speed each and each port uses four 50G lanes. This means a total of 400G speed and 8 lanes.

You must select an option that corresponds to the speed of pluggables on the connected switch.

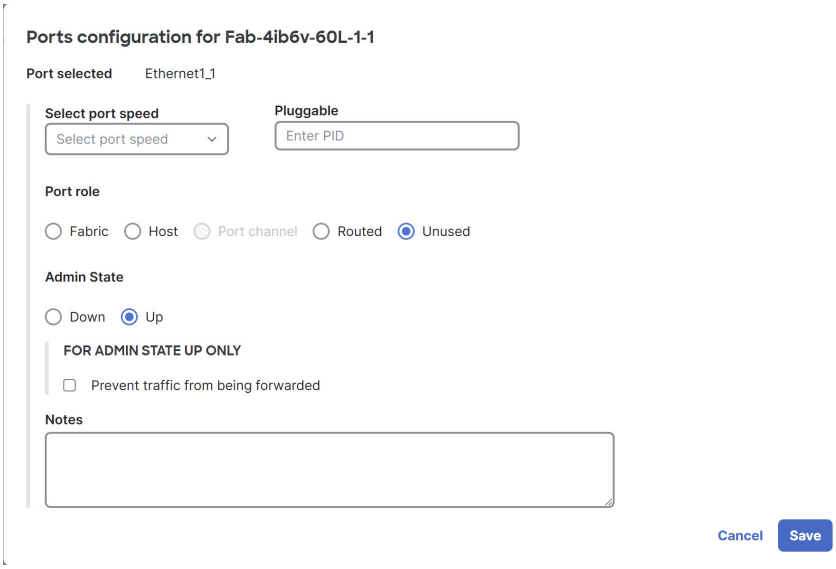
b) If you enabled breakout, select the cable type and optic type.

c) Select **Save**.

Step 5 In the **Port configurations** table, under the **Action** column, select edit ( ) for the port that you want to configure.

A dialog box opens for configuring the port.

Figure 7: Configuring a port



a) For **Select port speed**, select the speed of the port.

The value in parentheses indicates the speed of each lane and the number of lanes. For example, the value "200G(50G x 4 lanes)" indicates that the port has a total speed of 200G and uses four 50G lanes.

By default, a port transmits and receives data at its highest speed. You can reduce the speed if necessary. You cannot change the speed of a breakout port.

To use a non-default speed for your switch fabric ports, you must select the appropriate speed. For example, for QSFP-DD ports, you can use 40G, 100G, or 200G instead of the default of 400G.

b) For **Pluggable**, enter the product ID (PID) of the optic for the port.

Forwarding error correction (FEC) is automatically set based on the specified pluggable. You cannot configure FEC manually.

c) For **Port role**, select **Fabric** for a port that connects to another leaf or spine switch in the fabric.

For other port types, select the role according to the purpose of the port.

- **Fabric**: Provides connectivity between fabric switches and allows for automatic discovery by peer switches. Select this for a port that connects to another leaf or spine switch in the fabric.
- **Host**: Provides a Layer 2 connection to a server or other general network device.
- **Port channel**: The port is a member of a port channel, also known as a link aggregation group (LAG). When you configure a port to be a member of a port channel, its role is automatically changed from **Unused** to **Port channel**. The **Port channel** role cannot be changed in this menu unless the port is first removed from the port channel.
- **Routed**: Provides a Layer 3 connection to a router or other network device.

  When you configure a routed port on a switch, you can select to enable VLAN tagging, which requires you to configure 802.1Q VLAN sub-interfaces on a Layer 3 interface to forward IPv4 and IPv6 packets to another device using static or dynamic routing protocols. You can use Layer 3 interfaces for IP routing. If you leave VLAN tagging disabled, select a route table (VRF) and specify IPv4 or IPv6 addresses for the routes. Refer to the succeeding substeps for the additional routed port configuration steps.

- **Unused**: The port does not forward or receive traffic.

d) If you chose **Routed** for the port role and you want to use VLAN tagging, **Enable VLAN tagging** and follow these additional steps.

1. Select **Add a sub interface**.
2. In **VLAN tag**, enter the tag.

   Do not assign the same VLAN tag to multiple sub-interfaces on the same physical interface.
3. For **Select VRF**, select a VRF instance.
4. For the **IP address** fields, enter an IPv4 address and mask, an IPv6 address, or both.

e) If you chose **Routed** for the port role and you do not want to use VLAN tagging, disable **Enable VLAN tagging** and perform these additional steps.

1. For **Select route table**, select the desired route table (VRF).
2. For the **IP addresses** fields, enter one or more IPv4 or IPv6 addresses and masks.

   If you enter multiple addresses, separate each one with a comma.

f) For **Admin state**, select **Up** to enable the port.

- **Up**: The port will receive and forward traffic. This is the default state.

  To set the port to be Up but not forwarding traffic, check **Prevent traffic from being forwarded**.
- **Down**: The port will not receive nor forward traffic.

g) Select **Save**.

Step 6 Repeat the previous step for every port on this switch that will connect to another leaf or spine switch or host in the fabric.

Step 7 Repeat this procedure for every leaf or spine switch in the fabric.

Configure multiple ports at the switch level

You can configure various properties of multiple ports of a switch simultaneously, such as the role and speed. The port role specifies the type of connection provided by a port.

Some property choices reveal additional properties that you cannot set with this procedure. To configure these additional properties, you must configure the ports individually. The descriptions in each step indicate which properties have this limitation.

Follow these steps to configure multiple ports at the switch level.

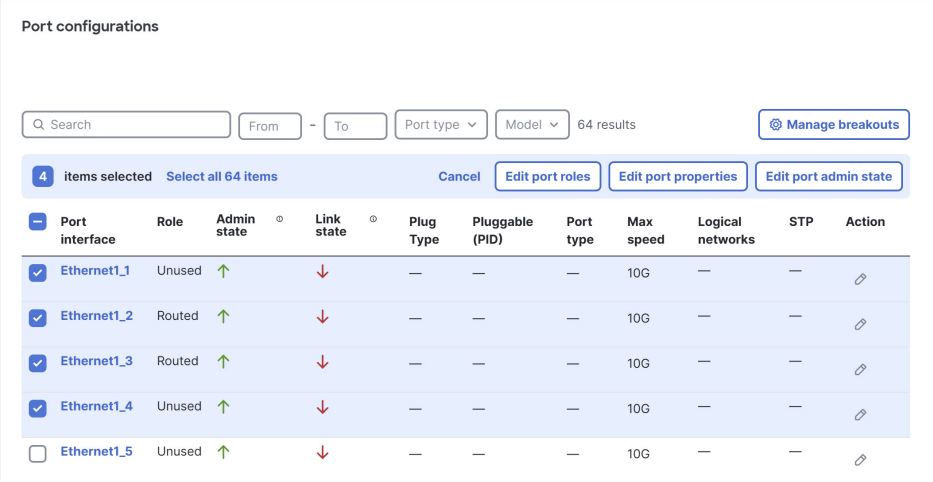
Step 1 Choose **Fabrics**, then select the fabric that contains the switch.

a) If the fabric is not in the edit mode, select **Switch to edit mode** ( ).

Step 2 In the **Topology** area, select the switch you want to configure, then select the switch name.

Step 3 In the **Configure** area, select **Port configurations**, then select the ports that you want to configure by putting a check in the box next to each port ID.

Figure 8: Multiple ports selected



Step 4 To change the role of the ports, select **Edit port roles**.

a) For **Port role**, select the desired role.

- **Fabric**: Provides connectivity between fabric switches and allows for automatic discovery by peer switches. Select this for a port that connects to another leaf or spine switch in the fabric.
- **Host**: Provides a Layer 2 connection to a server or other general network device. If you enabled global spanning tree, you cannot select the STP blocking modes using this procedure.
- **Port channel**: The port is a member of a port channel, also known as a link aggregation group (LAG). When you configure a port to be a member of a port channel, its role is automatically changed from **Unused** to **Port channel**. The **Port channel** role cannot be changed in this menu unless the port is first removed from the port channel.
- **Routed**: Provides a Layer 3 connection to a router or other network device. You cannot enable nor disable VLAN, add sub interfaces, select a route table, nor specify IP addresses using this procedure.
- **Unused**: The port does not forward nor receive traffic. If you intend to configure the port to be a port channel member, the role must initially be **Unused**.

b) Select **Save**.

Step 5 To change port properties, select **Edit port properties**. These properties include port speed and pluggable options.

a) From the **Configure port speed** drop-down list, select the speed of the port.

The value in parentheses indicates the speed of each lane and the number of lanes. For example, the value "200G(50G x 4 lanes)" indicates that the port has a total speed of 200G and uses four 50G lanes.

By default, a port transmits and receives data at its highest speed. You can reduce the speed if necessary. You cannot change the speed of a breakout port.

To use a non-default speed for your switch fabric ports, select the appropriate speed. For example, for QSFP-DD ports, you can use 40G, 100G, or 200G instead of the default of 400G.

b) For **Pluggable**, enter the product ID (PID) of the optic for the port or select the pluggable from the table. You can enter search and filter criteria in the fields below to find a specific PID.

c) Select **Save**.

Step 6 To change the administrative state, select **Edit port admin state**.

a) For **Admin state**, select the desired state.

- **Up**: The port will receive and forward traffic. This is the default state.

  To set the port to be Up but not forwarding traffic, check **Prevent traffic from being forwarded**.

- **Down**: The port will not receive or forward traffic.

b) Select **Save**.

Configure a switch port at the fabric level

You can configure any port of any switch in a fabric by going to the **Physical topology** area of that fabric. This page makes it easy to configure the ports of different switches.

Follow these steps to configure a port at the fabric level.

Step 1 Choose **Fabrics**, then select the fabric that contains the switch for which you want to configure its ports.

a) If the fabric is not in the edit mode, select **Switch to edit mode** ( ).

Step 2 In the **Physical topology** area, select **Port configurations**.

The **Port configurations** table lists all ports of all switches in the fabric.

Step 3 Above the **Port configurations** table, you can use the various fields and drop-down lists to filter the table.

- **Search**: The table displays only the ports whose IDs contain the string that you enter.
- **From**: The table displays only the ports whose port number is equal to or greater than the integer that you enter.
- **To**: The table displays only the ports whose port number is equal to or less than the integer that you enter.
- **Port type**: The table displays only the ports with the type that you select.
- **Model**: The table displays only the ports with the model that you select.

Step 4 In the **Action** column, select edit ( ) for the port that you want to configure.

a) From the **Select port speed** drop-down list, select the speed of the port.

The value in parentheses indicates the speed of each lane and the number of lanes. For example, the value "200G(50G x 4 lanes)" indicates that the port has a total speed of 200G and uses four 50G lanes.

By default, a port transmits and receives data at its highest speed. You can reduce the speed if necessary. You cannot change the speed of a breakout port.

To use a non-default speed for your switch fabric ports, you must select the appropriate speed. For example, for QSFP-DD ports, you can use 40G, 100G, or 200G instead of the default of 400G.

b) For **Pluggable**, enter the product ID (PID) of the optic for the port.

Forwarding error correction (FEC) is automatically set based on the specified pluggable. You cannot configure FEC manually.

c) For **Port role**, select **Fabric** for a port that connects to another leaf or spine switch in the fabric.

For other port types, select the role according to the purpose of the port.

- **Fabric**: Provides connectivity between fabric switches and allows for automatic discovery by peer switches. Select this for a port that connects to another leaf or spine switch in the fabric.
- **Host**: Provides a Layer 2 connection to a server or other general network device.
- **Port channel**: The port is a member of a port channel, also known as a link aggregation group (LAG). When you configure a port to be a member of a port channel, its role is automatically changed from **Unused** to **Port channel**: The **Port channel** role cannot be changed in this menu unless the port is first removed from the port channel.
- **Routed**: Provides a Layer 3 connection to a router or other network device.

  When you configure a routed port on a switch, you can select to enable VLAN tagging, which requires you to configure 802.1Q VLAN sub-interfaces on a Layer 3 interface to forward IPv4 and IPv6 packets to another device using static or dynamic routing protocols. You can use Layer 3 interfaces for IP routing. If you leave VLAN tagging disabled, select a route table (VRF) and specify IPv4 or IPv6 addresses for the routes. Refer to the succeeding substeps for the additional routed port configuration steps.
- **Unused**: The port does not forward or receive traffic.

d) If you chose **Routed** for the port role and you want to use VLAN tagging, toggle on **Enable VLAN tagging** and follow these additional steps.

1. Select **Add a sub interface**.
2. In **VLAN tag**, enter the tag.

   Do not assign the same VLAN tag to multiple sub-interfaces on the same physical interface.
3. For **Select VRF**, select a VRF instance.
4. For the **IP address** fields, enter an IPv4 address and mask, an IPv6 address, or both.

e) If you chose **Routed** for the port role and you do not want to use VLAN tagging, leave **Enable VLAN tagging** as toggled off and perform these additional steps.

1. For **Select route table**, select the desired route table (VRF).
2. For the **IP addresses** fields, enter one or more IPv4 or IPv6 addresses and masks.

   If you enter multiple addresses, separate each one with a comma.

f) For **Admin state**, select **Up** to enable the port.

- **Up**: The port will receive and forward traffic. This is the default state.

  To set the port to be Up but not forwarding traffic, check **Prevent traffic from being forwarded**.
- **Down**: The port will not receive nor forward traffic.

g) Select **Save**.

Configure multiple switch ports at the fabric level

You can configure any switch port in a fabric by going to the **Physical topology** area of that fabric. This page makes it easy to configure the ports of multiple switches simultaneously.

Some property choices reveal additional properties that you cannot set with this procedure. To configure these additional properties, you must configure the ports individually. The descriptions in each step indicate which properties have this limitation.

Follow these steps to configure multiple ports at the fabric level.

Step 1 Choose **Fabrics**, then select the fabric that contains the switches for which you want to configure their ports.

a) If the fabric is not in the edit mode, select **Switch to edit mode** ( ).

Step 2 In **Physical topology**, select **Port configurations**.

The **Port configurations** table lists all ports of all switches in the fabric.

Step 3 Above the **Port configurations** table, you can use the various fields and drop-down lists to filter the table.

- **Search**: The table displays only the ports whose IDs contain the string that you enter.
- **From**: The table displays only the ports whose port number is equal to or greater than the integer that you enter.
- **To**: The table displays only the ports whose port number is equal to or less than the integer that you enter.
- **Port type**: The table displays only the ports with the type that you select.
- **Model**: The table displays only the ports with the model that you select.

Step 4 Select the ports that you want to configure by checking the box next to each port ID. The ports can belong to different switches.

Step 5 To change the role of the ports, select **Edit port roles**.

a) For **Port role**, select the desired role.

- **Fabric**: Provides connectivity between fabric switches and allows for automatic discovery by peer switches. Select this for a port that connects to another leaf or spine switch in the fabric.
- **Host**: Provides a Layer 2 connection to a server or other general network device. If you enabled global spanning tree, you cannot select the STP blocking modes using this procedure.
- **Port channel**: The port is a member of a port channel, also known as a link aggregation group (LAG). When you configure a port to be a member of a port channel, its role is automatically changed from **Unused** to **Port channel**. The **Port channel** role cannot be changed in this menu unless the port is first removed from the port channel.
- **Routed**: Provides a Layer 3 connection to a router or other network device. You cannot enable nor disable VLAN, add sub interfaces, select a route table, nor specify IP addresses using this procedure.
- **Unused**: The port does not forward nor receive traffic. If you intend to configure the port to be a port channel member, the role must initially be **Unused**.

b) Select **Save**.

Step 6 To change the port properties, select **Edit port properties**. The properties include the port speed and pluggable.

a) For **Configure port speed**, select the speed of the port.

The value in parentheses indicates the speed of each lane and the number of lanes. For example, the value "200G(50G x 4 lanes)" indicates that the port has a total speed of 200G and uses four 50G lanes.

By default, a port transmits and receives data at its highest speed. You can reduce the speed if necessary. You cannot change the speed of a breakout port.

To use a non-default speed for your switch fabric ports, you must select the appropriate speed. For example, for QSFP-DD ports, you can use 40G, 100G, or 200G instead of the default of 400G.

b) For **Pluggable**, enter the product ID (PID) of the optic for the port or select the pluggable from the table. You can enter search and filter criteria in the fields below to find a specific PID.

c) Select **Save**.

Step 7 If you want to change the admininistrative state, select **Edit port admin state**.

a) For **Admin state**, select the desired state.

- **Up**: The port will receive and forward traffic. This is the default state.

  To set the port to be Up but not forwarding traffic, check **Prevent traffic from being forwarded**.

- **Down**: The port will not receive nor forward traffic.

b) Select **Save**.

Change the Craft console password of a switch

Follow these steps to change the Craft console password of a switch.

Step 1 Choose **Fabrics**, then click the fabric that contains the switch.

Step 2 In the **Topology** area, click the switch for which you want to change the Craft console password, then click the switch name.

Step 3 In the **Configure** area, click **Management port**, then click **Edit**.

Step 4 In the **Forgot your console password?** area, enter the new password in both fields, then click **Update**.

What to do next

Push the configuration change. For more information, see Finish and commit your changes, on page 17.

Cabling

Cisco Nexus Hyperfabric automatically determines how you should cable the devices in a fabric based on the number of connections per device pair and cabling strategy. The cabling strategy comprises two notions.

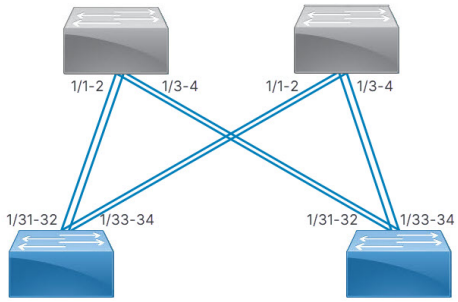
The first notion comprises these things:

- Strict: Cisco Nexus Hyperfabric connects each device pair with the number of connections that you selected. If a device does not have enough available ports for any reason, Cisco Nexus Hyperfabric indicates that there is an error with the cabling.
- Best effort: Cisco Nexus Hyperfabric connects each device pair with the number of connections that you selected if possible. However, if a device does not have enough available ports for any reason, Cisco Nexus Hyperfabric connects the devices in the pair to other devices in the same switch group that have enough available ports.

The second notion comprises these things:

- Dense: Cisco Nexus Hyperfabric connects consecutive ports of a device to consecutive ports of a paired device. The number of consecutive ports is equal to the number of connections per device pair that you selected. Cisco Nexus Hyperfabric repeats this connecting of consecutive ports for each successive paired device. This is used for switch-to-switch connections.

Figure 9: Switch-to-switch dense cabling strategy with two spine switches and two leaf switches

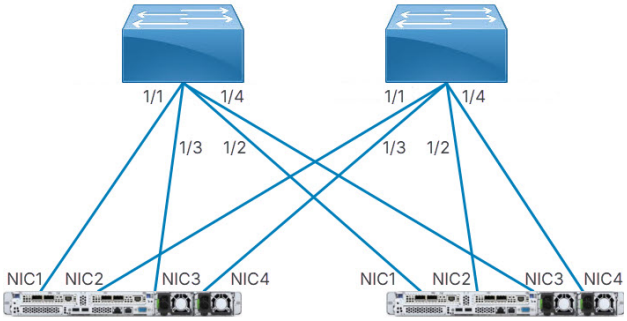


- Distributed: Cisco Nexus Hyperfabric connects the first network interface card (NIC) of the first server to the first port of the first switch, then the second NIC of the server to the first port of the second switch, and so on until the server has one connection to the first port of each switch. Cisco Nexus Hyperfabric repeats this process with the successive servers, but the NICs connect to the successive port of each switch. After each servers' NICs are connected to each switch, the entire process repeats until there are a number of connections from each server to each switch equal to the specified number of connections.

  Server port groups that have GPUs also use the distributed cabling strategy. In the rail group properties, these servers have width equal to the number of GPU servers and a count of 8.

  With or without GPUs, the same NIC number on each server connects to the same switch. Thus, for a fabric with two servers (server1 and server2) and two switches (switch1 and switch2), NIC1 of both servers connect to switch1 while NIC2 of both servers connect to switch2.

Figure 10: Switch-to-server distributed cabling strategy with two leaf switches and two servers



If your fabric includes connections between devices in the same switch group or server group, Cisco Nexus Hyperfabric allocates ports to device pairs in different groups first, then allocates ports to devices in the same group.

Server cabling

For uplink redundancy with servers that do not have GPUs, the cabling strategy is always distributed. The connections go from the ports of the switches to the ports of the network interface cards (NICs) of the servers.

Server port groups that have GPUs also use the distributed cabling strategy. In the rail group properties, these servers have width equal to the number of GPU servers and a count of 8.

In either case, the same NIC number on each server connects to the same switch. Thus, for a fabric with two servers (server1 and server2) and two switches (switch1 and switch2), NIC1 of both servers connect to switch1 while NIC2 of both servers connect to switch2.

Because switch-to-server connections are distributed, each subsequent NIC of a server connects to a different switch until all switches have one connection from different NICs of each server. After that, the next NIC of each server connects to the first switch and the process repeats.

Continuing the example, if you selected two connections, server1 has these connections:

- NIC1 to switch1
- NIC2 to switch2
- NIC3 to switch1
- NIC4 to switch2

Server2 has the same connections, but to different ports of the switches.

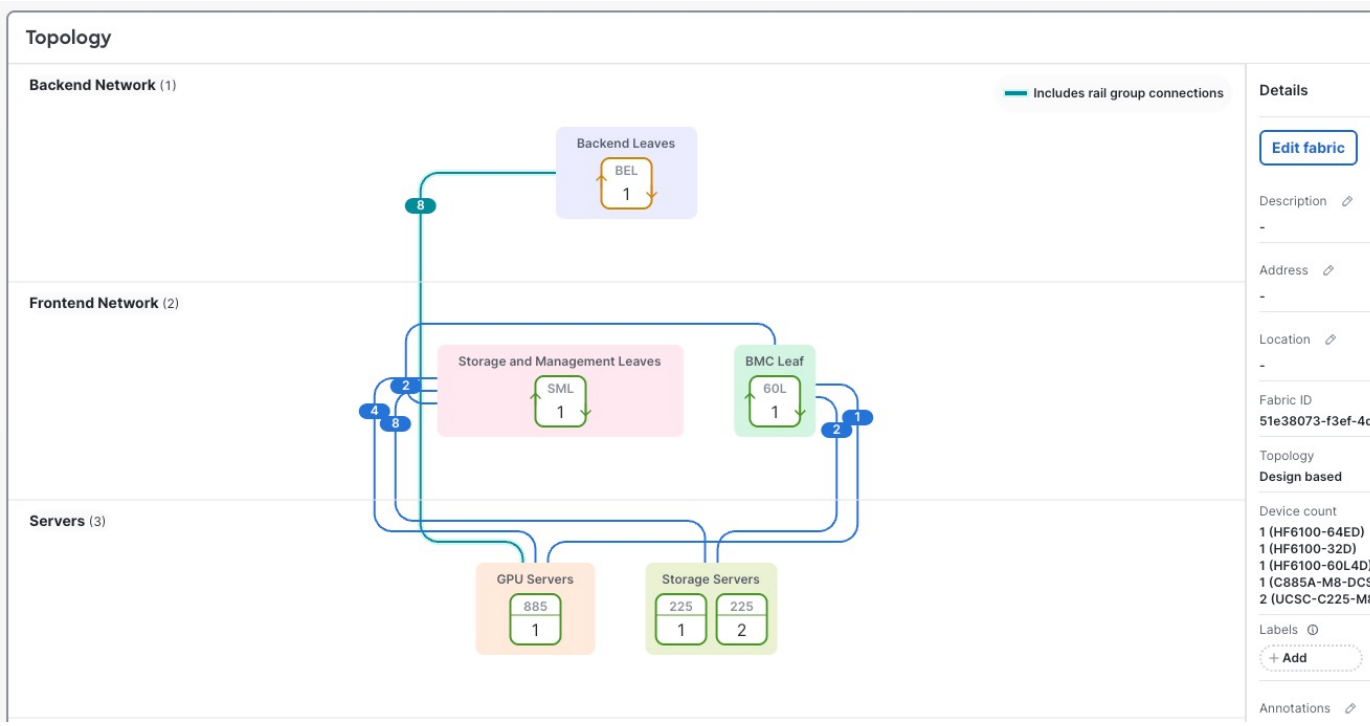
View cabling topology

You can view a graphical representation of the cables in a fabric and related device group connection information.

Step 1 Choose **Fabrics**, then select the fabric you want to view cabling for.

Cable connections appear as colored lines joining different device groups. Each number on a cable indicates the number of connections between device groups.

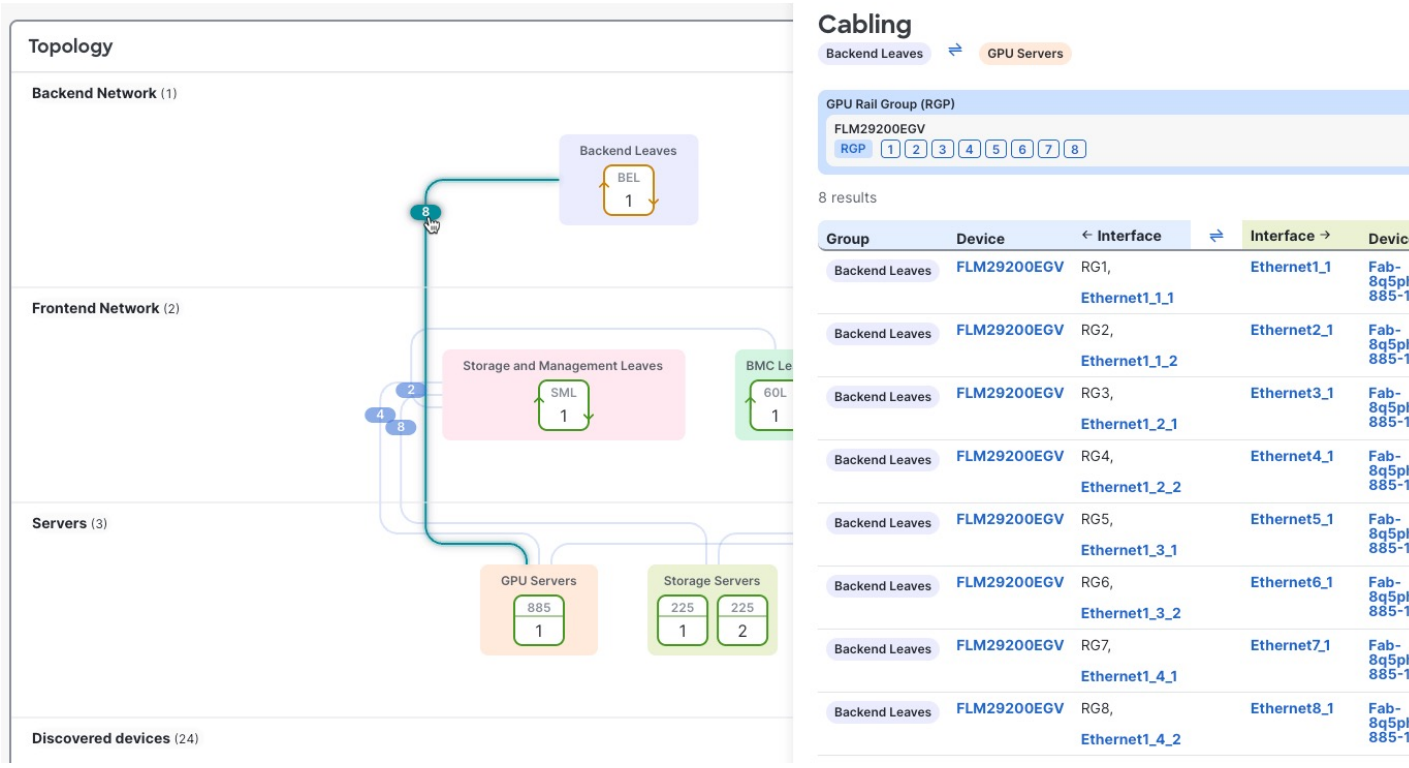
Figure 11: Cabling topology



Step 2 Select the number on a cable you are interested in.

The Cabling drawer lists all connections for the selected cable. Each connection shows the device group names, device names, and interfaces.

Figure 12: Device group connections



Modify cabling

Cabling is initially configured when creating a fabric. However, you may want to modify fabric connections to increase the number of links between switch pairs, add new switches, and so on.

You can view, export, and modify all fabric connections associated with a switch.

Step 1 Navigate to the switch you are interested in.

a) Choose **Fabrics**, then select the fabric that has the switch for which you want to modify cabling.

b) In the **Topology** area, select the switch you want to configure, then select the switch name.

Step 2 In the **Configure** area, select **Cabling**. The Cabling table lists all fabric connections associated with the switch. To export this list, select **Export CSV**.

Step 3 Select **Edit cabling**.

Step 4 To modify switch group connections:

a) For a selected switch group's connections, select edit (✎) and select the PID for the pluggables.

1. The first switch group radio button is selected by default. Select a PID from the table at the bottom for the first switch group. You can filter the table using the fields above the table.
2. Select **Cable** and select a PID.
3. Select the second switch group and select a PID.
4. Select **Select**.

b) For **Connections**, enter the number of connections per switch pair.

Note This is not the total number of connections between switch groups.

c) For **Cabling strategy**, choose how you want to cable the switches.

Step 5 To modify rail groups, you can add a new rail group or modify these fields:

- **Rail group name**: Enter a name.
- **Short name**: Enter a short name.
- **Count**: Currently, 8 is the only possible value.
- **Placement strategy**: Select the strategy. Currently, you can only select **Distributed** for servers.
- **Switch group**: Select the switch group to which this rail group will be connected.

Step 6 To modify server port group to switch group connections, you can edit these fields:

- **Pluggable PID**: Select edit (✎) and select the PID for the pluggables.
- **Cabling strategy**: Select a strategy.
- **Switch/rail group**: Select a switch group or rail group.

Step 7 Select **Save fabric blueprint** to save configuration changes.

Step 8 Select **Run cabling** to apply configuration changes.

Finish and commit your changes

Your changes are not applied to the fabric until you review, commit, and push them.

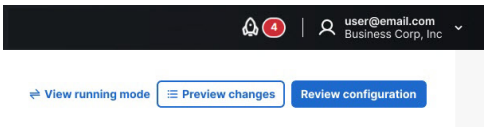
Note

For a more detailed description of this procedure, see "Modify a fabric" in *Cisco Nexus Hyperfabric—Getting Started*.

Follow these steps to finish and commit your changes.

Step 1 Select **Review configuration**.

Figure 13: Review configuration



Step 2 Verify your changes in the review list.

Step 3 Select **Comment and push**.

Step 4 In the **Comment before pushing configuration** dialog box, enter the reason for the change.

Step 5 Select **Push configuration**.