Revised: August 15, 2025

Cisco Nexus Hyperfabric — Configure Switches

Configuring the switches and ports

The Blueprint Designer in Cisco Nexus Hyperfabric provides the basic switch and port configuration for the fabric. During the fabric design phase or after fabric deployment, you can use Cisco Nexus Hyperfabric to perform general configuration tasks such as enabling host ports for your network devices, creating routed ports for external connections, or adding sub-interfaces.

Modify a switch configuration

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

Step 2 If the fabric is not in edit mode, click **Switch to edit mode**.

Step 3 In the **Topology** area, click the switch position that you want to configure, then click the switch name.

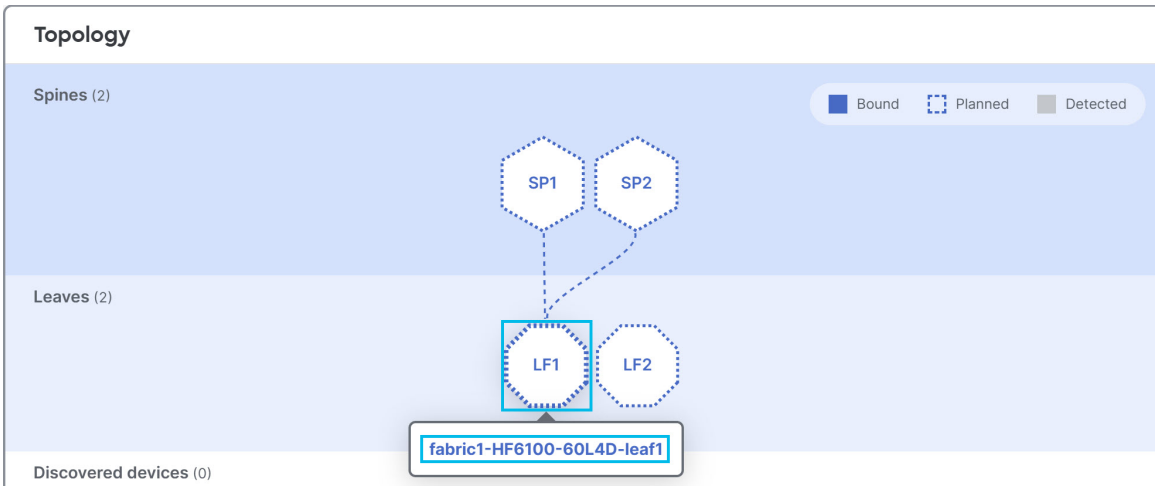

Step 4 In the **Configure** area, select the switch property you want to configure.

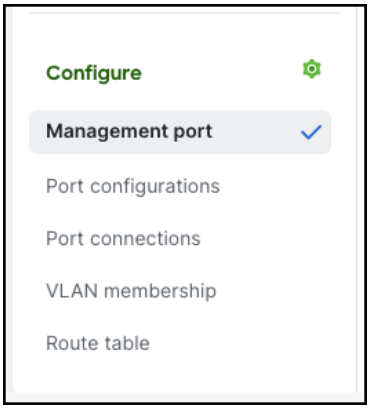

Step 5 Configure the selected switch property.

The configuration procedures are described in these links.

- Management port
- Port configurations
- Port connections
- VLAN membership
- Route table

Configure the management interface

You typically configure the management interface prior to onboarding, but you may configure or modify the management interface after onboarding.

Follow these steps to configure the management interface.

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

Step 2 If the fabric is not in the edit mode, click **Switch to edit mode**.

Step 3 In the **Topology** area, click the switch position that you want to configure, then click the switch name.

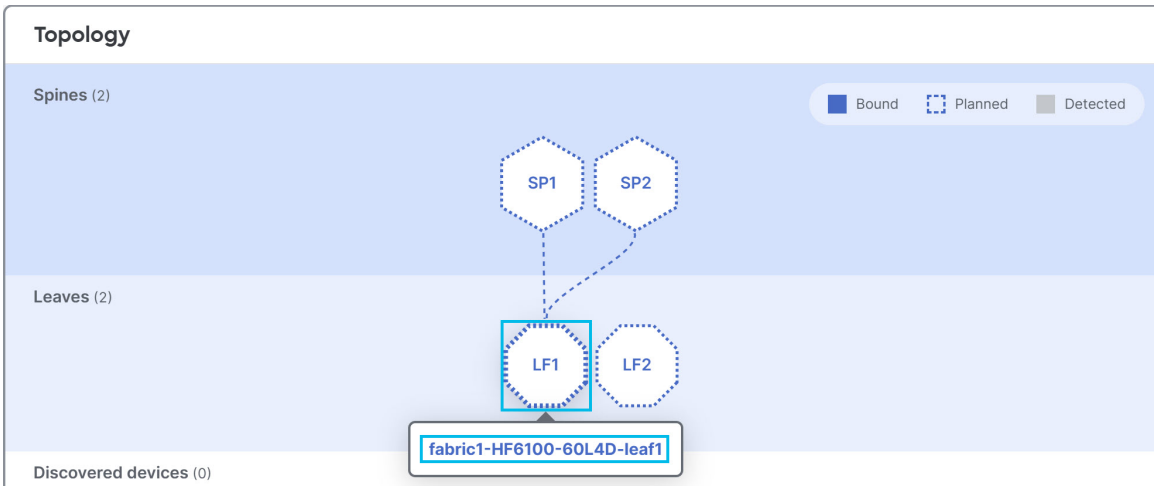

Step 4 In the **Configure** area, click **Management port**.

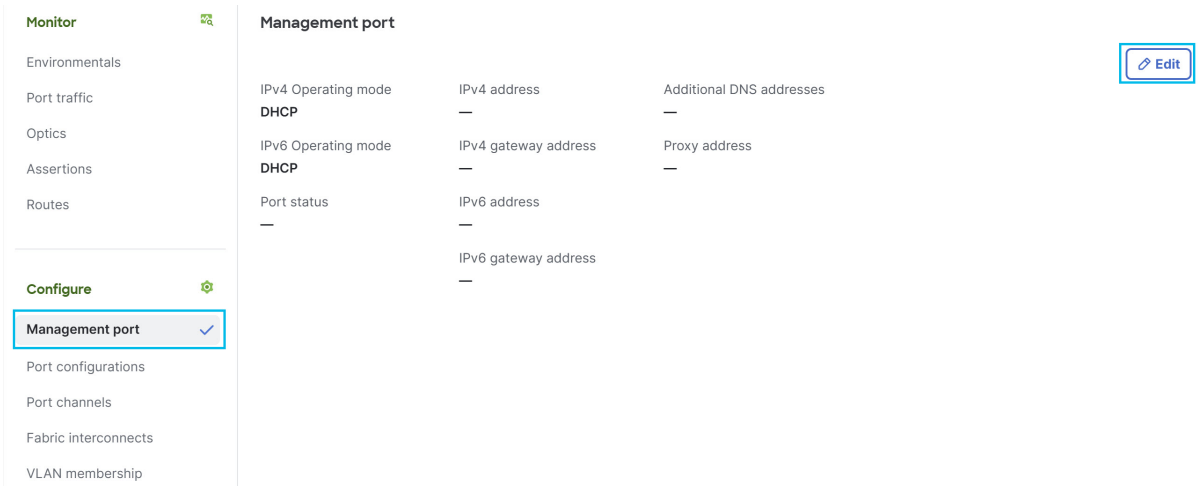

Step 5 Click **Edit** to open the **Edit management ports** dialog box.

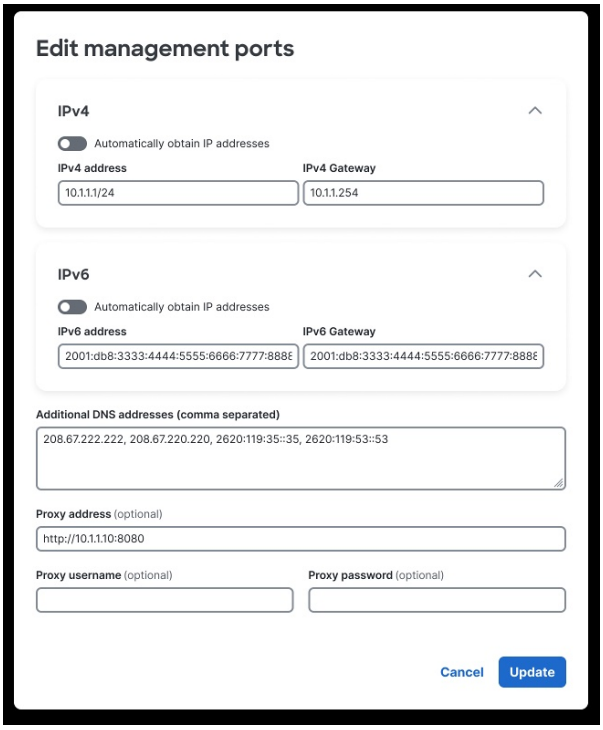

Step 6 In the **Edit management ports** dialog box, select how the management port obtains IP addresses.

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

Step 8 Click **Update**.

Configure fabric connections using auto-cabling

You can use the auto-cabling feature to configure the fabric connections quickly and easily. This feature simplifies the process and in most cases creates a desirable cabling plan.

Follow these steps to configure fabric connections using auto-cabling.

Note

Auto-cabling uses only QSFP-DD ports. If you want to use an SFP56 port (10G, 25G, or 50G) for a switch-to-switch connection, you must configure the connections manually. For the procedure, see Configure fabric connections manually, on page 4.

For your switch fabric ports, if you want to use a non-default speed such as 40G, 100G, or 200G instead of the default of 400G, you must select the appropriate speed when you configure the ports. For the procedure, see Configuring ports, on page 6.

Step 1 Select **Fabrics**, then click the fabric for which you want to use auto-cabling.

Step 2 If the fabric is not in the edit mode, click **Switch to edit mode**.

Step 3 In the **Physical topology** area, click **Fabric connects**.

The **Fabric connects** table lists all configured connections, if any.

Step 4 Click **Auto Cabling**.

The **Auto Cabling** dialog box opens.

a) For **Order**, enter the number of links that you want to have between each switch pair.

In most cases, you can use the default of 1.

b) For **Pluggable**, enter the product ID of the optic. Optionally, select one of the optics from the table at the bottom of the dialog box to fill in the field based on what you chose. You can use the **Search** field and **Plug type**, **Speed**, and **Cable type** drop-down lists to filter the table.

c) Click **Save** at the bottom of the dialog box.

The **Auto Cabling** dialog box closes and the **Fabric connects** table now shows how you should connect the switches.

Configure fabric connections manually

You can use the auto-cabling feature to configure the fabric connections quickly and easily. However, you can also configure the fabric connections manually for precise control over the device connections.

Follow these steps to configure fabric connections manually.

Note

For your switch fabric ports, if you want to use a non-default speed, you must select the appropriate speed when you configure the ports. For example, for QSFP-DD ports, you can use 40G, 100G, or 200G instead of the default of 400G. For the procedure, see Configuring ports, on page 6.

Step 1 Select **Fabrics**, then click the fabric that contains the switches that you want to connect.

Step 2 If the fabric is not in the edit mode, click **Switch to edit mode**.

Step 3 In the **Topology** area, click the switch position that you want to configure, then click the switch name.

Figure 1: Topology area

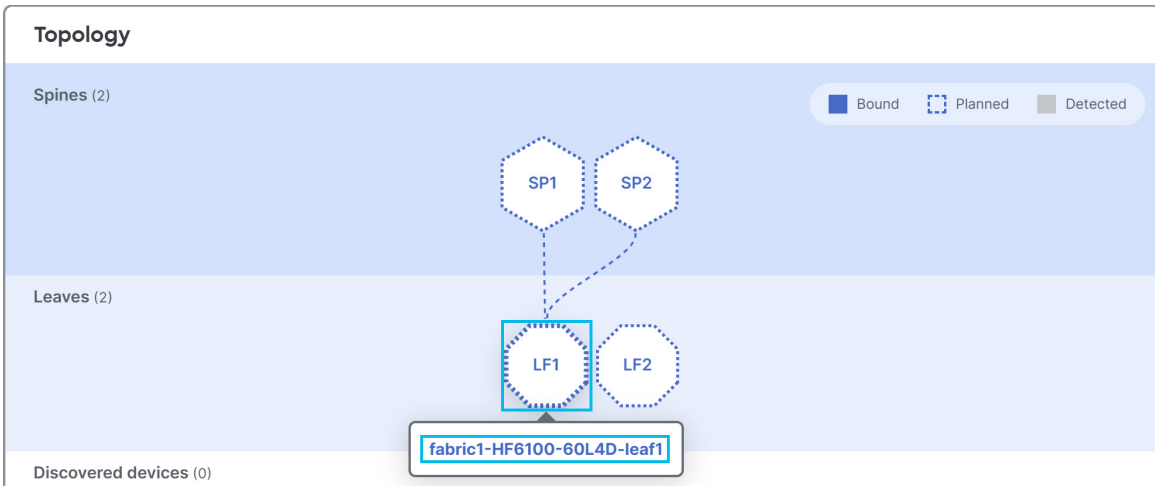

Step 4 In the **Configure** area, click **Fabric connects**.

The **Fabric connects** table lists all configured connections.

Step 5 Click **+ Add port connection**.

A dialog box opens for configuring a new connection.

a) For **Port interface**, select an interface.
b) For **Pluggable**, enter the product ID of the optic. Optionally, select one of the optics from the table at the bottom of the dialog box to fill in the field based on what you chose. You can use the **Search** field and drop-down lists to filter the table.
c) For **To switch**, select the destination switch.
d) For **Port interface** under **To switch**, select the destination switch's port.
e) Click **Save** at the bottom of the dialog box.

Dynamic port breakout

Dynamic port breakout enables you to break down a high-speed, channelized port on a network element into multiple low-speed ports, each going to multiple network elements, while fully utilizing the high-speed port's bandwidth. Breakout is suitable for very short links and offer a cost effective way to connect within racks and across adjacent racks. As an 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 the leaf switches and spine switches.

For how to configure breakout, see Configuring ports, on page 6.

Configuring ports

Configure a port at the switch level

You can configure various properties of a switch's port, such as the role and speed, and if it is broken out. 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 Select **Fabrics**, then click the fabric that contains the switch.

Step 2 If the fabric is not in the edit mode, click **Switch to edit mode**.

Step 3 In the **Topology** area, click the switch position that you want to configure, then click the switch name.

Figure 2: Topology area

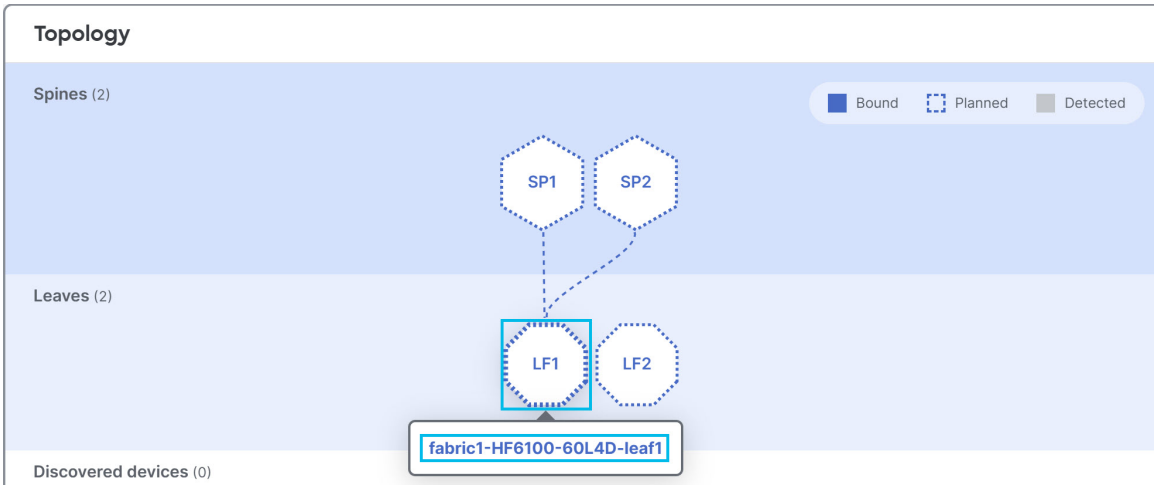

Step 4 In the **Configure** area, click **Port configurations**.

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

Figure 3: Port configurations

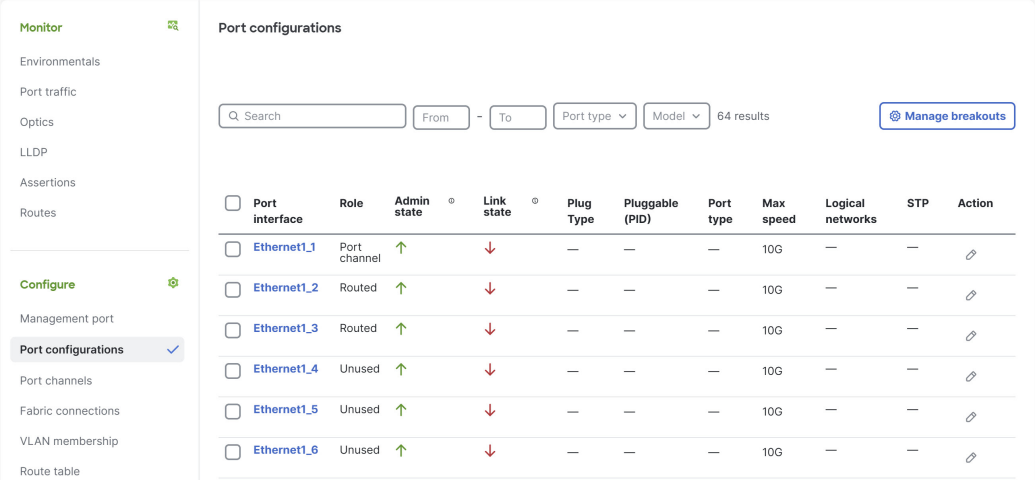

Step 5 (Optional) If you want to enable or disable breakout for any of the ports, click **Manage breakouts**.

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

The value in parenthesis indicates the number of lanes per port and the speed of each lane. For example, the value "2x200G (4x50G)" indicates 2 ports of 200G speed each and each port uses four 50G lanes for 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) (Optional) If you enabled breakout, select the cable type and optic type.

c) Click **Save**.

Step 6 In the **Port configurations** table, under the **Action** column, click the edit button (✎) for the port that you want to configure.

A dialog box opens for configuring the port.

Figure 4: Configuring a port

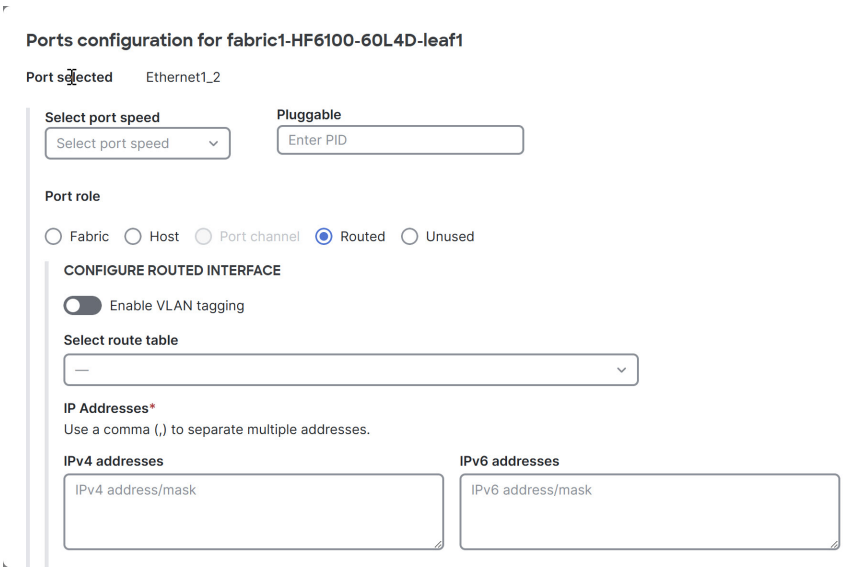

a) For **Select port speed**, select the speed of the port.

The value in parenthesis indicates the number of lanes and the speed of each lane. For example, the value "200G (4x50G)" 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 change the speed to something slower, if desired. You cannot change the speed of a broken out port.

For your switch fabric ports, if you want to use a non-default speed, 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. Alternatively, if you leave VLAN tagging disabled, you instead only select a route table (VRF) and specify IPv4 addresses or IPv6 addresses for the routes. See 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. Click **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) Click **Save**.

Step 7 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 8 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 properties choices expose additional properties that you cannot set using this procedure; you can set the additional properties only by configuring the ports individually. The descriptions in the steps specify which properties have this limitation.

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

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

Step 2 If the fabric is not in the edit mode, click **Switch to edit mode**.

Step 3 In the **Topology** area, click the switch position that you want to configure, then click the switch name.

Figure 5: Topology area

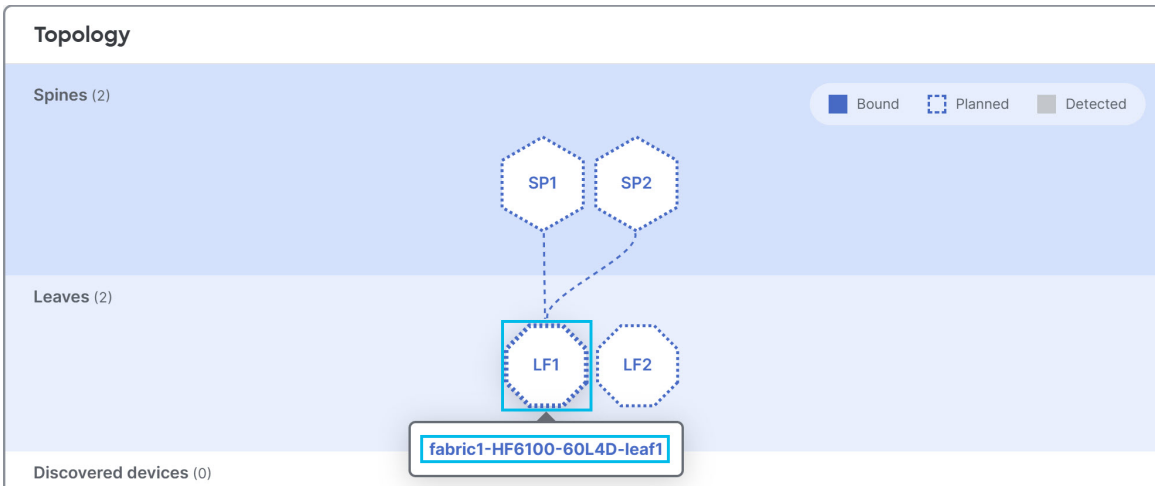

Step 4 In the **Configure** area, click **Port configurations**, then select the ports that you want to configure by putting a check in the box next to each port ID.

Figure 6: Multiple ports selected

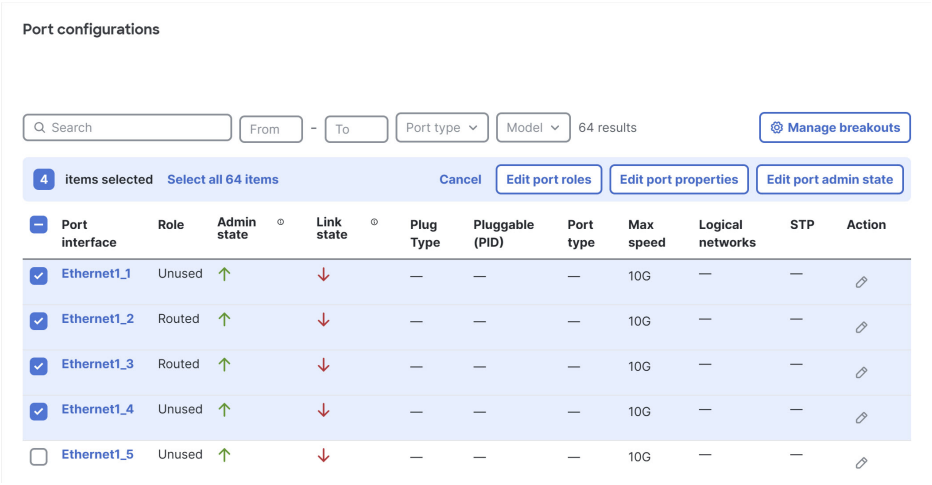

Step 5 If you want to change the role of the ports, click **Edit port roles**.

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) Click **Save**.

Step 6 If you want to change the port properties, click **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 parenthesis indicates the number of lanes and the speed of each lane. For example, the value "200G (4x50G)" 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 change the speed to something slower, if desired. You cannot change the speed of a broken out port.

For your switch fabric ports, if you want to use a non-default speed, 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.

c) Click **Save**.

Step 7 If you want to change the admininistrative state, click **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) Click **Save**.

Configure a 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 Select **Fabrics**, then click the fabric that contains the switch for which you want to configure its ports.

Step 2 If the fabric is not in the edit mode, click **Switch to edit mode**.

Step 3 In the **Physical topology** area, click **Port configurations**.

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

Step 4 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 5 In the **Action** column, click the edit button (✎) for the port that you want to configure.

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

The value in parenthesis indicates the number of lanes and the speed of each lane. For example, the value "200G (4x50G)" 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 change the speed to something slower, if desired. You cannot change the speed of a broken out port.

For your switch fabric ports, if you want to use a non-default speed, 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. Alternatively, if you leave VLAN tagging disabled, you instead only select a route table (VRF) and specify IPv4 addresses or IPv6 addresses for the routes. See 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. Click **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) Click **Save**.

Configure multiple ports 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 multiple switches simultaneously.

Some properties choices expose additional properties that you cannot set using this procedure; you can set the additional properties only by configuring the ports individually. The descriptions in the steps specify which properties have this limitation.

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

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

Step 2 If the fabric is not in the edit mode, click **Switch to edit mode**.

Step 3 In the **Physical topology** area, click **Port configurations**.

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

Step 4 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 5 Select the ports that you want to configure by putting a check in the box next to each port ID. The ports can belong to different switches.

Step 6 If you want to change the role of the ports, click **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) Click **Save**.

Step 7 If you want to change the port properties, click **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 parenthesis indicates the number of lanes and the speed of each lane. For example, the value "200G (4x50G)" 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 change the speed to something slower, if desired. You cannot change the speed of a broken out port.

For your switch fabric ports, if you want to use a non-default speed, 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.

c) Click **Save**.

Step 8 If you want to change the admininistrative state, click **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) Click **Save**.

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.

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. You can also see additional details about the switch.

Follow these steps to view the ports of a switch.

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

Step 2 If the fabric is not in the edit mode, click **Switch to edit mode**.

Step 3 In the **Topology** area, click the switch position whose ports you want to view, then click the switch name.

Figure 7: Topology area

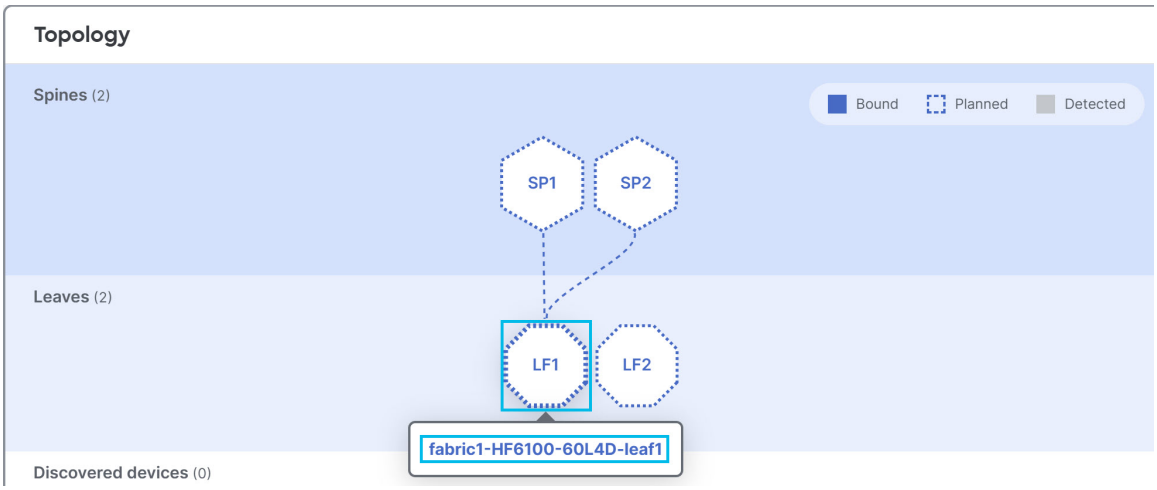

Step 4 In the **Ports** area, select **Port status** or **Face plate**.

Figure 8: Ports area

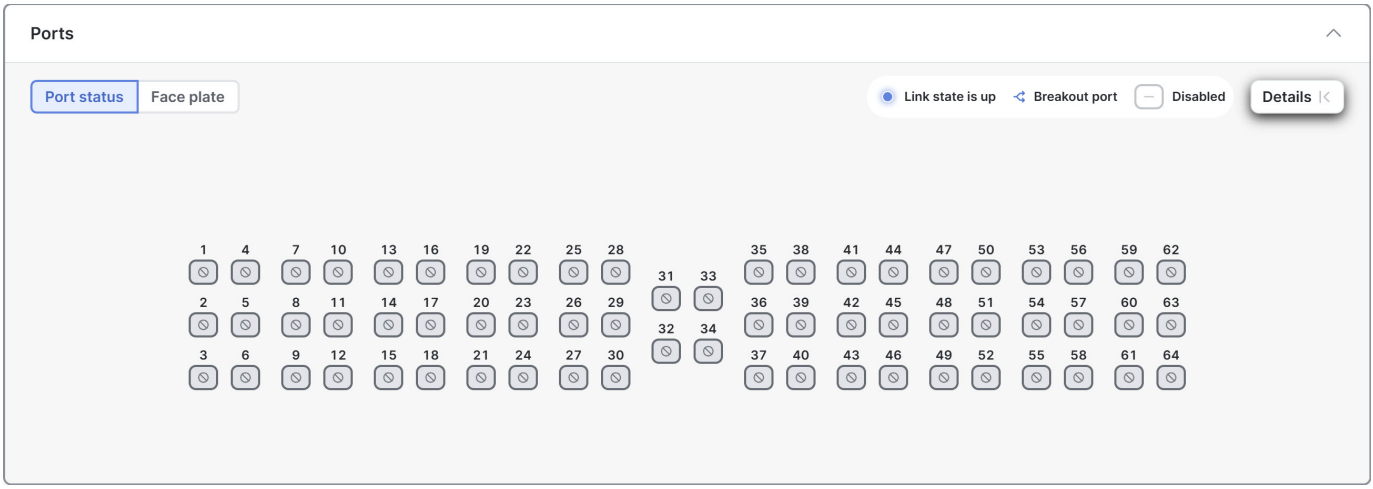

- **Port status**: This view shows all ports of the switch, as well as the status of each port. The icon next to the port number indicates whether the port is up or configured for breakout. The port image itself indicates whether the port is enabled or disabled. To understand the icons, see the legend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

  If a port is configured for breakout, the port shows the broken out ports as well as the color status for each broken out port. You can click the broken out port to see an enlarge view of the ports, which eanbles 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 5 In the **Port status** view, to see details about a specific port, 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.

View Inventory

In Inventory, you can view the all the devices that are claimed and bound to a fabric for an organization and also view the all the claimed devices that are not bound to a fabric.

Step 1 Choose **Inventory** > **Device Inventory**.

In the Device Inventory page you can view the number of claimed devices that are not bound to the fabric, the number of claimed devices that are bound to the fabric, and number of nodes as defined in the fabric blueprint.

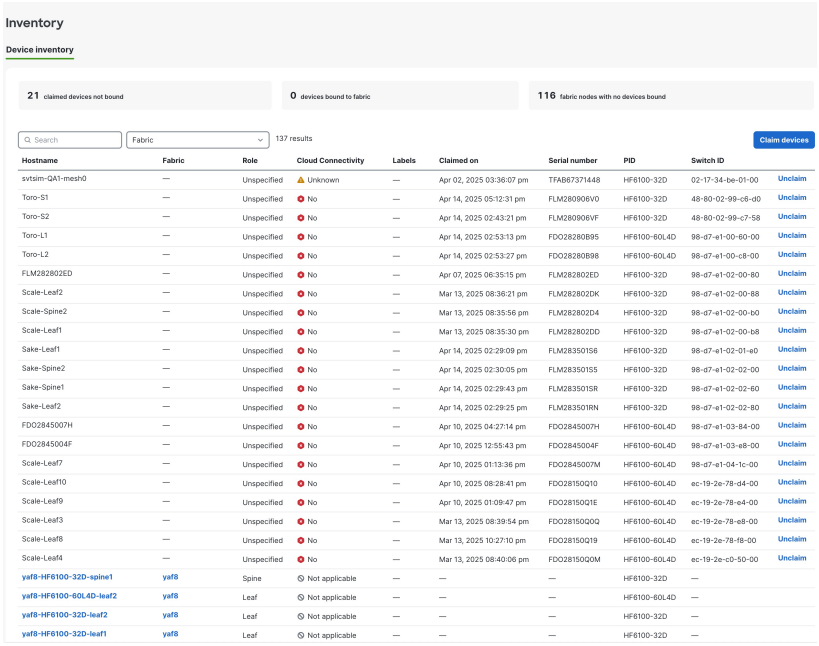

Step 2 Click **Claimed device not bound** to view the list of the claimed devices not bound to a fabric.

For each device you can view information such as hostname, fabric name, role, cloud connectivity status, claim date, serial number, PID, and switch ID.

a) Select a device and click **Unclaim** to unclaim a device.

If a device is bound to a fabric, you have to unbind the device first to enable **Unclaim** .

Step 3 Click **Device bound to fabric** to view the list of the claimed devices bound to a fabric.

For each device you can view information such as hostname, fabric name, role, cloud connectivity status, claim date, serial number, PID, and switch ID.

a) Select a device and click **Unclaim** to unclaim a device.

If a device is bound to a fabric, you have to unbind the device first to enable **Unclaim** .

Step 4 Click **Fabric nodes with no device bound** to view the list of the nodes such as spine or leaf switches as defined in the fabric blueprint.

For each node you can view information such as hostname, fabric name, role, and PID.

Step 5 Click **Claim devices** to claim one or more devices. See Claim a device.

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 "Workflow for making changes to the fabric" in *Cisco Nexus Hyperfabric—Getting Started*.

Follow these steps to finish and commit your changes.

Step 1 Click **Review configuration**

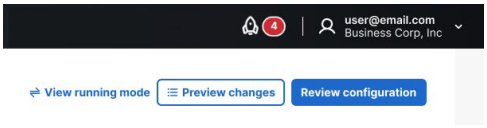

Step 2 Verify your changes in the review list.

Step 3 Click **Comment and push**.

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

Step 5 Click **Push configuration**.