Revised: April 23, 2026

Cisco Nexus Hyperfabric — Servers

Servers

Cisco UCS servers provide a combination of high-performance compute servers and scalable, high-speed storage servers that are required for Artificial Intelligence (AI) applications. These servers are monitored through Cisco Nexus Hyperfabric, which uses telemetry to collect real-time performance and health data from each device. Cisco Nexus Hyperfabric tracks server metrics, but does not allow any configuration changes.

View server information

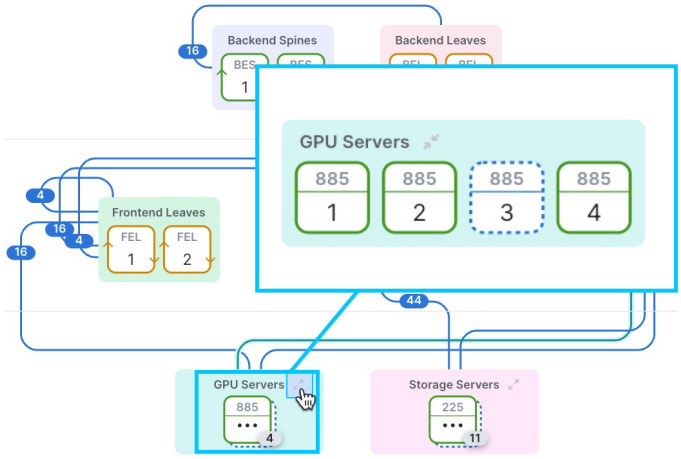
Viewing real-time server health, traffic, and state provides operational awareness. Monitoring diagnostic and performance metrics helps identify anomalies, detect issues early, and maintain network health.

Step 1 Navigate to the server you are interested in.

a) Choose **Fabrics**, then select the fabric that has the server you want to see property details for.

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

Figure 1: Expanded servers group



Step 2 In the **Monitor** area, select the server property you want to view.

- Information: Shows general and BMC server details.
- **Port information**: Lists all server ports and connected switch information, rail group membership, admin and link states, connection status, maximum port speed, MTU, and any VLAN tags.
- Assertions: Shows all issues on the server.
- Port statistics: Shows port statistics that can help with analyzing traffic patterns and congestion points..

- **Pluggables**: Shows expected pluggable details (PID, switch, and port) versus discovered details that are pulled live from telemetry. By comparing these details, you can validate actual design specifications or planned network diagrams. Any mismatches can highlight cabling mistakes, compliance issues, and so on.
- **Pluggable statistics**: Shows pluggable details, including lanes and diagnostic information which can help identify physical cabling issues.
- **LLDP neighbors**: Shows information about all LLDP neighbors that can help with verifying that the server is connected to the intended switch and port.
- **Server interface**: Lists all server interfaces and their IP address, MAC address, VLAN tag, and the protocol method used.
- **ARP/ND**: Lists all Address Resolution Protocol (ARP) mapping of IPv4 addresses to MAC addresses and Neighbor Discovery (ND) mapping of IPv6 addresses to MAC addresses. The list also includes interface, port, and status details of each mapping.
- **Routes**: Lists all IPv4 and IPv6 routes. The list includes route type, interface name, gateway, and gateway MAC address details.
- **Rail group**: For GPU servers, it shows all the rail groups the server is connected to.
- **Environmentals**: Shows the status of all server power supply units, fans, and temperature sensors.
- **VLAN membership**: Lists VLAN member switch connections and includes port, VLAN tag, VNI, and logical network details.

View BMC and general server details

Baseboard Management Controller (BMC) is a microcontroller embedded on a server's motherboard. BMC allows administrators to monitor and control the server even if it is powered down, crashed, or unresponsive.

Administrators can view BMC details to

- monitor and manage server Ethernet interfaces for connectivity and troubleshooting,
- handle server certificates to ensure secure authentication and communication, and
- track firmware inventory to update and maintain system security and stability.

Step 1 Navigate to the server you are interested in.

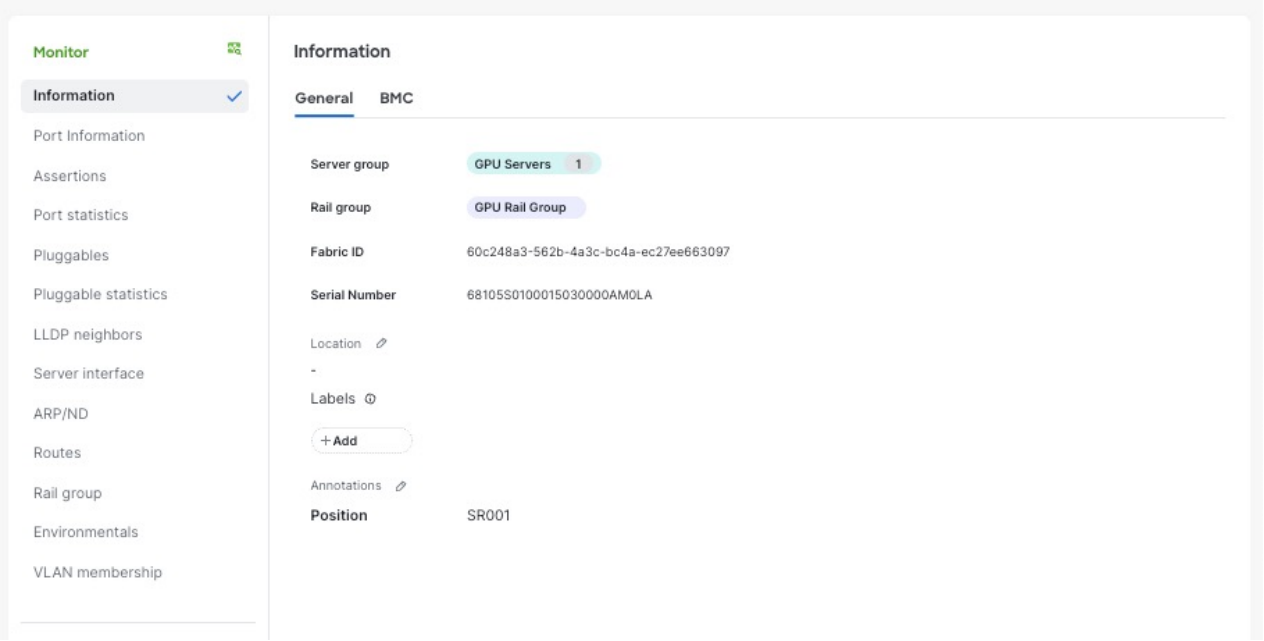
a) Choose **Fabrics**, then select the fabric that has the server.
b) In the **Topology** area, select the server you want to view, then select the server name. You may have to expand the server group in order to see the server.

Step 2 In the **Monitor** area, select **Information**.

Step 3 By default, the **General** server details are displayed. If available, server details include

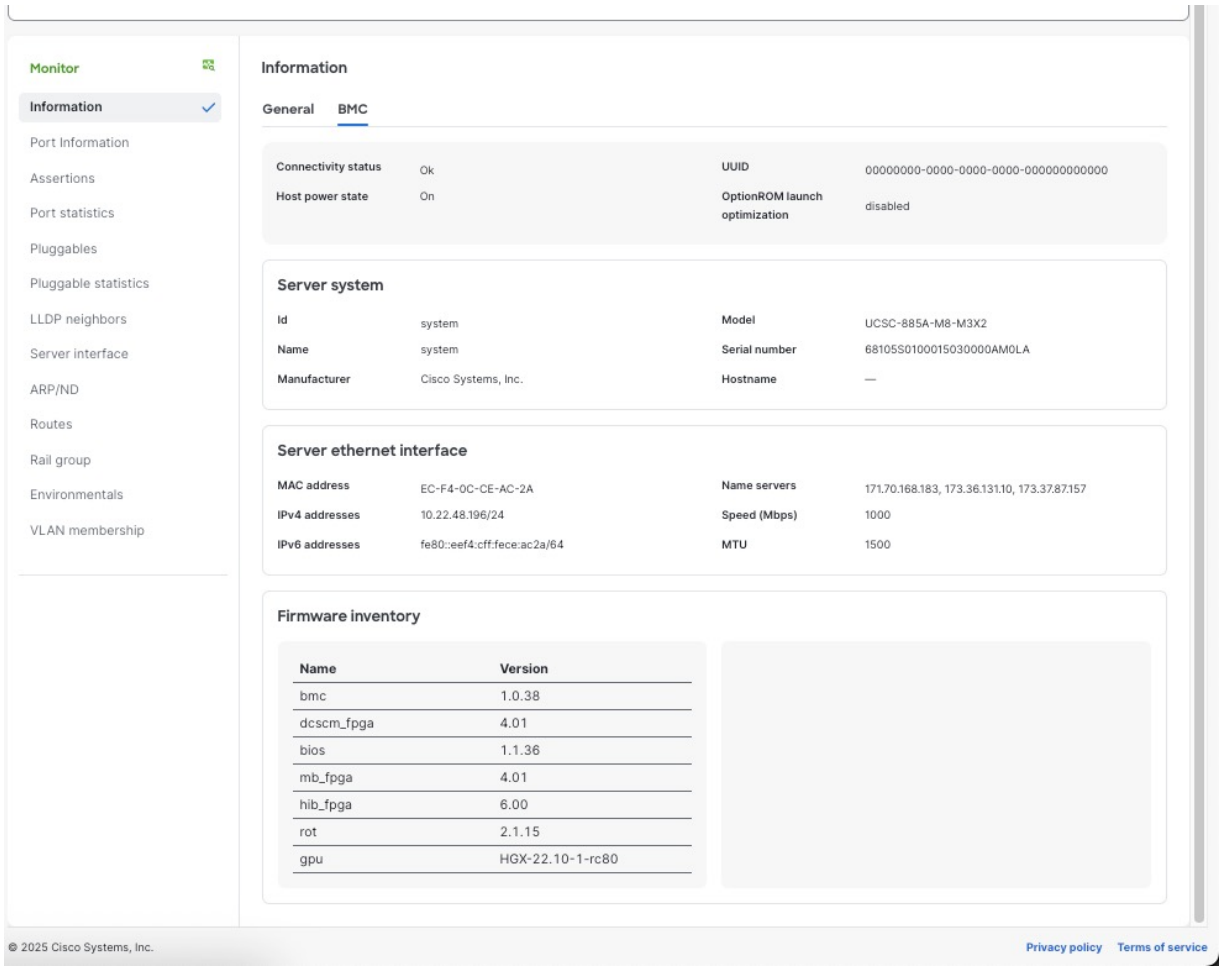
- binding status and option to bind or unbind the server,
- the server group, rail group, and fabric ID to which this server belongs,
- serial number,
- location, and

• labels and annotations that you can add.



Step 4 Click **BMC** to view

- connection state,
- host power state,
- UUID,
- OptionROM launch optimization state,
- server system information,
- server ethernet interface information,
- server certificates, and
- firmware inventory.



View the ports of a server

At a glance, you can see which server ports are up, not connected, or disabled. You can view a list of all ports in a server to quickly identify which switch port or rail group it connects to. This list can help isolate network problems and verify link status, speed, and VLAN configuration. You can also investigate details for specific server ports.

Step 1 Navigate to the server you are interested in.

a) Choose **Fabrics**, then select the fabric that has the server.

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

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

- **Port status**: This view shows all ports of the server, 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. For more information, see Assertions.

- **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 server, in the **Monitor** area, select **Port information**.

The **Port information** table lists all the ports on the server. It displays:

- switch and rail group connections,
- admin and link states,
- assertion status,
- max speed,
- MTU, and
- VLAN tag.

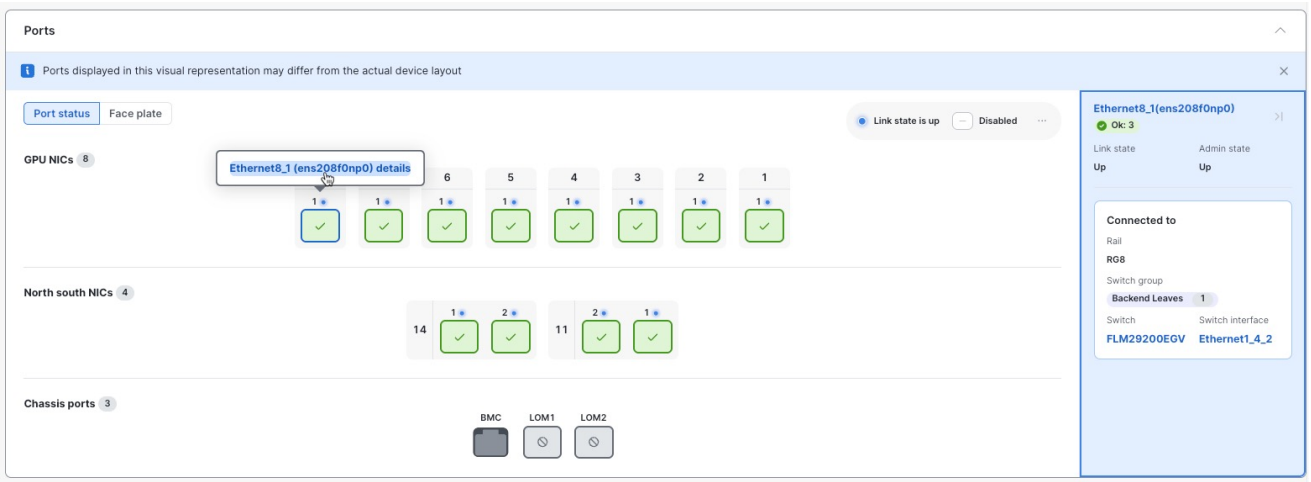
Step 4 To see details about a specific port:

a) In the **Port status** view, select that port.

The **Details** drawer opens, which contains more information about the port, including which rail group and switch it is connected to.

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

Figure 2: Server port details



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

View port statistics

Monitoring server port statistics helps you understand the operational health and performance of network ports. It also helps identify potential issues, including physical layer problems, congestion, traffic anomalies, and security concerns.

Step 1 Navigate to the server you are interested in.

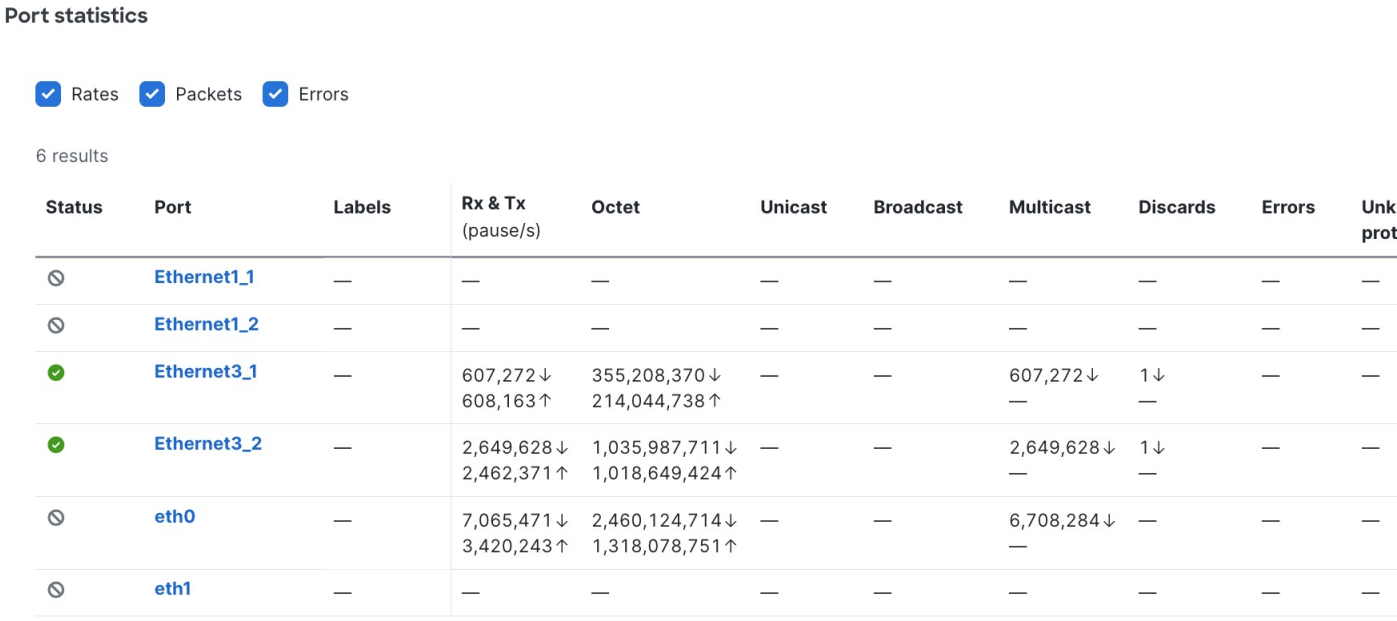
a) Choose **Fabrics**, then select the fabric that has the server you want to see port statistics for.

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

Step 2 In the **Monitor** area, select **Port statistics**.

Each metric in the **Port statistics** table offers insights.

- **Status**: Quickly identifies if there are issues.
- **Rx & Tx**, **Octet**: Shows the receive (Rx) and transmit (Tx) packet rates and the total octets (bytes) transferred. This information highlights current data throughput and volume, making it easier to assess network capacity and spot potential issues.
- **Unicast**, **Broadcast**, **Multicast**: Helps analyze traffic patterns to detect excessive broadcast or multicast activity, which could point to network loops or misconfigurations.
- **Discards**: Shows the number of packets that the port has dropped. This helps identify network congestion by showing packets dropped due to buffer overflows, directly impacting performance and requiring retransmissions.
- **Errors**: Shows the number of errors encountered on the port. This information helps diagnose physical layer problems (bad cables, duplex mismatches, and so on) that cause data corruption.
- **Unknown protos**: Shows the number of packets with unrecognized protocol type. This helps flag any suspicious or unauthorized traffic, making it easier to detect security anomalies and spot misconfigured devices.
- **FCS errors**: Shows the number of Frame Check Sequence (FCS) errors. FCS errors occur when a frame failed the integrity check due to corruption during transmission. FCS errors are typically caused by physical layer issues or timing mismatches in packet spacing.



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.

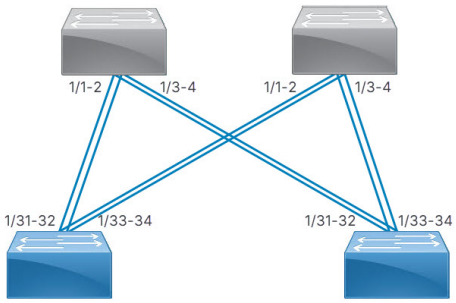
The first notion comprises these things:

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

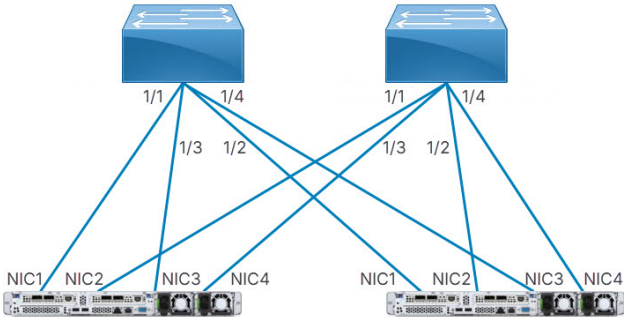
  

- Distributed: Cisco Nexus Hyperfabric connects the first network interface card (NIC) of the first server to the first port of the first switch, then the second NIC of the server to the first port of the second switch, and so on until the server has one connection to the first port of each switch. Cisco 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 4: 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.

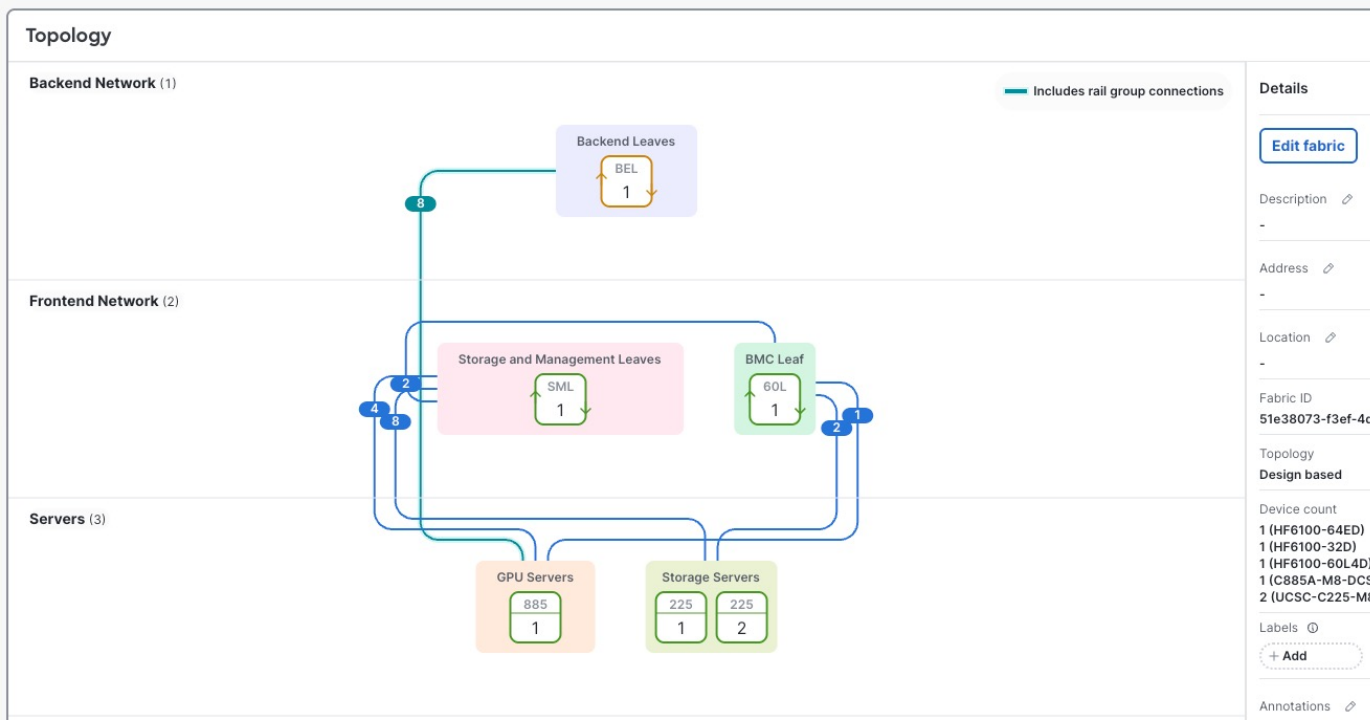
View cabling topology

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

Step 1 Choose **Fabrics**, then click 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.

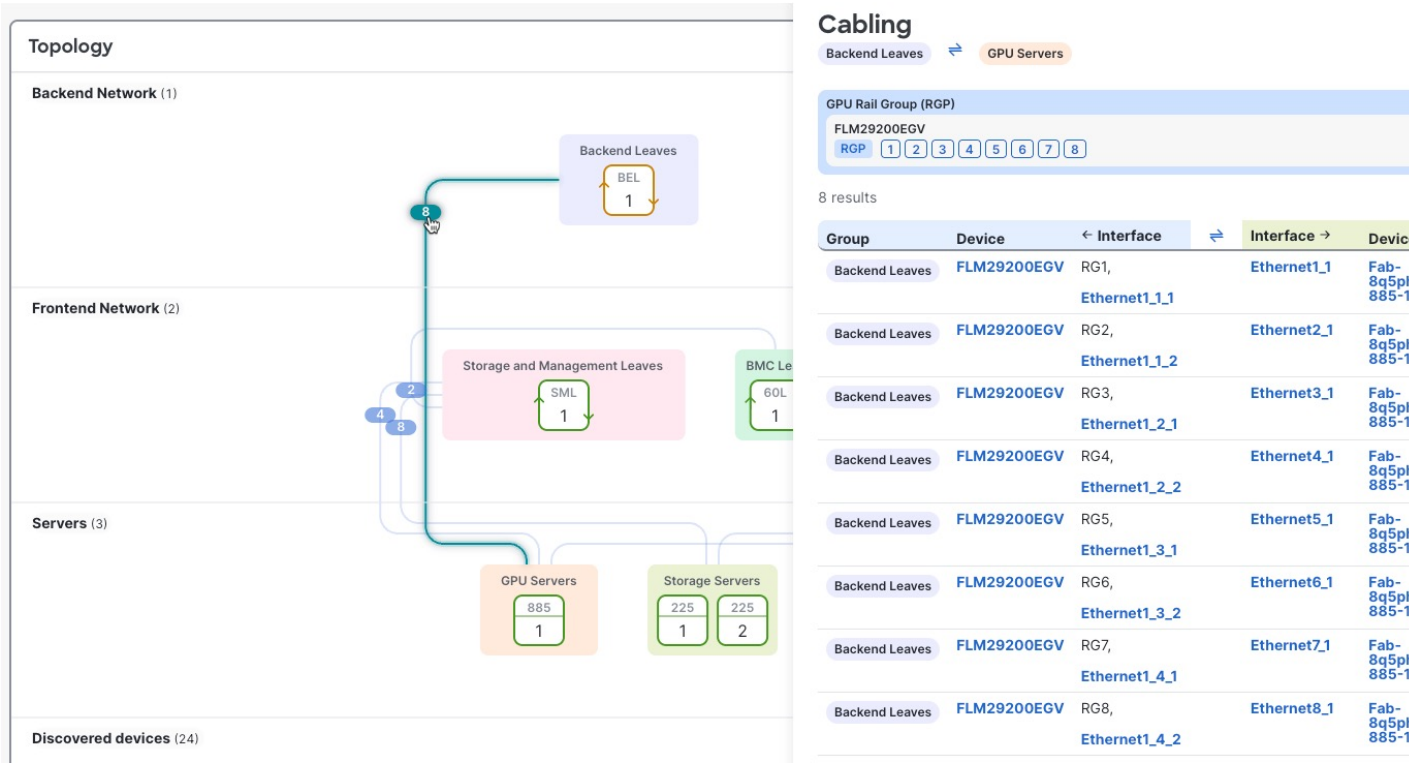
Figure 5: Cabling topology



Step 2 Click 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 6: 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 click 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 Click **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. Click 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.