Revised: October 31, 2026

Cisco Nexus Hyperfabric — Getting Started

Cisco Nexus Hyperfabric

The Cisco Nexus Hyperfabric is a fabric-as-a-service solution that allows you to design and build a physical network that is managed by a cloud-based service. Your network can start as small as a single switch and can grow to multiple locations, each with a fabric of many switches.

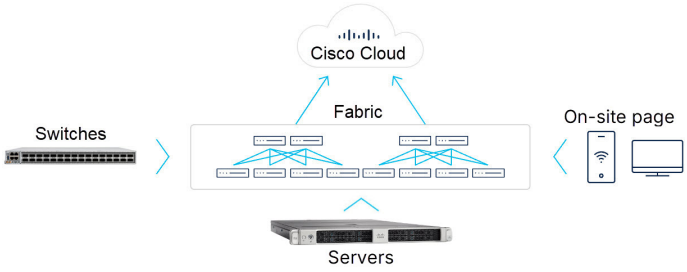
The design, deployment, and operation of the Cisco Nexus Hyperfabric is based on a shared responsibility model between you and Cisco. Cisco automation and operations staff own the proper functioning of the cloud-based service and the software upgrade process. Your staff maintains direct control of the timing of switch upgrades and of all interconnections to your applications, hosts, domestic network, and connectivity to the cloud-based service.

Components of a Cisco Nexus Hyperfabric deployment

These components comprise a Cisco Nexus Hyperfabric deployment:

- Cisco Nexus Hyperfabric service: a scalable, globally-distributed, multi-tenant cloud-based software-as-a-service (SaaS) network infrastructure management solution. The Cisco Nexus Hyperfabric service is hosted and operated by Cisco, and is the single point of provisioning and visibility for your fabrics. Organizations are created within a specific geographic region, which determines where device telemetry and services reside.
- Cisco Nexus Hyperfabric switches: cloud-managed Cisco 6000 switches on your premises or in a colocation environment. Bootstrapping and administration of the switches is handled by the cloud-based service.
- High-performance fabrics: mesh or spine-leaf topologies that use EVPN VXLAN, Layer 2 VLANs, or IPv4/IPv6 routing.
- Graphical user interface (GUI) and the REST API: the interfaces that you can use to perform tasks in Cisco Nexus Hyperfabric.
- On-site page: provides the network cabling technicians with step-by-step tasks to claim, bind, and cable the devices in your fabric. This page also provides real-time validation of the fabric topology.
- Cisco UCS servers: (Optional) cloud-monitored Cisco UCSC-C225-M8N-1P servers for storage on your premises or in a colocation environment. Bootstrapping and administration of the switches is handled by the cloud-based service.

Figure 1: Cisco Nexus Hyperfabric components



How Cisco Nexus Hyperfabric works

This is the overall process of how Cisco Nexus Hyperfabric works:

1. In collaboration with the Cisco Nexus Hyperfabric team, you design your fabric online using an automated design tool provided in the Cisco Nexus Hyperfabric GUI. The completed design is a blueprint for your new fabric.

2. When your fabric blueprint is complete, Cisco Nexus Hyperfabric generates these additional resources:
    - A detailed bill of materials (BoM) for purchasing the necessary fabric components, including servers, switches, and pluggables such as optics, Direct Attach Cables (DACs), and Active Optical Cables (AOCs)
    - A complete cabling plan for your installer to wire the fabric
3. When you receive your components, your installer physically wires them as directed by the **On-site** page.
4. Each new device contains an agent that allows it to contact Cisco Nexus Hyperfabric. You configure the management port of one switch and provide that port with an external Internet connection. That switch then contacts Cisco Nexus Hyperfabric. This Internet-connected switch is referred to as the "exit node".
5. The agent enables each device to discover and authenticate its neighbors. As each device discovers a path to an exit node in the fabric, the device contacts the service through the exit node.
6. When your devices are connected and ready, you register ("claim") each device online with Cisco, specifying the organization that owns the device.

Note

In some cases, your devices might be received already claimed by your organization.

7. In Cisco Nexus Hyperfabric, you assign ("bind") each physical device to its logical location ("node") in the fabric blueprint. When you bind a device to a node in the blueprint, the device can download its configuration from Cisco Nexus Hyperfabric.
8. After the fabric has been fully provisioned, you can connect your local hosts and servers to the fabric.

Device telemetry

Cisco Nexus Hyperfabric provides nearly real-time telemetry for these device stats:

- Switch or server: port status, assertions, and traffic stats
- Interface: port counters, link status, and network configuration
- Environmental: per switch sensor readings for fans, PSUs, and other components
- Optics and pluggables: digital optical monitoring (DOM) statistics

Nexus Hyperfabric use cases

Typical use cases for Cisco Nexus Hyperfabric include small to medium branch offices or remote colocation sites where you do not want on-premises management. The Cisco Nexus Hyperfabric is ideal for meeting needs such as

- deploying or extending new data center fabrics or easily building new data center fabrics
- deploying reliably and rapidly with basic network requirements and compute connectivity
- designing new data center networks easily and evolving them over time
- deploying small, spineless fabrics of one or two switches
- deploying new fabrics for AI clusters easily

- deploying seamlessly and extending distributed data centers to edge or co-locations, and
- starting small and expanding fabrics seamlessly without confined architectures.

Supported topologies

Cisco Nexus Hyperfabric supports these network topologies:

- Single switch or redundant switch pair. For more information, refer to Topology: single leaf switch or redundant switch pair, on page 3.
- Small mesh fabric. For more information, refer to Topology: small mesh fabric with two to six leaf switches, on page 4.
- Spine and leaf fabric, which can include a multi-tier fabric. For more information, refer to Topology: 2-tier or 3-tier fabric with one to four spine switches and one or more leaf switches, on page 6.
- Super spine fabric. For more information, refer to Topology: super spine fabric, on page 8.

A growing site could change the topology as needed by revisiting the design tool in Cisco Nexus Hyperfabric and creating an updated blueprint.

The illustrations in the topology subections show cable connections as colored lines joining different device groups. Each number on a cable indicates the number of connections between device groups.

Topology: single leaf switch or redundant switch pair

In a site with fewer than 50 hosts, a single leaf switch may be sufficient. For high availability, you can deploy a second switch in a redundant configuration.

Figure 2: Single leaf switch

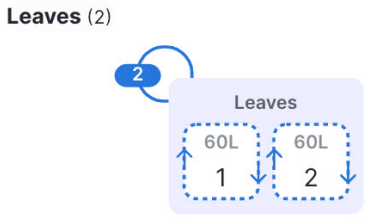


Overview of configuring a redundant switch pair

Figure 3: Diagram of a redundant switch pair



Figure 4: Redundant switch pair



This high-level list of steps provides an overview of configuring a redundant switch pair:

1. In the default switch tier, set the switch quantity to 2 in its group.
2. For **Switch groups** under **Group connections**, set the number of connections to 2.

Topology: small mesh fabric with two to six leaf switches

A small mesh fabric is intended for small-scale, high-performance environments such as edge data centers or artificial intelligence (AI) clusters. This fabric has two to six interconnected leaf switches without any spine switches, and in most cases you want to deploy the switches in pairs for equal redundancy across the fabric.

You can increase the scale depending on the number of uplink ports that you have. Beyond six leaf switches, you can add spine switches for a 2-tier fabric, or spine switches and another leaf switch tier for a 3-tier fabric. For more information about a multi-tier fabric, refer to Topology: 2-tier or 3-tier fabric with one to four spine switches and one or more leaf switches, on page 6.

Overview of configuring a small mesh fabric with two leaf switches

Figure 5: Diagram of a small mesh fabric with two leaf switches

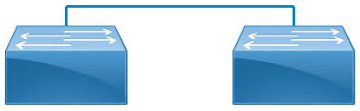
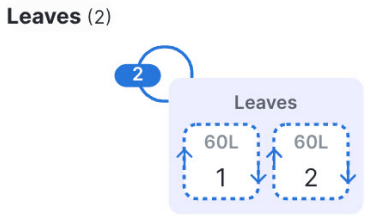


Figure 6: Small mesh fabric with two leaf switches



This high-level list of steps provides an overview of configuring a small mesh fabric with two leaf switches:

1. In the default switch tier, set the switch quantity to 2 in its group.
2. For **Switch groups** under **Group connections**, set the number of connections to 2.

Overview of configuring a small mesh fabric with four leaf switches

Figure 7: Diagram of a small mesh fabric with four leaf switches

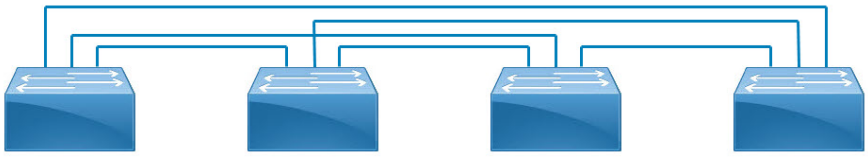


Figure 8: Small mesh fabric with four leaf switches



This high-level list of steps provides an overview of configuring a small mesh fabric with four leaf switches:

1. In the default switch tier, set the switch quantity to 4 in its group.
2. For **Switch groups** under **Group connections**, set the number of connections to 1.

Overview of configuring a small mesh fabric with six leaf switches

Figure 9: Diagram of a small mesh fabric with six leaf switches

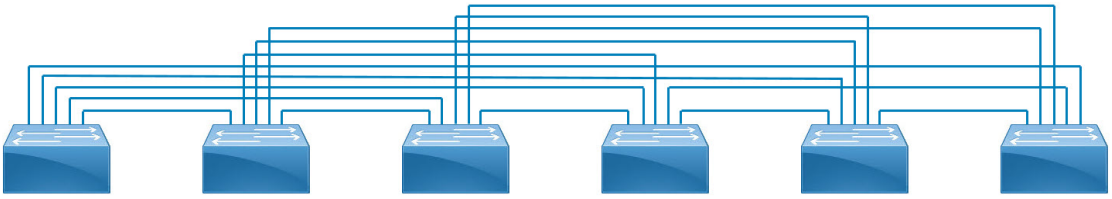
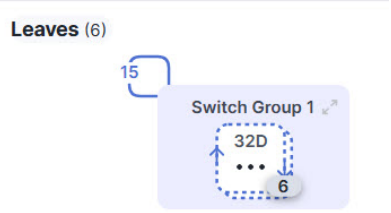


Figure 10: Small mesh fabric with six leaf switches



This high-level list of steps provides an overview of configuring a small mesh fabric with six leaf switches:

1. In the default switch tier, set the switch quantity to 6 in its group.
2. For **Switch groups** under **Group connections**, set the number of connections to 1.

Topology: 2-tier or 3-tier fabric with one to four spine switches and one or more leaf switches

A spine and leaf fabric can be a 2-tier topology that has one to four spine switches as the first tier and one or more leaf switches as the second tier. You can also have a 3-tier topology, in which you have additional leaf switches as the third tier connected to the second tier leaf switches.

Overview of configuring a 2-tier fabric with two spine switches and two leaf switches

Figure 11: Diagram of a 2-tier fabric with two spine switches and two leaf switches

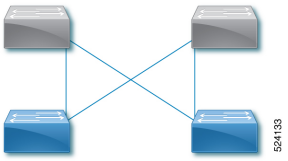
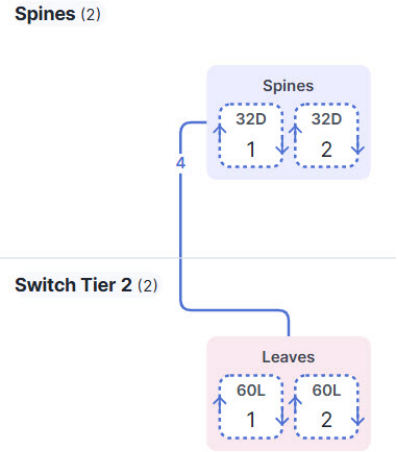


Figure 12: 2-tier fabric with two spine switches and two leaf switches



This high-level list of steps provides an overview of configuring a 2-tier fabric with two spine switches and two leaf switches:

1. For the spine switches, use the default switch tier and set the switch quantity to 2 in its group.
2. For the leaf switches, add a new switch tier and set the switch quantity to 2 in its group.
3. For **Switch groups** under **Group connections**, set the number of connections to 1 from the leaf switch group to the spine switch group.

Overview of configuring a 3-tier fabric with two spine switches in tier 1, four leaf switches in tier 2, and two leaf switches in tier 3

Figure 13: Diagram of a 3-tier fabric with two spine switches in tier 1, four leaf switches in tier 2, and two leaf switches in tier 3

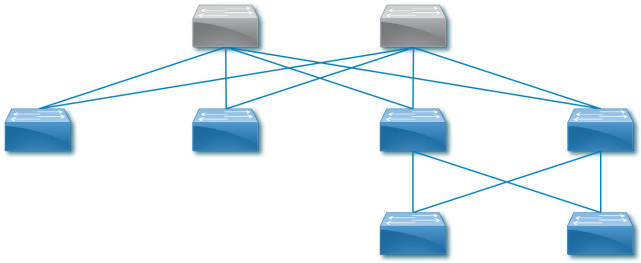
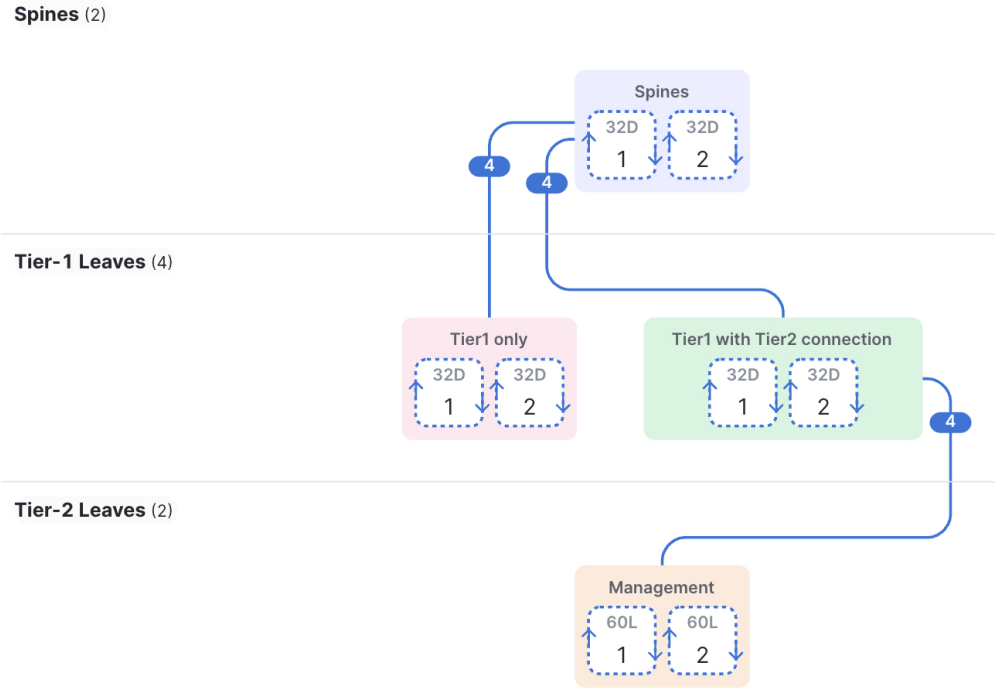


Figure 14: 3-tier fabric with two spine switches in tier 1, four leaf switches in tier 2, and two leaf switches in tier 3



This high-level list of steps provides an overview of configuring a 3-tier fabric with two spine switches in tier 1, four leaf switches in tier 2, and two leaf switches in tier 3:

1. For the tier 1 spine switches, use the default switch tier and set the switch quantity to 2 in its group.
2. For the tier 2 leaf switches, add a new switch tier and set the switch quantity to 2 in its group, then add another group and set the switch quantity to 2 in this group.
3. For the tier 3 leaf switches, add a new switch tier and set the switch quantity to 2 in its group.
4. For **Switch groups** under **Group connections**, set the number of connections as indicated:
    - 1 from each tier 2 leaf switch group to the spine switch group
    - 1 from the tier 3 leaf switch group to the second tier 2 leaf switch group

Topology: super spine fabric

A super spine switch topology is a three-tier architecture used to interconnect multiple spine-leaf PODs (points of delivery) within the same fabric. The super spine layer acts as an additional tier above the spine layer, providing extra interconnectivity and scalability.

Overview of configuring a super spine fabric

Figure 15: Diagram of a super spine fabric

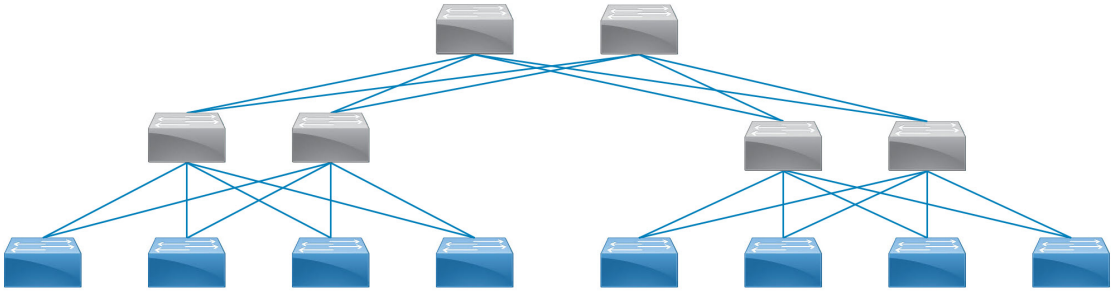
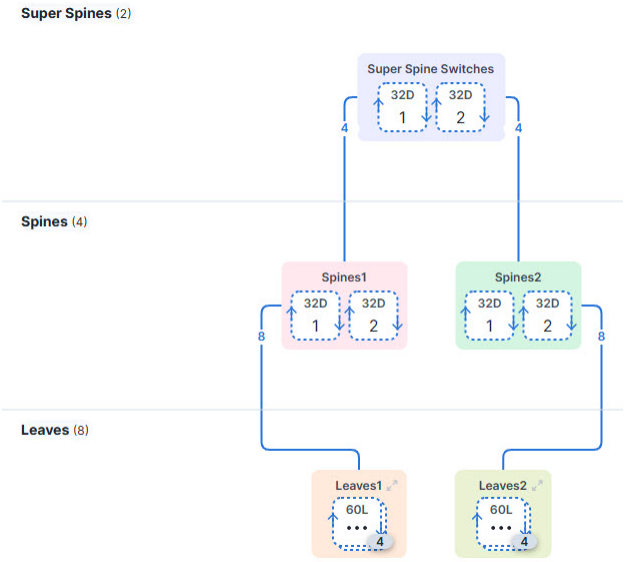


Figure 16: Super spine fabric with two super spine switches, four regular spine switches, and eight leaf switches



This high-level list of steps provides an overview of configuring a super spine fabric:

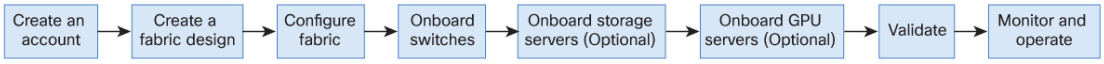
1. For the super spine switches, use the default switch tier and set the switch quantity to 2 in its group.
2. For the regular spine switches, add a new switch tier, add a second group in the tier, and set the switch quantity to 2 in both groups.
3. For the leaf switches, add a new switch tier, add a second group in the tier, and set the switch quantity to 4 in both groups.
4. For **Switch groups** under **Group connections**, set the number of connections as indicated:
    - 1 from each spine switch group to the super spine switch group
    - 1 from the first leaf switch group to the first spine switch group

- 1 from the second leaf switch group to the second spine switch group

The Cisco Nexus Hyperfabric lifecycle

To design, deploy, and maintain a network based on the Cisco Nexus Hyperfabric, you will work together with Cisco to follow the process illustrated in Figure 17: Cisco Nexus Hyperfabric lifecycle , on page 9.

Figure 17: Cisco Nexus Hyperfabric lifecycle



1. Log in to Cisco Nexus Hyperfabric.

   The first step in creating your cloud-managed network is logging in to the Cisco Nexus Hyperfabric portal and creating an organization in which you will design and operate your network fabric. Refer to Log into Cisco Nexus Hyperfabric, on page 9.

2. Create a fabric design.

   In Cisco Nexus Hyperfabric, you will design your fabric using the Blueprint Designer, selecting the devices and the topology to create a blueprint of your fabric. Refer to Designing the fabric blueprint, on page 26.

3. Configure (or modify) the fabric.

   In the Blueprint Designer, configure the interconnections between the fabric devices. Configure the pluggables such as optics, Direct Attach Cables (DACs), and Active Optical Cables (AOCs), and configure the remaining ports for connection to servers and other network devices. Refer to Modify a fabric, on page 29.

4. Onboard the switches.

   When the Nexus Hyperfabric switches have been physically installed and wired according to the blueprint, they must be "claimed" by the organization and "bound" to their logical positions (nodes) in the blueprint. The switches will become cloud-managed through the Cisco Nexus Hyperfabric portal. Refer to Onboard a fabric, on page 42.

5. (Optional) Onboard the servers.

   When the servers have been physically installed and wired according to the blueprint, they must be "claimed" by the organization and "bound" to their logical positions (nodes) in the blueprint. The servers will become cloud-monitored through the Cisco Nexus Hyperfabric portal. Refer to Onboard a fabric, on page 42.

6. Validate the design.

   After you bind the switches to the fabric blueprint, verify that traffic is being forwarded correctly.

7. Monitor and operate the fabric.

   Refer to *Cisco Nexus Hyperfabric — Configure Fabrics* .

Accessing Cisco Nexus Hyperfabric

Log into Cisco Nexus Hyperfabric

To get started in the design of your cloud-managed network, log into Cisco Nexus Hyperfabric.

You must have a cisco.com user account to log into Nexus Hyperfabric, which you can create at https://id.cisco.com/.

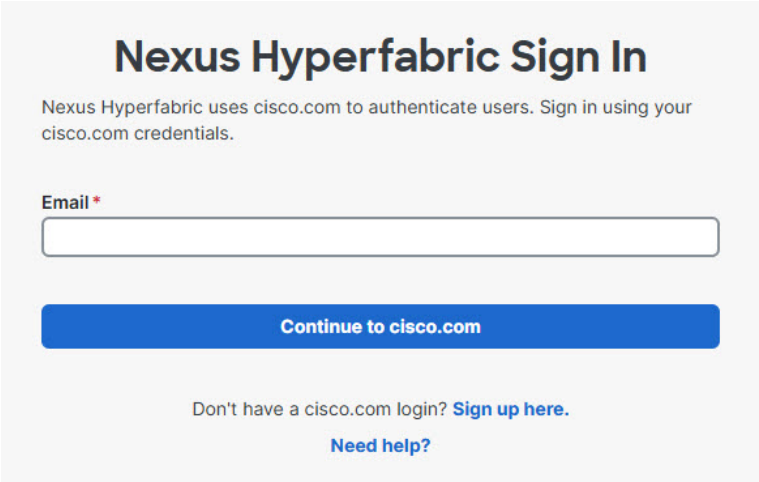
Follow these steps to log into Cisco Nexus Hyperfabric.

Step 1 Point your Web browser to https://hyperfabric.cisco.com/.

Step 2 For **Email**, enter the email address that you registered with https://id.cisco.com/, then select **Continue to cisco.com**.

You will be authenticated by Cisco's Single Sign-On (SSO) service. The Nexus Hyperfabric account is in addition to a cisco.com user account.

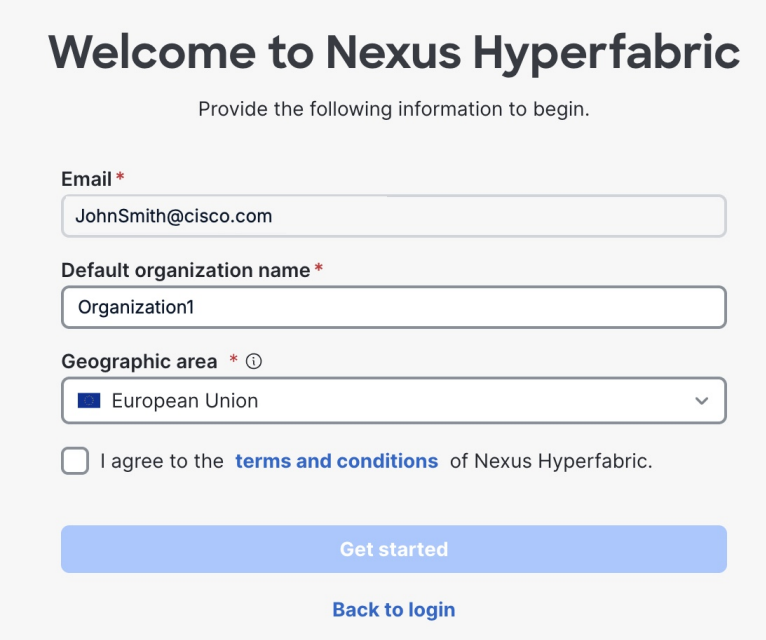
Figure 18: Cisco Nexus Hyperfabric sign in



If you already are an administrator in one or more Nexus Hyperfabric organizations, you will be taken the main page of the organization you logged into last.

a) If you do not have an email address that you registered with https://id.cisco.com/, select **Sign up here**, fill out the requested information, and repeat this procedure.

Step 3 If you are not the administrator for an existing organization, Nexus Hyperfabric prompts you to create one.



a) Enter the name of your organization.

   The default organization name is your email address plus the word "scratch." "scratch" only indicates that this is your first organization and does not otherwise have any special significance. You can change the name, and the name can contain a limited selection of special characters. For a list of allowed special characters, refer to Guidelines for creating an organization, on page 23.

b) Select the region in which the organization resides.

   You can select either European Union (EU) or United States (US). The region determines where the device telemetry and services for the organization are stored and processed. For more information, refer to Region-based organizations, on page 21.

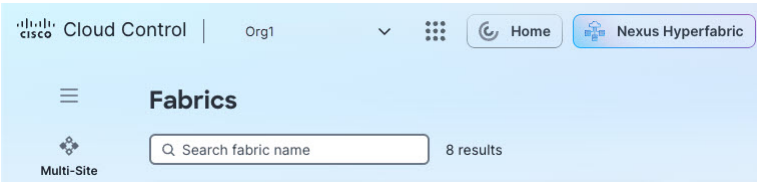
c) Check **I agree to the terms and conditions of Nexus Hyperfabric**

d) Select **Get started**.

   Nexus Hyperfabric displays the main page for the organization.

When you log into Nexus Hyperfabric, if you are a member of at least one organization, Nexus Hyperfabric selects one of your organizations automatically. If that organization is linked to Cisco Cloud Control, the page's header indicates this fact.

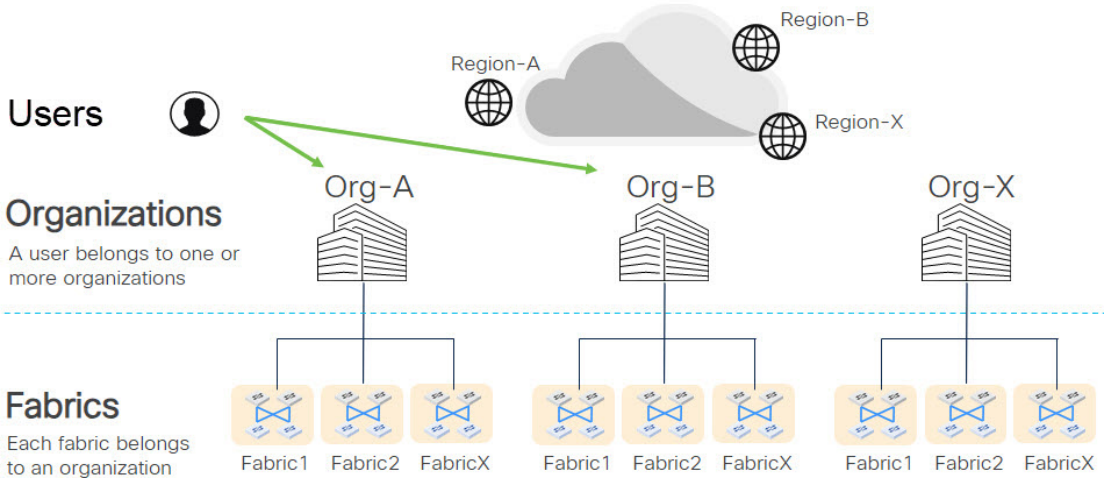
Figure 19: Cisco Cloud Control header



Add a new user to an organization

If you are an administrator for an organization, you can add a new user to that organization. A user can belong to multiple organizations. Figure 20: Users and organizations, on page 12 illustrates the association of users and organizations.

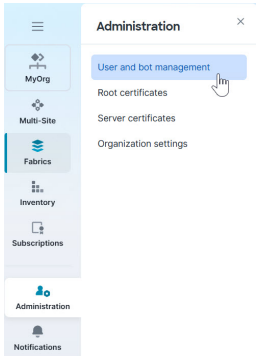
Figure 20: Users and organizations



Follow these steps to add a new user to an organization.

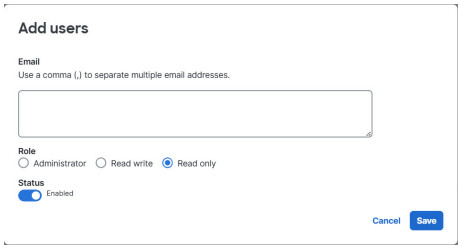
Step 1 Choose **Administration** > **User and bot management**.

Figure 21: User and bot management



Step 2 Select **+ Add user**.

Figure 22: Add users



Step 3 Enter the email addresses of the new users.

You can enter a comma-separated list of email addresses. The role and status settings you select will apply to all users in this list.

Step 4 Select a role: **Administrator**, **Read Write**, or **Read Only**.

Step 5 Select the status of the users: **Enabled** or **Disabled**.

Step 6 Select **Save**.

Identity provider federations

An identity provider (IdP) federation enables you to specify an IdP to use for single sign-on capability when logging into Cisco Nexus Hyperfabric. After you configure an IdP federation to use a specific domain, all users that are part of that domain can log into Nexus Hyperfabric using their regular (non-Cisco Connected Online) ID that they have with your company. These users have access to your company's Nexus Hyperfabric organization.

Limitations for identity provider federations

These limitations apply for identity provider (IdP) federations:

- A federated user cannot create an organization, regardless if the user is an administrator.
- A federated user cannot create an API bearer token. If you need a bearer token, contact a Cisco administrator within your organization to create one for you.

Configure an identity provider federation

Follow these steps to configure an identity provider federation.

Step 1 Add an authentication protocol.

Refer to the procedure for the authentication protocol that you want to add.

- Add SAML as an authentication protocol, on page 13
- Add OpenID Connect as an authentication protocol, on page 16

Step 2 Add a domain for an identity provider federation, on page 19.

Step 3 Add a routing rule for an identity provider federation, on page 20.

Repeat this step if you want to configure multiple routing rules.

Step 4 Log into your identity provider (IdP) and add the users whom you want to be able to log into Nexus Hyperfabric using this IdP.

Refer to the documentation for your IdP for how to add users.

Step 5 Verify the identity provider federation configuration, on page 20.

Add SAML as an authentication protocol

This procedure adds Security Assertion Markup Language (SAML) as an authentication protocol in Cisco Nexus Hyperfabric for use as an IdP federation. The procedure provides steps for using Microsoft Entra as the IdP as an example. You can use the IdP of your choice by substituting the Microsoft Entra-specific information as appropriate for your IdP. Refer to your IdP's documentation for the exact steps.

Note

The Microsoft Entra steps are subject to change without notice.

Follow these steps to add SAML as an authentication protocol.

Step 1 Collect the service provider details from Nexus Hyperfabric.

a) In Nexus Hyperfabric, choose **Administration** > **User management** > **External**.

b) Select **Launch admin portal**.

The **Cisco Identity** site opens in a new tab.

c) In Cisco Identity, go to **Settings** > **Manage IdPs** > **Service providers (SP)**.

The SAML service provider area contains this information that you must enter into the IdP. This list specifies the information that you need:

- **Entity ID**
- **Assertion consumer service URL**
- **Single sign-out URL** (only the first URL)

Step 2 In your IdP, create a new Enterprise application for Nexus Hyperfabric.

Follow these substeps for Microsoft Entra.

a) Point your browser to https://entra.microsoft.com.

b) Choose **Enterprise apps** > **New Application**.

c) Search for **Cisco Webex** in the application gallery and select the Cisco Webex app.

d) Enter a name for the application, fill out any other fields, and select **Create**.

Step 3 Set up single sign-on with SAML.

Follow these substeps for Microsoft Entra.

a) Choose **Enterprise apps** > **All applications** > ***app_name***.

app_name is the name of the application that you created.

b) Choose **Single sign-on** and select **SAML** as the single sign-on method.

c) Copy the values from the Cisco Identity fields listed in the table and paste the values into the equivalent fields for Microsoft Entra.

Table 1: Cisco Identity-to-Microsoft Entra mapping

Cisco Identity	Microsoft Entra
Entity ID	**Identifier (Entity ID)**
Assertion consumer service URL	**Reply URL (Assertion Consumer Service URL)**
Assertion consumer service URL	**Sign on URL**
Single sign-out URL	**Logout URL**

Step 4 Configure user attributes and claims.

Follow these substeps for Microsoft Entra.

a) From **Set up Single Sign-On with SAML**, select **Edit** in the **Attributes & Claims** area
b) Ensure that you configured the claims specified in the table.

Table 2: Required claims

Claim	Value
emailaddress	`user.mail`
uid	`user.userprincipalname`

c) Select **Add a group claim** and configure the claim as specified in the table.

Table 3: Group claim properties

Property	Value
Which groups associated with the user should be returned in the claim?	**Groups assigned to the application**
Source attribute	**Cloud-only group display names**

d) Save the claim name that your IdP generated.

You will use this claim name in step 8.h, on page 16 of this procedure.

Step 5 Set up user groups.

Follow these substeps for Microsoft Entra.

a) From the **Microsoft Entra admin center** home page, choose **Groups** > **New Group**.
b) Enter this information specified in the table.

Table 4: User group properties

Property	Value
Group type	**Security**
Group name	**ADMIN**
Group description	**Hyperfabric Admin Group**
Membership type	**Assigned**

c) Select **Create**.
d) If you must create additional user groups for different Nexus Hyperfabric roles, repeat this step once for each group.

Step 6 Assign the user groups to the Nexus Hyperfabric application.

Follow these substeps for Microsoft Entra.

a) Choose **Enterprise apps** > **All Applications** > ***app_name***.

app_name is the name of the application that you created.

b) Select **Users and groups**.
c) Select **Add user/group**.
d) Select the groups that you want to add, then select **Assign**.

At the least, you must add the ADMIN group that you created.

Step 7 Download the federation metadata XML certificate file from your IdP.

Follow these substeps for Microsoft Entra.

a) Choose **Enterprise apps** > **All Applications** > ***app_name***.

app_name is the name of the application that you created.

b) Select **Single sign-on**.
c) Under **SAML Certificates**, select the **Download** icon that is next to **Federation Metadata XML**.

Step 8 In Cisco Identity, go to **Settings** > **Manage IdPs** > **Identity providers (IdPs)**.

a) Select **+ Add an IdP**.
b) Select **SAML**, then select **Next**.
c) For **Select a method to connect your identity provider (IdP)**, select **Upload your IdP's metadata**.
d) For **Upload your IdP's metadata**, select **Not signed, self-signed or private CA signed IdP metadata file**.
e) Using **Click or drag file to this area to upload**, upload the federation metadata XML SAML certificate file that you previously downloaded.
f) Select **Next**.
g) Expand **Configure custom claims, JIT settings, and linking attributes (optional)** by selecting the down arrow next to this title.
h) Select **+ Add custom claim**, enter the claim that you saved in step 4.d, on page 15, and press Enter.

If you do not press Enter, the claim will not be added.

i) Select **+ Add custom claim**, enter `http://schemas.microsoft.com/ws/2008/06/identity/claims/groups`, and press Enter.

This custom claim ensures that the application roles you configured in step 4.d, on page 15 are included in the token for Nexus Hyperfabric role-based access control (RBAC) to work properly.

j) For the **Just-In-Time (JIT)** settings, check **Create or activate users** and **Update user with SAML attributes**.

These settings enable the federation to provision records for the users as they log in and the remote IdP authenticates them.

k) Ignore the remaining fields and select **Add IdP**.

Add OpenID Connect as an authentication protocol

This procedure adds OpenID Connect as an authentication protocol in Cisco Nexus Hyperfabric for use as an identity provider (IdP) federation. The procedure provides steps for using Microsoft Entra as the IdP as an example. You can use the IdP of your choice by substituting the Microsoft Entra-specific information as appropriate for your IdP. Refer to your IdP's documentation for the exact steps.

Note

The Microsoft Entra steps are subject to change without notice.

Before following this procedure, configure OpenID Connect for use with Nexus Hyperfabric. Refer to your IdP documentation for how to configure it.

Follow these steps to add OpenID Connect as an authentication protocol.

Step 1 Collect the service provider details from Nexus Hyperfabric.

a) In Nexus Hyperfabric, choose **Administration** > **User management** > **External**.
b) Select **Launch admin portal**.

The **Cisco Identity** site opens in a new tab.

c) In Cisco Identity, go to **Settings** > **Manage IdPs** > **Service providers (SP)**.

Collect the **Redirect URI** from this page.

Step 2 In your IdP, create a new Enterprise application for Nexus Hyperfabric.

Follow these substeps for Microsoft Entra.

a) Point your browser to https://entra.microsoft.com.
b) Choose **Enterprise apps** > **New Application**.
c) Select **Create your own application**.
d) Enter a name for the application, fill out any other fields, and select **Create**.

Step 3 Collect information about the application that you created.

Collect the client ID and discovery URL.

Follow these substeps for Microsoft Entra.

a) Choose **Enterprise registrations** > **All applications** > ***app_name***.

app_name is the name of the application that you created.

b) Collect the application (client) ID.
c) Select **Endpoints** and collect the **OpenID Connect metadata document** URL.

The OpenID Connect metadata document URL is also known as the Discovery URL.

Step 4 Configure the client secret.

Follow these substeps for Microsoft Entra.

a) From the application overview page, choose **Certificates & secrets**.
b) Select **New client secret**.
c) Enter a description and select an expiration date, then select **Add**.
d) Collect the generated client secret value.

You will use this value as the **Client Secret** during the IdP setup in the IdP federation. Ensure that you get the value from the **Value** column, not the **Secret ID** column. You must collect the secret now, as the value becomes hidden after you navigate away from this page.

Step 5 Configure the redirect URI.

Follow these substeps for Microsoft Entra.

a) From the application overview page, choose **Authentication**.
b) Select **Add Redirect URI**.

c) Select **Web** and enter the redirect URI that you collected from step 1.c, on page 17.
d) Select **Configure**.

Step 6 Configure the application roles.

Follow these substeps for Microsoft Entra.

a) From the application overview page, choose **App roles**.
b) Select **Create app role**.
c) Fill out the fields.
   - **Display name**: Enter `Hyperfabric Admin`.
   - **Allowed member types**: Select **Users/Groups**.
   - **Value**: Enter `ADMIN`.
   - **Description**: Enter `Hyperfabric Admin role with full access to Hyperfabric features.`
d) Select **Apply**.
e) Repeat this step if you want to create additional application roles.

Step 7 Assign users or groups to the application roles.

Follow these substeps for Microsoft Entra.

a) From the **Microsoft Entra admin center** home page, choose **Enterprise registrations** > **All applications** > ***app_name***.

   app_name is the name of the application that you created.

b) Choose **Users and groups**.
c) Select **Add user/group**.
d) Select the desired users or groups and assign them to the appropriate application roles, such as `Hyperfabric Admin`.

Step 8 Configure applications roles to be included in the token.

This step is not required for Microsoft Entra, as Microsoft Entra includes application roles in the `roles` claim by default. However, if you are using a different IdP, you may need to add a custom/optional claim to ensure the application roles or user group names are included in the token. Refer to your IdP's documentation for the exact steps.

a) Save the claim name that your IdP generated, which you use in step 12.j, on page 19.

Step 9 Configure API permissions.

Follow these substeps for Microsoft Entra.

a) From the **Microsoft Entra admin center** home page, choose **Enterprise registrations** > **All applications** > ***app_name***.

   app_name is the name of the application that you created.

b) Choose **API permissions**.
c) Select **Add a permission**.
d) Select **Microsoft Graph** > **Delegate permissions**.
e) Check the **email** and **openid** permissions, then select **Add permissions**.

Step 10 In Nexus Hyperfabric, choose **Administration** > **User management** > **External**.

Step 11 Select **Launch admin portal**.

The **Cisco Identity** site opens in a new tab.

Step 12 In Cisco Identity, go to **Settings** > **Manage IdPs** > **Identity providers (IdPs)**.

a) Select **+ Add an IdP**.
b) Select **OpenID Connect**, then select **Next**.
c) For **IdP name**, enter the name for your remote IdP.
d) For **Client ID**, enter the client ID that you collected from step 3.b, on page 17.
e) For **Client secret**, enter the client secret that you collected from step 4.d, on page 17.
f) For **Endpoints**, select **Use the discovery URL**.
g) For **Discovery URL**, enter the URL of the app.
h) Select **Next**.
i) Expand **Configure custom claims, JIT settings, and linking attributes (optional)** by selecting the down arrow next to this title.
j) Select **+ Add custom claim**, enter the claim that you collected from step 8.a, on page 18, and press Enter.

   If you are using Microsoft Entra, enter "roles" as the claim.

   If you do not press Enter, the claim will not be added.

k) For **Just-In-Time (JIT) settings**, check **Create or activate users** and **Update user with OIDC claims**.

   These settings enable the federation to provision records for the users as they log in and the remote IdP authenticates them.

l) Ignore the remaining fields and select **Add IdP**.

Add a domain for an identity provider federation

Follow these substeps to add a domain for an identity provider federation.

Step 1 If you do not already have **Cisco Identity** open in your browser, follow these substeps.

a) In Cisco Nexus Hyperfabric, choose **Organization** > ***org_name***.

   org_name is the same organization for which you added an authentication protocol.

b) Choose **Administration** > **User management** > **External**.
c) Select **Launch admin portal**.

   The **Cisco Identity** site opens in a new tab.

Step 2 In Cisco Identity, go to **Settings** > **Manage Domains**.

a) Select **Add a domain**.
b) For **Domain name**, enter the name of your domain, then select **Add domain**.

   - Users who are part of the specified domain will get authenticated using the remote identity provider (IdP) that you configured for the IdP federation.
   - Nexus Hyperfabric uses an exact match of the domain name that you entered. If you entered a parent domain, subdomains are not included.

c) For **DNS verification** token, select **Copy**.
d) Go to the DNS records of the DNS server and paste the token in the TXT record section.

   This verifies that you own the domain.

e) In Cisco Identity, select **Done**.
f) On the **Domains** page in the row with the domain that you added, select **…** > **Verify domain**.

After Cisco Identity verifies the domain, the status for the domain changes from "Pending" to "Verified". You no longer require the TXT record and can remove the verification token from your DNS server.

If the Cisco Identity fails to verify the domain, your DNS server caches the error. Your DNS server clears the cache after the length of time specified in the **Time To Live (TTL)** setting passes. At that time, you can add the verification token again and request Cisco Identity to verify the domain.

g) Select **Skip for now**.

Add a routing rule for an identity provider federation

Follow these steps to add a routing rule for an identity provider federation.

Step 1 If you do not already have **Cisco Identity** open in your browser, follow these substeps.
a) In Cisco Nexus Hyperfabric, choose **Administration** > **User management** > **External**.
b) Select **Launch admin portal**.

The **Cisco Identity** site opens in a new tab.

Step 2 In Cisco Identity, go to **Settings** > **Manage IdPs** > **Routing rules**.
a) Select **+ Add a routing rule**.
b) For **Rule name**, enter a name.
c) For **Select the routing type**, select **Domain**.
d) For **If these are your domains**, select the domain that you added.

If you added multiple domains, you can select more than one.

e) For **Then use this identity provider**, select the IdP that you added.
f) Select **Add**. The **Routing Rules** page displays again.
g) In the **Routing rules** table, for the row with the routing rule that you added, select **...** > **Activate**.

The status for the rule changes from "Inactive" to "Active".

Verify the identity provider federation configuration

After you have completed configuring your desired identity provider (IdP) and Cisco Identity for your new IdP federation, you can run a test to ensure that your federation works properly.

Follow these steps to verify the identity provider federation configuration.

Step 1 In Cisco Identity, go to **Settings** > **Manage IdPs** > **Identity providers (IdPs)**.
Step 2 Select the IdP that you added.
Step 3 In the pane that appears on the right, select **Actions** > **Test SSO Setup**.

The test will inform you if there are any issues with the configuration.

Common troubleshooting scenarios for identity provider federations

These are some common troubleshooting scenarios for identity provider (IdP) federations.

After creating an IdP federation, I can no longer log into Nexus Hyperfabric

Most likely, you misconfigured something, such as either forgetting to add the custom claim or entering the claim incorrectly. You can log into Nexus Hyperfabric using your cisco.com email address to bypass the misconfigured IdP federation by prefixing your email address with "cisco/". For example, "cisco/myEmailAddress@cisco.com". After logging in, fix the configuration.

All Nexus Hyperfabric administrators got removed from the remote IdP or the only administrator for your Nexus Hyperfabric organization is no longer an active user

You must have at least one Nexus Hyperfabric administrator in your organization. To resolve this issue, log into your remote IdP and either add a new administrator or upgrade an existing user to the administrator role.

Organizations

An organization represents a business entity, such as a company or business unit, within the Nexus Hyperfabric platform. It is the highest-level tenant construct designed to securely partition all data, ensuring that data is inaccessible to other organizations in Cisco Nexus Hyperfabric. Each organization contains fabrics, which represent physical presences within data centers composed of switches and servers. Users are assigned roles within organizations, with access controls applied on an organization-by-organization basis.

Region-based organizations

A region-based organization is a Nexus Hyperfabric organization that

- is bound to a single geographic region at the time of creation
- stores and processes its device telemetry and services within that region only, and
- helps you align with regional data handling requirements.

Cisco Nexus Hyperfabric supports these regions for organization creation:

- European Union (EU)
- United States (US)

How to identify an organization's region

The organizations list shows a region icon next to each organization name. Use the icon to quickly tell apart which region an organization resides in.

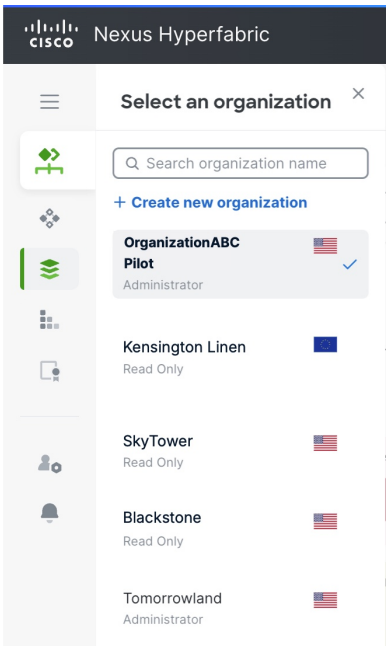


Table 5: Region icons

Region icon	Region
	European Union
	United States

Create a new organization

Follow these steps to create an organization.

Step 1 Choose **Organization** > **Create new organization**.



Step 2 Enter the name of the new organization.

Step 3 Select the region in which the organization resides.

You can select either European Union (EU) or United States (US). The region determines where the device telemetry and services for the organization are stored and processed.

Step 4 Select **Save**.

Guidelines for creating an organization

Follow these guidelines when you create an organization in Cisco Nexus Hyperfabric.

Organization name guidelines

The organization name must meet these requirements:

- The name must be between 3 and 120 characters.
- You can include these special characters in the name, as long as they are not the first or last character:
  - -
  - _
  - +
  - (space)
  - /
  - @
  - .
  - &

  You cannot include any other special characters.

Organization region guidelines

Follow these guidelines when you select the region for an organization:

- Select the region based on where you want the organization's telemetry and service data to reside.

- Cisco Nexus Hyperfabric supports identity provider federation only for organizations in the US region. If you create an organization in the EU region, users must log in with their cisco.com credentials.

Caution

You cannot change the region after you create the organization.

The region you select applies to the organization permanently. To have your fabrics, devices, or telemetry data reside in a different region, you must create a new organization in that region.

Organization limits and ownership

These limits and ownership rules apply to organizations:

- You can create up to 10 organizations.
- You cannot delete an organization after you create it.
- You are by default an administrator of any organization that you create.

Select an organization

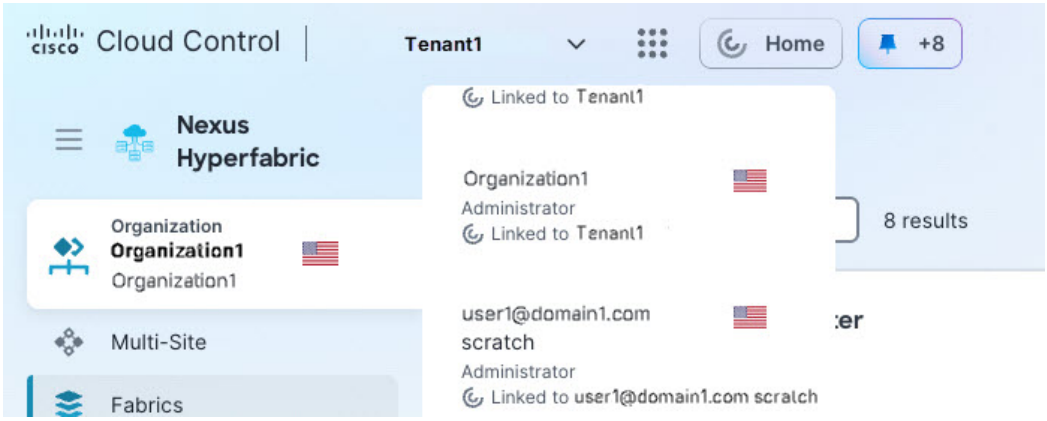
When you log into Cisco Nexus Hyperfabric, if you are a member of at least one organization, Nexus Hyperfabric selects one of your organizations automatically. If desired, you can select a different organization.

Follow these steps to select an organization.

Step 1 Select **Organization**.

The **Select an organization** drawer opens.

Figure 23: Select an organization

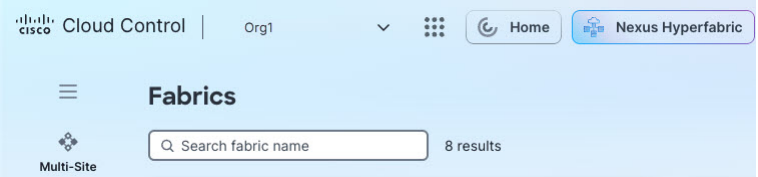


Step 2 (Optional) Use the **Search organization name** field to filter the list of organizations.

Step 3 Select an organization.

If the organization that you selected is linked to Cisco Cloud Control, the page's header indicates this fact.

Figure 24: Cisco Cloud Control header



User settings

Explains how to use bearer tokens for REST API authorization and change the default time zone and time format for a user.

Bearer tokens

Bearer tokens grant you authorization to use the Cisco Nexus Hyperfabric REST API calls. For more information about bearer tokens, as well as how to create them, refer to Cisco Nexus Hyperfabric API Authentication on Cisco DevNet.

Change the default time zone for user

You can choose the preferred time format and time zone that Cisco Nexus Hyperfabric uses. By default, Nexus Hyperfabric displays timestamps based on your browser's local time zone.

Step 1 Select your username in the upper right of the GUI.

Step 2 From the **Timezone** drop-down list, select the desired the time zone.

Step 3 Set the time format to **12hr** or **24hr**.

All timestamps in the GUI display the new time zone and format.

Data sharing with other applications

Describes how to share basic device inventory, telemetry, and notification data from Nexus Hyperfabric with other applications.

Share data with Cisco Cloud Control

Cisco Cloud Control is an AI-native platform that provides centralized management, automation, and enhanced visibility into your Cisco ecosystem. After you onboard Nexus Hyperfabric to Cisco Cloud Control, you must configure your organization to share data with Cisco Cloud Control. The shared data includes basic device inventory, telemetry, and notification data. For more information about onboarding Nexus Hyperfabric to Cisco Cloud Control, refer to Cisco Cloud Control Getting Started.

Follow these steps to share data with Cisco Cloud Control.

Step 1 Choose **Administration** > **Organization settings**.

Step 2 Toggle on **Enable data sharing**.

Designing the fabric blueprint

When you have created a new organization or need to add a fabric to an existing organization, you can use the Blueprint Designer in the Cisco Nexus Hyperfabric to create a blueprint for the new fabric.

For supported fabric topologies, refer to Nexus Hyperfabric use cases.

The procedures for additional configuration, including port channels, logical networks, and routing tables, are described in other Cisco Nexus Hyperfabric articles.

Create a fabric using a template

You can create a fabric using a template, which has a predefined topology. You select a template based on the size of the fabric that you want to create.

Follow these steps to create a fabric using a template.

Step 1 Choose **Fabrics**.

Step 2 Select **+ Add new fabric** and select **New from template**.

Step 3 On the **Templates** page, select **AI clusters** or **Switch fabrics** as appropriate.

Step 4 Select the template with the topology that you want.

Step 5 If you want to change the default fabric name, select edit (✎) next to the name, enter a new name, and select **Update**.

Step 6 Modify the topology as appropriate, then select **Save fabric blueprint**.

Create a fabric without using a template

Follow these steps to create a fabric without using a template.

Step 1 Choose **Fabrics**.

Step 2 Select **+ Add new fabric** and select **New fabric**.

a) If you want to change the fabric name, select edit (✎) next to the default name, modify the name, and select **Update**.

The name must be unique and can have from 3 to 120 characters. The name can contain a hyphen (-), but no other special characters.

By default, the fabric has one switch tier.

b) If you want to upgrade the fabric entitlement level, select **Upgrade**.

The fabric's current entitlement level displays to the right of the fabric name. For more information about entitlement levels, refer to Cisco Nexus Hyperfabric - Subscriptions.

Step 3 Follow these substeps to configure the switch tiers.

a) If you want to change the default switch tier name, select edit (✎) next to the default name, modify the name, and select **Update**.

By default, the switch tier has one switch group.

b) If you want to change the default group name and short name, select edit (✎) next to the default name, modify the name and short name, and select **Update**.

c) From the drop-down list at the top of the group, select the switch model.

d) Enter a name for the hardware set and fill out the other fields.

   Cisco Nexus Hyperfabric uses the hardware set name as part of the device name.

   By default, the group has one hardware set. If you want to add another hardware set, select **Add hardware set**, then repeat this substep.

e) If you want to add another switch group, select **Add group**, then repeat substep 3.b, on page 26 and the succeeding substeps.

f) If you want to add another switch tier, select **Add tier**, select **Switches**, and repeat this entire step.

To design a single-switch fabric or a mesh fabric with a switch model, create only one switch group. If you need to add another switch model, add another switch group. For example, you can have two switch groups—one for spine switches and one for leaf switches—to create a spine and leaf topology.

As you make your switch choices, Cisco Nexus Hyperfabric updates the topology.

Step 4 In the **Switch groups** area of **Group connections**, select a box to configure the connections between the corresponding switch groups.

The connections between the switch groups are displayed in a matrix of boxes. The matrix defines how the switches are connected to one another. The top and side labels show the names of the switch groups; the labels match on both axes.

Each box shows the number of connections per switch group pair and is labelled with a letter such as A, B, or C. The **Connections** property specifies the number of connections, which you can change as indicated in the substeps. Each letter indicates a unique pair of switch PIDs, which can help you to identify outliers or misconfigurations quickly.

The switches of the switch group specified by the left label are connected under the switches of the switch group indicated by the top label. The exception is if the switch group is the same on both axes, in which case the switches of that group are interconnected.

As an example, assume you have two switches in a switch group named "Spine", four switches in another switch group named "Leaf", and you specify "2" for **Connections** between the groups. This means you have eight pairs of switches. Because you specified two connections, you have 16 total connections between the groups.

After you select a box, the right side of the **Switch groups** area shows the properties of the group connections. You can select multiple boxes to review each of those group connections at the same time.

For a selected switch group's connections, perform these substeps.

a) Select the PID for the pluggables by selecting edit (✎).

   The **Select PIDS for <switch_group> to <switch_group>** dialog appears. Select the PID from the table at the bottom for the first switch group, select **Cable** and select a PID, then select the radio button for the second switch group and select a PID. You can filter the table using the fields above the table.

   Select **Select** after you are done.

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

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

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

d) If you want to override the automatic cabling of the switches, in the **Manual Overrides** area, select **Add override** and select **Connection override**.

   The radio button on the left is selected by default. In the left tile, select the device and port from the drop-down lists, and select the pluggable from the table.

   Select **Cable** and select a PID.

Select the radio button on the right. In the right tile, select the device and port from the drop-down lists, and select the pluggable from the table.

Select **Select**.

Step 5 If your fabric will contain servers, follow these substeps.

a) Select **Add tier** and select **Server**.

By default, the server tier has one server group.

b) If you want to change the default server tier name, select edit (✎) next to the default name, modify the name, and select **Update**.

c) If you want to change the default group name and short name, select edit (✎) next to the default name, modify the name and short name, and select **Update**.

d) From the drop-down list at the top of the group, select the server model.

e) Enter a name for the hardware set and fill out the other fields.

Cisco Nexus Hyperfabric uses the hardware set name as part of the device name.

By default, the group has one hardware set. If you want to add another hardware set, select **Add hardware set**, then repeat this substep.

f) If you want to add another server group, select **Add group**, then repeat substep 5.c, on page 28 and the subsequent substeps.

g) If you want to add another server tier, repeat this entire step.

h) If your fabric includes devices with GPUs, then in the **Manage rail groups for East-West server port connections** area, select **Add rail group** and fill out the fields for the group.

- **Rail group name**: Enter a name.
- **Short name**: Enter a short name.
- **Width**: Enter the number of GPU servers.
- **Count**: Enter 8. Currently, this 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.

i) In the **Server port groups to switch groups** area, configure these things:

- **Pluggable PID <–**: Select edit (✎) and select the pluggable for the group server. You can select the pluggable for **Pluggable PID <–** (group server), **Cable PID**, and **Pluggable PID –>** (switch group or rail group) in the same dialog.
- **Connections**: This is the maximum number of possible connections. You cannot change this value. If you do not want any connections to the server port group, remove the value from **Switch/rail group**, if any.
- **Cabling strategy**: Select a strategy.
- **Cable PID**: Select edit (✎) and select the pluggable for the cable. You can select the pluggable for **Pluggable PID <–** (group server), **Cable PID**, and **Pluggable PID –>** (switch group or rail group) in the same dialog.
- **Pluggable PID –>**: Select edit (✎) and select the pluggable for the switch group. You can select the pluggable for **Pluggable PID <–** (group server), **Cable PID**, and **Pluggable PID –>** (switch group or rail group) in the same dialog.

- **Switch/rail group**: Select a switch group or rail group.

Step 6 Select **Save fabric blueprint**.

The **Saving fabric blueprint** dialog appears and informs you that Cisco Nexus Hyperfabric generated the blueprint.

Step 7 If you want Cisco Nexus Hyperfabric to create the cabling plan now, select **Run cabling**. Otherwise, select **Continue to topology**.

Modifying fabrics

If you have administrator or read-write access to the organization, you can modify the design of an existing fabric on the **Fabrics** page. You can make changes to the fabric blueprint at any stage: from an undeployed fabric design to an installed and running fabric.

The **Fabrics** blueprint page operates in one of two modes:

- Edit mode: This is the default mode for a user with administrator or read-write access. All fabric edit options are exposed.
- Running mode: This is the default mode for a user with read-only access. Settings and status are displayed, but no edit options are exposed. You can switch to edit mode only if you have administrator or read-write access.

All configuration procedures in this document assume that you are logged in with administrator or read-write access, and that the menus are in edit mode.

Modify a fabric

This procedure provides the typical workflow for making fabric configuration changes.

Follow these steps to modify a fabric.

Step 1 Select **Fabrics**, then select the fabric that you want to configure.

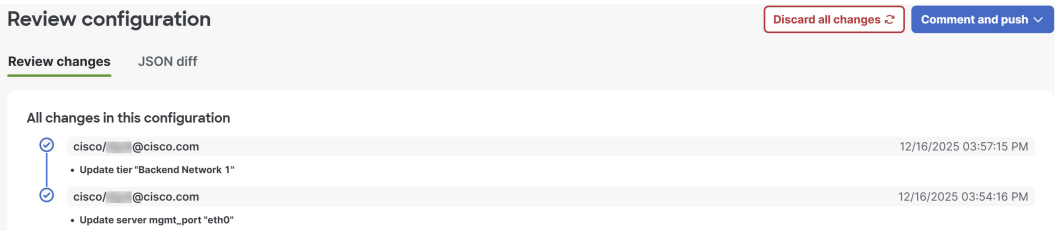
a) If the fabric is not in the edit mode, select **Switch to edit mode** ( ).

Step 2 Make your changes.

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

Step 3 When you are ready to apply your changes, select **Review configuration**.

Figure 25: Review configuration



In the **Review configuration** page, you can review a list of brief descriptions of the changes since the last push, including the author's ID and the date and time when the change was saved.

Step 4 When you have reviewed the changes and are ready to apply them, select **Comment and push**.

Step 5 In the **Comment before pushing configuration** dialog box, enter the reason for the change.

Step 6 Select **Push configuration**.

Step 7 Review the status of your changes.

On the **Fabrics** blueprint page, the fabric card presents the status of the pushed changes. Your changes may take some time to apply or they might generate red assertions. For example, if you push changes to an installed and running fabric, and the changes conflict with the physical cabling of the fabric, an assertion is raised.

Deploying your fabric on site

The Cisco Nexus Hyperfabric greatly simplifies cabling operations of the switches in a fabric by providing features to guide tasks that you or a network cabling technician must perform on site. A network administrator uses Cisco Nexus Hyperfabric to create a blueprint that includes a cabling plan for the switches. You can review and follow tasks to set up the physical fabric topology or track the progress by going to a dedicated on-site page.

Alternatively, the network administrator can give a URL for the Cisco Nexus Hyperfabric on-site mobile-friendly site to the network cabling technicians. The technicians then open the URL on a mobile device to review the blueprint and follow that to set up the physical fabric topology or to review the status of the fabric.

For more information about deploying your fabric on site, refer to *Cisco Nexus Hyperfabric — Deploying Your Fabric on Site*.

Get the shareable link to the Cisco Nexus Hyperfabric on-site mobile-friendly site

Follow these steps to get the shareable link to the Cisco Nexus Hyperfabric on-site mobile-friendly site.

Step 1 In the Cisco Nexus Hyperfabric GUI, choose **Fabrics**, then select the fabric for which you want to get the link.

Step 2 Select **On-site**

Step 3 Select **Shareable link**.

Step 4 Go to the on-site mobile friendly site using one of these methods.

- Scan the QR code with a mobile device to open the on-site page on that device.
- Select **Copy** to copy the URL, then send the URL to a network cabling technician who can then go to the URL using a mobile device.

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.

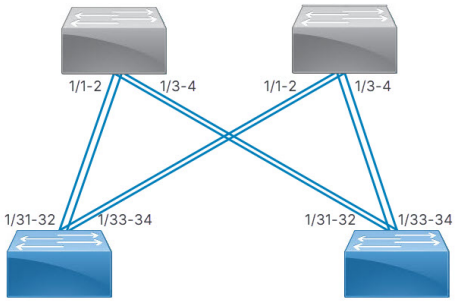
The first notion comprises these things:

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

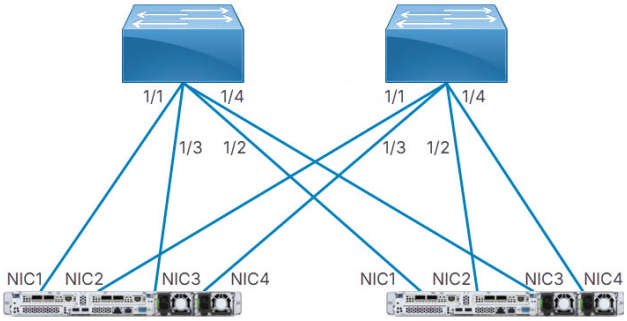


- Distributed: Cisco Nexus Hyperfabric connects the first network interface card (NIC) of the first server to the first port of the first switch, then the second NIC of the server to the first port of the second switch, and so on until the server has one connection to the first port of each switch. Cisco 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 27: 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.

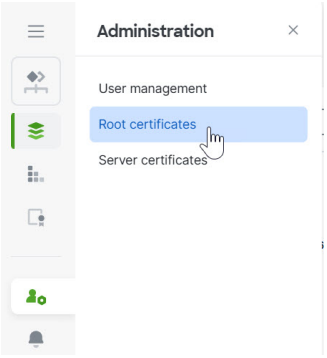
Add a root certificate for secure Transport Layer Security connectivity

You can configure switches to trust root certificates issued by their organization for establishing a secure Transport Layer Security (TLS) connection to Cisco Nexus Hyperfabric.

Follow these steps to add a root certificate for secure Transport Layer Security connectivity.

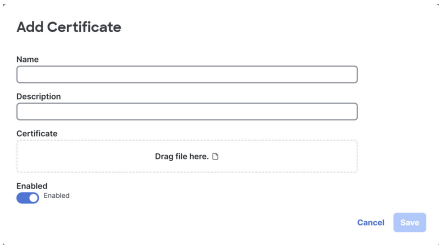
Step 1 Choose **Administration** > **Root certificates**.

Figure 28: Root certificates



Step 2 Select + **Add Certificate**.

Figure 29: Add Certificate dialog



Step 3 Fill out the fields, then select **Save**.

For **Certificate**, drag the file with the certificate to the box, or select the box, navigate to the file, select the file, and select **Open**.

Configuring switch 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 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.

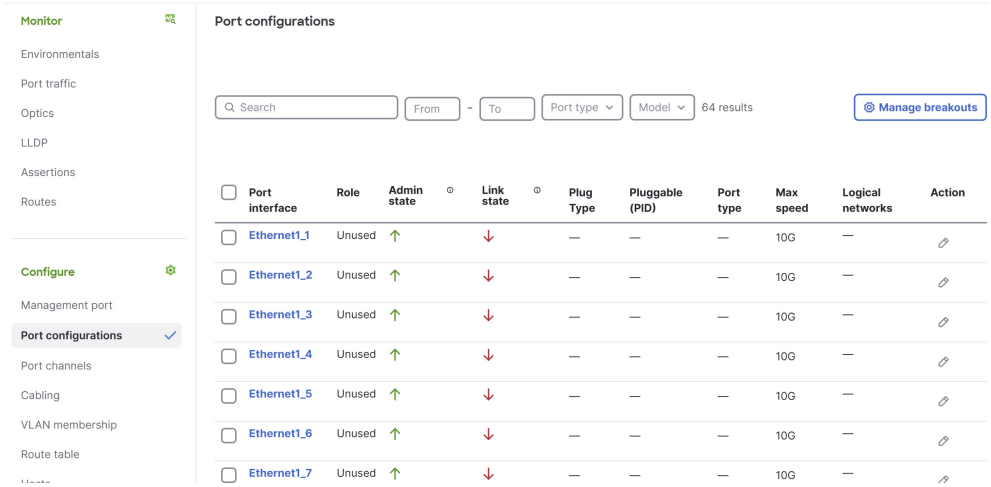
a) If the fabric is not in the edit mode, select **Switch to edit mode** ( ).

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 30: 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.

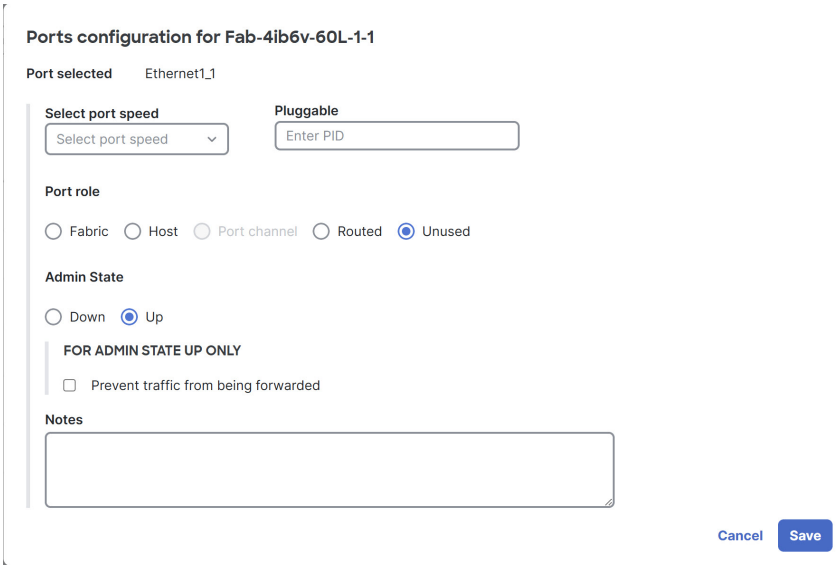
b) If you enabled breakout, select the cable type and optic type.

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 31: 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.

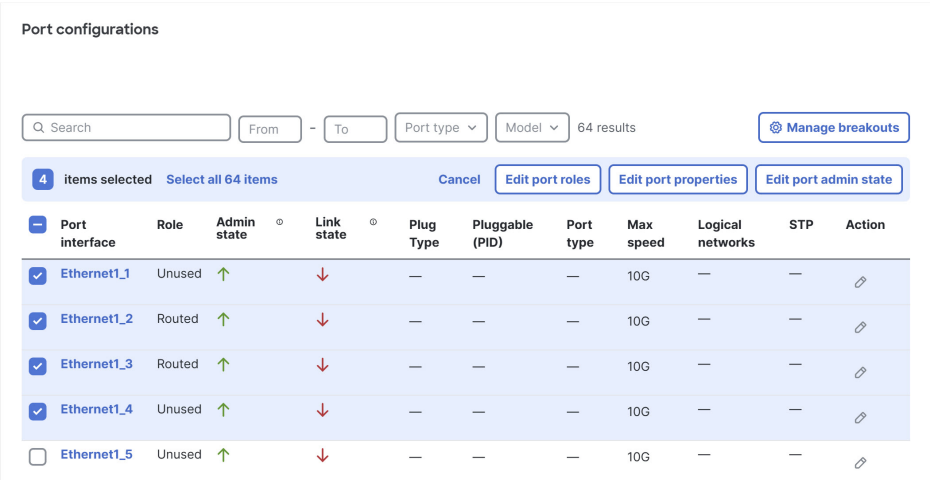
Step 1 Choose **Fabrics**, then select the fabric that contains the switch.

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 32: 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**.

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.

Finish and commit the design

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

Follow these steps to finish and commit the design.

Step 1 Select **Review configuration**.

Step 2 Verify your changes in the review list.

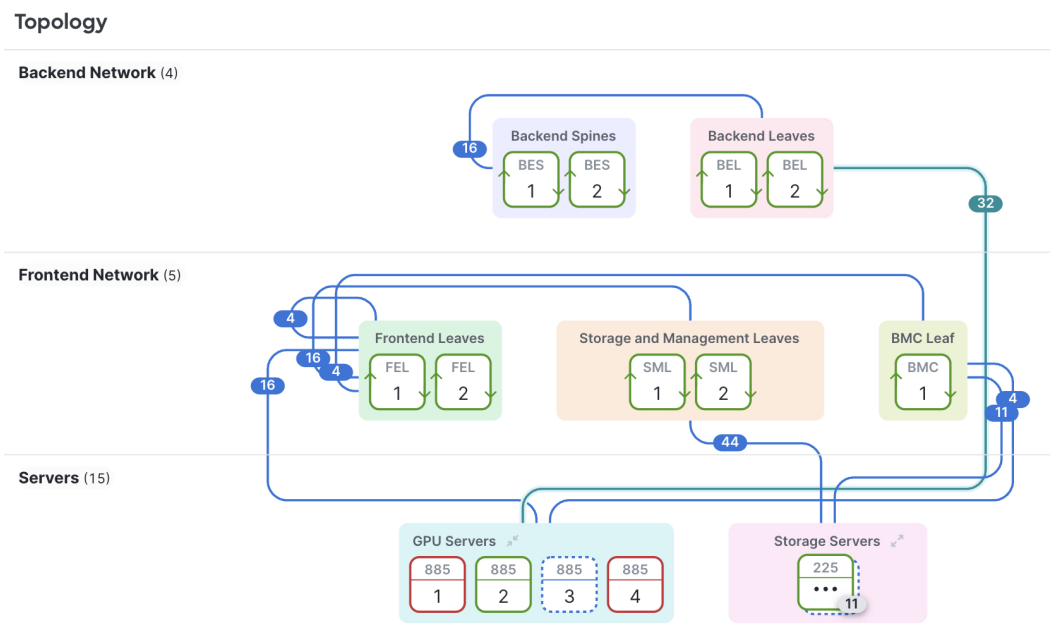
Step 3 Select **Comment and push**.

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

Step 5 Select **Push configuration**.

If you bound any devices as part of this push, the outline color of the device icons indicates the assertions for the devices. The devices should have a green outline, which indicates that there are no issues. A yellow outline indicates the device has an unexpected condition and there are issues detected, but the issues should not cause operation disruption. A red outline indicates the device has an unexpected condition and there are issues detected that can cause operation disruption; you should investigate the issues.

Figure 33: Fabric topology after pushing a configuration



Bill of materials and ordering

Deployment

Deployment provides information about the assembly list and bill of materials (BOM) calculations and specifies if you need any parts for your fabric. You can export a CSV file of the assembly list and BOM for reference outside of Cisco Nexus Hyperfabric.

Review or export deployment information

You can review deployment information or also export the deployment information to a CSV file to use as reference. You can review or export this information:

- A bill of materials for ordering products to build the fabric according to the blueprint design

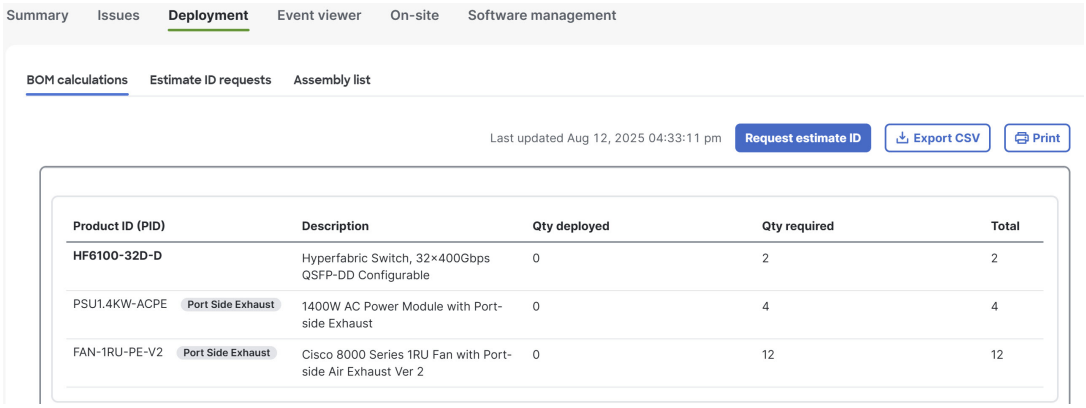
- An assembly list to check the fulfillment of the necessary accessories, such as fans, power supplies, and optics required for each device

Follow these steps to review or export deployment information.

Step 1 Choose **Fabrics**, then select the desired fabric.

Step 2 Select **Deployment** > **BOM calculations**.

Figure 34: BOM calculations

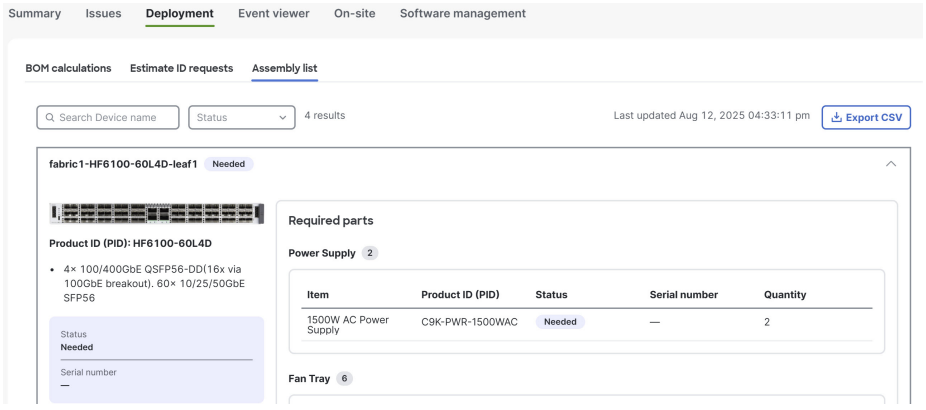


This page displays a bill of materials (BOM) of product IDs (PIDs), which specify the devices and accessories needed for this fabric design, as well as the optics based on the port connection configuration. The table shows how many of each PID are required and how many have been received and deployed.

Step 3 (Optional) To export the BOM as a comma-separated value (CSV) spreadsheet, select **Export CSV**.

Step 4 Select the **Assembly list**.

Figure 35: Assembly list



This page displays a list of accessories, such as fans, power supplies, and optics required for each device. The table indicates whether the parts have been deployed and, if available, the serial numbers of the deployed parts.

Step 5 (Optional) To export the assembly list as a comma-separated value (CSV) spreadsheet, select **Export CSV**.

Request an estimate ID for a bill of materials

When you request an estimate ID for a bill of materials (BOM), Cisco Nexus Hyperfabric communicates with the Cisco Commerce Workspace (CCW) to create the estimate and sends you an email with the estimate ID. You can then work with your partner or sales team contacts to discuss the estimate and order the parts.

Cisco Nexus Hyperfabric retains estimate IDs for 90 days. From the **Estimate ID requests** page, you can download any of the BOMs to save for a longer amount of time. If an estimate ID expires, you can request a new estimate ID.

Follow these steps to request an estimate ID for a bill of materials.

Step 1 Choose **Fabrics**.

Step 2 Select the fabric that has the BOM for which you want the estimate ID.

Step 3 Select **Deployment** > **BOM calculations**.

Step 4 Select **Request estimate ID** and perform these substeps.

a) In the **Request estimate ID** dialog box, select **Submit request**.

The dialog box closes and Cisco Nexus Hyperfabric sends you an email from "Cisco Hyperfabric <noreply@hyperfabric.cisco.com>". The body of the email contains the estimate ID and the estimate name.

b) (Optional) If you want to refer to information about this estimate ID and any past estimate IDs, select **Deployment** > **Estimate ID requests**.

From this page, you can download a BOM in the format of a CSV file by selecting the download button in the **Action** column.

Step 5 Provide your Cisco partner or sales team contacts with the estimate ID.

The estimate created by Cisco Nexus Hyperfabric is read-only. Your contacts must create a clone of the estimate, which will have a new estimate ID. Your contacts will then assist you with the cost estimate and complete your order using the new ID.

Onboarding the fabric

Onboard a fabric

This procedure provides a high-level overview of the tasks you must perform to onboard a fabric.

Before you can onboard a fabric, you must physically install your fabric and wire the fabric to match the Cisco Nexus Hyperfabric blueprint.

Follow these steps to onboard a fabric.

Step 1 Allow the switches to contact the Cisco Nexus Hyperfabric cloud service.

Step 2 Claim the switches to your organization.

For more information, refer to Claiming devices, on page 57.

Step 3 Bind each switch to its position in your fabric blueprint, allowing it to receive its specific configuration from the Cisco Nexus Hyperfabric.

For more information, refer to Bind the devices, on page 59.

Step 4 If you are deploying AI clusters, follow these substeps.

a) Claim the Cisco UCSC-C225-M8N-1P servers for storage.

For more information, refer to Claiming devices, on page 57.

b) Bind the UCSC-C225-M8N-1P servers.

For more information, refer to Bind the devices, on page 59.

c) Install the Cisco Nexus Hyperfabric agent on the servers.

For more information, refer to Install the Cisco Nexus Hyperfabric agent on a server.

d) For C880 servers, add a certificate to each server.

For more information, refer to one of these procedures:

- Add a certificate to a server using a device onboard token and the Craft console, on page 47
- Add a certificate to a server using a device onboard token and a CSR enrollment script, on page 49
- Add a certificate to a server using a bearer token and the Craft console, on page 51
- Add a certificate to a server using a bearer token and a CSR enrollment script, on page 55

e) Claim the C880 servers for the GPUs.

For more information, refer to Claiming devices, on page 57.

f) Bind the C880 servers.

For more information, refer to Bind the devices, on page 59.

Prerequisites for onboarding

Before you can onboard your installed fabric, you must have the following:

- A working Internet connection to which a switch can connect for HTTPS (port 443) traffic
- An available DNS server
- A configured proxy server, if needed to reach the cloud
- At least one management IPv4 or IPv6 address for a switch in your fabric

Note

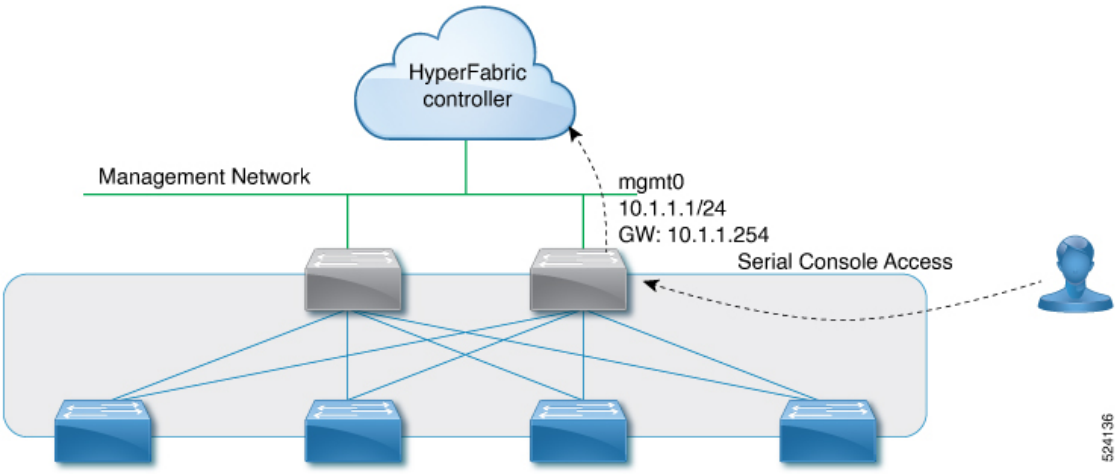
We recommend that you connect all switches in your fabric to a management network.

Connect the first switch to the Cisco Nexus Hyperfabric cloud service

Upon delivery, a new Nexus Hyperfabric switch contains an agent that allows it to discover its neighbors and to contact Cisco Nexus Hyperfabric.

The following diagram shows an example of configuring the first switch for out-of-band (OOB) management node connection to the Cisco Nexus Hyperfabric. The management node configuration can be achieved through DHCP, if available, or by serial port configuration if you need a static IP address.

Figure 36: Configuring the first switch for out-of-band management node connection to Cisco Nexus Hyperfabric



Typically, you will need to connect physically only one switch to the cloud. This switch automatically contacts Cisco Nexus Hyperfabric. When discovered by its neighbor switches in the fabric, this switch reports its availability as the "exit node" switch. Neighboring switches that lack OOB management connectivity can connect to Cisco Nexus Hyperfabric through the exit node switch.

Follow these steps to connect the first switch to the Cisco Nexus Hyperfabric cloud service.

Step 1 At the physical site, select a switch that will have the initial Internet connection. This cloud-connected switch will become the exit node.

Step 2 Connect the switch's OOB management interface to an Internet connection.

Step 3 If a DHCP server is present, you can skip the remaining steps in this procedure.

The switch will receive its DHCP configuration and will then contact Cisco Nexus Hyperfabric through the Internet connection.

Step 4 If you must configure a static IP address or if you must configure a proxy, connect a terminal to the serial console port of the switch using these connection parameters:

- 115200 baud
- no parity
- 8 data bits
- 1 stop bit

Step 5 Log in to the serial console using "ttg-admin" as the username and "hfabric2025" as the password.

You are asked whether you want to change the default password, which we recommend that you do.

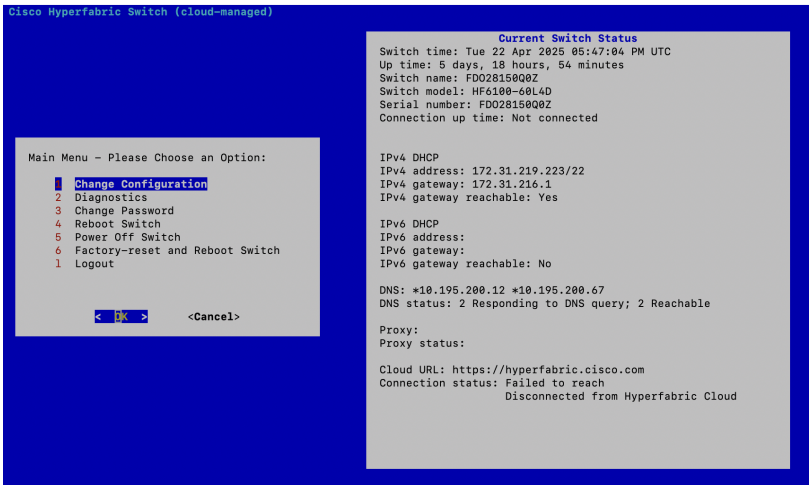
a) In the **Change default password** dialog, select **Yes**.
b) Enter the current password ("hfabric2025"), then enter your new password twice.

```
Changing password for ttg-admin.
                                    Current password:
                                    New password:
                                    Retype new password:
                                    passwd: password updated successfully
                                    Press any key to continue
```

c) In the **Default password is changed successfully** dialog, select **OK**.

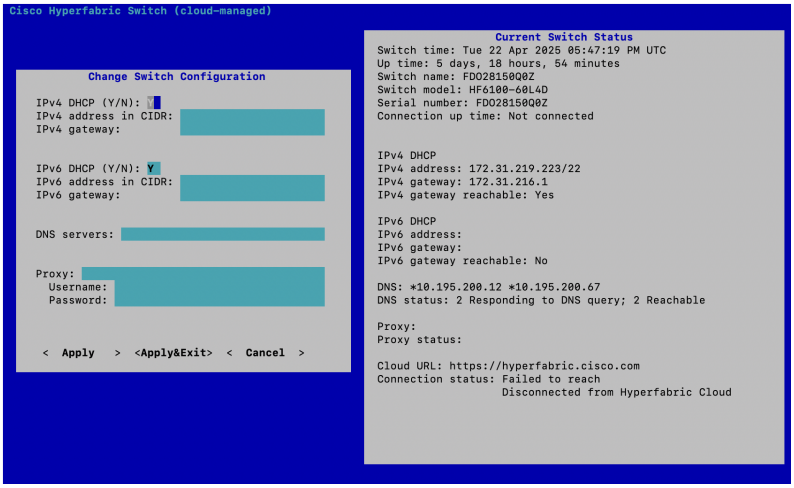
Step 6 In the console menu, select **Change Configuration**.

Figure 37: Change configuration



a) Configure at least one management IPv4 or IPv6 address, or you can select DHCP if a DHCP server is present.

Figure 38: Change switch configuration



b) Configure one or more DNS addresses.

If you chose DHCP, keep these things in mind:

- You might not need any DNS addresses, although you can still configure DNS addresses if desired.
- You cannot have both a static DNS server configuration and a dynamic DNS server configuration. Cisco Nexus Hyperfabric overwrites the static configuration with the dynamic configuration.

c) Configure a proxy setting, if needed.

Step 7 Select **Apply**.

a) Using the console, ensure that the **Connection status** shows that the switch is connected to the Cisco Nexus Hyperfabric cloud.

Neighboring switches can also now discover and access the Cisco Nexus Hyperfabric connection through this switch.

If the switch is not connected, you must troubleshoot the issue before you can proceed.

Tip In the console, you can check the reachability of the gateway, the DNS server, and Cisco Nexus Hyperfabric.

Preparing servers

Install the Cisco Nexus Hyperfabric agent on a server

You must install the Cisco Nexus Hyperfabric agent on a server to enable Nexus Hyperfabric to control the server.

Follow these steps to install the Cisco Nexus Hyperfabric agent on a server.

Step 1 Use SSH to connect to the server.

Step 2 Configure the NVIDIA BlueField-3 (BF3) data processing unit to use the NIC mode.

a) On the server, create a privileged pod.

This example uses "ec-e7-a7-11-f7-9c" as the node name.

```
admin@cluster3-centos-vm4:~$ oc get node
NAME                STATUS   ROLES                         AGE   VERSION
ec-e7-a7-11-f7-9c   Ready    control-plane,master,worker   58m   v1.32.9

admin@cluster3-centos-vm4:~$ oc debug node/ec-e7-a7-11-f7-9c --image=centos:10
Starting pod/ec-e7-a7-11-f7-9c-debug ...
To use host binaries, run `chroot /host`
Pod IP: 172.27.127.212
sh-5.2#
```

You might need to press Enter to receive the prompt after you run this command.

b) Install **pciutils** and **mstflint**.

```
sh-5.2# yum install -y pciutils mstflint
```

c) If you need to use a proxy server for connecting to the Internet, edit `/etc/yum.conf`.

This example shows a proxy configuration in `/etc/yum.conf`.

```
sh-5.2# cat /etc/yum.conf
[main]
gpgcheck=1
installonly_limit=3
clean_requirements_on_remove=True
best=True
skip_if_unavailable=False
proxy=http://proxy.example.com:80
```

d) Change the BF3 to the NIC mode.

```
sh-5.2# lspci | grep ConnectX-7 | grep -v "\.1" | awk '{print $1}' | xargs -I {} mstconfig -y \
-d {} set INTERNAL_CPU_OFFLOAD_ENGINE=1
```

e) Exit the pod.

```
sh-5.2# exit
```

Step 3 Power cycle the server.

a) Connect to the baseboard management controller (BMC) of the server.
b) Power down the server and wait 30 seconds.
c) Power on the server.

Step 4 Install the Nexus Hyperfabric agent on the server.

a) Download `hf_agent_deploy.yaml`.

Download the file from here:

https://cdn.hyperfabric.cisco.com/server/releases/agent/rhos/latest/hf_agent_deploy.yaml

b) Copy `hf_agent_deploy.yaml` to the server.
c) Run the **apply** command on `hf_agent_deploy.yaml`.

```
admin@cluster3-centos-vm4:~$ oc apply -f hf_agent_deploy.yaml
namespace/cisco-hf created
serviceaccount/cisco-hf-agents-sa created
clusterrolebinding.rbac.authorization.k8s.io/cisco-hf-agents-privileged created
persistentvolume/cisco-hf-agents-pv created
persistentvolumeclaim/cisco-hf-agents-pvc created
daemonset.apps/cisco-hf-agents created
```

d) Verify that the agent installed successfully.

```
admin@cluster3-centos-vm4:~$ oc get pods -n cisco-hf
NAME                    READY   STATUS    RESTARTS   AGE
cisco-hf-agents-8pz6b   1/1     Running   0          13s

admin@cluster3-centos-vm4:~$ oc describe  pods -n cisco-hf
Name:         cisco-hf-agents-8pz6b
Namespace:    cisco-hf
Priority:     0
Node:         ec-e7-a7-11-f7-9c/172.27.127.212
Start Time:   Thu, 16 Oct 2025 16:27:46 -0700
<--- skipped --->
Events:
  Type     Reason            Age                    From               Message
  ----     ------            ----                   ----               -------
  Warning  FailedScheduling  4m48s (x2 over 4m50s)  default-scheduler  0/1 nodes are available:
pod has unbound immediate PersistentVolumeClaims. preemption: 0/1 nodes are available: 1 Preemption
 is not helpful for scheduling.
  Normal   Scheduled         4m46s                  default-scheduler  Successfully assigned
cisco-hf/cisco-hf-agents-8pz6b to ec-e7-a7-11-f7-9c
  Normal   Pulling           4m45s                  kubelet            Pulling image
"public.ecr.aws/z9f4t0k6/hyperfabric/cisco_hf_agents:0.1.10"
  Normal   Pulled            4m39s                  kubelet            Successfully pulled image
 "public.ecr.aws/z9f4t0k6/hyperfabric/cisco_hf_agents:0.1.10" in 6.851s (6.851s including waiting).
 Image size: 347159297 bytes.
  Normal   Created          4m38s                kubelet           Created container: hf-agents

  Normal   Started          4m38s                 kubelet           Started container hf-agents
```

Add a certificate to a server using a device onboard token and the Craft console

After you install the Cisco Nexus Hyperfabric agent on a C880 server, you must then add a certificate to the server. You can use either a device onboard token or a bearer token, and can use the Craft console or a CSR enrollment script. We recommend that you use a device onboard token.

Follow these steps to add a certificate to a server using a device onboard token and the Craft console.

Step 1 Generate a device onboard token.

You must be an admin user to generate the token.

a) In the Nexus Hyperfabric GUI, select the user avatar (), then select **Generate tokens**.
b) Select **Device tokens**, then select **Generate device token**.
c) Select **Copy** to copy the device onboard token.

Step 2 Start the Craft console.

a) Use SSH to connect to the management IP address of the server as admin and enter the admin password when prompted.

```
$ ssh admin@management_ip_address
```

b) Run an interactive terminal session inside the cisco_hf_agent Docker container.

```
$ docker exec -ti cisco_hf_agent bash
```

c) Run `main.sh`.

```
# /opt/cisco/bin/tortuga/craft-console/main.sh
```

Step 3 Select **Delete Device Onboard Token** then select **Yes**.

You must delete an existing token before you can add a new one. If a token does not already exist, this message displays:

```
token file not found: /run/drake/crypto/device-onboard-token
Failed to delete device onboard token.
Press any key to continue
```

Press any key and proceed with the next step.

Step 4 Program the device onboard token in the server to add the certificate.

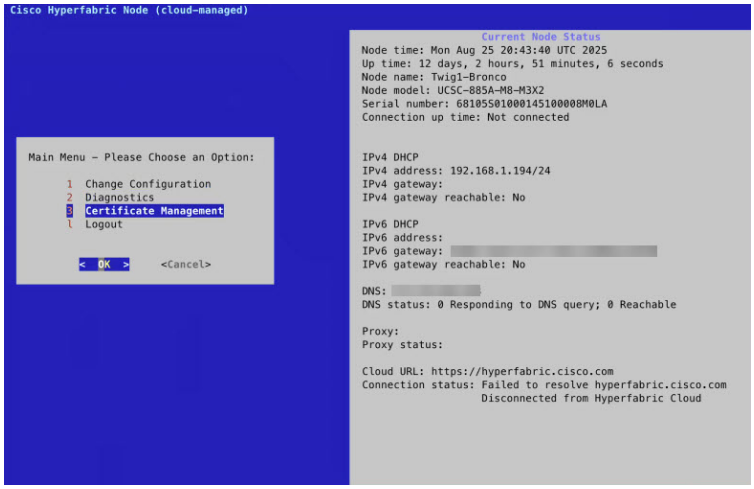
a) In the Craft console, select **Program Device Onboard Token**.
b) Paste the copied token, press Enter, then type Ctrl+D.
c) When prompted, press any key.

The console returns to the **Cert Enrollment Menu**.

Step 5 Verify that the certificate exists.

a) From the **Main Menu**, select **Certificate Management**.

Figure 39: Craft console main menu



b) Select **Check Certificate Status** to verify that the certificate now exists.

The console displays text similar to this example:

```
A valid certificate exists.
Press any key to continue
```

c) Press any key.

The console returns to the **Cert Enrollment Menu**.

Step 6 Select **Delete Device Onboard Token** then select **Yes**.

You must delete the token after you have added certificate.

Step 7 If you have a cluster, repeat this procedure for each server in the cluster.

Add a certificate to a server using a device onboard token and a CSR enrollment script

After you install the Cisco Nexus Hyperfabric agent on a C880 server, you must then add a certificate to the server. You can use either a device onboard token or a bearer token, and can use the Craft console or a CSR enrollment script. We recommend that you use a device onboard token.

Even though this procedure is for using a CSR enrollment script, you use the Craft console for some of the steps.

Follow these steps to add a certificate to a server using a device onboard token and a CSR enrollment script.

Step 1 Start the Craft console.

a) Use SSH to connect to the management IP address of the server as admin and enter the admin password when prompted.

```
$ ssh admin@management_ip_address
```

b) Run an interactive terminal session inside the cisco_hf_agent Docker container.

```
$ docker exec -ti cisco_hf_agent bash
```

c) Run `main.sh`.

```
# /opt/cisco/bin/tortuga/craft-console/main.sh
```

Step 2 Select **Delete Device Onboard Token** then select **Yes**.

You must delete an existing token before you can add a new one. If a token does not already exist, this message displays:

```
token file not found: /run/drake/crypto/device-onboard-token
Failed to delete device onboard token.
Press any key to continue
```

Press any key and proceed with the next step.

Step 3 Generate a device onboard token.

You must be an admin user to generate the token.

a) In the Nexus Hyperfabric GUI, select the user avatar (), then select **Generate tokens**.
b) Select **Device tokens**, then select **Generate device token**.
c) Select **Copy** to copy the device onboard token, then select **Done**.

Step 4 In the server's CLI, export the device onboard token.

a) If you are not already connected to the server, use SSH to connect to the server.
b) Export the device onboard token.

```
admin@cluster3-centos-vm4:~$ export DEVICE_ONBOARD_TOKEN="token"
```

Paste the device onboard token that you copied in place of *token*.

Step 5 Use a CSR enrollment script to add a certificate to the server.

a) Download one of the CSR enrollment scripts, depending on your setup.

If you have a standalone server that has Ubuntu or Red Hat, download the `server-csr-enroll.sh` script from here:

https://cdn.hyperfabric.cisco.com/server/releases/agent/utils/server-csr-enroll.sh

If you have a Kubernetes or OpenShift cluster, download the `cluster-csr-enroll.sh` script from here:

https://cdn.hyperfabric.cisco.com/server/releases/agent/utils/cluster-csr-enroll.sh

b) Run the CSR enrollment script that you downloaded to sign a certificate signing request (CSR) for the server or all servers in the cluster, as appropriate.

Run one of these commands:

```
admin@cluster3-centos-vm4:~$ ./server-csr-enroll.sh [--force]
```

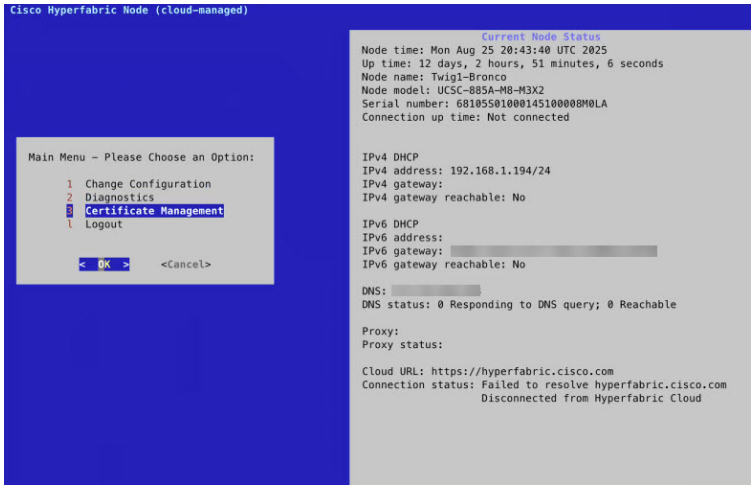
```
admin@cluster3-centos-vm4:~$ ./cluster-csr-enroll.sh [--force]
```

If the server or cluster already has a certificate, include **--force** to overwrite the certificate.

Step 6 Verify that the certificate exists.

a) In the Craft console, from the **Main Menu**, select **Certificate Management**.

Figure 40: Craft console main menu



b) Select **Check Certificate Status** to verify that the certificate now exists.

The console displays text similar to this example:

```
A valid certificate exists.
                  name: 000b5043bca227a1bc8d3e72e05abc6d073e40341337ee67ab401a1db66db8df123
                  Press any key to continue
```

c) Press any key.

The console returns to the **Cert Enrollment Menu**.

Step 7 If you have a cluster, repeat this procedure for each server in the cluster.

Add a certificate to a server using a bearer token and the Craft console

After you install the Cisco Nexus Hyperfabric agent on a C880 server, you must then add a certificate to the server. You can use either a device onboard token or a bearer token, and can use the Craft console or a CSR enrollment script. We recommend that you use a device onboard token.

Follow these steps to add a certificate to a server using a bearer token and the Craft console.

Step 1 Start the Craft console.

a) Use SSH to connect to the management IP address of the server as admin and enter the admin password when prompted.

```
$ ssh admin@management_ip_address
```

b) Run an interactive terminal session inside the cisco_hf_agent Docker container.

```
$ docker exec -ti cisco_hf_agent bash
```

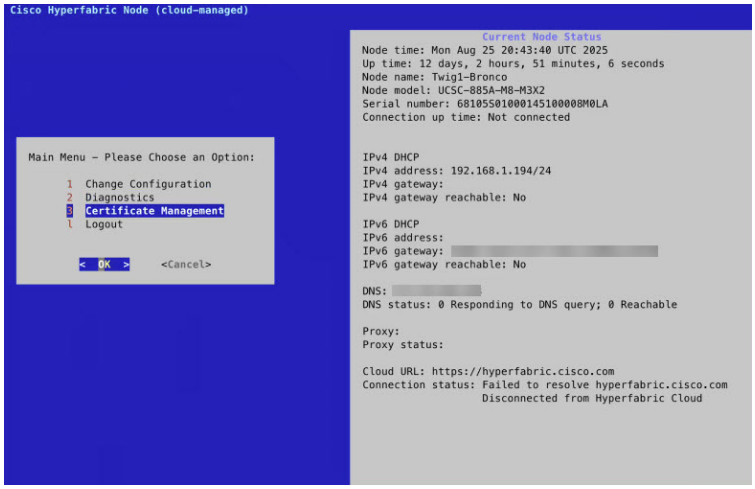
c) Run `main.sh`.

```
# /opt/cisco/bin/tortuga/craft-console/main.sh
```

Step 2 Verify that a certificate does not exist.

a) From the **Main Menu**, select **Certificate Management**.

Figure 41: Craft console main menu



b) Select **Check Certificate Status** to verify that the certificate now exists.

The console should display this text:

```
Certificate does not exist; please generate a CSR, and import certificate.

                Failed to read TPM keys from nvram cannot execute /bin/bash: exit status 3
                Press any key to continue
```

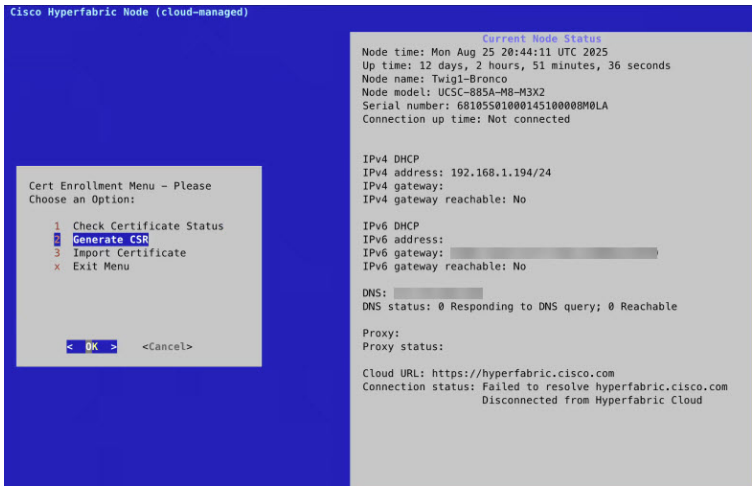
c) Press any key.

The console returns to the **Cert Enrollment Menu**.

Step 3 Generate a certificate.

a) Select **Generate CSR**.

Figure 42: Craft console cert enrollment menu



After a brief delay, the console displays the generated certificate.

b) Select the text of the certificate, including the "BEGIN CERTIFICATE REQUEST" and "END CERTIFICATE REQUEST" lines, and copy the text.

This is an example certificate:

```
-----BEGIN CERTIFICATE-----
                pG8jZQ3XkDYTq5mF2sH9rVb1Lw6uJcN0eA7vXoBzS4iTnPfRxWlE3yKjG6hCvM1q
                LdF2tUzN8mB7pXoS5cJaQkV9gWlR6yEpTzU4iFbM0hCxN2vL3sYwQ7uJdK8nV1rX
                zW4bLmP5tGqS9jAfE6hR0xUoI2vYlZ3cNwQ8pJ/K1mV7rLxT4eFiP9bHzW6uJqD5
                yCkN0tRvL3sXwA8mBpS2fJgQ7hUeK4nVzT1wL5uJrD9cHnM2pXqV.yLkS6tFbE0z
                W8vJqN7xLrC5mUdG1hPzS4tWfKjQ6VbY4kN2XcP9mEzA3wUoF7sTnH1lD8rSgL0p
                JvB5yMZxR2uKfW6qCjI3hOaT7eVnS9lQwG5dXpF1cJzK8bNvU4oRsL0yMhP3kTqV
                6xWjF2nZcD7lBaE9mPtG1vRzS4wXoN5jQkH8uCzI2pLrY7sTnV0fJbA3dMpE6qQx
                W9kJcO1hLwS8vPzF5tNfU2yGmB7rXoL0nD4uKjT3cVaE6pQwS9hM1/FzL7bJvN8x
                WfR2lCoG5tUeK0n/jT4yVqS6dLwA1xMzB9uFrN2vCjH3kUoS5tWpQ8rLzEZ3qBf8
                uH6rP.j1xW5Lk2VnRzG9aD7tSmX0C4wEoTgJvN5bF1yKsU2hQlM8pO6dXcS3nLrA
                0zVjI7uG4kP9yTqWbM6cF2vJoS1hE8xLwR3pNfK5dU0mJzA7tYgQnV4lB9wXfC2s
                H5rP6uKjT1oD8qS0vN3yLpI7bZxR4nWfM6tU2cG5aH9jEoT0pV3lDsK8wQmJ1yLz
                S7uN4xFbR6qC2nP5oH0kVtG3dJ8yM1vWzA7rQjL4sXnE9pU6oS2hT5cKfNkF8vN2
                yR7jQ1tXpZ6wLsH3oD9uB5eV4aGmT0lJcW8zM2/SnY7qK1pXrI6hV3dU9bS5oL4j
                A0nGvT8wE2cFzM7tQpN1yJkR6lW3sD9uB5a/4xH0mJfC8oS2nY7qK1vXrT6pE3zI
                9bL5jW4hU0lGmA8dF2yR7tQ.N1oKcV6wS3nJ9uB5aL4xG0mHfT8vN2yR7jQ1tXpZ
                6wLsH3oD9uB5eV4aGmT0lJcW8zM2fSnY7qK1pXrI6hV3dU9bS5oL4jA0nGvT8wE2
                cFzM7tQpN1yJkR6lW3sDpG8jZQ3XkDYTq5mF2sH9rVb1Lw6uJcN0eA7vXoBzS4iT
                nPfRxWlE3yKjG6hCvM1qLdF2tUzN8mB7pXoS5cJaQkV9gWlR6yEpTzU4iFbM0hCx
                N2vL3sYwQ7uJdK8nV1rXzW4bLmP5tGqS9jAfE6hR0xUoI2vYlZ3cNwQ8pJsK1mV7
                rLxT4eFiP9bHzW6uJqD5yCkN0tRvL3sXwA8mBpS2fJgQ7hUeK4nVzT1wL5uJrD9c
                HnM2pXq/3yLkS6tFbE0zW8vJqN7xLrC5mUdG1hPzS4tW.KjQ6VbY4kN2XcP9mEzA
                3wUoF7sTnH1lD8rSgL0pJvB5yMZxR2uKfW6qCjI3hOaT7eVnS9lQwG5dXpF1cJzK
                8bNvU4oRsL0yMhP3kTqV6xWjF2nZcD7lBaE9mPtG/vRzS4wXoN5jQkH8uCzI2pLr
                Y7sTnV0fJbA3dMpE6qQxW9kJcO1hLwS8vPzF5tNfU2yGmB7rXoL0nD4uKjT3cVaE
                6pQwS9hM1yFzL7bJvN8xWfR2lCoG5tUeK0nPjT4yVqS6dLwA1xMzB9uFrN2vCjH3
                kUoS5tWpQ8rLzEZ3qBf8uH6rPYj1xW5Lk2VnRzG9aD7tSmX0C4wEoTgJvN5bF1yK
                sU2hQlM8pO6dXcS3nLrA0zVjI7uG4kP9yTqWbM6cF2vJoS1hE8xLwR3pNfK5dU0m
                JzA7tYgQnV4lB9wXfC2sH5rP6uKjT1oD8qS0vN3yL/I7bZxR4nWfM6tU2cG5aH.j
                EoT0pV3lDsK8wQmJ1yLzS7uN4xFbR6qC2nP5oH0kVtG3dJ8yM1vWzA7rQjL4sXnE
                9pU6oS2hT5cKfNkF8vN2yR7jQ1tXpZ6wLsH3oD9uB5eV4aGmT0lJcW8zM2fSnY7q
                K1pXrI6hV3dU9bS5oL4jA0nGvT8wE2cFzM7tQpN1yJkR6lW3sD9/B5aV4xH0mJfC
                8oS2nY7qK1vXrT6pE3zI9bL5jW4hU0lGmA8dF2yR7tQpN1oKcV6wS3nJ9uB5aL4x
                W3.1X39e83zEKT2yskG=
                -----END CERTIFICATE-----
```

c) In the Cisco Nexus Hyperfabric GUI, choose **Administration** > **Server certificates**.

Figure 43: New server Certificate Signing Request (CSR)



d) Paste the certificate in the left field, then select **Submit**.

The certificate now appears in the right field.

Figure 44: Issued a new server Certificate Signing Request (CSR)



e) Select **Copy** next to the right field.

Step 4 Import the certificate into the server.

a) In the Craft console, select **Import Certificate**.

The console displays this text:

```
Paste your certificate text below.
                When you are done, press Enter, then type Ctrl-D.
```

b) Paste the certificate, press Enter, then press Ctrl+D.

The process can take approximately one minute to complete, after which the console displays text similar to this example:

```
-----BEGIN CERTIFICATE-----
                pG8jZQ3XkDYTq5mF2sH9rVb1Lw6uJcN0eA7vXoBzS4iTnPfRxWlE3yKjG6hCvM1q
                LdF2tUzN8mB7pXoS5cJaQkV9gWlR6yEpTzU4iFbM0hCxN2vL3sYwQ7uJdK8nV1rX
                zW4bLmP5tGqS9jAfE6hR0xUoI2vYlZ3cNwQ8pJ/K1mV7rLxT4eFiP9bHzW6uJqD5
                yCkN0tRvL3sXwA8mBpS2fJgQ7hUeK4nVzT1wL5uJrD9cHnM2pXqV.yLkS6tFbE0z
                W8vJqN7xLrC5mUdG1hPzS4tWfKjQ6VbY4kN2XcP9mEzA3wUoF7sTnH1lD8rSgL0p
                JvB5yMZxR2uKfW6qCjI3hOaT7eVnS9lQwG5dXpF1cJzK8bNvU4oRsL0yMhP3kTqV
                6xWjF2nZcD7lBaE9mPtG1vRzS4wXoN5jQkH8uCzI2pLrY7sTnV0fJbA3dMpE6qQx
                W9kJcO1hLwS8vPzF5tNfU2yGmB7rXoL0nD4uKjT3cVaE6pQwS9hM1/FzL7bJvN8x
                WfR2lCoG5tUeK0n/jT4yVqS6dLwA1xMzB9uFrN2vCjH3kUoS5tWpQ8rLzEZ3qBf8
                uH6rP.j1xW5Lk2VnRzG9aD7tSmX0C4wEoTgJvN5bF1yKsU2hQlM8pO6dXcS3nLrA
                0zVjI7uG4kP9yTqWbM6cF2vJoS1hE8xLwR3pNfK5dU0mJzA7tYgQnV4lB9wXfC2s
                H5rP6uKjT1oD8qS0vN3yLpI7bZxR4nWfM6tU2cG5aH9jEoT0pV3lDsK8wQmJ1yLz
                S7uN4xFbR6qC2nP5oH0kVtG3dJ8yM1vWzA7rQjL4sXnE9pU6oS2hT5cKfNkF8vN2
                yR7jQ1tXpZ6wLsH3oD9uB5eV4aGmT0lJcW8zM2/SnY7qK1pXrI6hV3dU9bS5oL4j
                A0nGvT8wE2cFzM7tQpN1yJkR6lW3sD9uB5a/4xH0mJfC8oS2nY7qK1vXrT6pE3zI
                9bL5jW4hU0lGmA8dF2yR7tQ.N1oKcV6wS3nJ9uB5aL4xG0mHfT8vN2yR7jQ1tXpZ
                6wLsH3oD9uB5eV4aGmT0lJcW8zM2fSnY7qK1pXrI6hV3dU9bS5oL4jA0nGvT8wE2
                cFzM7tQpN1yJkR6lW3sDpG8jZQ3XkDYTq5mF2sH9rVb1Lw6uJcN0eA7vXoBzS4iT
                nPfRxWlE3yKjG6hCvM1qLdF2tUzN8mB7pXoS5cJaQkV9gWlR6yEpTzU4iFbM0hCx
                N2vL3sYwQ7uJdK8nV1rXzW4bLmP5tGqS9jAfE6hR0xUoI2vYlZ3cNwQ8pJsK1mV7
                rLxT4eFiP9bHzW6uJqD5yCkN0tRvL3sXwA8mBpS2fJgQ7hUeK4nVzT1wL5uJrD9c
                HnM2pXq/3yLkS6tFbE0zW8vJqN7xLrC5mUdG1hPzS4tW.KjQ6VbY4kN2XcP9mEzA
                3wUoF7sTnH1lD8rSgL0pJvB5yMZxR2uKfW6qCjI3hOaT7eVnS9lQwG5dXpF1cJzK
                8bNvU4oRsL0yMhP3kTqV6xWjF2nZcD7lBaE9mPtG/vRzS4wXoN5jQkH8uCzI2pLr
                Y7sTnV0fJbA3dMpE6qQxW9kJcO1hLwS8vPzF5tNfU2yGmB7rXoL0nD4uKjT3cVaE
                6pQwS9hM1yFzL7bJvN8xWfR2lCoG5tUeK0nPjT4yVqS6dLwA1xMzB9uFrN2vCjH3
                kUoS5tWpQ8rLzEZ3qBf8uH6rPYj1xW5Lk2VnRzG9aD7tSmX0C4wEoTgJvN5bF1yK
                sU2hQlM8pO6dXcS3nLrA0zVjI7uG4kP9yTqWbM6cF2vJoS1hE8xLwR3pNfK5dU0m
                JzA7tYgQnV4lB9wXfC2sH5rP6uKjT1oD8qS0vN3yL/I7bZxR4nWfM6tU2cG5aH.j
                EoT0pV3lDsK8wQmJ1yLzS7uN4xFbR6qC2nP5oH0kVtG3dJ8yM1vWzA7rQjL4sXnE
                9pU6oS2hT5cKfNkF8vN2yR7jQ1tXpZ6wLsH3oD9uB5eV4aGmT0lJcW8zM2fSnY7q
                K1pXrI6hV3dU9bS5oL4jA0nGvT8wE2cFzM7tQpN1yJkR6lW3sD9/B5aV4xH0mJfC
                8oS2nY7qK1vXrT6pE3zI9bL5jW4hU0lGmA8dF2yR7tQpN1oKcV6wS3nJ9uB5aL4x
                W3.1X39e83zEKT2yskG=
                -----END CERTIFICATE-----
```

```
Certificate imported successfully
drake: stopped
drake: started
Press any key to continue
```

c) Press any key.

The console returns to the **Cert Enrollment Menu**.

Step 5 Verify that the certificate exists.

a) Select **Check Certificate Status** to verify that the certificate now exists.

The console displays text similar to this example:

```
A valid certificate exists.
                name: 000b5043bca227a1bc8d3e72e05abc6d073e40341337ee67ab401a1db66db8df123
                Press any key to continue
```

b) Press any key.

The console returns to the **Cert Enrollment Menu**.

Step 6 If you have a cluster, repeat all of these substeps for each server in the cluster.

Add a certificate to a server using a bearer token and a CSR enrollment script

After you install the Cisco Nexus Hyperfabric agent on a C880 server, you must then add a certificate to the server. You can use either a device onboard token or a bearer token, and can use the Craft console or a CSR enrollment script. We recommend that you use a device onboard token.

Even though this procedure is for using a CSR enrollment script, you use the Craft console for some of the steps.

Follow these steps to add a certificate to a server using a bearer token and a CSR enrollment script.

Step 1 Generate a bearer token.

You must be an admin user to generate the token. For more information about bearer tokens, refer to Authentication on Cisco DevNet.

a) In the Nexus Hyperfabric GUI, select the user avatar (), then select **Generate tokens**.
b) Select **API bearer tokens**, then select **Add a bearer token**.
c) Fill out the fields, then select **Generate Token**.
d) Select **Copy** to copy the device token, then select **Done**.

Step 2 In the server's CLI, export the device onboard token.

a) If you are not already connected to the server, use SSH to connect to the server.
b) Export the device onboard token.

```
admin@cluster3-centos-vm4:~$ export DEVICE_ONBOARD_TOKEN="token"
```

Paste the device token that you copied in place of *token*.

Step 3 Use a CSR enrollment script to add a certificate to the server.

a) Download one of the CSR enrollment scripts, depending on your setup.

If you have a standalone server that has Ubuntu or Red Hat, download the `server-csr-enroll.sh` script from here:

https://cdn.hyperfabric.cisco.com/server/releases/agent/utils/server-csr-enroll.sh

If you have a Kubernetes or OpenShift cluster, download the `cluster-csr-enroll.sh` script from here:

https://cdn.hyperfabric.cisco.com/server/releases/agent/utils/cluster-csr-enroll.sh

b) Run the CSR enrollment script that you downloaded to sign a certificate signing request (CSR) for the server or all servers in the cluster, as appropriate.

Run one of these commands:

```
admin@cluster3-centos-vm4:~$ ./server-csr-enroll.sh [--force] [-p http://proxy.example.com:80]
admin@cluster3-centos-vm4:~$ ./cluster-csr-enroll.sh [--force] [-p http://proxy.example.com:80]
```

If the server or cluster already has a certificate, include **--force** to overwrite the certificate.

If you are using a proxy, include **-p**.

Step 4 Start the Craft console.

a) Use SSH to connect to the management IP address of the server as admin and enter the admin password when prompted.

```
$ ssh admin@management_ip_address
```

b) Run an interactive terminal session inside the cisco_hf_agent Docker container.

```
$ docker exec -ti cisco_hf_agent bash
```

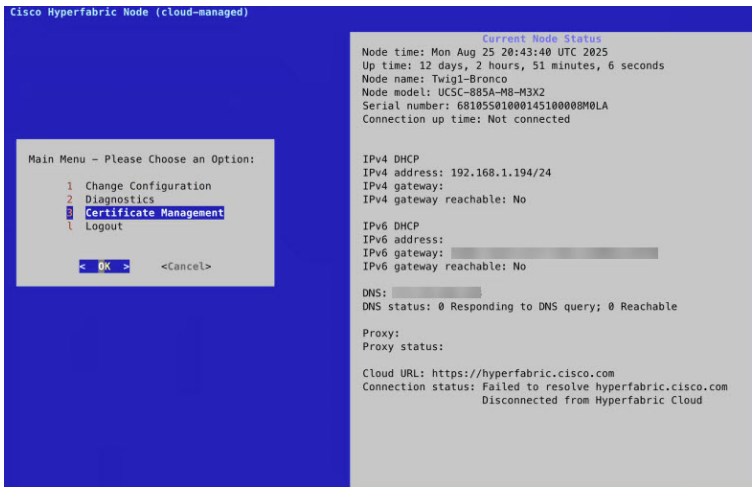
c) Run `main.sh`.

```
# /opt/cisco/bin/tortuga/craft-console/main.sh
```

Step 5 Verify that the certificate exists.

a) From the **Main Menu**, select **Certificate Management**.

Figure 45: Craft console main menu



b) Select **Check Certificate Status** to verify that the certificate now exists.

The console displays text similar to this example:

```
A valid certificate exists.
                  name: 000b5043bca227a1bc8d3e72e05abc6d073e40341337ee67ab401a1db66db8df123
                  Press any key to continue
```

c) Press any key.

The console returns to the **Cert Enrollment Menu**.

Step 6 If you have a cluster, repeat this procedure for each server in the cluster.

Considerations for installing the Cisco Nexus Hyperfabric agent

These considerations apply for installing the Cisco Nexus Hyperfabric agent:

- You can install the agent on a server only after you have installed OpenShift CoreOS as the host operating system on that server.
- Connect a server on which you installed the Nexus Hyperfabric agent to the Nexus Hyperfabric cloud controller for configuration management and telemetry. The agent can use any port, depending on the server routing table that is available in its kernel, regardless of how ports are used. For example, the ports can be used for east-west traffic, north-south traffic, and management. If the communication must be through management interfaces, you must configure a routing table accordingly. If the agent cannot reach the cloud controller directly, the agent relies on an exit node in the fabric to reach the controller.

Cisco Nexus Hyperfabric server agents and ISO images

This table provides the download location of the agents and ISO images that you need to use servers with Cisco Nexus Hyperfabric.

Table 6: Server agents and ISO images

Agent or ISO image	Download location
UCS agent installer	Download the appropriate package. RPM: https://cdn.hyperfabric.cisco.com/server/releases/agent/latest/ucs_agent_installer_latest.rpm DEB: https://cdn.hyperfabric.cisco.com/server/releases/agent/latest/ucs_agent_installer_latest.deb
Nexus Hyperfabric deployment agent For installation information, refer to Install the Cisco Nexus Hyperfabric agent on a server, on page 46.	https://cdn.hyperfabric.cisco.com/server/releases/agent/latest/hf_agent_deploy.yaml

Claiming devices

Upon delivery, a new device is not associated with any owner or organization. Before deploying the device, you must claim it through Cisco Nexus Hyperfabric and add it to the organization that owns the fabric in which the device will be deployed.

Note

In some cases, your devices might already be claimed for your organization before delivery. If so, you can skip the claim procedures.

To claim the device, you must have physical and administrative access to the device.

Limitations for claiming devices

These limitations apply for claiming devices:

- You must claim the first device of your fabric using a claim code.
- When claiming devices without using claim codes, you can claim only switches that Cisco Nexus Hyperfabric has discovered. The discovery process is limited to switches that are directly connected to an already-claimed Nexus Hyperfabric switch. Only these directly connected devices will appear in the discovered list and be available for claiming.

Claim a switch using a claim code

You must claim the first switch of your fabric using a claim code. You can also claim subsequent switches using claim codes or you can claim discovered switches without using the claim codes. For information about claiming a switch without using a claim code, refer to Claim a discovered device without using a claim code, on page 58.

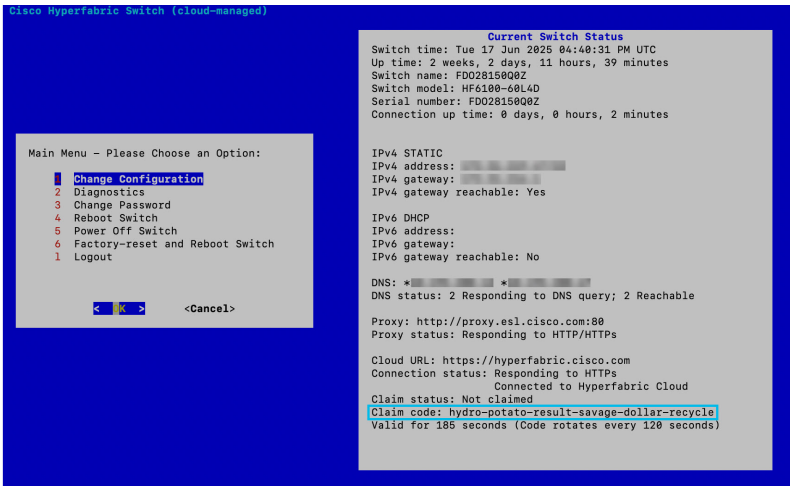
Follow these steps to claim a switch using a claim code.

Step 1 Use the serial console to log in to the switch and get the claim code that is shown in the serial console menu.

If you want to claim multiple switches, then get the code from every switch that you want to claim.

The claim code that displays in the console changes every two minutes, although a code is valid for four minutes.

Figure 46: Example of a claim code



Step 2 In the Cisco Nexus Hyperfabric GUI, choose **Inventory** > **Device Inventory**.

Step 3 Select **Claim devices**.

Step 4 Enter the claim code, then select **Claim**.

If you want to claim multiple switches, enter the codes separated by commas or new lines.

The switch then becomes registered to the current organization as claimed. As the other fabric switches are installed and connected, unclaimed switches connected to the first switch will be discovered.

Claim a discovered device without using a claim code

You must claim the first device of your fabric using a claim code, which means you cannot use this procedure for the first device. However, you can claim additional discovered devices without using the claim codes. For information about claiming the first device or a subsequent device using a claim code, refer to Claim a switch using a claim code, on page 58.

You cannot use a baseboard management controller (BMC) port to claim a server.

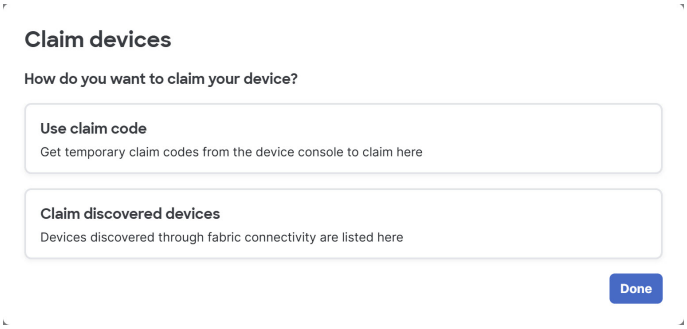
Follow these steps to claim a discovered device without using a claim code.

Step 1 Select **Inventory** > **Device Inventory**.

Step 2 Click **Claim devices**.

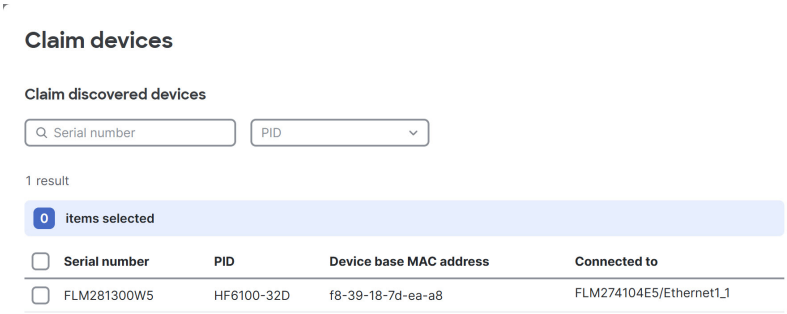
Step 3 Click **Claim discovered devices**.

Figure 47: Claim devices



Step 4 Select one or more switches, then click **Claim**.

Figure 48: Claim devices



The devices then become registered to the current organization as claimed.

Optionally, you can find the device with a specific serial number using the **Serial number** field and filter for switches with the product ID that you select from the **PID** drop-down.

Bind the devices

In Cisco Nexus Hyperfabric, you must assign each physical device ("bind") to its logical position (node) in the fabric blueprint. When a device has been bound to a node in the blueprint, it can download its configuration from Cisco Nexus Hyperfabric.

A device that has been claimed but not yet bound to an organization's fabric is considered to be bound to the parking lot—a "null" fabric—until it is bound to an actual fabric.

Note

When you bind the first device to the fabric, Cisco Nexus Hyperfabric sets the fabric version to the software version running on that device. Because devices in the parking lot are automatically upgraded to the latest version, the fabric version is set to the latest release.

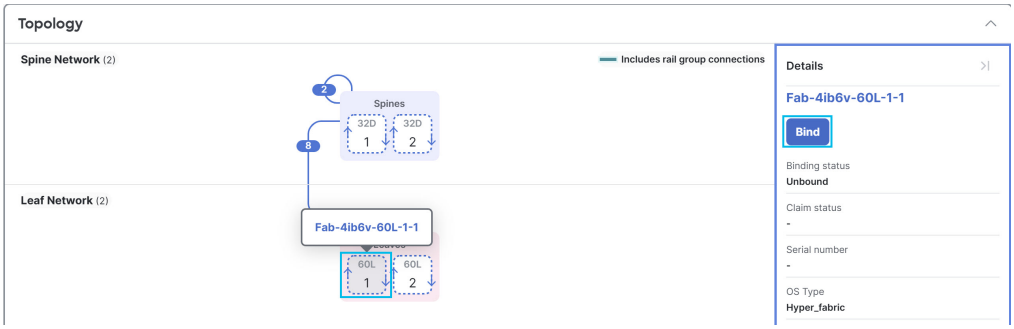
Follow these steps to bind the devices.

Step 1 Select **Fabrics**, then select the fabric that contains the devices.

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

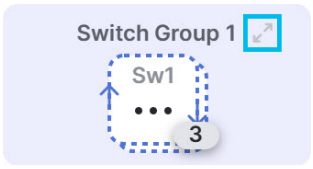
Step 2 In the **Topology** area, select the position of the device that you want to bind, then select **Bind** in the device information pane.

Figure 49: Topology



If a group has more than two devices, you must expand the group before you can select the position.

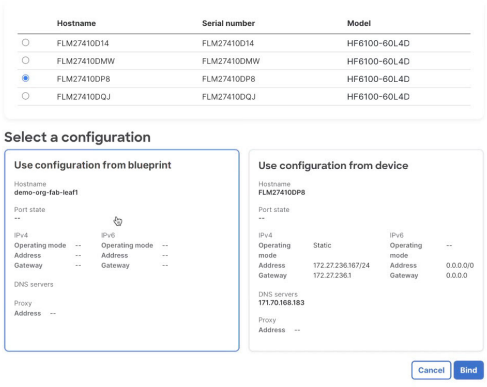
Figure 50: Expand a switch group



Step 3 Select the physical device to assign to this position in the topology.

The resulting dialog displays two options for the device configuration.

Figure 51: Select a configuration



Step 4 Select the configuration.

- **Use configuration from blueprint**: uses the configuration that is defined in the blueprint. If you select this, ensure that your device can still connect to the cloud with the configuration defined in the blueprint. Refer to the "Configure the management interface" procedure in the *Cisco Nexus Hyperfabric — Configure Switches* document.
- **Use configuration from device**: uses the current device settings that were configured through the serial port. Cisco Nexus Hyperfabric updates the blueprint with the device's settings.

Step 5 To bind the device to this logical location in the fabric blueprint, select **Bind**.

When the binding operation for a device has completed successfully, the device icon in the topology area is blue.

Step 6 Repeat these steps to bind the remaining devices.

Step 7 Verify the cloud connectivity status of the devices.

a) Go back to the fabric page by selecting the name of the fabric near the top left of the page.

b) In the **Physical Topology** area, select **Cloud connectivity**.

On this page, these are some properties to observe:

- **Source IP:Port**: the source IP address and port.
- **Destination IP:Port**: the destination IP address and port.
- **Connection path**: the path through which the device is connected to the cloud. These are the possible values:
  - **Management port**: the path is through the management port of the switch.
  - ***switch_name***: the path is through the specified switch, which is an exit node.
  - ***port_ID***: the path is through the in-band management port of the switch. The port's ID is shown.
- **Certificate chain**: specifies whether the switch's transport layer security (TLS) certificate is valid or invalid. If the certificate is invalid, select **Invalid** to receive information about the issue.

Review switch details by scanning the QR code

The Cisco 6000 series switches have stickers that have a unique QR code. The QR code provides a launch URL that contains the serial number of the switch. You can use a mobile device to scan the code, then open the URL to go to the switch details page in the Cisco Nexus Hyperfabric GUI.

For more information about the QR code, refer to the Cisco Nexus Hyperfabric Series – QR Code document.

Follow these steps to review switch details by scanning the QR code.

Step 1 Use a mobile device to scan the QR code.

Step 2 Go to the URL that the code provides.

One of these results occurs:

- If you are not logged into the Cisco Nexus Hyperfabric GUI, you are redirected to the login page. After you log in, the process continues as if you scanned the QR code while you were already logged in. Refer to the other possible results.
- If you are logged into the Cisco Nexus Hyperfabric GUI and the switch is claimed by the organization that you are logged into, you access the GUI page that shows the details of the switch.

- If you are logged into the Cisco Nexus Hyperfabric GUI but the switch is unclaimed or the switch is claimed by an organization that you are not logged into, you cannot access the switch details.

Finish and commit the design

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

Follow these steps to finish and commit the design.

Step 1 Select **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**.

If you bound any devices as part of this push, the outline color of the device icons indicates the assertions for the devices. The devices should have a green outline, which indicates that there are no issues. A yellow outline indicates the device has an unexpected condition and there are issues detected, but the issues should not cause operation disruption. A red outline indicates the device has an unexpected condition and there are issues detected that can cause operation disruption; you should investigate the issues.

Figure 52: Fabric topology after pushing a configuration

