

Revised: June 10, 2025

# Cisco Nexus Hyperfabric — Configure Logical Networks

## Logical networks

A logical network is a virtual construct that allows for the creation of network segments that are independent of the underlying physical network infrastructure. Each logical network in a Cisco Nexus Hyperfabric is assigned a virtual network identifier (VNI) that is unique within the fabric.

Logical networks can be categorized into Layer 2 and Layer 3 networks, each serving different purposes and providing different functionalities.

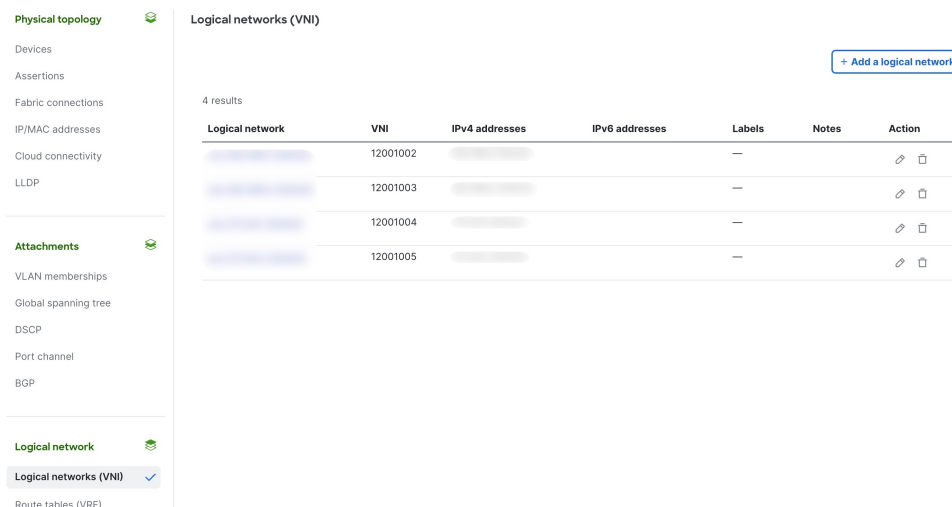
- Layer 2 logical networks, such as VLANs and VXLANs, allow devices to communicate as if they are on the same physical network, even if they are not. A Layer 2 logical network and its member ports are identified by a unique Layer 2 VNI.
- Layer 3 logical networks provide logical Layer 3 segmentation, allowing for routing and communication between different Layer 2 networks. A Layer 3 logical network and its member ports are identified by a unique Layer-3 VNI, which is associated with a unique tenant VRF (Virtual Routing and Forwarding).

## Create a logical network

A Layer 2 or Layer 3 logical network requires a logical network name. A Layer 3 logical network requires an IP address and a route table. If you are creating a Layer 3 logical network, you must select an existing route table or create one if it does not exist already. A default route table is created automatically when an IP interface such as Layer 3 logical network or routed interface is created. .

Follow these steps to create a Layer 2 logical network for intranetwork traffic or Layer 3 logical network for intranetwork or internetwork traffic.

- Step 1** Select **Fabrics**, then click the fabric to configure.
- Step 2** If the fabric is not in the edit mode, click **Switch to edit mode**.
- Step 3** In the **Logical network** area, click **Logical networks (VNI)**.



The screenshot shows the Cisco Nexus Hyperfabric configuration interface. On the left, there is a sidebar with a tree view containing sections: Physical topology, Attachments, and Logical network. The Logical network section is expanded, showing 'Logical networks (VNI)' with a checkmark and 'Route tables (VRF)'. The main area displays 'Logical networks (VNI)' with a '+ Add a logical network' button. Below the button, it says '4 results' and shows a table with columns: Logical network, VNI, IPv4 addresses, IPv6 addresses, Labels, Notes, and Action. The table contains four rows of data, each with a blue bar representing the logical network name, a VNI value (12001002, 12001003, 12001004, 12001005), and empty fields for addresses, labels, and notes. Each row has edit and delete icons in the Action column.

Logical network	VNI	IPv4 addresses	IPv6 addresses	Labels	Notes	Action
[Redacted]	12001002	[Redacted]	[Redacted]	—		[Edit] [Delete]
[Redacted]	12001003	[Redacted]	[Redacted]	—		[Edit] [Delete]
[Redacted]	12001004	[Redacted]	[Redacted]	—		[Edit] [Delete]
[Redacted]	12001005	[Redacted]	[Redacted]	—		[Edit] [Delete]

- Step 4** Click + Add a logical network.

**Add logical network**

Logical networks are useful for intra network and inter network connection. Let's create one here.

Logical network name \*  VNI   
A VNI will be autogenerated if left blank

Select routing table

**IP Addresses**  
 Use a comma (,) to separate multiple addresses.

IPv4 addresses  IPv6 addresses

Notes

Labels

**Step 5** Follow these substeps to create a logical network.

- Enter a logical network name.
- By default, a virtual network interface (VNI) number will be generated automatically if you leave this field blank. Enter a VNI number to manually assign one.
- For Layer 3 logical network, select a routing table from the drop-down list.
- To create a new route table, select **Create new route table**. Enter the route table name and L3VNI number greater than 4096. The L3VNI number is optional.
- For Layer 3 logical network, enter IPv4 and/or IPv6 IP addresses. Use a comma to separate multiple addresses.
- Optionally, you can add notes and labels for the logical network. To add a label, click + **Add**, enter the text and press **Enter**.
- Click **Save**.

**Step 6** The logical network (VNI) is displayed in the Logical network table. For a Layer 2 logical network, you can view the name and VNI number. For a Layer 3 logical network you can view the name, VNI number, and IPv4 and/or IPv6 IP addresses.

Physical topology	Logical networks (VNI)
Devices	<a href="#">+ Add a logical network</a>
Assertions	3 results
Fabric connections	
IP/MAC addresses	
Cloud connectivity	
LLDP	
Attachments	
VLAN memberships	
Global spanning tree	
DSCP	
Port channel	
BGP	
Logical network	
Logical networks (VNI)	

Logical network	VNI	IPv4 addresses	IPv6 addresses	Labels	Notes	Action
net1	12001001			--		<a href="#">edit</a> <a href="#">delete</a>
net2	12001002			--		<a href="#">edit</a> <a href="#">delete</a>
net3	12001003			--		<a href="#">edit</a> <a href="#">delete</a>

For Layer 3 logical network, IP address is also displayed in the SVI table at the VRF level. Follow the following steps to view the SVI table.

- In the **Logical network** area, click **Route tables**.
- Select a route table (VRF).

- c) In the Configurations area, click **SVIs**.  
The IP address is displayed in the SVI table.

## Add a VLAN member to a logical network

A logical network is deployed by adding a VLAN member. A VLAN member maps a logical network to switch(s), port interface(s), and VLAN (tagged or untagged). After you add a VLAN member, the traffic coming from the switch, interface, and VLAN will be mapped to the logical network.

Follow these steps to add a VLAN member to an existing logical network.

- Step 1** Select **Fabrics**, then click the fabric to configure.
- Step 2** If the fabric is not in the edit mode, click **Switch to edit mode**.
- Step 3** In the **Attachments** area, click **VLAN memberships**.
- Step 4** Click + **Add a VLAN membership**.

The screenshot shows the 'Add VLAN membership' form. It includes a 'Logical network' dropdown menu, a 'VNI' field, radio buttons for 'Ports' (selected) and 'Port channels', a 'Switches' dropdown menu, a 'Port interface' dropdown menu, and a 'VLAN' section with a selected 'VLAN tag' (with a text input 'Must be between 2-3600') and an 'Untagged' option. At the bottom right are 'Cancel' and 'Save' buttons.

- Step 5** Select an existing logical network from the **Logical network** drop-down list.  
The VNI number for the selected network appears under **VNI**.
- Step 6** Select an individual switch or select \* to select all switches from the **Select switch** drop-down list.
- Step 7** Select an individual port or port channel from the **Port interface** drop down list.
- Step 8** Select **VLAN tag** and enter a VLAN number between 2 and 3600, or select **Untagged**.

Nexus Hyperfabric allows you to configure a port as both native (untagged) as well as tagged (VLAN IDs 2-3600). When traffic with the configured tag arrives on the selected interface on the selected switch, it will be mapped to this logical network.

- VLAN 3601 and above are reserved by the system.
- One logical network is associated with one VLAN. Different VLAN tags cannot be mapped to the same logical network.
- If you select different switches, VLAN tagged and untagged can be mapped to the same logical network.
- If you select different port interfaces, VLAN untagged can be mapped to the different logical networks.

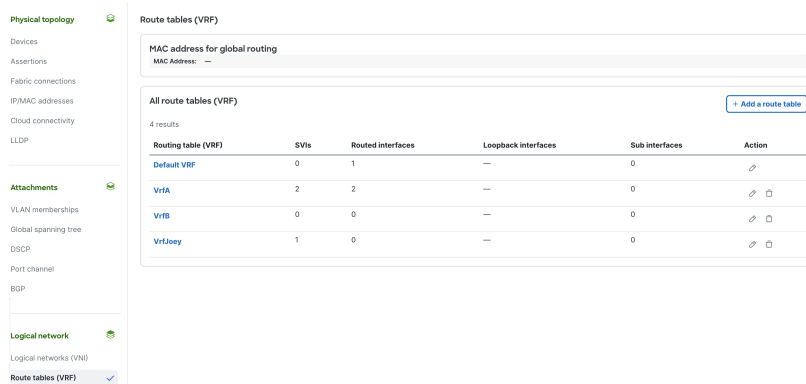
- Step 9** Click **Save**.  
The VLAN members are displayed in the VLAN membership table.

# Create a route table (VRF)

A route table is a data table stored in a Layer 3 switch that lists the routes to particular network destinations. The route table contains information about the topology of the network immediately around it. It is used to determine the best path for forwarding packets to their destination.

Virtual Route Forwarding (VRF) is a technology that allows having more than one route table on a single router or Layer 3 switch. In Nexus Hyperfabric, you can create multiple route tables on a single fabric.

- Step 1** Select **Fabrics**, then click the fabric to configure.
- Step 2** If the fabric is not in the edit mode, click **Switch to edit mode**.
- Step 3** In the **Logical network** area, click **Route tables (VRF)**.



- Step 4** Click + **Add a route table**.

Add route table

Routing tables will generate VRFs to help you with virtual routing.

Route table name (VRF) \*

Vrf

VNI

An VNI will be autogenerated if left blank

Notes

Labels

+ Add

Cancel Add

- a) Enter a name for the route table.
- b) By default, a virtual network interface (VNI) number will be generated automatically if you leave this field blank. Enter a VNI number to manually assign one. VNI number should be greater than 4096.  
If you leave this field blank, a number will be generated automatically.
- c) Optionally, you can add tags to this route table. To add a label, click + **Add**, enter the text and press **Enter**.

- Step 5** Click **Add**.

The route table is displayed in the All route tables (VRF) area.

# Add a static route

Routers forward packets using route information from route table entries that you manually configure or from the route information that is calculated using dynamic routing algorithms.

Static routes, which define explicit paths between multiple one or more next hops, must be manually configured when network changes occur. Static routes use less bandwidth than dynamic routes. No CPU cycles are used to calculate and analyze routing updates.

Follow these steps to add a static route.

**Step 1** Select **Fabrics**, then click the fabric to configure.

**Step 2** In the **Logical network** area, click **Route tables**.

**Step 3** Select a route table (VRF).

**Step 4** In the Configurations area, click **Routes**.

IP prefix/mask	Route type	Route distance	Next-hop IP address	Next-hop MAC address	Status	Next-hop location	Action
	Directly connected	0	---	---	Up	TME-fab2-HF8100-80L4D-leaf3 (Vlan1011)	
	Directly connected	0	---	---	Up	TME-fab2-HF8100-80L4D-leaf3 (Vlan1012)	
	Directly connected	0	---	---	Up	TME-fab2-HF8100-80L4D-leaf1 (Vlan1011)	
	Directly connected	0	---	---	Up	TME-fab2-HF8100-80L4D-leaf1 (Vlan1012)	
	Directly connected	0	---	---	Up	TME-fab2-HF8100-80L4D-leaf2 (Vlan1011)	
	Directly connected	0	---	---	Up	TME-fab2-HF8100-80L4D-leaf2 (Vlan1012)	

**Step 5** Click **Add a static route**.

**Add a static route**

IP prefix/mask \*

IPv4 or IPv6

☒ Next hop address ☐ Discard

IPv4 or IPv6

Traffic sent to the destination IP prefix/mask will be discarded.

Set route distance (1-255)

Cancel Add

- Enter an IPv4 or IPv6 prefix/mask.
- Select **Next hop address** and enter the IP address.
- Select **Discard** to discard the traffic sent to the destination IP prefix/mask.
- By default, the route distance is set to 1.
- Click **Add**.

The static routes are displayed in the Routes area.

Configurations

- Routes
- SVIs
- Routed interfaces
- Loopback interfaces
- BGP peers
- DHCP relay profiles

Routes

+ Add a static route

IPv4 routes 14 IPv6 routes 8

IP prefix/mask	Route type	Route distance	Next-hop IP address	Next-hop MAC address	Status	Next-hop location	Action
	Static	1		00-81-c4-14-9c-bf	Up	TME-fab1-HF6100-60L4D-leaf1 (Ethernet1_4)	
	Directly connected	0		—	Up	TME-fab1-HF6100-60L4D-leaf1 (Ethernet1_4)	
	BGP	20		ec-19-2e-c0-4c-00	Up	TME-fab1-HF6100-60L4D-leaf1 (Vlan3801)	
	Directly connected	0		—	Up	TME-fab1-HF6100-60L4D-leaf1 (Vlan1011)	
	BGP	20		ec-19-2e-c0-4c-00	Up	TME-fab1-HF6100-60L4D-leaf1 (Vlan3801)	
	Directly connected	0		—	Up	TME-fab1-HF6100-60L4D-leaf1 (Vlan1012)	
	BGP	20		ec-19-2e-c0-4c-00	Up	TME-fab1-HF6100-60L4D-leaf1 (Vlan3801)	
	BGP	20		ec-19-2e-c0-4c-00	Up	TME-fab1-HF6100-60L4D-leaf1 (Vlan3801)	
	BGP	20		ec-19-2e-c0-68-00	Up	TME-fab1-HF6100-60L4D-leaf2 (Vlan3801)	

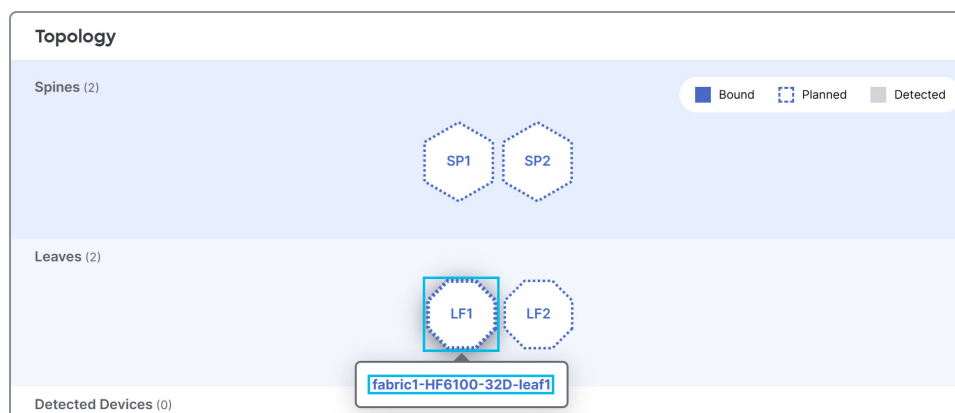
## Add a routed interface

When you configure a routed port (interface) on a switch, you can choose 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. Alternatively, if you leave VLAN tagging disabled, you can choose a route table (VRF) and specify IPv4 addresses or IPv6 addresses for the routed interface.

This procedure configures a routed port at the switch level, but you can also configure a routed port at the fabric level.

Follow these steps to configure a routed port.

- Step 1** Select **Fabrics**, then click the fabric that contains the switch.
- Step 2** If the fabric is not in the edit mode, click **Switch to edit mode**.
- Step 3** In the Topology area, click the switch position that you want to configure, then click the switch name.



- Step 4** In the **Configure** area, click **Port configurations**.
- The **Port configurations** table lists all ports of the switch.
- Step 5** In the **Action** column, click (ⓘ) for the port that you want to configure.
- Step 6** In the **Port configuration for switch\_name**, select **Routed** as the port role.
- Step 7** To use VLAN tagging, enable **Enable VLAN tagging**.
- The **SUB-INTERFACES** configuration area appears. Perform these substeps for each sub-interface.
- a) Click **Add a sub interface**.

- b) Enter the VLAN tag.



Do not assign the same VLAN tag to multiple sub-interfaces on the same physical interface.

**Note**

- c) Select a VRF instance from the **Select VRF** drop-down list.  
d) Enter an IPv4 address and mask, an IPv6 address, or both.

**Step 8**

If you do not want to use VLAN tagging, disable **Enable VLAN tagging** and perform these substeps.

- a) Select a route table (VRF) from the drop-down list.  
b) Enter one or more IPv4 or IPv6 addresses and masks.

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

**Step 9**

For **Admin state**, select the desired administrative state.

**Step 10**

Click **Save**.

## Create a DHCP relay profile

A DHCP relay allows DHCP communication between hosts and remote DHCP servers that are not on the same network. When a host sends a DHCP broadcast for an IP address, the DHCP relay agent forwards the request to the subnet of the remote DHCP server. DHCP server dynamically assigns IP addresses.

Nexus Hyperfabric relays the request from the host to the DHCP server and relays the offer from the DHCP server to the host. To allow for this communication, you have to create a DHCP relay profile in Nexus Hyperfabric.

The DHCP relay supports:

- only IPv4 address family
- DHCP relay is applicable to one or more logical networks with SVIs
- DHCP server in the same VRF as the relay
- DHCP server can be inside or outside of the fabric



**Note**

---

Because a DHCP relay profile is defined within a route table (VRF), a route table must exist before you can create a DHCP relay profile.

---

Follow these steps to configure a DHCP relay profile.

**Step 1**

Select **Fabrics**, then click the fabric to configure.

**Step 2**

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

**Step 3**

In the **Logical network** area, click **Route tables**.

**Step 4**


In the list of route tables (VRFs), click the route table name where you will add a DHCP relay profile.

**Step 5**

In the Configurations area, click **DHCP profiles**.

If DHCP profiles have already been created in this VRF, a list of existing profiles is displayed.

**Step 6**

In the Loopback IPv4 range area, click  to enter the IPv4 range in the CIDR format and click **Save**.

This is the loopback IP address range for Hyperfabric switches that will forward DHCP requests. A Unique loopback IP address is automatically allocated to Hyperfabric switches that have the logical network with the DHCP profile associated. The loopback IP addresses should be reachable from the DHCP server. For example, if the DHCP server is outside the fabric behind an external router, then the loopback IP address range needs to be advertised to the external network. Otherwise, the return traffic from the DHCP server will be dropped before coming back to the fabric.



If you use BGP to advertise fabric subnets to the external network, make sure that the loopback IP addresses are permitted in your BGP export policy.

#### Note

### Step 7 Click + Add a DHCP relay profile.

- Enter a name for the new DHCP relay profile.
- Enter the DHCP server IPv4 addresses.
- Select the logical network **VLANs** from the drop-down list.
- Click **Add**.

## Associate a Layer 3 logical network to an existing DHCP relay profile

The switch virtual interface (SVI) represents a logical interface between the bridging function and the routing function of a VLAN in the device. SVI can have members that are physical ports or port channels. The SVI logical interface is associated with VLANs, and the VLANs have port membership.

In Nexus Hyperfabric associating a Layer 3 logical network to a DHCP relay profile to allows the host to talk to DHCP server using SVI.



#### Note

You must create DHCP relay profile before associating with a Layer 3 logical network.

Follow these steps to associate a Layer 3 logical network to an existing DHCP relay profile.

- Step 1** Select **Fabrics**, then click the fabric to configure.
- Step 2** If the fabric is not in the edit mode, click **Switch to edit mode**.
- Step 3** In the **Logical network** area, click **Route tables**.
- Step 4** Select a route table (VRF).
- Step 5** In the **Configurations** area, click **SVIs**.
- Step 6** Click **Add a SVI**.



**Edit SVI**

Logical network

IPv4 addresses

IPv6 addresses


DHCP relay

DHCP relay profile

IP relay

☐ IPv4 relay enabled



Cancel Add

Optionally, in the **Action** column, click  for the logical network that you want to edit for an SVI.

- a) Select a Layer 3 logical network from the drop-down list.
- b) Enter an IPv4 address.
- c) Select the DHCP relay profile from the drop-down list.
- d) Enable IPv4 relay.
- e) Click **Add**.



## Manage logical networks

Follow these steps to edit or delete a logical network.

- Step 1** Select **Fabrics**, then click the fabric to configure.
- Step 2** If the fabric is not in the edit mode, click **Switch to edit mode**.
- Step 3** In the **Logical network** area, click **Logical networks (VNI)**.  
The **Logical networks (VNI)** table lists all the logical networks of the fabric.
- Step 4** In the **Action** column, click  for the logical network that you want to edit.
- Step 5** Edit the configuration and click **Save**.
- Step 6** In the **Action** column, click  to delete the logical network.



## Manage VLAN members

Follow these steps to edit or delete a VLAN member.

- Step 1** Select **Fabrics**, then click the fabric to configure.
- Step 2** If the fabric is not in the edit mode, click **Switch to edit mode**.
- Step 3** In the **Logical network** area, click **Logical networks (VNI)**.
- Step 4** Select a logical network and in the **Monitor** area, click **VLAN membership**.  
The **VLAN membership** table lists all the VLAN members.
- Step 5** In the **Action** column, click  for the VLAN members that you want to edit.
- Step 6** Edit the configuration and click **Save**.
- Step 7** In the **Action** column, click  to delete the VLAN member.


## Manage route tables

Follow these steps to edit or delete a route table (VRF).

- Step 1** Select **Fabrics**, then click the fabric to configure.
- Step 2** If the fabric is not in the edit mode, click **Switch to edit mode**.
- Step 3** In the **Logical network** area, click **Route tables (VRF)**.  
The **All route tables (VRF)** table lists all the route tables.
- Step 4** In the **Action** column, click  for the route table that you want to edit.
- Step 5** Edit the configuration and click **Save**.
- Step 6** In the **Action** column, click  to delete the route table.



## Manage routed interfaces

Follow these steps to edit a routed interface.

- Step 1** Select **Fabrics**, then click the fabric to configure.
- Step 2** If the fabric is not in the edit mode, click **Switch to edit mode**.
- Step 3** In the **Logical network** area, click **Route tables (VRF)**.
- Step 4** Select a route table and in the **Configurations** area, click **Routed interfaces**.
- Step 5** In the **Action** column, click  for the routed interface that you want to edit.
- Step 6** Edit the configuration and click **Save**.


## Manage DHCP relay profiles

Follow these steps to edit or delete a DHCP relay profile.

- Step 1** Select **Fabrics**, then click the fabric to configure.
- Step 2** If the fabric is not in the edit mode, click **Switch to edit mode**.
- Step 3** In the **Logical network** area, click **Route tables (VRF)**.
- Step 4** Select a route table and in the **Configurations** area, click **DHCP relay profiles**.
- Step 5** In the **Action** column, click  for the DHCP relay profile that you want to edit.
- Step 6** Edit the configuration and click **Save**.
- Step 7** In the **Action** column, click  to delete the DHCP relay profile.

## Manage SVIs

Follow these steps to edit a SVI.

- Step 1** Select **Fabrics**, then click the fabric to configure.
- Step 2** If the fabric is not in the edit mode, click **Switch to edit mode**.
- Step 3** In the **Logical network** area, click **Route tables (VRF)**.
- Step 4** Select a route table and in the **Configurations** area, click **SVIs**.
- Step 5** In the **Action** column, click  for the logical network that you want to edit.

**Step 6** Edit the configuration and click **Save**.

## Manage anycast gateway MAC addresses

The MAC address is the fabric wide value used by all of switch virtual interfaces (SVIs) owned by the fabric.

Follow these steps to manage anycast gateway MAC addresses.

- Step 1** Select **Fabrics**, then click the fabric to configure.
- Step 2** If the fabric is not in the edit mode, click **Switch to edit mode**.
- Step 3** In the **Logical network** area, click **Logical networks (VNI)**.
- Step 4** Next to **Anycast gateway MAC address**, click  $\emptyset$ .
- Step 5** Enter a new MAC address and click **Save**.

### Limitations for anycast gateway MAC addresses

These limitations apply for anycast gateway MAC addresses:

- The MAC address must not be a broadcast or multicast address.
- The MAC address must not be from a reserved range. These are the reserved MAC address ranges:
  - ae-00
  - ae-10
  - de-ad

## Spanning tree protocol

Spanning tree protocol (STP) prevents loops from being formed when switches or bridges are interconnected using multiple paths. Spanning tree protocol implements the 802.1D IEEE algorithm by exchanging Bridge Protocol Data Unit (BPDU) messages with other switches to detect loops, and then removes the loop by shutting down selected bridge interfaces. This algorithm guarantees that there is one and only one active path between two network devices.

By default, global STP is disabled on Nexus Hyperfabric. If global STP is disabled, host ports will drop BPDU messages. If global STP is enabled, Nexus Hyperfabric uses per-VLAN spanning tree plus (PVST+) and uses the same bridge ID for all the Nexus Hyperfabric switches.

### STP recommendations

Follow these recommendations if you enabled global STP:

- Nexus Hyperfabric is the root bridge.
- Enable root guard on all host ports.
- Enable STP on all host ports with STP PortFast enabled and BPDU guard disabled.

Regardless of global STP status, non-host ports never transmit BPDU messages. As a result, BPDU messages are not transmitted to Nexus Hyperfabric overlay network. We recommend that you use Nexus Hyperfabric as the root bridge instead of relying on an external root bridge for loop prevention.

Since root guard supports per VLAN and auto-recovery, we recommend that you enable root guard to protect against an unexpected external root.

Even with STP PortFast enabled, if the port receives BPDU messages and BPDU guard is disabled, PortFast is disabled automatically and the BPDU messages will go through the regular STP learning state.

Because BPDU guard does not support per VLAN or auto-recovery, we recommend that you disable it. If a port is in the **err-disable** state because of BPDU guard, you must disable and then enable the port to recover it, which affects all VLANs on the port.

## Enable spanning tree protocol on a fabric

Configuring spanning tree protocol (STP) involves enabling STP on the entire fabric and configuring the STP options for any ports with which you want to use STP.

Follow these steps to enable spanning tree protocol on a fabric.

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

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

**Step 3** In the **Attachments** area, click **Global spanning tree**.

**Step 4** Click **Configure** and follow these substeps.

a) Check **Enable spanning tree protocol (STP) functionality on the fabric** to enable STP on the entire fabric.

By default, global STP is disabled in Nexus Hyperfabric. If you disabled global STP, host ports will drop bridge protocol data units (BPDUs). If you enabled global STP, Nexus Hyperfabric uses per-VLAN spanning tree plus (PVST+) and uses the same bridge ID for all Nexus Hyperfabric switches.

b) Enter the **MAC address** and **Bridge priority** to configure the bridge ID.

Nexus Hyperfabric combines the bridge priority and the MAC address to form the bridge ID. If a switch has a bridge priority that is lower than all the other switches, then Nexus Hyperfabric uses the switch as the root bridge. If more than one switch has the same lowest bridge priority, then Nexus Hyperfabric uses the switch with the lowest MAC address as the root bridge.

The bridge priority value forms the bridge ID for a switch in STP. It is a value between 0 and 61440 in increments of 4096.

We recommend that you use Nexus Hyperfabric as the root bridge. We recommend that you use the default value of "00-00-00-00-00-01" for the MAC address and "0" for the bridge priority.

c) Click **Save**.

## Configure spanning tree protocol at the fabric level

To configure spanning tree protocol (STP), you must first enable STP on the entire fabric and then configure the STP options for each port. This procedure configures STP at the fabric level, but you can also configure STP at the switch level.

Follow these steps to configure spanning tree protocol at the switch level.

**Step 1** Enable STP on the fabric.

See [Enable spanning tree protocol on a fabric, on page 12](#).

**Step 2** In the **Physical Topology** area of the fabric's page, click **Port configurations**.

**Step 3** In the **Action** column, click the edit button (✎) for the port that you want to configure.

- a) In the **Port role** area, select **Host** to configure STP on the port.
- b) Check **Enabled** to enable STP on the port.  
We recommend that Nexus Hyperfabric is the root bridge.
- c) Check **Enable Guard** to allow the port to enter a blocking state if the port receives a superior BPDU.  
We recommend that you enable root guard.
- d) Check **Enable BPDU Guard** to change the port status to an error disabled state if the port receives a BPDU.  
We recommend that you disable BPDU guard.
- e) Check **Enable PortFast** to allow port to directly enter forwarding state by bypassing the listening and learning states.  
We recommend that you enable STP PortFast.

**Step 4** Click **Save**.

## Finish and commit your changes

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




### Note

For a more detailed description of this procedure, see "Workflow for making changes to the fabric" in *Cisco Nexus Hyperfabric—Getting Started*.

Follow these steps to finish and commit your changes.

**Step 1** Click **Review configuration**



**Step 2** Verify your changes in the review list.

**Step 3** Click **Comment and push**.

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

**Step 5** Click **Push configuration**.