Revised: June 2, 2025

Cisco Nexus Hyperfabric — Configure BGP

Border Gateway Protocol (BGP)

BGP is a dynamic routing protocol for routing traffic between different autonomous systems (AS) on the internet. BGP exchanges full routing information when a peer connection is first established and only sends updates for routes that have changed. BGP is highly scalable and can adapt to changes in the network topology, making it suitable for large, complex networks.

For medium to large networks, BGP is a better routing solution than static routing, which is simpler and requires less overhead than BGP but lacks the scalability and adaptability of BGP. With static routing, routes are fixed and do not change unless manually updated by a network administrator.

BGP Policy

BGP import and export policies are sets of rules and configurations that control the selection, advertisement, and acceptance of routing information between BGP peers.

BGP import policy is used for the routes imported from the external BGP peers and BGP export policy is used for the routes exported to the external BGP peers.

- Default import and export policies are predefined and not configurable. The default policy is a template containing rules.
- If you use both default import and export policies, these routes are not advertised to the external peer.
    - Externally learned routes
    - Default route
    - Internal endpoint IP addresses

To view the default import or export policies, navigate to the Fabrics page. In the Attachments area, choose **BGP** > **Import policies** or **BGP** > **Export policies**.

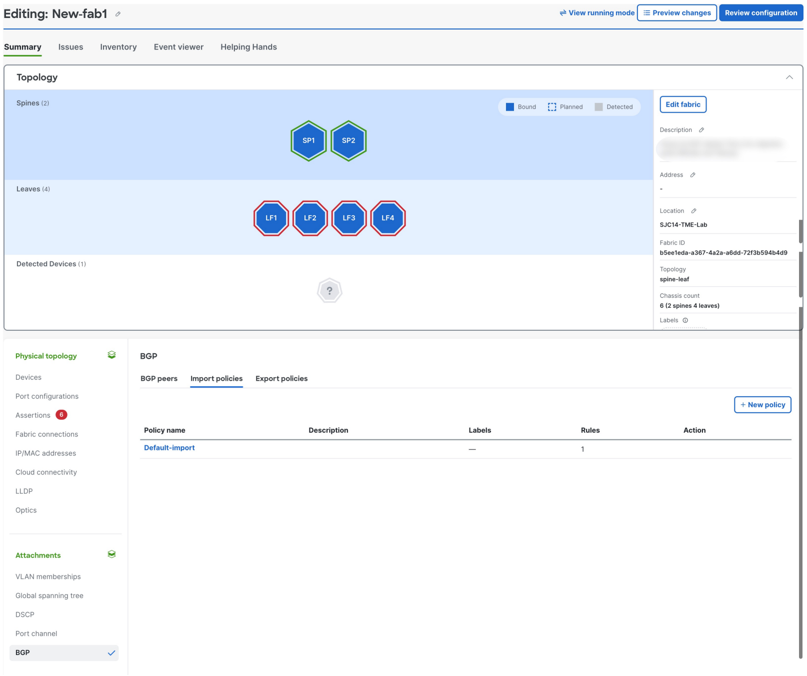

Description
[redacted]

Address
-

Location
SJC14-TME-Lab

Fabric ID
b5ee1eda-a367-4a2a-a6dd-72f3b594b4d9

Topology
spine-leaf

Chassis count
6 (2 spines 4 leaves)

Labels

BGP

Policy name	Description	Labels	Rules	Action
Default-import		—	1	

Default import policy

Table 1: Default Import Policy

Match	Set	Action
All	Community 64511:99	Permit

Default export Policy

Table 2: Default Export Policy

Match	Set	Action
Community 64511:99	-	Deny
Route tag Black	-	Deny
IPv4 Prefix 0.0.0.0/0 (Exact)	-	Deny
IPv6 Prefix ::/0 (Exact)	-	Deny

Match	Set	Action
Community 64510:*	-	Permit
IPv4 Prefix 0.0.0.0/0 (Exact or longer) 32	-	Deny
IPv6 Prefix ::/0 (Exact or longer) 128	-	Deny
All All conditions are matched	-	Permit

- When you add a static route, the **Discard** option in the Add Static Route page, matches **Match Route tag Black, Deny** rule in the default export policy. See Add a static route.
- The rule **Community: 64510:*** is used to advertise one or more static routes to an external network.
- In the default export policy, these rules are used to prevent advertising the default route and internal endpoint IP addresses (host routes) to the external network.
  - Match 0.0.0.0/0 (Exact), Deny
  - Match ::/0 (Exact), Deny
  - Match 0.0.0.0/0 (Exact or longer) 32, Deny
  - Match ::/0 (Exact or longer) 128, Deny
- Because the route imported from the external BGP peer has the rule **Community: 64511:99** set by the default import policy, the route is not exported to a BGP peer because of the **Deny** rule in the default export policy.

Guidelines

- For BGP peers that support both IPv4 and IPv6, you must create a separate BGP peer configuration for each address family.
- BGP peering requires a routed interface or routed sub-interface. An SVI cannot be used.
- When configuring a BGP peer, your selection of available interfaces in the VRF depends on your chosen time to live (TTL) value.
  - With TTL set to 1, you are presented with only the IP interface that is in the same subnet as the BGP peer.
  - With TTL set to greater than 1, you are shown only the routed or routed sub-interface IP interfaces in the VRF.

Create a BGP peer

Create a BGP peer, also known as a BGP neighbor, to establish a BGP session with another switch to exchange routing information. Follow these steps to create a BGP peer.

Step 1 Choose **Fabrics**.

Step 2 On the **Fabrics** page, click the fabric you want to configure.

Step 3 In the Logical Networks area, click **Route tables (VRF)**.

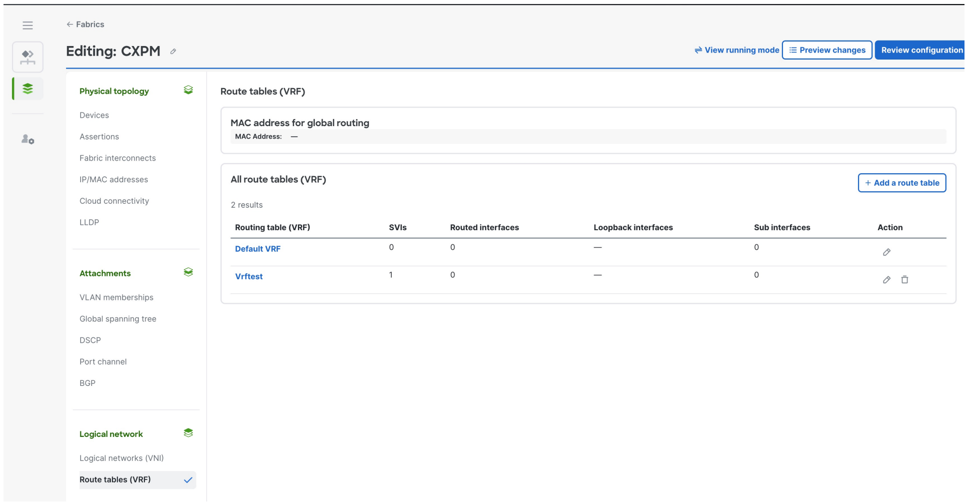

Route tables (VRF)

MAC address for global routing

MAC Address: —

All route tables (VRF)

2 results

Routing table (VRF)	SVIs	Routed interfaces	Loopback interfaces	Sub interfaces	Action
Default VRF	0	0	—	0	
Vrftest	1	0	—	0	

Step 4 In the All route tables (VRFs), click the route table name where you will add a BGP peer.

Step 5 In the Configurations area, click **BGP peers**.

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

Step 6 In the BGP peers area, click the pencil icon and enter the VRF Local AS. Click **Save**.

By default, the VRF Local AS will be used in the BGP peer unless a new value is provided in the Local AS override. See 8.e, on page 5.

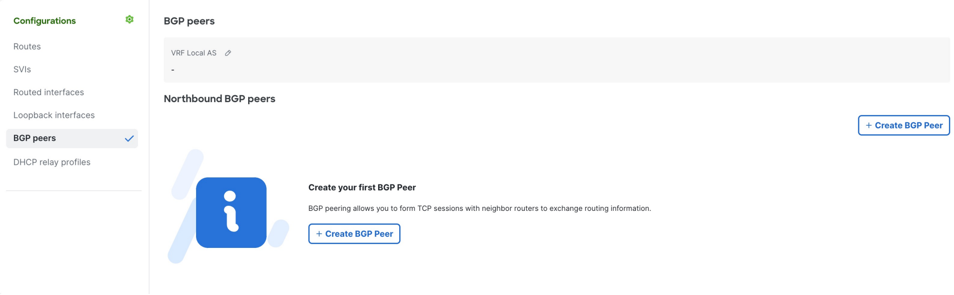

BGP peers

VRF Local AS

-

Northbound BGP peers

Create your first BGP Peer

BGP peering allows you to form TCP sessions with neighbor routers to exchange routing information.

Step 7 Click **Create BGP peer**.

Step 8 Complete the following fields for **Create BGP peer**.

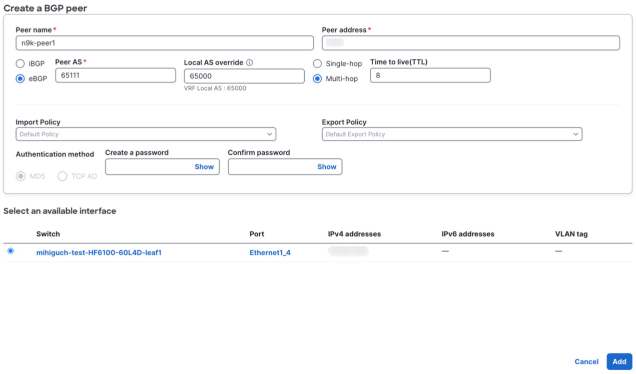

Create a BGP peer

Peer name *: n9k-peer1

Peer address *: [redacted]

○ iBGP ● eBGP

Peer AS *: 65111

Local AS override: 65000

VRF Local AS : 65000

○ Single-hop ● Multi-hop

Time to live(TTL): 8

Import Policy: Default Policy

Export Policy: Default Export Policy

Authentication method: ● MD5 ○ TCP AO

Create a password: Show

Confirm password: Show

Select an available interface

	Switch	Port	IPv4 addresses	IPv6 addresses	VLAN tag
●	mihiguch-test-HF6100-60L4D-leaf1	Ethernet1_4	[redacted]	—	—

a) Enter a name for the new BGP peer.

b) Enter an IP address for the peer.

For IPv4 BGP route exchange, enter an IPv4 address.

For IPv6 BGP route exchange, enter an IPv6 address.

Note For BGP peers that support both IPv4 and IPv6, you must create a separate BGP peer configuration for each address family.

c) Select the BGP type.

d) Enter the autonomous system (AS) number of the peer.

e) (Optional) Enter the **Local AS override** number.

Entering an override local AS number allows a switch to replace the AS number of the originating switch with the AS number of the sending BGP switch, preventing the receiving switch from dropping packets from an originator with the same AS as the receiver. By default, the VRF Local AS will be used unless a new value is provided.

f) For eBGP, select single-hop or multi-hop

g) Enter the time to live (TTL) in number of hops for a peering session.

As a security check, BGP will establish or maintain a session only if the TTL value in the received IP packet header is equal to or greater than the TTL value configured here for the peering session.

h) Select an interface from the list.

i) Select an **Import policy** and an **Export policy** from the drop-down list or keep the defaults.

These policies are defined in Create a BGP import policy, on page 6 and Create a BGP export policy, on page 7.

j) (Optional) Select an **Authentication method** and enter the credentials.

Step 9 Select an available interface.

In the case of a local BGP peer, the available interfaces based on peer address are displayed.

Step 10 Click **Save**.

View BGP peer details

Follow these steps to view BGP peer details.

Step 1 Choose **Fabrics**.

Step 2 On the **Fabrics** page, click the fabric you want to configure.

Step 3 In the Attachments area, choose **BGP** > **BGP peers**.

a) In the BGP peers table, click the peer name.

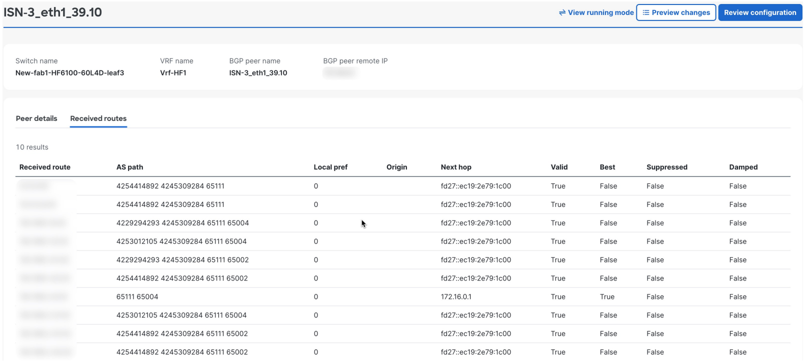

ISN-3_eth1_39.10

Switch name	VRF name	BGP peer name	BGP peer remote IP
New-fab1-HF6100-60L4D-leaf3	Vrf-HF1	ISN-3_eth1_39.10	[redacted]

10 results

Received route	AS path	Local pref	Origin	Next hop	Valid	Best	Suppressed	Damped
[redacted]	4254414892 4245309284 65111	0		fd27::ec19:2e79:1c00	True	False	False	False
[redacted]	4254414892 4245309284 65111	0		fd27::ec19:2e79:1c00	True	False	False	False
[redacted]	4229294293 4245309284 65111 65004	0		fd27::ec19:2e79:1c00	True	False	False	False
[redacted]	4253012105 4245309284 65111 65004	0		fd27::ec19:2e79:1c00	True	False	False	False
[redacted]	4229294293 4245309284 65111 65002	0		fd27::ec19:2e79:1c00	True	False	False	False
[redacted]	4254414892 4245309284 65111 65002	0		fd27::ec19:2e79:1c00	True	False	False	False
[redacted]	65111 65004	0		172.16.0.1	True	True	False	False
[redacted]	4253012105 4245309284 65111 65004	0		fd27::ec19:2e79:1c00	True	False	False	False
[redacted]	4254414892 4245309284 65111 65002	0		fd27::ec19:2e79:1c00	True	False	False	False
[redacted]	4254414892 4245309284 65111 65002	0		fd27::ec19:2e79:1c00	True	False	False	False

b) Click **Peer details** to view peer details.

c) Click **Received routes** to view route information.

Step 4 You can also access the BGP peer details from the VRF page. In the Logical Networks area, click **Route tables (VRF)**.

a) In the All route tables (VRFs), click the route table name.

b) In the Configurations area, click **BGP peers**.

c) In the BGP peers table, click the peer name.

d) Click **Peer details** to view peer details.

e) Click **Received routes** to view route information.

Create a BGP import policy

Follow these steps to create a BGP import policy.

Step 1 Choose **Fabrics**.

Step 2 On the **Fabrics** page, click the fabric you want to configure.

Step 3 In the Attachments area, choose **BGP** > **Import policies**.

A list of BGP import policies appears. If no policies have been added, the list contains only the Default import policy.

Step 4 Click **+ New Policy**.

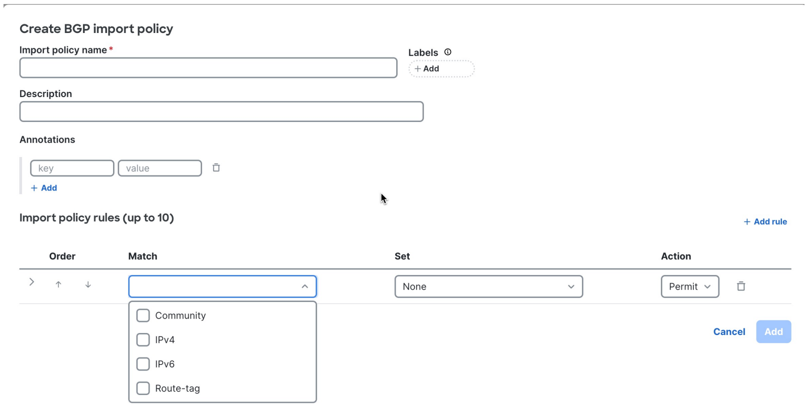

Create BGP import policy

Import policy name *

Labels ⓘ

+ Add

Description

Annotations

key | value

+ Add

Import policy rules (up to 10)

+ Add rule

Order	Match	Set	Action
› ↑ ↓	^	None ⌄	Permit ⌄

□ Community

□ IPv4

□ IPv6

□ Route-tag

Cancel | Add

Step 5 Follow these substeps to create a BGP import policy.

a) Enter a name and description for the new BGP import policy.

b) Optionally, you can add labels for the BGP import policy. To add a label, click **+ Add**, enter the text and press **Enter**.

c) In the Annotations area, click **+ Add** and enter the key value pair to add annotations.

d) In the Import policy rules area, click **+ Add rule** to add a rule. You can add up to 10 rules.

e) In the Order column, use the arrow key to assign an order to the rules in the policy.

f) Select a match for the rule.

Available types are:

- IPv4 — Provide IPv4 prefixes, match conditions, and route origins.
- IPv6 — Provide IPv6 prefixes, match conditions, and route origins.
- Community — Provide community numbers. Wildcards are supported.
- Route-tag — Provide route tags.

g) Select a Set option from the drop-down list.

h) Select an Action (permit or deny) to be taken for packets that match the rule from the drop-down list.

i) Click (›) to add match values for each rule. Up to 5 comma-separated values are supported for each rule (except route tags).

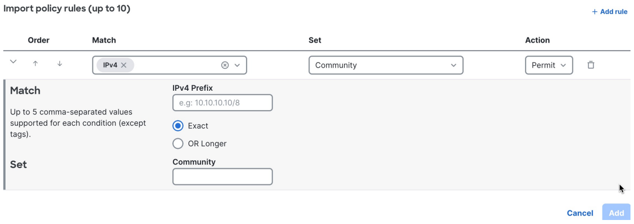

Import policy rules (up to 10)

+ Add rule

Order	Match	Set	Action
⌄ ↑ ↓	IPv4 ×	Community ⌄	Permit ⌄

Match

Up to 5 comma-separated values supported for each condition (except tags).

IPv4 Prefix

e.g: 10.10.10.10/8

● Exact

○ OR Longer

Set

Community

Cancel | Add

j) Click **Add**.

Create a BGP export policy

Follow these steps to create a BGP export policy.

Step 1 Choose **Fabrics**.

Step 2 On the **Fabrics** page, click the fabric you want to configure.

Step 3 In the Attachments area, choose **BGP** > **Export policies**.

A list of BGP export policies appears. If no policies have been added, the list contains only the Default export policy.

Step 4 Click **+ New Policy**.

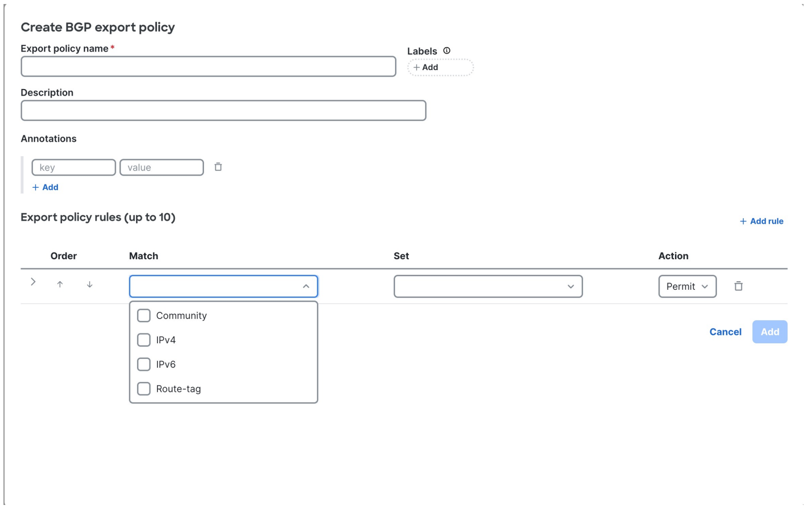

Step 5 Follow these substeps to create a BGP export policy.

a) Enter a name and description for the new BGP export policy.

b) Optionally, you can add labels for the BGP export policy. To add a label, click **+ Add**, enter the text and press **Enter**.

c) In the Annotations area, click **+ Add** and enter the key value pair to add annotations.

d) In the Export policy rules area, click **+ Add rule** to add a rule. You can add up to 10 rules.

e) In the Order column, use the arrow key to assign an order to the rule.

f) Select a match for the rule.

Available types are:

- Community — Provide community numbers. Wildcards are supported.
- IPv4 — Provide IPv4 prefixes, match conditions, and route origins.
- IPv6 — Provide IPv6 prefixes, match conditions, and route origins.
- Route-tag — Provide route tags.

g) Select a Set option from the drop-down list.

Available types are:

- AS path pad — Provide AS path list.
- Community — Provide community numbers.
- Next-hop IPv4 — Provide IPv4 address of the connection between the peers.
- Next-hop IPv6 — Provide IPv6 address of the connection between the peers.

h) Select an Action (permit or deny) to be taken for packets that match the rule from the drop-down list.

i) Click (>) to add match values and set values for each rule. Up to 5 comma-separated values are supported for each rule (except route tags).

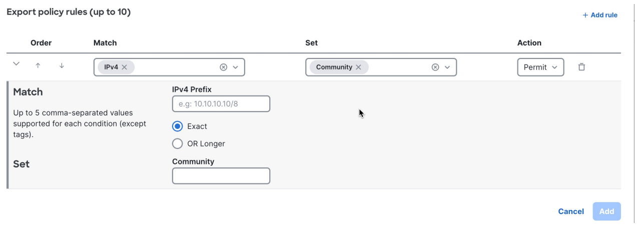

Export policy rules (up to 10)

Order	Match	Set	Action
	IPv4	Community	Permit

Match

Up to 5 comma-separated values supported for each condition (except tags).

IPv4 Prefix: e.g: 10.10.10.10/8

● Exact

○ OR Longer

Set

Community:

j) Click **Add**.

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**

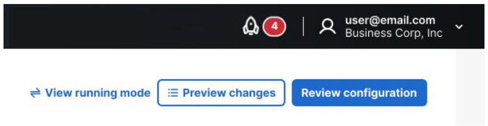

Step 2 Verify your changes in the review list.

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**.