Configuring Inband Management and Out-of-Band PnP, Release 4.3.1

Table of Contents

New and changed information . . . 1
Inband management in External fabrics and Classic LAN fabrics . . . 2
- Inband management . . . 2
- Prerequisites . . . 3

Inband PnP for Catalyst 9200-CX ToR switches . . . 5
- Key aspects of inband PnP for ToR switches . . . 5
- Prerequisites . . . 7
- Guidelines and limitations . . . 8
- Configuration templates used . . . 8
- DHCP and PnP traffic flow . . . 9
- Configure Campus VXLAN EVPN fabric . . . 9
  - Create management VRF and SVI . . . 9
  - Configure VRF leaking on BGW spines . . . 11
  - Apply the PnP VLAN template on leaf-to-ToR interfaces . . . 12
- Onboard Catalyst 9200-CX ToR switches . . . 13
  - Create the external fabric . . . 13
  - Enable inband management and bootstrap . . . 14
  - Add ToR switches using bootstrap . . . 14
  - Verify onboarding . . . 15
- Troubleshoot ToR onboarding issues . . . 16

Inband POAP management in External fabrics and Classic LAN fabrics . . . 17
- Inband POAP . . . 17
- Prerequisites . . . 17
- Guidelines and limitations . . . 18
- Enable inband management and POAP on External fabrics and Classic LAN fabrics . . . 18
- Add switches . . . 19
  - Pre-provision switches to a fabric . . . 19
  - Import switches using the bootstrap mechanism . . . 20
- Add an interface . . . 20
- Add a policy to a fabric . . . 21
- Recalculate and deploy configurations on a switch . . . 21

Inband management and inband POAP in Data Center VXLAN EVPN fabrics . . . 22
- Prerequisites for inband and out-of-band management . . . 23
- Guidelines and limitations . . . 24
- Enable inband POAP on Data Center VXLAN EVPN fabrics . . . 25
- Import switches in a Brownfield deployment . . . 26
- Pre-provision switches through inband POAP . . . 26
- Add a policy for Data Center VXLAN EVPN fabrics . . . 27
- Change the fabric management mode . . . 27

Secure POAP . . . 29
- Prerequisites for secure POAP . . . 29

Inband PnP in Campus VXLAN EVPN fabrics . . . 30
About Inband PnP for Cisco Catalyst switches . . . 30
Key aspects of Inband PnP . . . 30
Prerequisites for Inband PnP . . . 30
Guidelines and Limitations . . . 30
Configure N9k DC border switch . . . 31
Configure DC BGW spines . . . 32
Configure Campus fabric . . . 34
Enable Inband PnP . . . 35
Import switches using bootstrap method . . . 37
Out-of-Band PnP in Campus VXLAN EVPN fabrics . . . 41
About Out-of-Band PnP in Campus VXLAN EVPN fabrics . . . 41
Enable Out-of-Band PnP . . . 41
Import switches using the bootstrap method . . . 42
Enable DHCP for PnP . . . 42
Copyright . . . 44

New and changed information

The following table provides an overview of the significant changes up to this current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

Release Version	Feature	Description
Nexus Dashboard 4.3.1	Support for inband plug and play (PnP) for Catalyst 9200-CX ToR switches	Beginning with Nexus Dashboard release 4.3.1, inband PnP supports the bootstrap and management of Catalyst 9200-CX switches acting as Layer-2 Top-of-Rack (ToR) switches in an external fabric. This feature enables zero-touch deployment for access switches connected to an inband-managed Campus fabric. For more information, see Inband PnP for Catalyst 9200-CX ToR switches.
Nexus Dashboard 4.3.1	Cisco Nexus 9000 Series name change	Cisco Nexus 9000 Series is now the Cisco N9000 Series.

Inband management in External fabrics and Classic LAN fabrics

Starting with the Nexus Dashboard 4.1 release, **Classic LAN** is considered an upper-level category, under which you can find these fabric types:

- Enhanced Classic LAN (both for creating new deployments and onboarding existing networks)
- Legacy Classic LAN (for onboarding existing networks)

See Editing Classic LAN Fabric Settings for more information.

Inband management

Cisco Nexus devices have dedicated out-of-band (OOB) management ports (mgmt0) to manage devices via telnet or SSH connections.

You can manage Cisco Nexus devices through inband using front panel ports either by assigning management IP addresses on one of the ports or using loopback or SVI. By default, (mgmt0) interface is part of management VRF.

In Nexus Dashboard by default, VRF is used for inband management, you can use other defined VRFs for inband management for nexus devices. Inband Management is the ability to administer a network through LAN connection.

You can import or discover switches with inband connectivity for External and Classic LAN fabrics in Brownfield deployments only. Enable inband management per fabric, while configuring or editing the Fabric settings. You cannot import or discover switches with inband connectivity using POAP.

After configuration, the Fabric tries to discover switches based on the VRF of the inband management. The fabric template determines the VRF of the inband switch using seed IP. If there are multiple VRFs for the same seed IP, then no intent will be learned for seed interfaces. You must create intent or configuration manually.

After configuring or editing the Fabric settings, you must Deploy Config. You cannot change the inband management settings after you import inband managed switches to the Fabric. If you uncheck the check box, the following error message is generated.

```
Inband IP <<IP Address>> cannot be used to import the switch,
please enable Inband Mgmt in fabric settings and retry.
```

After the switches are imported to the Fabric, you must manage the interfaces to create intent. Create the intent for the interfaces that you are importing the switch. Edit\update the Interface configuration. When you try to change the Interface IP, for this inband managed switch, an error message is generated:

```
Interface <<interface_name>> is used as seed or next-hop egress interface
```

for switch import in inband mode.
IP/Netmask Length/VRF changes are not allowed for this interface.

While managing the interfaces, for switches imported using inband management, you cannot change the seed IP for the switch. The following error will be generated:

<<switch-name>>: Mgmt0 IP Address (<ip-address>) cannot be changed,
when is it used as seed IP to discover the switch.

Create a policy for next-hop interfaces. Routes to Nexus Dashboard from 3rd party devices can contain multiple interfaces, which are known as ECMP routes. Find the next-hop interface and create an intent for the switch. Interface IP and VRF changes are not allowed.

If inband management is enabled, during Image management, the data interface of nexus dashboard is used to copy images on the switch, in ISSU, EPLD, RPM & SMU installation flows.

If you import the switches using inband connectivity in the fabric and later disable the inband Mgmt in the Fabric settings after deployment, the following error message is generated:

The fabric <<fabric name>> was updated with below message:
Fabric Settings cannot be changed for Inband Mgmt when switches are already imported using inband Ip.
Please remove the existing switches imported using Inband IP from the fabric, then change the Fabric Settings.

However, the same fabric can contain switches imported using both inband and out-of-band connectivity.

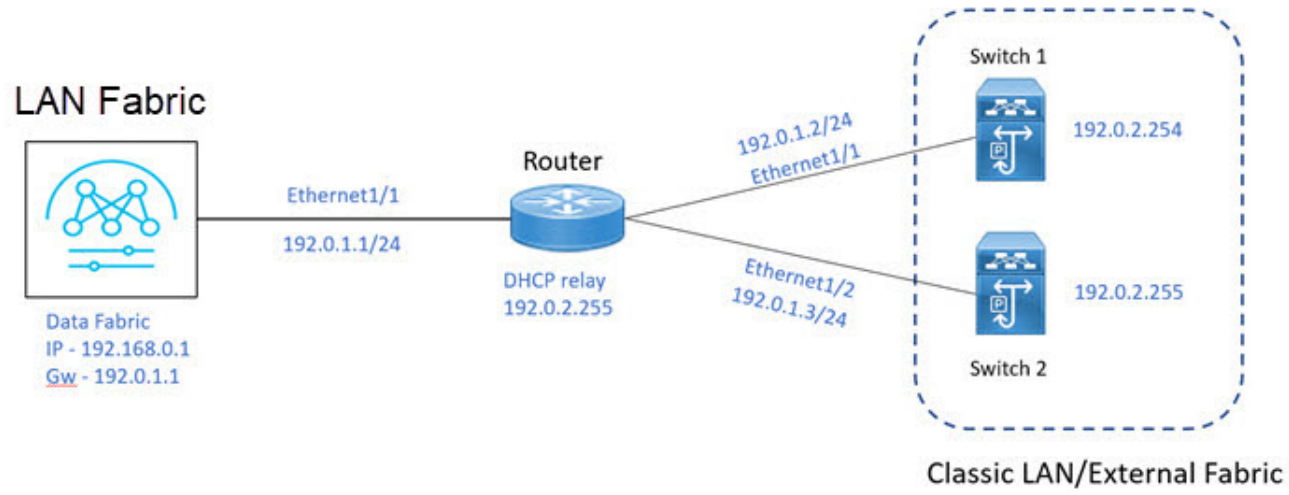

Prerequisites

The following are the prerequisites for using inband management:

- Configure the appropriate data network routes for reachability to the switch inband interfaces. Choose **Admin > System Settings > General**, locate the **Routes** area, then click **Edit** to enter route IP addresses in the **Data network routes** area.
- On the Nexus Dashboard Web UI, navigate to **Admin > System Settings > Fabric management > Advanced settings > Admin** and choose **Data** from the **LAN Device Management Connectivity** drop-down list to manage Data Center VXLAN EVPN fabrics through inband management, or an error message is displayed. If you choose **Data**, ensure that the required data service IP addresses are available in the Nexus Dashboard **External Service Pools** area (**Admin > System Settings > General > External pools**).

When server settings changed from **Data** to **Management** or vice-versa, allow some time for syslog or poap functionalities to be online and ensure that the IP addresses in Cluster configuration are moved to the appropriate pool.

Inband PnP for Catalyst 9200-CX ToR switches

Beginning with release 4.3.1, Nexus Dashboard supports the inband bootstrap and management of Catalyst 9000 series switches (specifically the 9200-CX) acting as Layer-2 Top-of-Rack (ToR) switches. You can connect these ToR switches to Catalyst 9000 leaf switches that you have already onboarded to an inband-managed Campus VXLAN EVPN fabric.

Unlike Catalyst leaf or border switches (which are added directly to the Campus fabric), you import Catalyst 9200-CX ToR switches into a separate external fabric configured for inband management. The Campus fabric provides the transport for DHCP and PnP signaling between the ToR and Nexus Dashboard.

This feature delivers a zero-touch deployment experience for access-layer switches connected to a Catalyst VXLAN EVPN fabric.

Key aspects of inband PnP for ToR switches

- **Two-fabric design**—The Campus VXLAN EVPN fabric hosts the Border Gateway (BGW) spines and leaf switches; the external fabric hosts the ToRs. You must configure both fabrics for inband management.
- **Inband management network for ToRs**—You can create a Layer 3 network with an SVI on the Campus fabric leaf switches that provides connectivity to ToRs. The SVI provides the anycast gateway and DHCP relay function for ToR onboarding.
- **VRF leaking on the BGW spines**—The BGW spines leak routes between the VRF (Layer 3) that contains the management SVI and the default VRF, where Nexus Dashboard and the DHCP server reside.
- **PnP VLAN signaling**—When a new ToR powers on for the first time, it sends PnP discovery traffic on its default VLAN. The management SVI for ToR onboarding typically uses a different VLAN, so the traffic cannot reach it. To fix this, the upstream leaf instructs the ToR which VLAN to use. The pnp_vlan template handles this automatically.

  

  The VLAN ID you specify in the pnp_vlan template must match the VLAN ID of the management SVI that you created on the Campus fabric. You must also configure this same VLAN ID on the **Bootstrap** page in the fabric settings of External fabrics. If these VLAN IDs do not match, the ToR cannot reach the management SVI, and the onboarding fails.

- **Automatic ToR configuration**—Once the ToR establishes a PnP channel with Nexus Dashboard via DHCP, Nexus Dashboard uses this channel to automatically push the management SVI and default route configuration to the switch.
- **Standard external fabric management**—After onboarding, you can manage the ToR like any other device in the external fabric.

The following diagram illustrates the two-fabric architecture and the PnP channel traffic flow required for successful ToR onboarding.

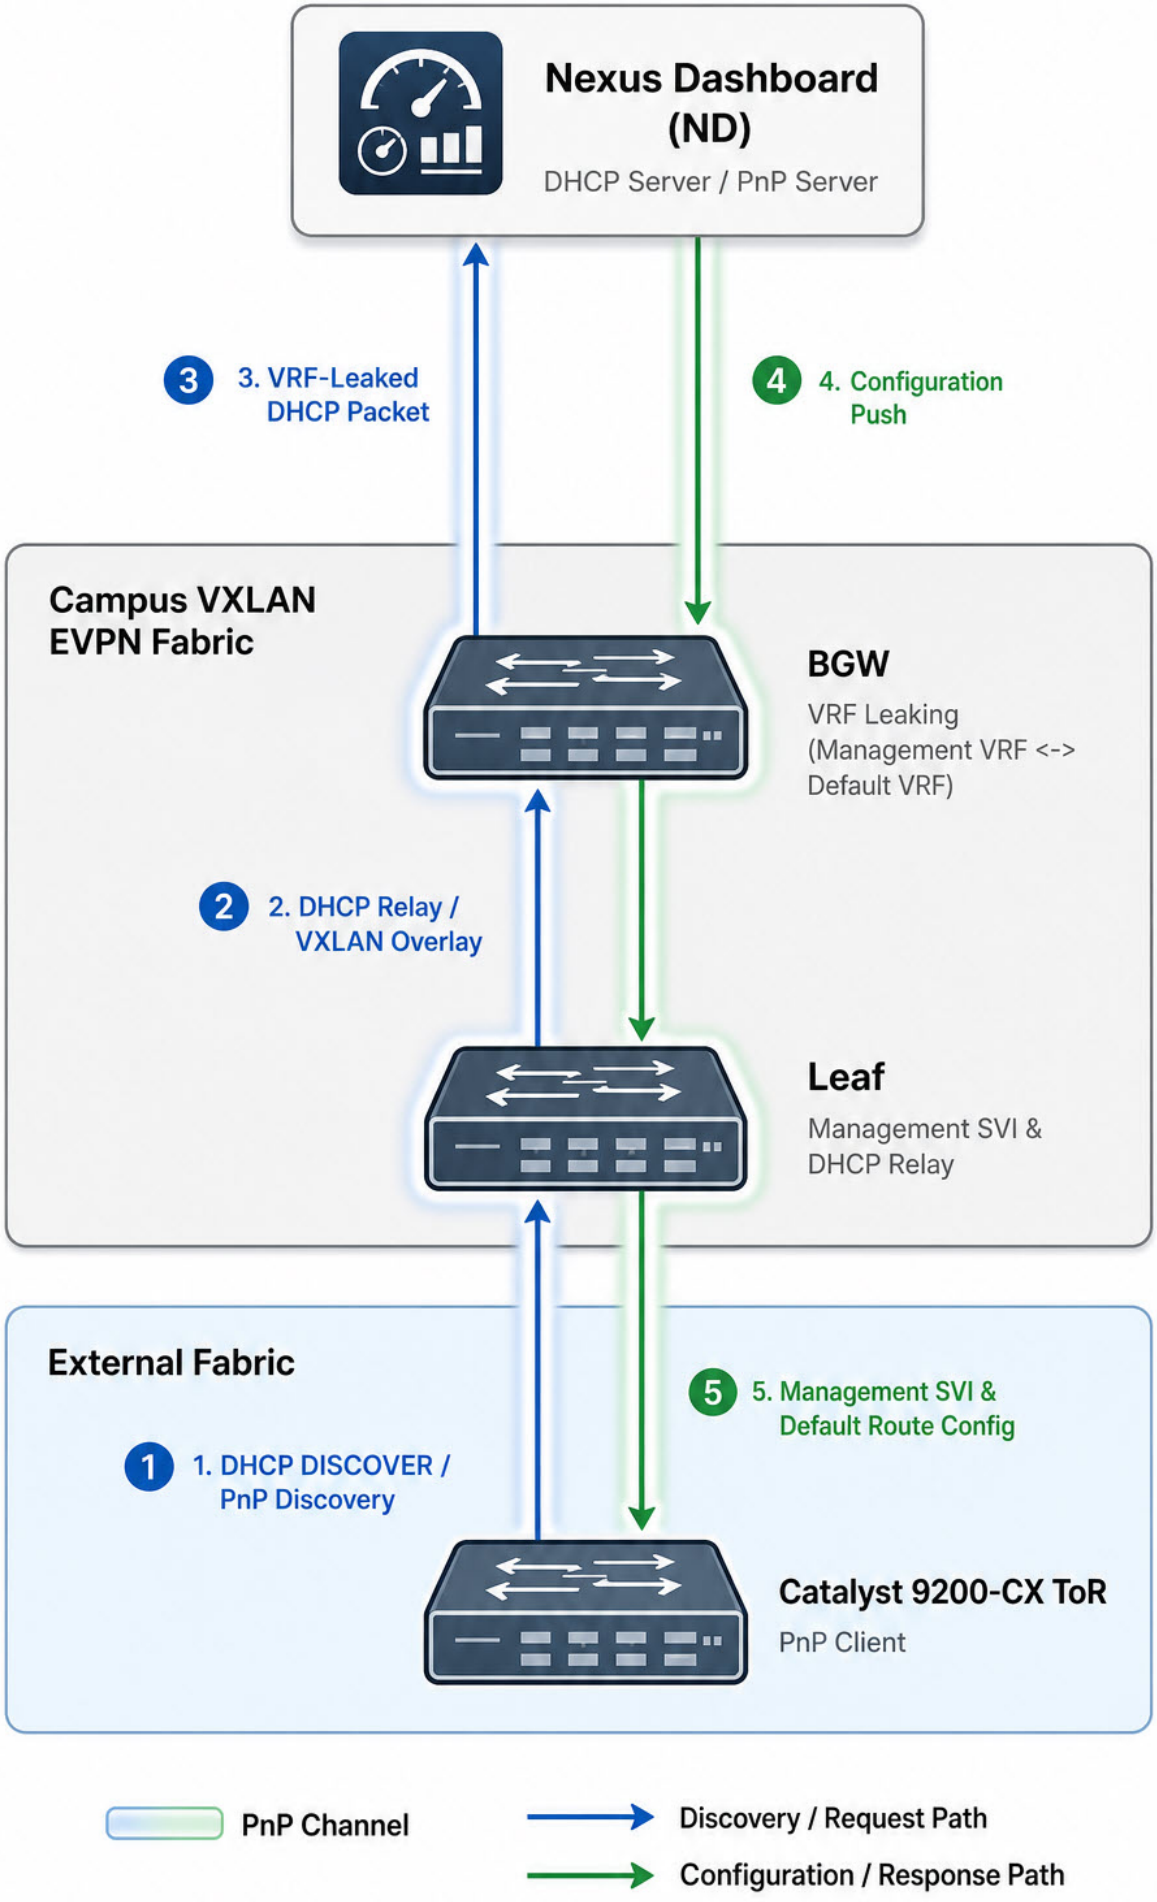
Nexus Dashboard (ND)
DHCP Server / PnP Server
3. VRF-Leaked DHCP Packet
4. Configuration Push
Campus VXLAN EVPN Fabric
BGW
VRF Leaking (Management VRF <-> Default VRF)
2. DHCP Relay / VXLAN Overlay
Leaf
Management SVI & DHCP Relay
External Fabric
5. Management SVI & Default Route Config
1. DHCP DISCOVER / PnP Discovery
Catalyst 9200-CX ToR
PnP Client
PnP Channel
Discovery / Request Path
Configuration / Response Path

Prerequisites

Before you onboard Catalyst 9200-CX ToR switches, complete the following prerequisites.

Fabric and topology prerequisites

- Confirm that the Campus VXLAN EVPN fabric exists and uses inband management.
- Onboard the N9K BGW spine to the Campus fabric using inband bootstrap.
- Onboard the Catalyst 9000 leaf switches that connect to the ToRs to the Campus fabric using inband bootstrap.
- Ensure that the Catalyst 9200-CX ToR switches meet the following conditions:
  - Powered on
  - Connected to the Catalyst 9000 leaf switches
  - Remain in a PnP-ready state (no startup configuration present)
- Verify that the Nexus Dashboard cluster (vND or pND) can reach the inband management network.

Configuration prerequisites on the Campus fabric

- Create a VRF (new or existing) on the Campus fabric to host the ToR management network.
- Create a network SVI in this VRF and attach it to every leaf interface that connects to a ToR.
- Enable DHCP relay on the SVI, pointing to the Nexus Dashboard DHCP server or an external DHCP server.
- Configure the DHCP server with a scope that matches the SVI subnet and serves the appropriate VRF.
- Apply the Vrf_leaking_BGW template on the BGW spine to import and export routes between the SVI's VRF and the default VRF.
- Apply the pnp_vlan template on the leaf interfaces facing the ToRs, with the VLAN ID matching the management SVI VLAN.

Connectivity and routing prerequisites

- Ensure that both the Campus fabric and the external fabric use inband management.
- If an external router connects the BGW spine and Nexus Dashboard, configure one of the following to ensure proper route exchange:
  - Establish a BGP EVPN session between the BGW spine and the external router to share routes learned on the external router with the BGW spine.
  - Configure static routes for all relevant subnets and redistribute them into BGP.
- In Nexus Dashboard, add routes for the subnets that the ToRs use, as well as the subnet that assigns Loopback0 addresses to the leaf switches.

  Follow these steps to add subnets

  1. Navigate to **Admin > System Settings > General**.
  2. In the **Routes** area, click **Edit**.

     The **Routes** page appears.

3. Under **Data network routes**, click **Add data network route**.

   The **Add data network route** page appears.

4. In the **Data network route** field, add the following subnets:
   - The DHCP subnets that the ToR use.
   - The subnet that assigns Loopback0 addresses to the leaf switches.
   - The underlay subnet IP range.
5. Click **Save**.

Guidelines and limitations

The following guidelines and limitations apply to inband PnP for Catalyst 9200-CX ToR switches.

- You can deploy and manage Catalyst 9200-CX ToR switches only in an external fabric.
- Nexus Dashboard does not provide automated network-attachment configuration for these ToR switches. After onboarding, you manage them as standard devices within the external fabric.
- During bootstrap, you must explicitly set the **Switch Role** to **ToR**.
- The IP address you specify when adding the switch using bootstrap must match the IP address that DHCP allocates to the switch.
- Nexus Dashboard does not push telemetry to Catalyst ToR switches, even when you enable telemetry on the external fabric.
  - Endpoint views originate at the L3 leaf boundaries, not at the ToR.
  - Traffic analytics flows originate and terminate at leaf nodes, not at ToRs.
- This feature supports DHCP for IPv4 only; IPv6 is not supported.
- The ToR must have no startup configuration so that Nexus Dashboard can discover it through PnP.

Configuration templates used

Template	Applied on	Purpose
Vrf_leaking_BGW	BGW switches in the Campus fabric	Imports and exports routes between the VRF containing the ToR management SVI and the default VRF, where Nexus Dashboard and the DHCP server reside.
pnp_vlan	Leaf interfaces in the Campus fabric facing the ToRs	Configures pnp start-up vlan <vlan-id> on the upstream leaf so that the ToR sends DHCP DISCOVER on the VLAN that matches the management SVI.

PnP uses the default VLAN when the ToR first powers on. Because the management SVI typically uses a different VLAN, you must apply the pnp_vlan template so that the ToR's DHCP DISCOVER reaches the correct SVI.

DHCP and PnP traffic flow

When a Catalyst 9200-CX ToR powers on without a startup configuration, the following sequence occurs:

1. The ToR enters PnP mode and sends a DHCP DISCOVER on the VLAN that the upstream leaf signals (configured by the pnp_vlan template).
2. The DHCP DISCOVER reaches the management SVI on the leaf, which acts as a DHCP relay.
3. The leaf forwards the relayed DHCP packet across the VXLAN overlay to the BGW spine.
4. The BGW spine leaks the packet into the default VRF (configured by the Vrf_leaking_BGW template), where Nexus Dashboard and the DHCP server reside.
5. The DHCP server responds with an IP address and PnP server information. The response traverses the reverse path back to the ToR.
6. The ToR connects to the Nexus Dashboard PnP server and downloads its configuration, which includes the management SVI and a default route pointing to the anycast gateway on the leaf.

Configure Campus VXLAN EVPN fabric

Before you onboard ToRs in the external fabric, ensure that you complete the following steps on the Campus VXLAN EVPN fabric.

Create management VRF and SVI

Follow these steps to create the management VRF and SVI.

1. Navigate to **Manage > Fabrics**, then click the Campus VXLAN EVPN fabric.

   The **Fabric Overview** page displays.

2. Click the **Segmentation and security** tab.
3. Click the **VRFs** tab.
4. From the **Actions** drop-down list, choose **Create** to create a new VRF (for example, MyVRF).

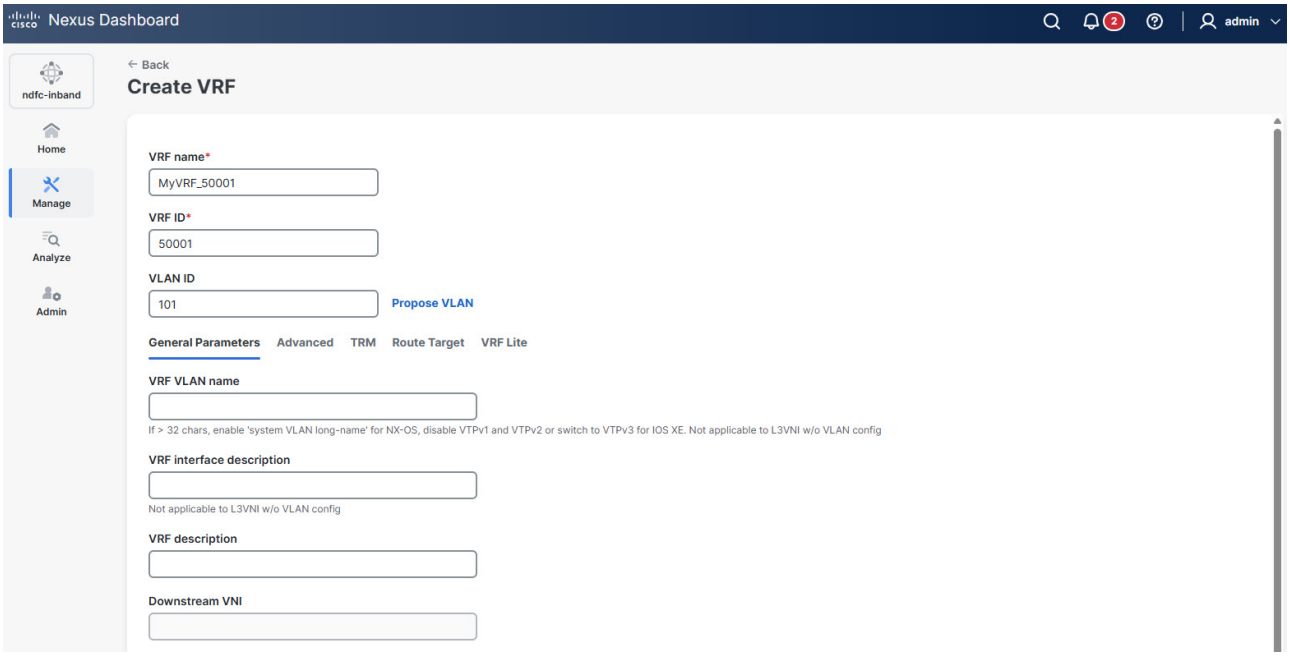

5. After you create a new VRF, navigate to **Segmentation and security** tab and click the **Networks** tab.
6. From the **Actions** drop-down list, choose **Create** to create a new network SVI in the VRF that you created (for example, MyVRF).

   The **Create Network** page displays.

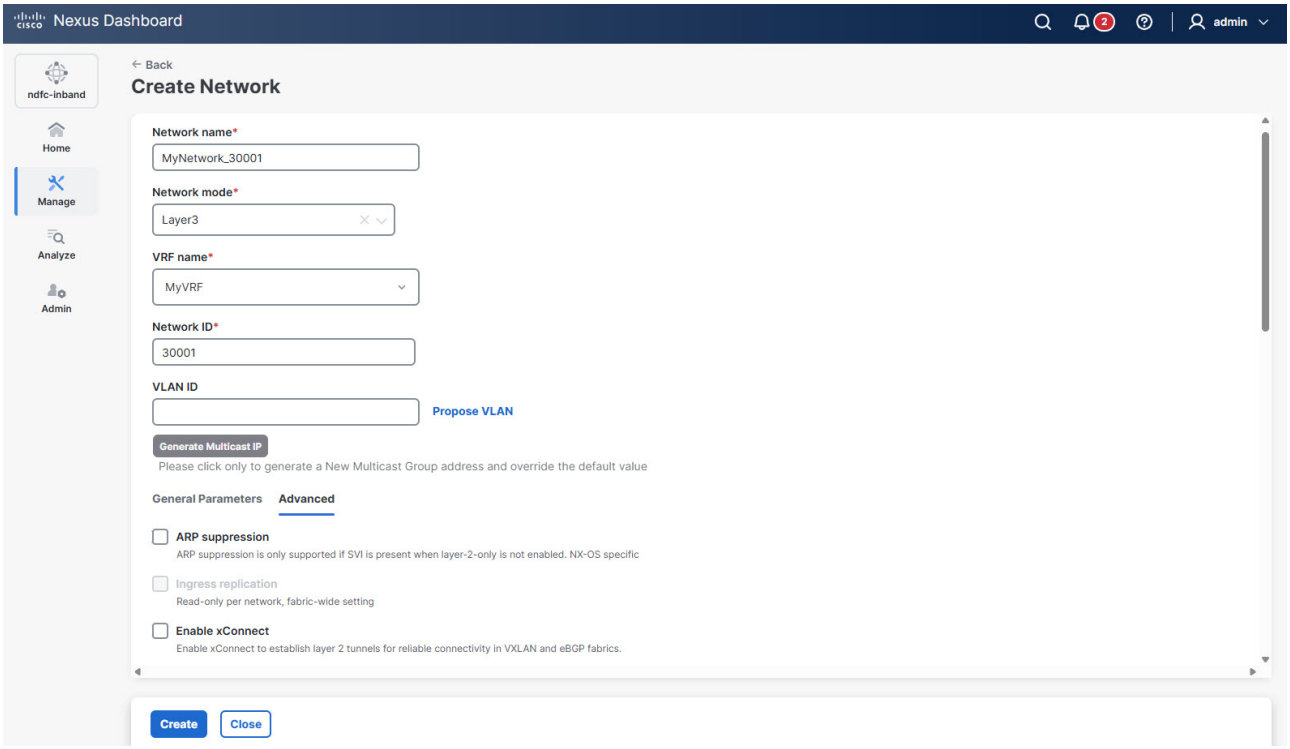

7. Provide the network configuration details.
   - From the **VRF name** drop-down list, choose the VRF that you previously created (for example, MyVRF).
   - Under **General parameters**, provide the **IPv4 gateway/netmask**.
8. Click **Advanced** and scroll to the **DHCP relay server information** area.

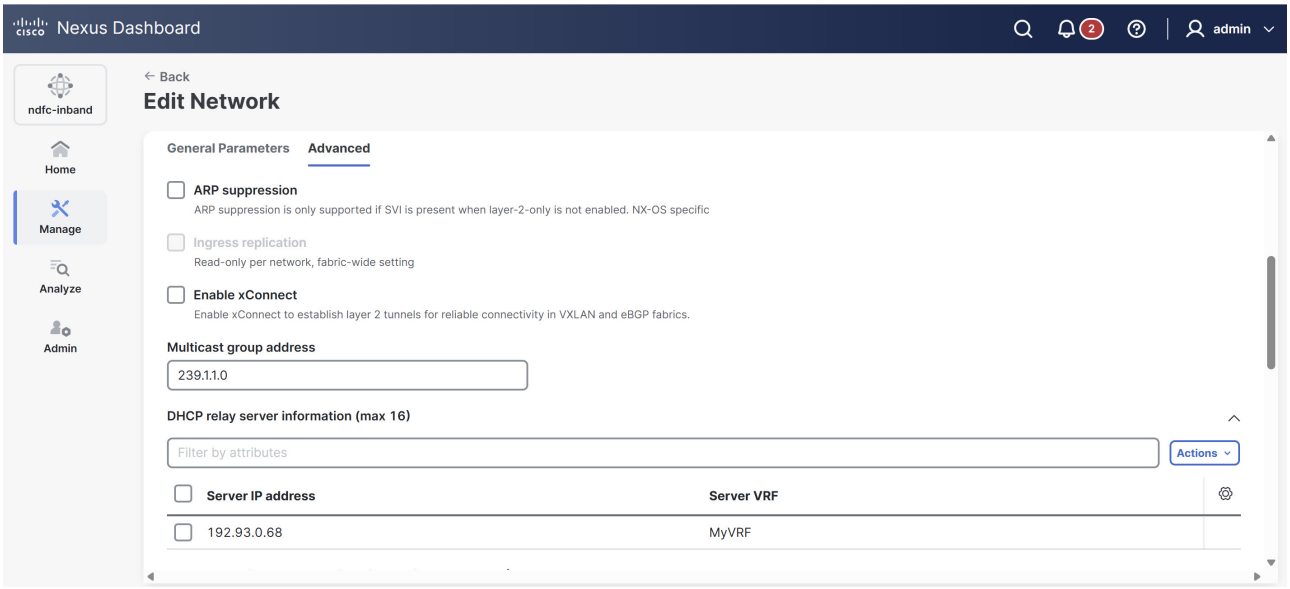

9. In the **DHCP relay server information** area, go to **Actions > Add**.
10. In the **Add Item** dialog box, provide the **Server IP address** and the **Server VRF**.

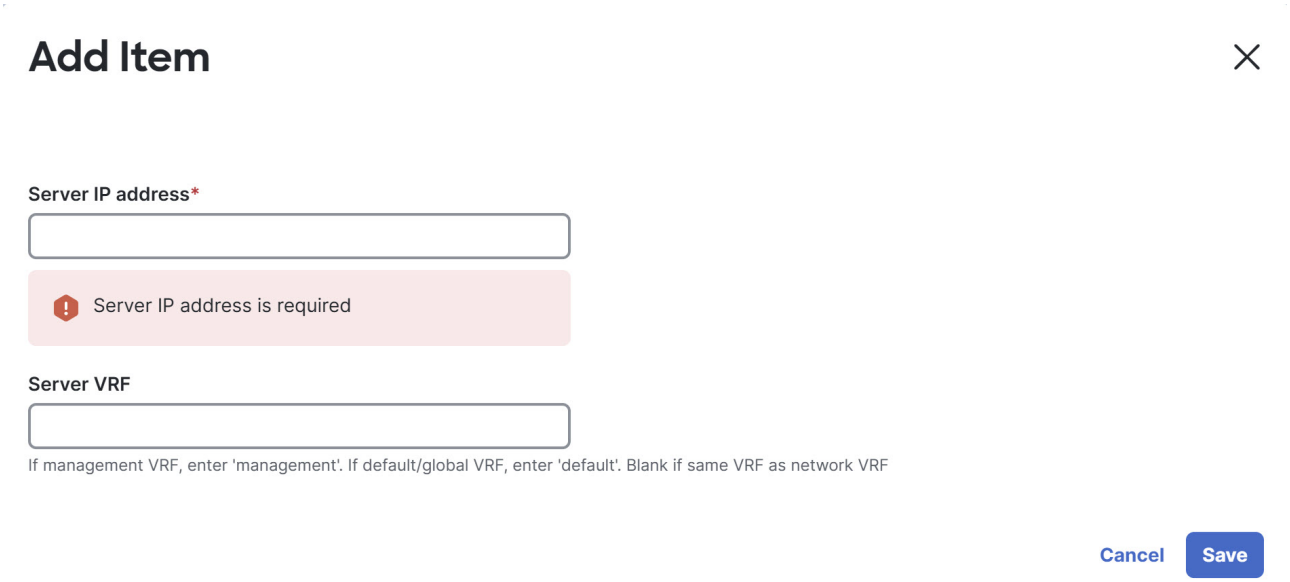

11. Click **Save** to save the network configuration.
12. From the **Networks** table, choose the network you created.
13. From the **Actions** drop-down list, choose **Multi-attach** to attach the network to every leaf interface that connects to a ToR.
14. Navigate to **Fabric Overview** page, from the **Actions** drop-down list, choose **Recalculate and deploy**.

Configure VRF leaking on BGW spines

Follow these steps to configure VRF leaking on BGW spines.

1. Navigate to **Manage > Fabrics**, then click the Campus VXLAN EVPN fabric.

   The **Fabric Overview** page displays.

2. Navigate to **Configuration Policies > Policies**.
3. From the **Actions** drop-down list, choose **Add policy**.

   The **Add policy** page displays.

4. Under **Select switches**, choose the BGW switches and click **Next**.
5. Under **Configure policy**, click **Select Policy Template**.

   The **Select Policy Template** page displays.

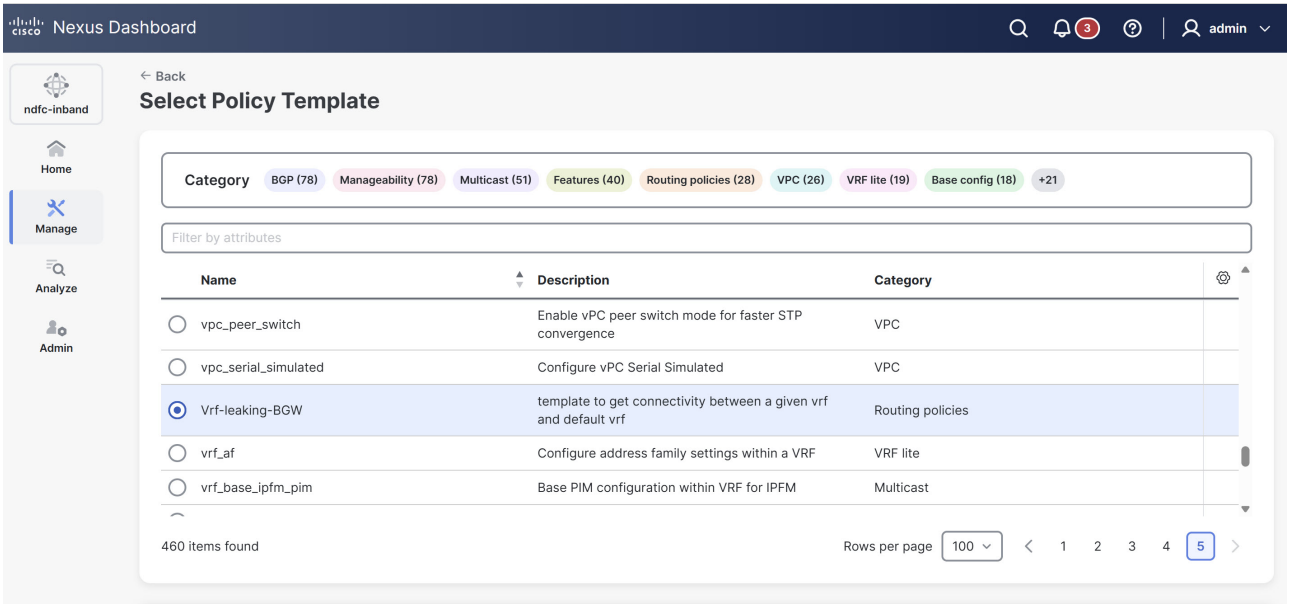

6. Choose the Vrf_leaking_BGW template and click **Select**.
7. Provide the source and destination VRFs (typically the SVI's VRF and the default VRF).
8. Save the configuration and deploy.

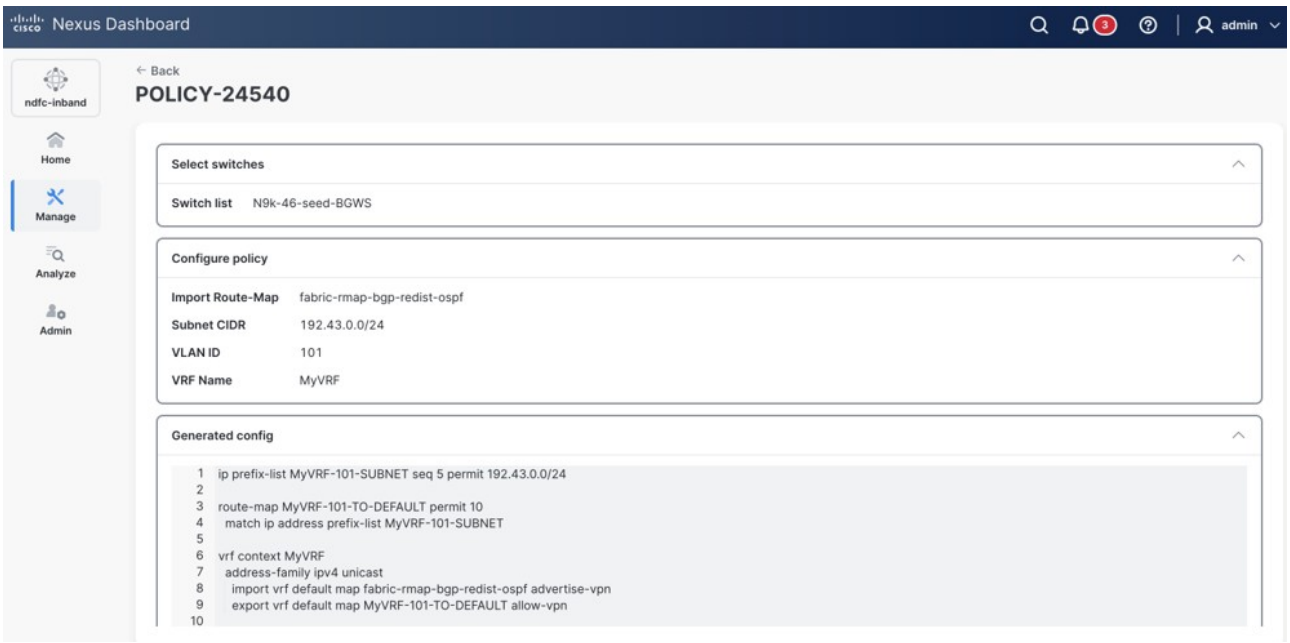

Apply the PnP VLAN template on leaf-to-ToR interfaces

Follow these steps to apply the PnP VLAN template.

1. Navigate to **Manage > Fabrics**, then click the Campus VXLAN EVPN fabric.

   The **Fabric Overview** page displays.

2. Navigate to **Connectivity > Interfaces**.
3. From the **Interfaces** table, choose all leaf interfaces that connect to the ToRs.
4. From the **Actions** drop-down list, choose **Edit configuration**, choose the pnp_vlan template, and enter the VLAN ID that matches the management SVI.

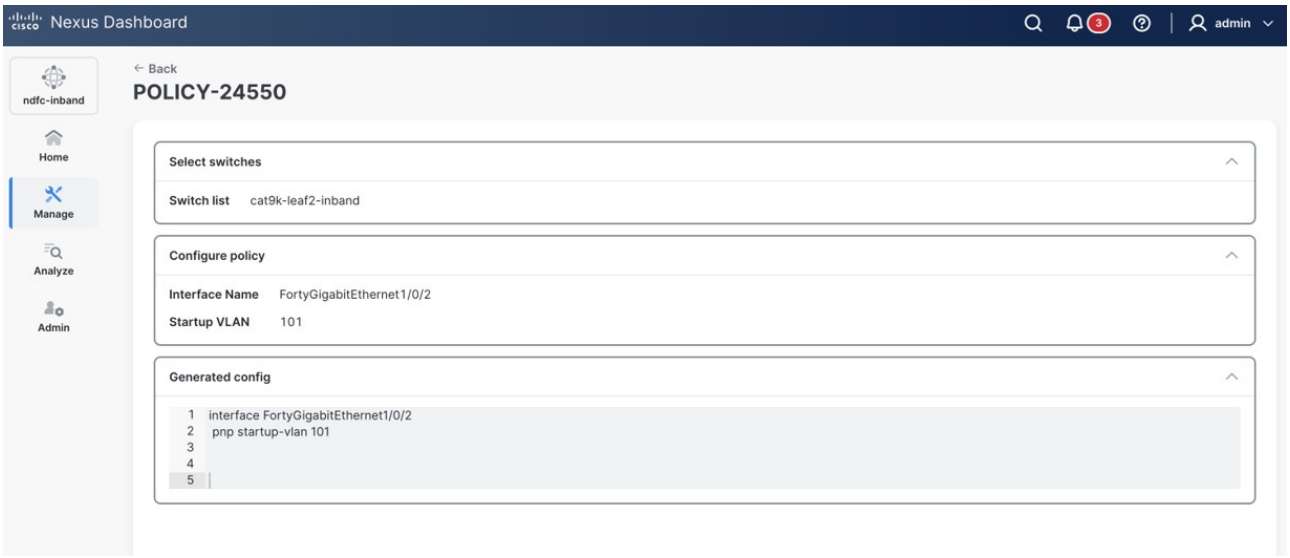

5. Save the configuration and deploy.

Onboard Catalyst 9200-CX ToR switches

After you prepare the Campus fabric, follow these steps to onboard Catalyst 9200-CX ToR switches into an external fabric.

Create the external fabric

Follow these steps to create the external fabric for ToR onboarding.

1. Navigate to **Manage > Fabrics**.
2. Click the **Fabrics** tab.
3. From the **Actions** drop-down list, choose **Create fabric**.

   The **Create/Onboard Fabric** page displays.

4. Under **Select a category**, choose **Create new LAN fabric** and click **Next**.
5. Under **Select a type**, click **External and Inter-Fabric Connectivity**, then click **Next**.
6. Under **Settings**, choose **Advanced**.
7. Provide the fabric name, location, and BGP ASN.
8. Choose the appropriate **License tier for fabric**.
9. Clear the **Fabric monitor mode** check box.
10. Click **Save** to create the fabric.

Enable inband management and bootstrap

Follow these steps to enable inband management and bootstrap on the external fabric.

1. Navigate to **Manage > Fabrics**.
2. Choose the external fabric that you created.
3. From the **Actions** drop-down list, choose **Edit fabric settings**.

   The **Edit *fabric-name* settings** page displays.

4. Click **Fabric Management > Advanced** and check the **Inband Mgmt** check box.
5. Click the **Bootstrap** tab.
6. Check the **Enable Bootstrap** check box.
7. Check the **Enable Local DHCP Server** check box.
8. In the **Catalyst ToR Bootstrap VLAN** field, enter the same VLAN ID that you assigned to the management SVI on the Campus fabric (for example, VLAN 101).
9. Configure the DHCP scope to match the SVI subnet that you defined on the Campus fabric.

   Nexus Dashboard allocates IP addresses from this subnet to the ToRs.

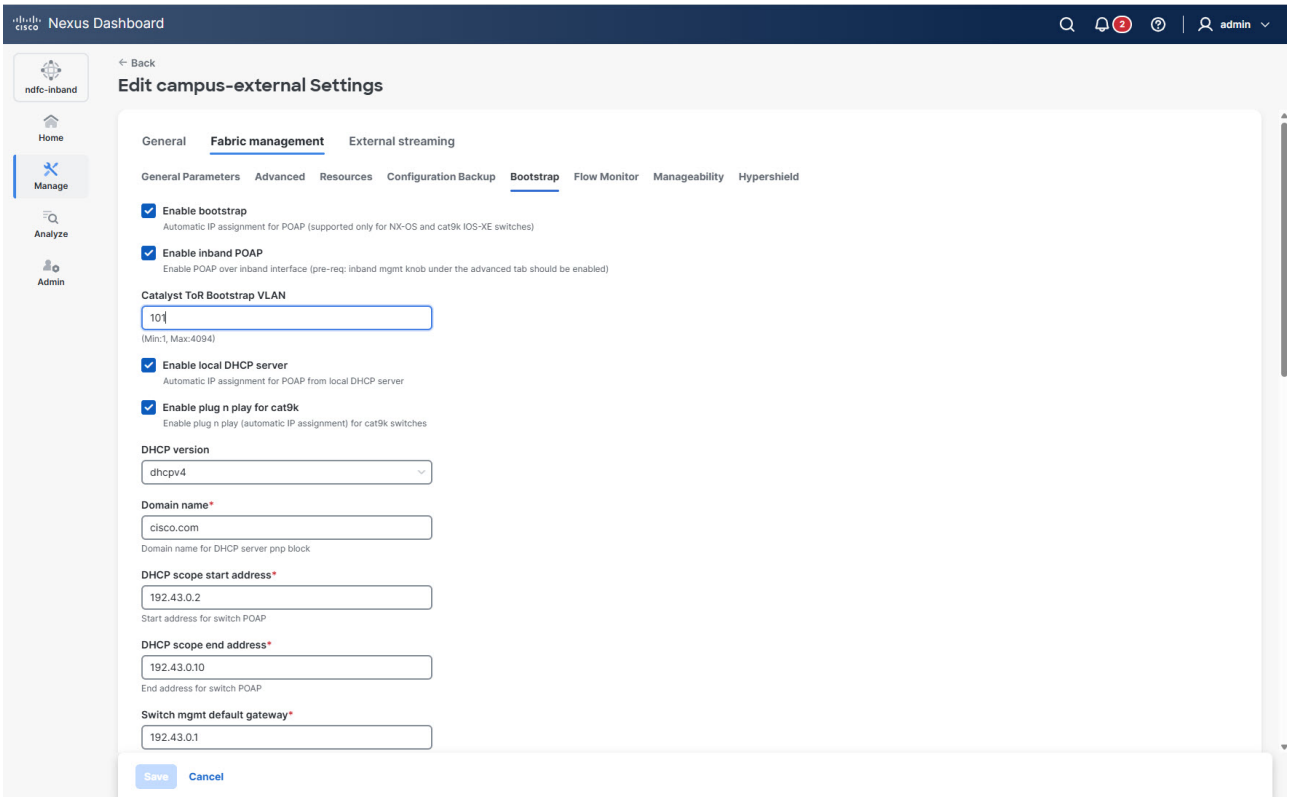

10. Click **Save**.

Add ToR switches using bootstrap

Follow these steps to add ToR switches using bootstrap.

1. Navigate to **Manage > Fabrics**.
2. Click the external fabric that you created.

The **Fabric Overview** page displays.

3. From the **Actions** drop-down list, choose **Add switches**.
4. Under **Switch addition mechanism**, click **Bootstrap**.

   The **Bootstrap** option appears only if you enable Bootstrap in the fabric settings.

5. In the **Admin password** field, provide the admin password.

   Nexus Dashboard detects the ToR switches in PnP-ready state and displays them in the **Switches to Bootstrap** table.

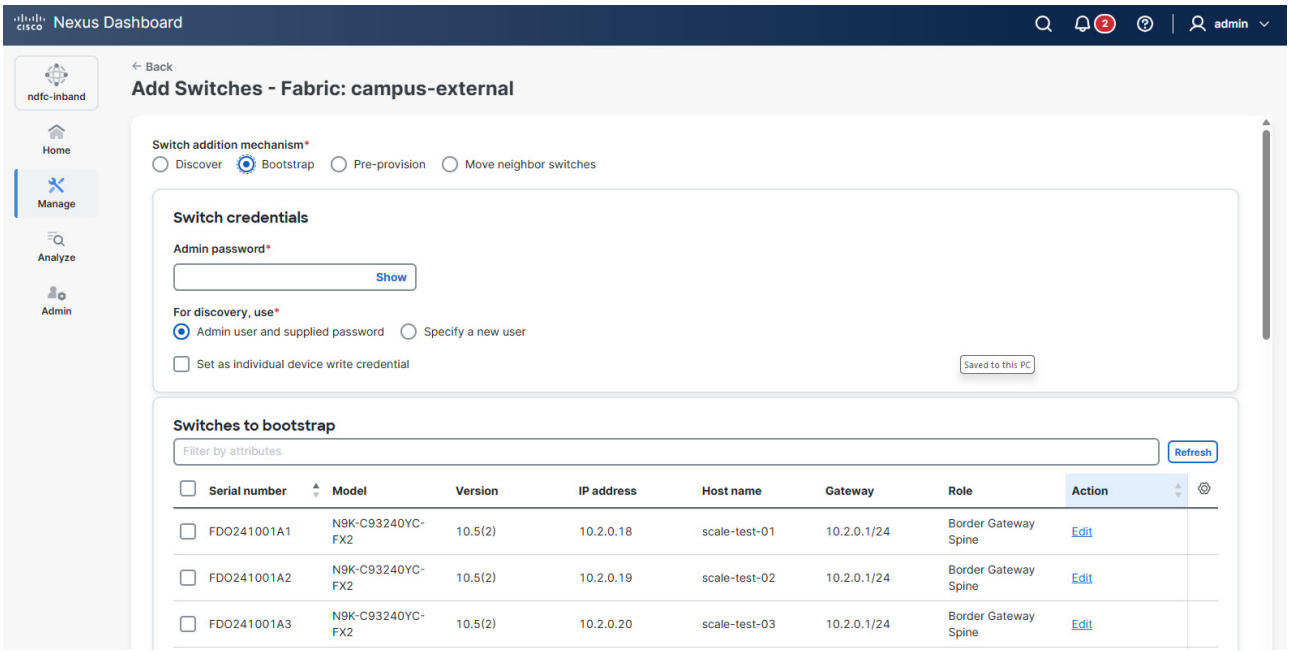

6. Choose the ToR switches you want to onboard.
7. For each chosen switch, click **Edit** under **Actions** column in the **Switches to Bootstrap** table.

   The **Edit bootstrap switch** dialog box appears.

8. In the **Edit bootstrap switch** dialog box, provide the following details.
   a. In the **Switch IP** field, assign an IP address that is outside the DHCP scope that you specified in the **External fabric bootstrap** settings.
   b. Set the **Switch Role** to **ToR**.
   c. Click **Save**.
9. Click **Import Selected Switches**.

Verify onboarding

After the import process completes, verify that Nexus Dashboard has successfully onboarded the ToRs.

1. In Nexus Dashboard, navigate to **Manage > Inventory** and confirm the following for each ToR in the external fabric.
   - **Role** is set to **ToR**.

- **Discovery Status** displays **Ok**.
- **Config Status** transitions to **In-Sync**.

2. On each ToR, confirm that Nexus Dashboard has pushed the management SVI and a default route. The default route points to the anycast gateway of the leaf's management network.

Troubleshoot ToR onboarding issues

If a Catalyst 9200-CX ToR switch fails to onboard, use the following checks to find the problem.

- Confirm that both the Campus fabric and the external fabric are configured for inband management. This onboarding flow requires inband management on both fabrics.
- If an external router connects the BGW spine and Nexus Dashboard, verify that either a BGP EVPN session or static routes to the relevant subnets are in place and are being distributed by the BGP on BGW spine.
- In Nexus Dashboard, verify that routes to the DHCP subnets and the loopback subnet used by the ToRs exist under **Admin > System Settings > General > Routes**. Add any missing routes.
- Verify that the pnp_vlan template is applied on the leaf interfaces facing the ToRs. Without it, the ToR cannot reach the management SVI.
- Verify that the Vrf_leaking_BGW template is applied on the BGW spines. Without it, DHCP traffic cannot reach the default VRF where the DHCP server resides.
- If the ToR does not appear in the **Switches to Bootstrap** table, check the DHCP server logs to confirm that the ToR is sending a DHCP DISCOVER and receiving a DHCP ACK.
- During onboarding, a Catalyst 9200-CX ToR switch may display an SSH session error in the UI until the switch finishes booting. This is expected behavior for ToR switches in the external fabric. After the switch is fully operational, the status changes to **Discovered**. No action is required.

Inband POAP management in External fabrics and Classic LAN fabrics

Inband POAP

Power On Auto Provisioning (POAP) automates the process of upgrading software images and installing configuration files on devices that are deployed on the network for the first time. POAP allows devices to bring up without performing any manual configuration.

When a POAP feature enabled device boots and does not find the startup configuration, the device enters POAP mode, locates a DHCP server, and bootstraps itself with its interface IP address, gateway, and DNS server IP addresses. The device obtains the IP address of a TFTP server and downloads a configuration script that enables the switch to download and install the appropriate software image and configuration file.

By using the POAP (Power On Auto Provisioning) feature of Nexus switches, Nexus Dashboard can automate the deployment of new datacenters reducing overall time and effort.

External Fabrics and Classic LAN fabrics support adding switches through POAP from inband interfaces.

The inband POAP is supported for all the roles for fabrics with External and Classic LAN templates.

Prerequisites

The following are the prerequisites for using inband poap:

- Configure the appropriate data network routes for reachability to the switch inband interfaces. Choose **Admin > System Settings > General**, locate the **Routes** area, then click **Edit** to enter route IP addresses in the **Data network routes** area.
- On the Nexus Dashboard Web UI, navigate to **Admin > System Settings > Fabric management > Advanced settings > Admin** and choose **Data** from the **LAN Device Management Connectivity** drop-down list to manage Data Center VXLAN EVPN fabrics through inband management, or an error message is displayed. If you choose **Data**, ensure that the required data service IP addresses are available in the Nexus Dashboard **External Service Pools** area (**Admin > System Settings > General > External pools**).

  

  When server settings changed from **Data** to **Management** or vice-versa, allow some time for syslog or poap functionalities to be online and ensure that the IP addresses in Cluster configuration are moved to the appropriate pool.

- Inband POAP on the Bootstrap tab is supported only when inband management is enabled on Advanced tab in the Fabric settings.

  Each subnet for the defined DHCP subnet scope that is mentioned in fabric settings must have a valid route for reverse traffic.

  Ensure that the DHCP relay functionality is set on intermediate routers.

Guidelines and limitations

- These are the guidelines and limitations for inband management:
  - Both inband and out-of-band switches in the same fabric is not supported.
  - When you add switches to fabric, ensure that the switches are not in maintenance mode.
- These are the guidelines and limitations for inband POAP:
  - Inband POAP is supported for NX-OS switches only.
  - You can enable inband POAP with Nexus Dashboard as a Local DHCP Server or on External DHCP Servers.
  - Inband POAP supports Multi Subnet scope.
  - Inband POAP requires the external router connected seed switches to have the following capabilities:
    - DHCP relay functionality
    - eBGP peering
- Enabling inband management and POAP is not supported in these situations:
  - When creating a new Classic LAN (Enhanced Classic LAN) fabric
  - When onboarding an existing Classic LAN (Enhanced) fabric
- However, enabling inband management and POAP is supported when onboarding an existing Classic LAN (Legacy) fabric.

Nexus Dashboard configures the SNMP username and password on a switch only when you import the switch through POAP, which creates the corresponding Point of Truth Instance (PTI). If you add a switch without POAP, Nexus Dashboard does not configure the SNMP credentials.

This can also occur after a switch that was onboarded through POAP undergoes a write erase followed by a reload from the switch CLI. The switch re-enters POAP mode and re-runs the POAP script, but may complete bootup without the SNMPv3 user that Nexus Dashboard uses for discovery. When this occurs, the switch shows as unreachable in Nexus Dashboard, and a manual rediscover fails with an SNMP unknown user/password error.

This behavior applies to all NX-OS switches and Catalyst switches. If AAA is enabled, this error does not occur during configuration.

To confirm the condition, run the show snmp user command on the switch. If no SNMP user is listed, the PTI for the SNMP credentials is missing. To restore it, configure the SNMP username and password manually on the switch.

Enable inband management and POAP on External fabrics and Classic LAN fabrics

For Classic LAN fabrics, enabling inband management and POAP is supported only when onboarding an existing Classic LAN (Legacy) fabric. For more information, see

> Inband management in External fabrics and Classic LAN fabrics and Guidelines and limitations.

To enable inband POAP on a fabric:

1. Navigate to **Manage > Fabrics**, then choose the appropriate fabric.
2. Click **Actions > Edit Fabric Settings**.
3. Click **Fabric Management**.
4. Click **Advanced**, then check the box in **Inband Mgmt**.
5. Click **Bootstrap**, then configure these settings:
   - Check the box in **Enable Bootstrap**.
   - Check the box in **Enable Local DHCP Server** and enter appropriate IP addresses in the required fields.

Add switches

To add or discover switches through inband POAP, follow these steps:

1. Pre-provision switches to a fabric
2. Add an interface
3. Add a policy to a fabric
4. Import switches using the bootstrap mechanism

Pre-provision switches to a fabric

To add switches to fabric:

1. Navigate to **Manage > Fabrics**, then click the appropriate fabric.

   The **Overview** window for that fabric appears.

2. Choose **Inventory > Switches** and click **Actions > Add Switches**.

   The **Add Switches** window appears.

3. Choose **Pre-provision** radio button.
4. Click **Actions** and add switches.

   You can add switches one at a time using the Add option or add multiple switches at the same time using the **Import** option.

   If you use the **Add** option, ensure you enter all the required details.

5. Choose a switch.
6. Enter the password in the **Admin password** field.
7. Click **Pre-provision**.

   The pre-provisioned switch is added.

For pre-provisioned switches, you can add dummy values for the serial number. After configuring the network successfully, you can change serial number with the appropriate number of the switch on the **Switches** tab. See the "Change Serial Number" section in "Performing Actions on Switches" in Working with Inventory in Your Nexus Dashboard LAN or IPFM Fabrics.

Import switches using the bootstrap mechanism

Ensure that you have pre-provisioned switches, added interface, and policy before importing the switches using the bootstrap mechanism.

To import switches using the bootstrap mechanism:

1. Navigate to **Manage > Fabrics**, then click the appropriate fabric.

   The **Overview** window for that fabric appears.

2. Choose **Inventory > Switches** and click **Actions > Add Switches**.

   The **Add Switches** window appears.

   You can view the existing added switches in the **Switches to Bootstrap** area.

3. Choose **Bootstrap (POAP)** radio button and enter a password in **Admin password** field.
4. Choose the required switches and click **Import Selected Switches** to bootstrap switches.

Add an interface

Add an interface to configure the interface IP addresses on the required switch.

Before you begin:

Ensure that you have added the required configurations on the switches such as IP addresses and static routes.

To add an interface:

1. Navigate to **Manage > Fabrics**, then click the appropriate fabric.

   The **Overview** window for that fabric appears.

2. Choose **Connectivity > Interfaces** and click **Actions > Create interface**.

   The **Create New Interface** window appears.

3. Choose **Ethernet** from the **Type** drop-down list.
4. Choose the required switch from the **Select a device** drop-down list and enter a name in the **Interface name** field.
5. Select the **int_routed_host** policy from the list of policies .
6. Enter the required configuration details in the **Interface IP** and **IP Netmask Length** fields.
7. Enter appropriate details in all the mandatory fields and ensure that you check **Enable Interface**

check box and then click **Save**.

Add a policy to a fabric

You can add a freeform policy to define external routes in the switch. To add a policy:

1. Navigate to **Manage > Fabrics**, then click the appropriate fabric.

   The **Overview** window for that fabric appears.

2. Choose **Configuration Policies > Policies** and click **Actions > Add policy**.
3. Choose an appropriate switch from the **Switch List** drop-down list and click **Choose Template**.
4. On **Select Policy Template** window, select the **switch_freeform** template and click **Select**.

   The **switch_freeform** policy type allows you to add configurations in CLI format.

Recalculate and deploy configurations on a switch

To push pending configurations to switches:

1. Navigate to **Manage > Fabrics**, then click the appropriate fabric.

   The **Overview** window for that fabric appears.

2. Choose **Inventory > Switches**.

   You can see that the **Config status** column displays **Pending** status.

3. Choose **Actions > Recalculate and deploy**.

   The **Deploy Configuration** window appears. It displays the configuration status of the switches. You can also view the pending configurations by clicking the respective link in the **Pending Config** column.

   The **Pending Config** window appears. The **Pending Config** tab on this window displays the pending configurations on the switch. The **Side-by-Side Comparison** tab displays the running configuration and the expected configuration side-by-side.

4. Close the **Pending Config** window. When all the pending configuration is complete, the **Config status** column displays **In-Sync**.

Inband management and inband POAP in Data Center VXLAN EVPN fabrics

You can manage switches with inband connectivity and inband POAP for Data Center VXLAN EVPN fabrics. For inband management, the Loopback0 interface of the devices is used in the Fabric Settings.

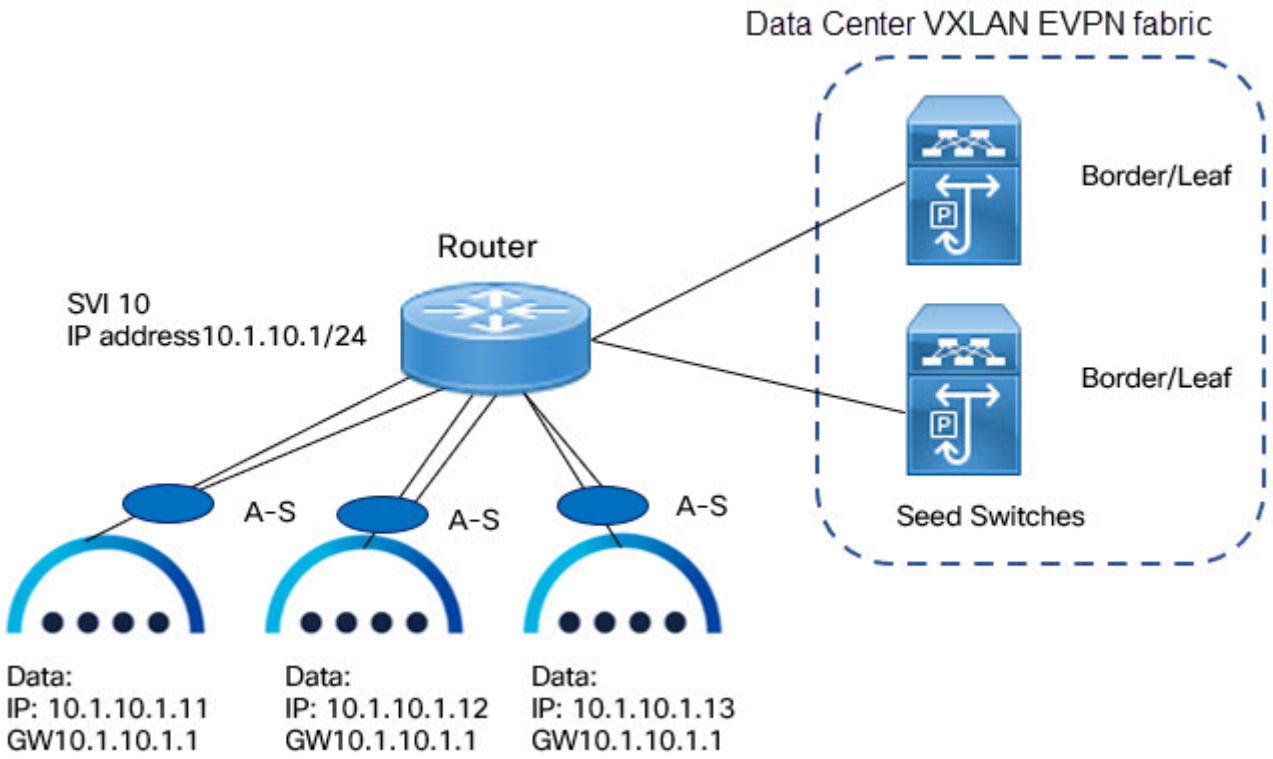

If you want POAP Layer-3 adjacency to the switches, you must add the Nexus Dashboard node IP address as DHCP Relay address:

1. Navigate to **Admin > System Settings**.
2. Click **Fabric management**.
3. In the **Advanced settings** area, click **Admin**.
4. Note the setting in the **LAN Device Management Connectivity** field.
    - If the default value is set to **Management** in the **LAN Device Management Connectivity** field, then the DHCP Relay address must be set to the management interface IP (bond1br) in all Nexus Dashboard nodes.
    - If the default value is set to **Data** in the **LAN Device Management Connectivity** field, then the DHCP Relay address must be set to the data interface IP (bond0br) in all Nexus Dashboard nodes.

You can add switches with inband management enabled for Data Center VXLAN EVPN fabrics either in Greenfield or brownfield deployment with inband POAP or pre-provision and inband POAP.

- For Brownfield deployment, check **Preserve Config** check box.
- For Greenfield deployment, uncheck **Preserve Config** check box.

Importing switches with the **Preserve Config** option enabled will not work under the following conditions:

- You have more than 100 VRFs and 500 networks, and
- You are running on either of the following NX-OS images:
  - 7.0(3)I7(9)
  - 7.0(3)I7(10)

In these conditions, the switches will stay in migration mode. Upgrade to NX-OS release 9.x or later to resolve this issue.

The seed switches connect the external routers, and it provides management connectivity to the other switches in the fabric. Switches connected to external routers to provide connectivity to the fabrics are termed as seed switches. The interfaces on the seed switches which connects to the external routers are termed as bootstrap interfaces.

Prerequisites for inband and out-of-band management

1. Navigate to **Admin > System Settings**.
2. Click **Fabric management**.
3. In the **Advanced settings** area, click **Admin**.
4. In the **LAN Device Management Connectivity** field, choose **Data** to manage Data Center VXLAN EVPN fabrics through inband management.

   If you choose **Data**, ensure that the required persistent data IPs are available in the **External pools** area under **Admin > System Settings > General**.

When server settings changed from **Data** to **Management** or vice-versa, allow some time for syslog or poap functionalities to be online and ensure that the IP addresses in Cluster configuration are moved to the appropriate pool.

This server setting is required for both inband and out-of-band connectivity. Configure below static routes over data interface in Cisco Nexus Dashboard:

Enter static routes IP address required for external route and route over data interface in Cisco Nexus Dashboard.

Inband POAP requires the external router IP address connected to the seed switches to have the following capabilities:

- Routes for External router
- Route for Routing Loopback subnet range for Data Center VXLAN EVPN fabric
- Route for Underlay Routing subnet range for Data Center VXLAN EVPN fabric

Inband POAP requires the external router connected seed switches to have the following capabilities:

- DHCP relay functionality
- eBGP peering

To add switches for inband management and inband POAP, see the section "Discovering New Switches" in Working with Inventory in Your Nexus Dashboard LAN or IPFM Fabrics.

Guidelines and limitations

The following are the guidelines and limitations for inband management:

- Inband management is not supported when discovering or onboarding brownfield Data Center VXLAN EVPN fabrics.
- Ensure that the **Inband Management** is enabled for the inband interface. Both inband and out-of-band switches for a same fabric is not supported.
- It is supported only for IPv4 underlay and OSPF routing protocol.
- You can change switch management from inband to out-of-band and conversely after creating a fabric.
- For the inband managed switches, the following roles are supported:
  - Spine
  - Leaf
  - Border
  - Border Spine
  - Border Gateway
  - Border Gateway Spine
- Inband management is supported for both numbered and unnumbered fabric interface numbering
- Ensure that the same role switches are assigned as seed switches. If spine role switch is assigned as a seed switch, all the spine role switches in that fabric must be assigned as seed switches. It is recommended to assign switch as seed switches.
- When you add switches to fabric, ensure that the switches are not in maintenance mode.
- You can add switches in Brownfield deployment (check **Preserve Config** check box) only when the fabric is created. To add more switches, use inband POAP with import switches option.
- Set **vPC Peer Keep Alive** option to loopback if the vPC switches mgmt0 interfaces are not configured.

The following are the guidelines and limitations for inband POAP:

- Inband POAP for a fabric can be enabled only if **Inband Management** is enabled.
- Inband POAP requires the fabric or core facing interfaces to be cabled consistently for seed switches and spine switches.
- All spine switches in fabric must use same set of fabric interface numbers.
- If a fabric has set of leaf switches which are seed switches, then the switches must use same fabric interface number.
- The seed switches must have eBGP peering with the external router. Therefore, the external router must have the required eBGP route peering capabilities and display the configuration for External router for DHCP relay and Static routes configured for the Subnets used in Data Center VXLAN EVPN fabrics.

- DHCP relay must be configured on external routers interface which connects the seed switch in inband interfaces. Ensure that the DHCP relay destination configured is same for all cluster node data interface on Cisco Nexus Dashboard.
- DHCP server can be internal Nexus Dashboard or the external server.

Enable inband POAP on Data Center VXLAN EVPN fabrics

To enable inband POAP on Data Center VXLAN EVPN fabrics, perform the following steps:

1. Navigate to **Manage > Fabrics**, then click the appropriate fabric.

   The **Overview** window for that fabric appears.

2. Click **Actions > Edit Fabric Settings**.
3. Click **Fabric Management**.
4. Click **Manageability** and check the **Inband Management** box.
5. Click **Bootstrap**, then configure these settings:
   a. Check the **Enable Bootstrap** check box.
   b. Check the **Enable Local DHCP Server** checkbox to assign Nexus Dashboard as DHCP Server and enter the DHCP scopes for all the fabric seed switches bootstrap interfaces.

      If you choose **Enable Local DHCP Server**, and choose unnumbered in Fabric Interface Numbering drop-down list in the General Parameters tab, add details for:

      - Bootstrap Seed Switch Loopback Interface ID
      - Switch Loopback DHCP Scope Start Address
      - Switch Loopback DHCP Scope End Address

   c. Check the **External DHCP Server IP Addresses** check box to provide connectivity to Nexus Dashboard from the external router.

      If you choose **External DHCP Server IP Addresses**, you can add a maximum of three IPv4 addresses with a comma separated list.

      

      To have eBGP peering between seeds and an external router, add bootstrap seed switch loopback interface IP address, this IP must be a subset of the loopback id range.

   d. Enter Seed Switch interface in **Seed Switch Fabric Interfaces** text field.
   e. Enter Spine Switch interface in **Spine Switch Fabric Interfaces** text field.

      

      If the Spine switches are the seed switches, then the lists must be consistent in **Seed Switch Fabric Interfaces** text field.

6. For fabrics with unnumbered interface, do the following:
   a. Click **General Parameters**, then choose **unnumbered** from **Fabric Interface Numbering** drop-down list.

b. Click **Bootstrap** and configure these settings:

- **Bootstrap Seed Switch Loopback Interface ID** the loopback ID is the default router IP for the fabric. This loopback ID must not overlap with any of the existing fabric loopback IDs.
- **Switch Loopback DHCP Scope Start Address** this IP address is start address of the DHCP pool of the routing loopback addresses range to assign to the bootstrapping switch. This IP address must not overlap with any of the existing IP addresses of **Underlay Routing Loopback IP Range**.
- **Switch Loopback DHCP Scope End Address** is the end address of the DHCP pool.

Import switches in a Brownfield deployment

Before you begin:

Make sure that you follow prerequisites procedure before adding switches.

1. Create a fabric using the **Data Center VXLAN EVPN** template.

   See Creating Fabrics and Fabric Groups.

   Ensure that you add switches in the order of Seed switches, Spine switches, and other switches. You can add spine switches as the seed switches.

2. Add the switches to the fabric with the **Preserve Config** check box.
3. Enter **hostname**, **Role**, enable **Seed Switch**, and enter appropriate IP address.
4. Enter the IP addresses for all the seed switches, click **Import Selected Switches** to add them to the fabric.
5. Navigate to **Policy** tab, click **Actions > Add policy**. Choose **ext_bgp_neighbor** policy so the seed switches establish eBGP peering. Enter the required details, and click **Save**.
6. Assign the appropriate switch roles.

   For more instructions, see the section " Adding Switches Using Bootstrap Mechanism" in Working with Inventory in Your Nexus Dashboard LAN or IPFM Fabrics.

Pre-provision switches through inband POAP

1. Navigate to **Manage > Fabrics**, then click the appropriate fabric.

   The **Overview** window for that fabric appears.

2. Choose **Inventory > Switches** and click **Actions > Add Switches**.

   The **Add Switches** window appears.

3. Choose **Pre-provision** radio button.
4. On **Switches to Pre-provision** table, click **Actions> Add**.

   The **Pre-provision a switch** window appears.

5. Enter appropriate details such as Serial Number, Model, IP Address, and click **Add**.

6. Enter single switch at once and enter the required information. If you have multiple switches.
7. Click **Import Switches to Fabric** to add switches.

Add a policy for Data Center VXLAN EVPN fabrics

1. Navigate to **Manage > Fabrics**, then click the appropriate fabric.

   The **Overview** window for that fabric appears.

2. Choose **Configuration Policies > Policies** and click **Actions > Add policy**.
3. Choose the appropriate switch from the **Switch** window and click **Choose Template**.
4. Choose **ext_bgp_neighbor policy** and click **Select**.

   The **Create Policy** window appears.

5. Click **Actions > Add policy**.

   The **Create Policy** window appears.

6. Enter the appropriate details in the window and click **Save**.
7. On **Fabric Overview** window, click **Actions > Recalculate and deploy**.

Change the fabric management mode

You can change the fabric from out-of-band to inband management and vice versa.

1. To change fabric management from out-of-band to inband management:
   a. Ensure that you followed the prerequisites for inband management.

      See Prerequisites for inband and out-of-band management.

   b. Navigate to **Manage > Fabrics**, then choose the appropriate fabric.

      The **Overview** window for that fabric appears.

   c. Click **Actions > Edit Fabric Settings**.
   d. Click **Fabric Management**.
   e. Click **Advanced**.
   f. Click the box for **Inband Mgmt** and click **Save**.
   g. Choose **Inventory > Switches**.
   h. Locate the switches that display **Migration** in the **Mode** column.
   i. Choose those switches, then click **Actions > Recalculate and deploy.**
      - The discovery IP address of the switches changes to the BGP routing loopback IP.
      - The discovery VRF displays default and discovery interface is updated to BGP routing loopback interface.
      - An error is generated displaying switch discovery is pending:

The discovery modes for switches have been updated but, discovery may not have completed. Please check to make sure Discovery Status is Ok and retry Recalculate & Deploy.

j. Click **OK**.

k. Ensure that the **Discovery Status** column displays **Ok** as the status, then click **Actions > Recalculate and deploy**.

2. To change fabric management from inband management to out-of-band:

a. Ensure that you followed the prerequisites for out-of-band management.

See Prerequisites for inband and out-of-band management.

b. Configure out-of-band IP addresses on the switch.

These IP addresses must be reachable from the Nexus Dashboard data or management interface.

c. Navigate to **Manage > Fabrics**, then choose the appropriate external/inter-fabric connectivity fabric.

The **Overview** window for that fabric appears.

d. Click **Actions > Edit Fabric Settings**.

e. Click **Fabric Management**.

f. Click **Advanced**.

g. Remove the check from the box in the **Inband Mgmt** field and click **Save**.

h. Choose **Inventory > Switches**.

i. Locate the switches that display **Migration** in the **Mode** column.

j. Choose those switches, then click **Actions > Recalculate and deploy.**

- The discovery IP address of the switches will be changed to the mgmt0 IP.
- The discovery VRF displays management and discovery interface will be updated to mgmt0. An error is generated displaying switch discovery is pending: The discovery modes for switches have been updated but, discovery may not have completed. Please check to make sure that Discovery Status is Ok and retry Recalculate & Deploy.

k. Click **OK**.

l. Ensure that the **Discovery Status** column displays **Ok** as the status, then click **Actions > Recalculate and deploy**.

Secure POAP

When you import switches through bootstrap or POAP in Nexus Dashboard, it locates a DHCP protocol and bootstraps with interface IP address, gateway, DNS server IP address, and POAP script path.

POAP uses an HTTPS server which is a secure protocol to encrypt traffic and validate Nexus Dashboard for network connection. You must configure Bench Router (BR) to host root Certificate Authority (CA), which is a signed server certificate of POAP server that is hosted on Nexus Dashboard. In the DHCP response, BR is identified which acts as a trust for a new switch.

Secure POAP is not supported for inband connectivity with bench routers.

See the "CA Certificates" and "Bootstrap Certificates" sections in Managing Certificates in your Nexus Dashboard to upload the appropriate certificates on Nexus Dashboard.

Prerequisites for secure POAP

- Secure POAP is supported from Cisco NX-OS 9000 Release 10.2.3 or higher version switches.
- Navigate to **Admin > System Settings**, click **Edit** in the **Switch bootstrap** area, then choose **https** or **http&https** from the drop-down list for **Bootstrap Script Download Protocol** field.
- For the **http** or **http&https** option, you must enter the IP address of the bench router (BR), port number, and name for certificate bundle in **Bench Router URL with port and certificate file name** filed. Ensure that the certificates are uploaded on Nexus Dashboard server for values to autopopulate in this field.
- By default, for the **http** or **http&https** options, the **Bench router URL with port** field in the **Switch bootstrap** window will be blank. After you install the Root CA Certificate bundle on Bench routers, this field will be autopopulated.

  If these fields are autopopulated, with default port number 29151 and URL https://*ip-address-_BR:_port/CA-Bundle-filename* (for example, https://10.10.10.11:29151/PoapCACertBundle.pem), you must configure this URL before you install the BR with the Root CA certificate bundle.

- Make sure that the Fabric is in managed mode before configuring the BR.
- Ensure that you configure the DHCP option if the DHCP server is used.
- You must upload CA signed POAP server certificate on Nexus Dashboard and upload the corresponding CA certificate bundle for the BR. On Nexus Dashboard, navigate to **Admin > Certificate Management > Fabric Certificates** to upload the relevant certificates.

Inband PnP in Campus VXLAN EVPN fabrics

About Inband PnP for Cisco Catalyst switches

Beginning with release 4.2.1, Nexus Dashboard provides support for the Inband Plug and Play (PnP) feature, which streamlines and automates onboarding of Cisco Catalyst 9000 Series switches (running OS version 16.3.1 and later) to deliver a zero-touch deployment experience. The inband PnP feature uses the switch's regular network ports to quickly connect, set up, and manage the switch without requiring additional steps.

Key aspects of Inband PnP

- Local DHCP server requirement—To create fabric and onboard devices using inband PnP, you must enable a local DHCP server. The local DHCP server assigns network parameters and provides the PnP server address to new switches during initial deployment, ensuring a successful onboarding process.
- Seed switches and topology—Seed switches act as foundational nodes in the network topology and play a crucial role in onboarding leaf switches. They enable Nexus Dashboard to automatically discover and integrate new devices into the fabric.
- Automatic subnet generation—Nexus Dashboard automatically generates subnets for onboarded switches. This automation reduces manual configuration steps and helps prevent IP conflicts.
- Pre-provisioning for efficiency—Pre-provisioning automatically adds switches to the fabric and bypasses manual verification steps.

Prerequisites for Inband PnP

Before you enable inband PnP within Nexus Dashboard, configure the network infrastructure on the switches that connect to the seed switches. This ensures that the seed switches can reach the DHCP server.

- Configure the switches connecting to the seed switches to ensure reachability. This is critical when the Campus fabric accesses the DHCP server through a DC fabric or using an external router.
- Ensure Cisco Nexus 9000 (N9k) borders/leafs and DC BGW switches are already added to the DC fabric (using OOB POAP or standard workflow) during greenfield deployment.
- Ensure that you complete the following specific switch configurations.
  - Configure N9k DC border switch
  - Configure DC BGW spines
  - Configure Campus fabric
- Upload required SSL certificates (CA-signed recommended, though self-signed are supported) to Nexus Dashboard before starting the import process.

Guidelines and Limitations

- Inband PnP is not supported for Cisco Catalyst 9000 Virtual (C9KV) switches.

- The inband PnP option is available only for Campus VXLAN EVPN fabrics.
- You must set the MTU value to 9198 for Catalyst fabrics. This setting is mandatory.
- This feature currently supports only DHCP for IPv4 addresses. IPv6 is not supported for inband PnP at this time.
- The seed switch must be an N9k device assigned to the role of BGW spine.
- When configuring Catalyst 9000 series (Cat9K) switches, you must choose either the leaf or border role; these are the only roles supported for Inband PnP.
- A local DHCP server must be configured within the fabric management system to provide IP addresses during the onboarding phase.

Configure N9k DC border switch

Follow these steps to configure the N9k DC border switch.

1. On N9K DC border, configure the DC-border to external router interfaces with IP address as regular routed interfaces using interface workflow.
2. Add dc_border_inband_setup policy with Nexus Dashboard data subnet and external router ASN/IP/hostName as input.

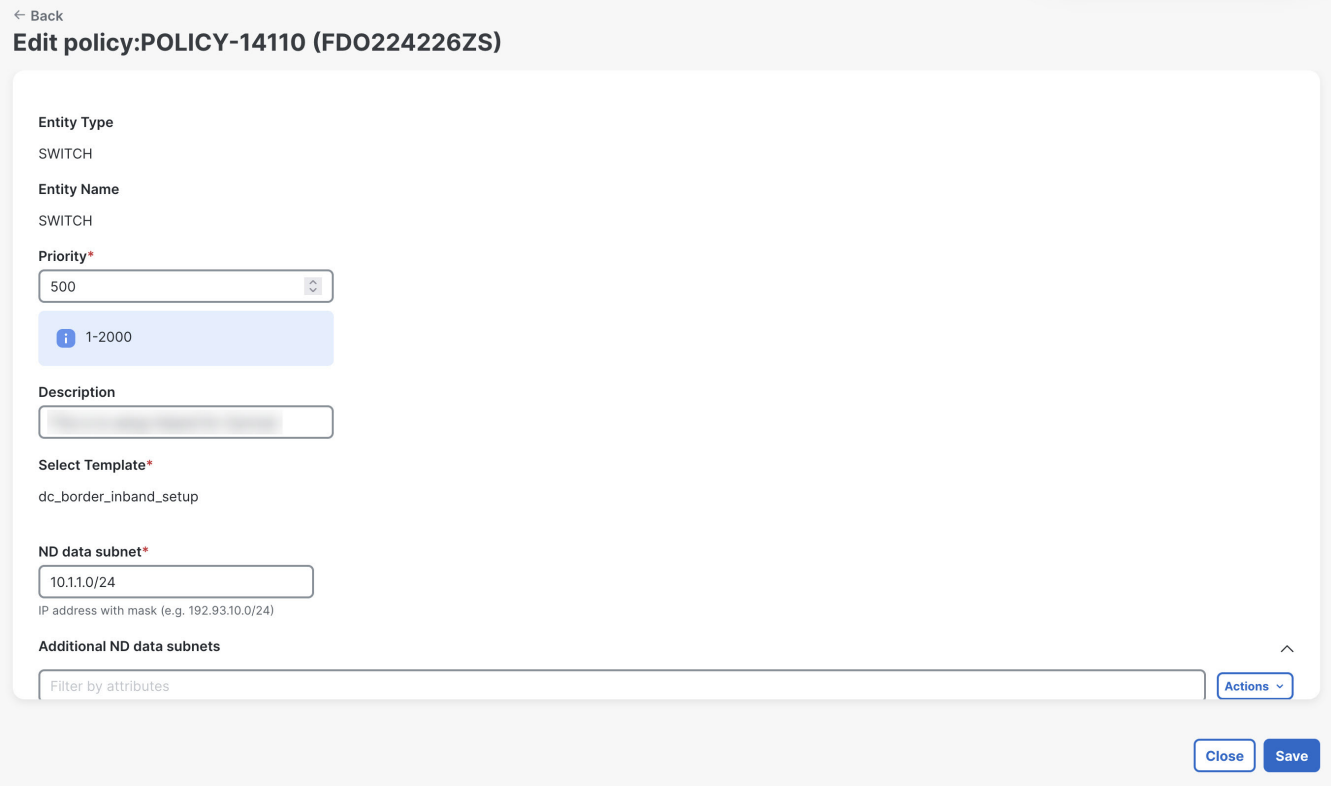

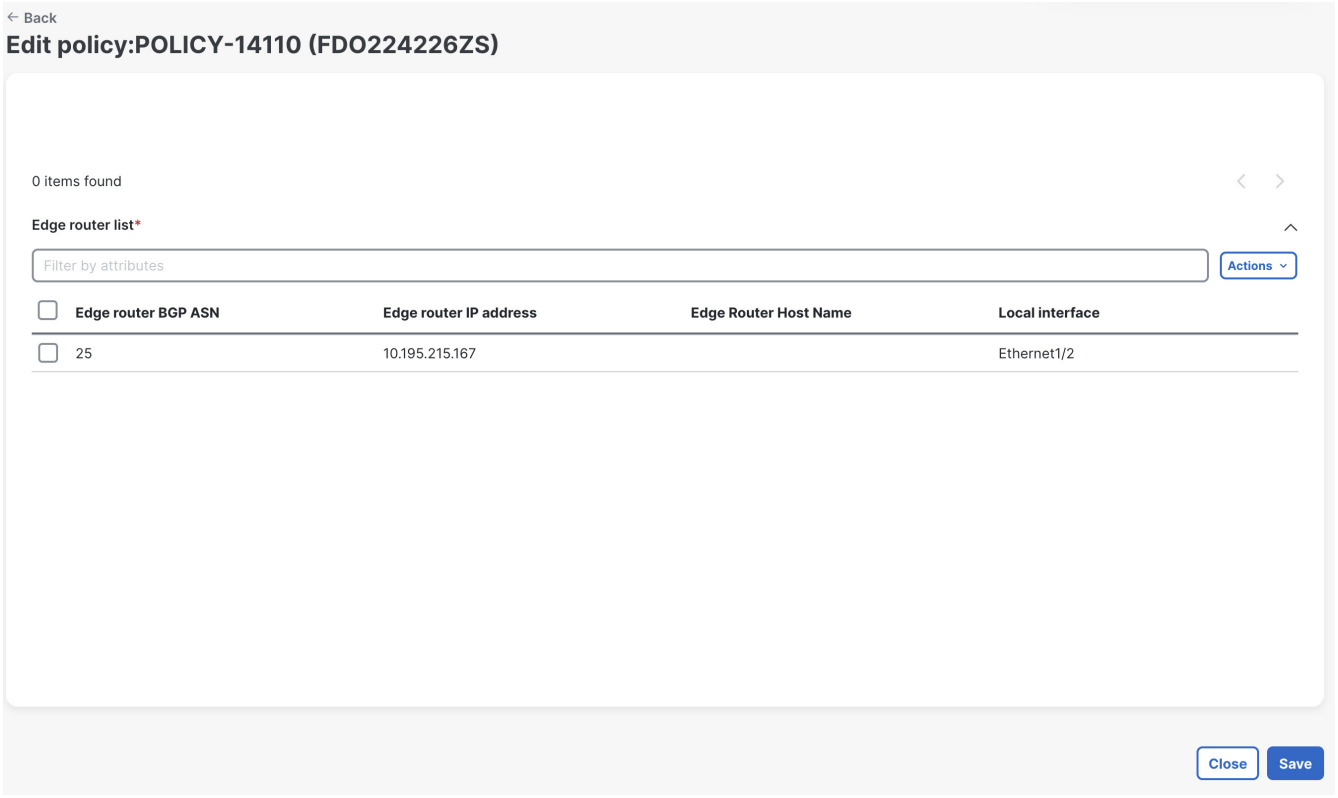

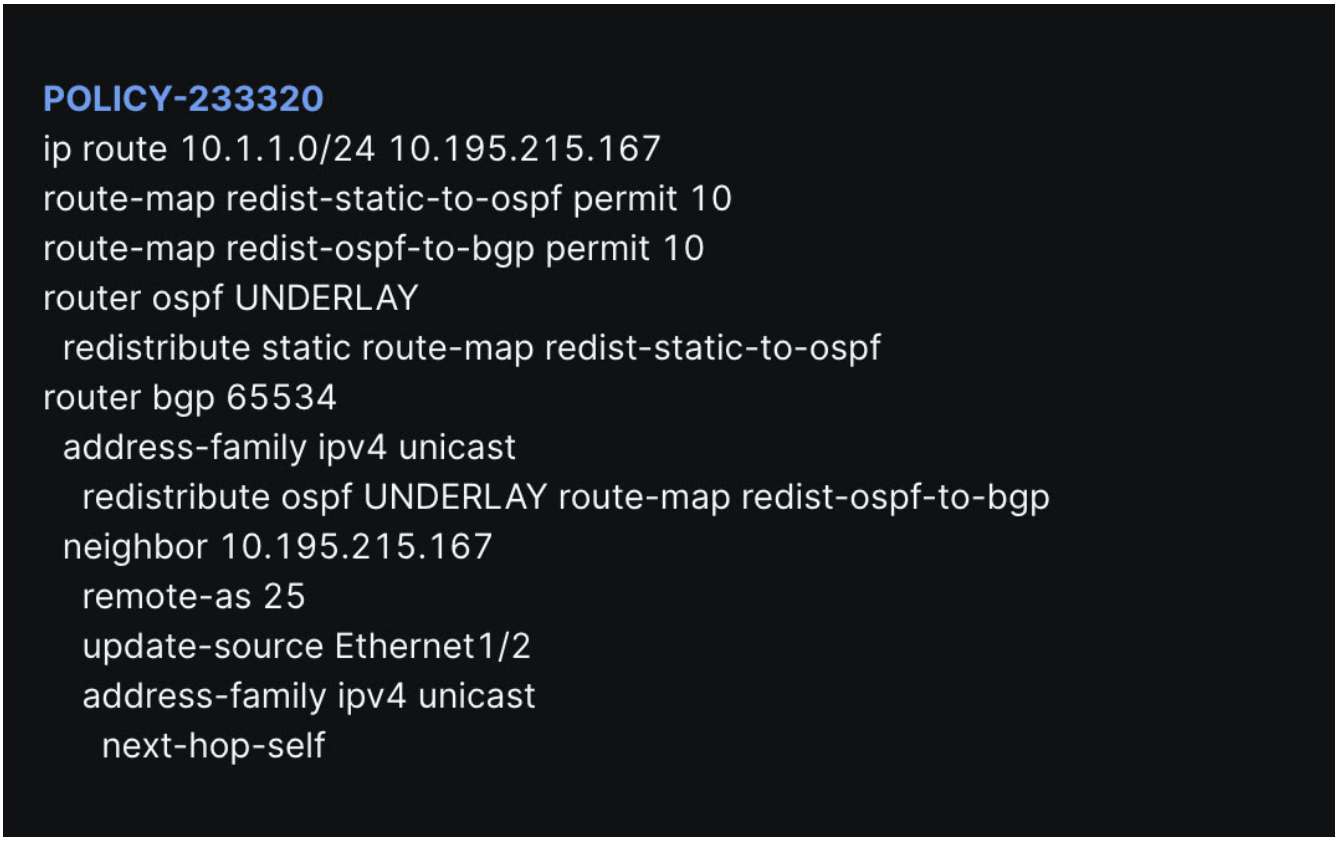

Configure DC BGW spines

Follow these steps to configure the DC BGW spines.

1. On DC BGW spines, configure the DC-BGW spine to Campus BGW spine facing interfaces as regular routed interfaces using interface workflow.
2. Configure the IP address within the DHCP subnet scope in the Campus fabric's fabric settings.

3. Add ip dhcp relay address 192.93.0.119 in the interface freeform field. Here 192.93.0.119 is the NDFC DHCP server IP address.
4. Add dc_border_gateway_inband_setup policy in each DC BGWS, with Nexus Dashboard data subnet and DCI interfaces and their IP addresses as input.

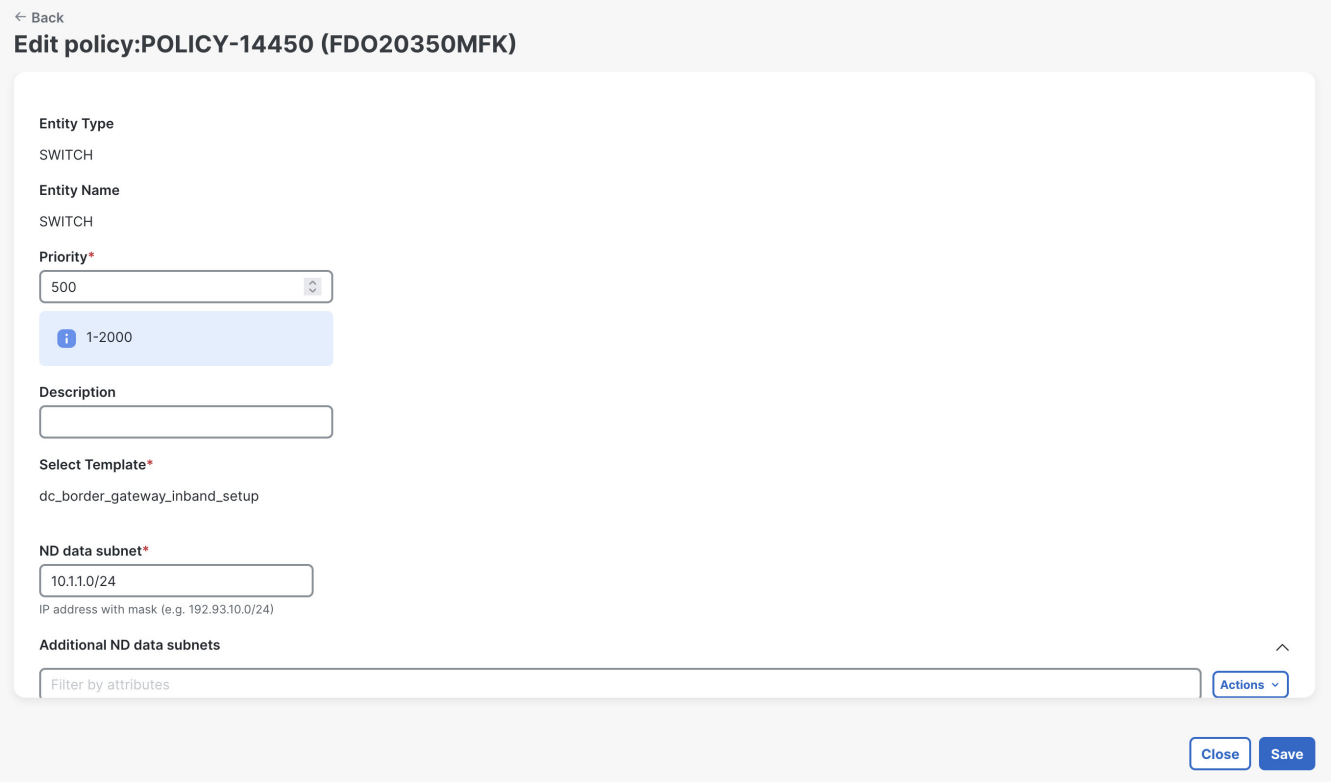

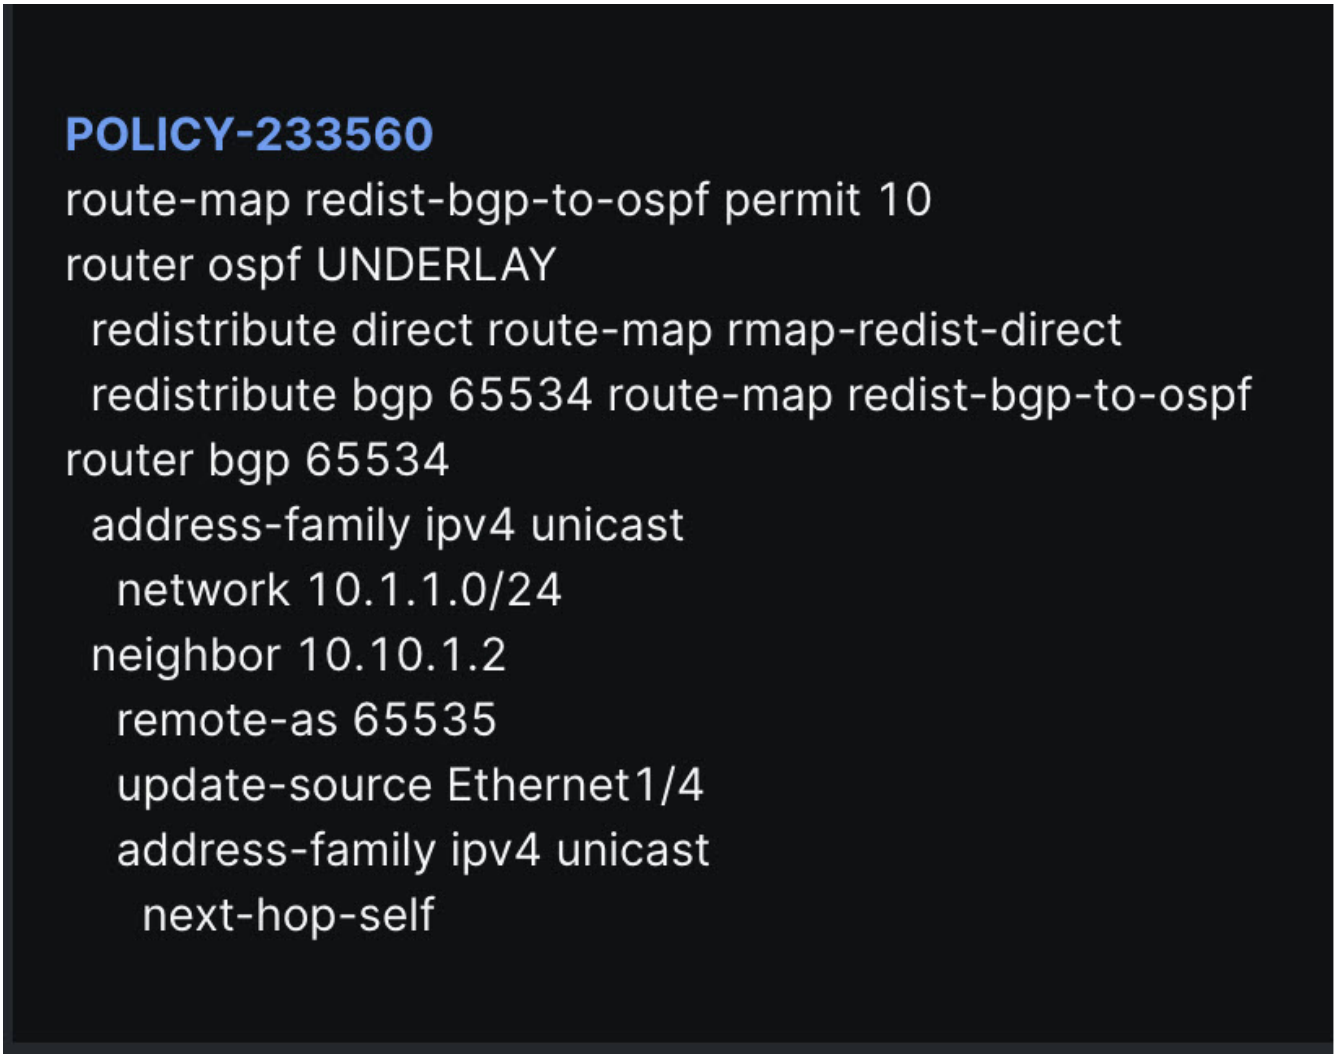

Configure Campus fabric

Follow these steps to configure the Campus fabric.

1. Create Campus VXLAN fabric, with inband-enabled and onboard Campus BGWS using inband POAP.
2. Add the Campus fabric to Multi-Site Domain (MSD).
3. Perform MSD R&D and build multisite underlay and overlay IFCs.

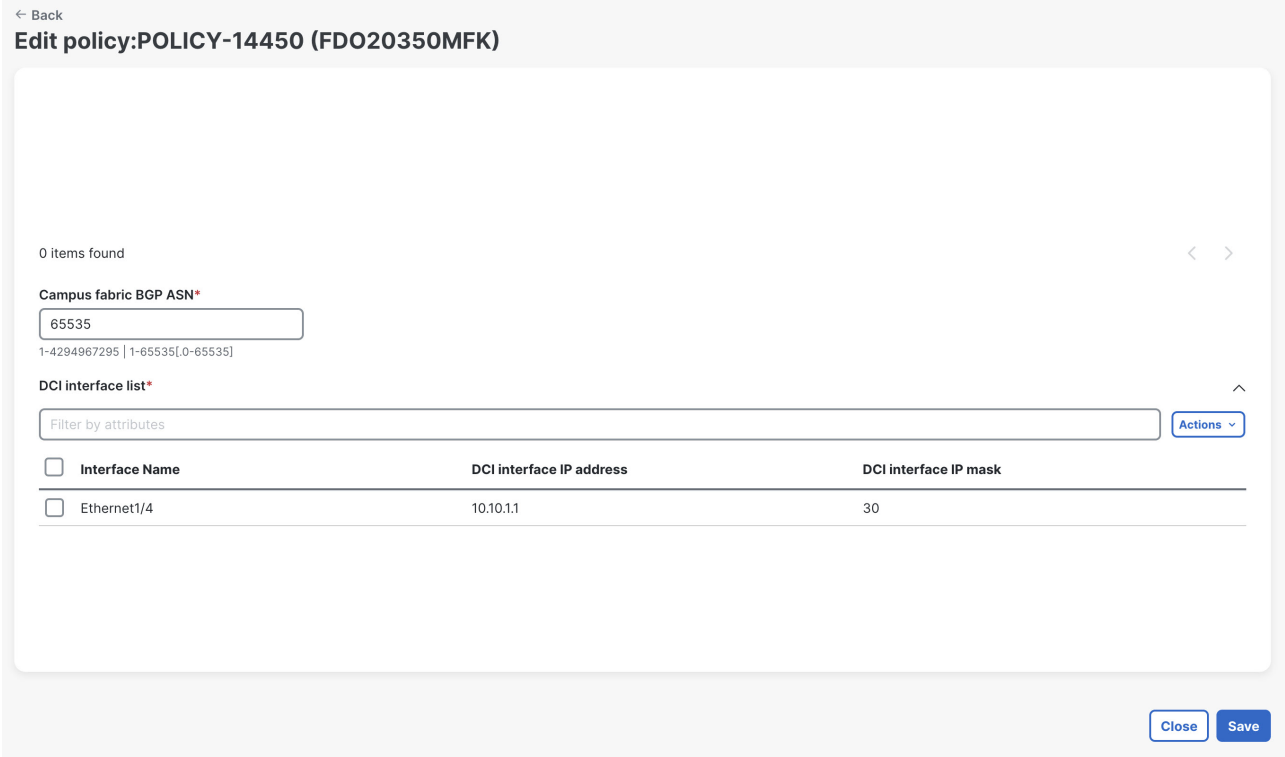

4. Onboard Catalyst leaf using inband PnP.

For detailed instructions on associating Campus fabrics under an MSD fabric, see Editing Fabric Settings for Fabric Groups.

Enable Inband PnP

Before you begin

- You can assign IP addresses to devices using either a local DHCP server (configured within your fabric management system) or an external DHCP server.
- This feature currently supports only DHCP for IPv4 addresses. IPv6 is not supported for inband PnP at this time.
- The inband PnP option is available only for Campus VXLAN EVPN fabrics.

Follow these steps to enable inband PnP.

1. Create a Campus VXLAN EVPN fabric.

   For more information, see Creating Fabrics and Fabric Groups.

2. Navigate to the **Fabrics** page.

   Go to Manage > Fabrics.

3. Choose the appropriate fabric.
4. From the **Actions** drop-down list, choose **Edit fabric settings**.

   The **Edit *fabric_name* Settings** page displays.

5. In the **Edit *fabric_name* Settings** page, click the **Fabric Management** tab.

6. Click the **Bootstrap** tab.

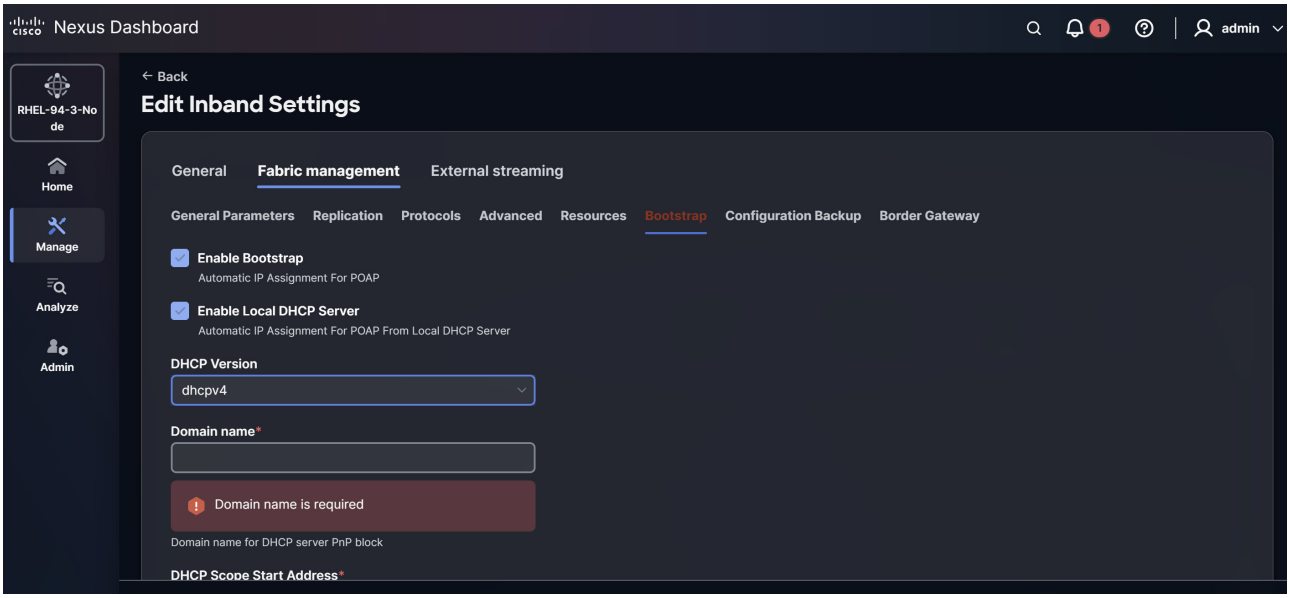

7. Check the **Enable Bootstrap** and **Enable Local DHCP Server** check boxes to activate the necessary services for inband PnP.
8. Provide the DHCP and management details.

Field	Description
DHCP Version	Choose the DHCP version. Only IPv4 is supported in this release.
Domain name	Specify the domain name that should be assigned to devices using DHCP.
DHCP Scope Start Address	Enter the starting IP address of the DHCP scope.
DHCP Scope End Address	Enter the ending IP address of the DHCP scope.
Switch Mgmt Default Gateway	Specify the default gateway IP address for the switch management network.
Switch Mgmt IP Subnet Prefix	Define the subnet mask in prefix notation. Set the prefix between 8 and 30. By default, Nexus Dashboard assigns 24.
Bootstrap Freeform Config (IOS-XE)	Enter any additional IOS-XE configuration CLIs to be applied during bootstrap.
Bootstrap Freeform Config (NXOS)	Enter any additional NXOS configuration CLIs to be applied during bootstrap.

9. Check the **Inband Management** check box to enable inband communication for device onboarding and management.
10. Check the **Restrict inband PnP DHCP to connected Catalyst** check box to limit the inband DHCP scopes to only the Catalyst switches that CDP discovers on the BGW seed interfaces. If a scope has no discovered serial numbers, Nexus Dashboard denies all discovery requests for that scope.

This option prevents loops that can occur when a ToR switch boots up early, before Nexus Dashboard onboards the leaf switches.

11. In the **Seed Switch Fabric Interfaces** field, provide the interfaces on your seed switch that connect to the downstream Catalyst switches.

The seed switch must be a Cisco N9000 device with the role border gateway spine.

12. Click **Save** to apply and activate the inband PnP configuration.

Import switches using bootstrap method

Before you begin

- Upload the required SSL certificates for your switches. You can use either CA-signed or self-signed certificates. Although Nexus Dashboard does not mandate CA-signed certificates, it is recommended to use CA-signed certificates for enhanced security and compliance with best practices.
- Verify that the local DHCP server is enabled and correctly configured to provide IP addresses to the switches during onboarding.
- Confirm that all switches are powered on, connected to the network, and in a PnP-ready state.

Follow these steps to import switches.

1. Navigate to the **Inventory** page.

   Go to Manage > Inventory.

2. From the **Actions** drop-down list, choose **Add switches**.

   The **Add switches** page opens.

3. Click **Choose Fabric**.

   The **Select Fabric** dialog box opens.

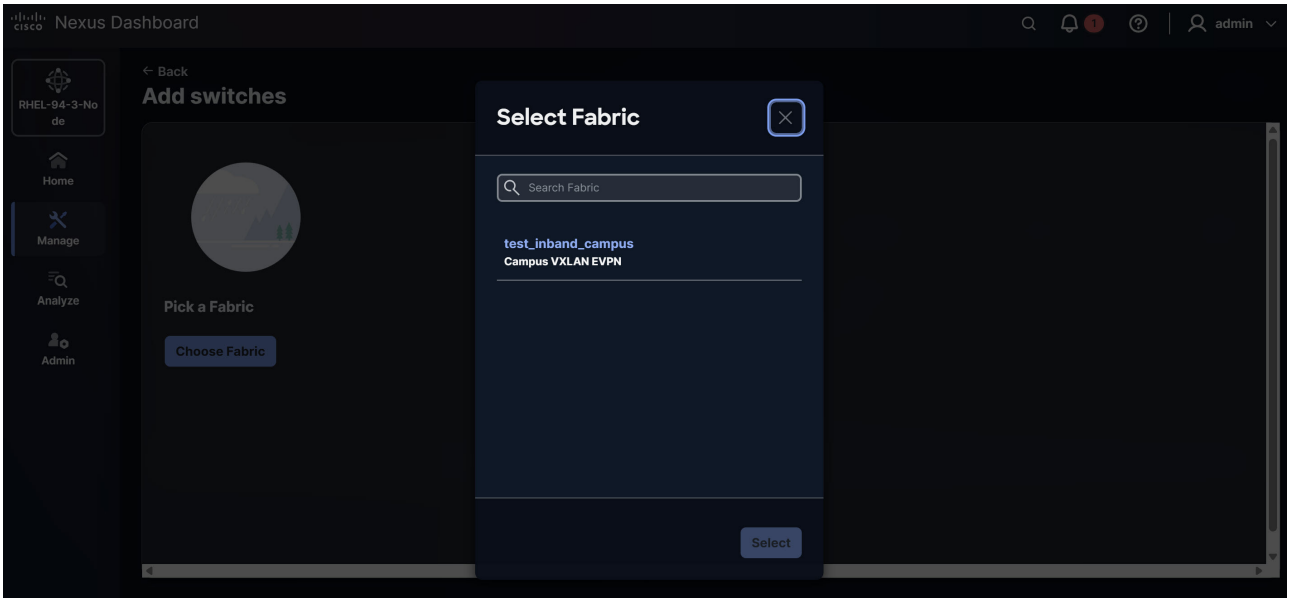

4. In the **Select Fabric** dialog box, choose the fabric to which you want to add switches.
5. Click **Select**.

   The **Add switches** page displays.

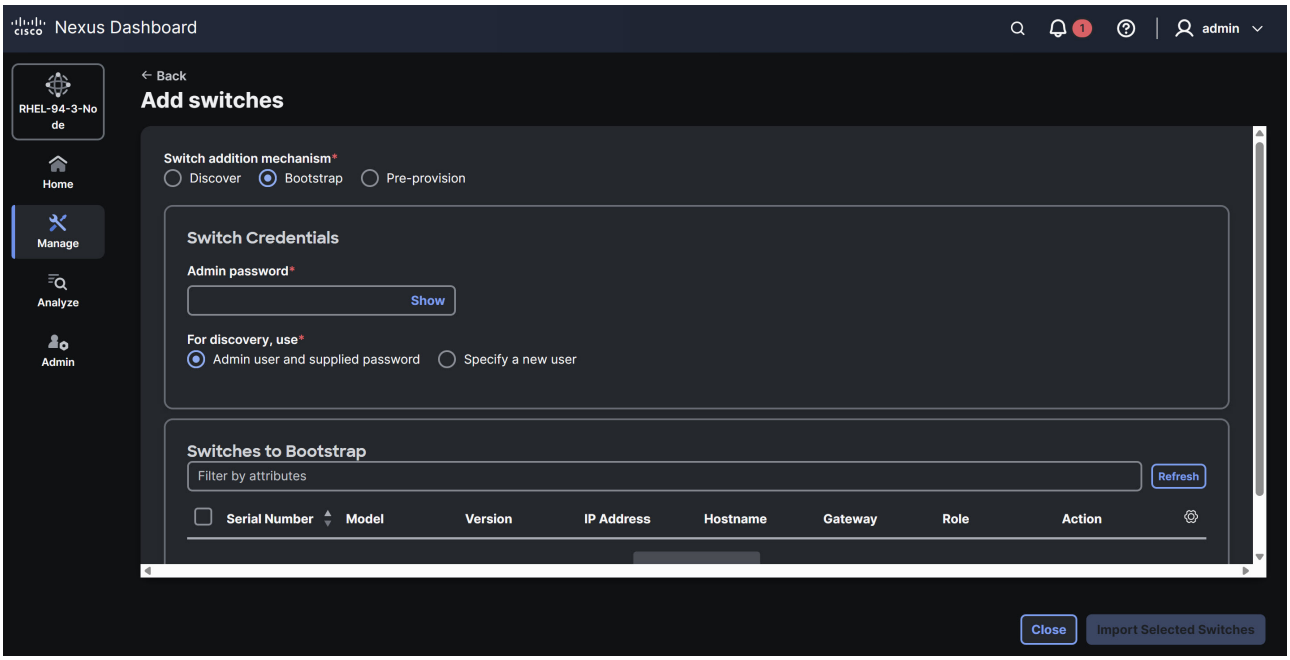

6. Under **Switch addition mechanism**, choose **Bootstrap**.
7. Under **Switch Credentials**, provide these details.
   - **Admin password**—enter the administrator password for the switches you want to import.
   - **For discovery, use**—choose one of the following options.
     - **Admin user and supplied password**—provide admin credentials.
     - **Specify a new user**—provide credentials for a different user account, if required.
8. Under **Switches to Bootstrap** section, review the list of switches that Nexus Dashboard detects or adds for import. Ensure all required switches are present and ready for bootstrapping.

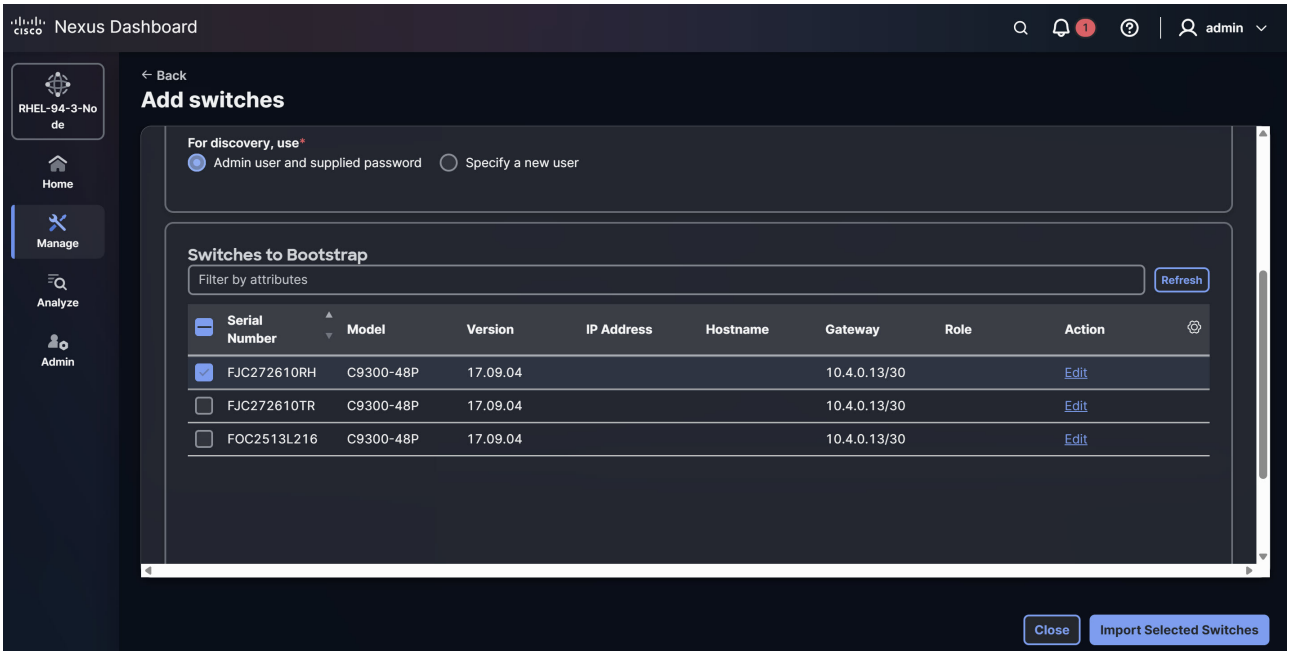

9. Choose the switches you wish to import from the list.

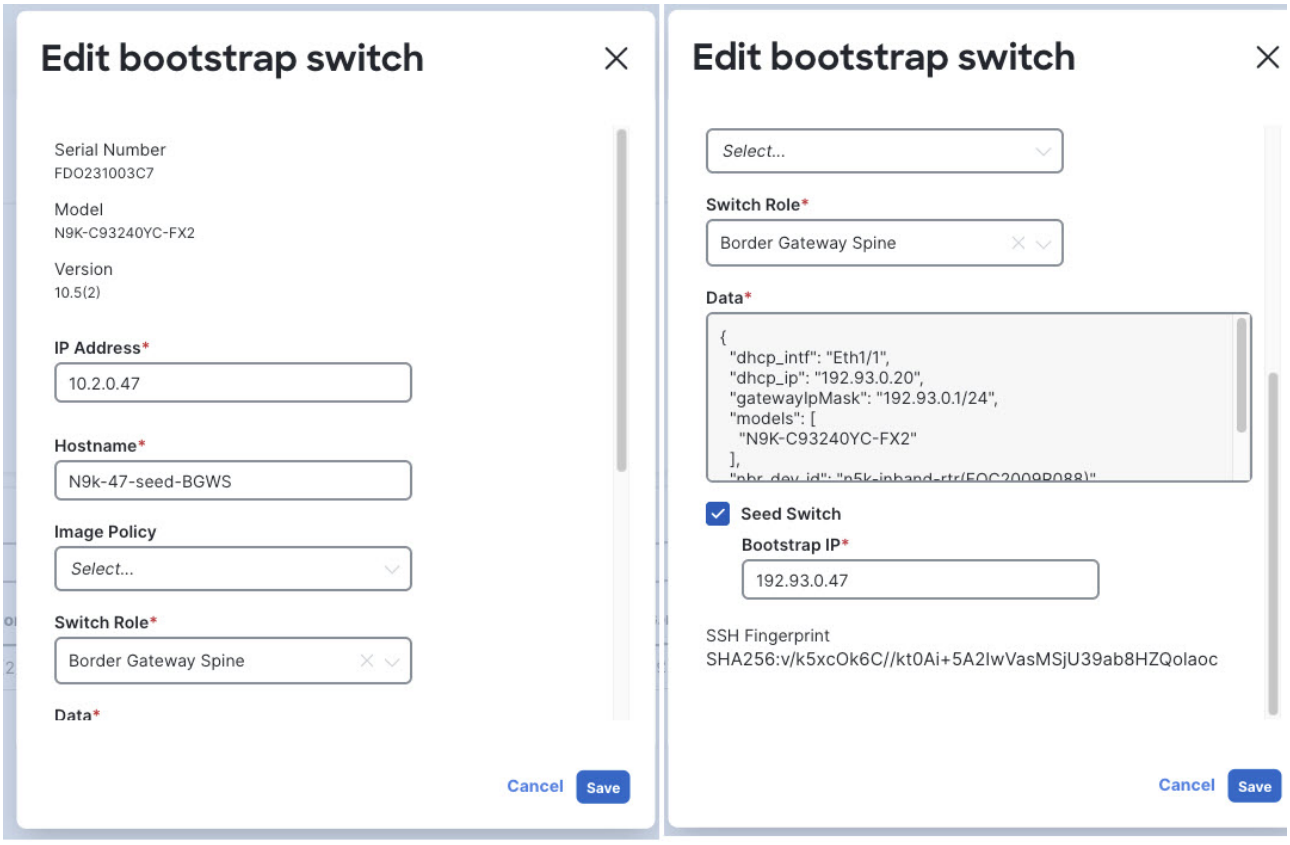

In the **Edit bootstrap switch** dialog box, choose either **leaf** or **border** for the **Switch Role** when you configure Catalyst 9000 series (Cat9K) switches, as these are the supported roles.

10. Click **Import Selected Switches** to initiate the import process.

Nexus Dashboard adds the switches to the chosen fabric.

When you add a Catalyst leaf switch to an IOS-XE Campus fabric that contains an N9k series border gateway spine, you may need to perform two **Recalculate and**

deploy cycles. After the first cycle, the Catalyst leaf transitions to **In-Sync** status, but the spine may transition to **Out-of-Sync** status. This occurs because Nexus Dashboard builds the spine configuration before the leaf configuration. Therefore, the fabric links and BGP configurations may not appear until you deploy and discover the leaf. You must perform **Recalculate and deploy** again to bring the spine to **In-Sync**.

Out-of-Band PnP in Campus VXLAN EVPN fabrics

About Out-of-Band PnP in Campus VXLAN EVPN fabrics

Nexus Dashboard provides support for the Out-of-Band (OOB) Plug and Play (PnP) feature, which simplifies the process of onboarding new devices with a zero-touch deployment experience. PnP automates the day-zero provisioning of Cisco Catalyst 9000 Series switches using Nexus Dashboard. With OOB connectivity, the switch provides a separate dedicated network for management traffic over GigabitEthernet 0/0 interface which is always placed in mgmt-vrf.

When a Cisco Catalyst switch powers up and does not find the startup configuration, the device enters PnP mode and sends out a DHCP request, which is served by a DHCP server. The DHCP offer contains the PnP server address and other required configuration for the PnP client. The PnP client on the switch uses this information to connect to the appropriate PnP sever to finish the configuration of the switch. This is a multi-step process, and to provide additional security, the system prompts you to approve the addition of this switch to the fabric. The device then downloads a configuration script and installs the appropriate software image and configuration file.

Enable Out-of-Band PnP

Before you begin:

- You can use a local or external DHCP Server for IP address assignment.
- This release only supports DHCP for IPv4.
- This feature is supported only on Campus VXLAN EVPN fabric.

To enable Out-of-Band PnP in a fabric:

1. When creating a Campus VXLAN EVPN fabric, navigate to the **Bootstrap** tab of the **Create Fabric** window.
2. Check the **Enable Bootstrap** and **Enable Local DHCP Server** check boxes.
3. Enter the domain name of the DHCP server in the **Domain name** field.
4. Enter the start address and the end address of your subnet in the **DHCP Scope Start Address** and **DHCP Scope End Address** fields.
5. Enter the default gateway for the management VRF on the switch in the **Switch Mgmt Default Gateway** field.
6. Enter the prefix for the management interface on the switch in the **Switch Mgmt IP Subnet Prefix** field.

   The prefix should be between 8 and 30. The system assigns 24, by default.

7. Enter any other additional configuration CLIs in the **Bootstrap Freeform Config** field, as required and click **Save**.

Import switches using the bootstrap method

Before you begin:

Ensure you have uploaded the required SSL certificates for the switches. You can use CA-signed certificates or self-signed certificates. Nexus Dashboard does not mandate the signing; however, the security guidelines suggest you use the CA-signed certificates.

1. In the **Fabric Overview** window for a fabric, choose **Actions > Add Switches**.
2. In the **Switch Addition Mechanism** area, click the **Bootstrap** radio button.
3. Enter the Admin password in the **Admin Password** field.

   The **Switches to Bootstrap** table lists all the discovered switches.
4. Select the switches you want to add to the fabric and click **Import Selected Switches**.

Enable DHCP for PnP

If you are using an external DHCP server and not the Nexus Dashboard DHCP server, ensure you perform the following steps on the DHCP server to facilitate PnP server discovery using DHCP.

1. Define the following two DHCP options:
   - pnpserver code 43
   - vrf code 194
2. Define option vendor-class-identifier to use the value " ciscopnp" .
3. Assign values to option 43 and option 194.
4. Assign option pnpserver " 5A1D;K4;B2;I<EXT-IP-NDFC>;J9666" ;

   For <EXT-IP-NDFC>, enter the external POAP IP address on Nexus Dashboard.
5. Assign option vrf to use the value " Mgmt-vrf" .

   The following example shows a sample configuration.

```
option pnpserver code 43 = text;

option vrf code 194 = text;

class " ciscopnp"  {

match if option vendor-class-identifier = " ciscopnp" ;

option vendor-class-identifier " ciscopnp" ;

option pnpserver " 5A1D; K4; B2; I<EXT-IP-NDFC>; J9666" ;

option vrf " Mgmt-vrf" ;

option domain-name " cisco.com" ;
```

```
}
```

Copyright

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706

USA
https://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883