Control

Fabrics

This section contains context-sensitive Online Help content for the **Control > Fabrics** tab. It has the following submenu:

VXLAN BGP EVPN Fabrics Provisioning

In DCNM 11.0, fabric creation is enhanced. In addition to overlay networks, you can also provision VXLAN BGP EVPN underlay network parameters to the fabric switches. Also, the concept of Multi-Site Domain (MSD) fabrics is introduced. The DCNM GUI is updated as follows:

Control > Fabric Builder menu option (under the **Fabrics** sub menu).

Fabric creation and updation:

- Create new standalone and MSD fabrics.
- Create an external fabric. The external network is representative of connections between the border devices of the fabric and the external fabric.
- View the list of fabrics that are already created and edit the overlay and underlay network ranges of the fabric, and the policy templates.

Device discovery and provisioning start-up configurations on new switches:

- Discover switches and the fabric topology. Also, provision start-up configurations and an IP address to a new switch through POAP configuration.
- Delete the fabrics.

Control > Interfaces menu option (under the **Fabrics** sub menu).

Underlay provisioning:

- Create, deploy, view, edit and delete a port-channel, vPC switch pair, straight through FEX, AA FEX, loopback, and sub interface.
- Create breakout and unbreakout ports.
- Shut down and bring up interfaces.
- Rediscover ports and view interface configuration history.
- Designate a switch interface as a routed port, trunk port, OSPF interface, and so on.

Control > Networks & VRFs menu option (under the **Fabrics** sub menu).

Overlay network provisioning.

- Create new overlay networks and VRFs (from the range specified in fabric creation).
- Provision the overlay networks and VRFs on the switches of the fabric.
- Undeploy the networks and VRFs from the switches.
- Remove the provisioning from the fabric in DCNM.

Control > Migration menu option (under the **Fabrics** sub menu).

NFM fabric migration to the VXLAN BGP EVPN fabric.

- In DCNM 10.4(2) release, Cisco Nexus Fabric Manager (NFM) fabric overlay migration to DCNM was introduced. In DCNM 11.0 release, NFM fabric underlay migration to DCNM has also been introduced.

This chapter covers all standalone fabric-related configurations. MSD fabric documentation is available in a separate chapter. Step by step configuration:

Create a New VXLAN BGP EVPN Fabric

1. Choose **Control > Fabric Builder**.

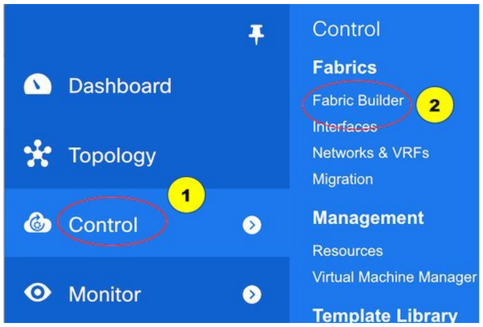

The **Fabric Builder** window appears. When you log in for the first time, the **Fabrics** section has no entries. After you create a fabric, it is displayed on the **Fabric Builder** window, wherein a rectangular box represents each fabric.

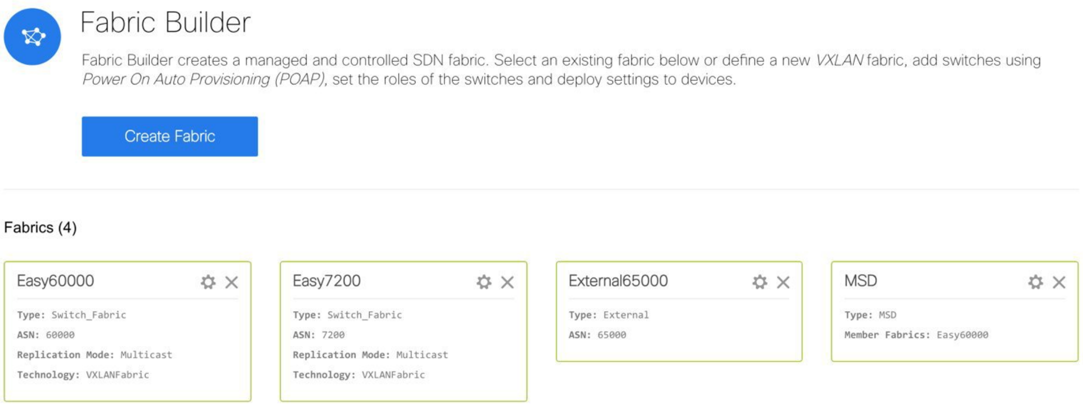

Fabric Builder

Fabric Builder creates a managed and controlled SDN fabric. Select an existing fabric below or define a new *VXLAN* fabric, add switches using *Power On Auto Provisioning (POAP)*, set the roles of the switches and deploy settings to devices.

Fabrics (4)

Easy60000
Type: Switch_Fabric
ASN: 60000
Replication Mode: Multicast
Technology: VXLANFabric

Easy7200
Type: Switch_Fabric
ASN: 7200
Replication Mode: Multicast
Technology: VXLANFabric

External65000
Type: External
ASN: 65000

MSD
Type: MSD
Member Fabrics: Easy60000

A standalone or member fabric contains **Switch_Fabric** in the Type field, the AS number in the **ASN** field, and mode of replication in the **Replication Mode** field.

2. Click **Create Fabric**. The **Add Fabric** window appears.

   Enter the name of the fabric in the **Fabric Name** field, and choose a template according to the type of fabric you want from the drop-down menu in **Fabric Template**.

   Choose **Easy_Fabric**. The fabric creation window for creating a standalone fabric comes up.

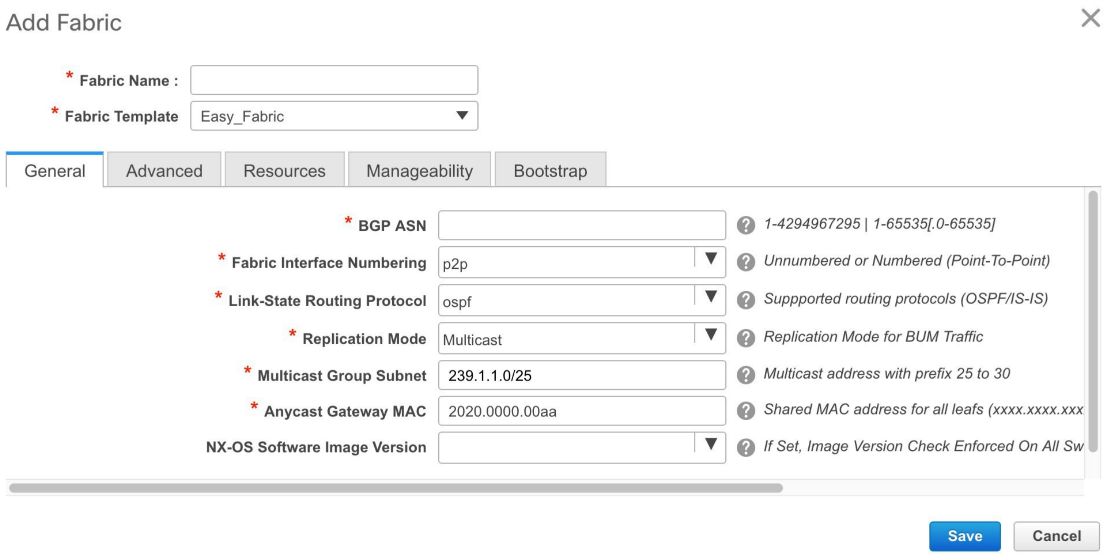

Add Fabric

* Fabric Name :
* Fabric Template: Easy_Fabric

* BGP ASN — 1-4294967295 | 1-65535[.0-65535]
* Fabric Interface Numbering: p2p — Unnumbered or Numbered (Point-To-Point)
* Link-State Routing Protocol: ospf — Suppported routing protocols (OSPF/IS-IS)
* Replication Mode: Multicast — Replication Mode for BUM Traffic
* Multicast Group Subnet: 239.1.1.0/25 — Multicast address with prefix 25 to 30
* Anycast Gateway MAC: 2020.0000.00aa — Shared MAC address for all leafs (xxxx.xxxx.xxx
NX-OS Software Image Version — If Set, Image Version Check Enforced On All Sw

The tabs and their fields in the screen are explained in the subsequent points. The overlay and underlay network parameters are included in these tabs.

Note If you are creating a standalone fabric as a potential member fabric of an MSD fabric (used for provisioning overlay networks for fabrics that are connected through EVPN Multi-Site technology), then browse through the overview of the MSD document before member fabric creation.

3. The **General** tab is displayed by default. The fields in this tab are:

BGP ASN: Enter the BGP AS number the fabric is associated with.

Fabric Interface Numbering : Specifies whether you want to use point-to-point or unnumbered networks.

Link-State Routing Protocol : The IGP used in the fabric, OSPF, or IS-IS.

Replication Mode : The mode of replication that is used in the fabric, Ingress Replication, or Multicast.

Multicast Group Subnet : Multicast group address of the network.

Anycast Gateway MAC : Anycast gateway MAC address.

NX-OS Software Image Version : Select an image from the list.

If you upload Cisco NX-OS software images through the image upload option, the uploaded images are listed in this field. If you select an image, the system checks if the switch has the selected version. If not, an error message is displayed. You can resolve the error by clicking on Resolve. The image management screen comes up and you can proceed with the ISSU option. Alternatively, you can delete the release number and save it later.

If you specify an image in this field, all switches in the fabric must run that image. If some devices do not run the image, a warning is prompted to perform an In-Service Software Upgrade (ISSU) to the specified image. Until all devices run the specified image, the deployment process is incomplete.

If you want to deploy more than one type of software image on the fabric switches, don’t specify any image. If an image is specified, delete it.

4. Click the **Advanced** tab. Most of the fields are auto generated. You can update the fields if needed.

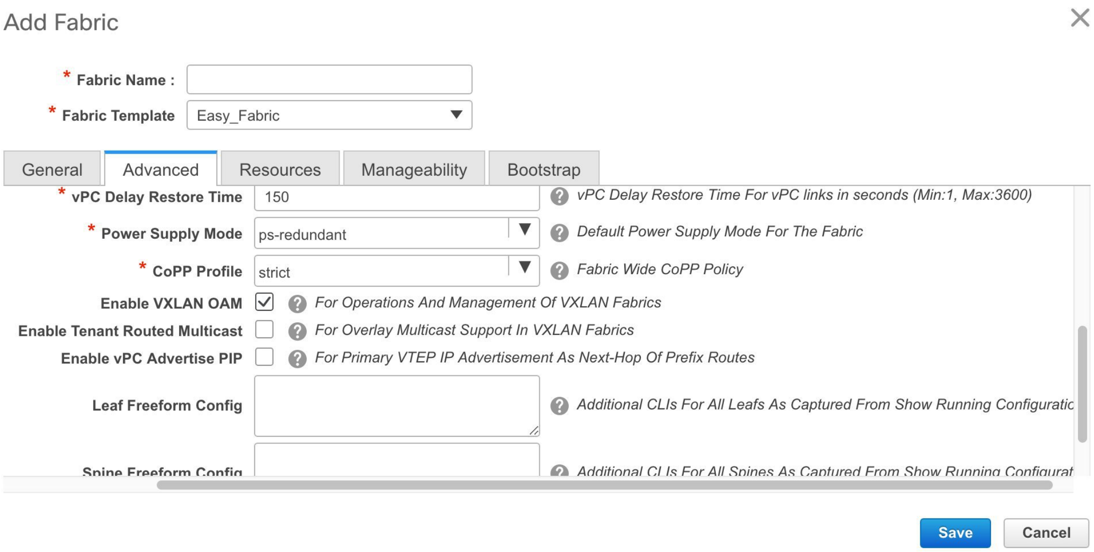

* Fabric Name :

* Fabric Template: Easy_Fabric

* vPC Delay Restore Time: 150 — vPC Delay Restore Time For vPC links in seconds (Min:1, Max:3600)

* Power Supply Mode: ps-redundant — Default Power Supply Mode For The Fabric

* CoPP Profile: strict — Fabric Wide CoPP Policy

Enable VXLAN OAM ☑ For Operations And Management Of VXLAN Fabrics

Enable Tenant Routed Multicast ☐ For Overlay Multicast Support In VXLAN Fabrics

Enable vPC Advertise PIP ☐ For Primary VTEP IP Advertisement As Next-Hop Of Prefix Routes

Leaf Freeform Config — Additional CLIs For All Leafs As Captured From Show Running Configuratic

Spine Freeform Config — Additional CLIs For All Spines As Captured From Show Running Configurat

VRF Template : Specifies the VRF template for the overlay networks.

Network Template : Specifies the network template for the overlay networks.

VRF Extension Template: Specifies the VRF extension template for extending the overlay networks to other fabrics.

Network Extension Template : Specifies the network extension template for extending the overlay networks to other fabrics.

Site ID : The ID for this fabric if you are moving this fabric within an MSD. The site ID is mandatory for a member fabric to be a part of an MSD. Each member fabric of an MSD has a unique site ID for identification.

Link-State Routing Protocol Tag : The tag defining the type of network.

vPC Peer Link VLAN : VLAN used for the vPC peer link SVI.

vPC Auto Recovery Time : Specifies the vPC auto recovery time-out period in seconds.

vPC Delay Restore Time - Specifies the vPC delay restore period in seconds.

Power Supply Mode - Choose the appropriate power supply mode.

CoPP Profile - Choose the appropriate Control Plane Policing (CoPP) profile policy for the fabric. By default, the strict option is populated.

Enable VXLAN OAM - Enables the VXLAM OAM function.

Note The VXLAN OAM feature in Cisco DCNM is only supported on a single fabric or site.

Enable Tenant Routed Multicast - Enables overlay multicast protocol support in the fabric.

Enable vPC Advertise PIP - Enables the Advertise PIP feature.

Freeform CLIs - Fabric level freeform CLIs can be added while creating or editing a fabric. They are applicable to switches across the fabric. You must add the configurations as displayed in the running configuration, without indentation. Switch level freeform configurations such as VLAN, SVI, and interface configurations should only be added on the switch. Refer the *Freeform Configurations on Fabric Switches* topic for a detailed explanation and examples.

Leaf Freeform Config - Add CLIs that should be added to switches that have the *Leaf*, *Border*, and *Border Gateway* roles.

Spine Freeform Config - Add CLIs that should be added to switches with a *Spine* role.

5. Click the **Resources** tab.

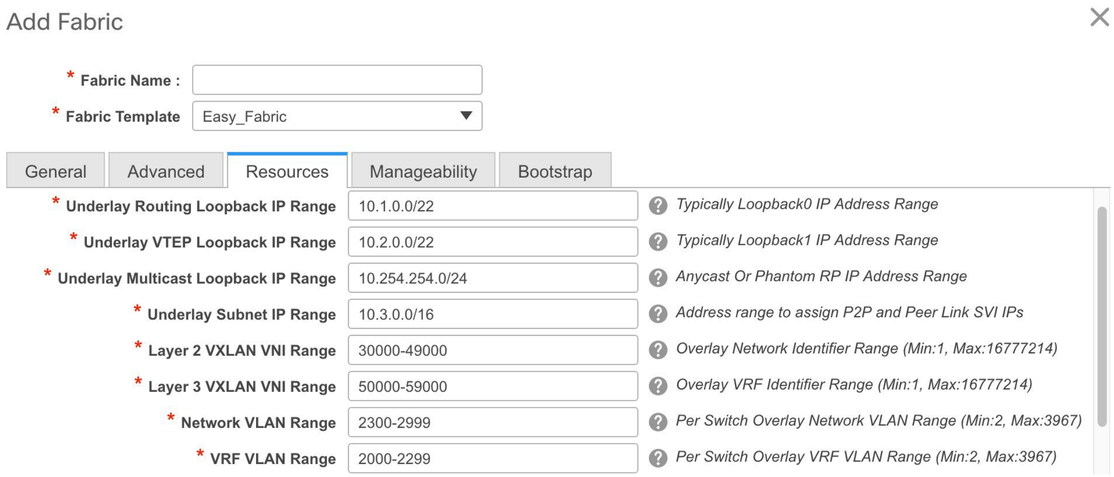

Add Fabric

* Fabric Name :

* Fabric Template: Easy_Fabric

* Underlay Routing Loopback IP Range	10.1.0.0/22	Typically Loopback0 IP Address Range
* Underlay VTEP Loopback IP Range	10.2.0.0/22	Typically Loopback1 IP Address Range
* Underlay Multicast Loopback IP Range	10.254.254.0/24	Anycast Or Phantom RP IP Address Range
* Underlay Subnet IP Range	10.3.0.0/16	Address range to assign P2P and Peer Link SVI IPs
* Layer 2 VXLAN VNI Range	30000-49000	Overlay Network Identifier Range (Min:1, Max:16777214)
* Layer 3 VXLAN VNI Range	50000-59000	Overlay VRF Identifier Range (Min:1, Max:16777214)
* Network VLAN Range	2300-2999	Per Switch Overlay Network VLAN Range (Min:2, Max:3967)
* VRF VLAN Range	2000-2299	Per Switch Overlay VRF VLAN Range (Min:2, Max:3967)

The fields in this tab are:

Underlay Routing Loopback IP Range - Specifies loopback IP addresses for the protocol peering.

Underlay VTEP Loopback IP Range - Specifies loopback IP addresses for VTEPs.

Underlay Multicast Loopback IP Range - Specifies loopback IP addresses for multicast routing.

Underlay Subnet IP Range - IP addresses for underlay P2P routing traffic between interfaces.

Layer 2 VXLAN VNI Range and **Layer 3 VXLAN VNI Range** - Specifies the VXLAN VNI IDs for the fabric.

Network VLAN Range and **VRF VLAN Range** - VLAN ranges for the Layer 3 VRF and overlay network.

Subinterface Dot1q Range - Specifies the subinterface range when L3 sub interfaces are used.

Note The values shown in the screen shot are automatically generated. If you want to update the IP address ranges, VXLAN Layer 2/Layer 3 network ID ranges or the VRF/Network VLAN ranges, ensure the following:

If you update a range of values, ensure that it does not overlap with other ranges.

You must only update one range of values at a time. If you want to update more than one range of values, do it in separate instances. For example, if you want to update L2 and L3 ranges, you should do the following.

1. Update the L2 range and click **Save**.
2. Click the Edit Fabric option again, update the L3 range and click **Save**.

6. Click the **Manageability** tab.

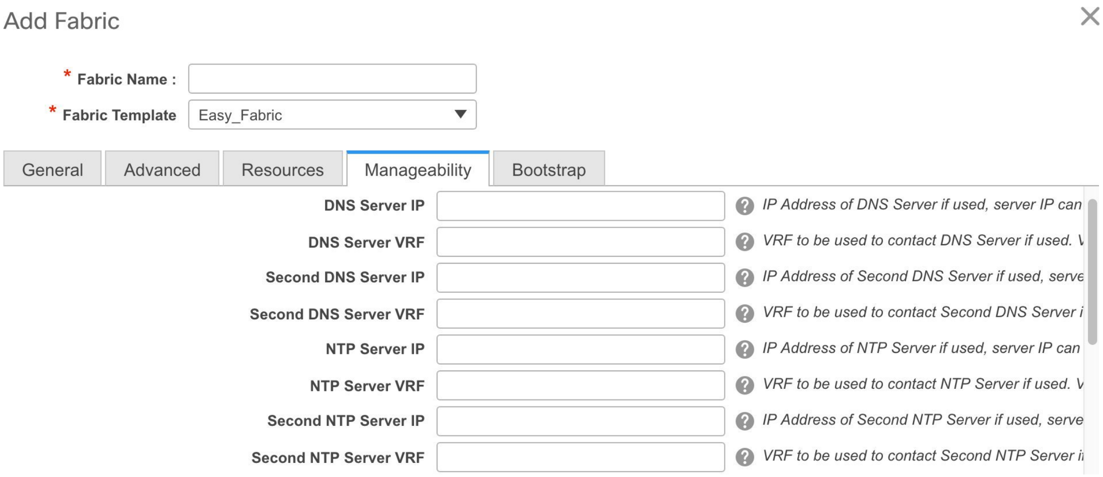

Add Fabric

* Fabric Name :

* Fabric Template: Easy_Fabric

DNS Server IP	*IP Address of DNS Server if used, server IP can*
DNS Server VRF	*VRF to be used to contact DNS Server if used. V*
Second DNS Server IP	*IP Address of Second DNS Server if used, serve*
Second DNS Server VRF	*VRF to be used to contact Second DNS Server i*
NTP Server IP	*IP Address of NTP Server if used, server IP can*
NTP Server VRF	*VRF to be used to contact NTP Server if used. V*
Second NTP Server IP	*IP Address of Second NTP Server if used, serve*
Second NTP Server VRF	*VRF to be used to contact Second NTP Server i*

The fields in this tab are:

DNS Server IP - Specifies the IP address of the DNS server, if you use a DNS server.

DNS Server VRF - Specifies the VRF to be used to contact the DNS server IP address.

Second DNS Server IP - Specifies the IP address of the second DNS server, if you use a second DNS server.

Second DNS Server VRF - Specifies the VRF to be used to contact the second DNS server IP address.

NTP Server IP - Specifies the IP address of the NTP server, if you use an NTP server.

NTP Server VRF - Specifies the VRF to be used to contact the NTP server IP address.

Second NTP Server IP - Specifies the IP address of the second NTP server, if you use a second NTP server.

Second NTP Server VRF - Specifies the VRF to be used to contact the second NTP server IP address.

AAA Server Type - Specifies the AAA server type. By default, no type is populated. You can select a radius or TACACS server.

AAA Server IP - Specifies the IP address of the AAA server, if you use a AAA server.

AAA Shared Secret - Specifies the shared secret of the AAA server, if used.

Note After fabric creation and discovery of switches, you must update the AAA server password on each fabric switch.

Second AAA Server IP - Specifies the IP address of the second AAA server, if you use a second AAA server.

Second AAA Shared Secret - Specifies the shared secret of the second AAA server, if used.

AAA Server VRF - Specifies the VRF to be used to contact the AAA server IP address.

7. Click the **Bootstrap** tab.

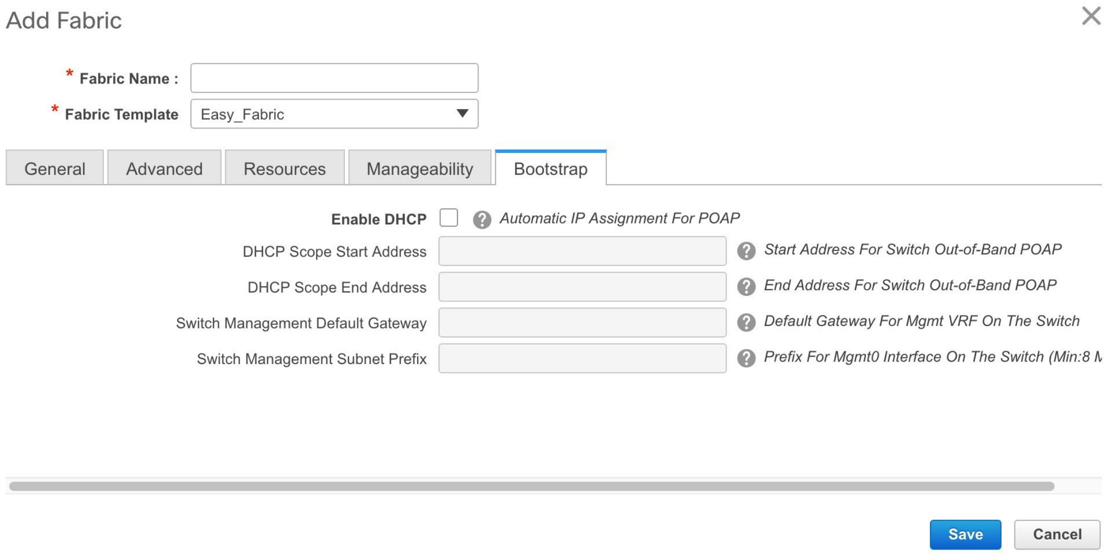

* Fabric Name :

* Fabric Template: Easy_Fabric

Enable DHCP □ Automatic IP Assignment For POAP

DHCP Scope Start Address — Start Address For Switch Out-of-Band POAP

DHCP Scope End Address — End Address For Switch Out-of-Band POAP

Switch Management Default Gateway — Default Gateway For Mgmt VRF On The Switch

Switch Management Subnet Prefix — Prefix For Mgmt0 Interface On The Switch (Min:8 M

The fields on this tab are:

Enable DHCP - Click this check box to initiate enabling of automatic IP address assignment through DHCP. When you click the check box, the other fields become editable. They are:

DHCP Scope Start Address and **DHCP Scope End Address** - Specifies the first and last IP addresses of the IP address range to be used for the switch out of band POAP.

Switch Management Default Gateway - Specifies the default gateway for the management VRF on the switch.

Switch Management Subnet Prefix - Specifies the prefix for the Mgmt0 interface on the switch. The prefix should be between 8 and 30.

DHCP scope and management default gateway IP address specification - If you specify the management default gateway IP address 10.0.1.0 and subnet mask 24, ensure that the DHCP scope is within the specified subnet, between 10.0.1.1 and 10.0.1.254.

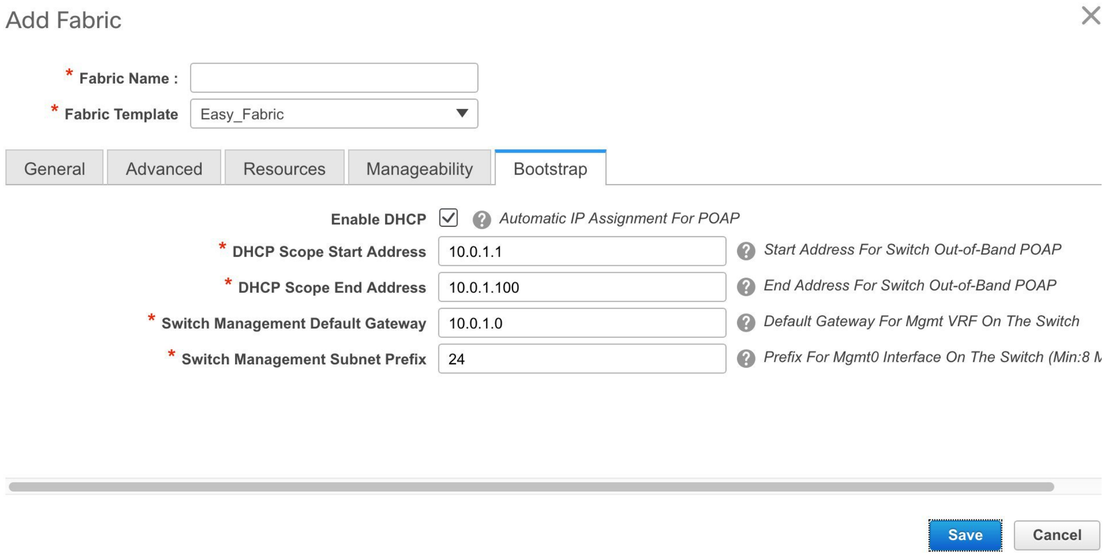

* Fabric Name :

* Fabric Template: Easy_Fabric

Enable DHCP ☑ Automatic IP Assignment For POAP

* DHCP Scope Start Address: 10.0.1.1 — Start Address For Switch Out-of-Band POAP

* DHCP Scope End Address: 10.0.1.100 — End Address For Switch Out-of-Band POAP

* Switch Management Default Gateway: 10.0.1.0 — Default Gateway For Mgmt VRF On The Switch

* Switch Management Subnet Prefix: 24 — Prefix For Mgmt0 Interface On The Switch (Min:8 M

8. Click **Save** after filling and updating relevant information. A note appears briefly at the bottom right part of the screen, indicating that the fabric is created. When a fabric is created, the fabric page comes up. The fabric name appears at the top left part of the screen.

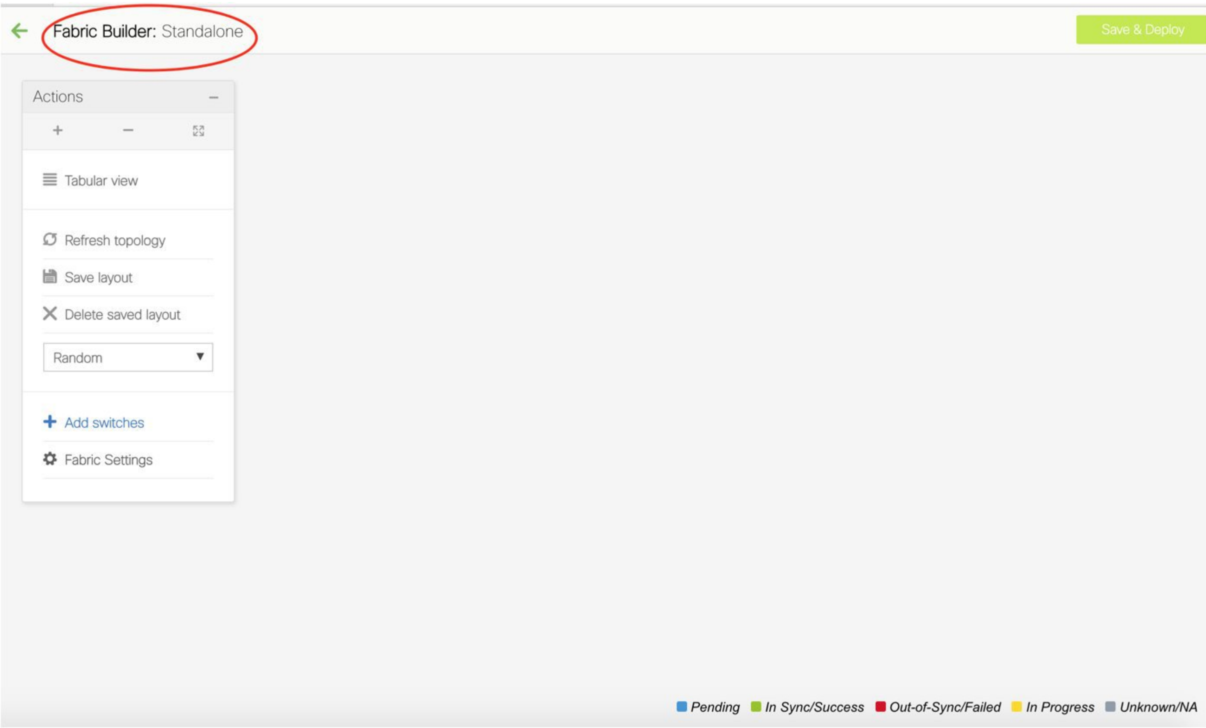

(At the same time, the newly created fabric instance appears on the Fabric Builder page. To go to the Fabric Builder page, click the left arrow (⟵) button above the Actions panel [to the left of the screen]).

The Actions panel at the left part of the screen allows you to perform various functions. One of them is the Add switches option to add switches to the fabric. After you create a fabric, you should add fabric devices. The other options are:

- **Tabular View** - By default, the switches are displayed in the topology view. Use this option to view switches in the tabular view.
- **Refresh topology** - Allows you to refresh the topology.
- You can choose between Hierarchical, Random and Custom saved layout display options.
- **Hierarchical** - Provides an architectural view of your topology. Various Switch Roles can be defined that draws the nodes on how you configure your CLOS topology.
- **Random** - Nodes are placed randomly on the screen. DCNM tries to make a guess and intelligently place nodes that belong together in close-proximity.
- **Custom saved layout** - You can drag nodes around to your liking. Once you have the positions as how you like, you can click Save Layout to remember the positions. Next time you come to the topology, DCNM will draw the nodes based on your last saved layout positions.
- **Save Layout** and **Delete saved layout** - Allows you to save the custom layout and remove the custom layout.

Delete a Fabric

Choose **Control > Fabric Builder**. On the Fabric Builder page, click **X** on the rectangular box that represents the fabric. Ensure the following before deleting a fabric.

- Fabric devices should not be in transition such as migration into or out of the fabric, ongoing network or VRF provisioning, and so on. Delete a fabric after the transition is complete.
- Remove devices that are still attached to the fabric. Remove non-Cisco Nexus 9000 Series switches first and then remove the 9000 Series switches.

Add Switch Instances to the Fabric

Networks and VRFs can be extended (and hence can be common) across fabrics. However, switches in each fabric are unique, and hence, each switch can only be added to one fabric.

Click the **Add Switches** option from the Actions panel to add switches to the fabric. The Inventory Management screen comes up. The screen contains two tabs, one for discovering existing switches and the other for discovering new switches. Both options are explained.

Discovering Existing Switches

1. Use the **Discover Existing Switches** tab to add an existing switch. In this case, a switch with known credentials is added to the Standalone fabric. The IP address (Seed IP), administrator username, and password (**Username** and **Password** fields) of the switch are keyed in.

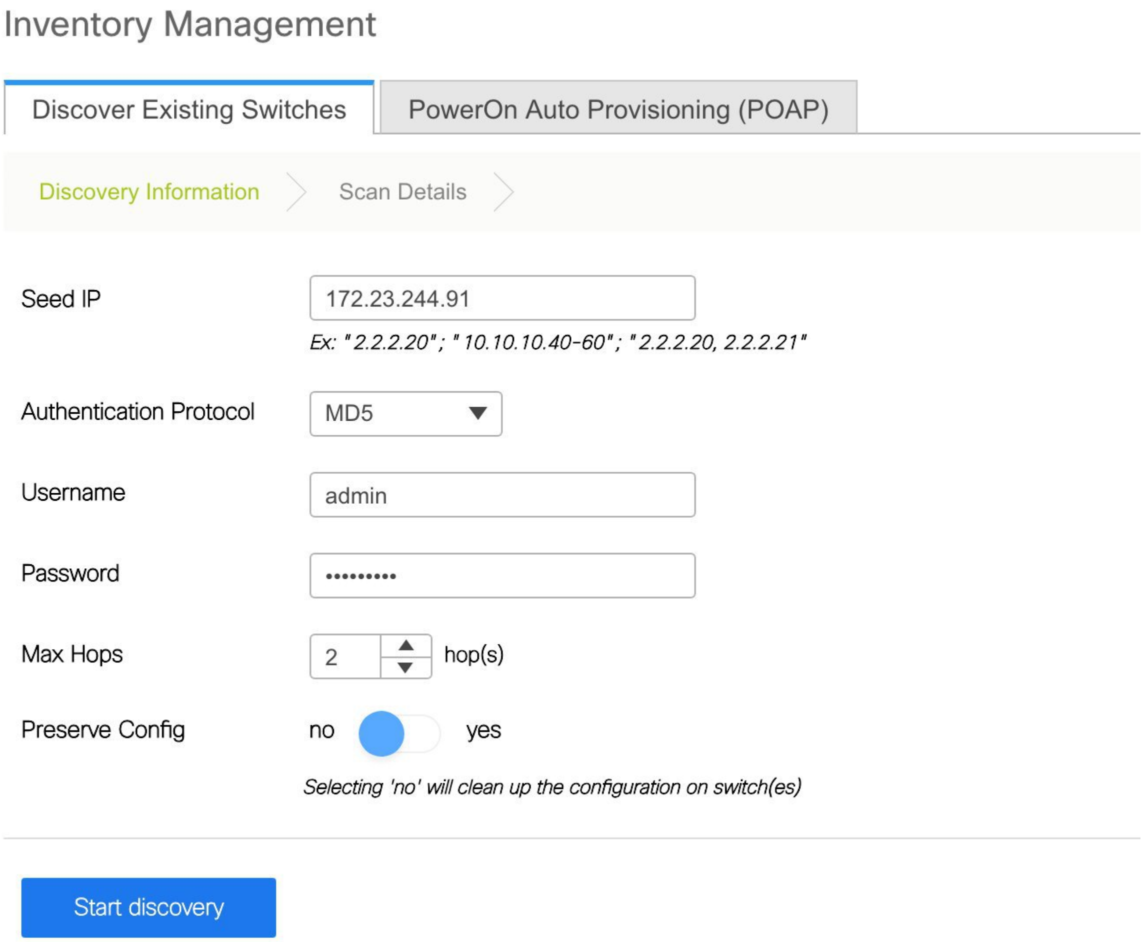

Inventory Management

Seed IP: 172.23.244.91

Ex: "2.2.2.20"; "10.10.10.40-60"; "2.2.2.20, 2.2.2.21"

Authentication Protocol: MD5

Username: admin

Password: •••••••••

Max Hops: 2 hop(s)

Preserve Config: no ● yes

Selecting 'no' will clean up the configuration on switch(es)

2. Click **Start discovery**. The Scan Details section comes up shortly. Since the **Max Hops** field was populated with 2, the switch with the specified IP address (leaf-91) and switches two hops from it are populated in the Scan Details section.

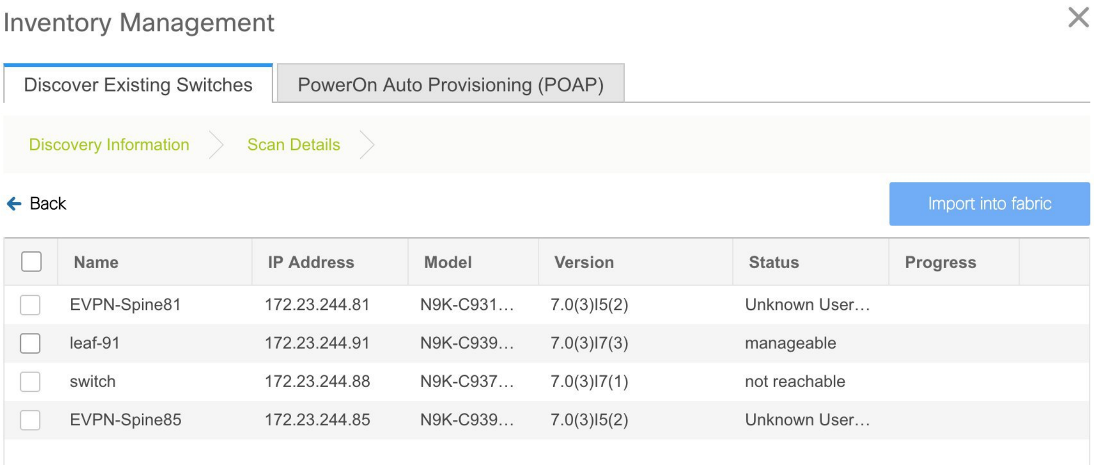

Inventory Management

□	Name	IP Address	Model	Version	Status	Progress
□	EVPN-Spine81	172.23.244.81	N9K-C931…	7.0(3)I5(2)	Unknown User…	
□	leaf-91	172.23.244.91	N9K-C939…	7.0(3)I7(3)	manageable	
□	switch	172.23.244.88	N9K-C937…	7.0(3)I7(1)	not reachable	
□	EVPN-Spine85	172.23.244.85	N9K-C939…	7.0(3)I5(2)	Unknown User…	

3. Select the check box next to the concerned switch and click **Import into fabric**.

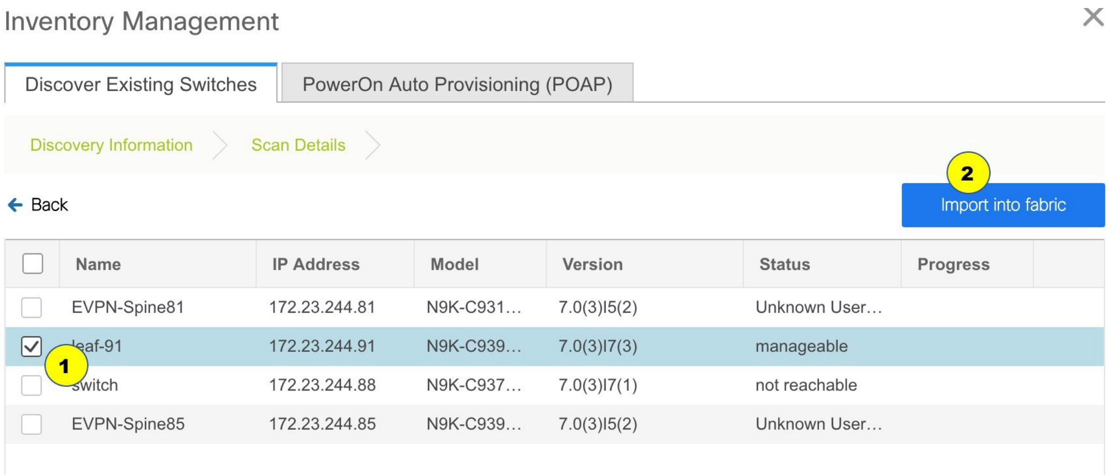

	Name	IP Address	Model	Version	Status	Progress
□	EVPN-Spine81	172.23.244.81	N9K-C931…	7.0(3)I5(2)	Unknown User…	
☑	leaf-91	172.23.244.91	N9K-C939…	7.0(3)I7(3)	manageable	
□	switch	172.23.244.88	N9K-C937…	7.0(3)I7(1)	not reachable	
□	EVPN-Spine85	172.23.244.85	N9K-C939…	7.0(3)I5(2)	Unknown User…	

This example describes the discovery of one switch. You can discover multiple switches at the same time. The switches must be properly cabled and connected to the DCNM server and the switch status must be *manageable*.

The switch discovery process is initiated. The **Progress** column displays the progress. After DCNM discovers the switch, the screen closes and the *Standalone* fabric screen comes up again. The switch icon can be seen at the center of the fabric page.

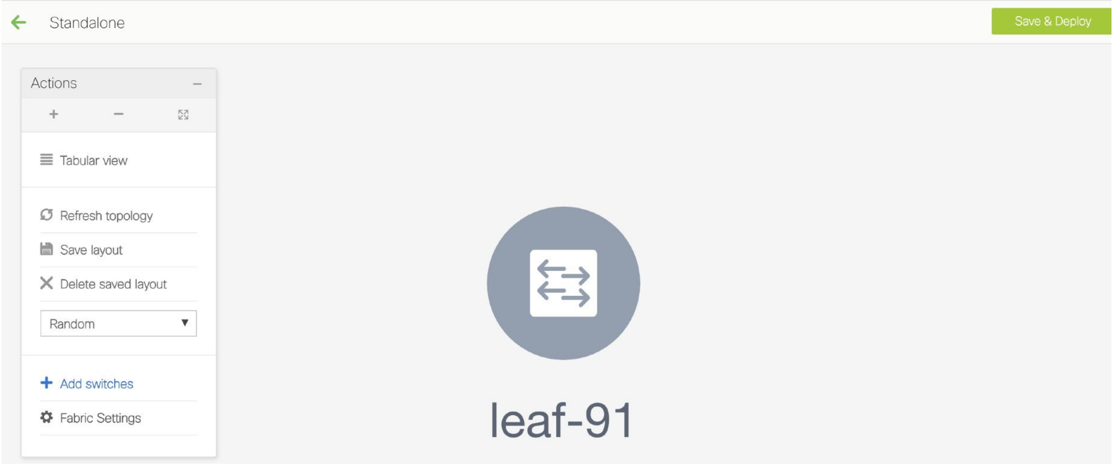

4. Click **Refresh topology** to view the latest topology view.

   When more switches are added and roles assigned to them (which is explained in the next point), the fabric topology looks like the following image:

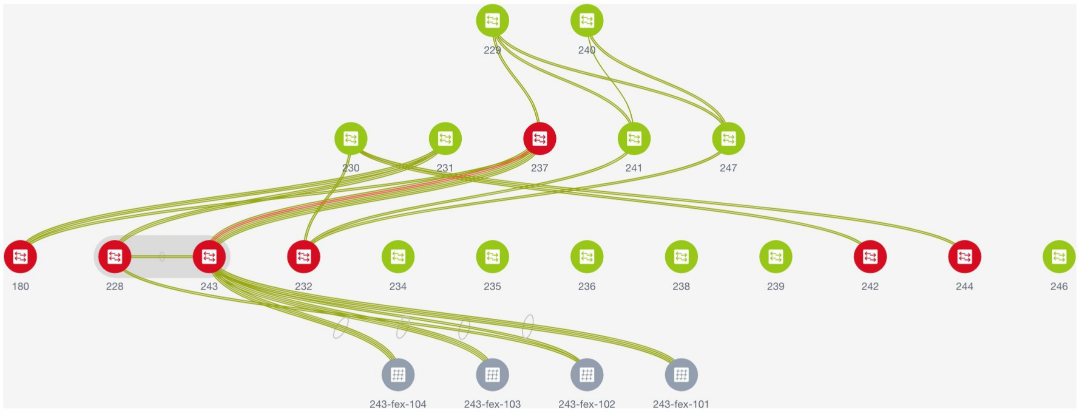

5. After discovering the switches, assign the fabric role to each switch. Since each switch is assigned the leaf role by default, assign the Border Gateway, Border (for a border leaf switch), and Spine roles. Right click the switch, and use the **Set role** option to set the appropriate role.

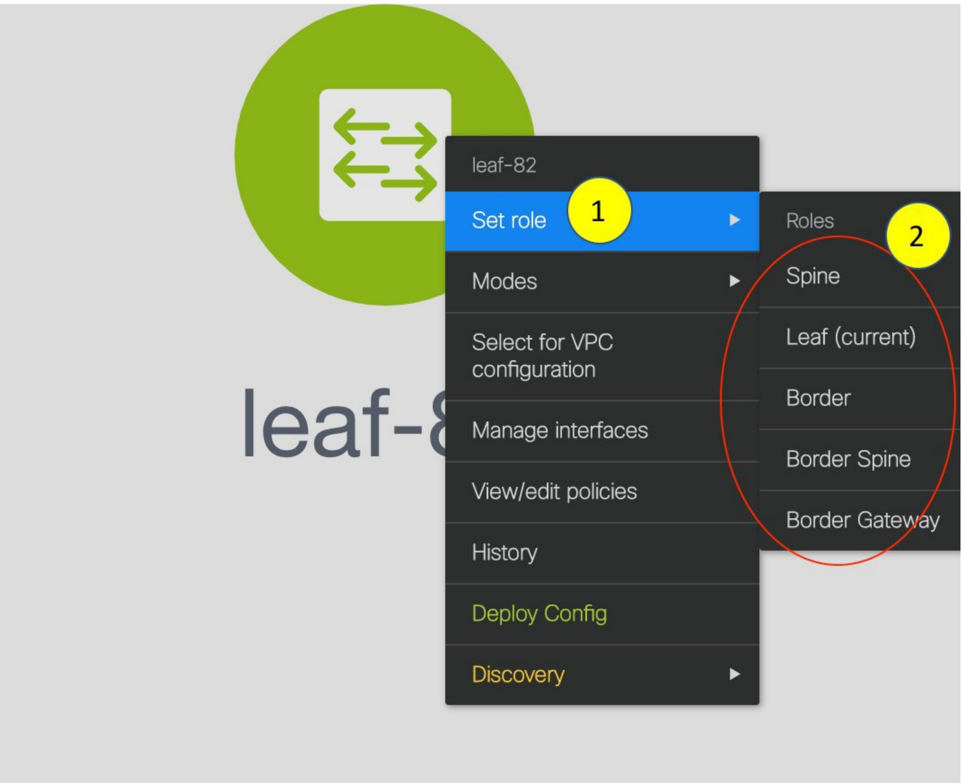

The topology automatically gets aligned as per role assignment, with the leaf switches at the bottom, the spine switches connected on top of them, and the BGW at the top.

Note To connect fabrics using the EVPN Multi-Site feature, you must change the role of the designated BGW to *Border Gateway*. To connect fabrics using the VRF Lite feature, you must change the role of the border leaf switch to *Border*. If you want to deploy VRF Lite and EVPN Multi-Site features in a fabric, you must set the device role to *Border Gateway* and provision VRF Lite and Multi-Site features. If you do not update border device roles correctly at this stage, then you will have to remove the device from the fabric and discover it again through DNCM using the POAP bootstrap option and reprovision the configurations for the device.

Assign vPC switch role - To designate a pair of switches as a vPC switch pair, right-click the switch and choose the vPC peer switch from the list of switches.

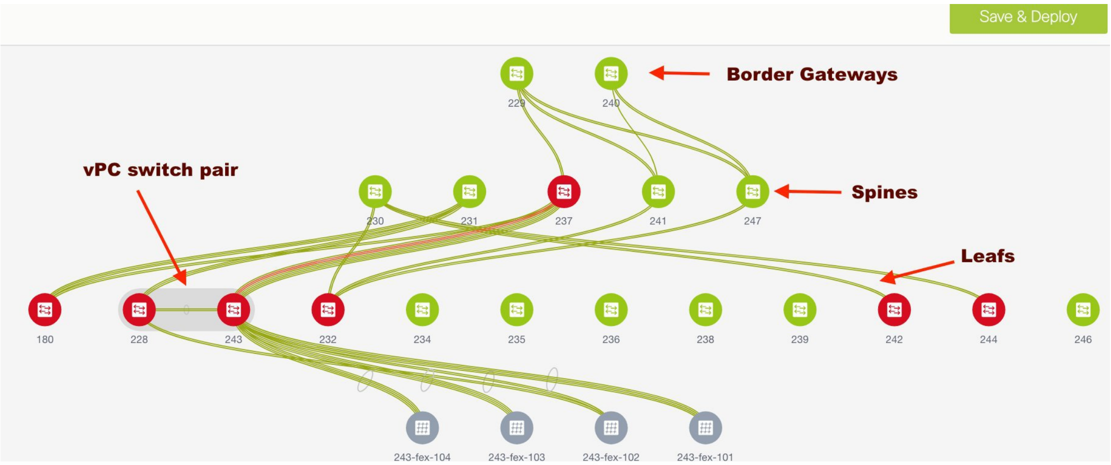

AAA server password - During fabric creation, if you have entered AAA server information (in the **Manageability** tab), you must update the AAA server password on each switch. Else, switch discovery fails.

6. Click **Save & Deploy** at the top right part of the screen. The template and interface configurations form the underlay network configuration provisioning on the switches.

   Also, freeform CLIs that were entered earlier are deployed.

   Configuration Compliance - If the provisioned configurations and switch configurations do not match, then the switch icon turns red, indicating an out of sync status. For example, if you enable a function on the switch manually through a CLI, then it results in a configuration mismatch.

   To ensure that the configurations that are provisioned from DCNM to the switch are accurate and detect any deviation from the intended configuration, DCNM recognizes and reports configuration deviation, and provides remediation configuration. Configuration compliance is supported for the fabric underlay and overlay deployments for Cisco Nexus 9000 Series switches.

7. When you click **Save & Deploy**, the Configuration Deployment Status section comes up.

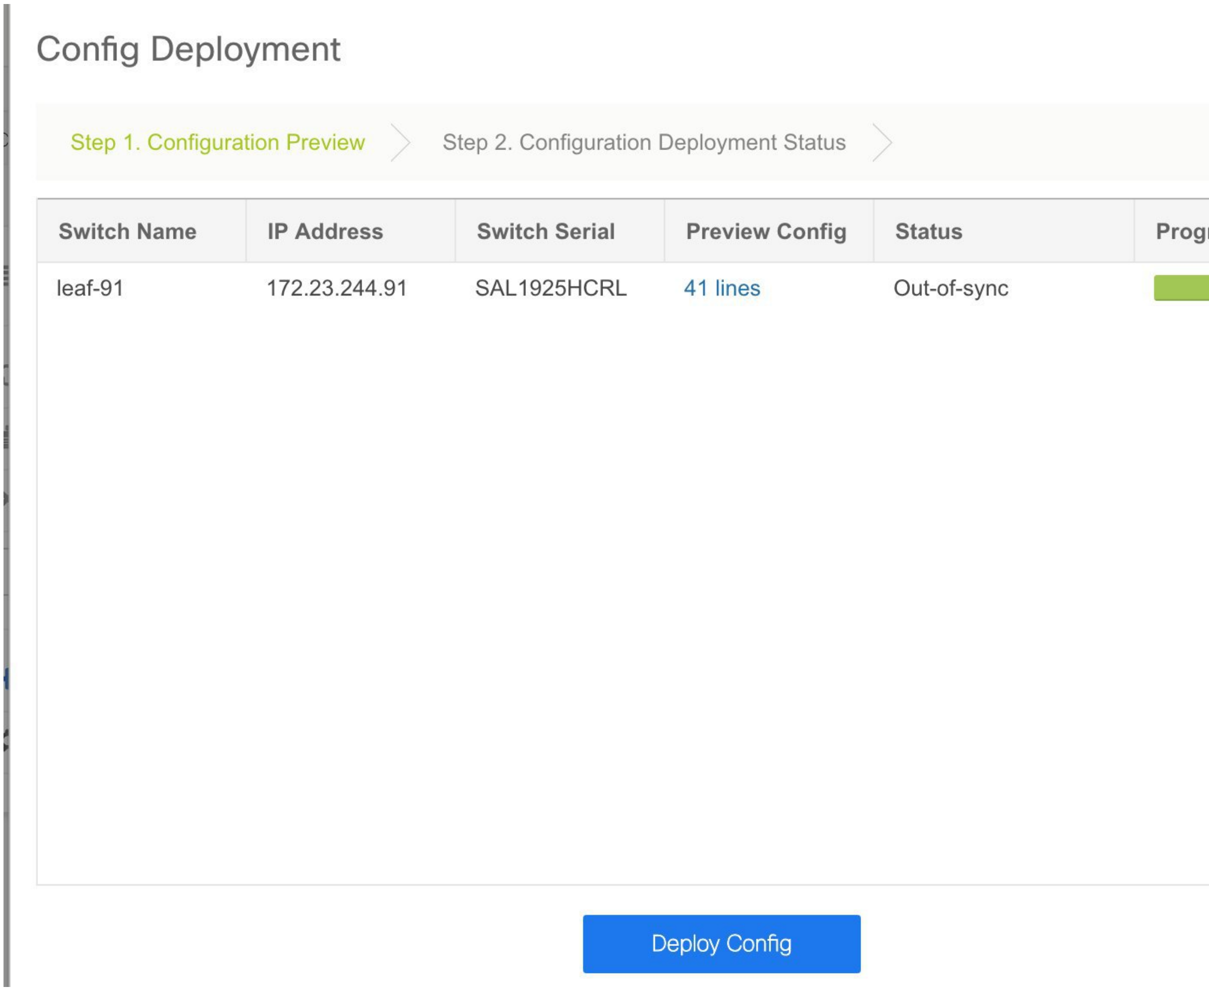

Config Deployment

Step 1. Configuration Preview
Step 2. Configuration Deployment Status

Switch Name	IP Address	Switch Serial	Preview Config	Status	Progr
leaf-91	172.23.244.91	SAL1925HCRL	41 lines	Out-of-sync	🟩

If the status is Out-of-sync, it suggests a compliance issue. Click the **Preview Config** column entry (updated with a specific number of lines). The **Config Preview** screen comes up.

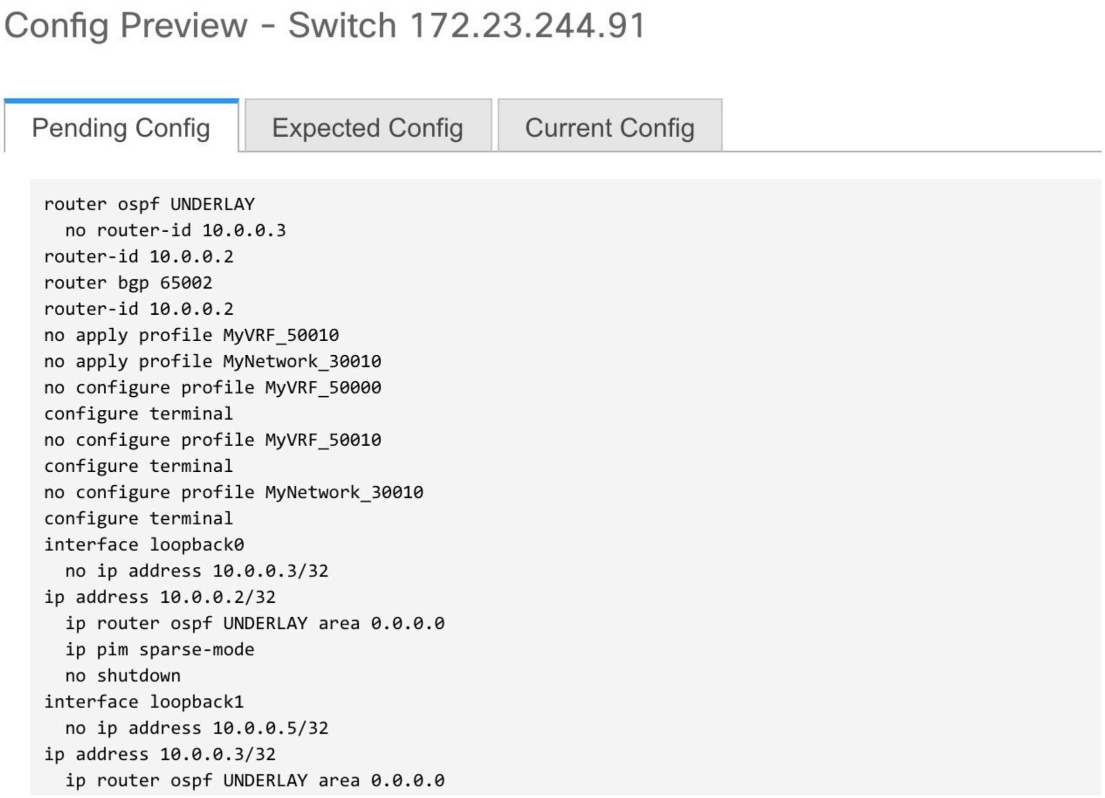

```
router ospf UNDERLAY
  no router-id 10.0.0.3
router-id 10.0.0.2
router bgp 65002
router-id 10.0.0.2
no apply profile MyVRF_50010
no apply profile MyNetwork_30010
no configure profile MyVRF_50000
configure terminal
no configure profile MyVRF_50010
configure terminal
no configure profile MyNetwork_30010
configure terminal
interface loopback0
  no ip address 10.0.0.3/32
ip address 10.0.0.2/32
  ip router ospf UNDERLAY area 0.0.0.0
  ip pim sparse-mode
  no shutdown
interface loopback1
  no ip address 10.0.0.5/32
ip address 10.0.0.3/32
  ip router ospf UNDERLAY area 0.0.0.0
```

The **Pending Config** tab displays the pending configurations for successful deployment. The other tabs display the expected and configured configurations.

8. Close the screen. In the Configuration Deployment screen, click **Deploy Config** at the bottom part of the screen to initiate pending configuration onto the switch. The Status column displays FAILED or SUCCESS state. For a FAILED status, investigate the reason for failure to address the issue.

   After correct provisioning and successful configuration compliance, close the screen. The switch icon colour turns to green, indicating successful configuration.

You can right click the switch icon and update switch related settings, as displayed in the image.

You can use **Save & Deploy** for single and multiple switches. Add switches and then click **Save & Deploy** to ensure configuration compliance. Whether discovering multiple switches at once or one by one, as a best practice, use **Save & Deploy** and not the **Deploy Config** option (accessible after right-clicking the switch icon).

When a leaf switch boots up after a switch reload or RMA operation, DCNM provisions configurations for the switch and FEX devices connected to it. Occasionally, FEX connectivity comes up after DCNM provisions FEX (host interface) configurations, resulting in a configuration mismatch. To resolve the mismatch, click **Save & Deploy** again in the fabric topology screen.

An example of the **Deploy Config** option usage is for switch-level freeform configurations. Refer the *Freeform Configurations on Fabric Switches* topic for details.

The Configuration Compliance function and principles are applicable for discovering existing and new switches. New switch discovery in DCNM (through a simplified POAP process) is explained next.

Discovering New Switches

1. Power on the new switch after ensuring that it is cabled to the DCNM server. Boot the Cisco NX-OS and setup switch credentials.

2. Execute the **write erase** and **reload** commands on the switch.

   Click *Yes* to both the CLI commands that prompt you to choose **Yes** or **No**.

3. Set the boot variable to the image that you want to POAP. DCNM uses this image to POAP. Also, DCNM injects an information script into the switch to collect the device onboarding information.

4. In the DCNM GUI, go to the *Standalone* fabric (Click **Control > Fabric Builder** and click the fabric *Standalone*). The fabric topology is displayed.

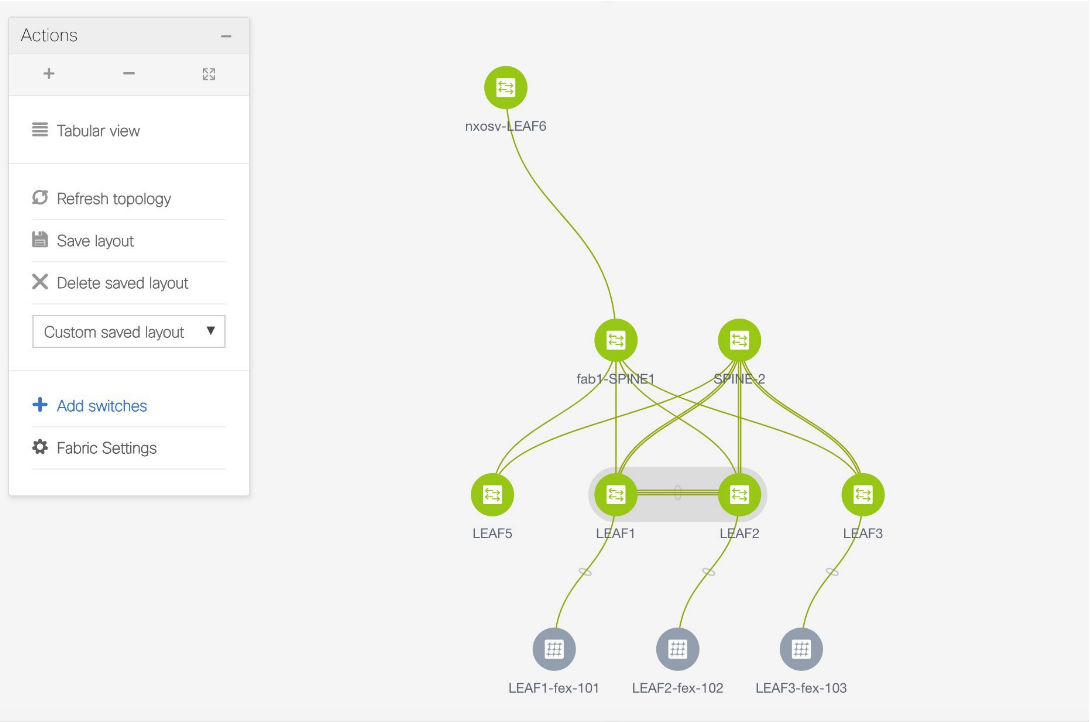

Note If you want to POAP with DHCP, make sure that DHCP is enabled on the fabric settings. Click **Fabric Settings** and edit the DHCP information in the **Bootstrap** tab.

5. Go to the fabric topology screen and click the **Add switches** option from the **Actions** panel. The Inventory Management screen comes up.

6. Click the **POAP** tab.

   In an earlier step, the **reload** command was executed on the switch. When the switch restarts to reboot, DCNM retrieves the serial number, model number, and version from the switch and displays them on the Inventory Management along screen. Also, an option to add the IP address, hostname, and password are made available. If the switch information is not retrieved, refresh the screen.

Note At the top left part of the screen, *export* and *import* options are provided to export and import the .csv file that contains the switch information.

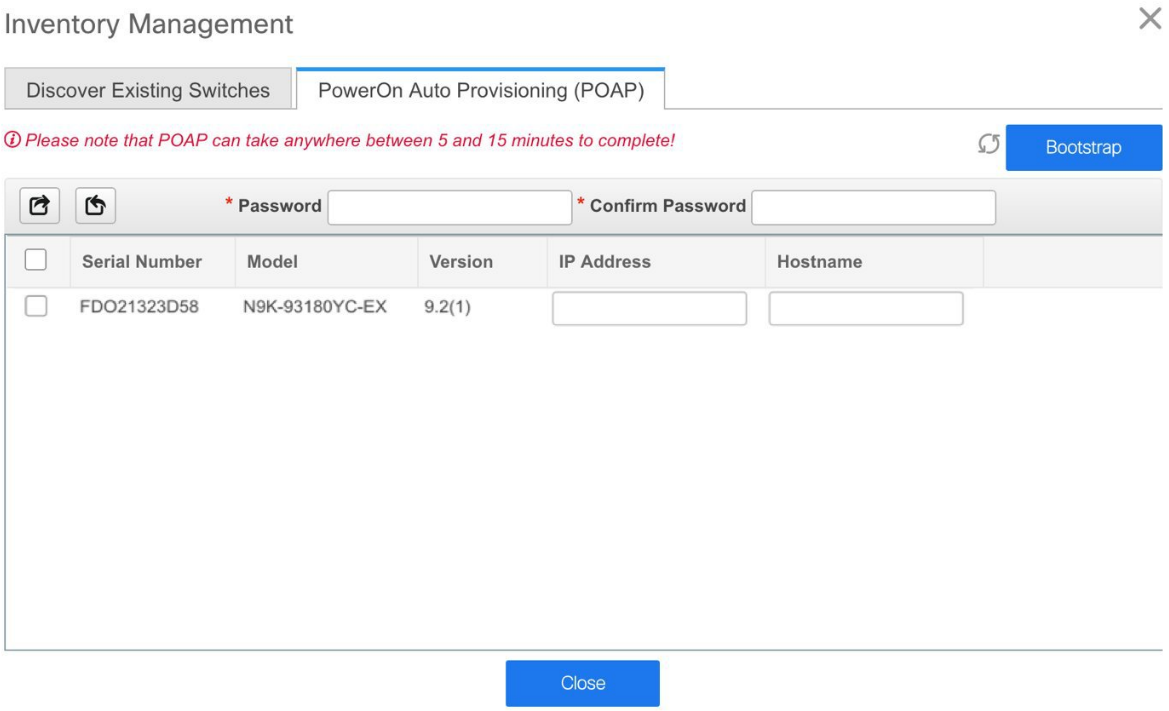

ⓘ *Please note that POAP can take anywhere between 5 and 15 minutes to complete!*

* Password
* Confirm Password

□	Serial Number	Model	Version	IP Address	Hostname
□	FDO21323D58	N9K-93180YC-EX	9.2(1)		

Select the checkbox next to the switch, add switch credentials (such as the IP address, host name and password), and click **Bootstrap** at the top right part of the screen. The fabric builder topology page appears.

DCNM provisions the management IP address and other credentials to the switch. In this simplified POAP process, all ports are opened up.

7. Click **Refresh Topology** to get updated information. The added switch goes through the POAP cycle. Monitor and check the switch for POAP completion.

8. After the added switch completes POAP, the fabric builder topology page is refreshed with the added switch with some physical connections. However, the switch icon is in red color indicating that the fabric is *Out-Of-Sync* and you must click **Save & Deploy** on the fabric builder topology to deploy pending configurations (such as template and interface configurations) onto the switches.

Note For any changes on the fabric that results in the out-of-sync, then you must deploy the changes. The process is the same as explained in the *Discovering Existing Switches* section.

During fabric creation, if you have entered AAA server information (in the **Manageability** tab), you must update the AAA server password on each switch. Else, switch discovery fails.

9. After the pending configurations are deployed, the **Progress** column displays 100% for all switches.

10. Click **Close** to return to the fabric builder topology.

11. Click **Refresh Topology** to view the update. All switches must be in green color indicating that they are functional.

12. The switch and the link are discovered in DCNM. Configurations are built based on various policies (such as fabric, topology, and switch generated policies). The switch image (and other required) configurations are enabled on the switch.

13. In the DCNM GUI, the discovered switches can be seen in the *Standalone* fabric topology. Up to this step, the POAP is completed with basic settings. All the interfaces are set to trunk ports. You must setup interfaces through the **Control > Interfaces** option for any additional configurations, but not limited to the following:

    - vPC pairing.
    - Breakout interfaces.
    - Port channels, and adding members to ports.

Note

- After discovering a switch (new or existing), at any point in time you can provision configurations on it again through the POAP process. The process removes existing configurations and provision new configurations. You can also deploy configurations incrementally without invoking POAP.
- In some instances, after new switches are discovered through POAP, a switch interface is displayed as connected to two interfaces. The additional, incorrect connection is displayed in red, similar to a failed connection.

  To resolve the issue, you must initiate the POAP process again for switches with incorrect interface connections. If still not resolved, perform a layer-by-layer switch discovery during the bring-up phase - spine switches first, then the leaf switches and then the border leaf switches.

You can right-click the switch to view various options:

- **Set Role** - Assign a role to the switch (Spine, Border Gateway, and so on).

  Note Changing of the switch role is allowed only before executing **Save & Deploy**.

- **Mode** - Maintenance and Active/Operational modes.
- **Select for vPC Configuration** - Select a switch for vPC and then select its peer.
- **Manage Interfaces** - Deploy configurations on the switch interfaces.
- **View/Edit Policies** - See switch policies and edit them as required.
- **History** - View per switch deployment history.
- **Deploy Config** - Deploy per switch configurations.
- **Discovery** - You can use this option to update the credentials of the switch, reload the switch, rediscover the switch, and remove the switch from the fabric.

The new fabric is created, the fabric switches are discovered in DCNM, the underlay networks provisioned on those switches, and the configurations between DCNM and the switches are synced. The remaining tasks are:

- Provision interface configurations such as vPCs, loopback interface, and subinterface configurations. [*Interfaces* topic].
- Create overlay networks and VRFs and deploy them on the switches. [*Networks and VRFs Creation and Deployment* section].

Return Material Authorization (RMA)

This section describes how to replace a physical switch in a Fabric when using Cisco DCNM Easy Fabric mode.

Prerequisites

- Fabric is assumed to be up and running, and minimal disruption is desired when replacing the switch. Also, the switch must be replaced with a switch of the same model (ASIC type) and physical port configuration.
- To use the POAP RMA flow, you must configure the fabric for bootstrap (POAP).
- To copy the FEX configurations for the RMA of switches which have FEX deployed, you may need to perform the Save and Deploy operation one or two times.

Guidelines and Limitations

- The switch must be replaced with a switch of the same model (ASIC type) and physical port configuration. If not, the old switch must be removed and a new switch (replacement) added as a new switch into the fabric.

POAP RMA Flow

Procedure

Step 1 Choose **Control > Fabric Builder**.

Step 2 Click the Fabric where you want to perform RMA.

Step 3 Move the device into maintenance mode. To move a device into maintenance mode, right-click on the device, and then choose **Modes > Maintenance Mode**.

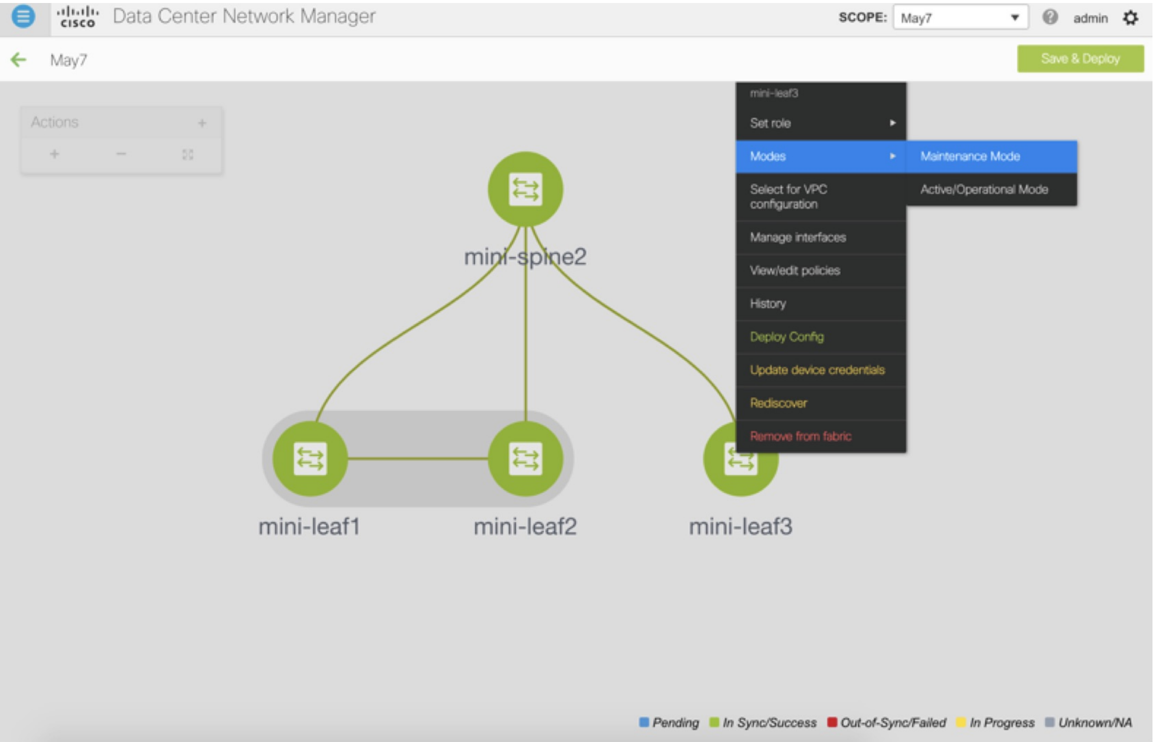

Step 4 Physically replace the device in the network. Physical connections should be made in the same place on the replacement switch as they existed on the original switch.

Step 5 Provision RMA flow and select the replacement device.

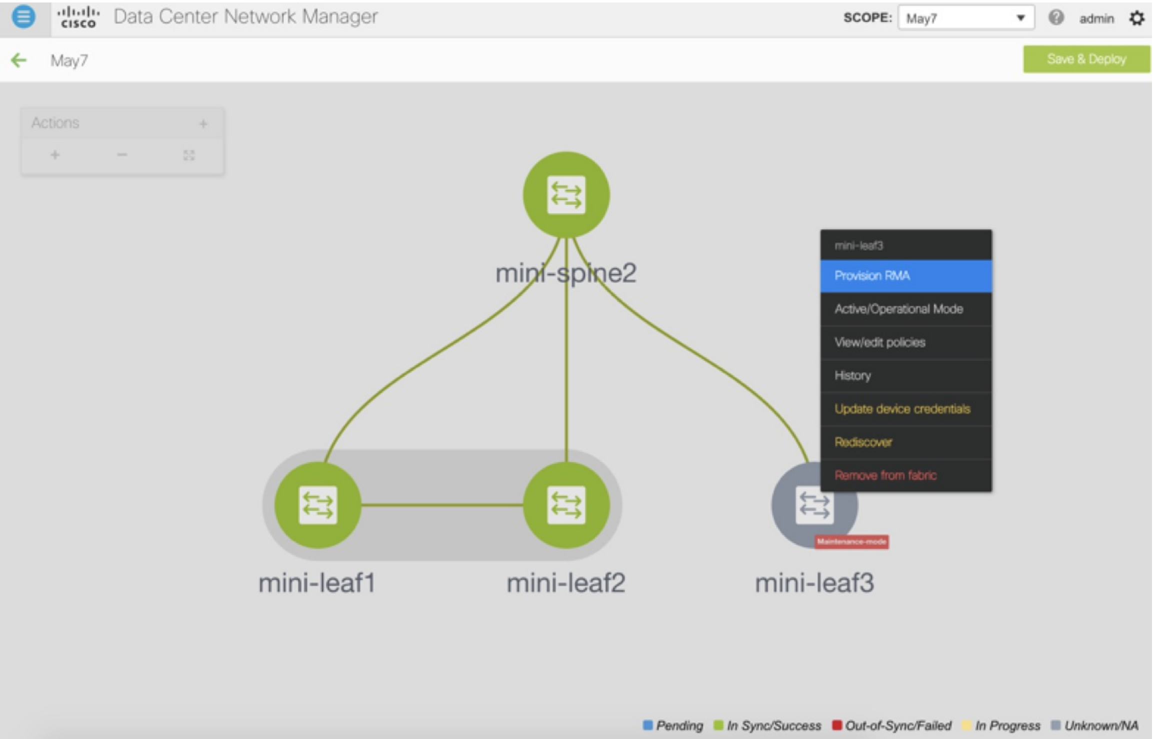

Step 6 The Provision RMA UI will show the replacement device 5-10 minutes after it is powered on.

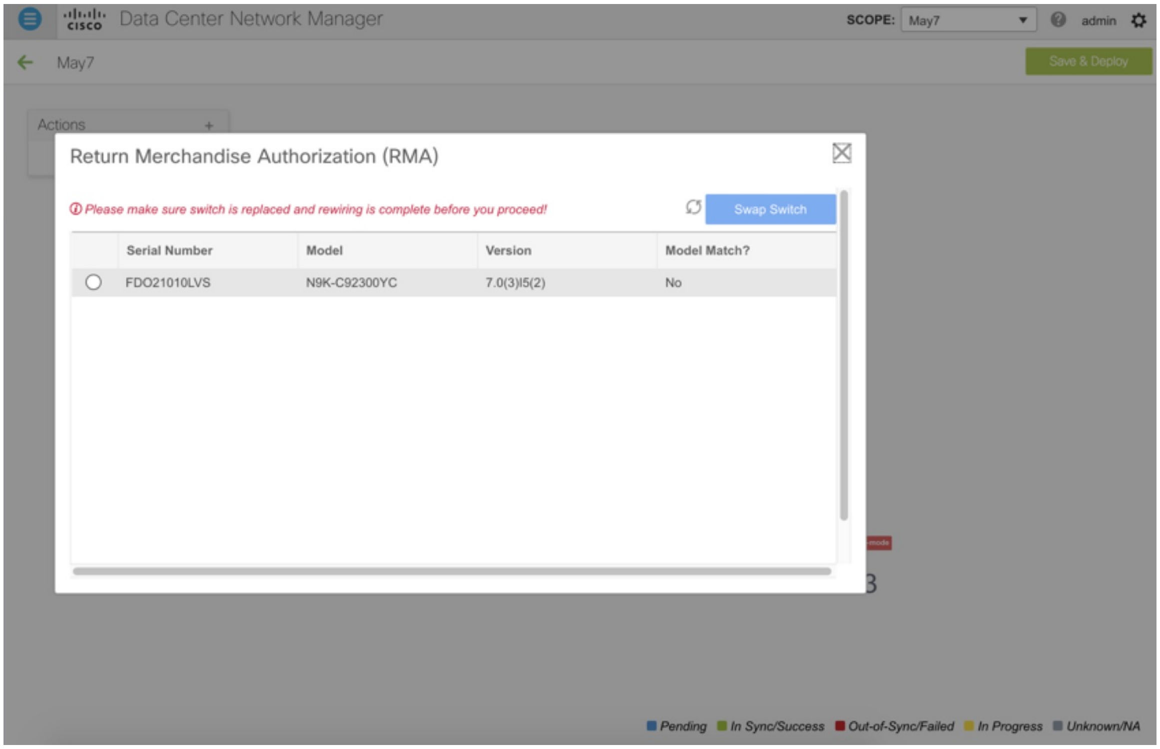

Return Merchandise Authorization (RMA)

ⓘ *Please make sure switch is replaced and rewiring is complete before you proceed!*

	Serial Number	Model	Version	Model Match?
○	FDO21010LVS	N9K-C92300YC	7.0(3)I5(2)	No

🟦 Pending 🟩 In Sync/Success 🟥 Out-of-Sync/Failed 🟨 In Progress ⬜ Unknown/NA

Step 7 Select the correct replacement device and click **Swap Switch**. This begins POAP with the full “expected” configuration for that device. Total POAP time is generally around 10-15 minutes.

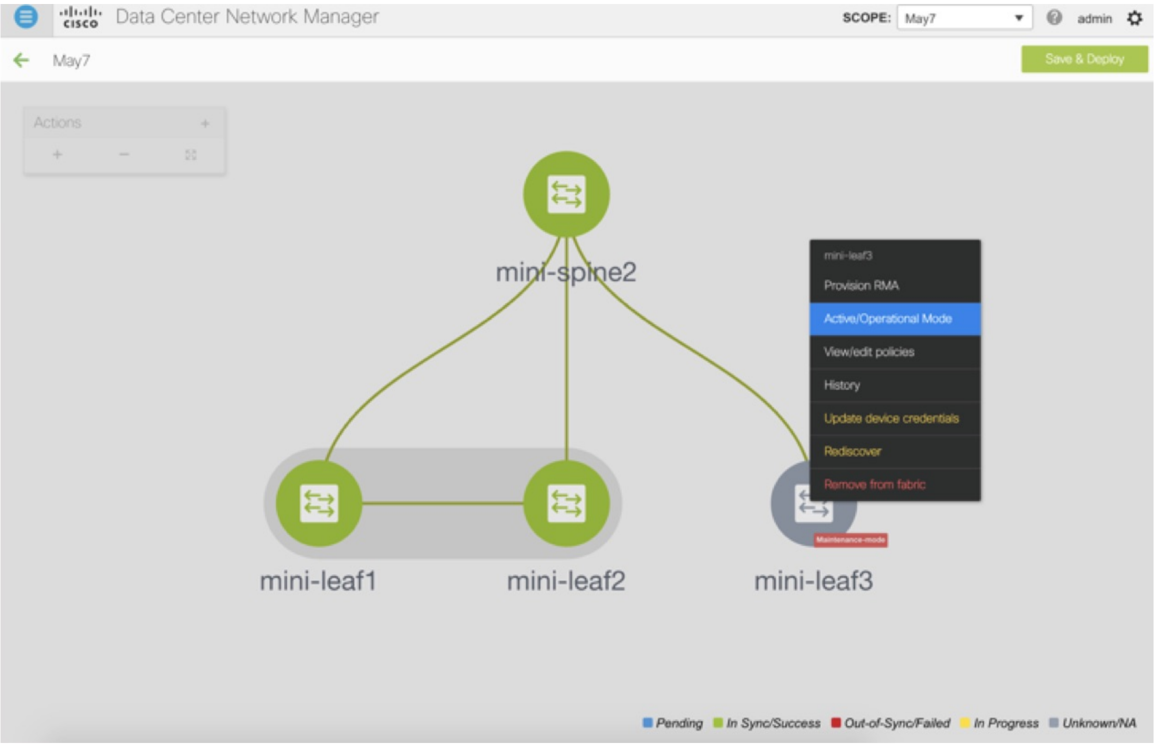

Manual RMA Flow

Use this flow when “Bootstrap” is not possible (or not desired), including cases that are *IPv6 only* for the initial Cisco DCNM 11.0(1) release.

Procedure

Step 1 Place the device in maintenance mode (optional).

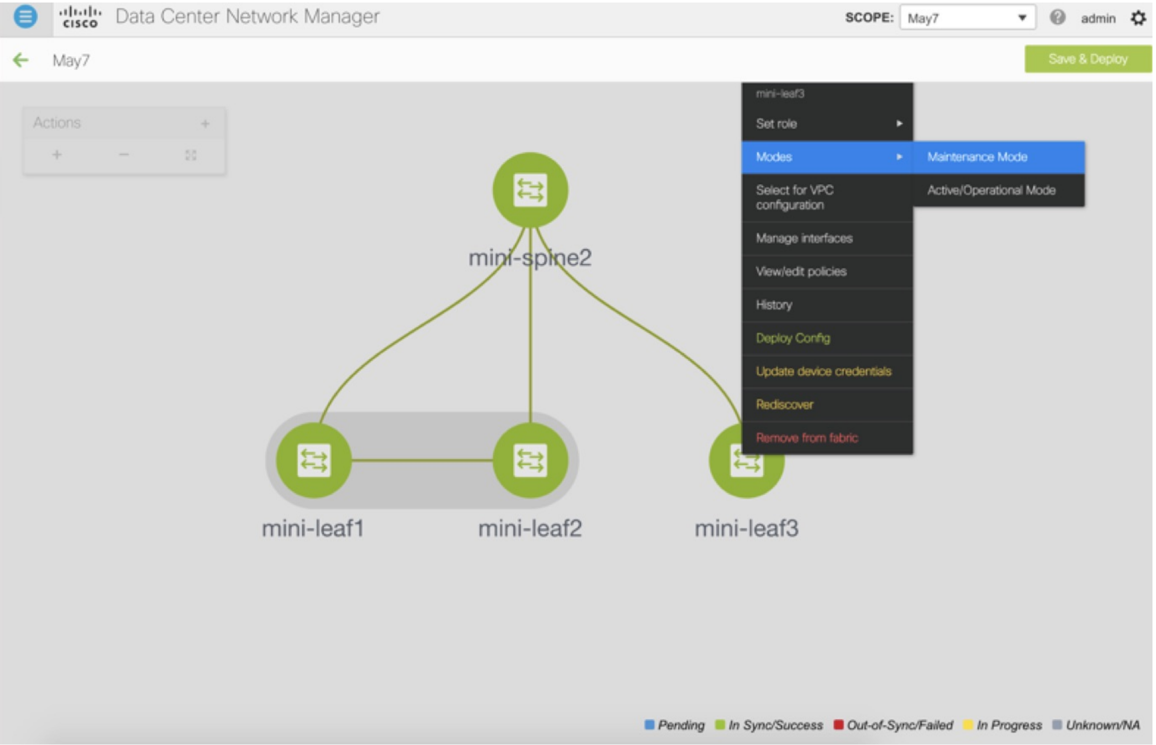

Step 2 Physically replace the device in the network.

Step 3 Log in through Console and set the Management IP and credentials.

Step 4 The Cisco DCNM rediscovers the new device (or you can manually choose **Discovery > Rediscover**).

Step 5 Deploy the expected configuration using **Deploy**.

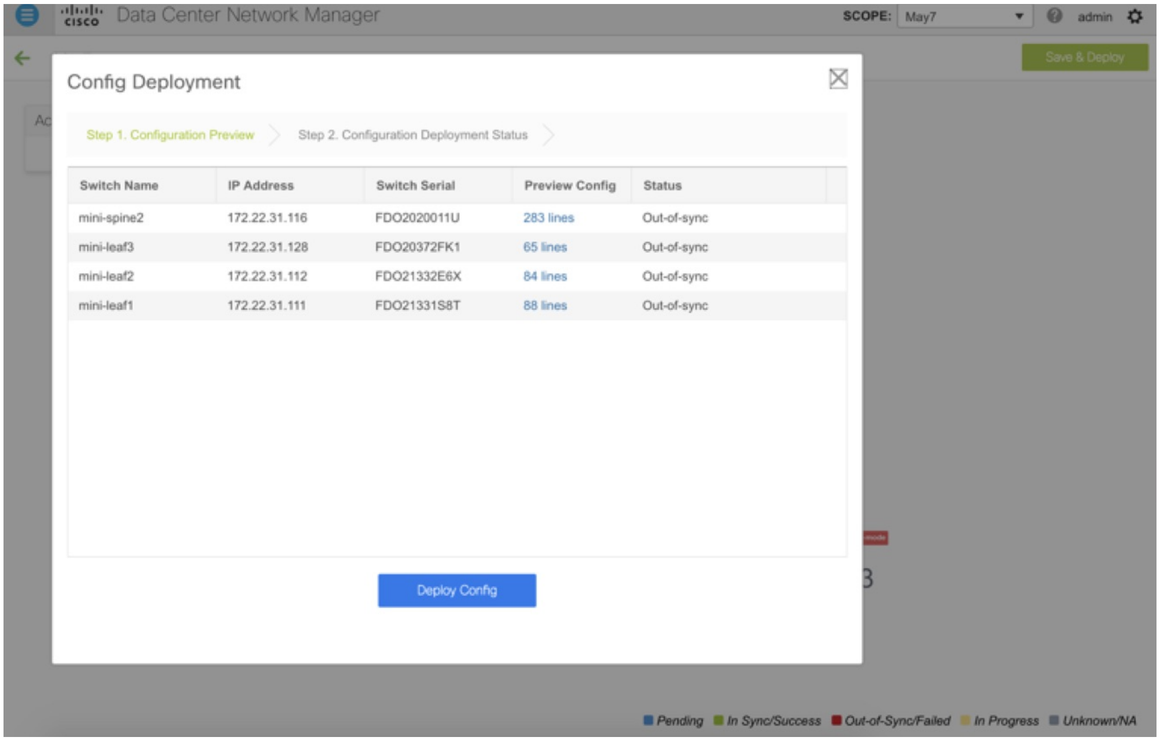

Switch Name	IP Address	Switch Serial	Preview Config	Status
mini-spine2	172.22.31.116	FDO2020011U	283 lines	Out-of-sync
mini-leaf3	172.22.31.128	FDO20372FK1	65 lines	Out-of-sync
mini-leaf2	172.22.31.112	FDO21332E6X	84 lines	Out-of-sync
mini-leaf1	172.22.31.111	FDO21331S8T	88 lines	Out-of-sync

Step 6 Depending on the configuration, if breakout ports or FEX ports are in use, you have to deploy again to completely restore the configuration.

Step 7 After a successful deployment, and the device is "In-Sync," you must move the device back to Normal Mode.

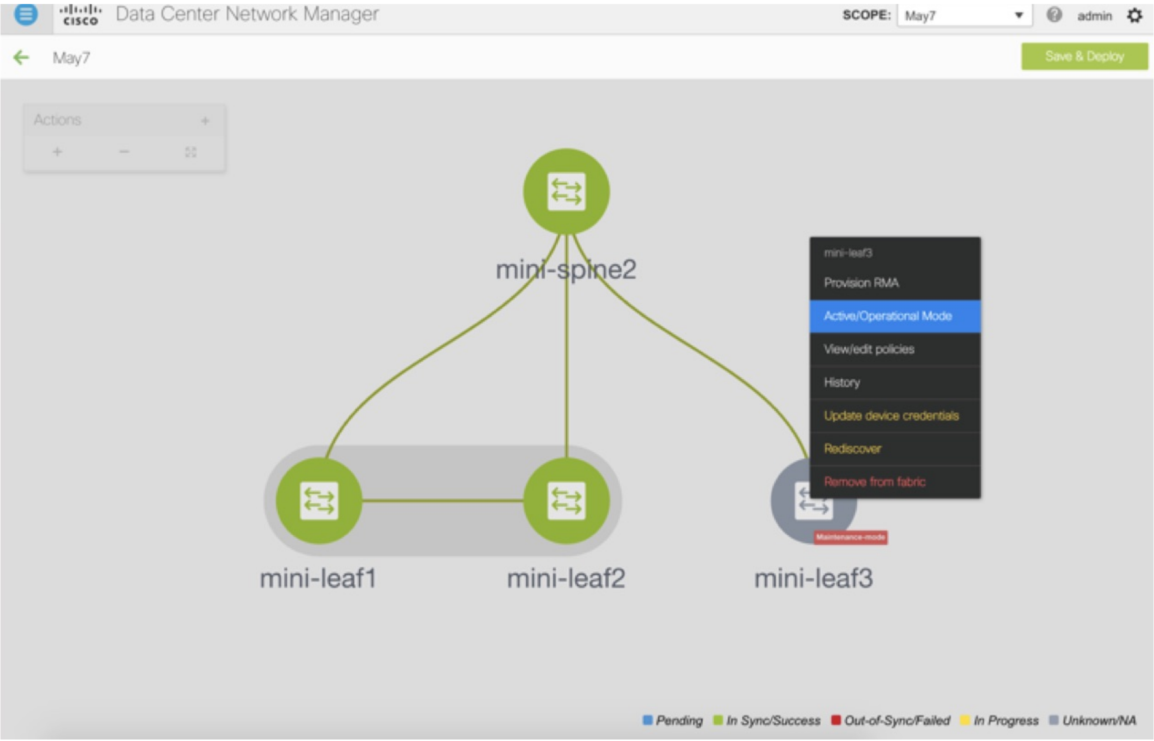

RMA for User with Local Authentication

Note This task is only applicable to non-POAP switches.

Use the following steps to perform RMA for a user with local authentication:

Procedure

Step 1 After the new switch comes online, SSH into the switch and reset the local user passwords with the cleartext password using the “username” command. Reset the local user passwords to resync the SNMP password. The password is stored in the configuration file in a nontransferable form.

Step 2 Wait for the RMA to complete.

Step 3 Update Cisco DCNM switch_snmp_user policy for the switch with the new SNMP MD5 key from the switch.

Interfaces

The Interfaces option displays all the interfaces that are discovered for the switch, Virtual Port Channels (vPCs), and intended interfaces missing on the device.

You can use the following functions:

- Create, deploy, view, edit and delete a port channel, vPC, Straight-through FEX, Active-Active FEX, loopback, and subinterface.
- Create breakout and unbreakout ports.
- Shut down and bring up interfaces.
- Rediscover ports and view interface configuration history.
- Apply host policies on interfaces and vPCs. For example, int_trunk_host_11_1, int_access_host_11_1, and so on.
- View interface information such as its admin status, operation status, reason, policy, speed, MTU, mode, VLANs, IP/Prefix, VRF, port channel, and the neighbor of the interface.

  Note The **Neighbor** column provides details of connected switches that are discovered, intent links, and Virtual Machine Manager (VMM) connectivity. You can navigate to the **Switch** dashboard of the corresponding switch by clicking it. However, intent links and VMM links are not hyperlinked and you cannot navigate to the corresponding dashboard.

  The **Status** column displays the following statuses of an interface:

  - Blue: Pending
  - Green: In Sync/Success

- Red: Out-of-Sync/Failed
- Yellow: In Progress
- Grey: Unknown/NA

You can filter and view information for any of the given fields (such as Device Name). The following table describes the buttons that appear on this page.

Note

- Ensure that appropriate configurations are deployed through the Fabric Builder option before deploying from the Interfaces option, including proper vPC pair configurations. If you add or edit an interface before fabric deployment, the configuration may fail on the device.
- You can also manage interfaces from the Fabric Builder topology screen. Right click the switch and on the Manage Interfaces option. You can manage the interfaces per switch. If the switch is part of a vPC Pair, then interfaces from both peers are displayed on the page.
- Deploy any underlays including vPC Pairing in the fabric before deploying any configurations from the interface manager.

Field	Description
Add	Allows you to add a logical interface such as a port channel, vPC, Straight-through FEX, Active-Active FEX, loopback and subinterface.
Breakout, Unbreakout	Allows you to *breakout* an interface or unbreakout interfaces that are in *breakout* state.
Edit	Allows you to edit and change policies that are associated with an interface.
Delete	Allows you to delete a logical interface that is created from the Interfaces screen. An interface having a policy that is attached from an overlay and underlay cannot be deleted.
No Shutdown	Allows you to enable an interface (no shutdown or admin up).
Shutdown	Allows you to shut down the interface.
Show	Allows you to display the interface show commands. A show command requires show templates in the template library.
Rediscover	Allows you to rediscover or recalculate the compliance status on the selected interfaces.
Interface History	Allows you to display the interface deployment history details.

Field	Description
Deploy	Allows you to deploy or redeploy saved interface configurations.

This section contains the following:

Adding Interfaces

To add the interfaces from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Interfaces**.

You see the **Scope** option at the top right part of the screen. If you want to view interfaces for a specific fabric, select the fabric window from the list.

External Fabric: On interfaces belonging to an external fabric, you cannot perform any operation except the *show* and *rediscovery* operations.

Step 2 Click **Add** to add a logical interface.

The **Add Interface** window appears.

Step 3 In the **Type** field, choose the type of the interface.

For example, port channel, Straight-through FEX, Active-Active FEX, vPC, loopback, and subinterface.

- When you create a port channel through DCNM, add interfaces of the same speed. A port channel that is created from interfaces of varying speeds will not come up. For example, a port channel with two *10 Gigabit Ethernet* ports is valid. However, a port channel with a *10-Gigabit Ethernet + 25-Gigabit Ethernet* port combination is not valid.
- To add vPC hosts, you must designate vPC switches in the fabric topology (through the Fabric Builder) and deploy the configurations. Once the vPC pair is created from Fabric Builder, it appears in the Select a vPC pair drop-down box.

  You can create a vPC using the **vpc_trunk_host** policy. However, you cannot associate a VLAN to the vPC (as an allowed VLAN) from the **Interfaces** option. When you deploy an overlay network on the fabric switches, you should associate the corresponding VLAN to port channels (of the vPC domain), as applicable. Refer the *Networks Deployment in the Standalone Fabric* topic for network deployment details.
- When adding a subinterface, you must select a routed interface from the interface table before clicking the Add button.

Step 4 In the **Select a Device** field, choose the device.

In the case of vPC or Active to Active FEX, select the vPC switch pair.

Step 5 In the **Number** field, on selection of Interface Type and device or vPC pair, this field is automatically populated from the Resource Manager.

You can override this value. The new value is used only if it is available in the Resource Manager pool. Else, it results in an error.

Step 6 In the **Policy** field, you can select the policy to be applied on an interface.

The field only lists the Interface Python Policy with tag *interface_edit_policy* and filtered based on the interface type.

You must not create a **_upg** interface policy. For example, you should not create a policy using the **vpc_trunk_host_upg**, **port_channel_aa_fex_upg**, **port_channel_trunk_host_upg**, and **trunk_host_upg** options.

Step 7 Click **Save** to save the configurations.

Only saved configurations are pushed to the device. While adding the interface, you can save the configuration only once. Successive saves results in the *Resource could not be allocated* error. Once saved, you can change the configurations by editing the interface.

Step 8 (Optional) Click the **Preview** option to preview the configurations to be deployed.

Step 9 Click **Deploy** to deploy the specified logical interface.

The newly added interface appears in the screen.

Breakout or Unbreakout: You can break out and unbreak out an interface by using the **breakout** option at the top left part of the screen.

Editing Interfaces

To edit the interfaces from the Cisco DCNM Web UI, perform the following steps:

Note You can edit the interface if it does not have an overlay or underlay policy attached. The **Edit Interface** allows you to change the policy and add or remove an interface from a port channel or vPC.

Procedure

Step 1 Choose **Control > Interfaces**.

You can break out and unbreak out an interface by using the breakout option at the top left part of the screen.

Step 2 Select the interface check box to edit an interface or vPC.

Select corresponding check boxes for editing multiple interfaces. You cannot edit multiple port channels and vPC. You cannot edit interfaces of different types at the same time.

Step 3 Click **Edit** to edit an interface.

The variables that are shown in the **Edit Configuration** window are based on the template and its policy. Select the appropriate policy. Preview the policy, save it and deploy the same. This window lists only Interface Python Policy with the tag *interface_edit_policy* and filtered based on the interface type.

In a vPC setup, the two switches are in the order the switch names are displayed in the edit window. For example, if Switch Name is displayed as *LEAF1:LEAF2*, then Leaf1 is peer switch one and Leaf2 is peer switch two.

Deleting Interfaces

To delete the interfaces from the Cisco DCNM Web UI, perform the following steps:

Note This option allows you to delete only logical ports, port channels, and vPCs. You can delete the interface if it does not have overlay or underlay policy attached.

When a port channel or vPC is removed, the corresponding member ports get the default policy associated. The Default Policy can be configured in `server.properties` file.

Procedure

Step 1 Choose **Control** > **Interfaces**.

Step 2 Select the interfaces.

Step 3 Click **Delete** to delete the interface.

Shutting Down and Bringing Up Interfaces

To shut down and bring up the interfaces from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control** > **Interfaces**.

Step 2 Select the interfaces that you want to shut down or bring up.

Step 3 Click **Shutdown** to disable the selected interfaces. For example, you may want to isolate a host from the network or a host that is not active in the network.

Step 4 Click **No Shutdown** to bring up the selected interfaces.

Viewing Interface Configuration

To view the interface configuration commands and execute them from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control** > **Interfaces**.

Select the interface whose configurations you want to view.

Step 2 In the **Interface Show Commands** window, select the action from the **Show** drop-down box and click **Execute**. The interface configurations are displayed in the **Output** section, at the right of the screen.

For Show commands, you must have corresponding *show* templates that are defined in the **Template Library**.

Rediscovering Interfaces

To rediscover the interfaces from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Interfaces**.

Step 2 Select the interfaces that you want to rediscover.

Step 3 Click **Rediscover** to rediscover the selected interfaces. For example, after you edit or enable an interface, you can rediscover the interface.

Viewing Interface History

To view the interface history from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Interfaces**.

Step 2 Select the interface.

Step 3 Click **Interface History** to view the configuration history on the interface.

Step 4 Click **Status** to view each command that is configured for that configuration instance.

Deploying Interface Configurations

To deploy the interface configuration from the Cisco DCNM Web UI, perform the following steps:

Procedure

Choose **Deploy** to deploy and redeploy configurations that are saved for an interface.

You can select multiple interfaces and deploy pending configurations.

Networks and VRFs Creation and Deployment in a Standalone Fabric

The steps for overlay networks and VRFs provisioning are:

1. Create networks and VRFs for the fabric.
2. Deploy the networks and VRFs on the fabric switches.

Note The undeployment and deletion of overlay networks and VRFs are explained after the explanation of deployment. Finally, creation of external fabrics and fabric extensions from VXLAN to external fabrics are documented.

The two steps are explained:

Create Networks for the Fabric

1. Click **Control > Networks & VRFs** (under Fabrics submenu). The LAN Fabric Provisioning page comes up.
2. Click **Continue**. The **Select a Fabric** page is displayed.

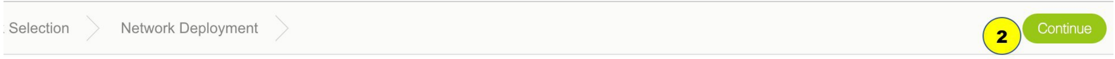

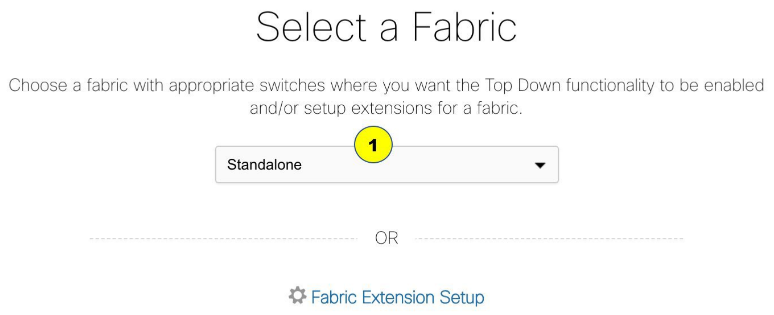
Select a Fabric
Choose a fabric with appropriate switches where you want the Top Down functionality to be enabled and/or setup extensions for a fabric.
1
Standalone
OR
Fabric Extension Setup

3. From the **Select a Fabric** drop-down list, select the fabric *Standalone*, and click **Continue** on the top right part of the screen. The Networks page is displayed. This page lists the networks that are created for the fabric. Initially, this page will not have any entries.

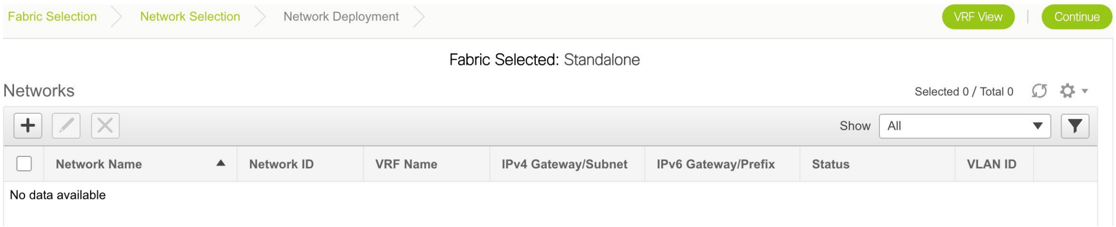

Fabric Selected: Standalone

Networks

Selected 0 / Total 0

Show All

☐	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
No data available							

4. Click the + button at the top left part of the screen (under **Networks**) to add networks to the fabric. The Create Network screen comes up. Most of the fields are autopopulated.

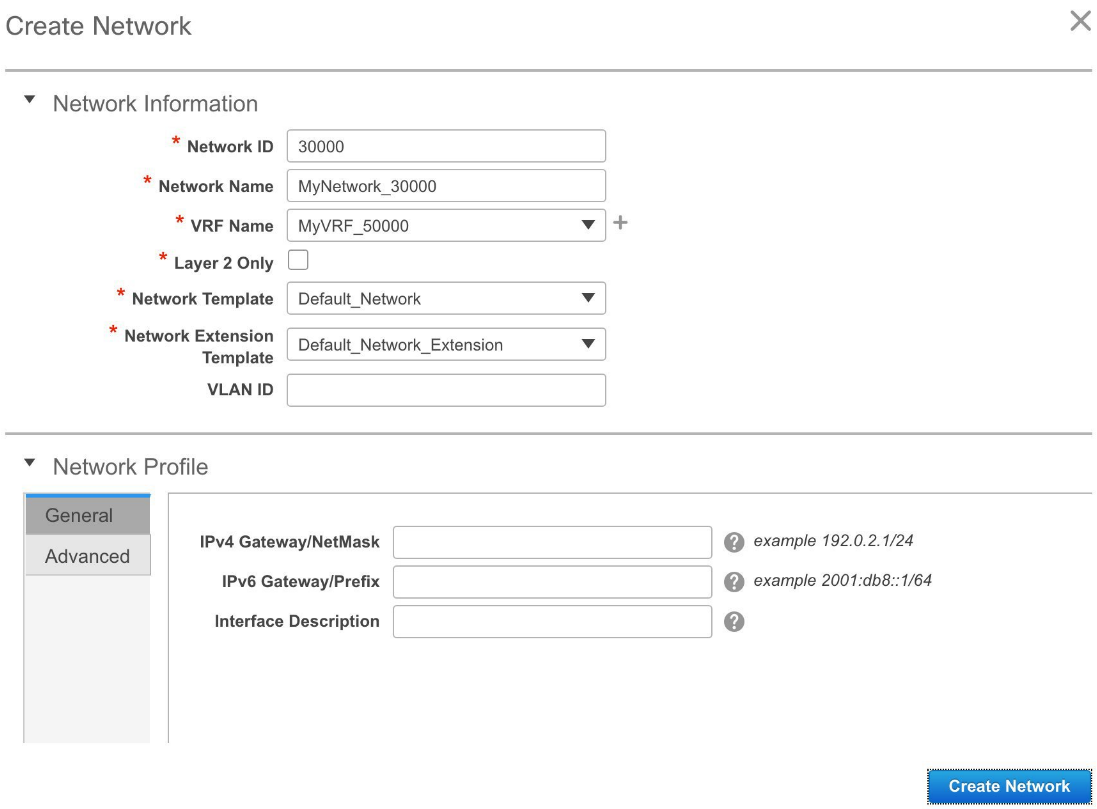

Create Network

Network Information

* Network ID	30000
* Network Name	MyNetwork_30000
* VRF Name	MyVRF_50000
* Layer 2 Only	☐
* Network Template	Default_Network
* Network Extension Template	Default_Network_Extension
VLAN ID	

Network Profile

IPv4 Gateway/NetMask		*example 192.0.2.1/24*
IPv6 Gateway/Prefix		*example 2001:db8::1/64*
Interface Description		

The fields in this screen are:

Network ID and **Network Name**: Specifies the Layer 2 VNI and name of the network. The network name should not contain any white spaces or special characters except underscore (_) and hyphen (-). The corresponding Layer 3 VNI (or VRF VNI) is generated along with VRF creation.

VRF Name: Allows you to select the Virtual Routing and Forwarding (VRF).

When no VRF is created, this field appears blank. If you want to create a new VRF, click the + button. The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:).

Layer 2 Only: Specifies whether the network is Layer 2 only.

Network Template: Allows you to select a network template, and is only applicable for leaf switches.

Network Extension Template: Allows you to extend this network to another fabric, based on the extension method that you select. The methods are *VRF Lite, Multi Site*, and so on. The template is applicable for border leaf switches and BGWs.

VLAN ID: Specifies the corresponding tenant VLAN ID for the network.

Network Profile section contains the *General* and *Advanced* tabs.

General tab

IPv4 Gateway/NetMask: Specifies the IPv4 address with subnet.

IPv6 Gateway/Prefix: Specifies the IPv6 address with subnet.

Specify the anycast gateway IP address for transporting the L3 traffic from a server belonging to MyNetwork_30000 and a server from another virtual network. By default the anycast gateway IP address is the same for MyNetwork_30000 on all switches of the fabric that have the presence of the network.

Interface Description: Specifies the description for the interface. This interface is a switch virtual interface (SVI).

Advanced tab: Optionally, specify the advanced profile settings by clicking the **Advanced** tab:

- ARP Suppression
- Ingress Replication

Note Ingress Replication is a read-only option in the Advanced tab. Changing the fabric setting updates the field.

- Multicast Group Address
- DHCPv4 Server
- DHCPv4 Server VRF
- MTU for the L3 interface

A sample of the Create Network page:

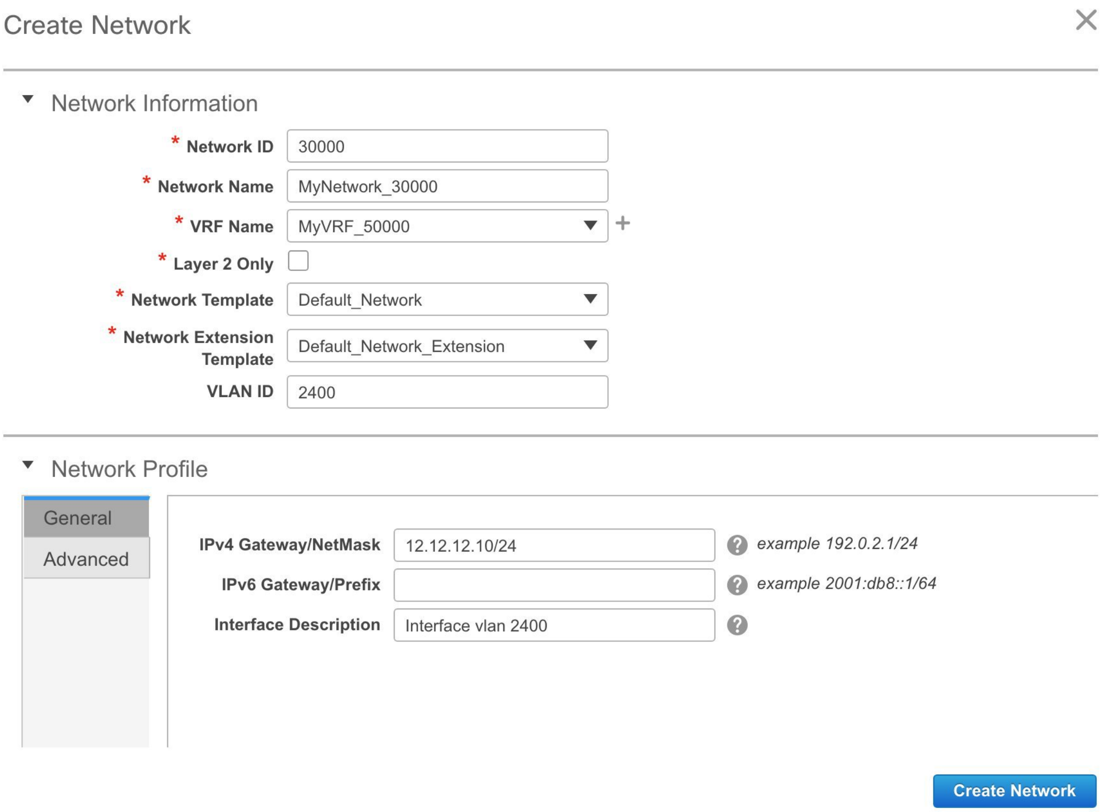

Create Network

Network Information

* Network ID	30000
* Network Name	MyNetwork_30000
* VRF Name	MyVRF_50000
* Layer 2 Only	□
* Network Template	Default_Network
* Network Extension Template	Default_Network_Extension
VLAN ID	2400

Network Profile

IPv4 Gateway/NetMask	12.12.12.10/24	example 192.0.2.1/24
IPv6 Gateway/Prefix		example 2001:db8::1/64
Interface Description	Interface vlan 2400	

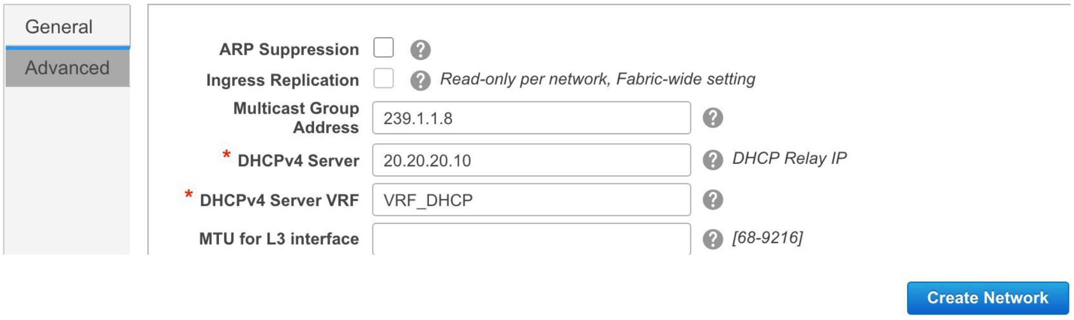

ARP Suppression	□	
Ingress Replication	□	Read-only per network, Fabric-wide setting
Multicast Group Address	239.1.1.8	
* DHCPv4 Server	20.20.20.10	DHCP Relay IP
* DHCPv4 Server VRF	VRF_DHCP	
MTU for L3 interface		[68-9216]

5. Click **Create Network**. A message appears at the bottom right part of the screen indicating that the network is created.

   The new network appears on the **Networks** page that comes up.

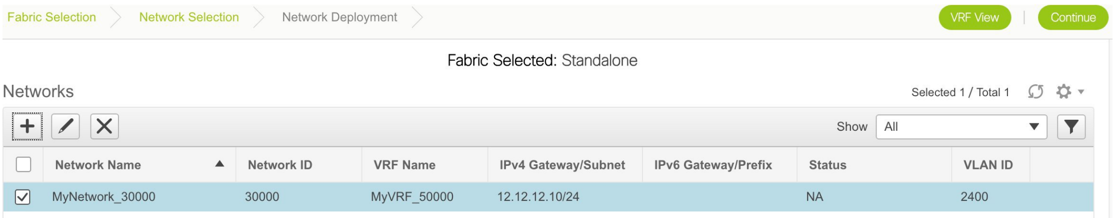

Fabric Selected: Standalone

Networks

Selected 1 / Total 1

Show All

	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
☑	MyNetwork_30000	30000	MyVRF_50000	12.12.12.10/24		NA	2400

The Status is *NA* since the network is created but not yet deployed on the switches. Now that the network is created, you can create more networks if needed and deploy the networks on the devices in the fabric.

Create VRFs for a Standalone Fabric

1. From the Networks page, click the **VRF View** button at the top right part of the screen to create VRFs.

   (If you have freshly logged in to DCNM, do the following:

   Click **Control > Networks & VRFs**.

   Click **Continue** in the LAN Fabric Provisioning page.

   Choose the fabric (*Standalone*) from the drop-down list and click **Continue** to reach the Networks page.

   Click **VRF View** at the top right part of the Networks page**)**.

   The VRFs page comes up. The page lists the list of VRFs created for the fabric. Initially, this page has no entries. One VRF is already created for this fabric. Let us create one more VRF.

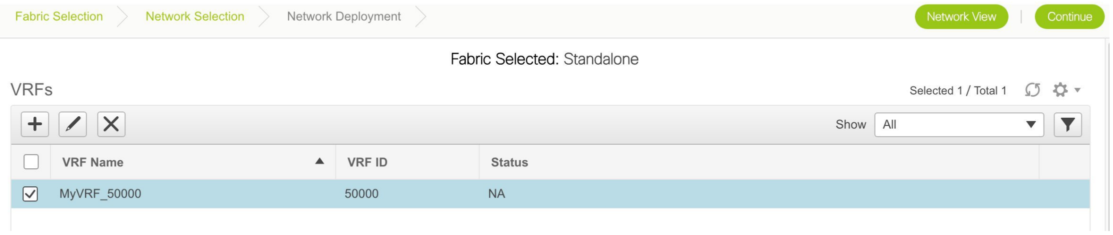

☐	VRF Name	VRF ID	Status
☑	MyVRF_50000	50000	NA

2. Click the + button to add VRFs to the *Standalone* fabric. The Create VRF screen comes up. Most of the fields are autopopulated.

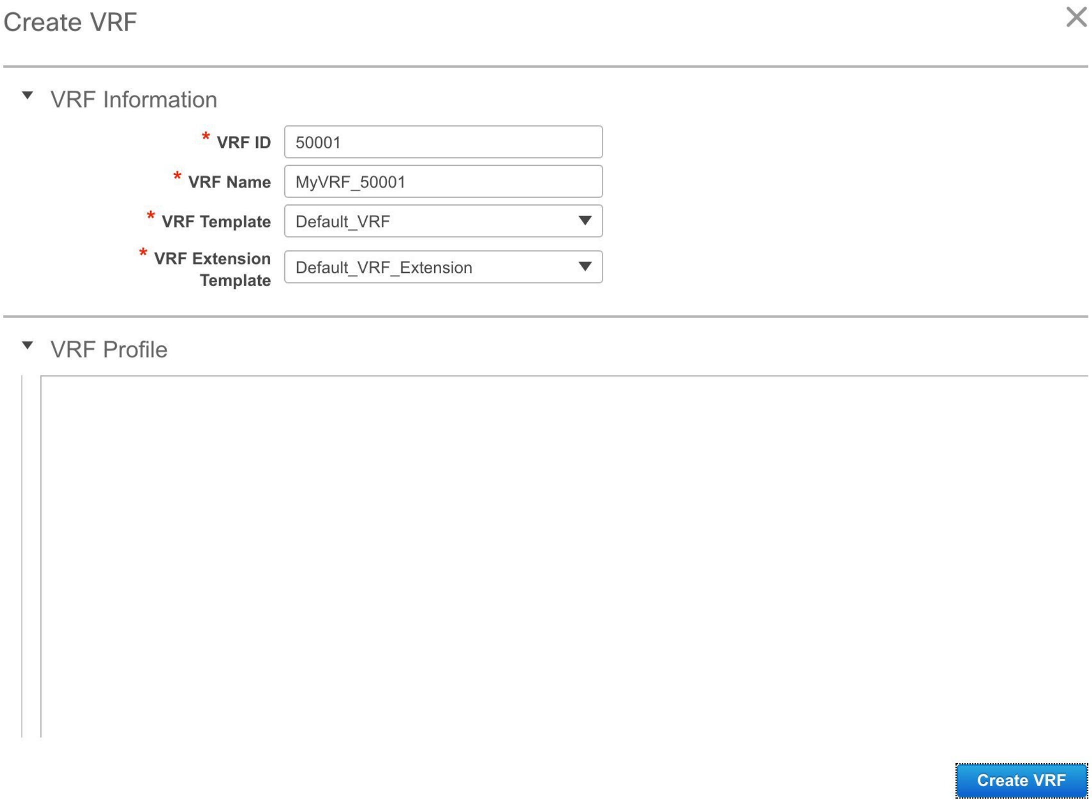

* VRF ID	50001
* VRF Name	MyVRF_50001
* VRF Template	Default_VRF
* VRF Extension Template	Default_VRF_Extension

The fields in this screen are:

VRF ID and **VRF Name**: The ID and name of the VRF.

Note For ease of use, the VRF creation option is also available while you create a network.

VRF Template: This template is applicable for VRF creation, and only applicable for leaf switches.

VRF Extension Template: The template is applicable when you extend the VRF to other fabrics, and is applicable for border leaf switches and border gateways.

3. Click **Create VRF**.

   The *MyVRF_50001* VRF is created and appears on the VRFs page.

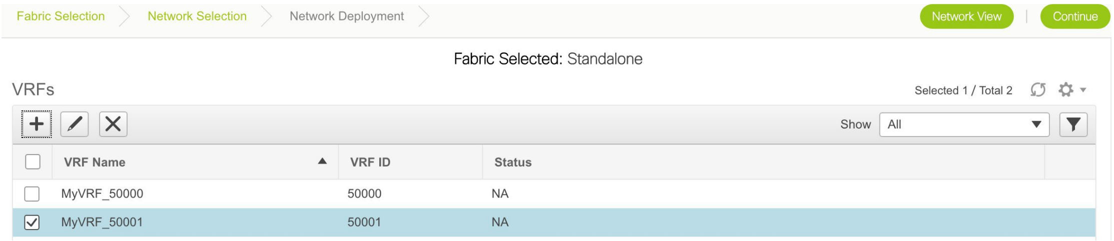

Fabric Selected: Standalone

☐	VRF Name ▲	VRF ID	Status
☐	MyVRF_50000	50000	NA
☑	MyVRF_50001	50001	NA

Networks Deployment in the Standalone Fabric

Before you begin: Ensure that you have created networks for the fabric.

1. Go to the Select a Fabric page.

   (To go to the Select a Fabric page do one of the following:

   Click **Fabric Selection** at the top left part of the screen.

   OR

   From the main menu, click **Control > Networks & VRFs** and click **Continue** in the LAN Fabric Provisioning page**)**.

2. Click *Standalone* from the drop-down list and click **Continue** on the top right part of the screen. The Networks page comes up.

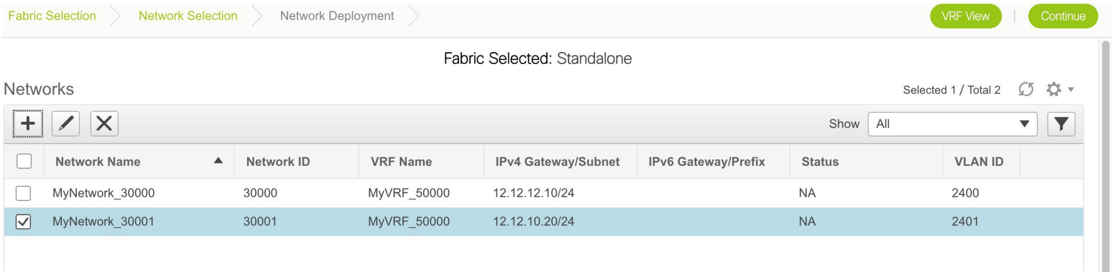

Fabric Selected: Standalone

☐	Network Name ▲	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
☐	MyNetwork_30000	30000	MyVRF_50000	12.12.12.10/24		NA	2400
☑	MyNetwork_30001	30001	MyVRF_50000	12.12.10.20/24		NA	2401

The list of networks in the fabric are displayed on the page. The network deployment status is *NA* since the networks have not been deployed on any switch.

Note You can edit or delete networks from this screen. You can only edit the **Network Profile** section at the bottom part of the screen.

3. Select networks that you want to deploy. In this case, select the checkboxes next to both the networks and click **Continue** at the top right part of the screen.

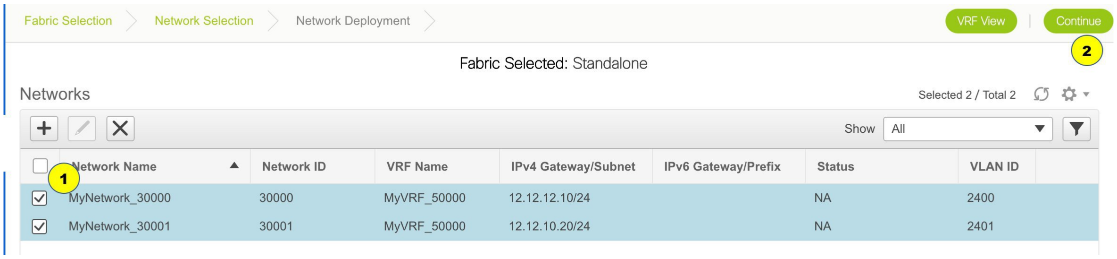

Fabric Selected: Standalone

☐	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
☑	MyNetwork_30000	30000	MyVRF_50000	12.12.12.10/24		NA	2400
☑	MyNetwork_30001	30001	MyVRF_50000	12.12.10.20/24		NA	2401

The Network Deployment page appears. On this page, you can see the network topology of the Standalone fabric.

You can deploy networks simultaneously on multiple switches. The selected devices should have the same role (Leaf, Border Gateway, and so on).

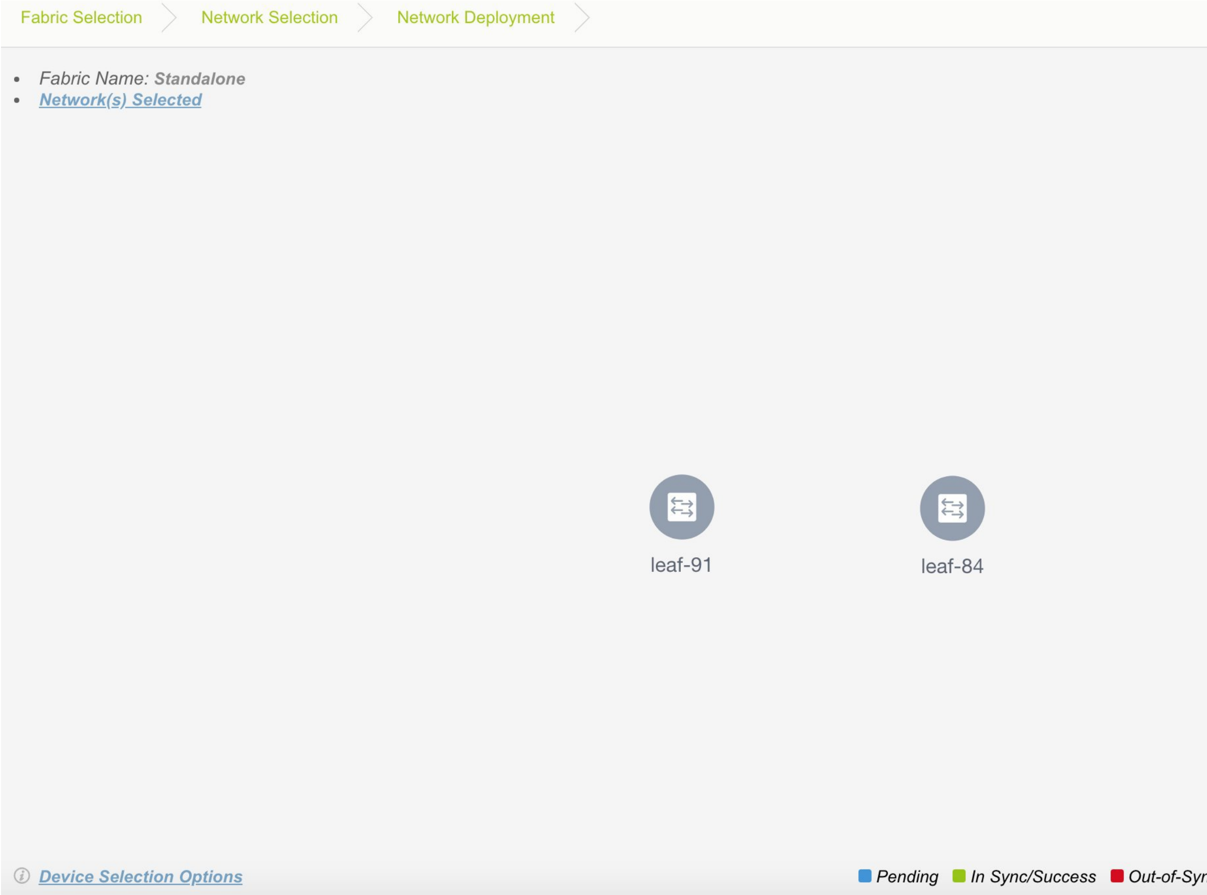

- *Fabric Name: Standalone*
- *Network(s) Selected*

At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly (Blue for *Pending* state, yellow for *In Progress* when the provisioning is in progress, green when successfully deployed, and so on).

The overlay networks (/VRFs) provisioning status is context-specific. It is a combination of networks that you chose for provisioning and the relevant switches in the topology. In this example, it means that the networks *MyNetwork_30000* and *MyNetwork_30001* are yet to be deployed on any switch in this fabric.

You can move the topology around the screen by clicking the left mouse button on the screen and moving it in the direction you desire. You can enlarge or shrink the switch icons proportionately by moving the cursor roller. You can also use corresponding alternatives on the touchpad.

4. Double-click a switch (or use the *Multi-Select* option) to deploy the networks on it. For deployment of networks on multiple switches (like in this case, deploying *MyNetwork_30000* and *MyNetwork_30001* on leaf switches leaf-84 and leaf-91), do the following:

    1. Click *Multi-Select* from the panel at the top right part of the screen. The topology freezes to a static state.

    2. Drag the cursor across the switches.

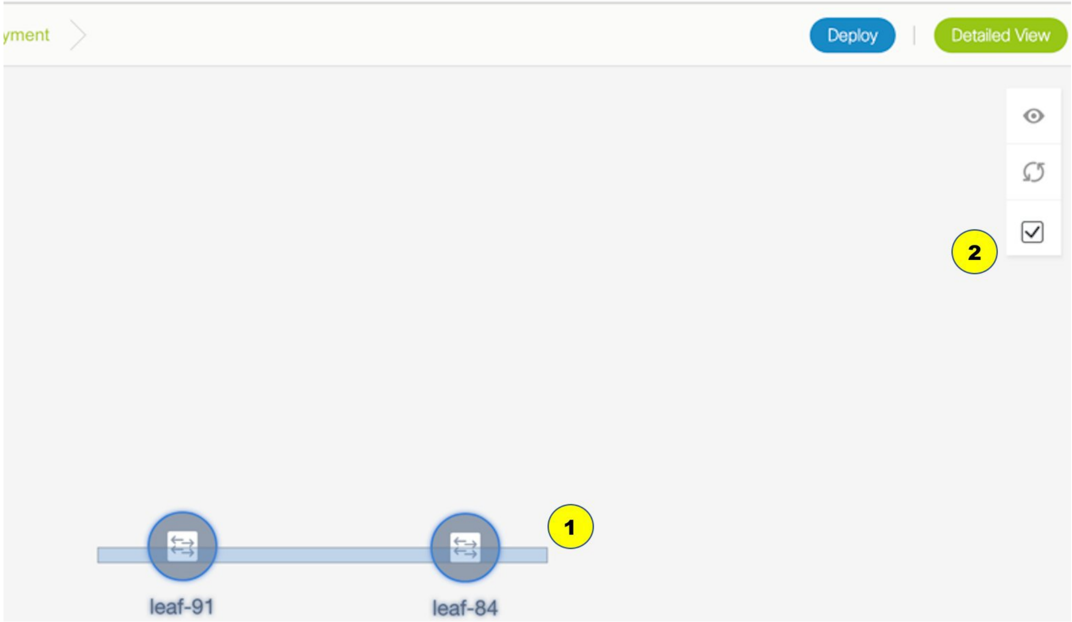

Immediately, the Switches Deploy screen (for networks) appears.

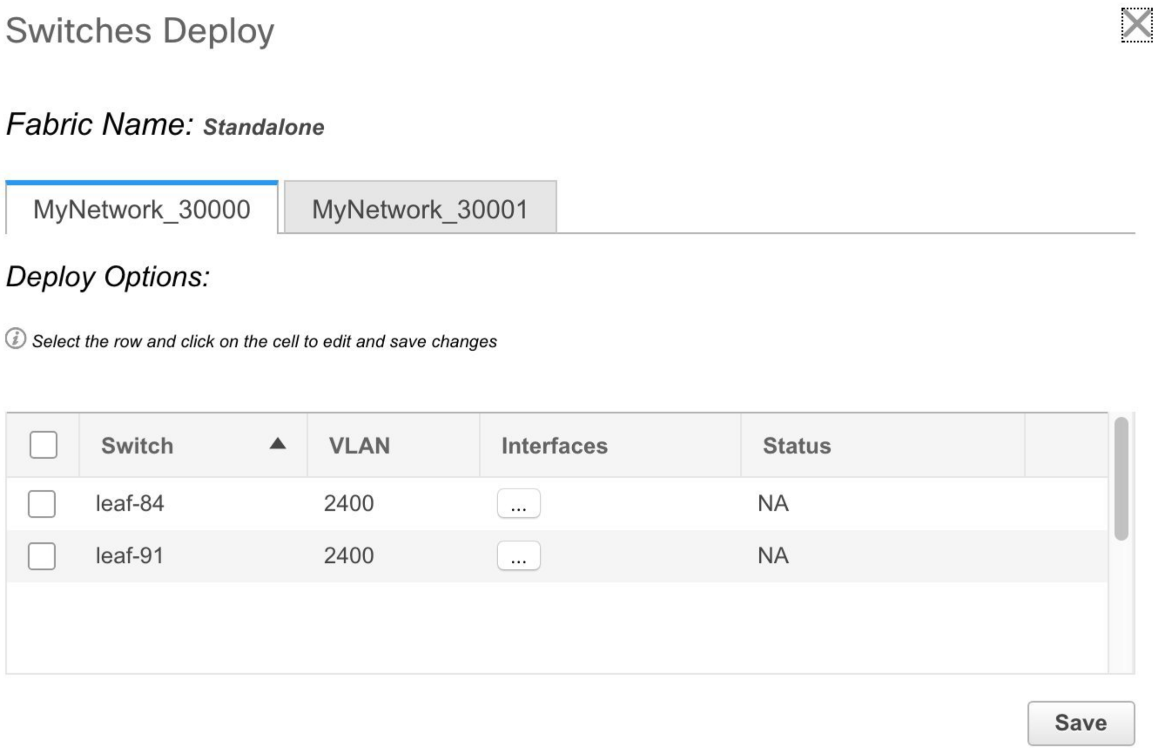

Switches Deploy

Fabric Name: Standalone

Deploy Options:

Select the row and click on the cell to edit and save changes

□	Switch ▲	VLAN	Interfaces	Status
□	leaf-84	2400	...	NA
□	leaf-91	2400	...	NA

A tab represents each network (the first network, *MyNetwork_30000*, is displayed by default) that is being deployed. In each network tab, the switches are displayed. Each row represents a switch.

Click the checkbox next to the **Switch** column to select the switches. Both the switch check boxes are selected automatically. The network *MyNetwork_30000* is ready to be provisioned on the switches leaf-84 and leaf-91.

Select the other network tab and make the same selections.

5. Click **Save** (at the bottom right part of your screen) to save the configurations.

Note Addition and removal of interfaces are displayed in the **Interfaces** column of the Switches Deploy screen. Though the interface-related updates (like addition or removal of trunk ports) are provisioned on the switches, the correct configurations will not reflect in the preview screen. When you add or remove a trunk or access port, the preview shows the addition or removal of configurations for the interface under that network.

The topology screen comes up again. Click *Refresh* in the vertical panel at the top right part of the screen. The blue color on the switch icons indicates that the deployment is pending.

Preview the configurations by clicking *Preview* (the eye icon above the Multi-Select option). Since *MyNetwork_30000* and *MyNetwork_30001* are networks of VRF *50000*, the configurations contain VRF configurations followed by the network configurations.

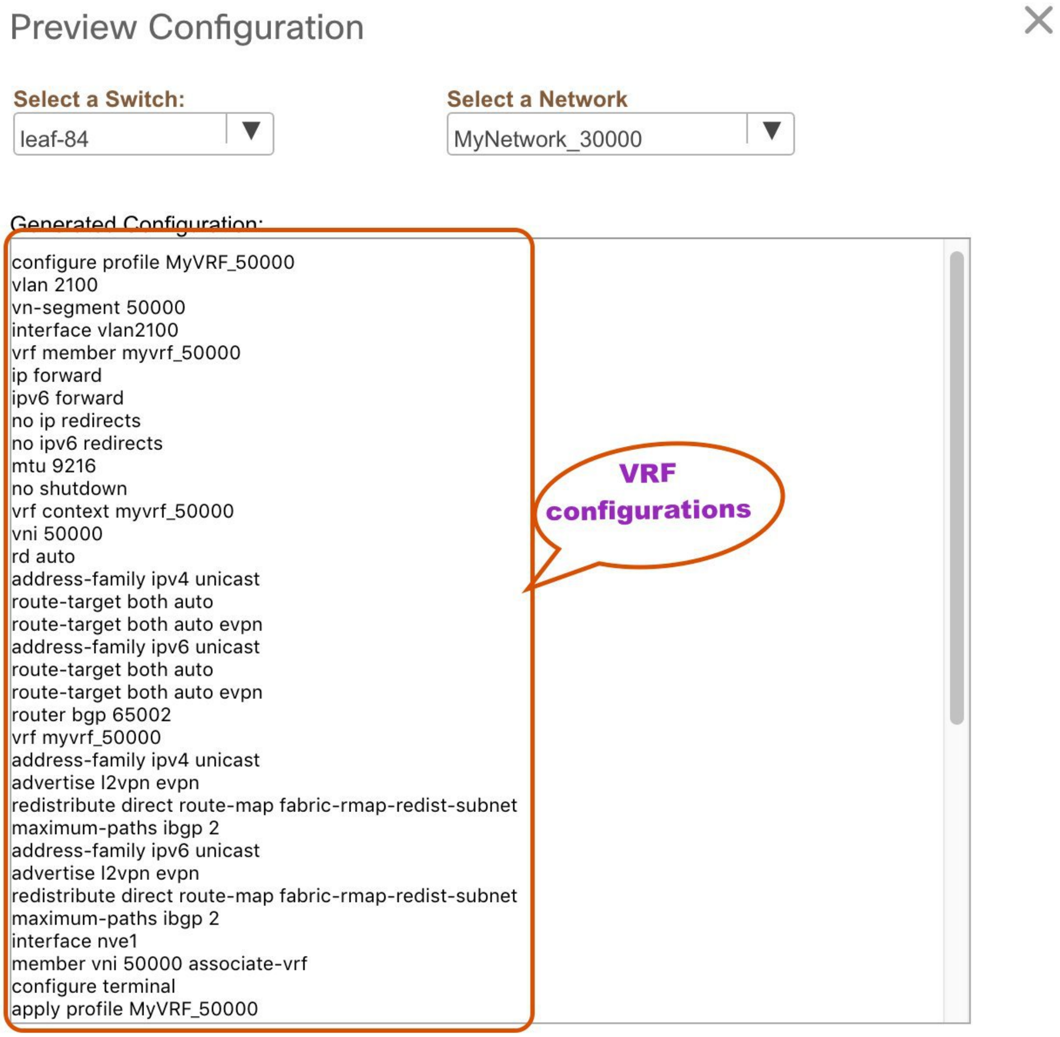

```
configure profile MyVRF_50000
vlan 2100
vn-segment 50000
interface vlan2100
vrf member myvrf_50000
ip forward
ipv6 forward
no ip redirects
no ipv6 redirects
mtu 9216
no shutdown
vrf context myvrf_50000
vni 50000
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
address-family ipv6 unicast
route-target both auto
route-target both auto evpn
router bgp 65002
vrf myvrf_50000
address-family ipv4 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redist-subnet
maximum-paths ibgp 2
address-family ipv6 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redist-subnet
maximum-paths ibgp 2
interface nve1
member vni 50000 associate-vrf
configure terminal
apply profile MyVRF_50000
```

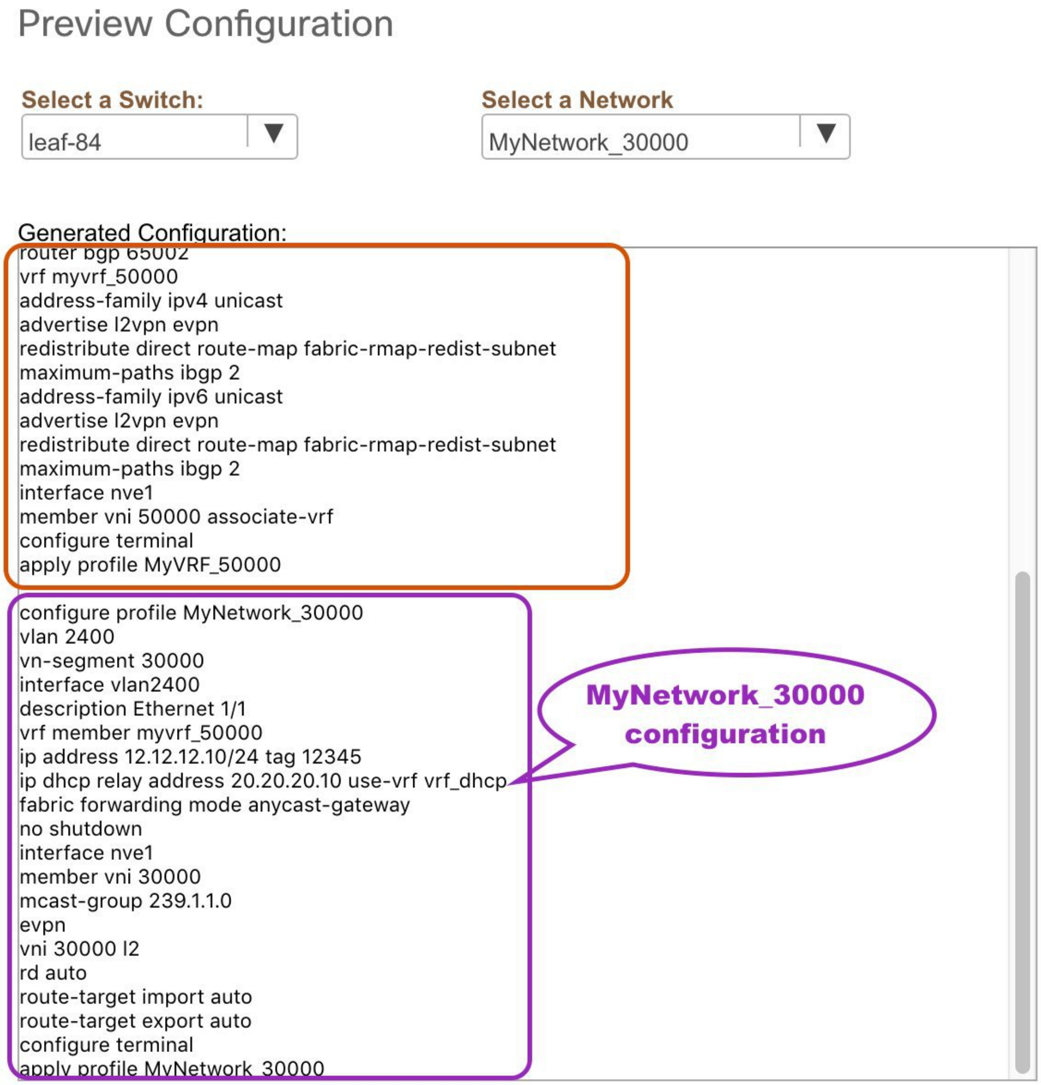

```
router bgp 65002
vrf myvrf_50000
address-family ipv4 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redist-subnet
maximum-paths ibgp 2
address-family ipv6 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redist-subnet
maximum-paths ibgp 2
interface nve1
member vni 50000 associate-vrf
configure terminal
apply profile MyVRF_50000
configure profile MyNetwork_30000
vlan 2400
vn-segment 30000
interface vlan2400
description Ethernet 1/1
vrf member myvrf_50000
ip address 12.12.12.10/24 tag 12345
ip dhcp relay address 20.20.20.10 use-vrf vrf_dhcp
fabric forwarding mode anycast-gateway
no shutdown
interface nve1
member vni 30000
mcast-group 239.1.1.0
evpn
vni 30000 l2
rd auto
route-target import auto
route-target export auto
configure terminal
apply profile MyNetwork_30000
```

On the preview screen, you can select from the **Select a switch** and **Select a network** drop-down boxes at the top of the screen to view other network configurations.

After checking the configurations, close the screen. The *Topology View* appears.

6. Click **Deploy** on the top right part of the screen. The color of the switch icons changes to yellow and a message appears at the bottom right part of the screen indicating that the deployment is in progress. After the networks' deployment is complete, the color of the switch icons changes to green, indicating successful deployment.

Note When you select multiple networks on the *Topology View* screen and proceed to the deployment screen, the switch color reflects the status of the first network in the selected list of networks. In this example, the switch color turns green when *MyNetwork_30000* is provisioned on the switch. Go to the Networks page to view the individual status for all networks.

You can also use the Detailed View option to deploy networks and VRFs. Click **Detailed View** at the top right part of the screen. The Detailed View screen comes up.

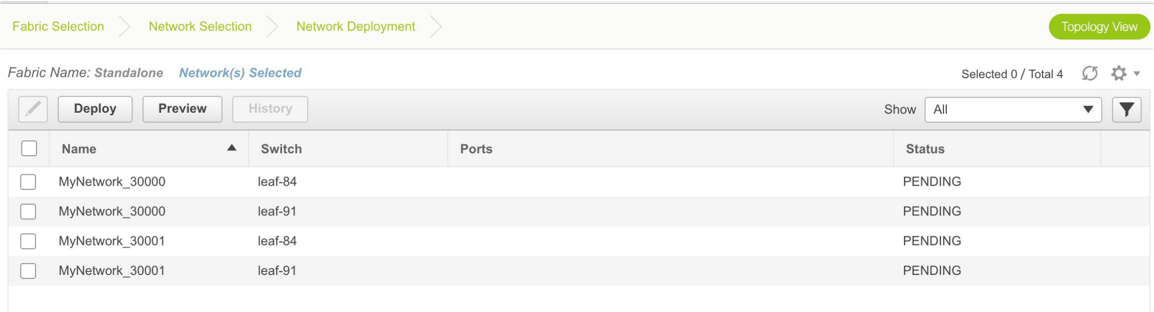

Name	Switch	Ports	Status
MyNetwork_30000	leaf-84		PENDING
MyNetwork_30000	leaf-91		PENDING
MyNetwork_30001	leaf-84		PENDING
MyNetwork_30001	leaf-91		PENDING

Similar to the *Topology View*, you can preview configurations and deploy networks/VRFs (using the **Preview** and **Deploy** buttons). The **Status** column indicates that the deployment is pending. Use the *Edit* option to edit the networks.

In addition, the **History** button allows you to view the previous configuration instances and status.

On the **Detailed View** page, the network profile configuration history is displayed. If you have associated specific trunk interfaces to that network, then the interface configuration is displayed as a separate configuration instance.

Note When you upgrade from an earlier release, such as DCNM 10.4(2) to the DCNM 11.0(1) release, overlay networks and VRFs deployment history information from the earlier DCNM release is not retained.

VRFs Deployment in the Standalone Fabric

1. From the Networks page, click **VRF View** at the top right part of the screen to deploy VRFs.

   (If you have freshly logged in to DCNM, do the following:

   Click **Control > Networks & VRFs**.

   Click **Continue** in the LAN Fabric Provisioning page.

   Choose *Standalone* from the drop-down list and click **Continue** to reach the Networks page.

   Click **VRF View** at the top right part of the Networks page**)**.

   The VRFs page comes up. The list of VRFs created for the *Standalone* fabric are displayed in this screen.

   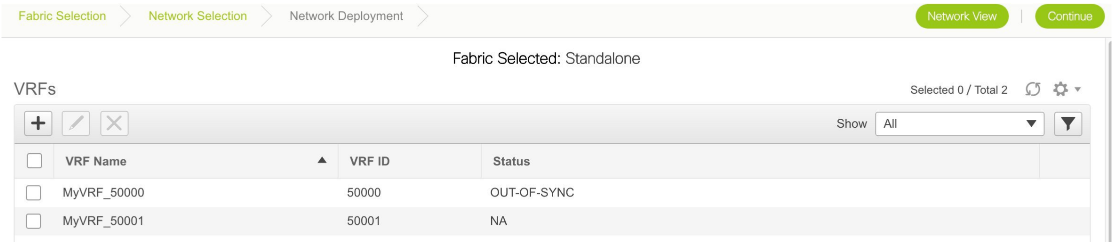

   Fabric Selected: Standalone

VRF Name	VRF ID	Status
MyVRF_50000	50000	OUT-OF-SYNC
MyVRF_50001	50001	NA

2. Select VRFs (by selecting corresponding check boxes) that you want to deploy and click **Continue** at the top right part of the screen.

The VRF Deployment page appears. On this page, you can see the topology of the *Standalone* fabric.

The following example shows you how to deploy the *MyVRF_50001* the VRF on the leaf switches leaf-84 and leaf-91. You can deploy VRFs simultaneously on multiple switches but of the same role (*Leaf*, *Border Gateway*, and so on).

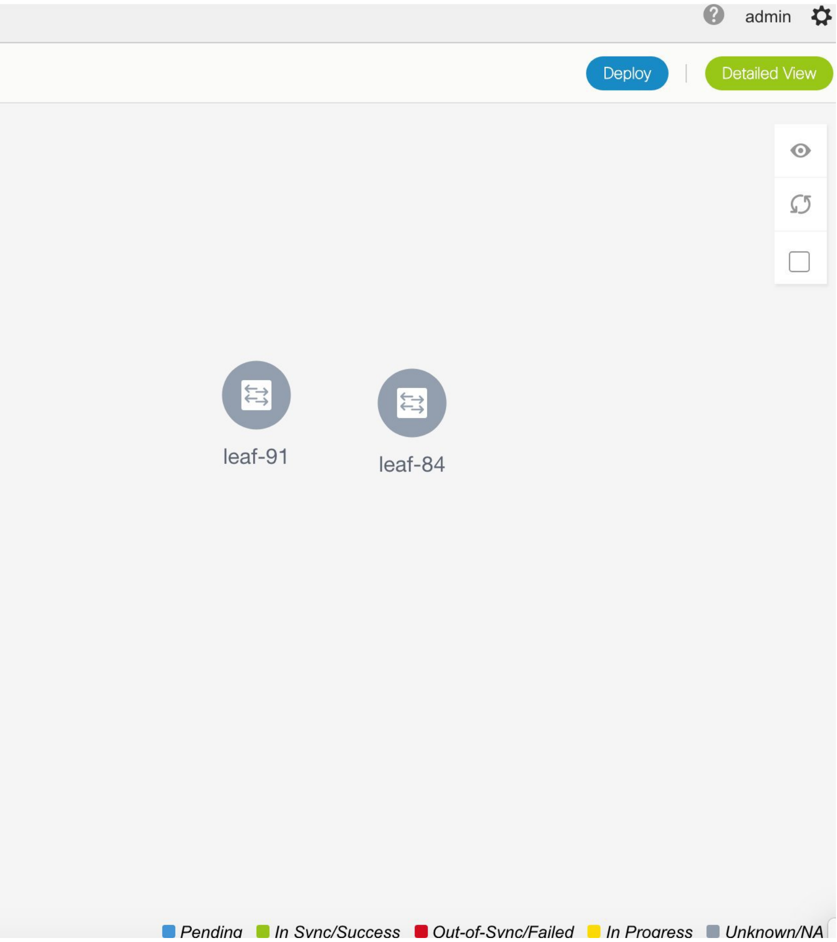

At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly (Blue for *Pending* state, yellow for *In Progress* state when the provisioning is in progress, red for failure state, green when successfully deployed, and so on).

The overlay networks (or VRFs) provisioning status is context-specific. It is a combination of VRFs that you chose for provisioning and the relevant switches in the topology. In this example, it means that the VRF *50001* is yet to be deployed on any switch in this fabric.

You can move the topology around the screen by clicking the left mouse button on the screen and moving it in the direction you desire. You can enlarge or shrink the switch icons proportionately by moving the cursor roller. You can also use corresponding alternatives on the touchpad.

3. Double-click a switch to deploy the VRF on it. For deployment of VRFs on multiple switches (like in this case, deploying VRF *50001* on leaf switches leaf-84 and leaf-91), do the following:
    1. Click the *Multi-Select* option from the panel at the top right part of the screen. This freezes the topology to a static state.

2. Drag the cursor across the switches.

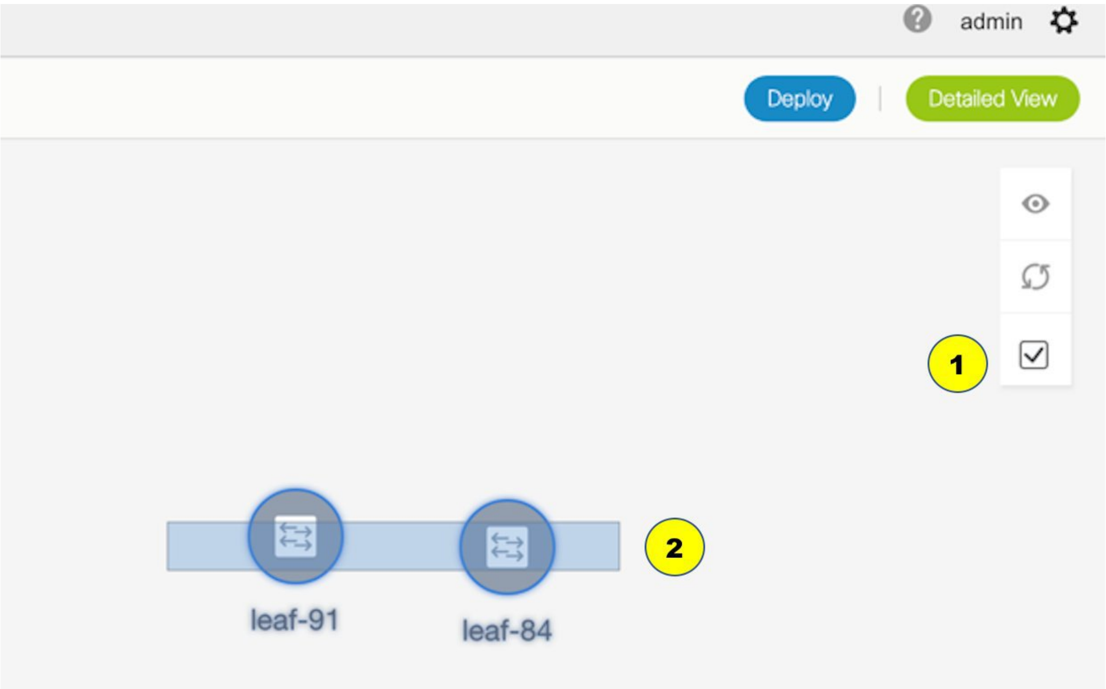

Immediately, the Switches Deploy screen (for VRFs) appears.

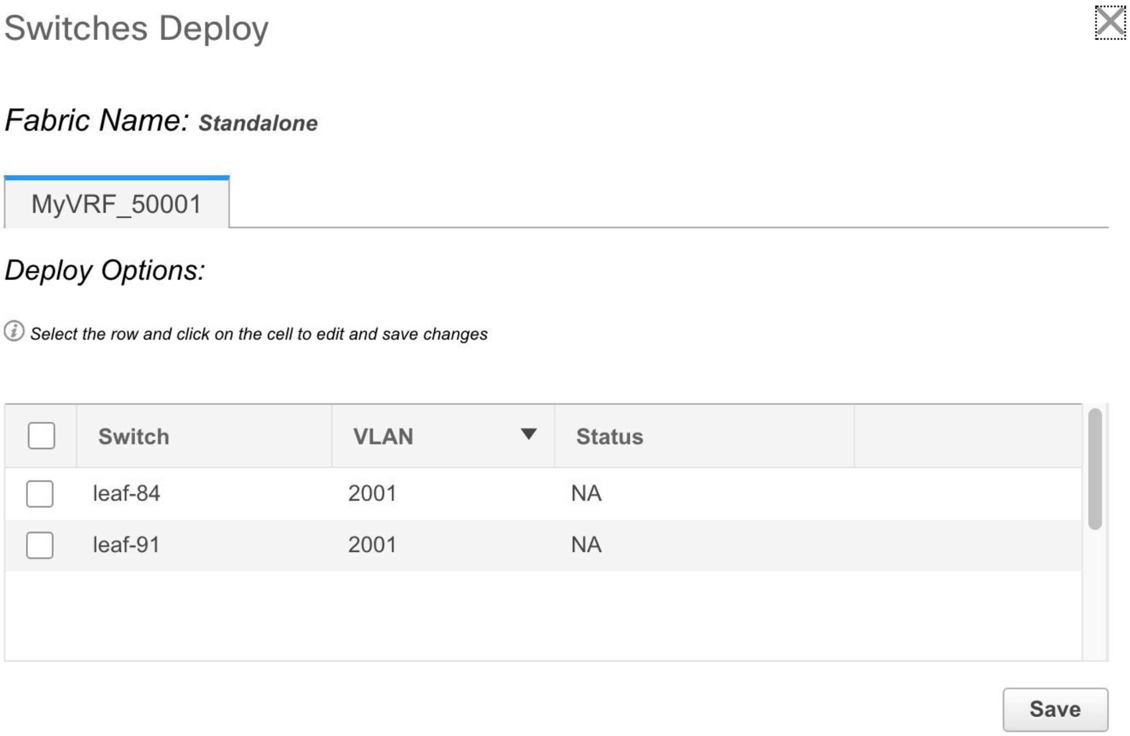

Switches Deploy

Fabric Name: **Standalone**

MyVRF_50001

Deploy Options:

Select the row and click on the cell to edit and save changes

□	Switch	VLAN	Status
□	leaf-84	2001	NA
□	leaf-91	2001	NA

Save

A tab represents each VRF (the first selected VRF is displayed by default) that is being deployed. In each VRF tab, the switches are displayed. Each row represents a switch.

Click the checkbox next to the **Switch** column to select the switches. Both the switch check boxes are selected automatically. VRF *50001* is ready to be provisioned on the switches leaf-84 and leaf-91.

Select the other VRF tab and make the same selections.

4. Click **Save** (at the bottom right part of your screen) to save VRF configurations.

   The topology screen comes up again. Click the *Refresh* button in the vertical panel at the top right part of the screen. The blue color on the switch icons indicates that the deployment is pending.

   Preview the configurations by clicking the *Preview* button (the eye icon above the *Multi-Select* option).

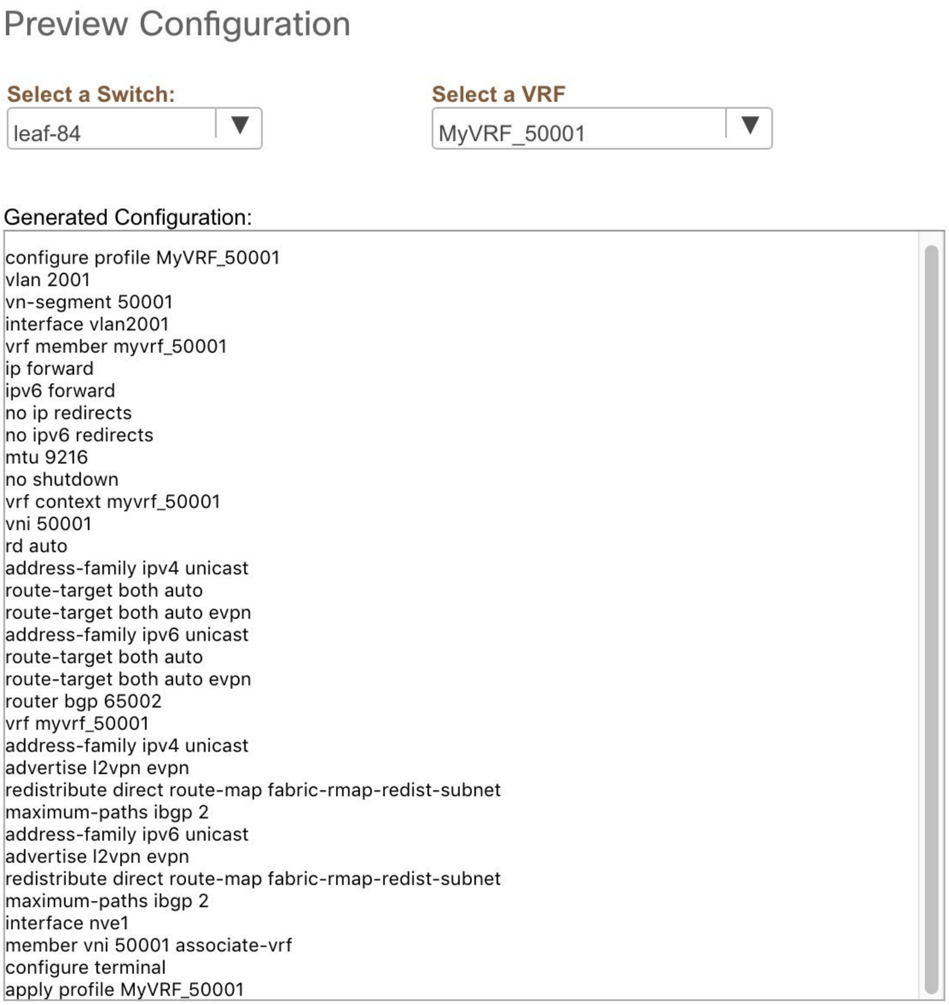

Preview Configuration

Select a Switch: leaf-84

Select a VRF: MyVRF_50001

Generated Configuration:

```
configure profile MyVRF_50001
vlan 2001
vn-segment 50001
interface vlan2001
vrf member myvrf_50001
ip forward
ipv6 forward
no ip redirects
no ipv6 redirects
mtu 9216
no shutdown
vrf context myvrf_50001
vni 50001
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
address-family ipv6 unicast
route-target both auto
route-target both auto evpn
router bgp 65002
vrf myvrf_50001
address-family ipv4 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redist-subnet
maximum-paths ibgp 2
address-family ipv6 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redist-subnet
maximum-paths ibgp 2
interface nve1
member vni 50001 associate-vrf
configure terminal
apply profile MyVRF_50001
```

   After checking the configurations, close the screen. The *Topology View* screen appears.

5. Click the **Deploy** button on the top right part of the screen. The color of the switch icons changes to yellow and a message appears at the bottom right part of the screen indicating that the deployment is in progress. After the VRF deployment is complete, the color of the switch icons changes to green, indicating successful deployment.

You can also use the **Detailed View** button to deploy networks and VRFs.

Click **Detailed View** at the top right part of the screen. The Detailed View screen comes up.

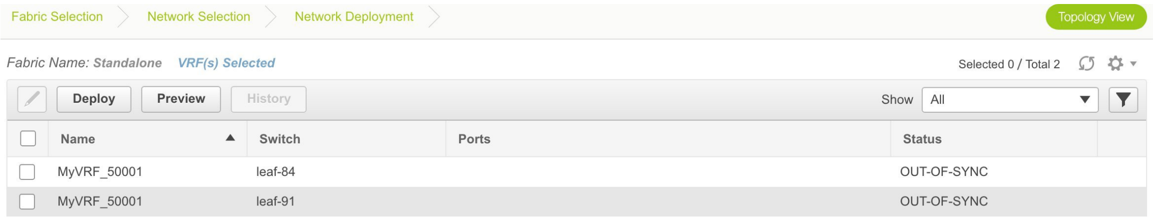

Name	Switch	Ports	Status
MyVRF_50001	leaf-84		OUT-OF-SYNC
MyVRF_50001	leaf-91		OUT-OF-SYNC

Similar to the *Topology View*, you can preview configurations and deploy networks/VRFs (from the **Preview** and **Deploy** buttons). The **Status** column indicates that the deployment is pending. Use the *Edit* option to edit the options.

In addition, the **History** button allows you to view the previous configuration instances and status.

Note When you upgrade from an earlier release, such as DCNM 10.4(2) to the DCNM 11.0(1) release, overlay networks and VRFs deployment history information from the earlier DCNM release is not retained.

Undeploying Networks

You can undeploy VRFs and networks from the *Topology View* page. The DCNM screen flow for undeployment is similar to the deployment process flow. Go to the Topology View page to undeploy networks:

1. Choose **Control > Networks and VRFs**.
2. In the **Select a Fabric** page, click **Continue** (at the top right part of the screen). The Networks page comes up.
3. Select the networks that you want to undeploy and click **Continue**. The Topology View page comes up.
4. On the **Topology View** page, select the **Multi-Select** button if you are undeploying the networks from multiple switches. The Switches Deploy screen comes up.

   (For a single switch, double-click the switch and the Switches Deploy screen comes up).
5. In the Switches Deploy screen, the **Status** column for the deployed networks is displayed as DEPLOYED. Unselect the check boxes next to the switches, as needed. Ensure that you repeat this on all tabs since each tab represents a network.
6. Click **Save** (at the bottom right part of the screen) to initiate the undeployment of the networks. The *Topology View* comes up again.

Note Alternatively, you can click the **Detailed View** button to undeploy networks.

7. Refresh the screen, preview configurations if needed and click **Deploy** to remove the network configurations on the switches. After the switch icons turn green, it indicates successful undeployment.
8. Go to the Networks page to verify if the networks have been undeployed.

Undeploying VRFs

You can undeploy VRFs and networks from the *Topology View* page. The DCNM screen flow for undeployment is similar to the deployment process flow.

1. Choose **Control > Networks and VRFs**.
2. In the Select a Fabric page, click **Continue** (at the top right part of the screen). The Networks page comes up.
3. Click the **VRF View** button (at the top right part of the screen) to go to the VRFs screen.
4. Select the VRFs that you want to undeploy and click **Continue**. The *Topology View* page comes up.
5. On the *Topology View* page, select the *Multi-Select* option if you are undeploying the VRFs from multiple switches. The Switches Deploy screen comes up.

   (For a single switch, double-click the switch and the Switches Deploy screen comes up).
6. In the Switches Deploy screen, the **Status** column for the deployed VRFs is displayed as DEPLOYED. Unselect the check boxes next to the switches, as needed. Ensure that you repeat this on all tabs since each tab represents a VRF.
7. Click **Save** (at the bottom right part of the screen) to initiate the undeployment of the VRFs. The *Topology View* comes up again.

Note Alternatively, you can click the **Detailed View** button to undeploy VRFs.

8. Refresh the screen, preview configurations if needed and click **Deploy** to remove the VRF configurations on the switches. After the switch icons turn green, it indicates successful undeployment.
9. Go to the VRFs page to verify if the networks have been undeployed.

Deleting Networks and VRFs in the MSD Fabric

If you want to delete networks and corresponding VRFs in the MSD fabric, follow this order:

1. Undeploy the networks, if not already done.
2. Delete the networks.
3. Undeploy the VRFs, if not already done.
4. Delete the VRFs.

Creating an External Fabric

You can create an external fabric in DCNM to depict a connection between the VXLAN and external fabrics in the DCNM GUI. After creating an external fabric, use the **Add switches** option to add switches to it. Some pointers:

- An external fabric is a monitor-only mode fabric.
- You can import, remove, and delete switches for an external fabric.

- For Inter-Fabric Connection (IFC) cases, you can choose Cisco 9000, 7000 and 5600 Series switches as destination switches in the external fabric.
- You can use non-existing switches as destination switches.
- The template that supports an external fabric is *External_Fabric.template*.
- On the Topology View screen, the VXLAN BGP EVPN and connected external fabrics can be viewed together.

Follow these steps to create an external fabric from Fabric Builder.

1. Click **Control > Fabric Builder**. The Fabric Builder page comes up.
2. Click the **Create Fabric** button. The Add Fabric screen comes up. The fields in this screen are:

   Fabric Name - Enter the name of the external fabric.

   Fabric Template - Choose *External_Fabric*.

   When you choose the fabric template, the fabric creation screen for creating an external fabric comes up.
3. Enter the BGP AS number and click **Save**.

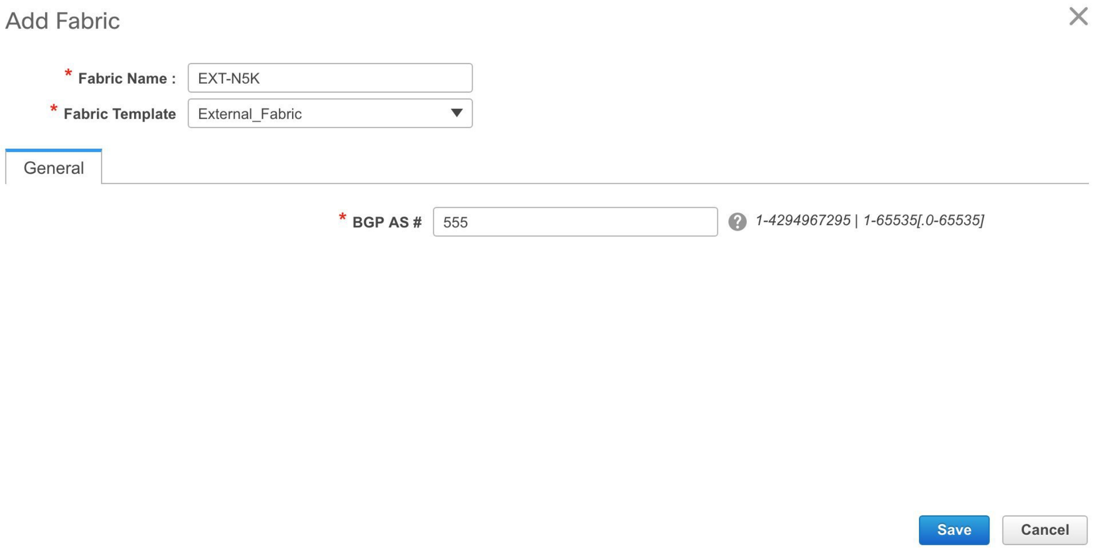

Add Fabric

* Fabric Name : EXT-N5K

* Fabric Template External_Fabric

General

* BGP AS # 555

When you create an Inter-Fabric Connection from a VXLAN fabric to this external fabric, the BGP AS number is referenced as the external or neighbor fabric AS Number.

After the external fabric is created, the external fabric topology page comes up.

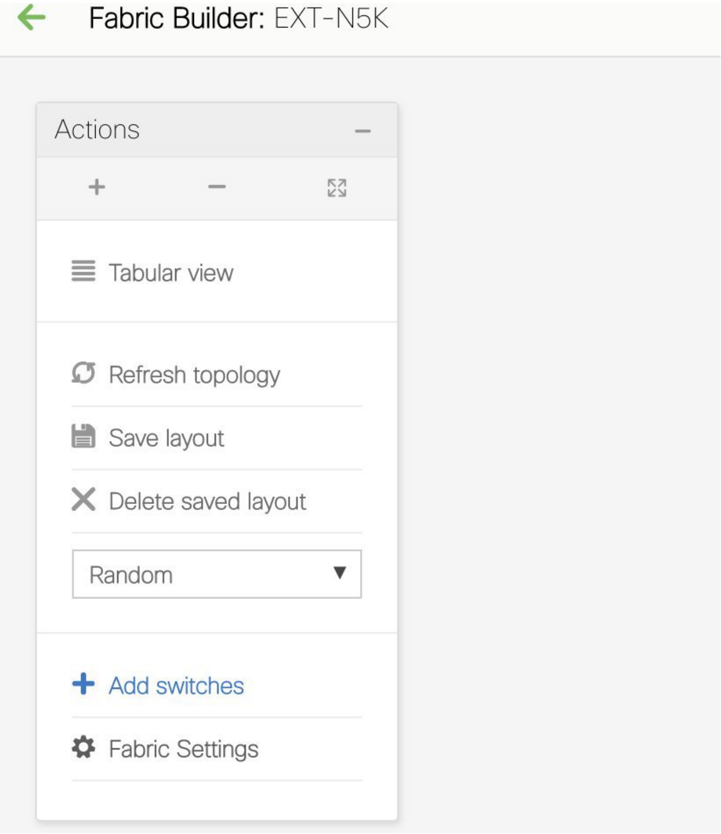

Note When you deploy networks or VRFs for the VXLAN fabric, the deployment page shows the VXLAN and external fabrics that are connected to each other.

A sample screenshot of the deployment page (*Topology View* screen) is shown. *Note that individual devices in the external fabric are not shown and only a cloud icon with the fabric name is displayed*.

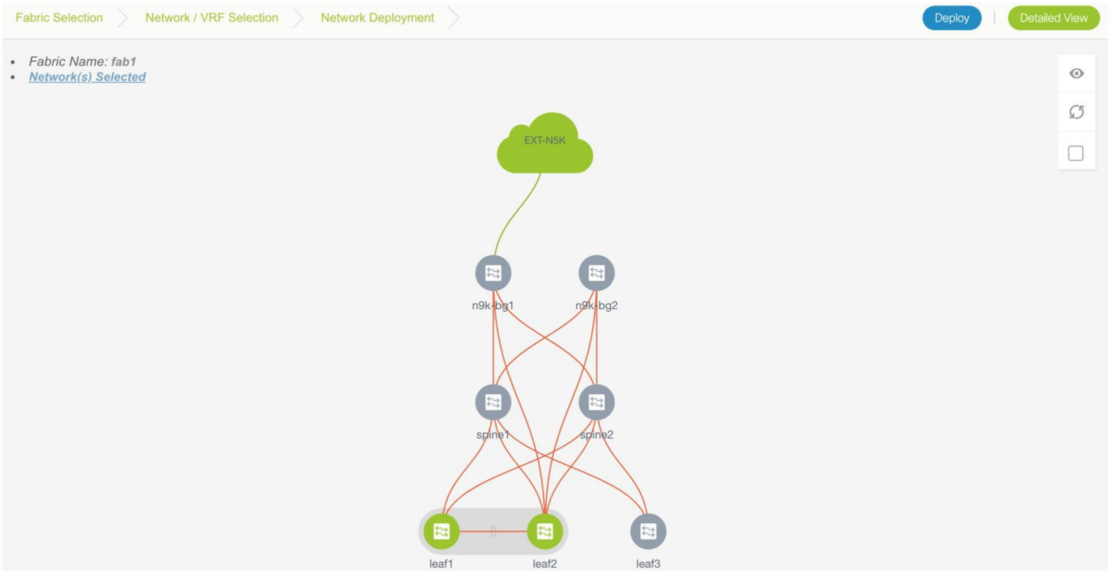

Adding Fabric Extensions

Before You Begin - In the fabric topology, the border switches should be set with an appropriate role (for example, Border Leaf or Border Gateway). The subsequent procedure describes how the inter-fabric connections between the border devices in the selected fabric and the external devices are defined.

1. Click **Control > Networks & VRFs** (under Fabrics submenu). The LAN Fabric Provisioning page comes up.

2. Click **Continue**. The Select a Fabric page is displayed. From the **Select a Fabric** drop down box, select the source fabric from which you want to connect to the other fabric.

3. Click **Fabric Extension Setup**.

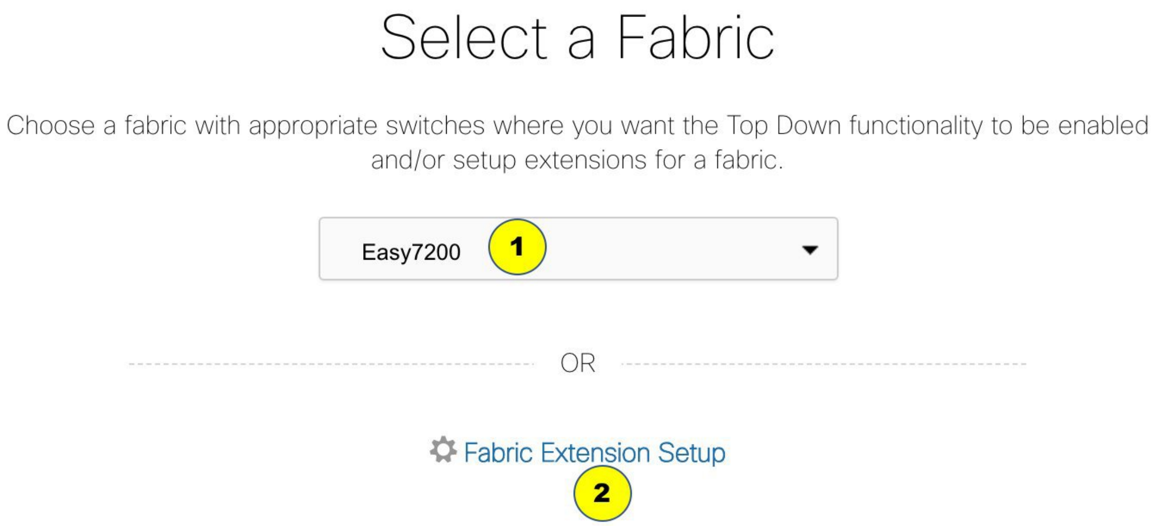
Select a Fabric

Choose a fabric with appropriate switches where you want the Top Down functionality to be enabled and/or setup extensions for a fabric.

Easy7200

OR

Fabric Extension Setup

The Fabric Extension screen comes up.

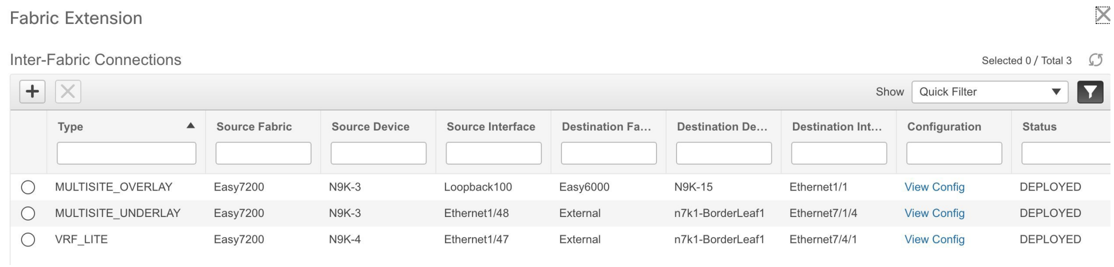

	Type	Source Fabric	Source Device	Source Interface	Destination Fa...	Destination De...	Destination Int...	Configuration	Status
○	MULTISITE_OVERLAY	Easy7200	N9K-3	Loopback100	Easy6000	N9K-15	Ethernet1/1	View Config	DEPLOYED
○	MULTISITE_UNDERLAY	Easy7200	N9K-3	Ethernet1/48	External	n7k1-BorderLeaf1	Ethernet7/1/4	View Config	DEPLOYED
○	VRF_LITE	Easy7200	N9K-4	Ethernet1/47	External	n7k1-BorderLeaf1	Ethernet7/4/1	View Config	DEPLOYED

The **Inter-Fabric Connections** section lists previously created external connections. Each line represents a physical or logical connection between a border node in the selected fabric and an external device in some other fabric. For each connection, the source fabric, source device, source interface, destination fabric, destination device, and destination interface are listed along with the type of external connectivity. This section is empty the first time you add an external connection. Two primary types of external connectivity are supported, *VRF Lite* and *EVPN Multi-Site*.

VRF Lite (VRF_LITE) - For each VRF, an external BGP (eBGP) peering session needs to be set up between the border node and the external device. As part of the connection setup, the eBGP peering session is established from the border node in the default VRF along with additional global configuration of route-maps for IPv4/IPv6 cases.

EVPN Multi-Site - This requires setting up the Border Gateway base configuration for enabling the Multi-Site feature and the underlay peering to the external devices (MULTISITE_UNDERLAY). This is followed by establishing overlay peering from the border gateway to appropriate external devices, either Border Gateways in other fabrics or Route Servers (MULTISITE_OVERLAY). Both the underlay and overlay peering are established over eBGP. Recall that Border Gateways are special devices that allow clear control and data plane segregation from one site to another while allowing for policy enforcement points for any inter-fabric traffic. They allow the same data plane (VXLAN) and control plane (BGP EVPN) to be employed both for inter-fabric and intra-fabric traffic.

Note If you extend the fabric through EVPN Multi-Site, you should first create an underlay extension (select MULTISITE_UNDERLAY in the **Extension Type** field) on the border gateway and then create overlay extensions (select MULTISITE_OVERLAY in the **Extension Type** field).

4. Click on the **Add** icon to add a new external connection. The Add Inter-Fabric Connections screen appears.

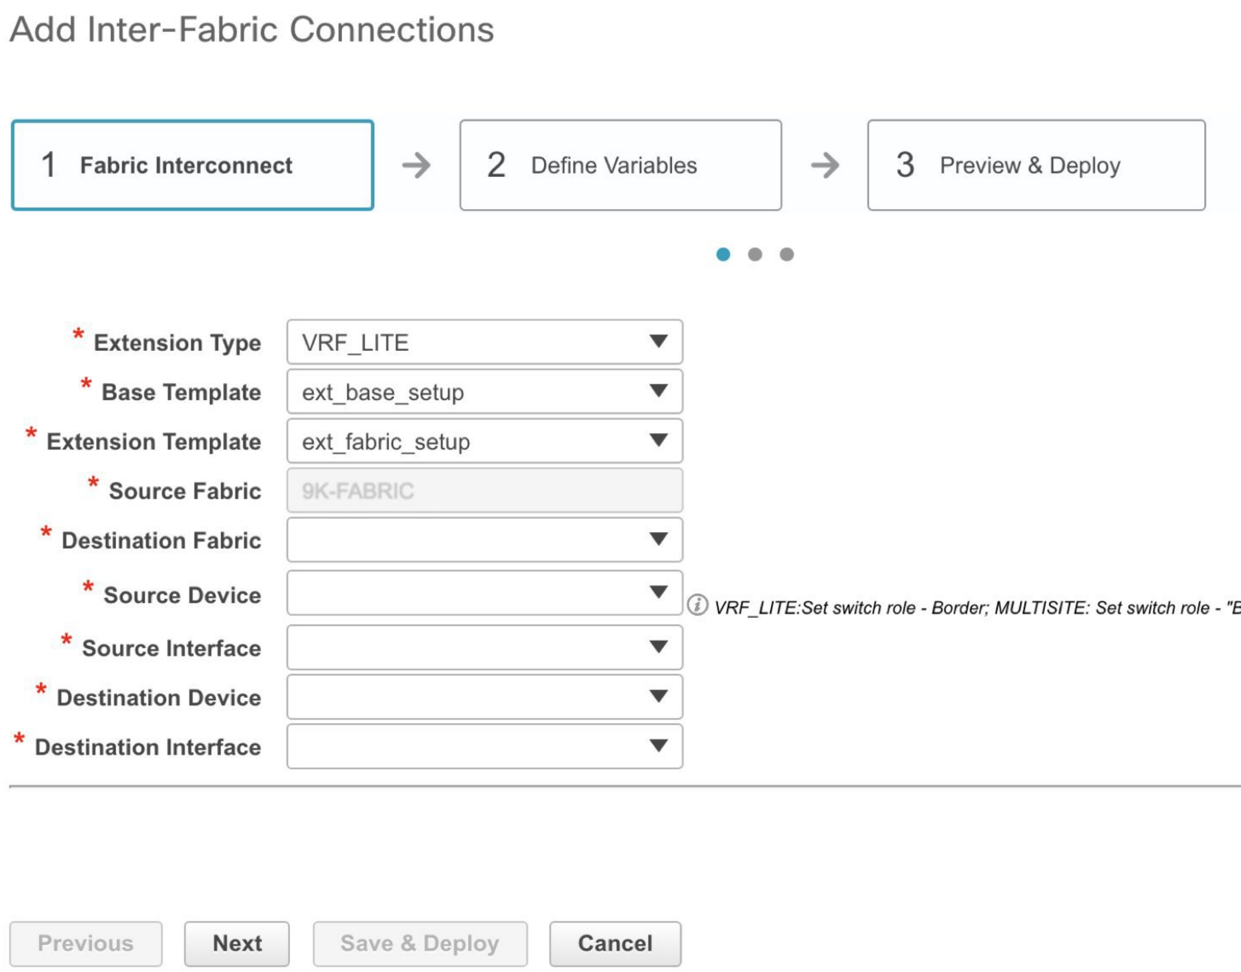

* Extension Type	VRF_LITE
* Base Template	ext_base_setup
* Extension Template	ext_fabric_setup
* Source Fabric	9K-FABRIC
* Destination Fabric	
* Source Device	
* Source Interface	
* Destination Device	
* Destination Interface	

VRF_LITE:Set switch role - Border; MULTISITE: Set switch role - "B

Fill up the fields on this page. The **Source Fabric** field is pre-populated in the Fabric Interconnect section. By default, the **Extension Type** is set to VRF_LITE. The **Base template** references the template that contains a one-time configuration pushed to border devices. The **Extension Template** references the setup template that contains the configuration that is generated and pushed to the border device to set up the corresponding inter-fabric connection. These templates are auto-populated with corresponding pre-packaged default templates based on user selections. The destination fabric that contains the external device peer must be selected. Note that based on the selection of the source device and source interface, the destination information is autopopulated based on CDP information if available. There is extra validation performed to ensure that the destination external device is indeed part of the destination fabric.

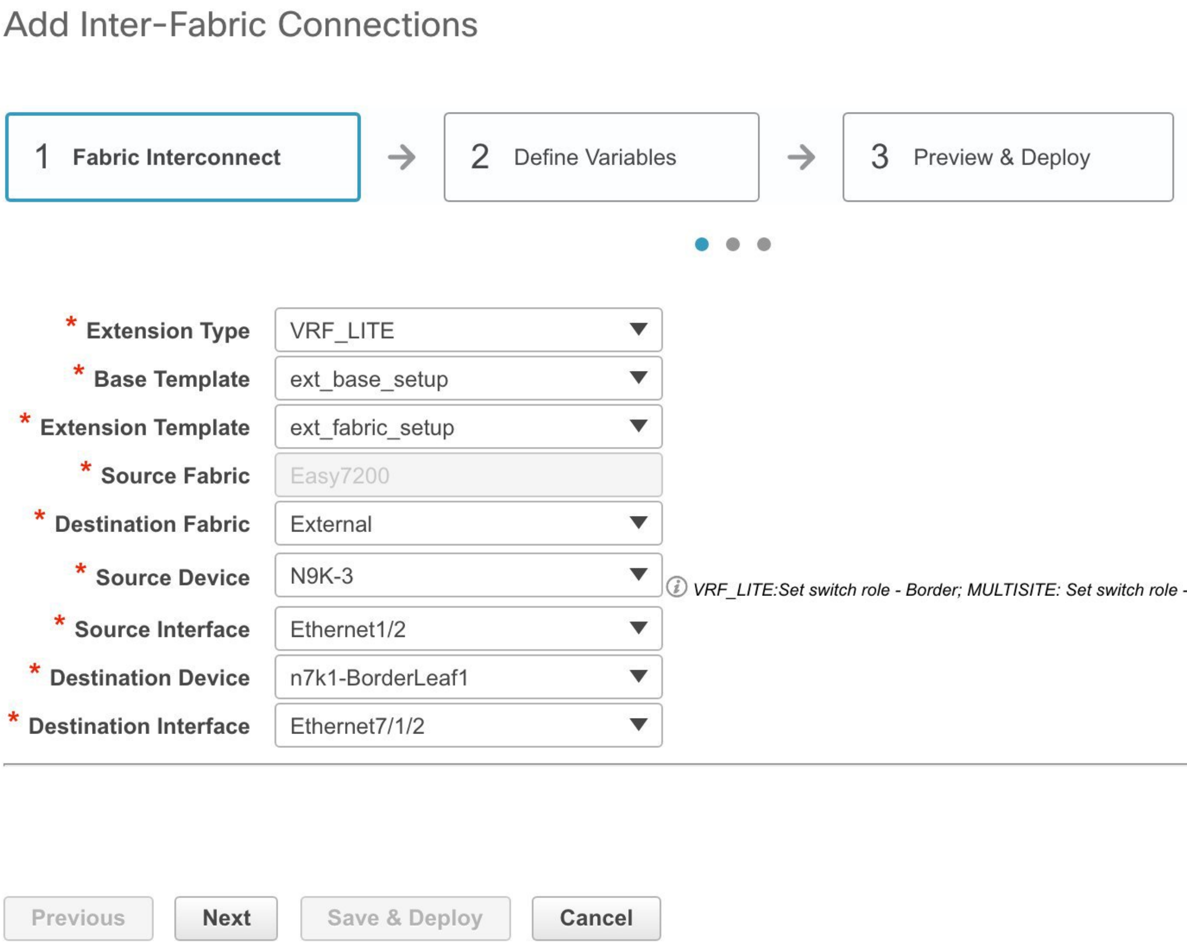

Add Inter-Fabric Connections

* Extension Type	VRF_LITE
* Base Template	ext_base_setup
* Extension Template	ext_fabric_setup
* Source Fabric	Easy7200
* Destination Fabric	External
* Source Device	N9K-3
* Source Interface	Ethernet1/2
* Destination Device	n7k1-BorderLeaf1
* Destination Interface	Ethernet7/1/2

VRF_LITE:Set switch role - Border; MULTISITE: Set switch role -

5. Click **Next** to go to the **Define Variables** section.

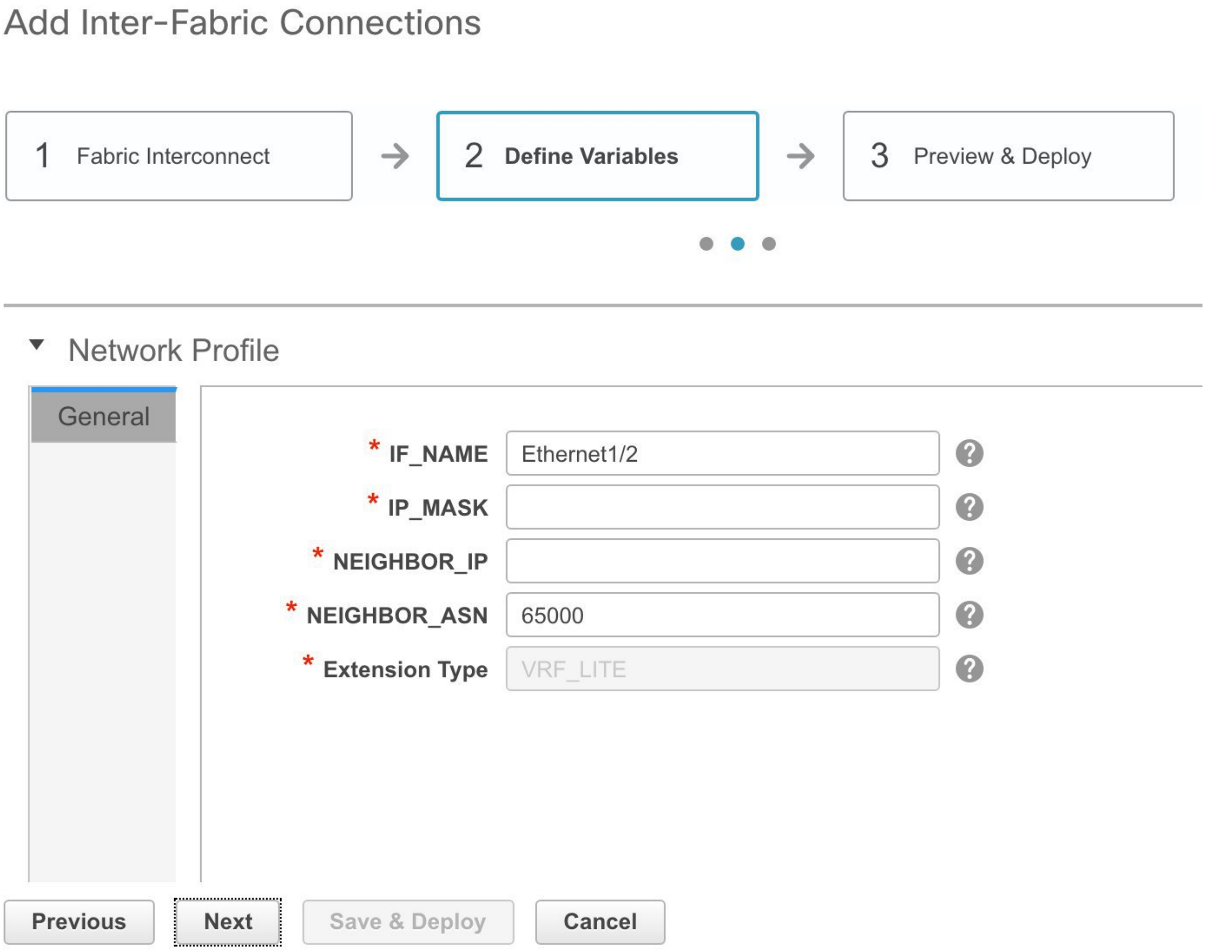

Add Inter-Fabric Connections

▼ Network Profile

* IF_NAME	Ethernet1/2
* IP_MASK	
* NEIGHBOR_IP	
* NEIGHBOR_ASN	65000
* Extension Type	VRF_LITE

Here, the source interface name, destination fabric ASN, and the extension type are autopopulated. The template variables are parsed from the templates that are selected in the previous step and displayed for user input. All mandatory parameters must be entered.

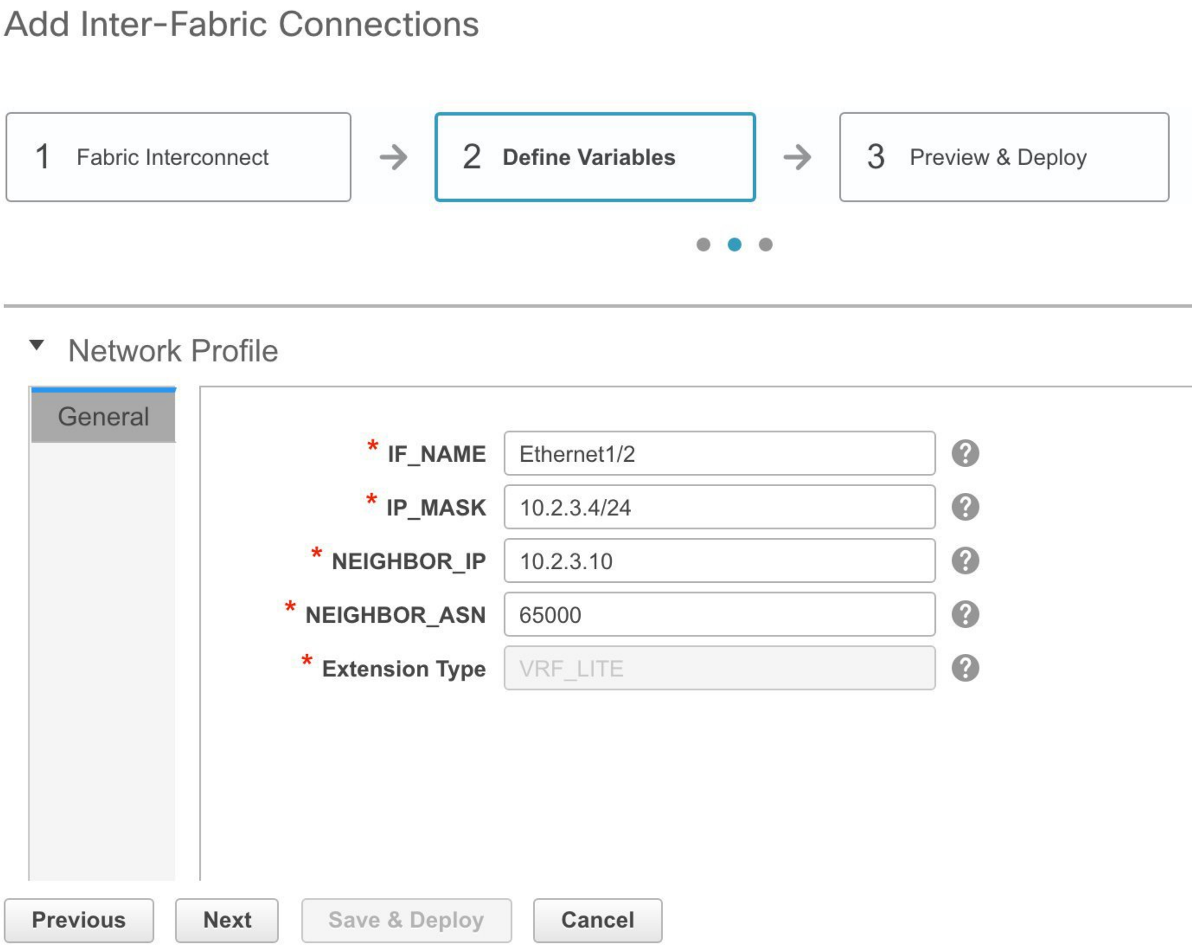

Add Inter-Fabric Connections

Network Profile

* IF_NAME	Ethernet1/2
* IP_MASK	10.2.3.4/24
* NEIGHBOR_IP	10.2.3.10
* NEIGHBOR_ASN	65000
* Extension Type	VRF_LITE

6. Click **Next** to go to the **Preview and Deploy** section.

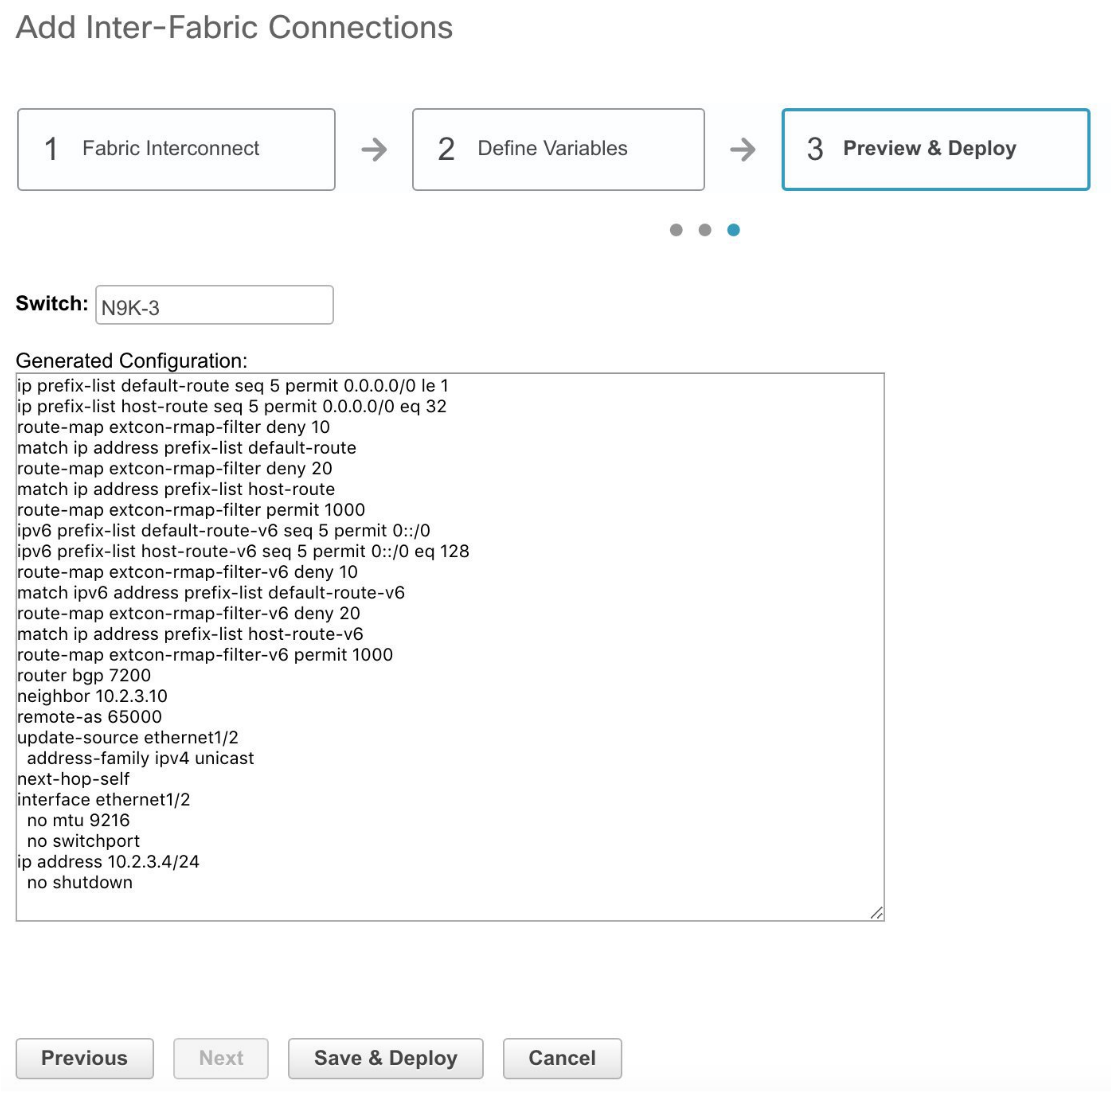

Switch: N9K-3

Generated Configuration:

```
ip prefix-list default-route seq 5 permit 0.0.0.0/0 le 1
ip prefix-list host-route seq 5 permit 0.0.0.0/0 eq 32
route-map extcon-rmap-filter deny 10
match ip address prefix-list default-route
route-map extcon-rmap-filter deny 20
match ip address prefix-list host-route
route-map extcon-rmap-filter permit 1000
ipv6 prefix-list default-route-v6 seq 5 permit 0::/0
ipv6 prefix-list host-route-v6 seq 5 permit 0::/0 eq 128
route-map extcon-rmap-filter-v6 deny 10
match ipv6 address prefix-list default-route-v6
route-map extcon-rmap-filter-v6 deny 20
match ip address prefix-list host-route-v6
route-map extcon-rmap-filter-v6 permit 1000
router bgp 7200
neighbor 10.2.3.10
remote-as 65000
update-source ethernet1/2
  address-family ipv4 unicast
next-hop-self
interface ethernet1/2
  no mtu 9216
  no switchport
ip address 10.2.3.4/24
  no shutdown
```

Here, you can preview the configuration that is deployed to the selected border device. Note that no configuration is pushed to the external device itself.

7. Click **Save and Deploy** to complete the task.

   This results in the configuration getting pushed to the appropriate border node. The external connection appears in the Fabric Extension screen.

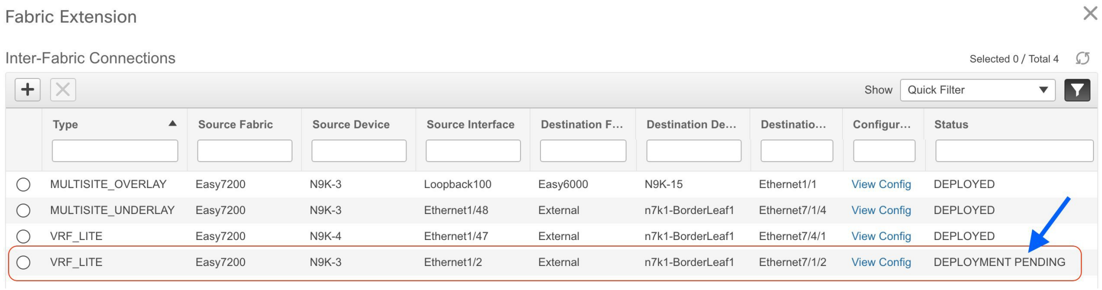

	Type	Source Fabric	Source Device	Source Interface	Destination F...	Destination De...	Destinatio...	Configur...	Status
○	MULTISITE_OVERLAY	Easy7200	N9K-3	Loopback100	Easy6000	N9K-15	Ethernet1/1	View Config	DEPLOYED
○	MULTISITE_UNDERLAY	Easy7200	N9K-3	Ethernet1/48	External	n7k1-BorderLeaf1	Ethernet7/1/4	View Config	DEPLOYED
○	VRF_LITE	Easy7200	N9K-4	Ethernet1/47	External	n7k1-BorderLeaf1	Ethernet7/4/1	View Config	DEPLOYED
○	VRF_LITE	Easy7200	N9K-3	Ethernet1/2	External	n7k1-BorderLeaf1	Ethernet7/1/2	View Config	DEPLOYMENT PENDING

You can check the status of the deployment (*Pending, Deployed, Failed* so on) in the **Status** column. In case of FAILED or UNDEPLOYMENT FAILED status, use the hyperlink in the **Status** column to check the error messages for failure.

In this case, the status will change to DEPLOYED after the screen refresh. The sample topology displays the external connection, including the border device being connected to the external fabric.

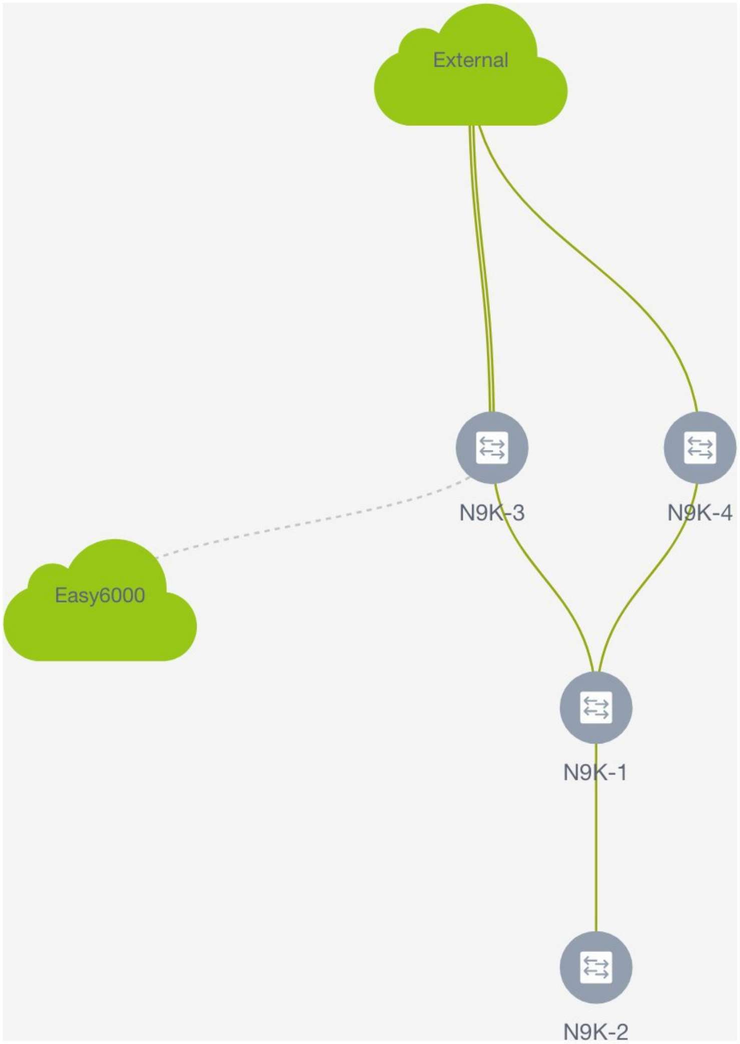

For additional inter-fabric connections, a similar set of steps is repeated. Note however, the base configuration to the border node is only pushed once, when the first inter-fabric connection is deployed for a given type. The connections can either be added or deleted, they cannot be updated or edited. On successful deployment of the inter-fabric connections, in the LAN Fabric provisioning topology view, each inter-fabric connection is displayed as an edge (solid for physical or dotted for logical) between the appropriate border node and the external fabric. *Note that individual devices in the external fabric are not shown and only a cloud icon with the fabric name is displayed*.

Note You can delete an IFC connection only if it is not attached to any network or VRF.

Post DCNM 10.4(2) to DCNM 11.0(1) Upgrade Procedure for VXLAN BGP EVPN Fabrics

This topic provides details on the procedure to gracefully on board a DCNM 10.4(2) managed VXLAN BGP EVPN fabric comprising Cisco Nexus 9000 switches, post upgrade to DCNM 11.0(1). The assumption is that the fabric was deployed with DCNM 10.4(2), including the underlay (via the DCNM published POAP templates) and the overlays including configuration on the border devices (optional). The DCNM provided POAP templates and the overlay profile templates themselves may have been customized for the desired deployment.

Before you begin - It is assumed that you have installed the Cisco DCNM 11.0(1) software. If not, follow the Upgrade process to upgrade from DCNM 10.4(2) to DCNM 11.0(1). After installation, follow the guidelines and start migrating devices to DCNM 11.0(1).

Note The term *upgrade* in this section refers to the actions of migrating the switches to the DCNM 11.0(1) release in the DCNM GUI and deployment of new configuration policies on the switches.

Guidelines and Limitations

- The assumption is that the fabric was operational and functional when it is being managed with DCNM 10.4(2). In other words, the underlay and overlays have been deployed to the switches in a consistent manner and the BGP sessions, VNIs, and so on, that are configured are part of a functional fabric.
- The switch roles (*leaf*, *border*, and so on) are retained from what they were set in DCNM 10.4.2 (prior DCNM). The assumption is that the roles were correctly set and hence the roles must not be changed during the migration process.
- As part of the migration process, DCNM reads the running configuration from every switch within the migrating fabric, and specifically for the VXLAN BGP EVPN underlay configuration, it does a match to reverse population of that state into the DCNM against the packaged best-practice policy templates. In other words, it *infers* the underlay intended state from the existing running configuration on the switches. The state of the overlay configuration from DCNM 10.4(2) is retained during the upgrade to DCNM 11.0(1).
- Configurations that are not supported in the upgrade or migration process are:
    - Manual VLAN and SVI (barring vPC peer link VLAN) configurations (that are not overlay related) - These are configurations that were not enabled as part of the DCNM 10.4(2) top down tenant configurations.
    - Loopback interface configurations other than loopback0, loopback1, and loopback254 interfaces. The assumption is that loopback0 is employed for BGP/IGP peering, loopback1 is employed for VTEPs, and loopback 254 is employed for the RP configuration on the spines (if applicable).
    - Subinterfaces (not provisioned via VRF-Lite extensions on the *Borders* via DCNM).

  After the upgrade is complete, you can add these configurations to the appropriate switches as needed using the **switch_freeform_config** policy (Refer *Freeform Configurations on Fabric Switches* for details). This ensures that the configuration is captured in DCNM as part of the intended configuration, hence, the configuration compliance module ensures that the intent is synchronized against the current running configuration with appropriate OUT-OF-SYNC/IN-SYNC status notification.

- vPC switches – Ensure that the following configurations are present on vPC switches as is expected for a typical functioning vPC pair in a VXLAN BGP EVPN fabric.
    - Secondary IP address on loopback1 (the loopback that is mapped to the NVE or VTEP interface).
    - vPC peer link port channel and member interfaces.
    - vPC peer link backup SVI and VLAN.

      If the switch is not a Cisco Nexus 9000 series switch with Cloud-scale ASICs, the peer link VLAN also needs to be specified in the **system nve infra-vlans** command.

  If the above configurations are missing, the upgrade will fail and the system will display an error message. To resolve the issue, you should enable correct vPC configurations and use the **Save and Deploy** option (explained during the upgrade process) to proceed with the upgrade.

- You can add more switch instances to the fabric after the upgrade process in the DCNM GUI. Refer the *Add Switch Instances in the Fabric* section for additional details.
- Policies created for the fabric underlay (for example, for fabric interfaces and routing) are created with the source set as *UNDERLAY*. These policies cannot be modified.

Upgrade Procedure in the DCNM GUI

1. Open a web browser and log on to the DCNM 11.0(1) Web UI https://<*DCNM-IP*> with the appropriate credentials.
2. Choose **Control > Fabric Builder**. The fabrics that were managed by DCNM 10.4(2) will be displayed in blue color. The blue color indicates that the fabric has been recognized as something that has been successfully imported from DCNM 10.4(2), but this fabric needs to be associated with an appropriate fabric template. In this screenshot, a single fabric is displayed.

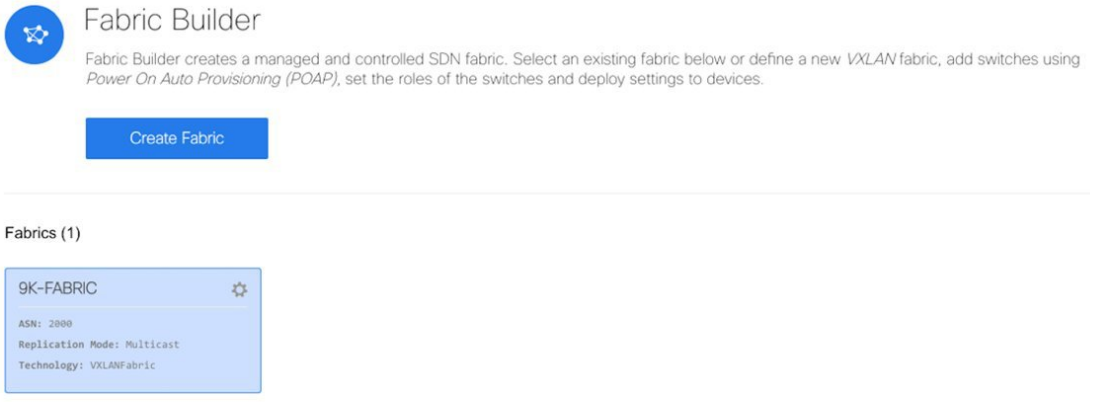

Fabric Builder

Fabric Builder creates a managed and controlled SDN fabric. Select an existing fabric below or define a new *VXLAN* fabric, add switches using *Power On Auto Provisioning (POAP)*, set the roles of the switches and deploy settings to devices.

Fabrics (1)

9K-FABRIC

ASN: 2000
Replication Mode: Multicast
Technology: VXLANFabric

3. Click the *wheel* icon of the fabric to associate it with an appropriate fabric template. The **Edit Fabric** screen comes up.
4. From the **Fabric Template** drop-down box, select **Easy_Fabric**.

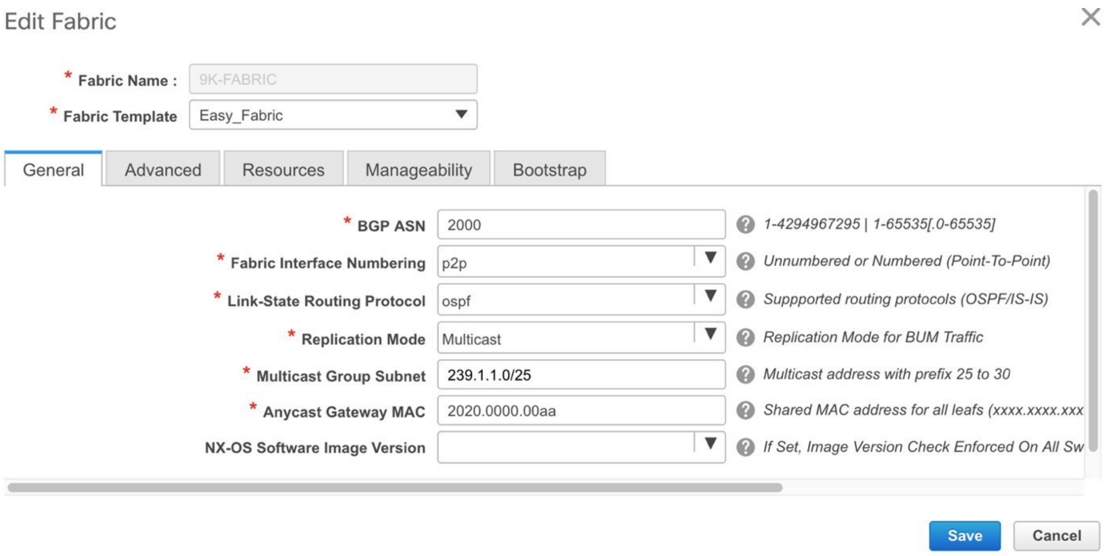

Edit Fabric

* Fabric Name : 9K-FABRIC

* Fabric Template Easy_Fabric

* BGP ASN 2000 — 1-4294967295 | 1-65535[.0-65535]

* Fabric Interface Numbering p2p — Unnumbered or Numbered (Point-To-Point)

* Link-State Routing Protocol ospf — Supported routing protocols (OSPF/IS-IS)

* Replication Mode Multicast — Replication Mode for BUM Traffic

* Multicast Group Subnet 239.1.1.0/25 — Multicast address with prefix 25 to 30

* Anycast Gateway MAC 2020.0000.00aa — Shared MAC address for all leafs (xxxx.xxxx.xxx

NX-OS Software Image Version — If Set, Image Version Check Enforced On All Sw

5. Update fabric parameters in accordance with the currently selected fabric. Recall that this is a functional fabric. The current support is present only for fabrics setup with underlay using IGP as IS-IS or OSPF. The BUM handling mechanism may be multicast or ingress-replication. The values entered should match the DCNM 10.4(2) fabric’s parameters.

   Specifically, ensure that the following values are the same as the switch configurations:

   - BGP AS Number.
   - Fabric underlay routing protocol (IS-IS or OSPF).
   - Replication mode (Multicast or Ingress Replication).
   - Fabric interface numbering (p2p or IP unnumbered).
   - vPC peer link VLAN, if vPC is present.
   - vPC delay restore time and other related parameters in the **Advanced** tab

   Manageability tab – To retain existing DNS, NTP and AAA configurations, clear the corresponding fields in this tab. Policies will be created using the source "". If you update any of the settings here, the settings will override corresponding switch configurations.

   You can also update the DNS, NTP and AAA parameters after the migration.

6. Click **Save** to save the updated settings.

   The topology screen comes up. This screen displays the existing devices and their connections. Since the devices are yet to be migrated to DCNM 11.0(1), the **Migration-mode** icon will be displayed on each switch. Validate that the roles have been appropriately retained from the DCNM 10.4(2) upgrade.

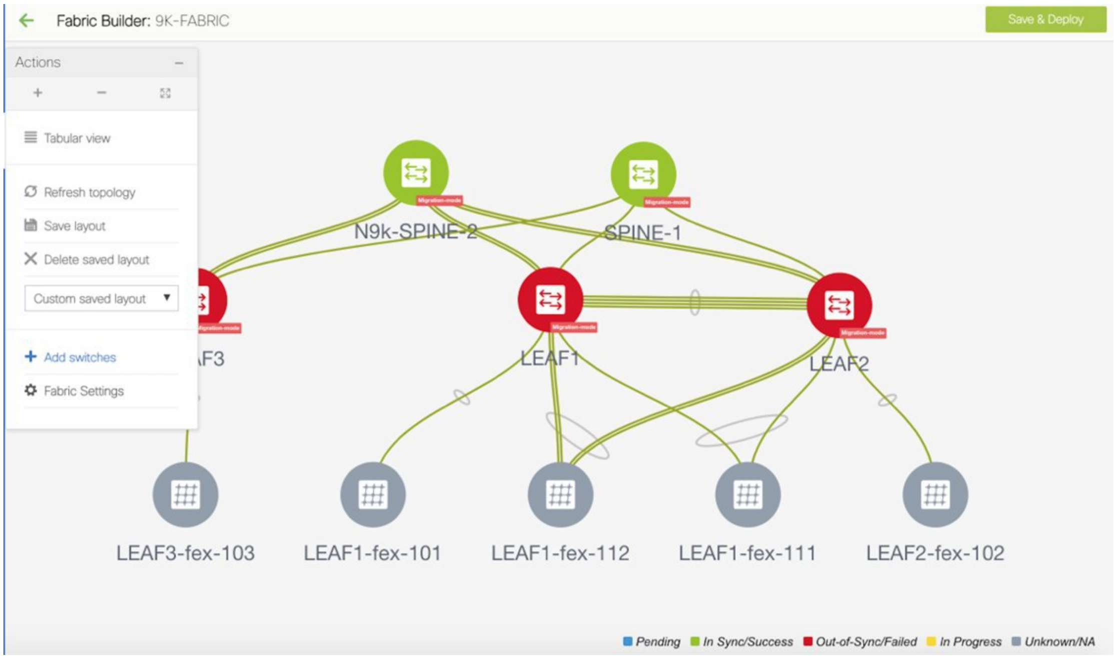

7. Click **Save & Deploy** at the top right part of the screen to start the migration process.

   Policy creation is initiated based on existing device configuration and how the devices are connected with each other. At this point, the policy creation in terms of the underlay intent is inferred from the running configuration of every device. In case there is a mismatch found between the switch configuration and the inputs provided in the Fabric Settings, an appropriate error will be reported. You must make appropriate changes to address the reported error before proceeding to execute “Save & Deploy” again. Addressing the error may involve making changes to the switch configuration on which the error was reported or making edits to the Fabric Settings or potentially customize policies to match the running configuration. You can see a message at the center of the screen indicating that the intended configuration for every switch in the fabric is being generated in the DCNM.

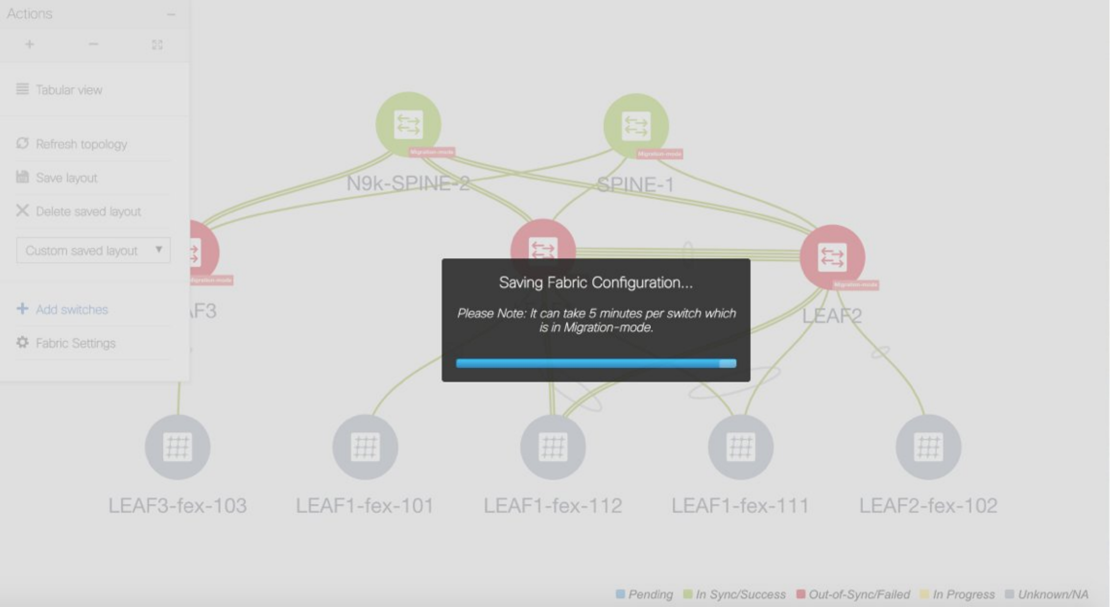
Saving Fabric Configuration...

Please Note: It can take 5 minutes per switch which is in Migration-mode.

Note that this process may take a while depending on the number of switches that are part of the fabric and the size of the running configuration, which is a function of the number of networks and VRFs deployed on the switches. Once this process has been successfully completed, next, the **Config Deployment** screen comes up as shown below.

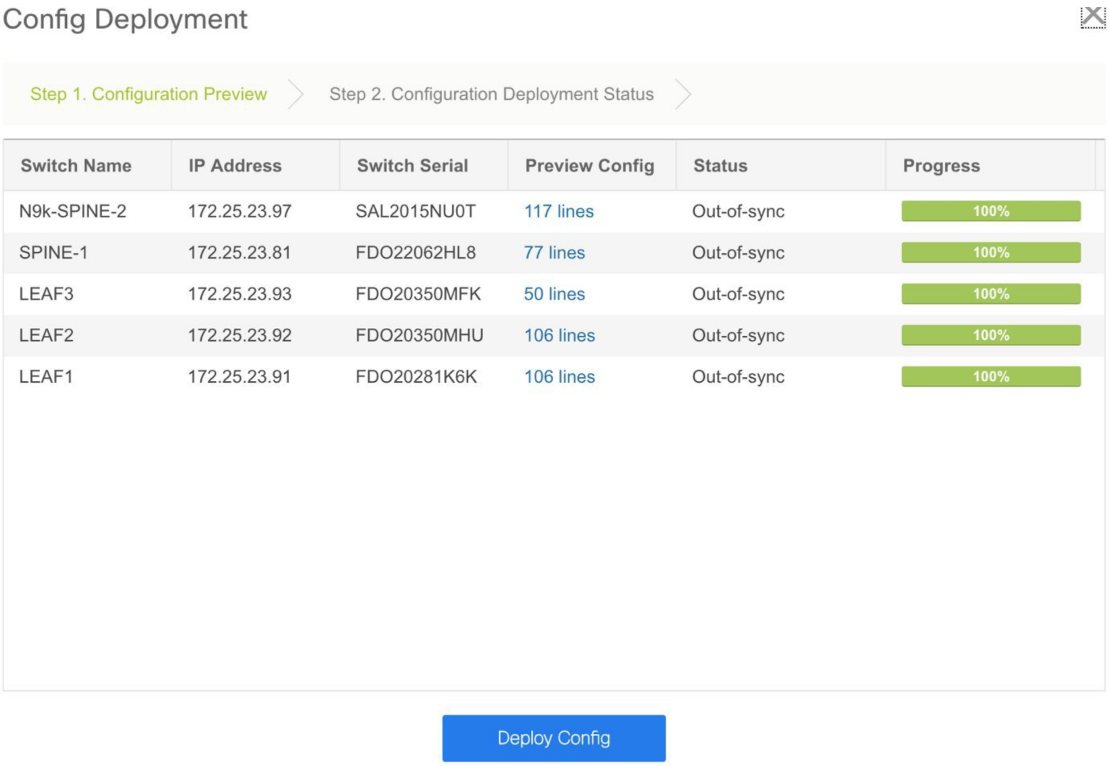

Config Deployment
Step 1. Configuration Preview
Step 2. Configuration Deployment Status

Switch Name	IP Address	Switch Serial	Preview Config	Status	Progress
N9k-SPINE-2	172.25.23.97	SAL2015NU0T	117 lines	Out-of-sync	100%
SPINE-1	172.25.23.81	FDO22062HL8	77 lines	Out-of-sync	100%
LEAF3	172.25.23.93	FDO20350MFK	50 lines	Out-of-sync	100%
LEAF2	172.25.23.92	FDO20350MHU	106 lines	Out-of-sync	100%
LEAF1	172.25.23.91	FDO20281K6K	106 lines	Out-of-sync	100%

This screen displays all the switches within the fabric with the **Status** column indicating whether the switches are IN-SYNC or OUT-OF-SYNC as per calculations from the Config Compliance module. You can click within the **Preview Config** column for a row that represents a specific switch, for more information. When you do so, the **Config Preview** screen comes up.

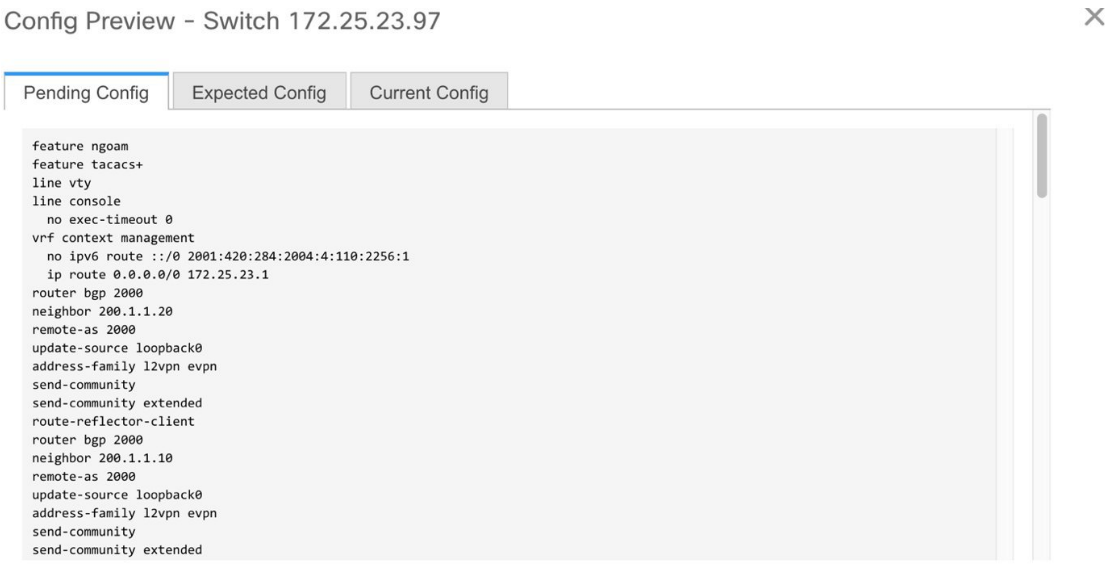

```
feature ngoam
feature tacacs+
line vty
line console
  no exec-timeout 0
vrf context management
  no ipv6 route ::/0 2001:420:284:2004:4:110:2256:1
  ip route 0.0.0.0/0 172.25.23.1
router bgp 2000
neighbor 200.1.1.20
remote-as 2000
update-source loopback0
address-family l2vpn evpn
send-community
send-community extended
route-reflector-client
router bgp 2000
neighbor 200.1.1.10
remote-as 2000
update-source loopback0
address-family l2vpn evpn
send-community
send-community extended
```

The **Pending Config** tab displays the set of configuration that needs to be deployed on the switch, to go from the current running configuration to the current expected/intended configuration. Note that the amount of configuration that shows up in the pending config tab needs to be carefully reviewed before deployment. Typically, if there is even a single line of difference in the configuration associated with a given policy associated with an ENTITY, be it a given interface or a given feature, the pending config will show the entire configuration associated with that policy.

The **Expected Config** and **Current Config** tabs display the expected and current configurations on the switch, respectively. After expected configurations are generated, the switches will be out of Migration-mode.

Close the screen after previewing it. The **Config Deployment** screen comes up again. Preview other switch configurations as needed.

8. Click **Deploy Config** at the bottom part of the **Config Deployment** screen to deploy pending configurations to the switches. This shows up Step 2 of the deployment process, where a per switch deployment status is depicted with an appropriate progress bar. In case there are any errors encountered during the deployment process, the deployment process for that particular switch, will be aborted with a "FAILED" status. The deployment on all the other switches continues to be executed in parallel. For the failure case, by clicking on the "FAILED" status, a pop-up will open up where the details of the configuration deployment history for the switch will be depicted. This in turn can be used to drill down into the exact error that was encountered during the deployment. After addressing the error, the deployment can be re-attempted.

   The **Progress** column displays the deployment progress on each switch.

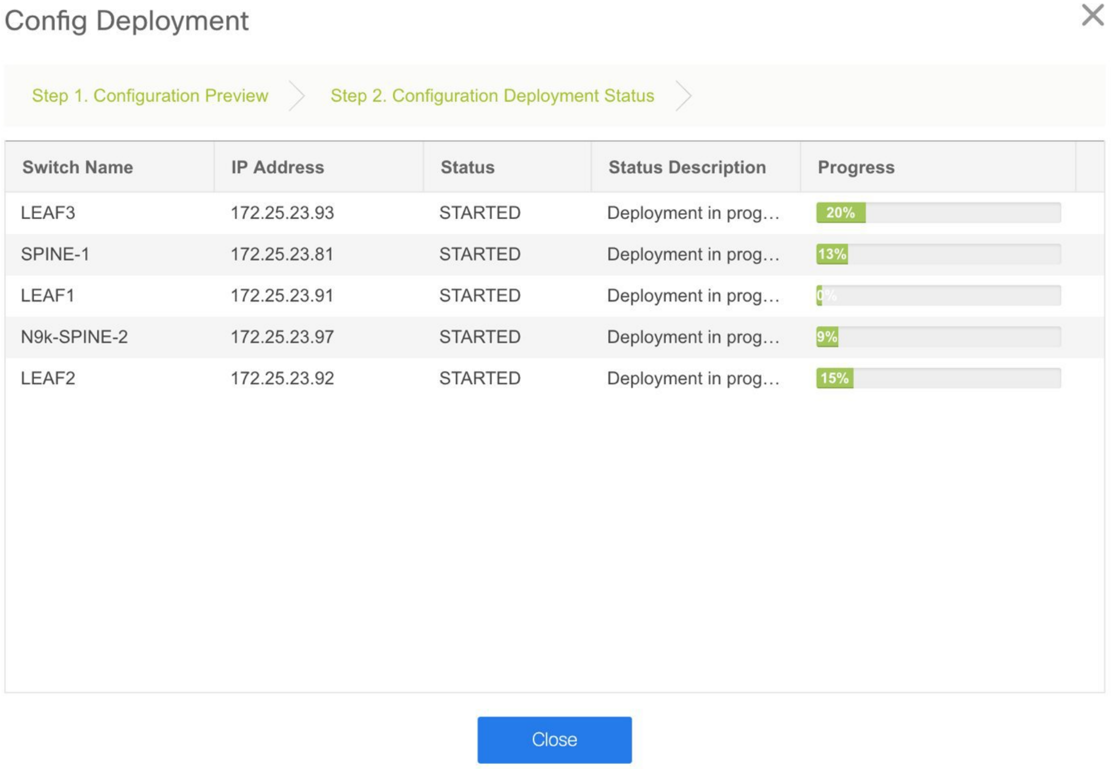

Config Deployment

×
Step 1. Configuration Preview
Step 2. Configuration Deployment Status
Close

Switch Name	IP Address	Status	Status Description	Progress
LEAF3	172.25.23.93	STARTED	Deployment in prog…	20%
SPINE-1	172.25.23.81	STARTED	Deployment in prog…	13%
LEAF1	172.25.23.91	STARTED	Deployment in prog…	0%
N9k-SPINE-2	172.25.23.97	STARTED	Deployment in prog…	9%
LEAF2	172.25.23.92	STARTED	Deployment in prog…	15%

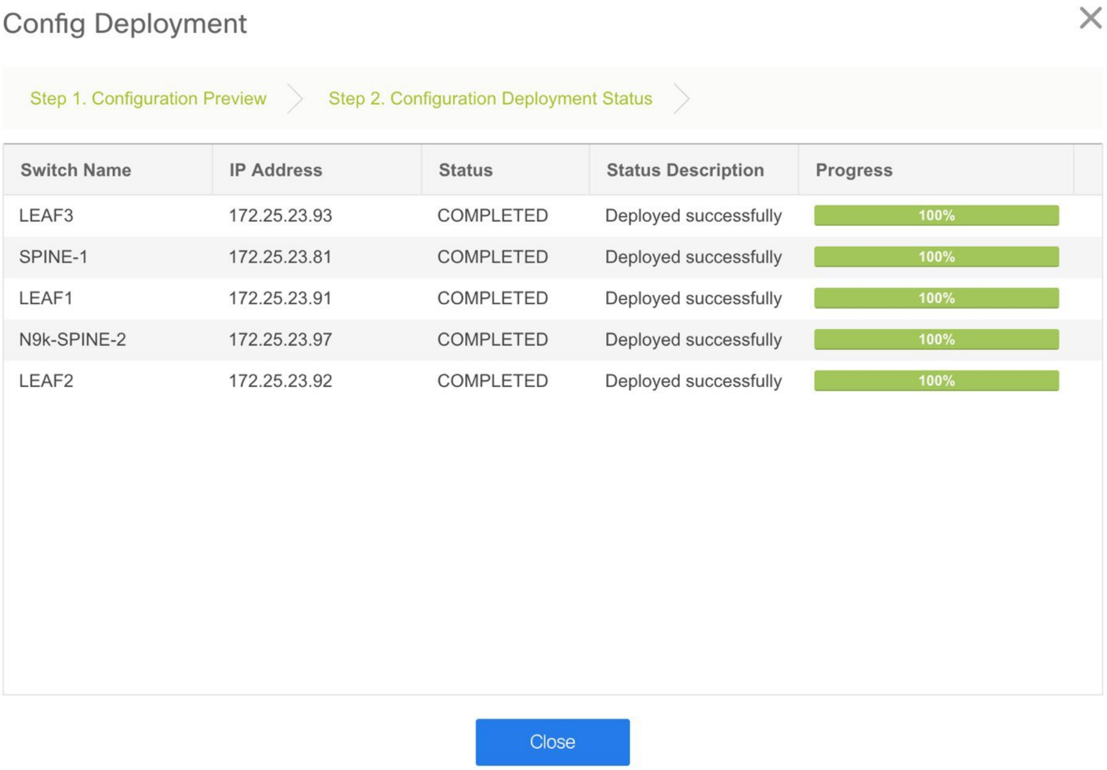

Config Deployment

Step 1. Configuration Preview › Step 2. Configuration Deployment Status

Switch Name	IP Address	Status	Status Description	Progress
LEAF3	172.25.23.93	COMPLETED	Deployed successfully	100%
SPINE-1	172.25.23.81	COMPLETED	Deployed successfully	100%
LEAF1	172.25.23.91	COMPLETED	Deployed successfully	100%
N9k-SPINE-2	172.25.23.97	COMPLETED	Deployed successfully	100%
LEAF2	172.25.23.92	COMPLETED	Deployed successfully	100%

For a successful deployment and an IN-SYNC status for the entire fabric, ensure that the progress column shows 100% for all switches.

9. Click **Close**.

The fabric topology will be displayed. You can see that the Migration-mode icon is no longer visible on the switches and the switch icons are in green color indicating an IN-SYNC status as regards to Configuration Compliance. In this way, the migration/onboarding of the fabric has been achieved.

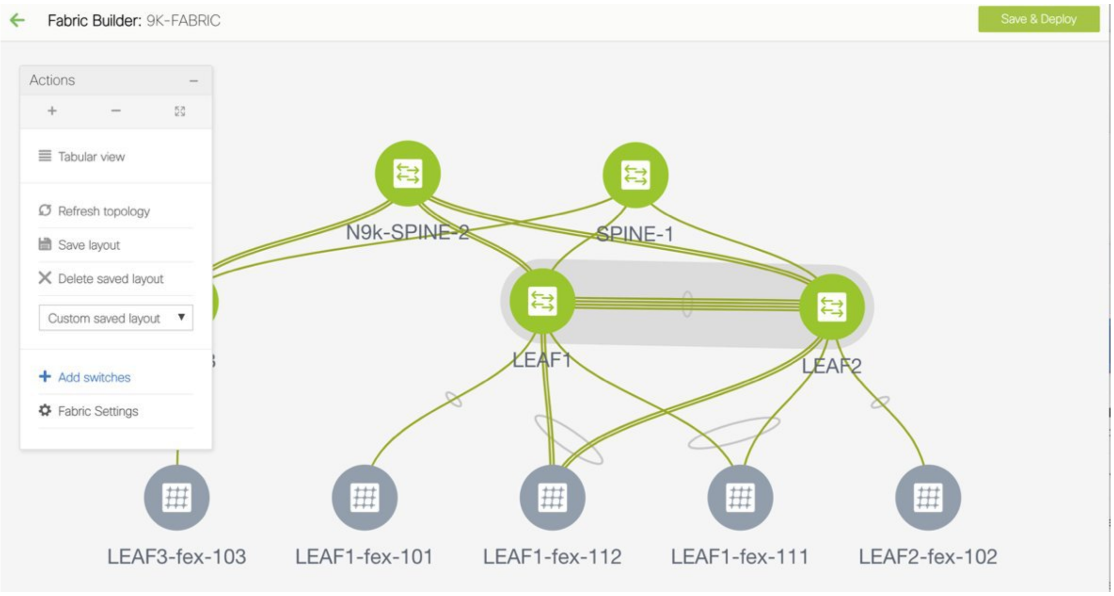

Multi-Site Domain for VXLAN BGP EVPN Fabrics

A Multi-Site Domain (MSD) is a multifabric container that is created to manage multiple member fabrics. An MSD is a single point of control for definition of overlay networks and VRFs that are shared across member fabrics. When you move fabrics (that are designated to be part of the multifabric overlay network domain) under the MSD as member fabrics, the member fabrics share the networks and VRFs created at the MSD-level. This way, you can consistently provision network and VRFs for different fabrics, at one go. It significantly reduces the time and complexity involving multiple fabric provisionings.

Since server networks and VRFs are shared across the member fabrics (as one stretched network), the new networks and VRFs provisioning function is provided at the MSD fabric level. Any new network and VRF creation is only allowed for the MSD. All member fabrics inherit any new network and VRF created for the MSD.

Note

- Network and VRF deployment is not applicable to the MSD fabric since it does not contain any switches, but only contains member fabrics.
- The VXLAN OAM feature in Cisco DCNM is only supported on a single fabric or site.

A few fabric-specific terms:

- **Standalone fabric**: A fabric that is not part of an MSD is referred as a standalone fabric from the MSD perspective. Before the MSD concept, all fabrics were considered standalone, though two or more such fabrics can be connected with each other.
- **Member fabrics**: Fabrics that are part of an MSD are called *member* fabrics or *members*. Create a standalone fabric (of the type *Easy_Fabric*) first and then move it within an MSD as a member fabric.

When a standalone fabric is added to the MSD, the following actions take place:

- The standalone fabric's relevant attributes and the network and VRF definitions are checked against that of the MSD. If there is a *conflict*, then the standalone fabric addition to the MSD fails. If there are no conflicts, then the standalone fabric becomes a member fabric for the MSD. If there is a conflict, the exact conflicts are logged in the pending errors log for the MSD fabric. You can remedy the conflicts and then attempt to add the standalone fabric to the MSD again.
- All the VRFs and networks definitions from the standalone fabric that do not have presence in the MSD are copied over to the MSD and in turn inherited to each of its other existing member fabrics.
- The VRFs (and their definitions) from the MSD (such as the MSD's VRF, and L2 and L3 VNI parameters that *do not* have presence in the standalone fabric) are inherited into the standalone fabric that just became a member.

Fabric and Switch Instance Variables

While the MSD provisions a global range of network and VRF values, some parameters are fabric-specific and some parameters are switch-specific. The parameters are called *fabric instance* and *switch instance* variables.

Fabric instance values can be edited in the fabric context. Specify fabric instance values for each fabric. For example, *multicast group subnet address*.

Switch instance values can be edited on deployment of the network on the switch. For example, *VLAN ID*.

MSD and Member Fabric Process Flow

An MSD has multiple sites (and hence, multiple member fabrics under an MSD). VRFs and networks are created for the MSD and get inherited by the member fabrics. For example, VRF-50000 (and L3 network with ID 50000), and L2 networks with IDs 30000 and 30001 are created for the MSD, in one go.

A high-level flow chart of the MSD and member fabric creation and MSD-to-member fabric inheritance process:

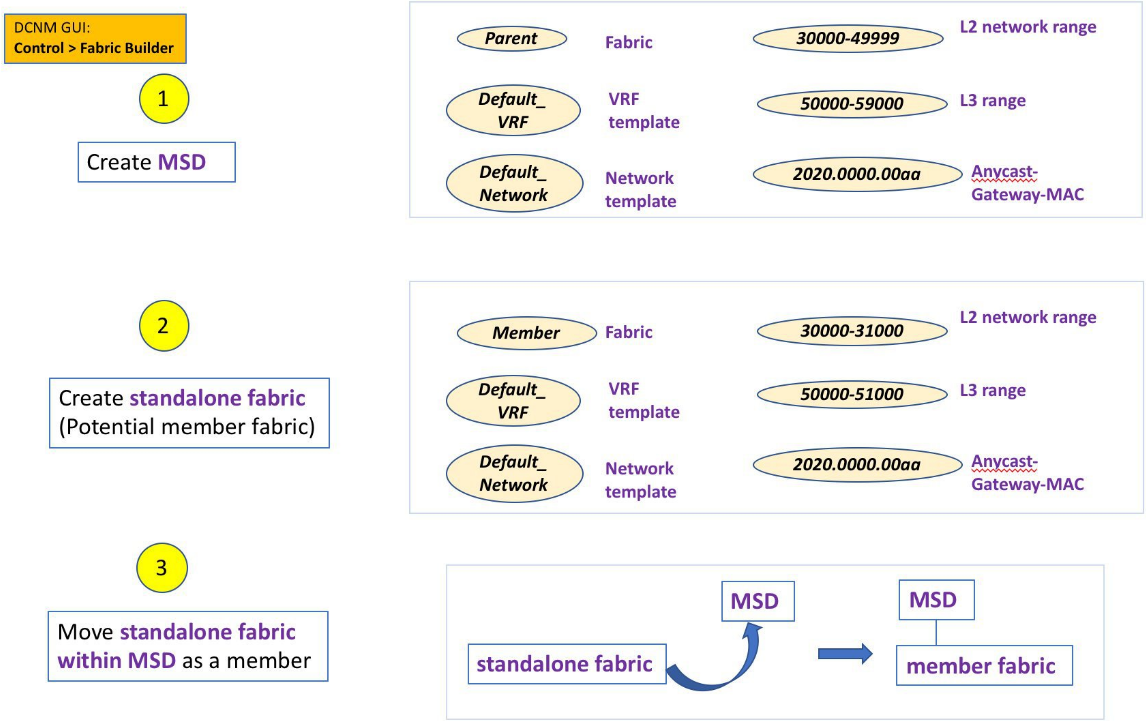

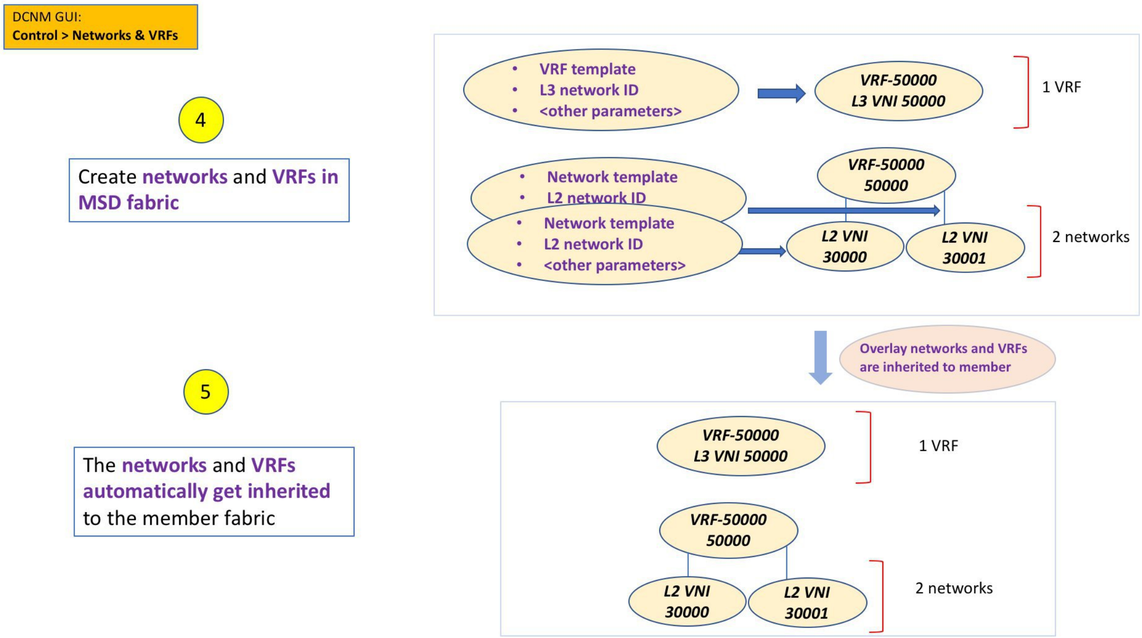

The sample flow explained the inheritance from the MSD to one member. An MSD has multiple sites (and hence, multiple member fabrics under an MSD). A sample flow from an MSD to multiple members:

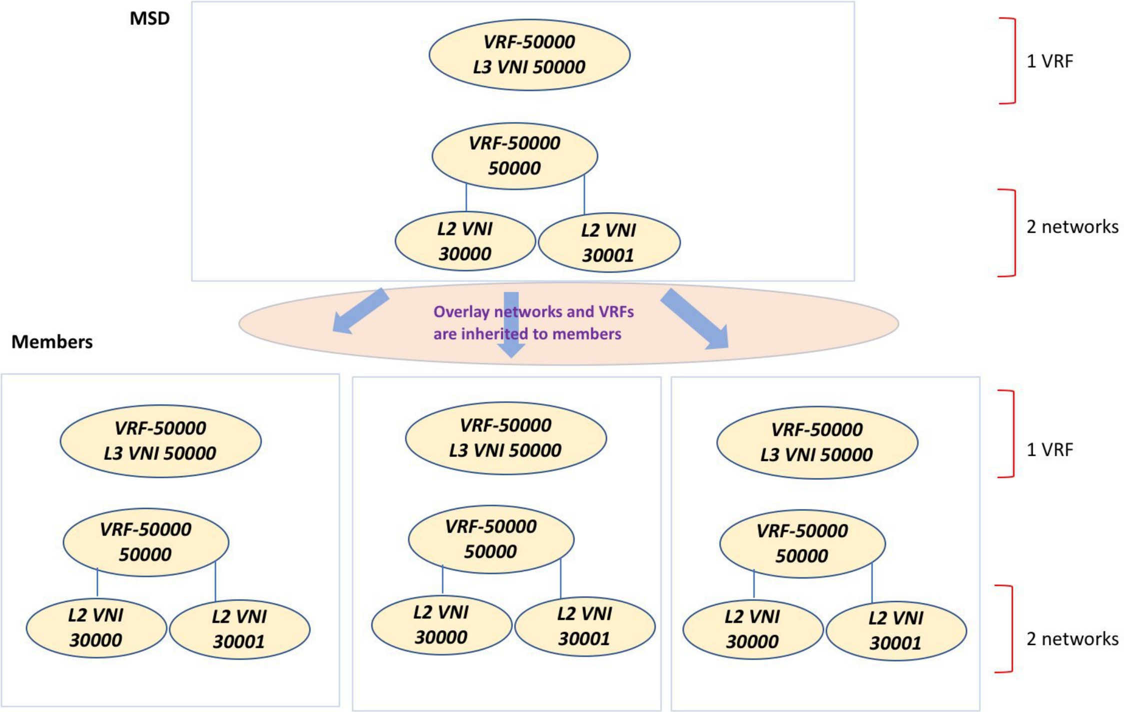

In this example, VRF-50000 (and L3 network with ID 50000), and L2 networks with IDs 30000 and 30001 are created in one go. Networks and VRFs are deployed on the member fabric switches, one after another, as depicted in the image.

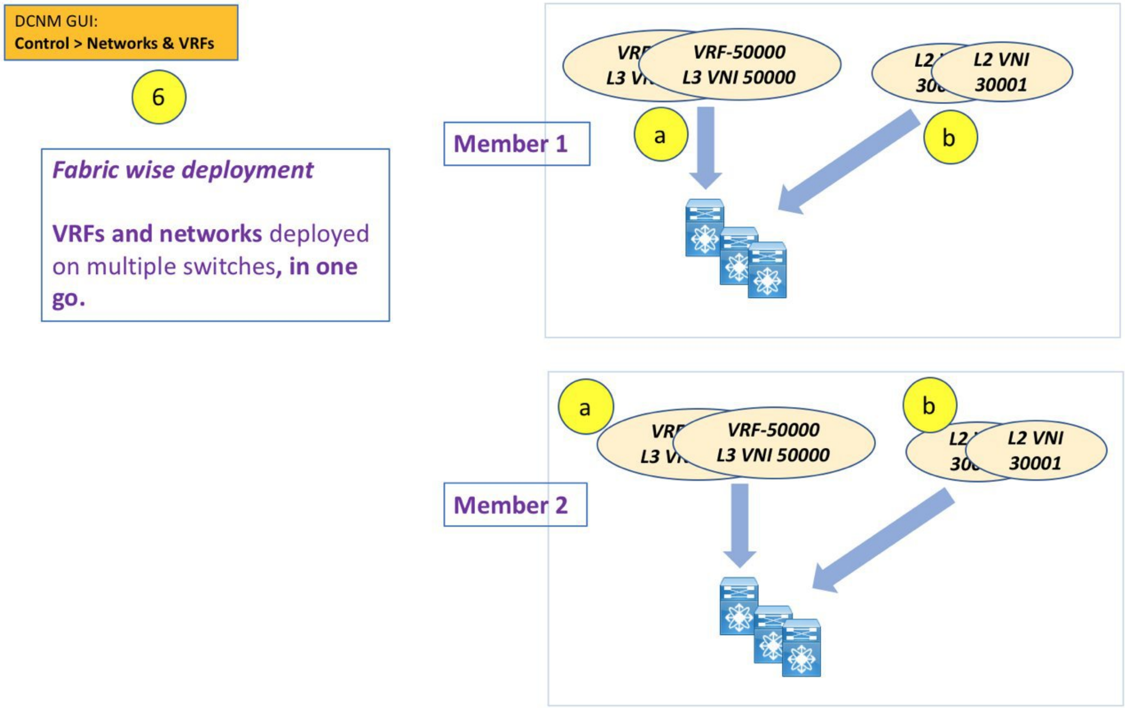

Note If you move a standalone fabric with existing networks and VRFs to an MSD, DCNM does appropriate validation. This is explained in detail in an upcoming section.

Upcoming sections in the document explain the following:

- Creation of an MSD fabric.
- Creation of a standalone fabric (as a potential member) and its movement under the MSD as a member.
- Creation of networks and VRFs in the MSD and their inheritance to the member fabrics.
- Deployment of networks and VRFs in a member fabric's switches.
- Other scenarios for fabric movement:
  - Standalone fabric with existing networks and VRFs to an MSD fabric.
  - Member fabric from one MSD to another.

Create an MSD Fabric and Associate Member Fabrics to It

The process is explained in two steps:

1. Create an MSD fabric.
2. Create a new standalone fabric and move it under the MSD fabric as a member fabric.

Create an MSD Fabric

1. Click **Control > Fabric Builder**.

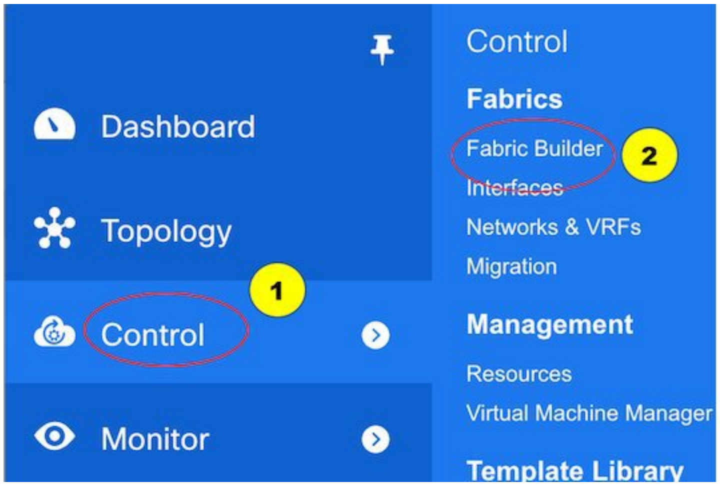

The Fabric Builder page comes up. When you enter for the first time, the **Fabrics** section has no entries. After you create a fabric, it is displayed on the Fabric Builder page, wherein a rectangular box represents each fabric.

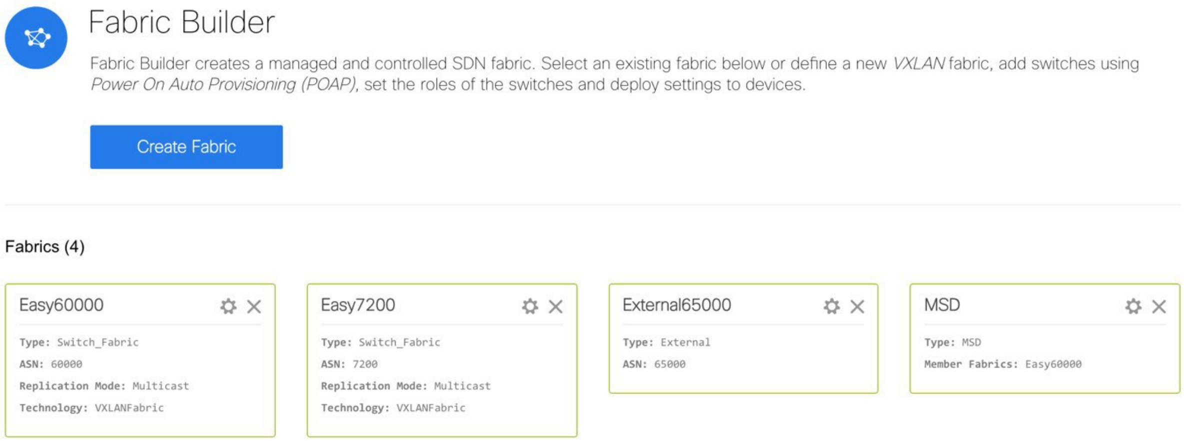

Fabric Builder

Fabric Builder creates a managed and controlled SDN fabric. Select an existing fabric below or define a new *VXLAN* fabric, add switches using *Power On Auto Provisioning (POAP)*, set the roles of the switches and deploy settings to devices.

Create Fabric

Fabrics (4)

Easy60000
Type: Switch_Fabric
ASN: 60000
Replication Mode: Multicast
Technology: VXLANFabric

Easy7200
Type: Switch_Fabric
ASN: 7200
Replication Mode: Multicast
Technology: VXLANFabric

External65000
Type: External
ASN: 65000

MSD
Type: MSD
Member Fabrics: Easy60000

A standalone or member fabric contains *Switch_Fabric* in the **Type** field, its AS number in the **ASN** field and mode of replication, *Multicast* or *Ingress Replication*, in the **Replication Mode** field. Since no device or network traffic is associated with an MSD fabric as it is a container, it does not have these fields.

2. Click the **Create Fabric** button. The Add Fabric screen comes up. The fields are:

   Fabric Name - Enter the name of the fabric.

   Fabric Template - This field has template options for creating specific types of fabric. Choose *MSD_Fabric*. The MSD screen comes up.

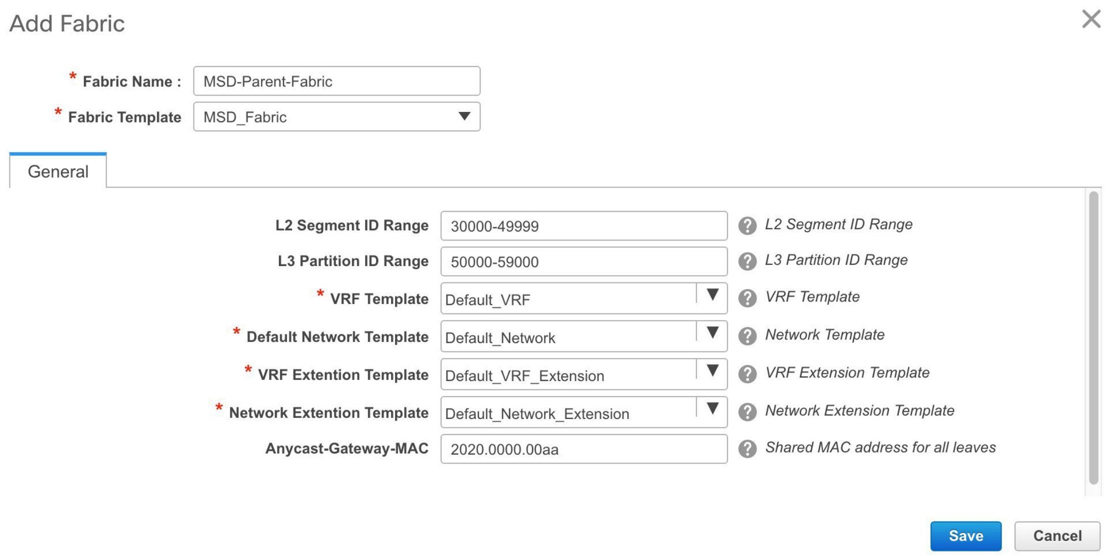

Add Fabric

* Fabric Name : MSD-Parent-Fabric

* Fabric Template MSD_Fabric

General

L2 Segment ID Range 30000-49999 — L2 Segment ID Range

L3 Partition ID Range 50000-59000 — L3 Partition ID Range

* VRF Template Default_VRF — VRF Template

* Default Network Template Default_Network — Network Template

* VRF Extention Template Default_VRF_Extension — VRF Extension Template

* Network Extention Template Default_Network_Extension — Network Extension Template

Anycast-Gateway-MAC 2020.0000.00aa — Shared MAC address for all leaves

The fields in the screen are explained:

In the **General** tab, all fields are autopopulated with data. The fields consist of the Layer 2 and Layer 3 VXLAN segment identifier range, the default network and VRF templates, and the anycast gateway MAC address. Update the relevant fields as needed.

L2 Segment ID Range - Layer 2 VXLAN segment identifier range.

L3 Partition ID Range - Layer 3 VXLAN segment identifier range.

VRF Template - Default VRF template.

Default Network Template - Default network template.

VRF Extension Template - Default VRF extension template.

Network Extension Template - Default network extension template.

Anycast-Gateway-MAC - Anycast gateway MAC address.

3. Click **Save**.

   A message appears briefly at the bottom right part of the screen, indicating that you have created a new MSD fabric. After fabric creation, the fabric page comes up. The fabric name *MSD-Parent-Fabric* appears at the top left part of the screen.

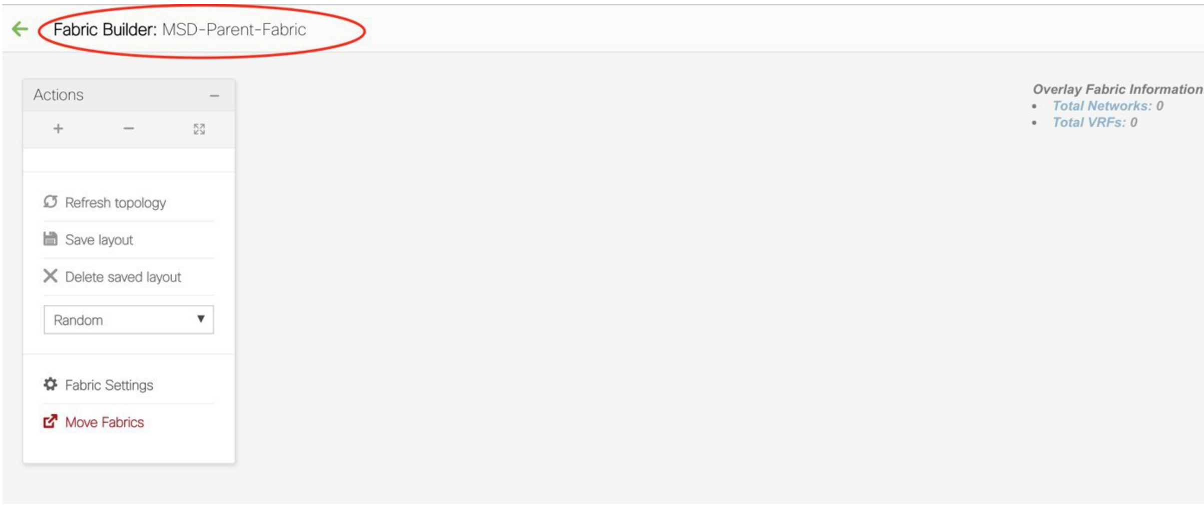

Since the MSD fabric is a container, you cannot add a switch to it. The **Add Switches** button that is available in the **Actions** panel for member and standalone fabrics is not available for the MSD fabric.

When a new MSD is created, the newly created MSD fabric instance appears (as a rectangular box) on the Fabric Builder page. To go to the Fabric Builder page, click the ← button at the top left part of the *MSD-Parent-Fabric* page.

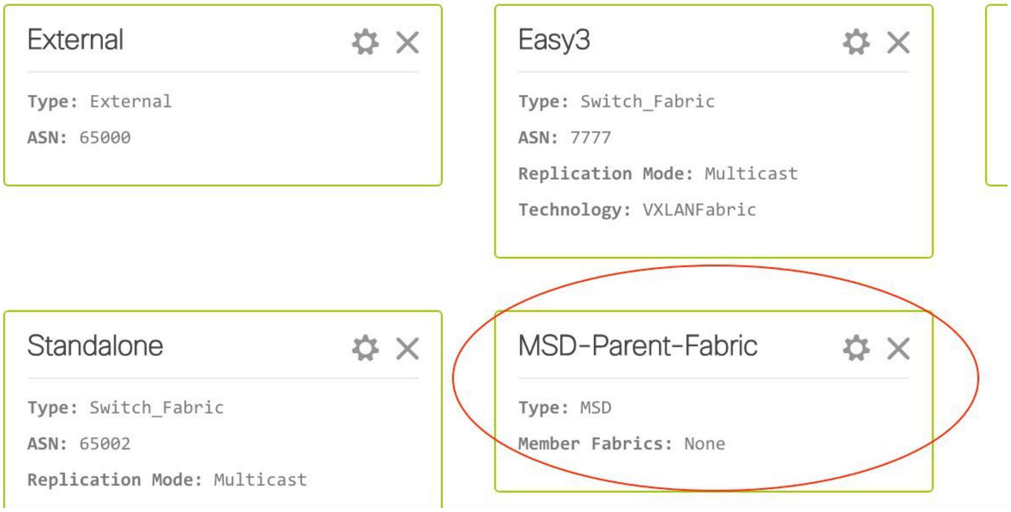
External

Type: External

ASN: 65000

Easy3

Type: Switch_Fabric

ASN: 7777

Replication Mode: Multicast

Technology: VXLANFabric

Standalone

Type: Switch_Fabric

ASN: 65002

Replication Mode: Multicast

MSD-Parent-Fabric

Type: MSD

Member Fabrics: None

An MSD fabric is displayed as *MSD* in the **Type** field, and it contains the member fabric names in the **Member Fabrics** field. When no member fabric is created, *None* is displayed.

The steps for creation of an MSD fabric and moving member fabrics under it are:

1. Create an MSD fabric.
2. **Create a new standalone fabric and move it under the MSD fabric as a member fabric.**

Step 1 is completed. Step 2 is explained in the next section.

Create and Move a New Fabric Under the MSD Fabric as a Member

A new fabric is created as a standalone fabric. After you create a new fabric, you can move it under an MSD as a member. As a best practice, when you create a new fabric that is a potential member fabric (of an MSD), do not add networks and VRFs to the fabric. Move the fabric under the MSD and then add networks and VRFs for the MSD. That way, there will not be any need for validation (or conflict resolution) between the member and MSD fabric network and VRF parameters.

New fabric creation is explained in the Easy Fabric creation process. In the MSD document, fabric movement is covered. The values that are displayed in the screen are automatically generated. The VXLAN VNI ID ranges (in the L2 Segment ID Range and L3 Partition ID Range fields) allocated for new network and VRF creation are values from the MSD fabric segment ID range. If you want to update the VXLAN VNI ranges or the VRF and Network VLAN ranges, ensure the following:

- If you update a range of values, ensure that it does not overlap with other ranges.
- You must update one range of values at a time. If you want to update more than one range of values, do it in separate instances. For example, if you want to update L2 and L3 ranges, you should do the following:
  1. Update the L2 range and click **Save**.
  2. Click the **Edit Fabric** option again, update the L3 range and click **Save**.

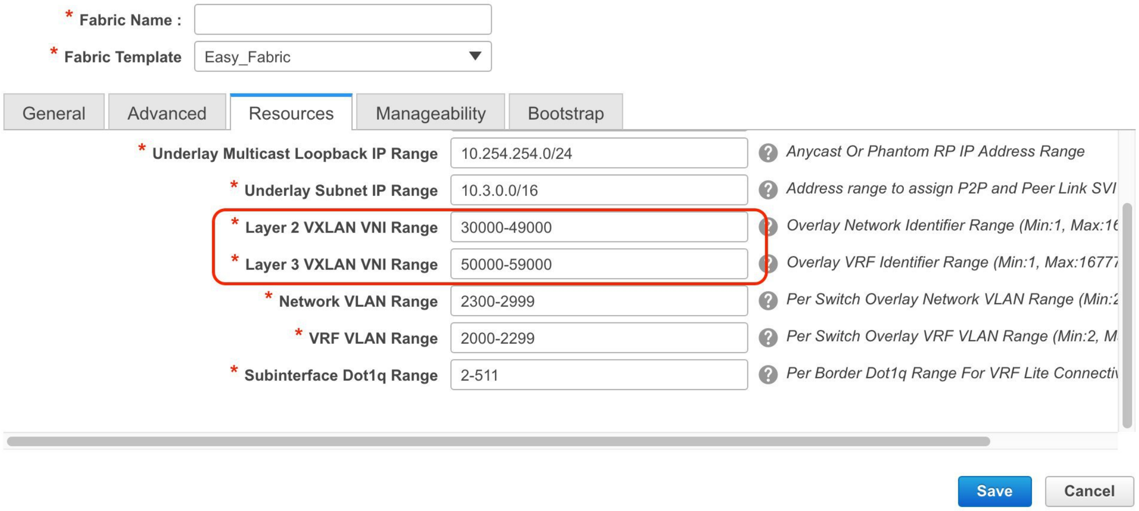

* Fabric Name :		
* Fabric Template	Easy_Fabric	
* Underlay Multicast Loopback IP Range	10.254.254.0/24	Anycast Or Phantom RP IP Address Range
* Underlay Subnet IP Range	10.3.0.0/16	Address range to assign P2P and Peer Link SVI
* Layer 2 VXLAN VNI Range	30000-49000	Overlay Network Identifier Range (Min:1, Max:16
* Layer 3 VXLAN VNI Range	50000-59000	Overlay VRF Identifier Range (Min:1, Max:16777
* Network VLAN Range	2300-2999	Per Switch Overlay Network VLAN Range (Min:2
* VRF VLAN Range	2000-2299	Per Switch Overlay VRF VLAN Range (Min:2, M
* Subinterface Dot1q Range	2-511	Per Border Dot1q Range For VRF Lite Connecti

Ensure that the **Anycast Gateway MAC**, the **Network Template** and the **VRF Template** field values are the same as the MSD fabric. Else, member fabric movement to the MSD fail.

Other pointers:

- The member fabric should have a Site ID configured and the Site ID must be unique among the members.
- The BGP AS number should be unique for a member fabric.
- The underlay subnet range for loopback0 should be unique.
- The underlay subnet range for loopback1 should be unique.

After you click **Save**, a note appears at the bottom right part of the screen indicating that the fabric is created. When a fabric is created, the fabric page comes up. The fabric name appears at the top left part of the screen. Simultaneously, the Fabric Builder page also displays the newly created *Member1* fabric.

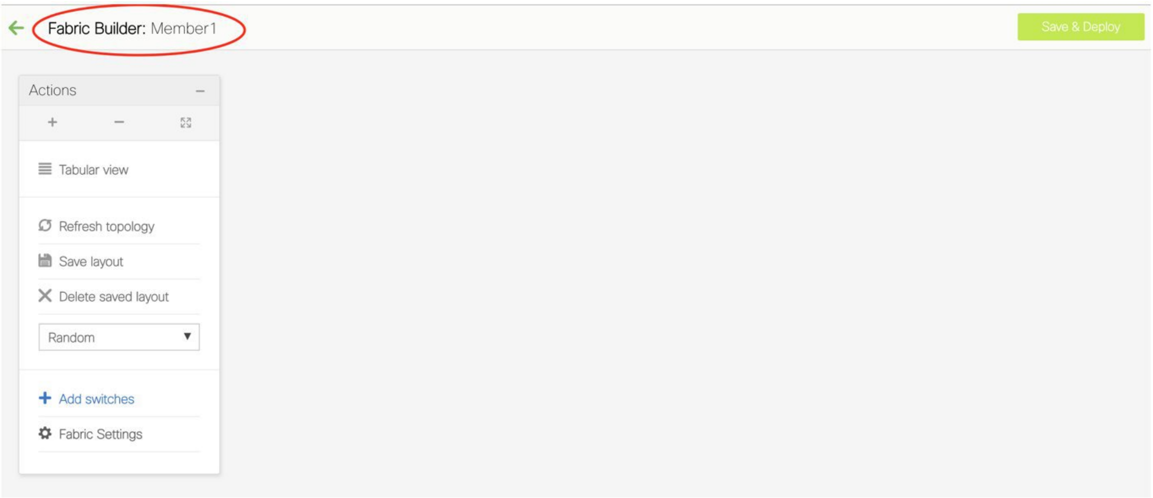

Simultaneously, the Fabric Builder page also displays the newly created fabric, *Member1*.

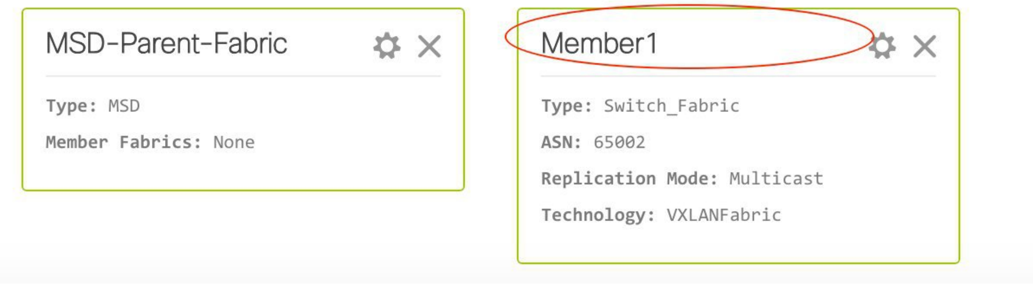

MSD-Parent-Fabric

Type: MSD

Member Fabrics: None

Member1

Type: Switch_Fabric

ASN: 65002

Replication Mode: Multicast

Technology: VXLANFabric

Move the Member1 Fabric Under MSD-Parent-Fabrics

You should go to the MSD fabric page to associate a member fabric under it.

If you are on the Fabric Builder page, click within the **MSD-Parent-Fabric** box to go to the MSD-Parent-Fabric page.

[If you are in the *Member1* fabric page, you should go to the MSD-Parent-Fabrics-Docs fabric page. Click **<-** above the **Actions** panel. You will reach the Fabric Builder page. Click within the **MSD-Parent-Fabric** box].

1. In the MSD-Parent-Fabric page, go to the **Actions** panel and click **Move Fabrics**.

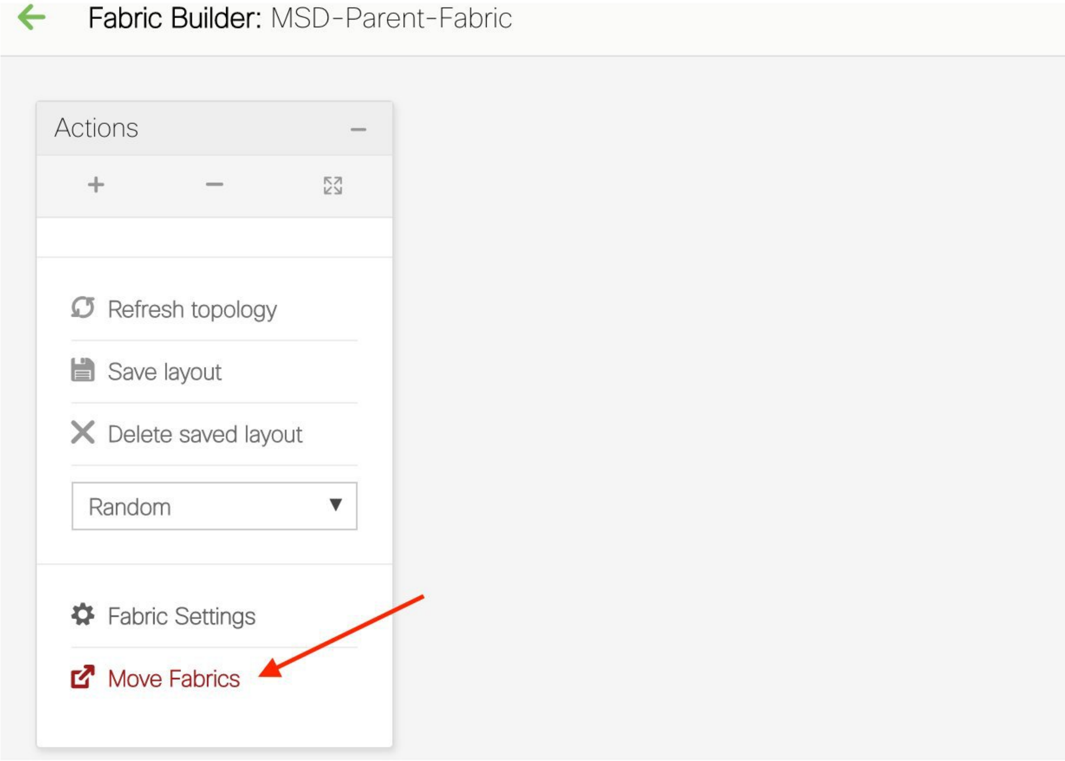

The Move Fabric screen comes up. It contains a list of fabrics.

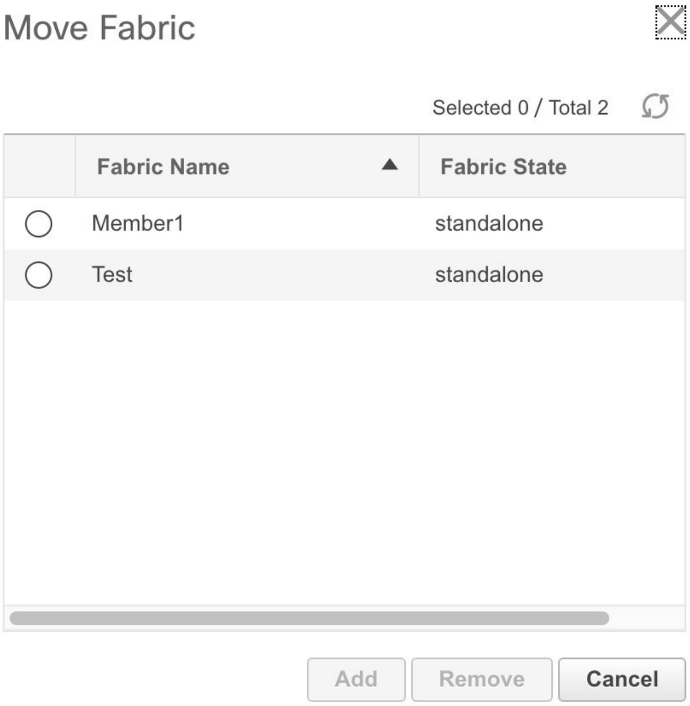

	Fabric Name	Fabric State
○	Member1	standalone
○	Test	standalone

Member fabrics of other MSD container fabrics will not be displayed here.

The *Member1* fabric is still a standalone fabric as seen in the image. A fabric is considered a member fabric of an MSD fabric only when you associate it with the MSD fabric. Also, each standalone fabric is a candidate for being an MSD fabric member, until you associate it to one of the MSD fabrics.

2. Since *Member1* fabric is to be associated with the MSD fabric, select the **Member1** radio button. The **Add** button is enabled.
3. Click **Add**.

   Immediately, a message appears at the top of the screen indicating that the *Member1* fabric is now associated with the MSD fabric *MSD-Parent-Fabric*. Now, the MSD-Parent-Fabric fabric page appears again.
4. Click the **Move Fabrics** option to check the fabric status. You can see that the fabric status has changed from standalone to member.

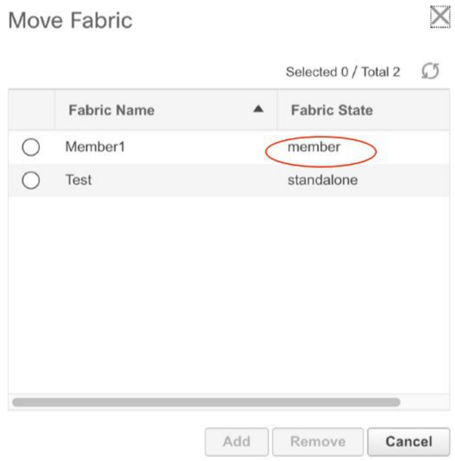

Move Fabric

Selected 0 / Total 2

	Fabric Name ▲	Fabric State
○	Member1	member
○	Test	standalone

5. Close this screen. Now, in the MSD-Parent-Fabric screen the member fabric icon is displayed.

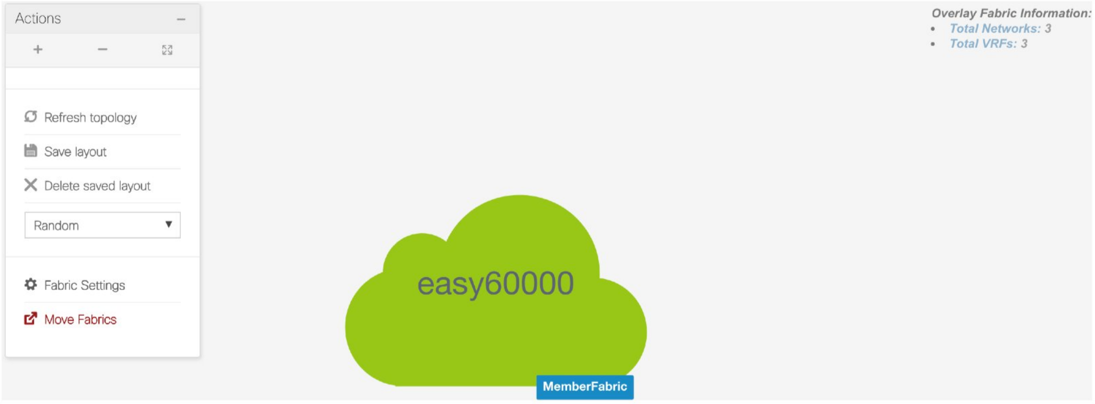

Overlay Fabric Information:
- Total Networks: 3
- Total VRFs: 3

6. Click ← above the Actions panel to go to the Fabric Builder page.

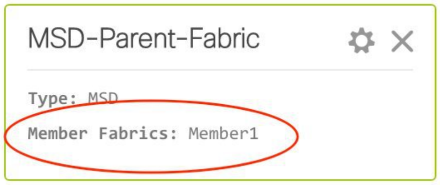

MSD-Parent-Fabric

Type: MSD

Member Fabrics: Member1

You can see that *Member1* is now added to MSD fabric and is displayed in the **Member Fabrics** field.

Networks and VRFs Creation and Deployment in an MSD Fabric

In standalone fabrics, networks and VRFs are created for each fabric. In an MSD fabric, networks and VRFs should be created at the MSD fabric level. The networks and VRFs are inherited by all the member networks. You cannot create or delete networks and VRFs for member fabrics. However, you can edit them.

For example, consider an MSD fabric with two member fabrics. If you create three networks in the MSD fabric, then all three networks will automatically be available for deployment in both the member fabrics.

Though member fabrics inherit the MSD fabric's networks and VRFs, you have to deploy the networks and VRFs distinctly, for each fabric.

Note Networks and VRFs are the common identifiers (represented across member fabrics) that servers (or end hosts) are grouped under so that traffic can be sent between the end hosts based on the network and VRF IDs, whether they reside in the same or different fabrics. Since they have common representation across member fabrics, networks and VRFs can be provisioned at one go. As the switches in different fabrics are physically and logically distinct, you have to deploy the same networks and VRFs separately for each fabric.

For example, if you create networks 30000 and 30001 for an MSD that contains two member fabrics, the networks are automatically created for the member fabrics and are available for deployment. But you have to deploy the networks 30000 and 30001 in one fabric, and then in the other.

Networks and VRFs are created in the MSD and deployed in the member fabrics. The steps are explained below:

1. Create networks and VRFs in the MSD fabric.
2. Deploy the networks and VRFs in the member fabric devices, one fabric at a time.

Create Networks in the MSD Fabric

1. Click **Control > Networks & VRFs** (under Fabrics submenu). The LAN Fabric Provisioning page comes up.

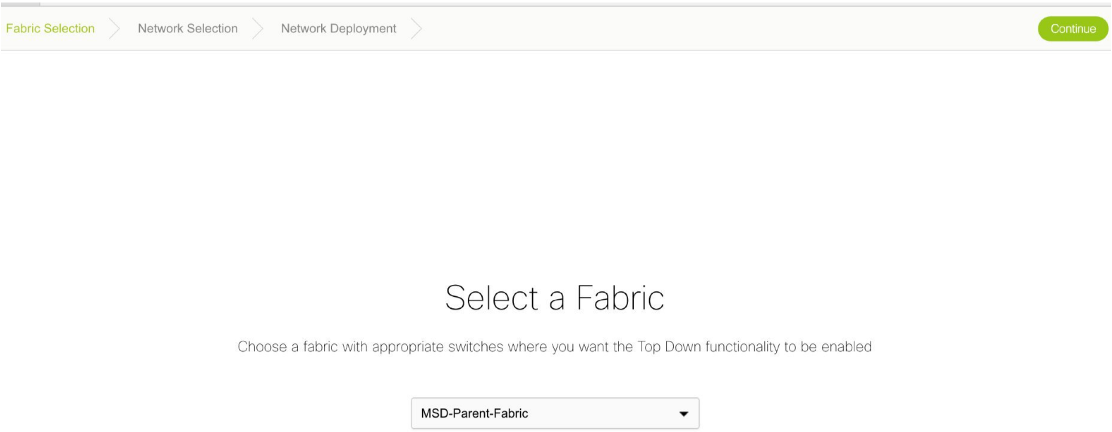

Select a Fabric

Choose a fabric with appropriate switches where you want the Top Down functionality to be enabled

MSD-Parent-Fabric

2. Click **Continue**. The Select a Fabric page comes up. Click the **Select a Fabric** drop-down box to see the list of fabrics.

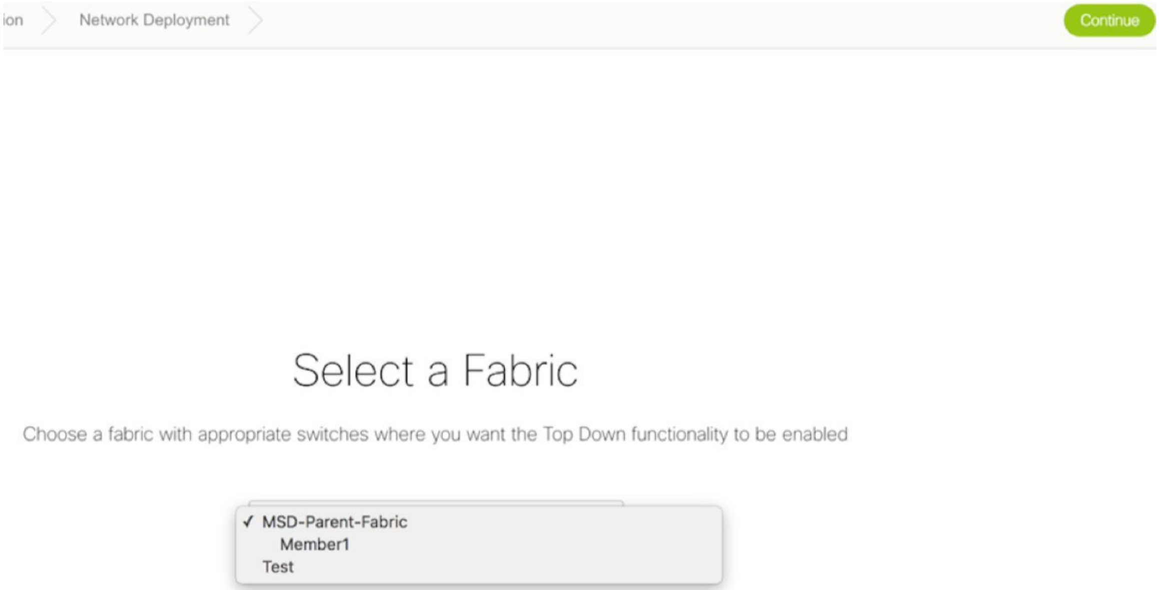

Select a Fabric

Choose a fabric with appropriate switches where you want the Top Down functionality to be enabled

✓ MSD-Parent-Fabric
  Member1
Test

The MSD fabric *MSD-Parent-Fabric* contains one member fabric, *Member1*. It is indented to the right, indicating that is a part of the MSD. All other standalone fabrics appear in the same indent level of the MSD.

Select *MSD-Parent-Fabric* from the list. The Select a Fabric screen for an MSD fabric comes up. Since this is a container of member fabrics and does not have any devices associated with it, associated device-relevant functions will not be seen in the GUI (for example, the Fabric Extension Setup option only appears for standalone and member fabrics).

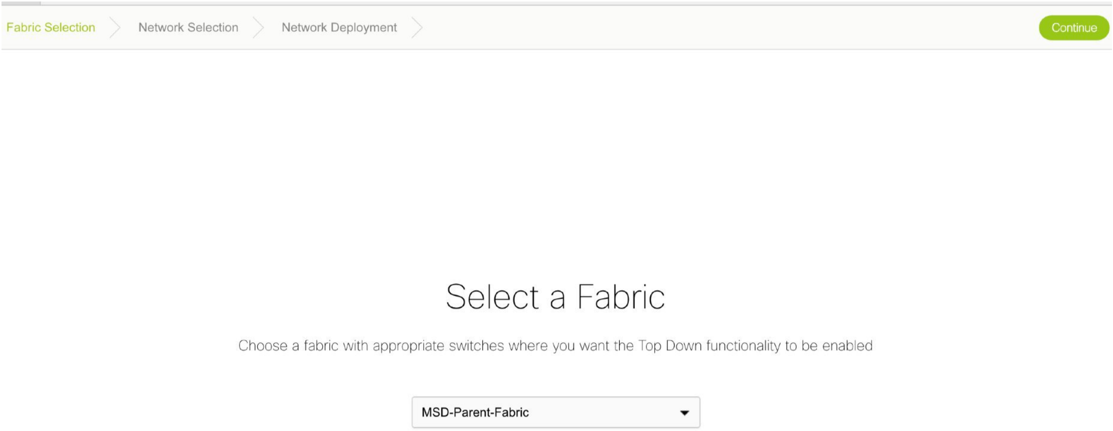

Select a Fabric

Choose a fabric with appropriate switches where you want the Top Down functionality to be enabled

MSD-Parent-Fabric

3. Click **Continue** on the top right part of the screen. The Networks page comes up. This lists the list of networks created for the MSD fabric. Initially, this screen has no entries.

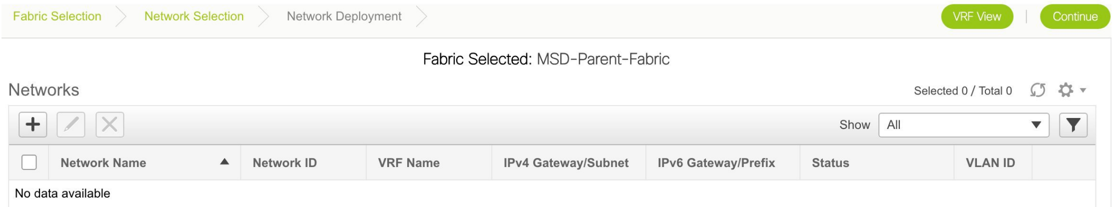

Fabric Selected: MSD-Parent-Fabric

□	Network Name ▲	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
No data available							

4. Click the + button at the top left part of the screen (under **Networks**) to add networks to the MSD fabric. The Create Network screen comes up. Most of the fields are autopopulated.

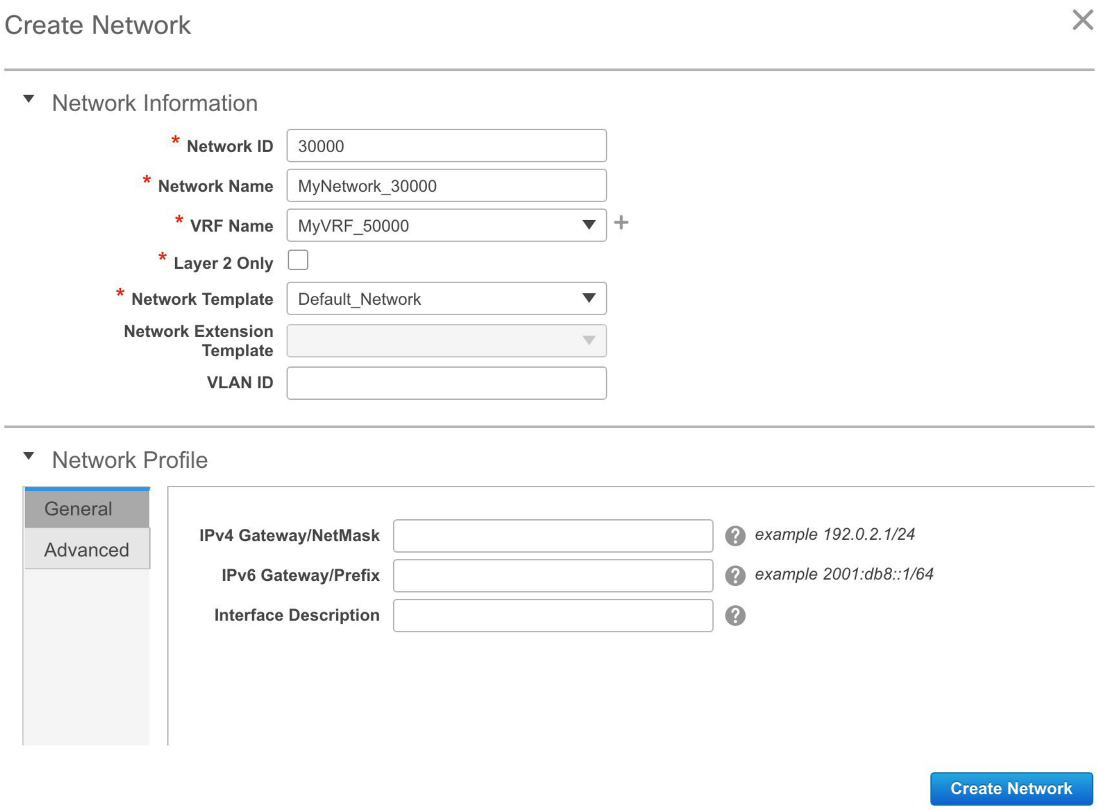

Create Network

Network Information

* Network ID: 30000
* Network Name: MyNetwork_30000
* VRF Name: MyVRF_50000
* Layer 2 Only: □
* Network Template: Default_Network
Network Extension Template:
VLAN ID:

Network Profile

IPv4 Gateway/NetMask: example 192.0.2.1/24
IPv6 Gateway/Prefix: example 2001:db8::1/64
Interface Description:

The fields in this screen are:

Network ID and **Network Name** - Specifies the Layer 2 VNI and name of the network. The network name should not contain any white spaces or special characters except underscore (_) and hyphen (-).

VRF Name - Allows you to select the Virtual Routing and Forwarding (VRF).

When no VRF is created, this field will be blank. If you want to create a new VRF, click the + button. The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:).

Note You can also create a VRF by clicking the VRF View button on the Networks page.

Layer 2 Only - Specifies whether the network is Layer 2 only.

Network Template - Allows you to select a network template.

Network Extension Template - This template allows you to extend the network between member fabrics.

VLAN ID - Specifies the corresponding tenant VLAN ID for the network.

Network Profile section contains the General and Advanced tabs, explained below.

General tab

IPv4 Gateway/NetMask - Specifies the IPv4 address with subnet.

IPv6 Gateway/Prefix - Specifies the IPv6 address with subnet.

Interface Description - Specifies the description for the interface.

Advanced tab - Optionally, specify the advanced profile settings by clicking the **Advanced** tab. The options are:

- ARP Suppression
- DHCPv4 Server
- DHCPv4 Server VRF
- MTU for the L3 interface

A sample of the Create Network screen is given below.

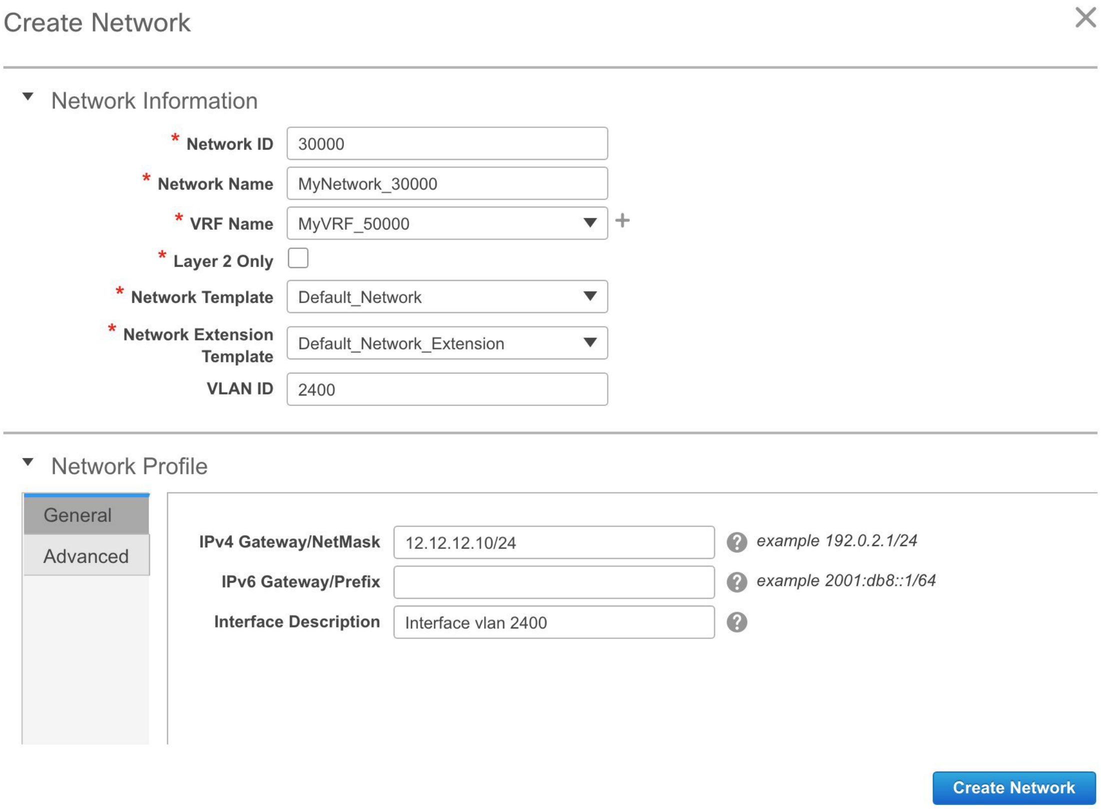

Create Network

Network Information

* Network ID: 30000

* Network Name: MyNetwork_30000

* VRF Name: MyVRF_50000

* Layer 2 Only: □

* Network Template: Default_Network

* Network Extension Template: Default_Network_Extension

VLAN ID: 2400

Network Profile

IPv4 Gateway/NetMask: 12.12.12.10/24 — *example 192.0.2.1/24*

IPv6 Gateway/Prefix: — *example 2001:db8::1/64*

Interface Description: Interface vlan 2400

Advanced tab:

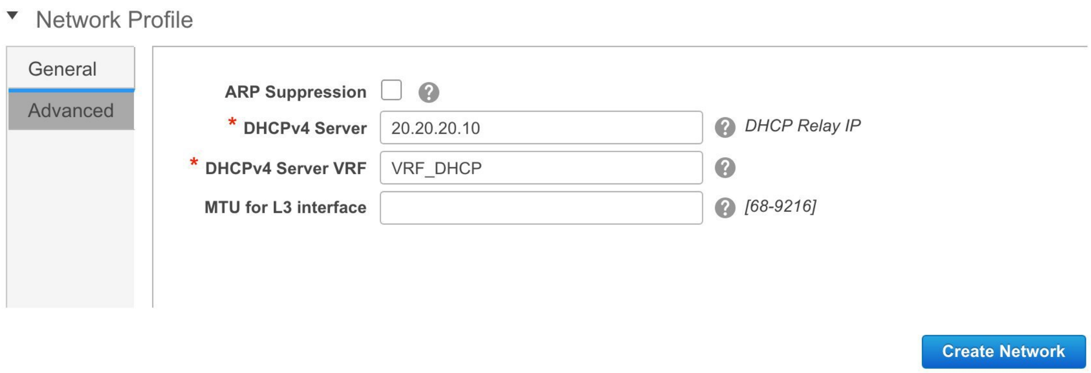

▾ Network Profile

ARP Suppression □

* DHCPv4 Server: 20.20.20.10 — DHCP Relay IP

* DHCPv4 Server VRF: VRF_DHCP

MTU for L3 interface: — [68-9216]

5. Click **Create Network**. A message appears at the bottom right part of the screen indicating that the network is created. The new network (*MyNetwork_30000*) appears on the Networks page that comes up.

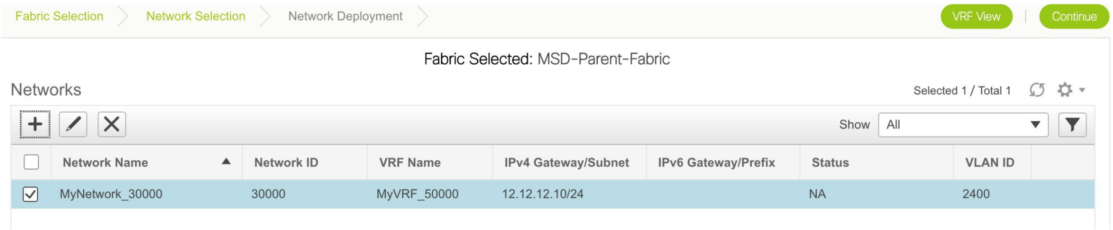

Fabric Selected: MSD-Parent-Fabric

Networks

Selected 1 / Total 1

Show: All

□	Network Name ▲	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
☑	MyNetwork_30000	30000	MyVRF_50000	12.12.12.10/24		NA	2400

Editing and Deleting Networks in the MSD Fabric

You can edit the **Network Profile** part (**General** and **Advanced** tabs) of the network, including the IPv4 gateway IP address, the DHCP information and the ARP suppression feature.

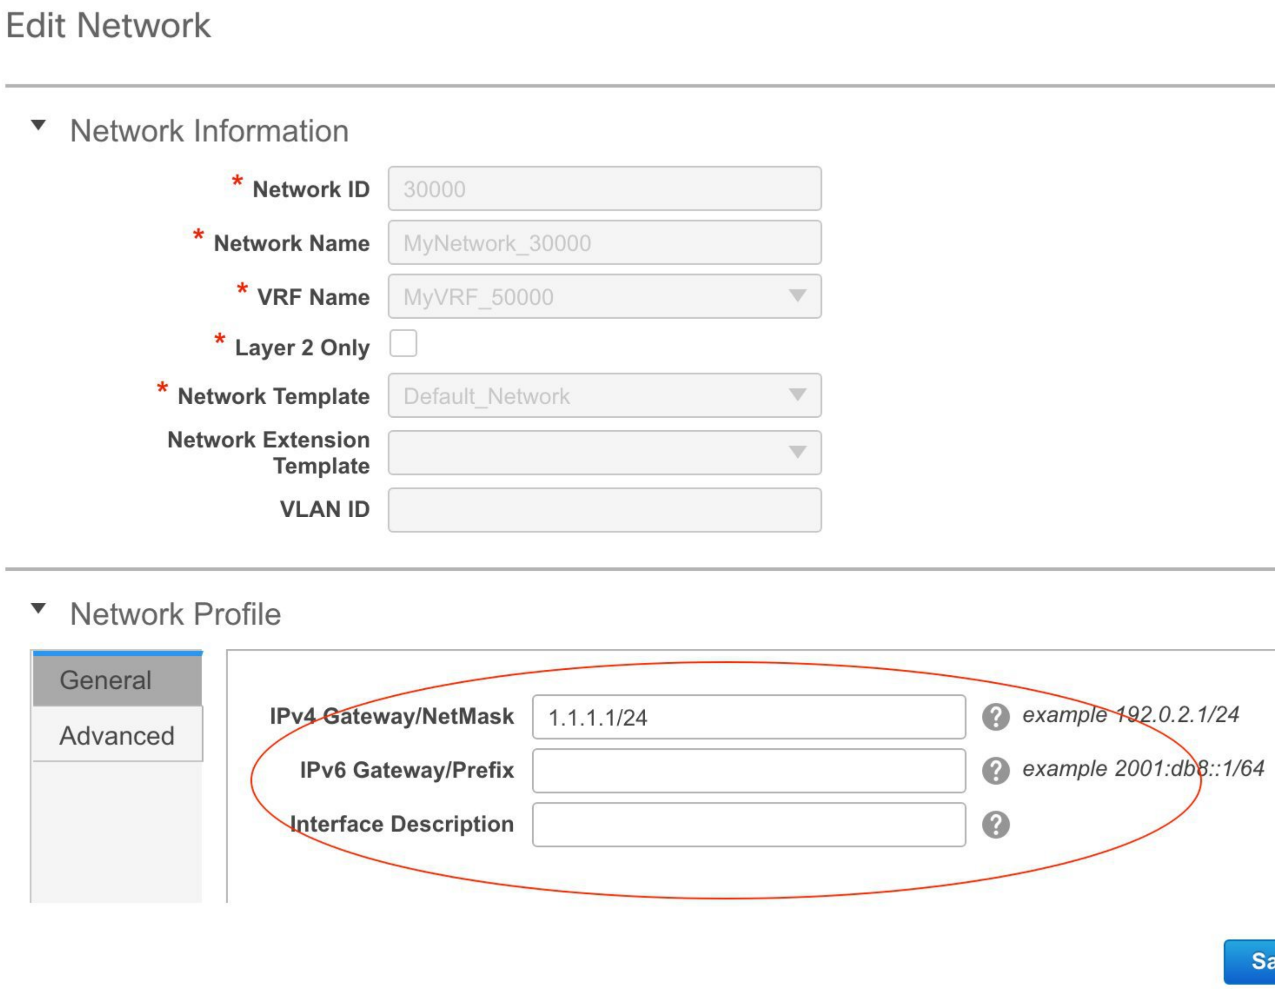

Edit Network

Network Information

* Network ID: 30000
* Network Name: MyNetwork_30000
* VRF Name: MyVRF_50000
* Layer 2 Only: □
* Network Template: Default_Network

Network Extension Template:

VLAN ID:

Network Profile

IPv4 Gateway/NetMask: 1.1.1.1/24 — example 192.0.2.1/24

IPv6 Gateway/Prefix: — example 2001:db8::1/64

Interface Description:

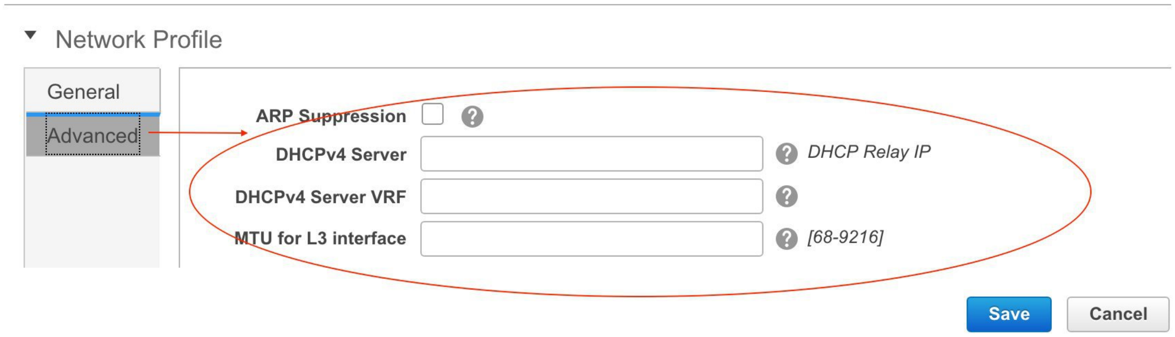

Network Profile

ARP Suppression □

DHCPv4 Server: — DHCP Relay IP

DHCPv4 Server VRF:

MTU for L3 interface: — [68-9216]

In a standalone fabric, you can proceed to deploy the networks on the fabric's devices. But since this is an MSD container fabric that has no physical devices associated with it, you should deploy the networks through the individual member fabric, for each fabric.

A network or VRF deployed in a member fabric cannot be deleted until all instances are undeployed.

Network Inheritance from MSD-Parent-Fabric to Member1

MSD-Parent-Fabric fabric contains one member fabric, *Member1*. Go to the Select a Fabric page to access the *Member1* fabric.

(To go to the Select a Fabric page do one of the following:

- Click the **Fabric Selection** button at the top left part of the screen.
- From the main menu, click **Control > Networks & VRFs** and click **Continue** in the LAN Fabric Provisioning page.

1. Click *Member1* from the drop-down box.
2. Click **Continue** on the top right part of the screen. The Networks page comes up. You can see that the network created for the MSD is inherited to its member.

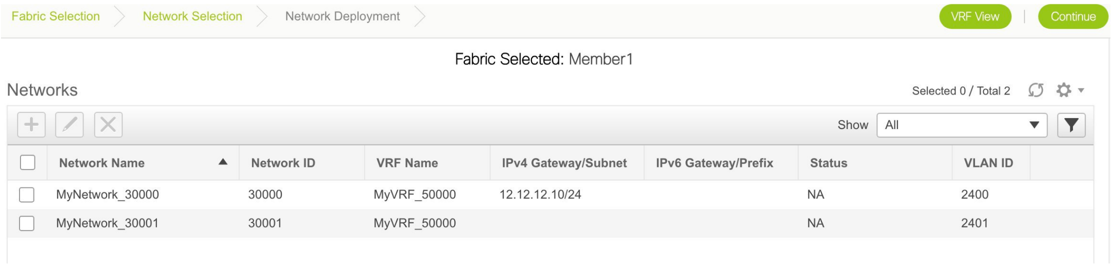

Fabric Selected: Member1

□	Network Name ▲	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
□	MyNetwork_30000	30000	MyVRF_50000	12.12.12.10/24		NA	2400
□	MyNetwork_30001	30001	MyVRF_50000			NA	2401

Editing Networks in the Member Fabric

You can only create and delete networks for the MSD fabric, and not for the member fabric. However, you can update a network's multicast group address since it is a fabric instance variable.

1. Select the network and click the *Edit* option at the top left part of the screen.

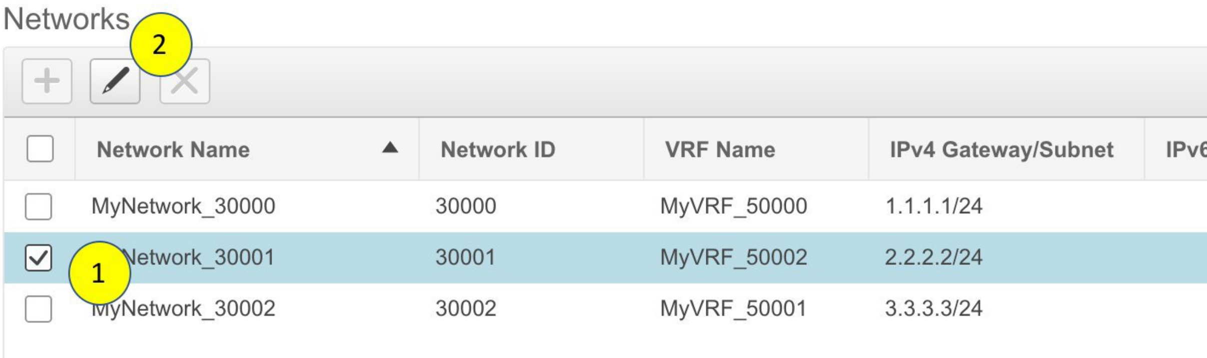

□	Network Name ▲	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6
□	MyNetwork_30000	30000	MyVRF_50000	1.1.1.1/24	
☑	MyNetwork_30001	30001	MyVRF_50002	2.2.2.2/24	
□	MyNetwork_30002	30002	MyVRF_50001	3.3.3.3/24	

2. In the Edit Networks screen that comes up, click the **Advanced** tab in the **Network Profile** section. Update the multicast group address and click **Save**.

   This option is only available for member fabrics and not MSD networks.

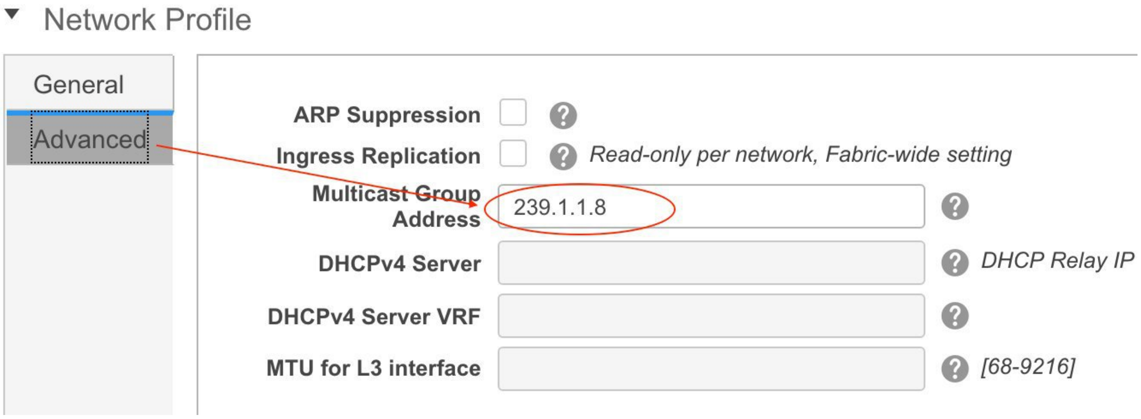

ARP Suppression □

Ingress Replication □ Read-only per network, Fabric-wide setting

Multicast Group Address: 239.1.1.8

DHCPv4 Server: DHCP Relay IP

DHCPv4 Server VRF:

MTU for L3 interface: [68-9216]

Create VRFs in the MSD Fabric

1. From the MSD fabric's Networks page, click the **VRF View** button at the top right part of the screen to create VRFs.

   [If you have freshly logged in to DCNM, do the following:

   Click **Control > Networks & VRFs**, click **Continue** in the LAN Fabric Provisioning page and choose the MSD fabric (*MSD-Parent-Fabric*) from the drop-down box.

   Click **Continue** to reach the Networks page and click **VRF View** at the top right part of the Networks page].

   The VRFs page comes up. This lists the list of VRFs created for the MSD fabric. Initially, this screen has no entries.

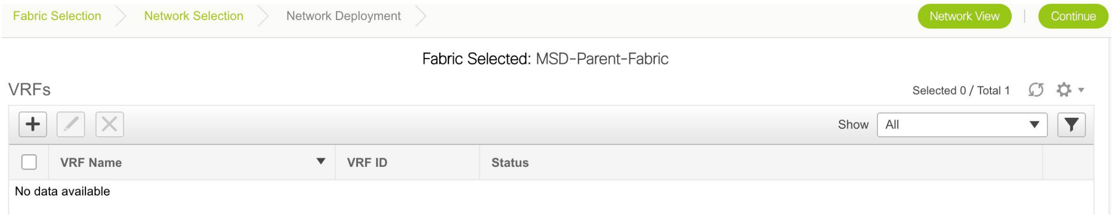

Fabric Selected: MSD-Parent-Fabric

□	VRF Name	VRF ID	Status
No data available			

2. Click the + button to add VRFs to the MSD fabric. The Create VRF screen comes up. Most of the fields are autopopulated.

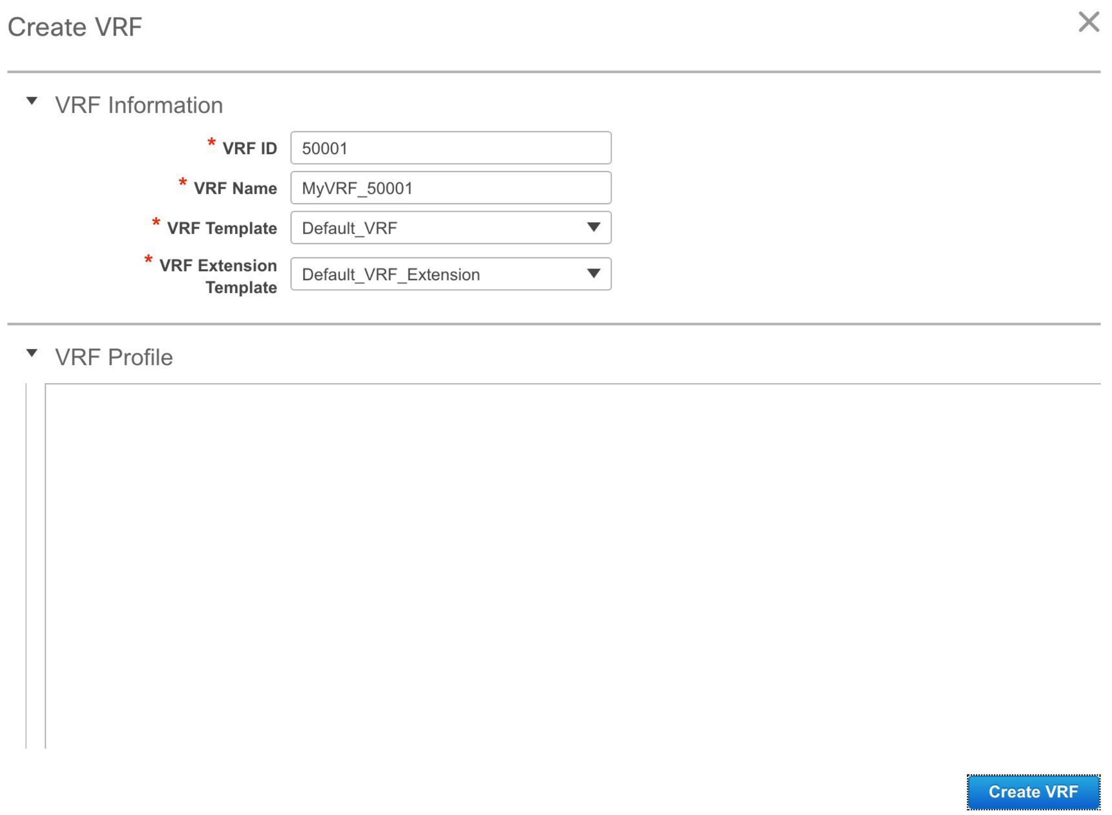

Create VRF

VRF Information

* VRF ID: 50001

* VRF Name: MyVRF_50001

* VRF Template: Default_VRF

* VRF Extension Template: Default_VRF_Extension

VRF Profile

The fields in this screen are:

VRF ID and **VRF Name** - The ID and name of the VRF.

The VRF ID is the VRF VNI or the L3 VNI of the tenant.

Note For ease of use, the VRF creation option is also available while you create a network.

VRF Template - This is populated with the *Default_VRF* template.

VRF Extension Template - This template allows you to extend the VRF between member fabrics.

3. Click **Create VRF**.

   The *MyVRF_50000* VRF is created and appears on the VRFs page.

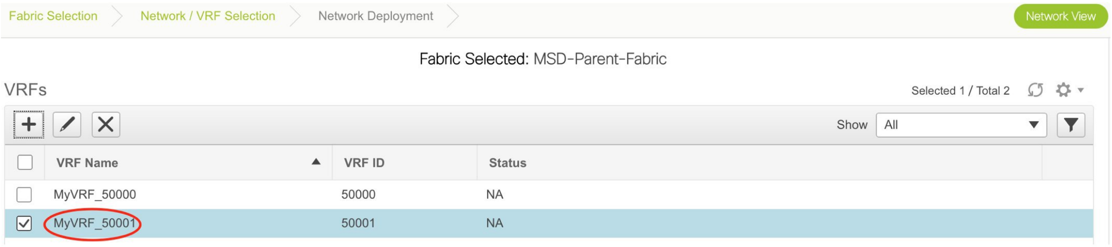

Fabric Selected: MSD-Parent-Fabric

VRFs

Selected 1 / Total 2

Show: All

☐	VRF Name	VRF ID	Status
☐	MyVRF_50000	50000	NA
☑	MyVRF_50001	50001	NA

Editing and Deleting VRFs in the MSD Fabric

To delete a VRF, use the delete (**X**) option at the top left part of the screen. You can delete multiple VRF instances by selecting them and clicking the delete button. You cannot edit VRF parameters after VRF creation.

A network or VRF deployed in a member fabric cannot be deleted until all instances are undeployed.

VRF Inheritance from MSD-Parent-Fabric to Member1

1. *MSD-Parent-Fabric* contains one member fabric, *Member1*. Go to the Select a Fabric page to access the Member1 fabric.

   [To go to the Select a Fabric page do one of the following:

   - Click the **Fabric Selection** button at the top left part of the screen.
   - From the main menu, click **Control > Networks & VRFs** and click **Continue** in the LAN Fabric Provisioning page].
   - Click *Member1* from the drop-down box.
   - Click **Continue** on the top right part of the screen. The Networks page comes up.
   - Click the **VRF View** button.

     On the VRFs page, you can see that the VRF created for the MSD is inherited to its member.

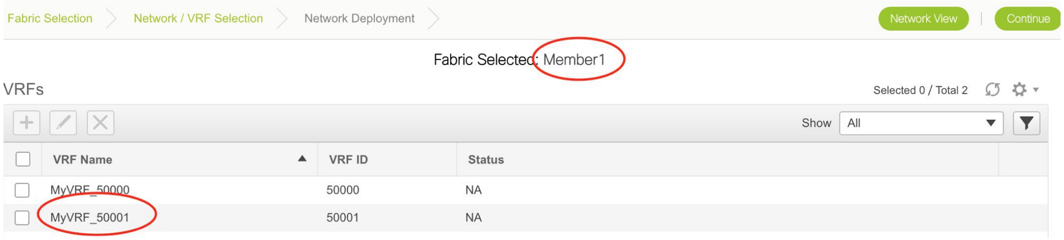

Fabric Selected: Member1

□	VRF Name	VRF ID	Status
□	MyVRF_50000	50000	NA
□	MyVRF_50001	50001	NA

Editing and Deleting VRFs in the Member Fabric

You cannot edit VRF parameters or delete a VRF at the member fabric level.

Step 1 of the following is explained. Step 2 information is mentioned in the next subsection.

1. Create networks and VRFs in the MSD fabric.
2. Deploy the networks and VRFs in the member fabric devices, one fabric at a time.

Deployment and Undeployment of Networks and VRFs in Member Fabrics

Before you begin, ensure that you have created networks at the MSD fabric level since the member fabric inherits networks and VRFs created for the MSD fabric.

Note The deployment (and undeployment) of networks and VRFs in member fabrics are the same as explained for standalone fabrics. Refer the standalone fabric documentation (*Networks Deployment* and *VRFs Deployment* sections in the *Networks and VRFs Creation and Deployment in a Standalone Fabric* topic).

Movement of a Standalone Fabric (With Existing Networks and VRFs) to an MSD Fabric

If you move a standalone fabric with existing networks and VRFs to an MSD fabric as a member, ensure that common networks (that is, L2 VNI and L3 VNI information), anycast gateway MAC, and VRF and network templates are the same across the fabric and the MSD. DCNM validates the standalone fabric (network and VRF information) against the (network and VRF information) of the MSD fabric to avoid duplicate entries. An example of duplicate entries is two common network names with a different network ID. After validation for any conflicts, the standalone fabric is moved to the MSD fabric as a member fabric. Details:

- The MSD fabric inherits the networks and VRFs of the standalone fabric that do not exist in the MSD fabric. These networks and VRFs are in turn inherited by the member fabrics.
- The newly created member fabric inherits the networks and VRFs of the MSD fabric (that do not exist in the newly created member fabric).
- If there are conflicts between the standalone and MSD fabrics, validation ensures that an error message is displayed. After the updation, when you move the member fabric to the MSD fabric, the move will be successful. A message comes up at the top of the page indicating that the move is successful.

If you move back a member fabric to standalone status, then the networks and VRFs remain as they are, but they remain relevant as in an independent fabric, outside the purview of an MSD fabric.

NFM Fabric Migration to a DCNM Fabric

NFM VXLAN fabric underlay and overlays can now be migrated and managed in DCNM 11.

Note DCNM 10.4(2) release only supported the NFM overlay migrations.

The migration involves processing the switch configurations and building the intent.

The two use cases involving NFM migration to DCNM are:

1. Migrate an NFM-managed VXLAN BGP EVPN fabric to DCNM 11. Here, the underlay and overlay networks are migrated.
2. Upgrade from DCNM 10.4(2) (or later) with NFM Overlay Migrations to DCNM 11. Here, the underlay is migrated.

Note This is only applicable to VXLAN BGP EVPN fabrics that were migrated from NFM to DCNM 10.4(2).

Both use cases are explained in this document.

Migrate an NFM-Managed VXLAN BGP EVPN Fabric to DCNM 11

The migration process involves creation of a new VXLAN BGP EVPN fabric through DCNM, adding switches to the fabric for underlay migration and migrating the VXLAN overlay networks from NFM to DCNM.

Prerequisites for NFM Fabric Migration to DCNM

- Install DCNM 11.0 release software. Refer the relevant Cisco DCNM Installation Guide for more details. Log in to DCNM and set the default LAN Credentials when prompted.
- Familiarity with the NFM configuration options and screen.

  (Go to **Switchpool > Settings > Edit**. Browse the **General** and **Underlay** tabs).
- Familiarity with the DCNM 11.0 fabric management and monitoring features before initiating the migration process.
- Familiarity with VXLAN BGP EVPN fabric concepts and functioning of the fabric from the DCNM perspective.
- Ensure that the NFM fabric switch nodes are operationally stable and functional:
  - All fabric links must be up.
  - vPC switches and the peer links must be up before the migration. Ensure that no configuration updates are in progress or pending changes from NFM.
- Create an inventory list of the switches in the fabric with their IP addresses and credentials. DCNM uses this information to connect to the switches.
- Open a console session to one of the leaf switches. The session is later used to collect some additional information directly from the switch.
- Shut down the Cisco NFM software so that it does not make any further configuration changes to the VXLAN fabric. Alternatively, disconnect the NFM network interfaces so that no changes are allowed on the switches.

Guidelines and Limitations

- Take a backup of the switch configurations and save them before the migration. These configurations can be used to restore the network if necessary.
- Before starting the process to migrate an NFM-managed VXLAN BGP EVPN fabric to DCNM 11, ensure that there are no configuration inconsistencies, such as inconsistencies in Switch Virtual Interfaces (SVI), VXLAN Network Identifiers (VNI), vPC port channels and so on, in the configurations applied to vPC pair devices.
- If the NFM-managed switch is not imported into DCNM due to an unknown username or password issue, log in to each switch and specify the username command using the plaintext password. This ensures that the SNMP credentials are set up correctly in NX-OS, and enables DCNM to discover the switch. For example, you can issue this CLI on the switch, where <*plaintext password*> is the placeholder for entering the plaintext password:

```
nfm-leaf: snmp-server user admin network-admin auth md5 <plaintext password>
```

- No configuration changes (unless instructed to do so in this document) must be made to the switches until the migration is completed. Else, significant network issues can occur.
- Cisco NFM to Cisco DCNM migration is only supported for Cisco Nexus 9000 switches.
- Before starting the process to migrate an NFM-managed VXLAN BGP EVPN fabric to DCNM 11, ensure that there are no configuration inconsistencies, such as inconsistencies in Switch Virtual Interfaces

(SVI), VXLAN Network Identifiers (VNI), vPC port channels and so on, in the configurations applied to vPC pair devices.

- Fabric point-to-point (P2P) port-channels (between leaf and spine switches) are supported in DCNM 11 only when the NFM fabric being migrated has them. When fabric port channel ports are present, the following guidelines are applicable:
    - Only a single fabric point-to-point port-channel must exist between a leaf switch and spine switch. Multiple fabric port-channels between a leaf switch and spine switch are not supported.
    - Adding or removing links between a leaf switch and spine switch updates the port channel membership automatically.
    - The fabric port channel is deleted when the last member is removed between a leaf switch and spine switch.
    - Adding links after the port channel is deleted makes them standalone point-to-point fabric interfaces.

Create a VXLAN BGP EVPN Fabric Through DCNM

A *fabric* defines a set of devices that makes up the physical fabric, their interconnectivity, configuration, and operational parameters.

1. Click **Control > Fabric Builder**.

   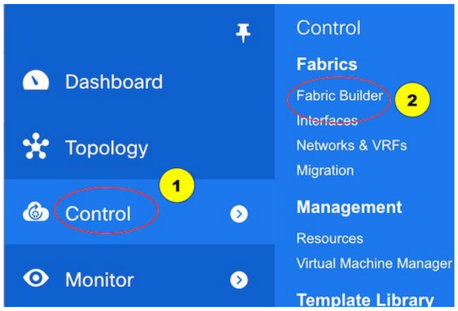

   The Fabric Builder page comes up.

2. Click the **Create Fabric** button. From the Add Fabric screen that comes up, select *NFM_Fabric* from the **Fabric Template** drop-down list.

Note The fabric requires several parameters to be set. Most of the parameters are prepopulated with default values. Carefully review each of the parameters and update them to match your specific fabric requirements.

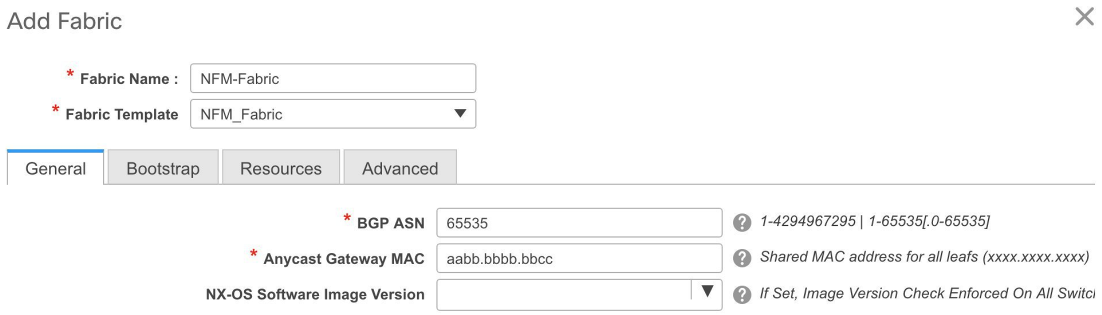
Add Fabric

* Fabric Name : NFM-Fabric

* Fabric Template NFM_Fabric

* BGP ASN 65535 — 1-4294967295 | 1-65535[.0-65535]

* Anycast Gateway MAC aabb.bbbb.bbcc — Shared MAC address for all leafs (xxxx.xxxx.xxxx)

NX-OS Software Image Version — If Set, Image Version Check Enforced On All Switc

General - The fields on this tab are specific to this fabric.

BGP ASN - Enter the BGP Autonomous System number of the fabric.

Anycast Gateway MAC - Enter the Anycast Gateway MAC address for the fabric.

Note The MAC address must be of the format *xxxx.xxxx.xxxx* (for example, ABCD.EF12:3456).

NX-OS Software Image Version - Select an image from the list.

If you upload Cisco NX-OS software images through the image upload option (**Control > Image Upload**), the uploaded images are listed in this field. If you select an image, the system checks if the switch has the selected version. If not, an error message is displayed. You can resolve the error by clicking on Resolve. The image management screen comes up and you can proceed with the ISSU option. Alternatively, you can delete the release number and save it later.

Bootstrap tab - The fields on this tab are specific to the DCHP settings for the fabric.

Click the **Enable DHCP** check box to initiate enabling of DHCP for automatic IP address assignment. When you click the check box, the other fields become editable.

Fill up the remaining fields for specifying a DHCP scope for allocating IP addresses to the device interfaces in the fabric. The fields are:

DHCP Scope Start Address and **DHCP Scope End Address** - The first and last IP addresses of the IP address range.

Switch Management Default Gateway and **Switch Management Subnet Prefix** - The management gateway IP address and the IP address subnet mask.

Note *DHCP scope and management gateway IP address specification* - If you specify the management gateway IP address 10.0.1.0 and subnet mask 24, ensure that the DHCP scope is within the specified subnet, between 10.0.1.1 and 10.0.1.254.

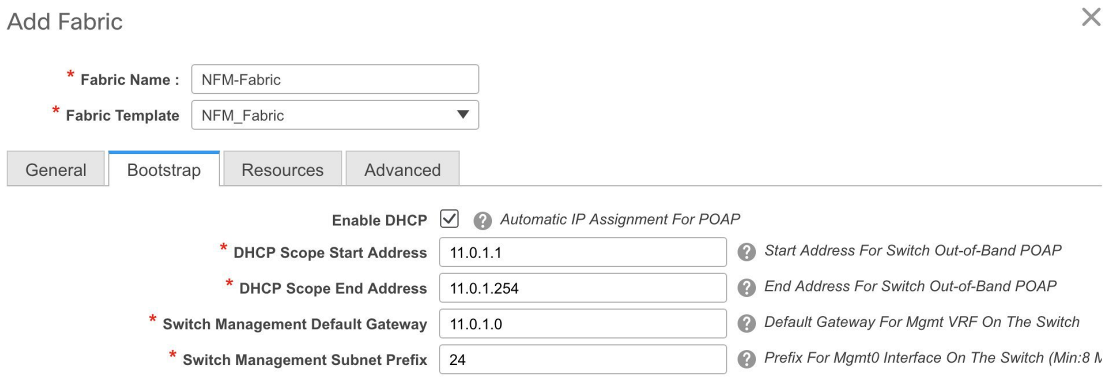

Add Fabric

* Fabric Name : NFM-Fabric

* Fabric Template NFM_Fabric

Enable DHCP ☑ Automatic IP Assignment For POAP

* DHCP Scope Start Address 11.0.1.1 Start Address For Switch Out-of-Band POAP

* DHCP Scope End Address 11.0.1.254 End Address For Switch Out-of-Band POAP

* Switch Management Default Gateway 11.0.1.0 Default Gateway For Mgmt VRF On The Switch

* Switch Management Subnet Prefix 24 Prefix For Mgmt0 Interface On The Switch (Min:8 M

Resources - This tab specifies the IP address, VXLAN VNI, VLAN, and subinterface ranges allocated for the fabric.

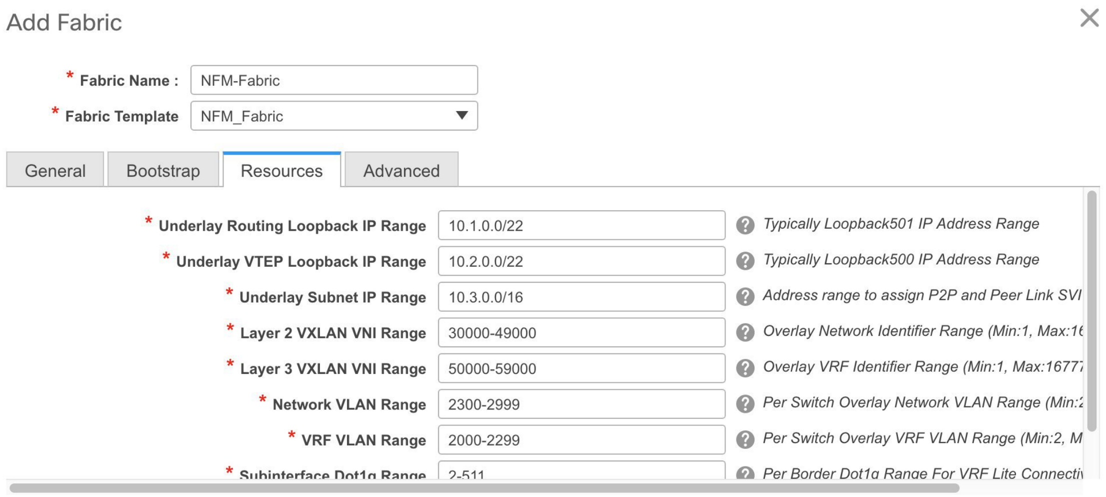

Add Fabric

* Fabric Name : NFM-Fabric

* Fabric Template NFM_Fabric

* Underlay Routing Loopback IP Range 10.1.0.0/22 Typically Loopback501 IP Address Range

* Underlay VTEP Loopback IP Range 10.2.0.0/22 Typically Loopback500 IP Address Range

* Underlay Subnet IP Range 10.3.0.0/16 Address range to assign P2P and Peer Link SVI

* Layer 2 VXLAN VNI Range 30000-49000 Overlay Network Identifier Range (Min:1, Max:16

* Layer 3 VXLAN VNI Range 50000-59000 Overlay VRF Identifier Range (Min:1, Max:16777

* Network VLAN Range 2300-2999 Per Switch Overlay Network VLAN Range (Min:2

* VRF VLAN Range 2000-2299 Per Switch Overlay VRF VLAN Range (Min:2, M

* Subinterface Dot1q Range 2-511 Per Border Dot1q Range For VRF Lite Connecti

Underlay Routing Loopback IP Range - Specifies loopback IP addresses for the protocol peering.

Underlay VTEP Loopback IP Range - Specifies loopback IP addresses for VTEPs.

Underlay Subnet IP Range - IP addresses for underlay P2P routing traffic between interfaces.

Note NFM uses a single IP underlay address pool. During the DCNM underlay migration, the IP addresses that are found on the switch are honored and retained. However, when any fresh IP address allocation is done after migration, the IP address is picked from the range that is specified here.

Layer 2 VXLAN VNI Range and **Layer 3 VXLAN VNI Range** - Specifies the VXLAN VNI IDs for the fabric.

Network VLAN Range and **VRF VLAN Range** - VLAN ranges for the Layer 3 VRF and overlay network.

Subinterface Dot1q Range - Specifies the subinterface range when L3 sub interfaces are used.

Note These values are defaults. If you want to update the IP address ranges, VXLAN Layer 2/Layer 3 network ID ranges or the VRF/Network VLAN ranges, ensure the following:

- If you update a range of values, ensure that it does not overlap with other ranges.
- Update one range of values (L2 Segment ID Range, for example) at a time. If you want to update more than one value, update a specific range, save the changes, and only then update another range of values.

Advanced tab.

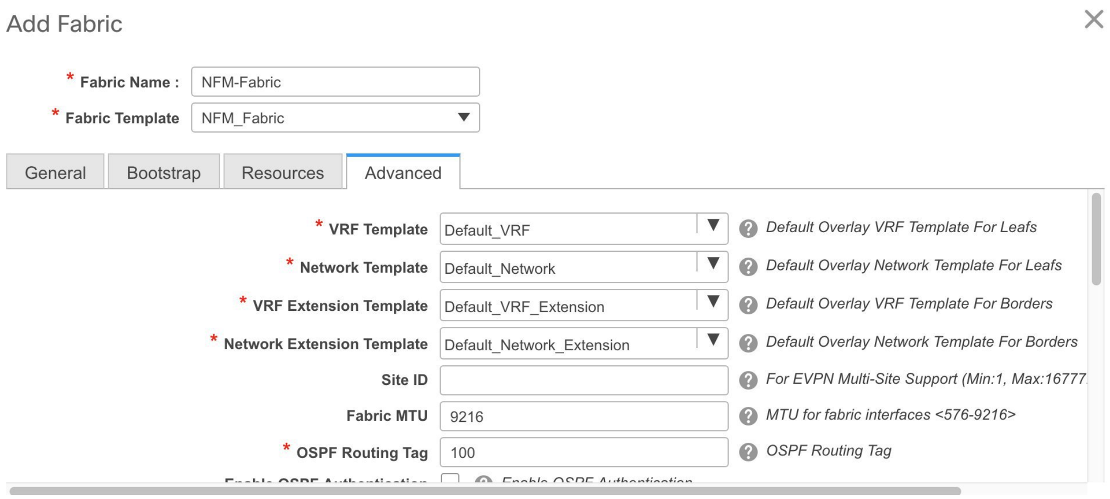

Add Fabric

* Fabric Name : NFM-Fabric

* Fabric Template NFM_Fabric

General | Bootstrap | Resources | Advanced

* VRF Template Default_VRF — Default Overlay VRF Template For Leafs

* Network Template Default_Network — Default Overlay Network Template For Leafs

* VRF Extension Template Default_VRF_Extension — Default Overlay VRF Template For Borders

* Network Extension Template Default_Network_Extension — Default Overlay Network Template For Borders

Site ID — For EVPN Multi-Site Support (Min:1, Max:16777

Fabric MTU 9216 — MTU for fabric interfaces <576-9216>

* OSPF Routing Tag 100 — OSPF Routing Tag

Enable OSPF Authentication □ Enable OSPF Authentication

The fields in this tab are:

VRF Template - Specifies the default VRF template for the overlay networks.

Network Template - Specifies the default Network template for the overlay networks.

VRF Extension Template - Specifies the default VRF extension template for extending the overlay networks to other fabrics.

Network Extension Template - Specifies the default Network extension template for extending the overlay networks to other fabrics.

Note NFM overlay migration supports *Default_Network* and *Default_VRF* templates only. Once the fabric has been successfully migrated into DCNM, any of the available templates can be used to deploy new overlay networks.

Site ID - The ID for this fabric if you are moving this fabric within an MSD.

The site ID is mandatory for a member fabric to be a part of an MSD. Each member fabric of an MSD has a unique site ID for identification.

Fabric MTU - Specifies the MTU for the fabric interfaces.

OSPF Routing Tag - Specifies the OSPF routing tag.

Enable OSPF Authentication - Select the check box to enable OSPF authentication. Deselect the check box to disable it.

If you enable this field, the OSPF Authentication Key ID and OSPF Authentication Key fields get enabled.

OSPF Authentication Key ID and **OSPF Authentication Key**.

Note The OSPF authentication key must be the 3DES key from the switch. Collect the key ID and the key from one of the leafs.

```
nfm-leaf# terminal width 300
nfm-leaf# show run ospf | grep message-digest-key
  ip ospf message-digest-key 127 md5 3 c7c83ec78f38f32f3d477519630faf7b
```

Enable BGP Authentication - Select the check box to enable BGP authentication. Deselect the check box to disable it. If you enable this field, the BGP Authentication Key field gets enabled.

BGP Authentication Key - Enter the 3DES key that is collected from the switch.

```
nfm-leaf# terminal width 300
nfm-leaf# show run bgp | grep password
password 3
9e39aa786319a7da1cd23e7dd933e80533b04208805b64077185ecebbcadaa25d791a1d353081e03
```

vPC Peer Link VLAN - VLAN used for the vPC peer link SVI.

For a vPC switch peer link SVI (vlan3966), you must configure these commands manually on each vPC switch. The import fails if you do not configure any of these CLIs or enable additional commands.

```
interface Vlan3966
  no shutdown
  bfd interval 50 min_rx 50 multiplier 3
  no bfd echo
  no ip redirects
  ip address 172.28.254.30/31
  no ipv6 redirects
  ip router ospf 1 area 0.0.0.0
  ip ospf bfd
```

vPC Delay Restore Time - Specifies the vPC delay restore period in seconds.

vPC Auto Recovery Time - Specifies the vPC auto recovery time-out period in seconds.

Power Supply Mode - Choose the appropriate power supply mode.

CoPP Profile - Choose the appropriate Control Plane Policing (CoPP) profile policy for the fabric.

Enable VXLAN OAM - Enables the VXLAM OAM function.

Note The VXLAN OAM feature in Cisco DCNM is only supported on a single fabric or site.

Enable vPC Advertise PIP - Enables the Advertise PIP feature.

Freeform CLIs - Fabric level freeform CLIs (such as AAA server parameters) can be added while creating or editing a fabric. They are applicable to switches across the fabric. You should add the configurations as displayed in the running configuration, without indentation. Switch level freeform configurations such as VLAN, SVI, and interface configurations should only be added on the switch.

Leaf Freeform Config - Add CLIs that should be added to switches that have the *Leaf, Border* and *Border Gateway* roles.

Spine Freeform Config - Add CLIs that should be added to switches with a *Spine* role.

3. Click **Save** after filing and updating relevant information.

Fabric Underlay Migration

The fabric is placed in a special *migration mode* when it is created. Several configuration restrictions are in place while the fabric is in this mode. Please ensure the following in this mode:

- Do not add or edit or delete an interface from the **Control > Interfaces** page.
- Do not update switch configurations through the Save & Deploy option (which appears at the top right part of the fabric page).
- Do not add a new switch (a switch that is not a part of the existing NFM fabric being migrated) through the Add switches or Bootstrap options.

The fabric is automatically taken out of the migration mode when both the underlay and overlay migrations are completed successfully.

Read the following guidelines and then refer the **Discovering existing switches** section in the *Add switches to the fabric* topic for detailed migration steps.

- In the fabric page, use the Add switches option in the Actions panel to add switches to the DCNM-managed fabric.
- When adding a switch, set **Preserve Switch Configuration** to *Yes*.

  Use the *No* setting to add *new* switches after the underlay and migration is complete.

  Note Adding switches with **Preserve Switch Configuration** set to *No* while the fabric is still in the migration state is not supported. Doing so reports an error and the switch is not added to the fabric without making any changes to the switch.

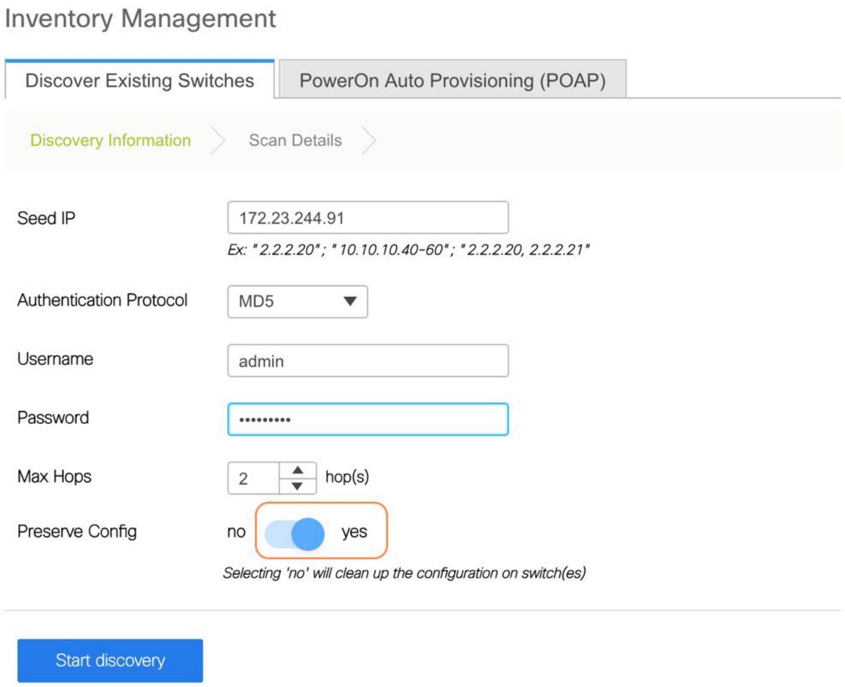
Inventory Management

Seed IP: 172.23.244.91

Ex: "2.2.2.20"; "10.10.10.40-60"; "2.2.2.20, 2.2.2.21"

Authentication Protocol: MD5

Username: admin

Password: •••••••••

Max Hops: 2 hop(s)

Preserve Config: no [✓ yes]

Selecting 'no' will clean up the configuration on switch(es)

- Click the **Start discovery** button and then select the set of switches to be imported from the Inventory Management page that shows up. A progress bar indicates the underlay migration status for each of the switches.

Note You should not close the Inventory Management page while there are active migrations.

- The migration workflow will analyse the configurations and the switch is added to the fabric after it passes a set of acceptance criteria. Errors and warnings are reported in the fabric **Pending Error** as appropriate.

Note Each switch has a *migration mode* to track the completion of its underlay migration. A switch in this mode is shown with a special **Migration Mode** tag in the topology view.

It is normal for a switch to be shown with the tag if an error is detected that prevents the underlay migration to complete. The error message will provide information on the nature of the error and suggested remedial action.

- Ensure that you add all the NFM fabric devices to the DCNM fabric to complete the underlay migration process. After the underlay networks' migration is complete, the topology is updated in the fabric page.
- Ensure that the interfaces in the **Control > Interfaces** screen show the correct policies and associated configurations.
- You can now proceed to completing the overlay migrations.

Fabric Overlay Migration

The Migration wizard will help you migrate over the NFM Overlay networks (or *broadcast domains* as known in the NFM). The migration has two phases, *Discovery* and *Migration*.

The Discovery phase is where the configurations that are on the switches are parsed and presented in the GUI for review. The networks, interfaces, and switches where the networks exist are displayed. Once you verify the information to be accurate, you can move to the Migration phase by selecting the networks and proceeding to deploy those networks. The GUI workflow tracks the status of the migrations for audit purposes. The migration is considered completed when all the networks are migrated.

Note It is important that no configuration or network changes are made to the switches until the migration is completed. Any out-of-band configuration changes can interfere with the migrations and can cause significant network issues.

It is important that you verify the discovered networks and data before you initiate a migration. Once the first network is migrated (Migration phase) it is not possible to go back to the Discovery phase to make changes.

Cisco NFM supports single fabric, whereas Cisco DCNM supports multiple fabrics, so the original NFM-deployed fabric becomes one fabric among all the Cisco DCNM-managed fabrics.

Note DCNM 11 currently allows only one active overlay migration to be in progress at a time.

Each overlay network migration consists of the following steps:

1. Preparing the switch for migration to DCNM Top-Down managed networks.
2. Preparing the Layer 3 network on the switch for migration to DCNM Top-Down managed networks.
3. Deploying the DCNM Top-Down networks configuration to the switch.
4. Removing the original configuration that existed on the switch before the deployment.

Follow these steps to migrate the NFM fabric overlay (networks, VRFs and other overlay parameters) to the DCNM fabric.

1. Click **Control > Migration**. The Select a Fabric page comes up. The newly created VXLAN fabric appears in the **Select a Fabric** drop down box.
2. Select the fabric and click **Continue** on the top right part of the screen.

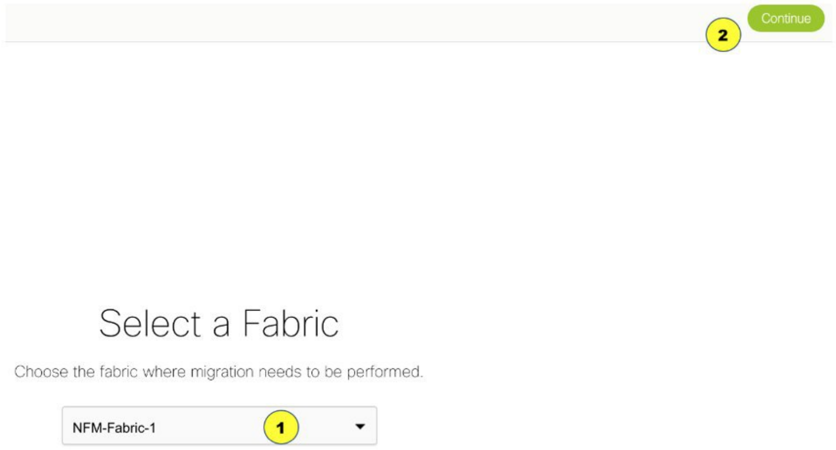

Select a Fabric

Choose the fabric where migration needs to be performed.

NFM-Fabric-1

The NFM fabric migration page comes up.

To start with, the DISCOVERY IN PROGRESS message appears at the top of the Migration screen. The discovery process auto-generates the network name of the form as *Auto_Net_VLANxxx_VNIyyyyy*.

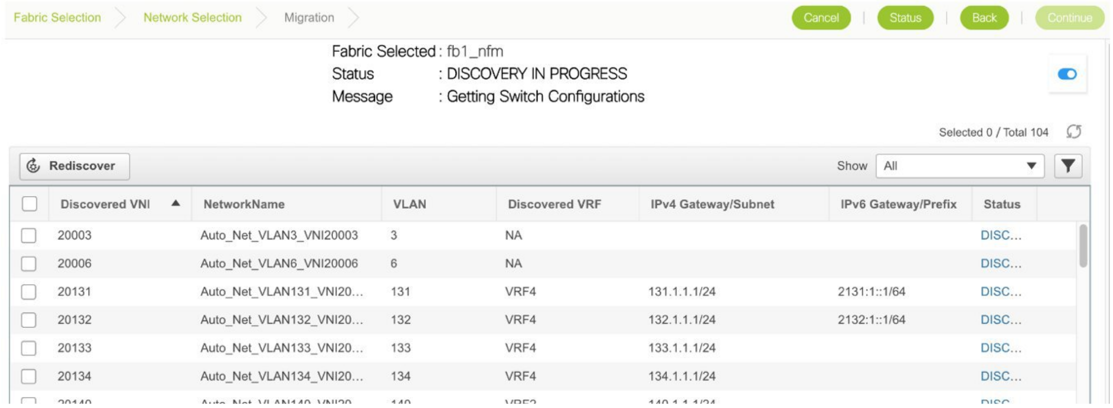

Fabric Selected : fb1_nfm
Status : DISCOVERY IN PROGRESS
Message : Getting Switch Configurations

Selected 0 / Total 104

□	Discovered VNI ▲	NetworkName	VLAN	Discovered VRF	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status
□	20003	Auto_Net_VLAN3_VNI20003	3	NA			DISC...
□	20006	Auto_Net_VLAN6_VNI20006	6	NA			DISC...
□	20131	Auto_Net_VLAN131_VNI20...	131	VRF4	131.1.1.1/24	2131:1::1/64	DISC...
□	20132	Auto_Net_VLAN132_VNI20...	132	VRF4	132.1.1.1/24	2132:1::1/64	DISC...
□	20133	Auto_Net_VLAN133_VNI20...	133	VRF4	133.1.1.1/24		DISC...
□	20134	Auto_Net_VLAN134_VNI20...	134	VRF4	134.1.1.1/24		DISC...
□	20140	Auto_Net_VLAN140_VNI20...	140	VRF3	140.1.1.1/24		DISC...

Cisco DCNM will retrieve the running configuration from the switches, parse the configurations to discover the VXLAN overlay data. At this point, the migration is considered to be in progress.

The parsing occurs in the background and the page refreshed with the discovered networks. You cannot proceed further until the discovery process is completed. The Continue button and the check boxes are disabled while discovery is in progress. The discovered networks are persisted until one of the following events occurs:

- Migration is completed (network is deployed and the original configuration CLIs are removed).
- Until you click the **Rediscover** button upon which the current list is discarded and configuration is parsed again. The **Rediscover** button will throw an error once the migration status is changed to MIGRATION IN PROGRESS. The only time a Rediscovery can be performed is when the status is DISCOVERY COMPLETED. The other states where the Rediscover can be triggered are DISCOVERY FAILED and DISCOVERY ABORTED.
- Until you cancel the migration.

After the discovery process is complete, the DISCOVERY COMPLETED message appears at the top of the screen.

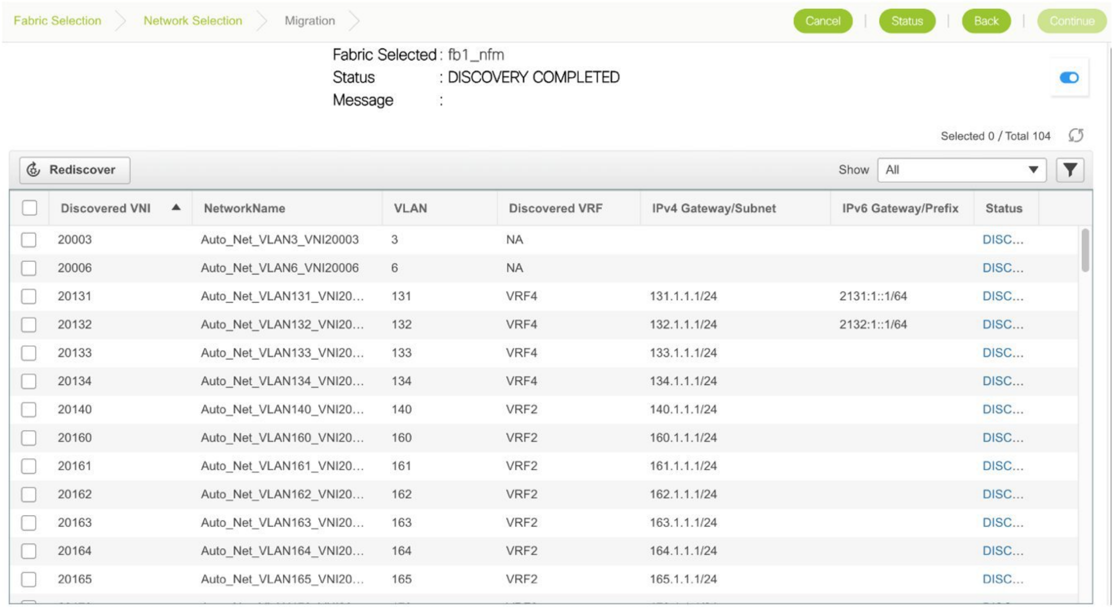

Fabric Selected : fb1_nfm
Status : DISCOVERY COMPLETED
Message :

Selected 0 / Total 104

□	Discovered VNI	NetworkName	VLAN	Discovered VRF	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status
□	20003	Auto_Net_VLAN3_VNI20003	3	NA			DISC...
□	20006	Auto_Net_VLAN6_VNI20006	6	NA			DISC...
□	20131	Auto_Net_VLAN131_VNI20...	131	VRF4	131.1.1.1/24	2131:1::1/64	DISC...
□	20132	Auto_Net_VLAN132_VNI20...	132	VRF4	132.1.1.1/24	2132:1::1/64	DISC...
□	20133	Auto_Net_VLAN133_VNI20...	133	VRF4	133.1.1.1/24		DISC...
□	20134	Auto_Net_VLAN134_VNI20...	134	VRF4	134.1.1.1/24		DISC...
□	20140	Auto_Net_VLAN140_VNI20...	140	VRF2	140.1.1.1/24		DISC...
□	20160	Auto_Net_VLAN160_VNI20...	160	VRF2	160.1.1.1/24		DISC...
□	20161	Auto_Net_VLAN161_VNI20...	161	VRF2	161.1.1.1/24		DISC...
□	20162	Auto_Net_VLAN162_VNI20...	162	VRF2	162.1.1.1/24		DISC...
□	20163	Auto_Net_VLAN163_VNI20...	163	VRF2	163.1.1.1/24		DISC...
□	20164	Auto_Net_VLAN164_VNI20...	164	VRF2	164.1.1.1/24		DISC...
□	20165	Auto_Net_VLAN165_VNI20...	165	VRF2	165.1.1.1/24		DISC...

Note It is important that the discovered networks and data is verified before a migration is attempted. Make necessary changes and click **Rediscover** to restart the discovery process.

At any point in time, click **Cancel** to cancel the discovery process that is in progress and click the **Status** button to view the status.

3. After the discovery process is completed, select the networks that you want to migrate to the DCNM fabric.

4. Click **Continue** at the top right part of the screen. The page that appears next has some additional options that allow you to preview existing configurations and the configurations that are going to be deployed on the switches.

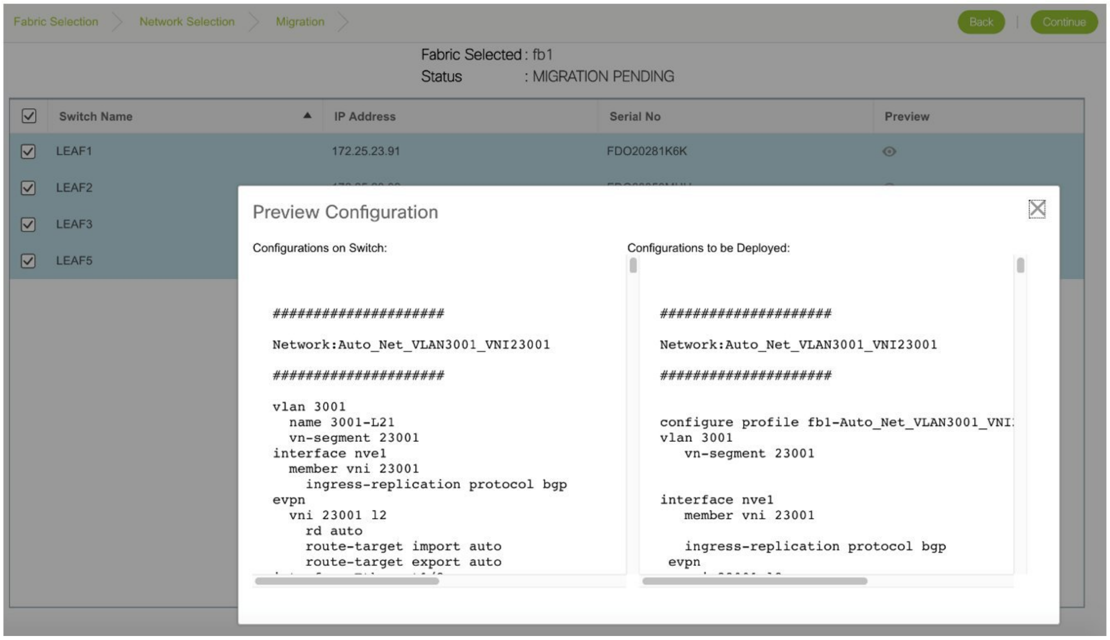

Fabric Selected : fb1
Status : MIGRATION PENDING

☑	Switch Name	IP Address	Serial No	Preview
☑	LEAF1	172.25.23.91	FDO20281K6K	👁
☑	LEAF2			
☑	LEAF3			
☑	LEAF5			

Preview Configuration

Configurations on Switch:

```
####################

Network:Auto_Net_VLAN3001_VNI23001

####################

vlan 3001
  name 3001-L21
  vn-segment 23001
interface nve1
  member vni 23001
    ingress-replication protocol bgp
evpn
  vni 23001 l2
    rd auto
    route-target import auto
    route-target export auto
```

Configurations to be Deployed:

```
####################

Network:Auto_Net_VLAN3001_VNI23001

####################


configure profile fb1-Auto_Net_VLAN3001_VNI
vlan 3001
   vn-segment 23001


interface nve1
   member vni 23001

   ingress-replication protocol bgp
 evpn
```

You can select the switch(es) where the networks needs to be migrated. It is however recommended to select all the switches for the migration. If only a subset of switches is selected, ensure that both the switches in the vPC pair are present.

After the overlay network migration is completed, a message MIGRATION COMPLETED is displayed at the top of the screen.

The fabric is moved out of the migration mode and the complete DCNM 11 fabric management functions are enabled.

Viewing Overlay Migration Status

In the Migration page, click the **Status** button. The page that appears reports the cumulative status of all migrations performed so far.

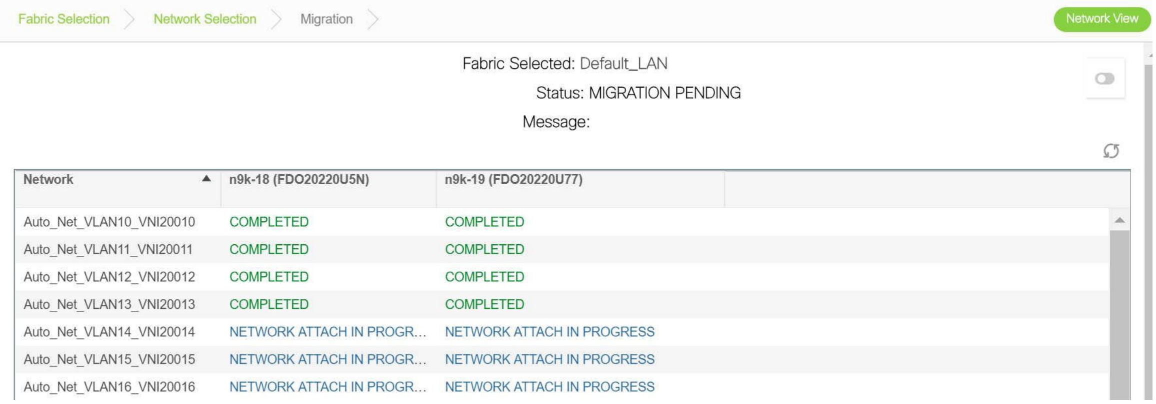

Fabric Selected: Default_LAN
Status: MIGRATION PENDING
Message:

Network	n9k-18 (FDO20220U5N)	n9k-19 (FDO20220U77)
Auto_Net_VLAN10_VNI20010	COMPLETED	COMPLETED
Auto_Net_VLAN11_VNI20011	COMPLETED	COMPLETED
Auto_Net_VLAN12_VNI20012	COMPLETED	COMPLETED
Auto_Net_VLAN13_VNI20013	COMPLETED	COMPLETED
Auto_Net_VLAN14_VNI20014	NETWORK ATTACH IN PROGR...	NETWORK ATTACH IN PROGRESS
Auto_Net_VLAN15_VNI20015	NETWORK ATTACH IN PROGR...	NETWORK ATTACH IN PROGRESS
Auto_Net_VLAN16_VNI20016	NETWORK ATTACH IN PROGR...	NETWORK ATTACH IN PROGRESS

You can click the hyperlinks to view migration history and status.

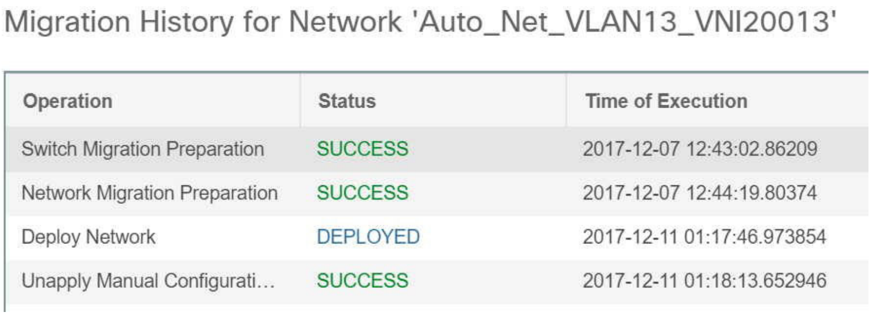

Migration History for Network 'Auto_Net_VLAN13_VNI20013'

Operation	Status	Time of Execution
Switch Migration Preparation	SUCCESS	2017-12-07 12:43:02.86209
Network Migration Preparation	SUCCESS	2017-12-07 12:44:19.80374
Deploy Network	DEPLOYED	2017-12-11 01:17:46.973854
Unapply Manual Configurati…	SUCCESS	2017-12-11 01:18:13.652946

Troubleshooting Cisco NFM to Cisco DCNM Migration

The Migration workflow involves multiple steps and some unexpected issues that you might encounter while migrating Cisco NFM to Cisco DCNM. Fabric underlay and overlay examples:

Fabric Underlay Troubleshooting

Errors and warnings reported during the underlay operations are reported in the fabric *Pending Errors*, at the top right part of the screen.

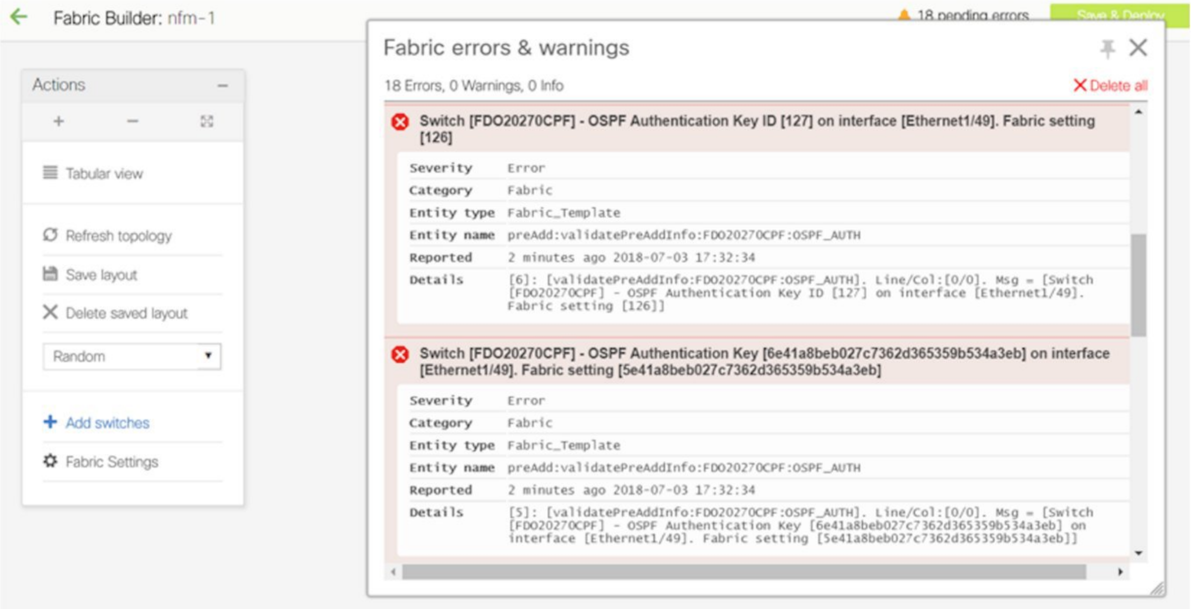

Fabric errors & warnings

18 Errors, 0 Warnings, 0 Info

Switch [FDO20270CPF] - OSPF Authentication Key ID [127] on interface [Ethernet1/49]. Fabric setting [126]

```
Severity     Error
Category     Fabric
Entity type  Fabric_Template
Entity name  preAdd:validatePreAddInfo:FDO20270CPF:OSPF_AUTH
Reported     2 minutes ago 2018-07-03 17:32:34
Details      [6]: [validatePreAddInfo:FDO20270CPF:OSPF_AUTH]. Line/Col:[0/0]. Msg = [Switch
             [FDO20270CPF] - OSPF Authentication Key ID [127] on interface [Ethernet1/49].
             Fabric setting [126]]
```

Switch [FDO20270CPF] - OSPF Authentication Key [6e41a8beb027c7362d365359b534a3eb] on interface [Ethernet1/49]. Fabric setting [5e41a8beb027c7362d365359b534a3eb]

```
Severity     Error
Category     Fabric
Entity type  Fabric_Template
Entity name  preAdd:validatePreAddInfo:FDO20270CPF:OSPF_AUTH
Reported     2 minutes ago 2018-07-03 17:32:34
Details      [5]: [validatePreAddInfo:FDO20270CPF:OSPF_AUTH]. Line/Col:[0/0]. Msg = [Switch
             [FDO20270CPF] - OSPF Authentication Key [6e41a8beb027c7362d365359b534a3eb] on
             interface [Ethernet1/49]. Fabric setting [5e41a8beb027c7362d365359b534a3eb]]
```

Fabric Overlay Troubleshooting

An issue encountered during the overlay migration will fail the process with an appropriate FAILED status and the Message field will indicate the failure.

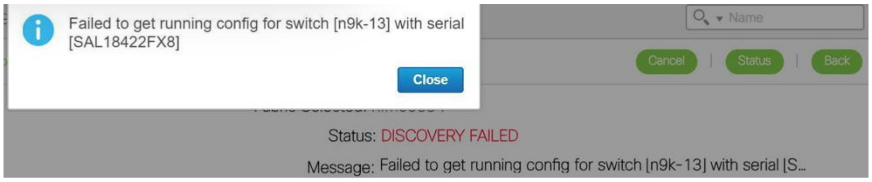

Failed to get running config for switch [n9k-13] with serial [SAL18422FX8]

Status: DISCOVERY FAILED

Message: Failed to get running config for switch [n9k-13] with serial [S...

Network Migration Failures

Go to the migration page, identify the network and switch that has encountered the failure and click the **Status** hyperlink. The resulting popup shows the status of each migration step.

Further details can be obtained by clicking the appropriate hyperlinks and additional details can be obtained by reviewing the log files.

Migration Workflow Failures

The migration status will indicate a FAILURE. Additional details can be obtained by reviewing the log files.

Migration Workflow Status Definitions

This section describes the various states for the discovery or migration workflow:

Discovery-related Status Definitions

DISCOVERY INITIATED - A discovery has been triggered and waiting to start.

DISCOVERY IN PROGRESS - The discovery is active.

DISCOVERY FAILED - The previous discovery failed.

DISCOVERY ABORT INITIATED - An attempt to cancel an active discovery has been initiated.

DISCOVERY ABORTED - The previous discovery has been canceled.

DISCOVERY COMPLETED - The discovery has been completed successfully.

Migration-related Status Definitions

MIGRATION INITIATED - Migration has been initiated for a set of networks.

MIGRATION IN PROGRESS - Migration is in progress for a set of networks.

MIGRATION FAILED - The previous migration failed.

MIGRATION ABORT INITIATED - An attempt to cancel an active migration has been initiated.

MIGRATION ABORTED - Migration has been canceled.

MIGRATION PENDING - There are more networks waiting to be migrated.

MIGRATION COMPLETED - All the networks have been migrated.

Network Migration Status Definitions

DISCOVERED - The network has been discovered from the switch configurations.

SWITCH MIGRATION PREPARATION IN PROGRESS - The switch where the network is present is being prepared.

SWITCH MIGRATION PREPARATION FAILED - The switch preparation step failed.

NETWORK MIGRATION PREPARATION IN PROGRESS - The L3 network is being prepared for migration.

NETWORK MIGRATION PREPARATION FAILED - The L3 network preparation step failed.

NETWORK CREATION IN PROGRESS - The LAN Fabric Provisioning Network entry is being created.

NETWORK CREATION FAILED - The LAN Fabric Provisioning Network entry creation failed.

NETWORK DEPLOYMENT IN PROGRESS - The LAN Fabric Provisioning Network deployment is in progress.

NETWORK DEPLOYMENT FAILED - The LAN Fabric Provisioning Network deployment failed.

ORIGINAL CONFIGURATION REMOVAL PENDING - The LAN Fabric Provisioning Network deployment is successful and waiting to remove the original NFM configured CLIs.

ORIGINAL CONFIGURATION REMOVAL IN PROGRESS - The removal of the original NFM configured CLIs is in progress.

ORIGINAL CONFIGURATION REMOVAL RECOVERABLE FAILURE - The removal of the original NFM configured CLIs failed, but, can be retried on a future attempt after fixing any underlying issues.

ORIGINAL CONFIGURATION REMOVAL FAILED - The removal of the original NFM configured CLIs failed. The failure reason must be reviewed and manual corrective action must be taken. Please review the nature of the failure(s). If some of the configuration CLIs were partially applied, please reapply the failed and rest of the CLIs manually on the switch(es).

COMPLETED - The network was migrated successfully.

Network Migration History Definitions

Switch Migration Preparation - Provides status of preparing the switch for the migration. This action is performed only once per switch, but, will show up in all network histories.

Network Migration Preparation - Provides status of the network migration preparations. This entry is only present for L3 networks.

Deploy Network - Provides status of the LAN Fabric Network provisioning.

Unapply Manual Configurations - Provides status of removing the network overlay CLIs configured by NFM. Note that this does not lead to any loss of configuration since LAN Fabric Provisioning uses configuration profiles.

Upgrade from DCNM 10.4(2) with NFM Overlay Migrations to DCNM 11

Note The explanation is only applicable to VXLAN fabrics that were migrated from NFM to DCNM 10.4(2) or later.

1. Follow the recommended DCNM upgrade procedure and upgrade to DCNM 11.
2. After DCNM is reachable, click **Control > Fabric Builder**.

   The fabrics are listed in a distinct color.
3. Identify the NFM fabric and click **Edit**. Select the NFM fabric from the **Fabric Template** drop-down box. Many of the fabric settings have default values. Review all the settings to make sure that they match your fabric. Refer to the *Create a VXLAN BGP EVPN Fabric Through DCNM* section for information on the fabric settings.
4. Click **Save** at the bottom right part of the screen. All the switches are displayed with the **Migration Mode** tag.
5. Click **Save & Deploy** to complete the migration of the underlay networks .
6. The overlay networks do not need any additional migration action.

Post Migration Operations

After completing the underlay and overlay migrations, follow these steps:

1. Navigate to **Control > Fabric Builder** in the DCNM GUI. On the page that comes up, click the fabric. The fabric topology page comes up.
2. Click **Save & Deploy**. This step implements the DNCM 11 VXLAN BGP EVPN fabric best practice of deploying all pending configurations on the fabric switches.

Note Review the configuration differences that show up, for accuracy, before deploying them to the switch.

Now the fabric is ready for use.

Updating Switch Level Settings

A few switch level settings can be updated using this procedure:

1. Navigate to **Control > Fabric Builder** and select the fabric.
2. Right click the switch to update its settings, and click the **View/Edit Policies** option and do the following:

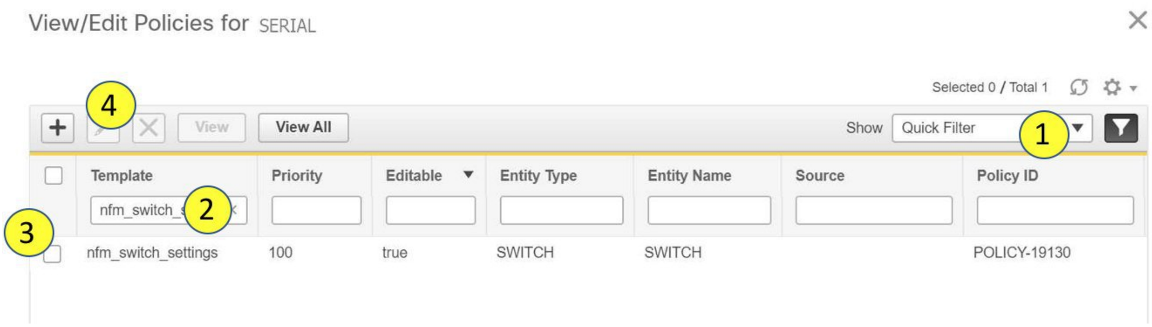

View/Edit Policies for SERIAL

Selected 0 / Total 1

Template	Priority	Editable	Entity Type	Entity Name	Source	Policy ID
nfm_switch_s						
nfm_switch_settings	100	true	SWITCH	SWITCH		POLICY-19130

   1. Enable the filtering option (at the top right part of the screen) and enter *nfm_switch_settings* in the **Template** field.
   2. Select the *nfm_switch_settings* policy and click Edit. The Edit Policy screen comes up.

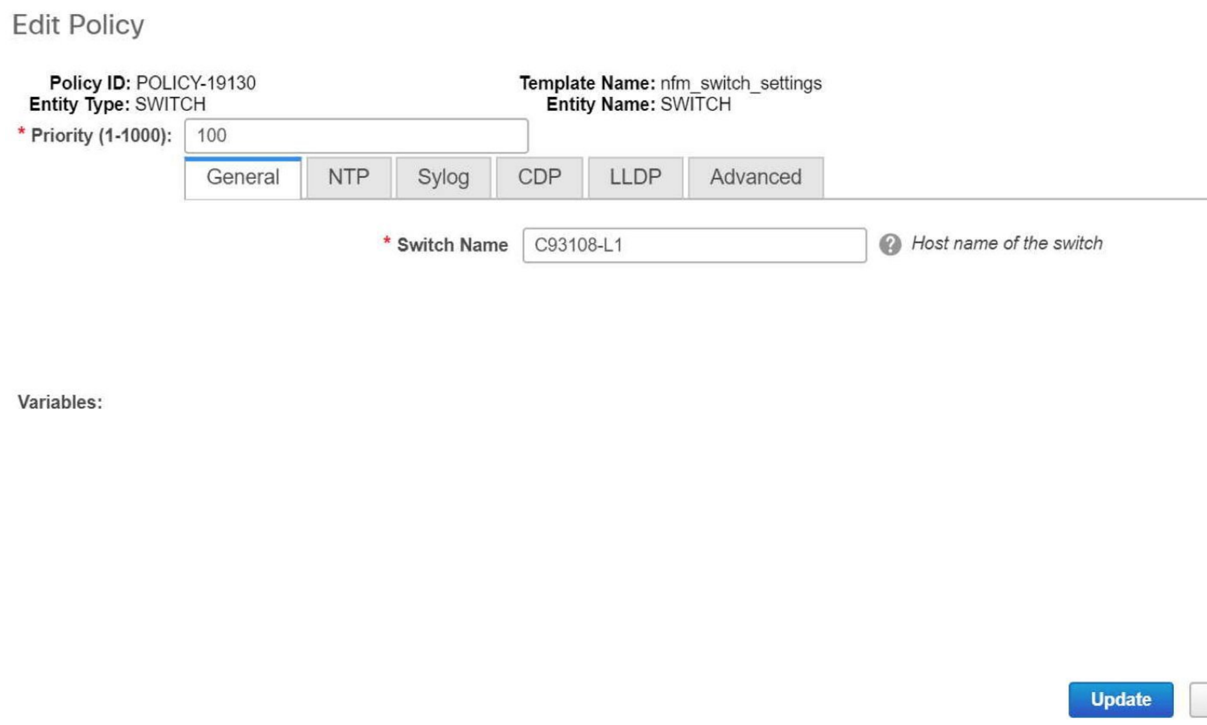

Edit Policy

Policy ID: POLICY-19130
Entity Type: SWITCH
Template Name: nfm_switch_settings
Entity Name: SWITCH
* Priority (1-1000): 100

* Switch Name: C93108-L1 — Host name of the switch

Variables:

3. Make changes and click **Update** to update the settings.

4. A **Save & Deploy** pushes these configuration changes to the switch.

Updating Fabric OSPF Authentication Parameters

Disabling OSPF Authentication

1. Navigate to **Control > Fabric Builder** and click the settings icon of the fabric. The Edit Fabric screen comes up.

2. Click the **Advanced** tab and deselect the **Enable OSPF Authentication** check box.

3. Click **Save**.

4. A **Save & Deploy** pushes these configuration changes to the switch.

Note The task can cause traffic disruption.

Enabling or Updating OSPF Authentication

1. Log in to one of the leaf switches in the fabric and collect the following information:

```
nfm-leaf(config)# interface loopback 999 [Pick a non-existent loopback id]

nfm-leaf(config-if)# ip ospf message-digest-key 127 md5 testPassword [Use the desired
key ID and password]


nfm-leaf(config-if)# show run interface lo999
interface loopback999
  ip ospf message-digest-key 127 md5 3 1afc85c3227850739fff5d727ad413f6


nfm-leaf(config-if)# no interface lo999 [delete the temporary loopback interface created
 earlier]
```

2. Navigate to **Control > Fabric Builder** and click the settings icon of the fabric. The Edit Fabric screen comes up.
3. Click the **Advanced** tab and select the **Enable OSPF Authentication** check box if not already selected.
4. From the information that is collected earlier, enter the key ID into the **OSPF Authentication Key ID** field and the 3DES key as-is into the **OSPF Authentication Key** field.
5. Click **Save**.
6. A Save & Deploy pushes these configuration changes to the switch.

Note The task can cause traffic disruption.

Updating Fabric BGP Authentication Parameters

Disabling BGP Authentication

1. Navigate to **Control > Fabric Builder** and click the settings icon of the fabric. The Edit Fabric screen comes up.
2. Click the **Advanced** tab and deselect the **Enable BGP Authentication** check box.
3. Click **Save**.
4. A **Save & Deploy** pushes these configuration changes to the switch.

Note The task can cause traffic disruption.

Enabling or Updating BGP Authentication

1. Log in to one of the leaf switches in the fabric and collect the following information:

```
nfm-leaf# conf t
nfm-leaf(config)# router bgp <bgp as #>      [BGP AS Number]
nfm-leaf(config-router)# neighbor 1.1.1.1    [A non existent BGP neighbor ID]

nfm-leaf(config-router-neighbor)# password testPassword [desired password in cleartext]



nfm-leaf(config-router-neighbor)# show run bgp
[snip]
router bgp <bgp as #>
[snip]
```

```
     neighbor 1.1.1.1
       password 3 f092f5f76d298504ca9b1ad0f1469ca8


   nfm-leaf(config-router-neighbor)# exit
   nfm-leaf(config-router)# no neighbor 1.1.1.1   [delete the neighbor created earlier
   ]
```

2. Navigate to **Control > Fabric Builder** and click the settings icon of the fabric. The Edit Fabric screen comes up.
3. Click the **Advanced** tab and select the **Enable BGP Authentication** check box, if already not selected.
4. From the information that is collected earlier, enter the highlighted 3DES key as-is into the **BGP Authentication Key** field.
5. Click **Save**.
6. A **Save & Deploy** pushes these configuration changes to the switch.

Note The task can cause traffic disruption.

Freeform Configurations on Fabric Switches

In DCNM, you can add custom configurations through freeform policies in the following ways:

1. Fabric-wide
    - On all leaf and border switches in the fabric, at once.
    - On all spine switches, at once.
2. On a specific switch.

Leaf switches are identified by the role *Leaf*, border switches by the role *Border* or *Border-Gateway* and spine switches by the role *Spine*.

Note You can deploy freeform CLIs when you create a fabric or when a fabric is already created. The following examples are for an existing fabric. However, you can use them as a reference for a new fabric.

Deploy Fabric-Wide Freeform CLIs on Leaf and Spine Switches

1. Click **Control > Fabric Builder**. The Fabric Builder screen comes up. A rectangular box represents each fabric.
2. Click the **Settings** icon (located on the top right part of the rectangular box) for adding custom configurations to an existing fabric. The **Edit Fabric** screen comes up.

    (If you are creating a fabric for the first time, click **Create Fabric**).
3. Click the **Advanced** tab and update the following fields:

Leaf Freeform Config – In this field, add configurations for all leaf and border switches in the fabric. For example, you can add NTP, TACAS, and AAA configurations in this field.

Don’t add VLAN, SVI, and interface-specific configurations.

Spine Freeform Config - In this field, add configurations for all spine switches in the fabric.

Note Copy-paste the intended configuration with correct indentation, as seen in the running configuration on the Nexus switches. For more information, see Resolving Freeform Config Errors in Switches, on page 113.

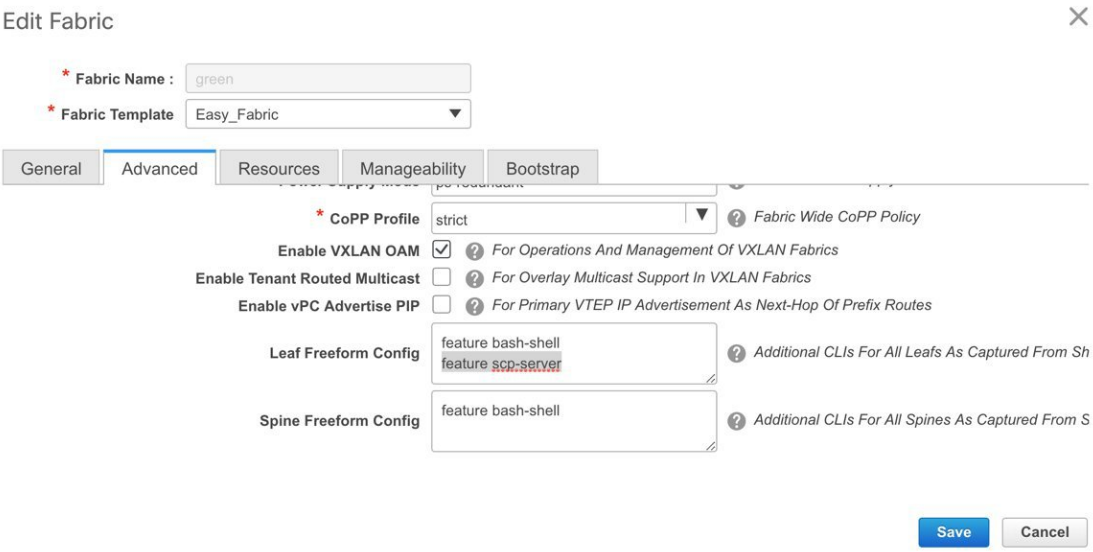

* Fabric Name : green

* Fabric Template Easy_Fabric

* CoPP Profile strict — Fabric Wide CoPP Policy

Enable VXLAN OAM ☑ For Operations And Management Of VXLAN Fabrics

Enable Tenant Routed Multicast □ For Overlay Multicast Support In VXLAN Fabrics

Enable vPC Advertise PIP □ For Primary VTEP IP Advertisement As Next-Hop Of Prefix Routes

Leaf Freeform Config
```
feature bash-shell
feature scp-server
```
Additional CLIs For All Leafs As Captured From Sh

Spine Freeform Config
```
feature bash-shell
```
Additional CLIs For All Spines As Captured From S

4. Click **Save**. The Fabric Builder screen comes up again.
5. Click within the box that represents the fabric. The Fabric Topology screen comes up.
6. Click **Save & Deploy** at the top right part of the screen to save and deploy configurations.

   Configuration Compliance functionality will ensure that that intended configuration as expressed by those CLIs are present on the switches and if they are removed or there is a mismatch, then it will flag it as a mismatch and indicate that the device is OUT-OF-SYNC.

Incomplete Configuration Compliance - On some Cisco Nexus 9000 Series switches, in spite of configuring pending switch configurations using the **Save & Deploy** option, configuration compliance is not successful. Add a **switch_freeform_config** policy to the affected switch (as explained in the *Deploy Freeform CLIs on a Specific Switch* section) to resolve the issue. For example, consider the following persistent pending configurations:

```
line vty
logout-warning 0
```

After adding the above configurations in a **switch_freeform_config** policy and saving the updates, click **Save and Deploy** in the topology screen to complete the deployment process.

Deploy Freeform CLIs on a Specific Switch

1. Click **Control > Fabric Builder**. The Fabric Builder screen comes up.

2. Click on the rectangular box that represents the fabric. The Fabric Topology screen comes up.

Note To provision freeform CLIs on a new fabric, you have to create a fabric, import switches into it, and then deploy freeform CLIs.

3. Right-click the switch icon and select the **View/edit policies** option.

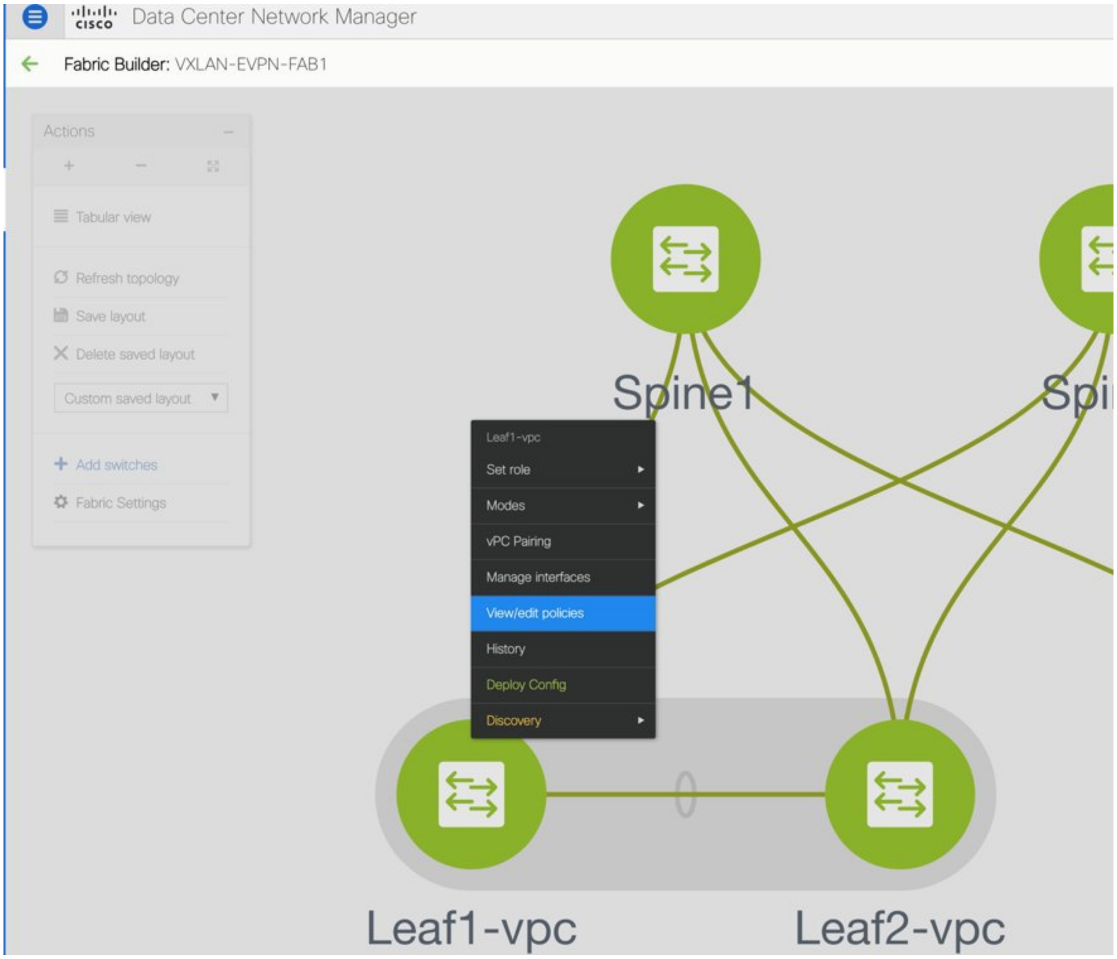

The **View/Edit Policies** screen comes up.

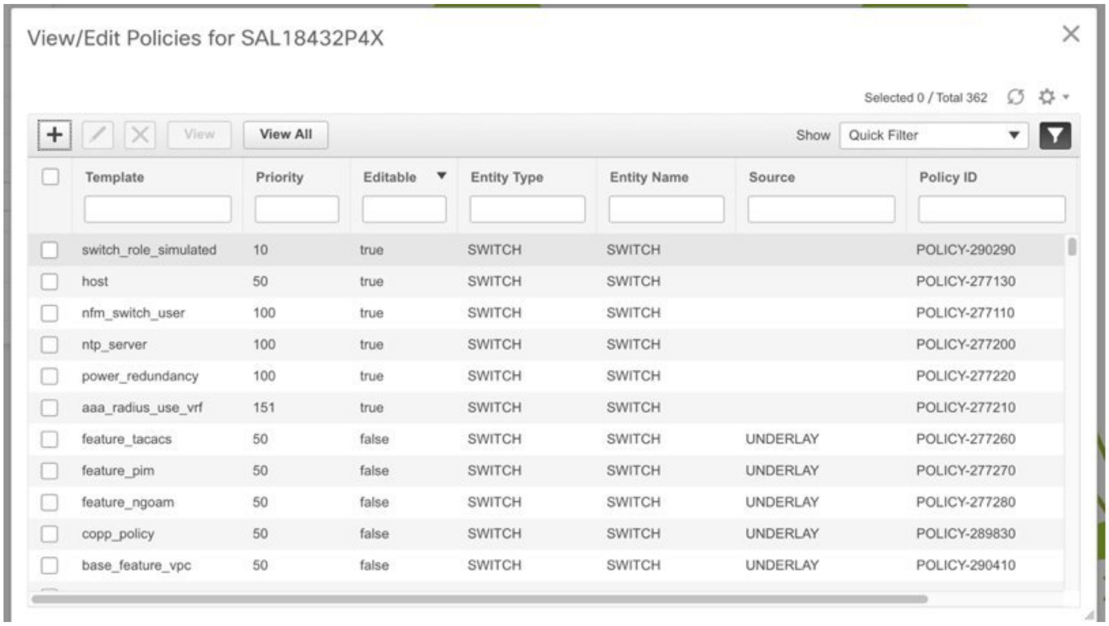

View/Edit Policies for SAL18432P4X

Selected 0 / Total 362

□	Template	Priority	Editable	Entity Type	Entity Name	Source	Policy ID
□	switch_role_simulated	10	true	SWITCH	SWITCH		POLICY-290290
□	host	50	true	SWITCH	SWITCH		POLICY-277130
□	nfm_switch_user	100	true	SWITCH	SWITCH		POLICY-277110
□	ntp_server	100	true	SWITCH	SWITCH		POLICY-277200
□	power_redundancy	100	true	SWITCH	SWITCH		POLICY-277220
□	aaa_radius_use_vrf	151	true	SWITCH	SWITCH		POLICY-277210
□	feature_tacacs	50	false	SWITCH	SWITCH	UNDERLAY	POLICY-277260
□	feature_pim	50	false	SWITCH	SWITCH	UNDERLAY	POLICY-277270
□	feature_ngoam	50	false	SWITCH	SWITCH	UNDERLAY	POLICY-277280
□	copp_policy	50	false	SWITCH	SWITCH	UNDERLAY	POLICY-289830
□	base_feature_vpc	50	false	SWITCH	SWITCH	UNDERLAY	POLICY-290410

4. Click +. The **Add Policy** screen comes up.

   In the **Priority** field, the priority is set to 500 by default. You can choose a higher priority (by specifying a lower number) for CLIs that need to appear higher up during deployment. For example, a command to enable a feature should appear earlier in the list of commands.

5. From the **Policy** field, select **switch_freeform_config**.

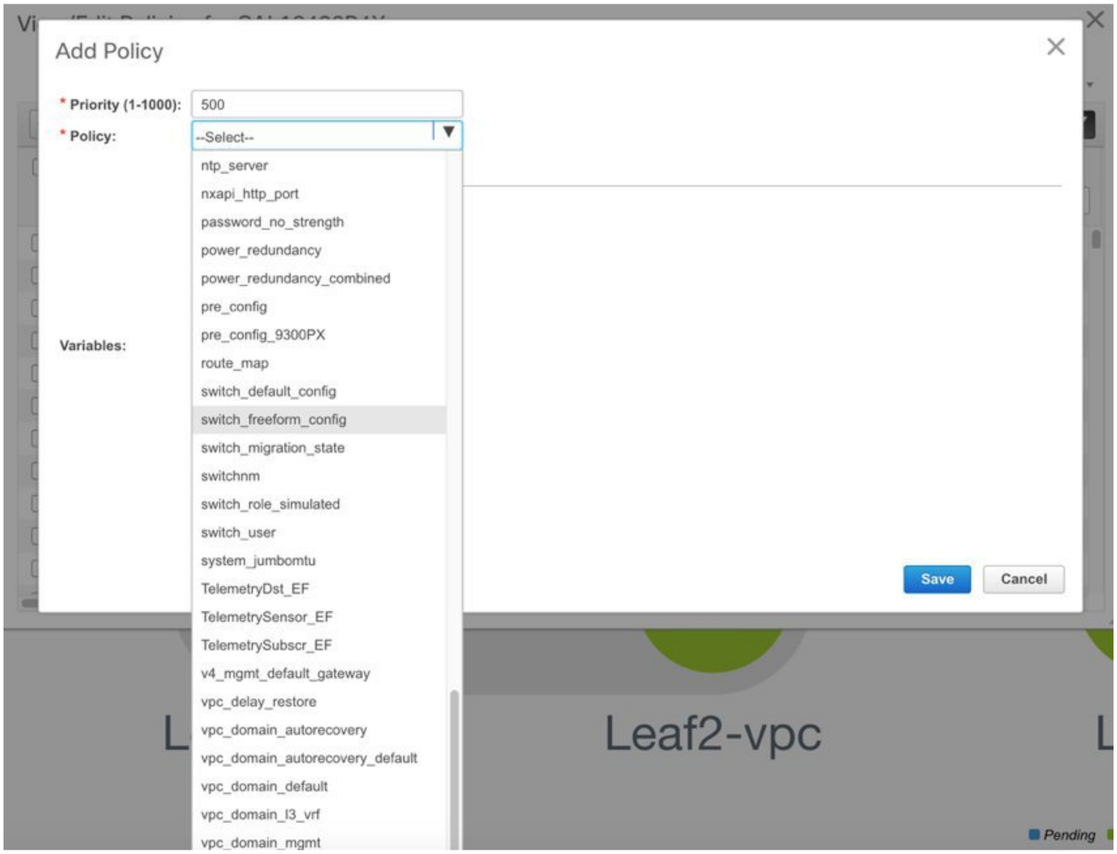

6. Add or update the CLIs in the **Freeform Config CLI** box.

   Copy-paste the intended configuration with correct indentation, as seen in the running configuration on the Nexus switches. For more information, see Resolving Freeform Config Errors in Switches, on page 113.

   A **switch_freeform_config** policy example for VLAN and corresponding SVI instantiation is given below.

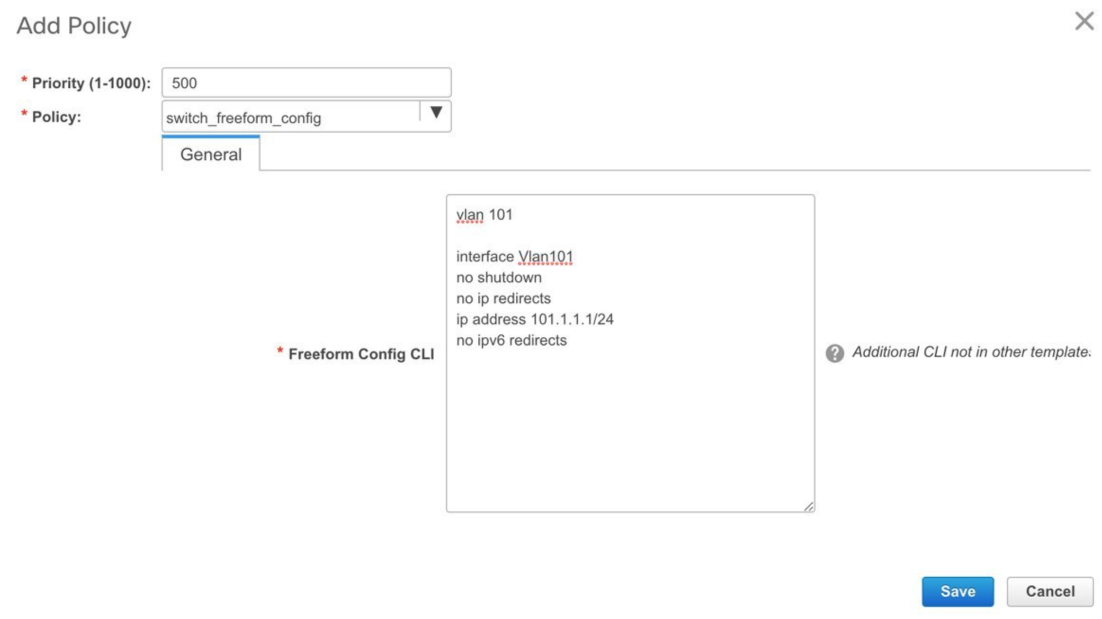

* Priority (1-1000): 500

* Policy: switch_freeform_config

General

* Freeform Config CLI

```
vlan 101

interface Vlan101
no shutdown
no ip redirects
ip address 101.1.1.1/24
no ipv6 redirects
```

Additional CLI not in other template

7. Click **Save**.

   After the policy is saved, it gets added to the intended configurations for that switch.

8. Close the policy screens. The Fabric Topology screen comes up again.

9. Right click the switch and click **Deploy Config**.

   The **Save & Deploy** option can also be used for deployment. However, the **Save & Deploy** option will identify mismatch between the intended and running configuration *across all* fabric switches.

Pointers for *switch_freeform_config* Policy Configuration:

- You can create multiple instances of the policy.
- You can add VLAN, SVI and other features. A specific VLAN and corresponding SVI instantiation should be configured through an individual **switch_freeform_config** policy.
- For a vPC switch pair, create consistent **switch_freeform_config** policies on both the vPC switches.
- Depending on the Cisco Nexus 9000 Series platform type (required for EX, FX, and FX2 platform types), you should include the **system nve infra-vlans 101** command in the policy.

Freeform CLI Configuration Examples

Note Refer the *Deploy Fabric-Wide Freeform CLIs on Leaf and Spine Switches* section and *Deploy Freeform CLIs on a Specific Switch* section for complete steps.

Console line configuration

This example involves deploying some fabric-wide freeform configurations (for all leaf, and spine switches), and individual switch configurations.

Fabric-wide session timeout configuration:

```
line console
  exec-timeout 1
```

Console speed configuration on a specific switch:

```
line console
  speed 115200
```

On the switch where the console speed was updated, both types of configurations are displayed:

```
N9k-switch # show run | b console

line console
  exec-timeout 0
  speed 115200
```

ACL configuration

ACL configurations are typically configured on specific switches and not fabric-wide (leaf/spine switches). When you configure ACLs as freeform CLIs on a switch, you should include sequence numbers. Else, there will be a mismatch between the switch and DCNM. A configuration sample with sequence numbers:

```
ip access-list ACL_VTY
  10 deny tcp 172.29.171.67/32 172.29.171.36/32
  20 permit ip any any
ip access-list vlan65-acl
  10 permit ip 69.1.1.201/32 65.1.1.11/32
  20 deny ip any any

interface Vlan65
  ip access-group vlan65-acl in
line vty
  access-class ACL_VTY in
```

If you have configured ACLs without sequence numbers in a **switch_freeform_config** policy, update the policy with sequence numbers *as displayed in the switch*. After updating, use the per switch **Deploy Config** option by right clicking the device. Alternatively, use the **Save and Deploy** option in the topology screen so that configuration compliance is triggered again and inconsistencies resolved.

Negotiation, speed and duplex port configuration

Consider the following commands configured on a leaf switch whose Ethernet1/10 interface is connected to a spine switch. The ethernet port speed, duplex mode and disabling of automatic negotiation of speed and duplex abilities over the link are configured for the interface.

```
interface Ethernet1/10
   speed 100000
   duplex full
   no negotiate auto
```

This can be configured as a **switch_freeform_config** policy on a switch.

If the above parameters are the same for all leaf switches (interface 1/10 on each leaf switch has the same settings and connected to a switch), then you can update fabric-wide CLIs for all leaf switches.

In the same way, you can configure all spine switches with the same port name and speed, duplex mode and negotiation settings.

Note If you are enabling freeform configurations on all leaf or spine switches, as a best practice, ensure that all switches are connected through the same type of cable. For example, Active Optical Cables or Direct Attach Copper cables.

Resolving Freeform Config Errors in Switches

Copy-paste the running-config to the freeform config with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. Otherwise, configuration compliance in DCNM marks switches as out-of-sync.

Let us see an example of the freeform config of a switch.

```
feature bash-shell
feature telemetry

clock timezone CET 1 0
# Daylight saving time is observed in Metropolitan France from the last Sunday in March
(02:00 CET) to the last Sunday in October (03:00 CEST)
clock summer-time CEST 5 Sunday March 02:00 5 Sunday October 03:00 60
clock protocol ntp

telemetry
  destination-profile
    use-vrf management
```

The highlighted line about the daylight saving time is a comment that is not displayed in the **show running config** command output. Therefore, configuration compliance marks the switch as out-of-sync because the intent does not match the running configuration.

Let us check the running config in the switch for the clock protocol.

```
spine1# show run all | grep "clock protocol"
clock protocol ntp vdc 1
```

You can see that **vdc 1** is missing from the freeform config.

In this example, let us copy-paste the running config to the freeform config.

Here is the updated freeform config:

```
feature bash-shell
feature telemetry

clock timezone CET 1 0
clock summer-time CEST 5 Sunday March 02:00 5 Sunday October 03:00 60
clock protocol ntp vdc 1

telemetry
```

```
destination-profile
  use-vrf management
```

After you copy-paste the running config and deploy, the switch will be in-sync. When you click **Save & Deploy**, the **Side-by-side Comparison** tab in the **Config Preview** window provides you information about the difference between the defined intent and the running config.

Management

The Management menu includes the following submenus:

Resources

Cisco DCNM allows you to manage the resources. The following table describes the fields that appear on this page.

Field	Description
Scope Type	Specifies the scope level at which the resources are managed. The scope types can be **Fabric**, **Device**, **DeviceInterface**, or **DevicePair**.
Scope	Specifies the resource usage scope. Valid values are the switch serial numbers or fabric names. Resources with serial numbers are unique, and can be used on the serial number of the switch only.
Allocated Resource	Specifies if the resources that are managed with device, device interface, or fabric. Valid values are ID type, subnet, or IP addresses.
Allocated To	Specifies the purpose of resource allocation.
Resource Type	Specifies the resource type. The valid values are **TOP_DOWN_VRF_LAN**, **TOP_DOWN_NETWORK_VLAN**, **LOOPBACK_ID**, **VPC_ID**, and so on.
Is Allocated?	Specifies if the resource is allocated or not. The value is set to **True** if the resource is permanently allocated to the given entity. The value is set to **False** if the resource is reserved for an entity and not permanently allocated.
Allocated On	Specifies the date and time of the resource allocation.

Adding, Editing, Re-Discovering and Removing VMware Servers

This section contains the following:

Adding a Virtual Center Server

You can add a virtual center server from Cisco DCNM.

Procedure

Step 1 Choose .

You see the list of VMware servers (if any) that are managed by Cisco DCNM-LAN in the table.

Step 2 Click **Add**.

You see the **Add VCenter** window.

Step 3 Enter the **Virtual Center Server** IP address for this VMware server.

Step 4 Enter the **User Name** and **Password** for this VMware server.

Step 5 Click **Add** to begin managing this VMware server.

Deleting a VMware Server

You can remove a VMware server from the Cisco DCNM.

Procedure

Step 1 Choose .

Step 2 Select the check box next to the VMware server that you want to remove and click **Delete** to discontinue data collection for that VMware server.

Editing a VMware Server

You can edit a VMware server from Cisco DCNM Web Client.

Procedure

Step 1 Choose .

Step 2 Check the check box next to the VMware server that you want to edit and click **Edit** virtual center icon.

You see the **Edit VCenter** dialog box.

Step 3 Enter a the **User Name** and **Password**.

Step 4 Select managed or unmanaged status.

Step 5 Click **Apply** to save the changes.

Rediscovering a VMware Server

You can rediscover a VMware server from Cisco DCNM.

Procedure

Step 1 Choose .

Step 2 Select the check box next to the VMware that you want to rediscover.

Step 3 Click **Rediscover**.
A dialog box with warning "Please wait for rediscovery operation to complete." appears.

Step 4 Click **OK** in the dialog box.

Template Library

You can add, edit, or delete templates that are configured across different Cisco Nexus and Cisco MDS platforms using Cisco DCNM Web client. From Cisco DCNM Web client home page, choose **Control > Template Library > Templates**. The following parameters are displayed for each template that is configured on Cisco DCNM Web client. Templates support JavaScript. You can use the JavaScript function in a template to perform arithmetic operations and string manipulations in the template syntax.

The following table describes the fields that appear on this page.

Table 1: Templates Operations

Field	Description
Add Template	Allows you to add a new template.
Modify/View Template	Allows you to view the template definition and modify as required.
Save Template As	Allows you to save the selected template in a different name. You can edit the template as required.
Delete Template	Allows you to delete a template
Import Template	Allows you to import a template from your local directory, one at a time.
Export template	Allows you to export the template configuration to a local directory location.
Import Template Zip File	Allows you to import `.zip` file, that contains more than one template that is bundled in a `.zip` format All the templates in the ZIP file are extracted and listed in the table as individual templates.

Table 2: Template Properties

Field	Description
Template Name	Displays the name of the configured template.

Field	Description
Template Description	Displays the description that is provided while configuring templates.
Tags	Displays the tag that is assigned for the template and aids to filter templates based on the tags.
Implements	Displays the abstract template to be implemented.
Dependencies	Specifies the specific feature of a switch.
Supported Platforms	Displays the supported Cisco Nexus platforms compatible with the template. Check the check box of platforms that are supported with the template. **Note** You can select multiple platforms.
Template Type	Displays the type of the template.
Template Sub Type	Specifies the sub type that is associated with the template.
Template Content Type	Specifies if it is Jython or Template CLI.
Published	Specifies if the template is published or not.
Imports	Specifies the base template for importing.

In addition, from the menu bar, choose **Control > Template Library > Templates** and you can also:

- Click **Show Filter** to filter the templates that is based on the headers.
- Click **Print** to print the list of templates.
- Click **Export to Excel** to export the list of template to a Microsoft Excel spreadsheet.

This section contains the following:

Template Structure

The configuration template content mainly consists of four parts. Click the **Help** icon next to the **Template Content** for information about editing the content of the template.

This section contains the following:

Template Format

This section describes the basic information of the template. The possible fields are as detailed in the table below.

Property Name	Description	Valid Values	Optional?
name	The name of the template	Text	No

Property Name	Description	Valid Values	Optional?
description	Brief description about the template	Text	Yes
userDefined	Indicates whether the user created the template. Value is ‘true’ if user created.	“true” or “false”	Yes
supportedPlatforms	List of device platforms supports this configuration template. Specify ‘All’ to support all platforms.	N1K, N3K, N3500, N4K, N5K, N5500, N5600, N6K, N7K, N9K, MDS, VDC, N9K-9000v, All list separated by comma.	No
templateType	Specifies the type of Template used.	• CLI • POAP **Note** POAP option is not applicable for Cisco DCNM LAN Fabric deployment. • POLICY • SHOW • PROFILE • FABRIC • ABSTRACT	Yes

Property Name	Description	Valid Values	Optional?
templateSubType	Specifies the sub type associated with the template.		

Property Name	Description	Valid Values	Optional?
		• CLI • N/A • POAP • N/A • VXLAN • FABRICPATH • VLAN • PMN **Note** POAP option is not applicable for Cisco DCNM LAN Fabric deployment. • POLICY • VLAN • INTERFACE_VLAN • INTERFACE_ETHERNET • INTERFACE_BD • INTERFACE_PORT_CHANNEL • INTERFACE_FC • INTERFACE_MGMT • INTERFACE_LOOPBACK • INTERFACE_NVE • INTERFACE_VFC • INTERFACE_SAN_PORT_CHANNEL • DEVICE • FEX • INTERFACE • SHOW • VLAN • INTERFACE_VLAN	

Property Name	Description	Valid Values	Optional?
		• INTERFACE_VPC • INTERFACE_ETHERNET • INTERFACE_BD • INTERFACE_PORT_CHANNEL • INTERFACE_FC • INTERFACE_MGMT • INTERFACE_LOOPBACK • INTERFACE_NVE • INTERFACE_VFC • INTERFACE_SAN_PORT_CHANNEL • DEVICE • FEX • INTERFACE • PROFILE • VXLAN • FABRIC • NA	

Property Name	Description	Valid Values	Optional?
		• ABSTRACT • VLAN • INTERFACE_VLAN • INTERFACE_VPC • INTERFACE_ETHERNET • INTERFACE_BD • INTERFACE_PORT_CHANNEL • INTERFACE_FC • INTERFACE_MGMT • INTERFACE_LOOPBACK • INTERFACE_NVE • INTERFACE_VFC • INTERFACE_SAN_PORT_CHANNEL • DEVICE • FEX • INTERFACE	

Property Name	Description	Valid Values	Optional?
contentType		• CLI • TEMPLATE_CLI • POAP • TEMPLATE_CLI **Note** POAP option is not applicable for Cisco DCNM LAN Fabric deployment. • POLICY • TEMPLATE_CLI • PYTHON • SHOW • TEMPLATE_CLI • PROFILE • TEMPLATE_CLI • PYTHON • FABRIC • PYTHON • ABSTRACT • TEMPLATE_CLI • PYTHON	Yes
implements	Used to implement the abstract template.	Text	Yes
dependencies	Used to select the specific feature of a switch.	Text	Yes
published	Used to Mark the template as read only and avoids changes to it.	“true” or “false”	Yes
timestamp	Shows the template modified time	Modified date and time in the format YYYY-MM-DD HH:MM:SS	Yes

Template Variables

This section contains declared variables, the data type, default values, and valid values conditions for the parameters that are used in the template. These declared variables are used for value substitution in the template content section during the dynamic command generation process. Also these variables are used in decision making and in iteration blocks in the template content section. Variables have predefined data types. You can also add a description about the variable. The following table describes the syntax and usage for the available datatypes.

Variable Type	Valid Value	Iterative?
boolean	true|false	No
enum	`Example: running-config, startup-config`	No
float	Floating number format	No
floatRange	`Example: 10.1,50.01`	Yes
Integer	Any number	No
integerRange	Contiguous numbers separated by "-" Discrete numbers separated by "," `Example: 1-10,15,18,20`	Yes
interface	Format: <if type><slot>[/<sub slot>]/<port> `Example: eth1/1, fa10/1/2 etc.`	No
interfaceRange	`Example: eth10/1/20-25, eth11/1-5`	Yes
ipAddress	IPv4 OR IPv6 address	No
ipAddressList	`Example: 172.22.31.97, 172.22.31.99, 172.22.31.105, 172.22.31.109`	Yes

ipAddressWithoutPrefix	`Example: 192.168.1.1`
	or
	`Example: 1:2:3:4:5:6:7:8`
		No

ipV4Address	IPv4 address	No
ipV4AddressWithSubnet	`Example: 192.168.1.1/24`	No
ipV6Address	IPv6 address	No

ipV6AddressWithPrefix	`Example: 1:2:3:4:5:6:7:8`
	`22`
		No

ipV6AddressWithSubnet	IPv6 Address with Subnet	No

<table>
<tr><th>Variable Type</th><th>Valid Value</th><th>Iterative?</th></tr>
<tr><td>ISISNetAddress</td><td>Example: 49.0001.00a0.c96b.c490.00</td><td>No</td></tr>
<tr><td>long</td><td>Example: 100</td><td>No</td></tr>
<tr><td>macAddress</td><td>14 or 17 character length MAC address format</td><td>No</td></tr>
<tr><td>string</td><td>Free text, for example, used for the description of a variable
Example:
string scheduledTime
{
    regularExpr=^([01]\d|2[0-3]):([0-5]\d)$;
}</td><td>No</td></tr>
<tr><td>string[]</td><td>Example: {a,b,c,str1,str2}</td><td>Yes</td></tr>
<tr><td>struct</td><td>Set of parameters that are bundled under a single variable.
struct <structure name declaration > {
<parameter type> <parameter 1>;
<parameter type> <parameter 2>;
…..
} [<structure_inst1>] [, <structure_inst2>]
[, <structure_array_inst3 []>];
struct interface_detail {
 string inf_name;
 string inf_description;
 ipAddress inf_host;
 enum duplex {
  validValues = auto, full, half;
 };
 }myInterface, myInterfaceArray[];</td><td>No
<b>Note</b> If the struct variable is declared as an array, the variable is iterative.</td></tr>
<tr><td>wwn
(Available only in Cisco DCNM Web Client)</td><td>Example: 20:01:00:08:02:11:05:03</td><td>No</td></tr>
</table>

Variable Meta Property

Each variable that is defined in the template variable section has a set of meta properties. The meta properties are mainly the validation rules that are defined for the variable.

The following table describes the various meta properties applicable for the available variable types.

Variable Type	Description	Variable Meta Property											
		default Value	valid Values	decimal Length	min	max	min Slot	max Slot	min Port	max Port	min Length	max Length	regular Expr
boolean	A boolean value. Example: `true`	Yes											
enum			Yes										
float	signed real number. Example: `75.56, -8.5`	Yes	Yes	Yes	Yes	Yes							
floatRange	range of signed real numbers Example: `50.5 - 54.75`	Yes	Yes	Yes	Yes	Yes							
integer	signed number Example: `50, -75`	Yes	Yes		Yes	Yes							
integerRange	Range of signed numbers Example: `50-65`	Yes	Yes		Yes	Yes							
interface	specifies interface/port Example: `Ethernet 5/10`	Yes	Yes				Yes	Yes	Yes	Yes			
interfaceRange		Yes	Yes				Yes	Yes	Yes	Yes			

Variable Type	Description	Variable Meta Property											
		default Value	valid Values	decimal Length	min	max	min Slot	max Slot	min Port	max Port	min Length	max Length	regular Expr
ipAddress	IP address in IPv4 or IPv6 format	Yes											
ipAddressList	Example: 192.10.2.10, 172.68.10.1 **Note** Separate the addresses in the list using commas and not hyphens.	Yes											
ipAddressWithPrefix	IPv4 or IPv6 Address (does not require prefix/subnet)												
ipV4Address	IPv4 address	Yes											
ipV4AddressWithSubnet	IPv4 Address with Subnet	Yes											
ipV6Address	IPv6 address	Yes											
ipV6AddressWithPrefix	IPv6 Address with prefix	Yes											

Variable Type	Description	Variable Meta Property											
		default Value	valid Values	decimal Length	min	max	min Slot	max Slot	min Port	max Port	min Length	max Length	regular Expr
ipV6AddressWithSubnet	IPv6 Address with Subnet	Yes											
ISISNetAddress	`Example:` `49.0001.0100.1001.0001.00`												
long	`Example: 100`	Yes			Yes	Yes							
macAddress	MAC address												
string	literal string `Example for string` `Regular` `expression:` `string` `scheduledTime` `{` `regularExpr=^([01]\d|2[0-3]):([0-5]\d)$;` `}`	Yes									Yes	Yes	Yes
string[]	string literals that are separated by a comma (,) `Example:` `{string1,` `string2}`	Yes											

Variable Type	Description	Variable Meta Property											
		default Value	valid Values	decimal Length	min	max	min Slot	max Slot	min Port	max Port	min Length	max Length	regular Expr
struct	Set of parameters that are bundled under a single variable. struct <structure name declaration > { <parameter type> <parameter 1>; <parameter type> <parameter 2>; ….. } [<structure_inst1>] [, <structure_inst2>] [, <structure_array_inst3 []>];												
wwn	WWN address												

Example: Meta Property Usage

```
##template variables

integer VLAN_ID {
min = 100;
max= 200;
};

string USER_NAME {
defaultValue = admin123;
minLength = 5;
};

struct interface_a{
 string inf_name;
 string inf_description;
 ipAddress inf_host;
 enum duplex {
  validValues = auto, full, half;
 };
}myInterface;

##
```

Variable Annotation

You can configure the variable properties marking the variables using annotations.

Note Variable Annotations are available for POAP only. However, the annotations do not impact on the template type ‘CLI’.

The following annotations can be used in the template variable section.

Annotation Key	Valid Values	Description
DataDepend	Text	
Description	Text	Description of the field appearing in the window
DisplayName	Text **Note** Enclose the text with quotes, if there is space.	Display name of the field appearing in the window
Enum	Text1, Text2, Text3, and so on	Lists the text or numeric values to select from
IsAlphaNumeric	“true” or “false”	Validates if the string is alphanumeric
IsAsn	“true” or “false”	
IsDestinationDevice	“true” or “false”	
IsDestinationFabric	“true” or “false”	
IsDestinationInterface	“true” or “false”	
IsDestinationSwitchName	“true” or “false”	
IsDeviceID	“true” or “false”	
IsDot1qId	“true” or “false”	
IsFEXID	“true” or “false”	

Annotation Key	Valid Values	Description
IsGateway	“true” or “false”	Validates if the IP address is a gateway
IsInternal	“true” or “false”	Makes the fields internal and does not display them on the window **Note** Use this annotation only for the ipAddress variable.
IsManagementIP	“true” or “false” **Note** This annotation must be marked only for variable “ipAddress”.	
IsMandatory	“true” or “false”	Validates if a value should be passed to the field mandatorily
IsMTU	“true” or “false”	
IsMultiCastGroupAddress	“true” or “false”	
IsMultiLineString	“true” or “false”	Converts a string field to multiline string text area
IsMultiplicity	“true” or “false”	
IsPassword	“true” or “false”	
IsPositive	“true” or “false”	Checks if the value is positive
IsReplicationMode	“true” or “false”	
IsSiteId	“true” or “false”	
IsSourceDevice	“true” or “false”	
IsSourceFabric	“true” or “false”	
IsSourceInterface	“true” or “false”	
IsSourceSwitchName	“true” or “false”	
IsSwitchName	“true” or “false”	
IsRMID	“true” or “false”	

Annotation Key	Valid Values	Description
IsVPCDomainID	“true” or “false”	
IsVPCID	“true” or “false”	
IsVPCPeerLinkPort	“true” or “false”	
IsVPCPeerLinkPortChannel	“true” or “false”	
IsVPCPortChannel	“true” or “false”	
Password	Text	Validates the password field
UsePool	“true” or “false”	
UseDNSReverseLookup		
Username	Text	Displays the username field on the window

Example: DisplayName Annotation

```
##template variables
@(DisplayName="Host Name", Description = "Description of the host")
String hostname;
@(DisplayName="Host Address", Description = " test description" IsManagementIP=true)
ipAddress hostAddress;
##
```

Example: IsMandatory Annotation

```
##template variables
@(IsMandatory="ipv6!=null")
ipV4Address ipv4;
@(IsMandatory="ipv4!=null")
ipV6Address ipv6;
##
```

Example: IsMultiLineString Annotation

```
##template variables
@(IsMultiLineString=true)
string EXTRA_CONF_SPINE;
##
```

IsShow Annotation

```
##template variables
boolean isVlan;
@(IsShow="isVlan==true")
integer vlanNo;
##

##template variables
boolean enableScheduledBackup;
@(IsShow="enableScheduledBackup==true",Description="Server time")
```

```
string scheduledTime;
##
The condition "enableScheduledBackup==true" evaluates to true/false

##template variables
@(Enum="Manual,Back2BackOnly,ToExternalOnly,Both")
string VRF_LITE_AUTOCONFIG;
@(IsShow="VRF_LITE_AUTOCONFIG!=Manual",  Description="Target Mask")
integer DCI_SUBNET_TARGET_MASK
##
The condition "VRF_LITE_AUTOCONFIG!=Manual" matches string comparison to evaluate to true
or false
```

Templates Content

This section includes the configuration commands and any parameters that you want to include in the template. These commands can include the variables declared in the template variables section. During the command generation process the variable values are substituted appropriately in the template content.

Note You must specify the commands that you include as if you were entering them in the global configuration command mode on any device. You must consider the command mode when you include commands.

Template content is governed by the usage of variables.

- Scalar variables: does not take a range or array of values which cannot be used for iteration (In the variable types table those marked iterate-able as 'No'). Scalar variables must be defined inside the template content.

  ```
  Syntax: $$<variable name>$$
  Example: $$USER_NAME$$
  ```

- Iterative variables: used for block iteration. These loop variable must be accessed as shown below inside the iteration block.

  ```
  Syntax:@<loop variable>
  Example:
  foreach val in $$INTEGER_RANGE_VALUE$$ {
  @val
  }
  ```

- Scalar Structure Variable: Structure member variables can be accessed inside the template content.

  ```
  Syntax: $$<structure instance name>.<member variable name>$$
  Example: $$myInterface.inf_name$$
  ```

- Array Structure Variable: Structure member variables can be accessed inside the template content.

  ```
  Syntax: $$<structure instance name>.<member variable name>$$
  Example: $$myInterface.inf_name$$
  ```

In addition to the template variables, you can use the conditional and iterative command generation using the following statements:

- if-else if-else Statement: makes a logical decision in inclusion/exclusion of set of configuration command based on the value assigned for the variable in it.

```
Syntax: if(<operand 1>  <logical operator>  <operand 2>){
command1 ..
command2..
..
}
else  if (<operand 3> <logical operator> <operand 4> )
{
Command3 ..
Command4..
..
}
else
{
Command5 ..
Command6..
..
}
Example: if-else if-else statement
if($$USER_NAME$$ == 'admin'){
Interface2/10
no shut
}
else {
Interface2/10
shut
}
```

- foreach Statement: used for iterating a block of commands. The iteration is performed based on the assigned loop variable value.

```
Syntax:
foreach <loop index variable> in $$<loop variable>$$ {
@<loop index variable> ..
}
Example: foreach Statement
foreach ports in $$MY_INF_RANGE$${
interface @ports
no shut
}
```

- Optional parameters: By default all parameters are mandatory. To make a parameter optional, you must annotate the parameter.

In the variable section, you can include the following command:

- **@(IsMandatory=false)**
- **Integer frequency;**

  In the template content section, a command can be excluded or included without using “if” condition check, by assigning a value to the parameter. The optional command can be framed as below:

- **probe icmp [frequency frequency-value] [timeout seconds] [retry-count retry-count-value]**

Advanced Features

The following are the advanced features available to configure templates.

- Assignment Operation

Config template supports assignment of variable values inside the template content section. The values are validated for the declared data type of the variable. If there is a mismatch, the value is not assigned.

Assignment operation can be used under the following guidelines:

- The operator on the left must be any of the template parameters or a for loop parameter.
- The operator on the right values can be any of the values from template parameters, for loop parameters, literal string values surrounded by quotes or simple string values.

If a statement does not follow these guidelines, or if it does not suit this format, it will not be considered as assignment operation. It is substituted during command generation like other normal lines.

```
Example: Template with assignment operation
##template properties
name =vlan creation;
userDefined= true;
supportedPlatforms = All;
templateType = CLI;
published = false;
##
##template variables
integerRange vlan_range;
@(internal=true)
integer vlanName;
##
##template content
foreach vlanID in $$vlan_range$${
vlan @vlanID
$$vlanName$$=@vlanID
name myvlan$$vlanName$$
}
##
```

- Evaluate methods

  Config template uses the Java runtime provided Java script environment to perform arithmetic operations (such as ADD, SUBTRACT, and so on), string manipulations, and so on.

  Locate the JavaScript file in the template repository path. This file contains primary set of arithmetic, string functions. You can also add custom JavaScript methods.

  These methods can be called from config template content section in below format:

```
Example1:
$$somevar$$ = evalscript(add, "100", $$anothervar$$)
```

  Also the *evalscript* can be called inside if conditions as below:

```
if($$range$$ > evalscript(sum, $$vlan_id$$,  -10)){
do something...
}
```

  You can call a method that is located at the backend of the Java script file.

- Dynamic decision

  Config template provides a special internal variable “LAST_CMD_RESPONSE”. This variable stores the last command response from the device during the execution of the command. This can be used in the config template content to make dynamic decisions to deliver the commands that are based on the device condition.

Note The if block must be followed by an else block in a new line, which can be empty.

An example use case to create a VLAN, if it is does not exist on the device.

```
Example: Create VLAN
##template content
show vlan id $$vlan_id$$
if($$LAST_CMD_RESPONSE$$ contains "not found"){
vlan $$vlan_id$$
}
else{
}
##
```

This special implicit variable can be used only in the “IF” blocks.

- Template referencing

  You can have a base template with all the variables defined. This base template can be imported to multiple templates. The base template content is substituted in the appropriate place of the extending template. The imported template parameters and the contents can be accessed inside the extending template.

```
Example: Template Referencing
Base template:
##template properties
 name =a vlan base;
 userDefined= true;
 supportedPlatforms = All;
 templateType = CLI;
 published = false;
 timestamp = 2015-07-14 16:07:52;
 imports = ;
##
##template variables
 integer vlan_id;
##
##template content
 vlan $$vlan_id$$
##

Derived Template:
##template properties
 name =a vlan extended;
 userDefined= true;
 supportedPlatforms = All;
 templateType = CLI;
 published = false;
 timestamp = 2015-07-14 16:07:52;
 imports = a vlan base,template2;
##
##template variables
 interface vlanInterface;
##
##template content
 <substitute a vlan base>
 interface $$vlanInterface$$
 <substitute a vlan base>
##
```

When you launch the extended template, the parameter inputs for the base template are also obtained. In addition, the substituted content is used for complete CLI command generation.

Adding a Template

To add user-defined templates and schedule jobs from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Template Library**.

The **Templates** window is displayed with the name of the template along with its description, supported platforms, and tags.

Step 2 Click **Add** to add a new template.

Step 3 Specify a **Template Name**, **Template Description**, **Tags**, **Implements**, and **Dependencies** for the new template. Specify a template name, description, tags, and supported platforms for the new template.

Step 4 Select the supported platforms that the template must support.

Step 5 Specify a **Template Type** for the template.

Step 6 Select a **Template Sub Type** and **Template Content Type** for the template. Select **Published** to make the template read-only. You cannot edit a published template.

Step 7 Click **Template Content** to edit the template syntax. For information about the structure of the Configuration Template, see the *Template Structure* section.

Step 8 From the **Imports > Template Name** list, check the template check box.

The base template content is displayed in the **Template Content** window. The base template displays the template properties, template variables, and template content. This template can be imported in to another template and the base template content is substituted in the appropriate place of the extending template. When you launch the extended template, the parameter inputs for the base template are also obtained. Also, the substituted content is used for complete CLI command generation.

Note The base templates are CLI templates.

Step 9 Click **Validate Template Syntax** to validate the template values.

If an error or a warning message appears, you can check the validation details in **Validation Table**.

Note You can continue to save the template if there are warnings only. However, if there is an error, you must edit the templates to fix the errors before you proceed.

Step 10 Click **Save** to save the template.

Step 11 Click **Save and Exit** to save the configuration and go back to the configuring templates screen.

Modifying a Template

You can edit the user-defined templates. However, the predefined templates and templates that are already published cannot be edited.

Procedure

Step 1 From **Control > Template Library**, select a template.

Step 2 Click **Modify/View template**.

Step 3 Edit the template description and tags.

The edited template content is displayed in a pane on the right.

Step 4 From the **Imports > Template Name** list, check the template check box.

The base template content is displayed in the **Template Content** window. You can edit the template content based on your requirement in the **Template Content** window. Click the help icon next to the **Template Content** window for information about editing the content of the template.

Step 5 Edit the supported platforms for the template.

Step 6 Click **Validate Template Syntax** to validate the template values.

Step 7 Click **Save** to save the template.

Step 8 Click **Save and Exit** to save the configuration and go back to the configuring templates screen.

Copying a Template

To copy a template from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Template Library**, and select a template.

Step 2 Click **Save Template As**.

Step 3 Edit the template name, description, tags, and other parameters.

The edited template content is displayed in the right-hand pane.

Step 4 From the **Imports > Template Name** list, check the template check box.

The base template content is displayed in the **Template Content** window. You can edit the template content that is based on your requirement in the **Template Content** window. Click the help icon next to the **Template Content** window for information about editing the content of the template.

Step 5 Edit the supported platforms for the template.

Step 6 Click **Validate Template Syntax** to validate the template values.

Step 7 Click **Save** to save the template.

Step 8 Click **Save and Exit** to save the configuration and go back to the configuring templates screen.

Deleting a Template

You can delete the user-defined templates. However, you cannot delete the predefined templates. From Cisco DCNM Release 11.0(1), you can delete multiple templates at once.

To delete a template from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Template Library**.

Step 2 Use the check box to select a template and click **Remove template** icon.

The template is deleted without any warning message.

What to do next

The template is deleted from the list of templates on the DCNM Web UI. When you restart the DCNM services, the deleted templates are displayed on the **Control > Template Library** page.

To delete the template permanently, delete the template that is located in your local directory: `Cisco Systems\dcm\dcnm\data\templates\`.

Importing a Template

To import a template from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Template Library** and click **Import Template**.

Step 2 Browse and select the template that is saved on your computer.

You can edit the template parameters, if necessary. For information, see Modifying a Template, on page 137.

Note The “\n” in the template is considered as a new line character when imported and edited, but it works fine when imported as a ZIP file.

Step 3 Click **Validate Template Syntax** to validate the template.

Step 4 Click **Save** to save the template or **Save and Exit** to save the template and exit.

Exporting a Template

To export a template from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Template Library**.

Step 2 Use the check box to select a template and click **Export Template**.

The browser requests you to open or save the template to your directory.

Image Management

The **Image Management** menu includes the following options:

Image Upload

This feature allows you to upload or delete images that are used during POAP and switch upgrade.

Deleting an Image

To delete an image from the repository from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 On the **Smart Image Management** window, select an existing image from the list, and click **Delete**.

Step 2 In the delete notification, click **Yes** to delete the image.

Note The default SCP Repository cannot be deleted.

Image Upload

To upload different types of images to the server from the Cisco DCNM Web UI, perform the following steps:

Note Devices use these images during POAP.

Procedure

Step 1 On the **Smart Image Management** window, check the server name check box to select the server for uploading images.

The **Select Image File** window appears.

Step 2 Click **Browse** to select the image file from the directory.

Step 3 From the **Platform** drop-down list, select the device to which you must upload this image.

Step 4 From the **Type** drop-down list, select the type of the image you are uploading to the device.

Step 5 Click **OK**.

The image is uploaded to the repository.

Install & Upgrade

The **Install & Upgrade** menu includes the following submenus:

Upgrade History

This feature enables you to upgrade the Cisco Nexus Platform Switches using In-Service Software Upgrade (ISSU). This upgrade procedure may be disruptive or non-disruptive based on the device configuration. You can select the Kickstart, System, or NX-OS images from image repository or the file system on the device. To select the images from the repository, the same needs to be uploaded from **Control > Image Management > Image upload** tab.

The following table describes the fields that appear on **Control > Image Management > Upgrade History**.

Field	Description
Task Id	Specifies the serial number of the task. The latest task will be listed in the top. **Note** If Failover is triggered in Native HA, the Task Id sequence number is incremented by 32.
Task Type	Specifies the type of task. • Compatibility • Upgrade
Owner	Based on the Role-Based Authentication Control (RBAC), specifies the owner who initiated this task.
Devices	Displays all the devices that were selected for this task.
Job Status	Specifies the status of the job. • Planned • In Progress • Completed • Completed with Exceptions
Created Time	Specifies the time when the task was created.

Field	Description
Scheduled At	Specifies the time when the task is specified to be executed. You can also choose to schedule a task to be executed at a later time.
Comment	Shows any comments that the Owner has added while performing the task.

Note After a fresh Cisco DCNM installation, this page will have no entries.

You can perform the following:

New Installation

To upgrade the devices that are discovered from the Cisco DCNM, perform the following steps:

Procedure

Step 1 Choose **Control > Image Management > Install & Upgrade > Upgrade History**, click **New Installation** to install, or upgrade the kickstart and the system images on the devices.

The devices with default VDCs are displayed in the **Select Switches** window.

Step 2 Select the check box to the left of the switch name.

You can select more than one device and move the devices to the right column.

Step 3 Click **Add** or **Remove** icons to include the appropriate switches for upgrade.

The selected switches appear in a column on the right.

Step 4 Click **Next** to navigate to the **Specify Software Images** window. This tab displays the switches that you selected in the previous screen and allows you to choose the images for upgrade.

- The **Auto File Selection** check box enables you to specify an image version, and a path where you can apply the upgraded image to the selected devices.
- **Select File Server** is disabled, and the default server is used.

- In the **Image Version** field, specify the image version as displayed in **Image Upload** screen.
- The **Path** field is disabled, and the default image path is used.

Step 5 Click **Select Image** in the **Kickstart image** column.

The **Software Image Browser** dialog box appears.

Note

- Cisco Nexus 9000 Series Switches require only the system image to load the Cisco NX-OS operating system. Therefore, the option to select kickstart images for these devices is disabled.

- If there is an issue in viewing the **Software Image Browser** dialog box, reduce the font size of your browser and retry.

Step 6 Click **Select Image** in the **System Image** column.

The **Software Image Browser** dialog box appears.

Step 7 On the **Software Image Browser** dialog box, you can choose the image from **File Server** or **Switch File System**.

If you choose **File Server**:

a) From the **Select the File server** list, choose the appropriate file server on which the image is stored.

b) From the **Select Image** list, choose the appropriate image. Check the check box to use the same image for all other selected devices of the same platform.

   Example: For platform types N9K-C93180YC-EX and N9K-C93108TC-EX, logic matches platform (N9K) and three characters (C93) from subplatform. The same logic is used across all platform switches.

c) Click **OK** to choose the kickstart image or **Cancel** to revert to the **Specify Software Images** window.

If you choose **Switch File System**:

a) From the **Select Image** list, choose the appropriate image that is located on the flash memory of the device.
b) Click **OK** to choose the kickstart image or **Cancel** to revert to the **Specify Software Images** dialog box.

Step 8 The **Vrf** column indicates the name of the virtual routing and forwarding (VRF).

Step 9 In the **Available Space** column, specify the available space for the **Primary Supervisor** and **Secondary Supervisor** modules of the switch.

Available Space column shows the available memory in MB on the switch (for less than 1 MB, it is shown and marked as KB).

Bootflash browser shows the filename, size, and last modified date for all the files and directories on the switch bootflash. You can delete files by selecting them and clicking **Delete** to increase the available space on the switch.

Step 10 **Selected Files Size** column shows the size of images that are selected from the server.

If the total size of selected images is greater than available space on a switch, the file size is marked in red. We recommend that you create more space on the switch to copy images to it and install.

Step 11 Drag and drop the switches to reorder the upgrade task sequence.

Step 12 Select **Skip Version Compatibility** if you are sure that the version of the Cisco NX-OS software on your device is compatible with the upgraded images that you have selected.

Step 13 Select **Select Parallel Line Card upgrade** to upgrade all the line cards at the same time.

Upgrading a parallel line card is not applicable for Cisco MDS devices.

Step 14 Select **Options** under the **Upgrade Options** column to choose the type of upgrade.

Upgrade Options window appears with two upgrade options. The drop-down list for **Upgrade Option 1** has the following options:

- **NA**
- **bios-force**
- **non-disruptive**

NA is the default value.

The drop-down list for **Upgrade Option 2** has the following options:

- **NA**
- **bios-force**

When **NA** is selected under **Upgrade Option 1**, **Upgrade Option 2** is disabled.

When **bios-force** is selected under **Upgrade Option 1**, **NA** is the only option under **Upgrade Option 2**

When **non-disruptive** is selected under **Upgrade Option 1**, you can choose **NA** or **bios-force** under **Upgrade Option 2**.

Check the **Use this Option for all other selected devices** check box to use the selected option for all the selected devices and click **OK**.

Note

- The upgrade options are applicable only for Cisco Nexus 3000 Series and 9000 Series switches.
- Selecting the non-disruptive option for upgrading does not ensure a non-disruptive upgrade. Perform a compatibility check to ensure that the device supports non-disruptive upgrade.

Step 15 Click **Next**.

If you did not select **Skip Version Compatibility**, the Cisco DCNM performs a compatibility check.

You can choose to wait until the check is complete or click **Finish Installation Later**.

The installation wizard is closed and a compatibility task is created in **Control > Image Management > Install & Upgrade > Upgrade History** tasks.

The time that is taken to check the image compatibility depends on the configuration and the load on the device.

The **Version Compatibility Verification** status column displays the status of verification.

If you skip the version compatibility check by choosing **Skip Version Compatibility**, Cisco DCNM displays only the name of the device, the **Current Action** column displays **Completed**, and the **Version Compatibility Verification** column displays **Skipped**.

Step 16 Click **Finish Installation Later** to perform the upgrade later.

Step 17 Click **Next**.

Step 18 Check the check box to save the running configuration to the startup configuration before upgrading the device.

Step 19 You can schedule the upgrade process to occur immediately or later.

1. Select **Deploy Now** to upgrade the device immediately.
2. Select **Choose time to Deploy** and specify the time in MMM/DD/YYYY HH:MM:SS format to perform the upgrade later.

   This value is relative to the server time. If the selected time to deploy is in the past, the job is executed immediately.

Step 20 You can choose the execution mode based on the devices and the line cards you have chosen to upgrade.

1. Select **Sequential** to upgrade the devices in the order in which they were chosen.
2. Select **Concurrent** to upgrade all the devices at the same time.

Step 21 Click **Finish** to begin the upgrade process.

The Installation wizard closes and a task to Upgrade is created on the **Control > Image Management > Install & Upgrade > Upgrade History** page.

What to do next

After you complete the ISSU on the switch, ensure that you wait for 20 minutes to allow the switch to reboot, and stabilize the SNMP agent. Cisco DCNM will discovery polling cycles in order to display the new version of the switch on the Cisco DCNM Web UI.

Finish Installation

You can choose to complete the installation for tasks which was completed on the **Compatibility Check** page. Perform the following task to complete the upgrade process on the devices.

Procedure

Step 1 Choose **Control > Image Management > Install & Upgrade > Upgrade History**, select a task for which the compatibility check is complete.

Select only one task at a time.

Step 2 Click **Finish Installation**.

Software Installation Wizard appears.

Step 3 Check the check box to save the running configuration to the startup configuration before upgrading the device.

Step 4 Check the check box to put a device in maintenance mode before upgrade. This option is valid only for the devices that support maintenance mode.

Step 5 You can schedule the upgrade process to occur immediately or later.

1. Select **Deploy Now** to upgrade the device immediately.
2. Select **Choose time to Deploy** and specify the time in DD/MM/YYYY HH:MM:SS format to perform the upgrade later.

Step 6 You can choose the execution mode that is based on the devices and the line cards that you have chosen to upgrade.

1. Select **Sequential** to upgrade the devices in the order in which they were chosen.
2. Select **Concurrent** to upgrade the devices at the same time.

Step 7 Click **Finish** to complete the upgrade process.

View

To view the image upgrade history from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Image Management > Install & Upgrade > Upgrade History**, check the task ID check box.

Select only one task at a time.

Step 2 Click **View**.

The **Installation Task Details** window is displayed.

Step 3 Click **Settings**. Select **Columns** and choose the column details options.

This window displays the location of the kickstart and system images, compatibility check status, installation status, descriptions, and logs.

Step 4 Select the device.

The detailed status of the task is displayed. For the completed tasks, the response from the device is displayed.

If the upgrade task is in progress, a live log of the installation process appears.

Note This table is refreshed every 30 secs for jobs in progress, when you are on this window.

Delete

To delete a task from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control** > **Image Management** > **Install & Upgrade** > **Upgrade History**, and check the **Task ID** check box.

Step 2 Click **Delete**.

Step 3 Click **OK** to confirm deletion of the job.

Switch Level History

You can view the history of the upgrade process at a switch level. You can view the current version of the switch and other details.

The following table describes the fields that appear on **Control > Image Management > Install & Upgrade > Switch Level History**.

Field	Description
Switch Name	Specifies the name of the switch
IP Address	Specifies the IP Address of the switch
Platform	Specifies the Cisco Nexus switch platform

Field	Description
Current Version	Specifies the current version on the switch software

Click the radio button next to a switch name to select the switch and view its upgrade history. Click **View** to view the upgrade task history for the selected switch.

The following table describes the fields that appear on **Control > Image Management > Install & Upgrade > Switch Level History > View Device Upgrade Tasks**:

Field	Description
Owner	Specifies the owner who initiated the upgrade.
Job Status	Specifies the status of the job. • Planned • In Progress • Completed
KickStart Image	Specifies the kickStart image that is used to upgrade the Switch.
System Image	Specifies the system image that is used to upgrade the switch.
Completed Time	Specifies the date and time at which the upgrade was successfully completed.
Status Description	Specifies the installation log information of the job.

Endpoint Locator

The Endpoint Locator menu includes the following submenus:

Endpoint Locator

The Endpoint Locator (EPL) feature allows real-time tracking of endpoints within a data center. The tracking includes tracing the network life history of an endpoint and getting insights into the trends that are associated with endpoint additions, removals, moves, and so on. An endpoint is anything with at least one IP address and MAC address. In that sense, an endpoint can be a virtual machine (VM), container, bare-metal server, service appliance and so on.

Important

- EPL is supported for VXLAN BGP EVPN fabric deployments only in the DCNM LAN fabric installation mode. The VXLAN BGP EVPN fabric can be deployed as Easy fabric, Easy eBGP fabric, or an External fabric (managed or monitored mode). EPL is not supported for 3-tier access-aggregation-core based network deployments.
- EPL displays endpoints that have at least one IP address (IPv4 and/or IPv6). Also, these endpoints must be residing in networks where the gateway or SVI is configured on the network switches within the VXLAN EVPN fabric. In other words, EPL cannot determine the identity (IPv4/IPv6 address) of the endpoints for networks that are deployed as Layer-2 Only within the fabric.

EPL relies on BGP updates to track endpoint information. Hence, typically the DCNM needs to peer with the BGP Route-Reflector (RR) to get these updates. For this purpose, IP reachability from the DCNM to the RR is required. This can be achieved over in-band network connection to the DCNM eth2 interface.

Some key highlights of the Endpoint Locator are:

- Support for dual-homed and dual-stacked (IPv4 + IPv6) endpoints
- Support for up to two BGP Route Reflectors
- Support real-time and historical search for all endpoints across various search filters such as VRF, Network, Layer-2 VNI, Layer-3 VNI, Switch, IP, MAC, port, VLAN, and so on.
- Support for real-time and historical dashboards for insights such as endpoint lifetime, network, endpoint, VRF daily views, and operational heat map.
- Support for high availability
- Support for endpoint data that is stored for up to 180 days, amounting to a maximum of 5 G storage space.
- Support for optional flush of the endpoint data in order to start afresh.
- Supported scale: 10K endpoints

For more information about EPL, refer to the following sections:

Configuring Endpoint Locator

The DCNM OVA or the ISO installation comes with 3 interfaces—eth0 interface for external access to the DCNM, eth1 interface that is used primarily for fabric management, and eth2 interface for in-band network connectivity to Cisco DCNM. In most deployments the eth1 interface is part of the same network on which the mgmt0 interfaces of the Cisco Nexus switches reside (Out-of-band or OOB network). This allows DCNM to perform out-of-band management of these devices including POAP.

BGP peering between the Cisco DCNM and the Route-Reflector is required for EPL. Since the BGP process on Nexus devices typically runs on the non-management VRF, specifically default VRF, it requires an in-band IP connectivity from the Cisco DCNM to the fabric. For this purpose, the eth2 interface can be configured using the **appmgr setup inband** command. The user will be prompted to specify an IP address, netmask and gateway IP. On the fabric side if the DCNM eth2 port is directly connected to one of the front-end interfaces on a switch then the front-end interface can be configured using the *epl_routed_intf* template.

After the in-band connectivity is established between the physical or virtual DCNM and the fabric, BGP peering can be established. There is a simple wizard for enabling Endpoint Locator.

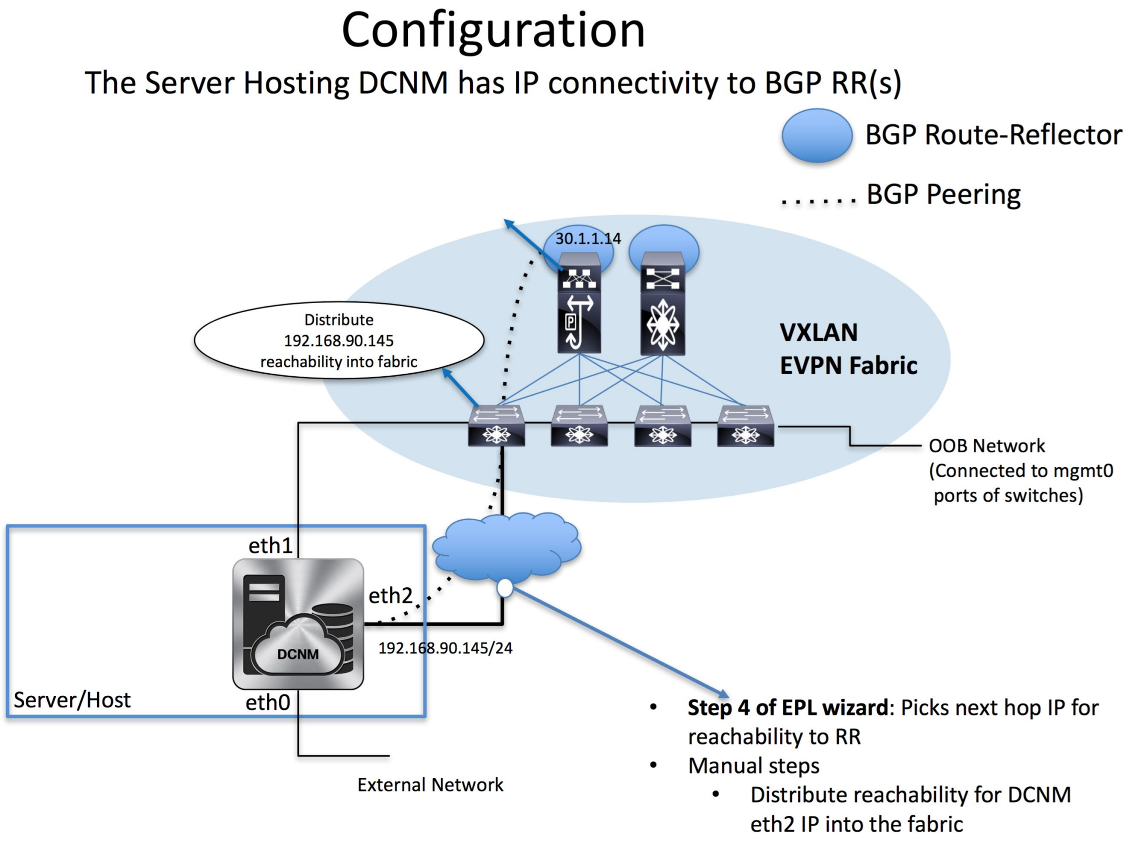

Configuration

The Server Hosting DCNM has IP connectivity to BGP RR(s)

- **Step 4 of EPL wizard**: Picks next hop IP for reachability to RR
- Manual steps
    - Distribute reachability for DCNM eth2 IP into the fabric

During the EPL configuration using the wizard, the route reflectors (RRs) are configured to accept DCNM as a BGP peer. During the same configuration, the DCNM is also configured by adding routes to the BGP loopback IP via the eth2 gateway. The DCNM can be directly attached to a ToR, or leaf, that in turn provides reachability to the RR. Also, DCNM can have simple IP connectivity via a gateway to the fabric in any case the gateway of eth2 should be appropriately configured when setting up the eth2 port on DCNM.

Note Cisco DCNM queries the BGP RR to glean information for establishment of the peering, like ASN, RR, IP, and so on.

It should be noted that it is very important to configure eth2 interface properly, if it is a native HA setup then eth2 on active and standby Cisco DCNMs must be in the same subnet, which means they should have the same gateway addresses.

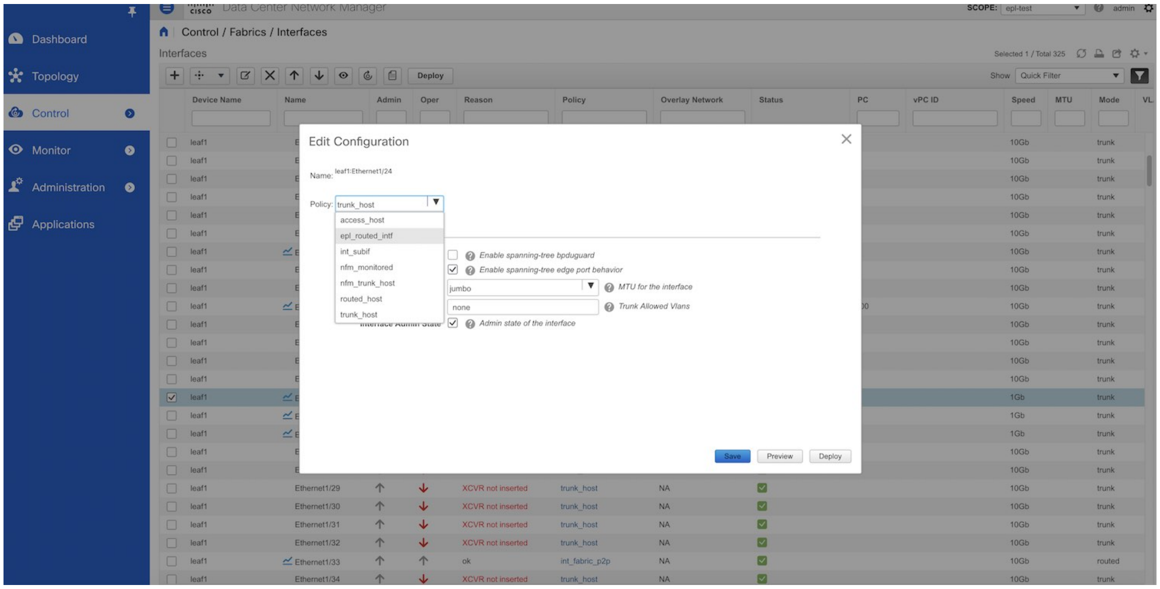

Control / Fabrics / Interfaces

Edit Configuration

Name: leaf1:Ethernet1/24

Policy: trunk_host

- access_host
- epl_routed_intf
- int_subif
- nfm_monitored
- nfm_trunk_host
- routed_host
- trunk_host

☐ Enable spanning-tree bpduguard
☑ Enable spanning-tree edge port behavior
jumbo — MTU for the interface
none — Trunk Allowed Vlans
Interface Admin State ☑ Admin state of the interface

Save Preview Deploy

Device Name	Name	Admin	Oper	Reason	Policy	Overlay Network	Status	Speed	Mode
leaf1	Ethernet1/29	↑	↓	XCVR not inserted	trunk_host	NA	✓	10Gb	trunk
leaf1	Ethernet1/30	↑	↓	XCVR not inserted	trunk_host	NA	✓	10Gb	trunk
leaf1	Ethernet1/31	↑	↓	XCVR not inserted	trunk_host	NA	✓	10Gb	trunk
leaf1	Ethernet1/32	↑	↓	XCVR not inserted	trunk_host	NA	✓	10Gb	trunk
leaf1	Ethernet1/33	↑	↑	ok	int_fabric_p2p	NA	✓	10Gb	routed
leaf1	Ethernet1/34	↑	↓	XCVR not inserted	trunk_host	NA	✓	10Gb	trunk

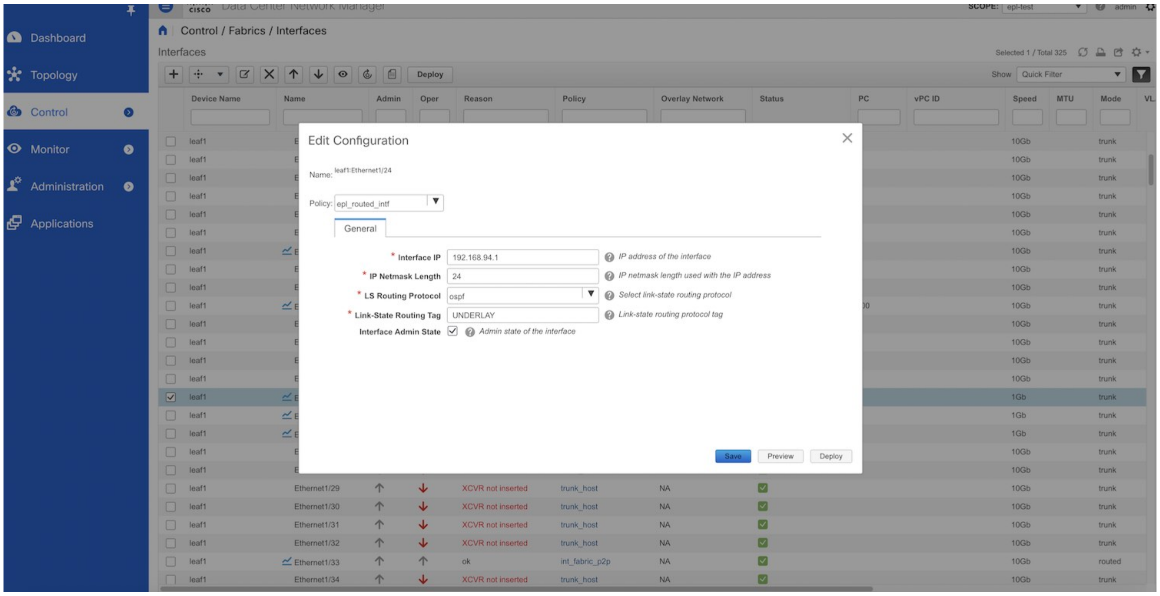

Control / Fabrics / Interfaces

Edit Configuration

Name: leaf1:Ethernet1/24

Policy: epl_routed_intf

General

* Interface IP: 192.168.94.1 — IP address of the interface
* IP Netmask Length: 24 — IP netmask length used with the IP address
* LS Routing Protocol: ospf — Select link-state routing protocol
* Link-State Routing Tag: UNDERLAY — Link-state routing protocol tag
Interface Admin State: ☑ Admin state of the interface

Save Preview Deploy

Procedure

Step 1 From the Cisco DCNM home page, choose **Control > Endpoint Locator > Configure**. The **Endpoint Locator** page appears with a **See how it works** help link.

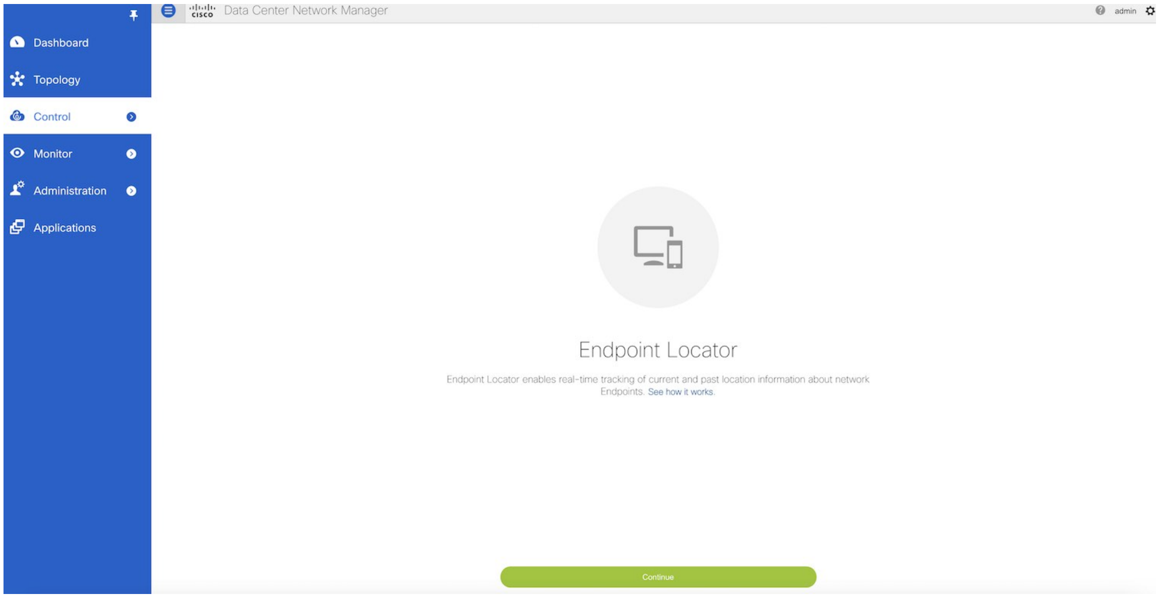

Endpoint Locator

Endpoint Locator enables real-time tracking of current and past location information about network Endpoints. See how it works.

Step 2 Click **Continue**.

Step 3 Select the appropriate fabric on which the endpoint locator feature should be enabled to track endpoint activity.

EPL can only be enabled for one fabric. It can be DFA or EVPN.

Step 4 Select the switches on the fabric hosting the RRs. Cisco DCNM will peer with the RRs.

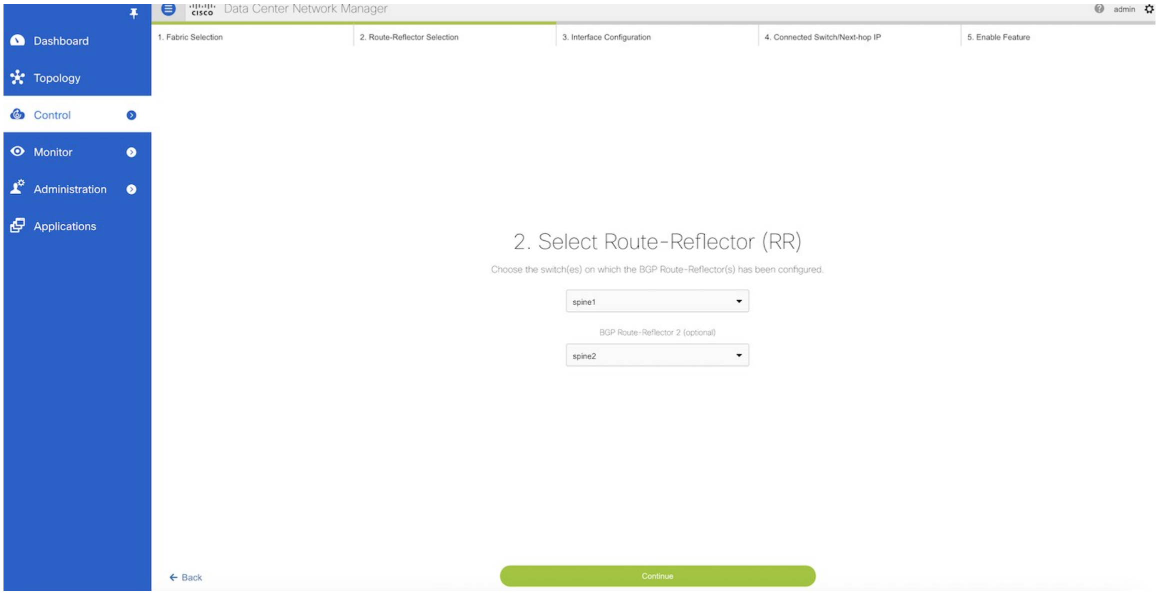
Data Center Network Manager
admin
Dashboard
Topology
Control
Monitor
Administration
Applications
1. Fabric Selection
2. Route-Reflector Selection
3. Interface Configuration
4. Connected Switch/Next-hop IP
5. Enable Feature
Back
Continue

2. Select Route-Reflector (RR)

Choose the switch(es) on which the BGP Route-Reflector(s) has been configured.

spine1

BGP Route-Reflector 2 (optional)

spine2

Step 5 Check DCNM eth2 configuration for IP reachability to the RR.

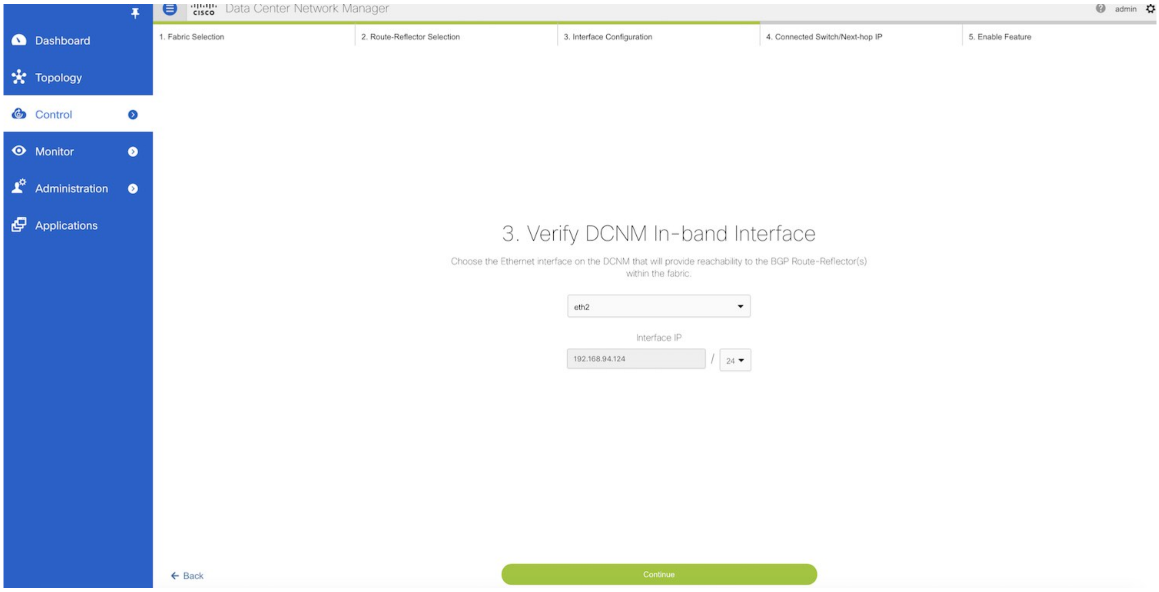

Data Center Network Manager
admin
Dashboard
Topology
Control
Monitor
Administration
Applications
1. Fabric Selection
2. Route-Reflector Selection
3. Interface Configuration
4. Connected Switch/Next-hop IP
5. Enable Feature
Back
Continue

3. Verify DCNM In-band Interface

Choose the Ethernet interface on the DCNM that will provide reachability to the BGP Route-Reflector(s) within the fabric.

eth2

Interface IP: 192.168.94.124 / 24

Step 6 Check Next-hop IP, and ensure the gateway IP is correct. If there is an error go to command line and reconfigure the eth2 port using the **appmgr setup inband** command.

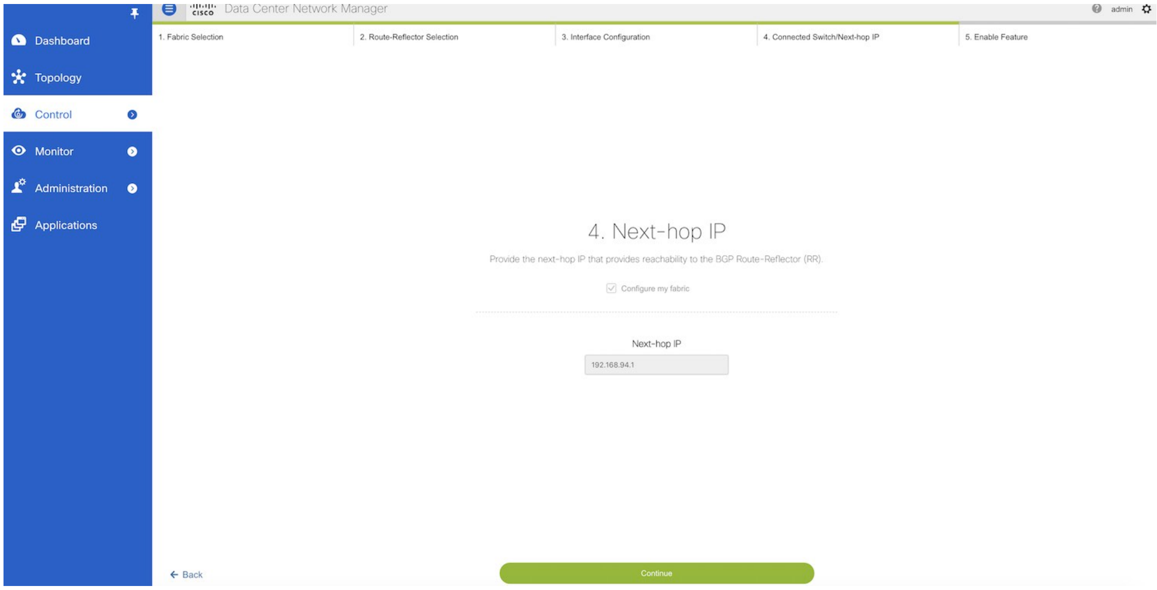

Data Center Network Manager
admin
Dashboard
Topology
Control
Monitor
Administration
Applications
1. Fabric Selection
2. Route-Reflector Selection
3. Interface Configuration
4. Connected Switch/Next-hop IP
5. Enable Feature
Back
Continue

4. Next-hop IP

Provide the next-hop IP that provides reachability to the BGP Route-Reflector (RR).

☑ Configure my fabric

Next-hop IP: 192.168.94.1

Step 7 The last step provides a summary of the information entered in the previous steps. The wizard view allows navigation to any particular step whereby one can make necessary changes/edits. In this step, one must specify whether additional information such as PORT, VLAN, VRF etc. is required when enabling the EPL feature. If the **No** option is selected, then this information will not be collected and reported by EPL.

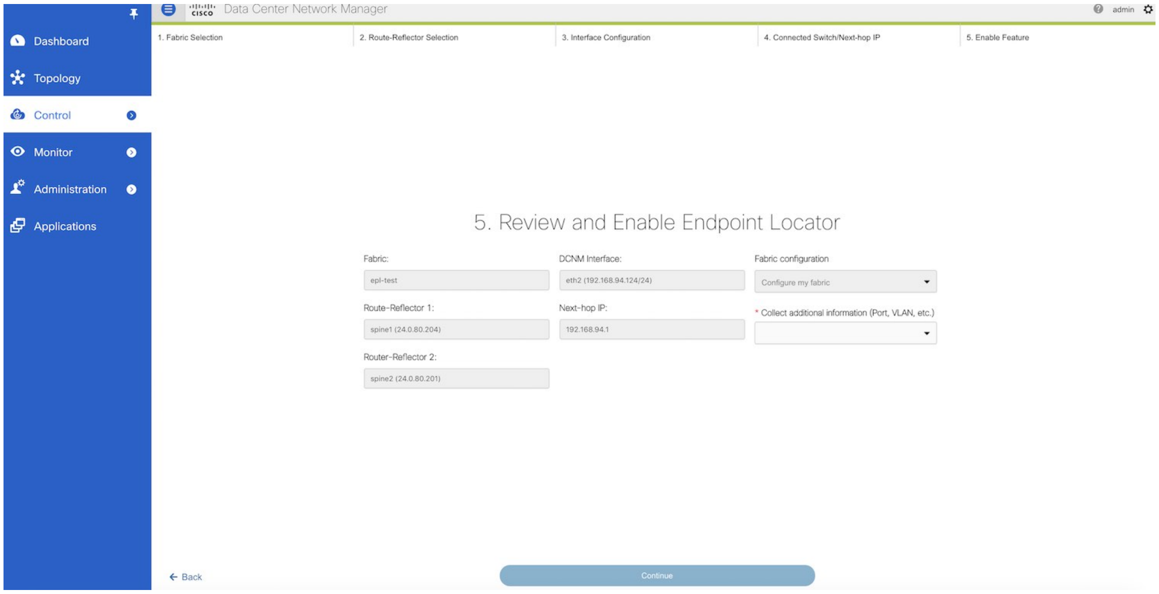
Data Center Network Manager
admin
Dashboard
Topology
Control
Monitor
Administration
Applications
1. Fabric Selection
2. Route-Reflector Selection
3. Interface Configuration
4. Connected Switch/Next-hop IP
5. Enable Feature
Back
Continue

5. Review and Enable Endpoint Locator

Fabric: epl-test

DCNM Interface: eth2 (192.168.94.124/24)

Fabric configuration: Configure my fabric

Route-Reflector 1: spine1 (24.0.80.204)

Next-hop IP: 192.168.94.1

* Collect additional information (Port, VLAN, etc.)

Router-Reflector 2: spine2 (24.0.80.201)

However, if the **Yes** option is selected in the drop down, a warning pop-up appears that feature NX-API must be supported and enabled on the switches/ToRs/leafs to gather this information. Otherwise this additional information cannot be fetched or reported.

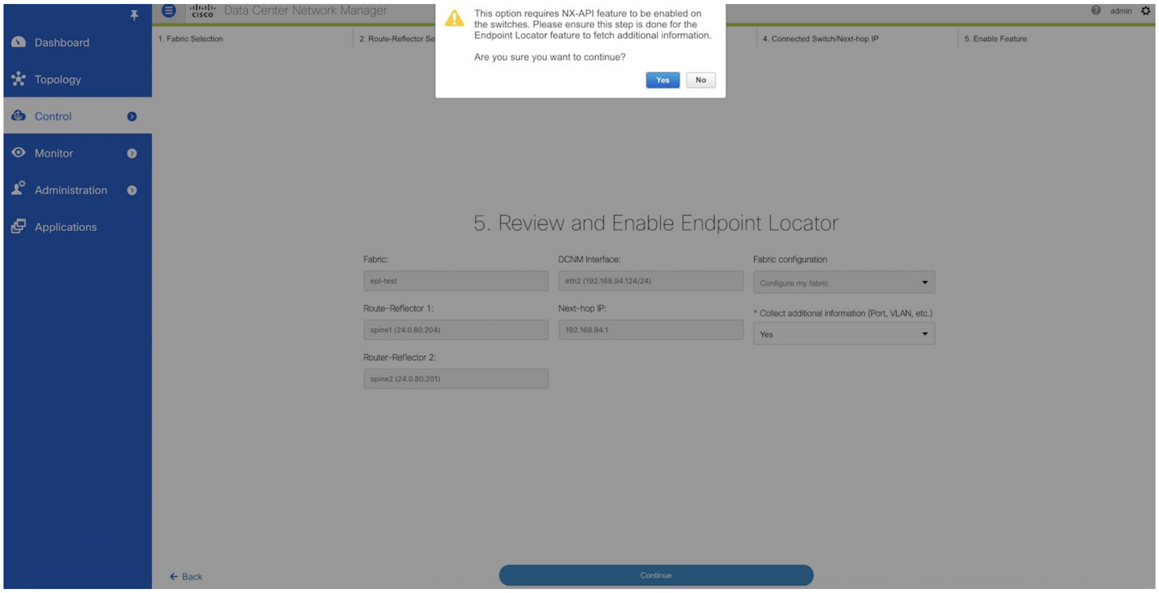

This option requires NX-API feature to be enabled on the switches. Please ensure this step is done for the Endpoint Locator feature to fetch additional information.

Are you sure you want to continue?

5. Review and Enable Endpoint Locator

Fabric: epl-test

DCNM Interface: eth2 (192.168.94.124/24)

Fabric configuration: Configure my fabric

Route-Reflector 1: spine1 (24.0.80.204)

Next-hop IP: 192.168.94.1

* Collect additional information (Port, VLAN, etc.): Yes

Router-Reflector 2: spine2 (24.0.80.201)

Step 8 Once the appropriate selections are made and various inputs have been reviewed, click **Continue** to enable EPL. A progress bar will appear indicating the status of the EPL feature enable process.

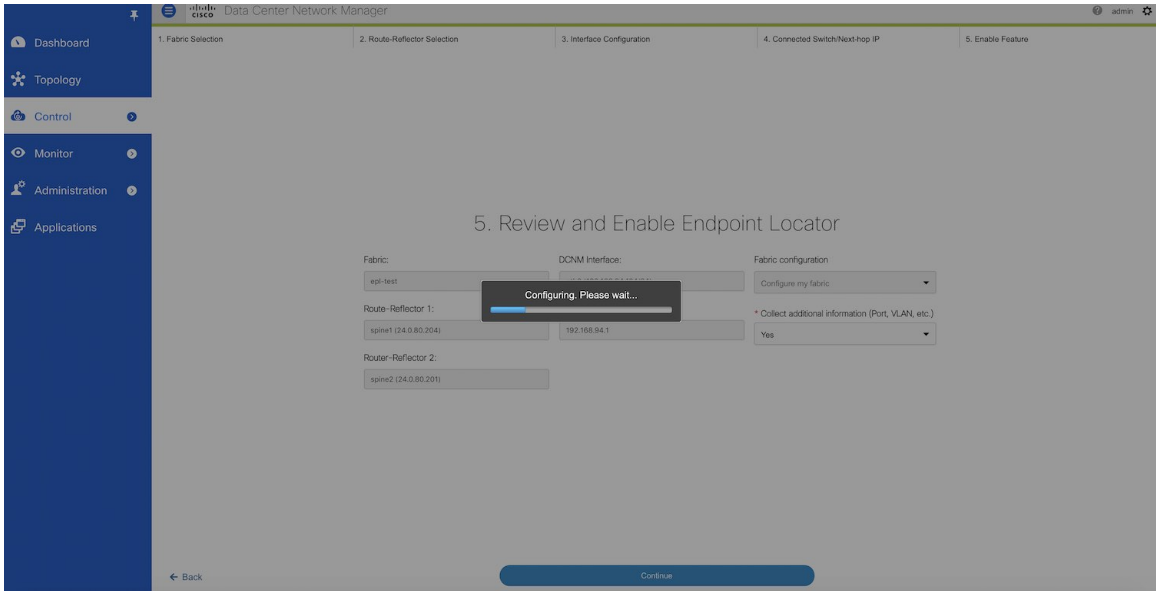
Data Center Network Manager
admin
Dashboard
Topology
Control
Monitor
Administration
Applications
1. Fabric Selection
2. Route-Reflector Selection
3. Interface Configuration
4. Connected Switch/Next-hop IP
5. Enable Feature
Back
Continue

5. Review and Enable Endpoint Locator

Fabric: epl-test

DCNM Interface:

Fabric configuration: Configure my fabric

Configuring. Please wait...

Route-Reflector 1: spine1 (24.0.80.204)

192.168.94.1

* Collect additional information (Port, VLAN, etc.): Yes

Router-Reflector 2: spine2 (24.0.80.201)

If there are any errors during the enablement, the enable process will abort and the appropriate error message will be displayed. Otherwise, EPL will be successfully enabled and on clicking **OK**, the screen will be automatically redirected to the EPL dashboard.

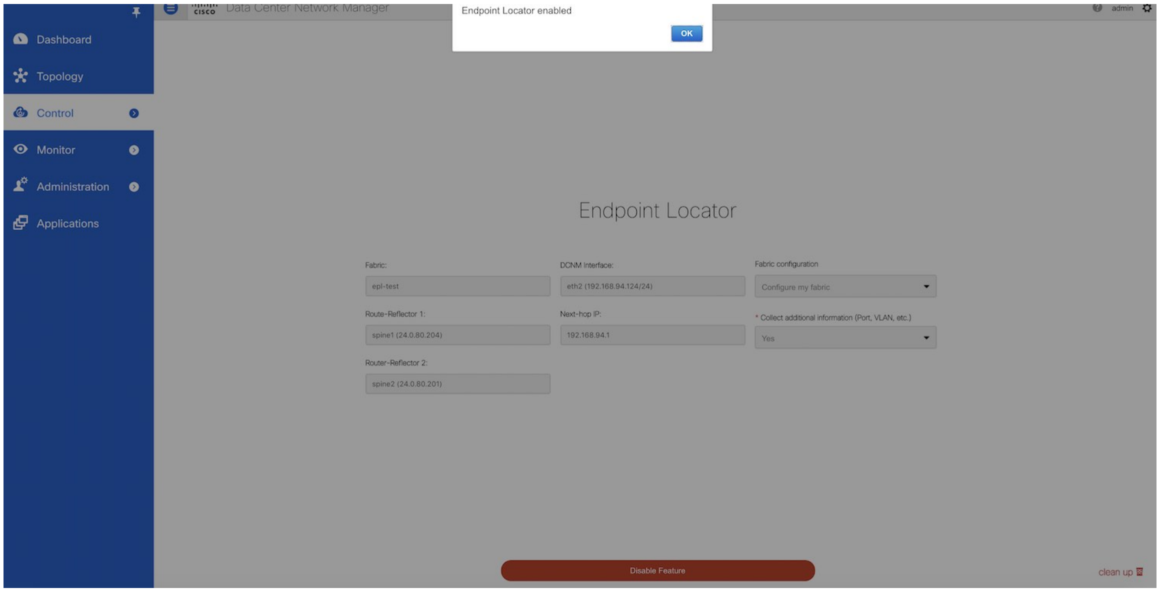

Endpoint Locator enabled

Endpoint Locator

Fabric: epl-test

DCNM Interface: eth2 (192.168.94.124/24)

Fabric configuration: Configure my fabric

Route-Reflector 1: spine1 (24.0.80.204)

Next-hop IP: 192.168.94.1

* Collect additional information (Port, VLAN, etc.): Yes

Router-Reflector 2: spine2 (24.0.80.201)

When the Endpoint Locator feature is enabled, there are a number of steps that occur in the background. DCNM will contact the selected RRs and determine the ASN, determine whether the fabric is L3VPN or EVPN enabled, and also determine the interface IP that is bound to the BGP process. Also, appropriate BGP neighbor statements are added on the RR(s), to get it ready to accept the BGP connection that will be initiated from the DCNM. The neighbor address will be the same as that of the eth2 interface shown in step 2. In order to provide reachability to the RR, a static route will be added to DCNM. This ensures that DCNM has connectivity to the RR. Once EPL is successfully enabled, the user is automatically redirected to the EPL dashboard that depicts operational and exploratory insights into the endpoints that are present in the fabric. For more information, refer to *Section Exploring Endpoint Locator Details*.

Flushing the Endpoint Database

To flush the all the Endpoint information, perform the following steps:

Procedure

Step 1 From Cisco DCNM home page, choose **Control > Endpoint Locator > Configure**, and then click the **clean up** link.

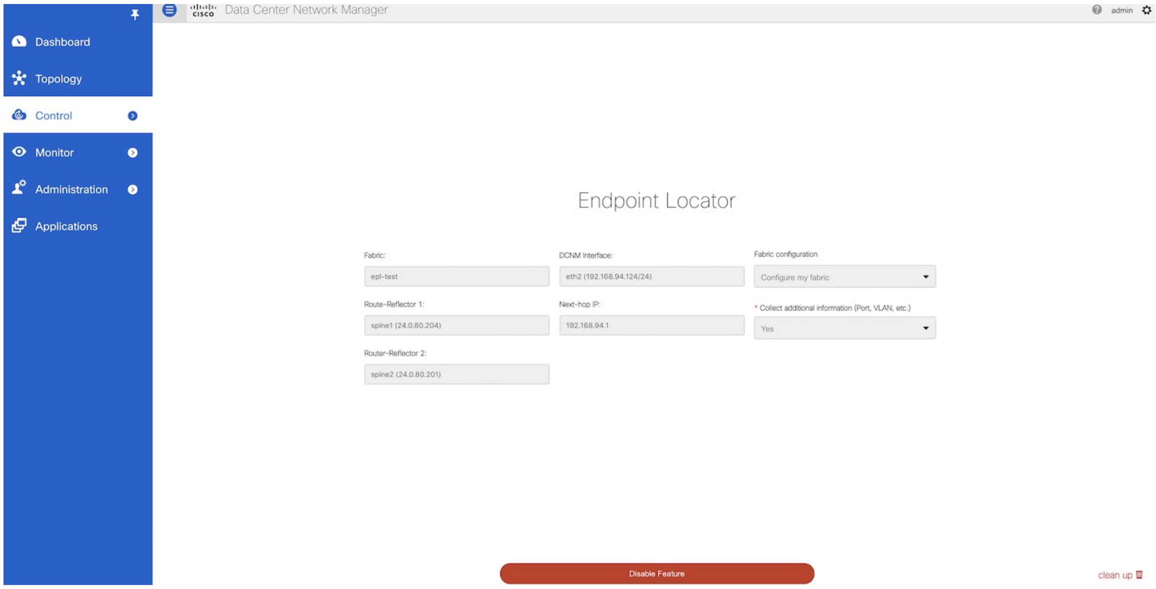

Endpoint Locator

Fabric: epl-test

DCNM Interface: eth2 (192.168.94.124/24)

Fabric configuration: Configure my fabric

Route-Reflector 1: spine1 (24.0.80.204)

Next-hop IP: 192.168.94.1

* Collect additional information (Port, VLAN, etc.): Yes

Router-Reflector 2: spine2 (24.0.80.201)

This shows a warning message indicating that all the endpoint information from the database will be flushed.

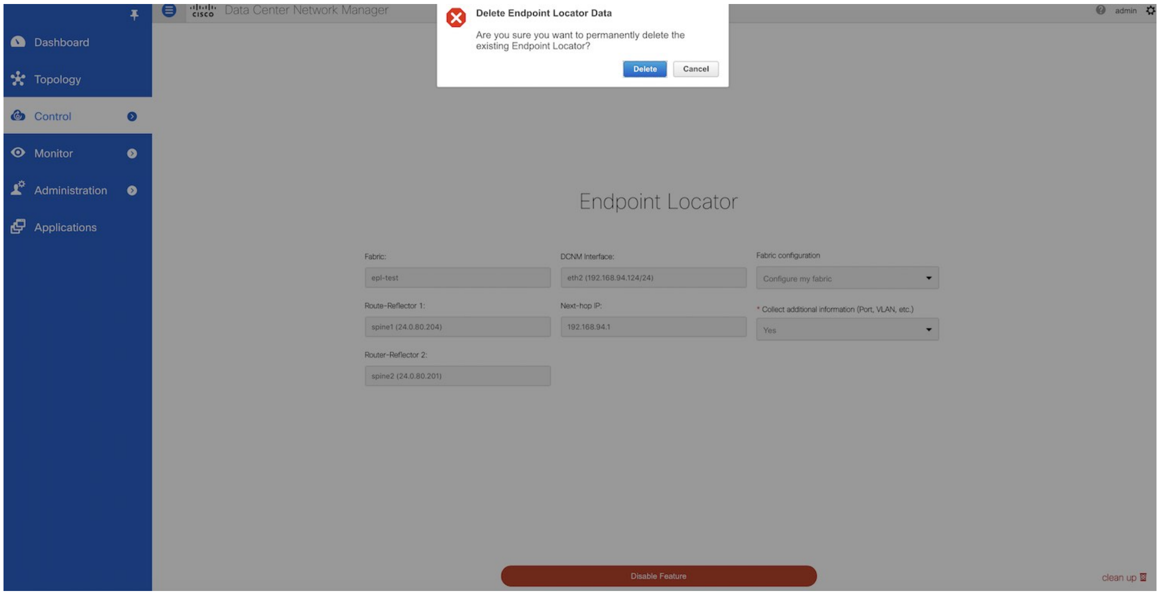

Delete Endpoint Locator Data

Are you sure you want to permanently delete the existing Endpoint Locator?

Delete Cancel

Endpoint Locator

Fabric: epl-test

DCNM Interface: eth2 (192.168.94.124/24)

Fabric configuration: Configure my fabric

Route-Reflector 1: spine1 (24.0.80.204)

Next-hop IP: 192.168.94.1

* Collect additional information (Port, VLAN, etc.): Yes

Router-Reflector 2: spine2 (24.0.80.201)

Step 2 Click **Delete** to continue or **Cancel** in case the user wants to abort.

Configuring the In-band Port

Procedure

To configure the in-band port in Cisco DCNM, enter the **appmgr setup inband** command. See the example below.

```
[root@localhost ~]# appmgr setup inband
Configuring Interface for InBand Connectivity...
Please enter the information as prompted:
InBand Physical IP [e.g. 2.2.2.69]: 192.168.94.124
InBand Network Mask [e.g. 255.255.255.0]: 255.255.255.0
InBand Gateway [e.g. 2.2.2.1]: 192.168.94.1
Validating Inputs ...
You have entered these values..
PIP=192.168.94.124
NETMASK=255.255.255.0
GATEWAY=192.168.94.1

Press 'y' to continue configuration, 'n' to discontinue [y] y
{"ResponseType":0,"Response":"Refreshed"}{"ResponseType":0,"Response":{"AfwServerEnabled":true,"AfwServerReady":true,"InbandSubnet":"192.168.94.0/24","InbandGateway":"192.168.94.1","OutbandSubnet":"192.168.126.0/24","OutbandGateway":"192.168.126.1","IsEmbeddedMode":true}}
Done.
```

Configuring Endpoint Locator in DCNM High Availability Mode

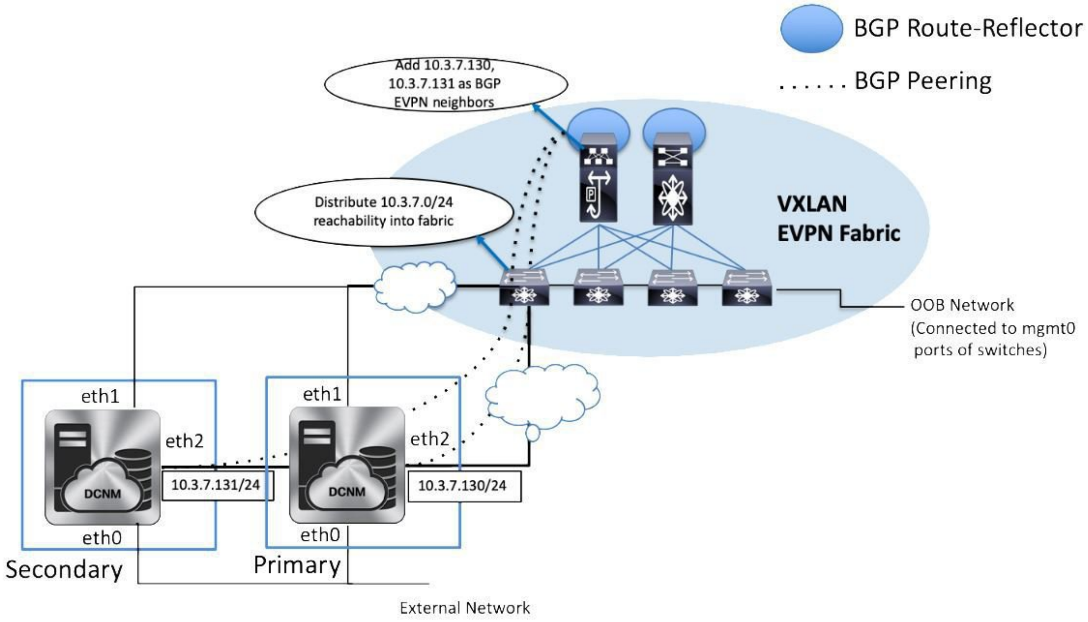
BGP Route-Reflector
...... BGP Peering
Add 10.3.7.130, 10.3.7.131 as BGP EVPN neighbors
Distribute 10.3.7.0/24 reachability into fabric
VXLAN EVPN Fabric
OOB Network (Connected to mgmt0 ports of switches)
eth1
eth2
DCNM
10.3.7.131/24
eth0
Secondary
eth1
eth2
DCNM
10.3.7.130/24
eth0
Primary
External Network

The following example shows a sample output for the **appmgr setup inband** command for a Cisco DCNM Native HA Appliance.

On Cisco DCNM Primary appliance:

```
[root@dcnm-primary]# appmgr setup inband
Configuring Interface for InBand Connectivity...
Please enter the information as prompted:
InBand Physical IP [e.g. 2.2.2.69]: 10.3.7.130
InBand Network Mask [e.g. 255.255.255.0]: 255.255.255.0
InBand Gateway [e.g. 2.2.2.1]: 10.3.7.1
InBand Virtual IP for HA setup [e.g. 2.2.2.60]: 10.3.7.254
InBand Virtual Network Mask [mandatory for HA setup] [e.g. 255.255.255.0]: 255.255.255.0
Peer Inband IP [mandatory for HA setup] [e.g. 2.2.2.59]: 10.3.7.131
Validating Inputs ...

You have entered these values..
PIP=10.3.7.130
NETMASK=255.255.255.0
GATEWAY=10.3.7.1
VIP=10.3.7.254
VIP_NETMASK=255.255.255.0
PEER_ETH2=10.3.7.131

Press 'y' to continue configuration, 'n' to discontinue [y] y

Done.
[root@dcnm-primary]#
```

On Cisco DCNM Secondary appliance:

```
[root@dcnm-secondary]# appmgr setup inband
```

```
Configuring Interface for InBand Connectivity...
Please enter the information as prompted:
InBand Physical IP [e.g. 2.2.2.69]: 10.3.7.131
InBand Network Mask [e.g. 255.255.255.0]: 255.255.255.0
InBand Gateway [e.g. 2.2.2.1]: 10.3.7.1
InBand Virtual IP for HA setup [e.g. 2.2.2.60]: 10.3.7.254
InBand Virtual Network Mask [mandatory for HA setup] [e.g. 255.255.255.0]: 255.255.255.0
Peer Inband IP [mandatory for HA setup] [e.g. 2.2.2.59]: 10.3.7.130
Validating Inputs ...

You have entered these values..
PIP=10.3.7.131
NETMASK=255.255.255.0
GATEWAY=10.3.7.1
VIP=10.3.7.254
VIP_NETMASK=255.255.255.0
PEER_ETH2=10.3.7.130

Press 'y' to continue configuration, 'n' to discontinue [y] y
HA Role is Active {"ResponseType":0,"Response":"Refreshed"}
Done.

[root@dcnm-secondary]#
```

After the in-band connectivity is established from both the Primary and Secondary nodes to the Fabric, to configure endpoint locator in DCNM HA mode from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Endpoint Locator > Configure**.

The **Endpoint Locator** window appears and the fabric configuration details are displayed.

Step 2 In the Select a fabric to configure endpoint locator in DCNM HA mode.

Step 3 Click **Continue**.

Step 4 Select one or two Route-Reflectors (RRs).

Step 5 Click **Continue**.

Step 6 Verify the Ethernet interfaces on both primary and standby DCNM nodes.

Step 7 Click **Continue**.

Step 8 Verify the next-hop IP address on the primary and standby DCNM.

Note that the next-hop IP corresponds to the eth2 gateway which should be the same on both the DCNMs.

Step 9 Click **Continue**.

Step 10 After selecting the NX-API enable or disable option and verifying the other information provided in the prior steps, click **Continue**.

What to do next

After you configure the Endpoint Locator in HA mode, you can view details such as Endpoint Activity and Endpoint History in the Endpoint dashboard. To view these details, navigate to **Monitor > Endpoint Locator > Explore**.

Adding High Availability Node to Endpoint Locator Configuration

A standalone DCNM setup can be converted into a native HA deployment at a subsequent time. If EPL is enabled on the standalone DCNM, you can enable EPL for Cisco DCNM Native HA deployment. To add a HA node to Endpoint Locator from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Endpoint Locator > Configure**.

The **Endpoint Locator** page appears and the fabric configuration details are displayed.

Step 2 Click the **Add HA node** link.

Step 3 In the **Configure Standby DCNM Interface** page, choose the Ethernet interface on DCNM that provides reachability to the BGP Route-Reflectors within the fabric.

Step 4 Click **Continue**.

Step 5 In the Next-Hop page check the value of the next-hop IP.

Step 6 Click **Configure HA Node**.

The configuration details are displayed on the Endpoint Locator page.

Configuring Endpoint Locator for External Fabrics

In addition to Easy fabrics, DCNM Release 11.2(1) allows you to enable EPL for VXLAN EVPN fabrics comprising of switches that are imported into the external fabric. The external fabric can be in managed mode or monitored mode, based on the selection of **Fabric Mode Monitor Flag** in the **External Fabric** Settings. In case the monitor or read-only fabric option is selected for the fabric, while enabling EPL, the **Configure my fabric** option must be unchecked; because, the EPL neighborship is added to the spines or RRs via some other means.

Disabling Endpoint Locator

To disable endpoint locator from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Endpoint Locator > Configure**.

The **Endpoint Locator** window appears and the fabric configuration details are displayed.

Step 2 Click **Disable Feature**.

Troubleshooting Endpoint Locator

There may be multiple reasons why enabling the Endpoint Locator feature may fail. Typically, if the appropriate devices are selected and the IP addresses to be used are correctly specified, the connectivity of the DCNM to the BGP RR may not be present due to which the feature cannot be enabled. This is a sanity check that is present to ensure that basic IP connectivity is available. The following image shows an example error scenario that was encountered during an attempt to enable the EPL feature.

The logs for EPL are located at the following location: /usr/local/cisco/dcm/fm/logs. The log that provides further details on what all occurred when the EPL feature is enabled or disabled, are present in the file epl.log. The following example provides a snapshot of the log that provides the user further information on when EPL enablement failed.

The following example helps you to understand why Endpoint Locator is unable to Connect to a switch.

```
#tail -f epl.log
2017.04.08 07:47:05 INFO [epl] Running script: [/sbin/appmgr, status, epls]
2017.04.08 07:47:05 INFO [epl] Received response:
2017.04.08 07:47:05 INFO [epl]
2017.04.08 07:47:05 INFO [epl] >>> Sat Apr 8 07:47:05 PDT 2017
2017.04.08 07:47:05 INFO [epl] appmgr status epls
2017.04.08 07:47:05 INFO [epl]
2017.04.08 07:47:05 INFO [epl] EPLS is stopped...
2017.04.08 07:47:08 INFO [epl] Running command: ifdown eth2
2017.04.08 07:47:08 INFO [epl] Received response:
2017.04.08 07:47:08 INFO [epl] EPL disabled succesfully
2017.04.08 08:00:06 INFO [epl] Enable End Point Locator
2017.04.08 08:00:13 ERROR [epl] Failed to connect to switch 192.169.6.2:java.lang.Exception:
 Authentication failed : Ssh/Telnet failed to connect with the switch
2017.04.08 08:00:13 INFO [epl] Failed to Enabled End Point Locator. Trying to removing
configuration
2017.04.08 08:00:13 INFO [epl] Disable EPL
2017.04.08 08:00:19 ERROR [epl] Failed to connect to switch 192.169.6.2:java.lang.Exception:
 Authentication failed : Ssh/Telnet failed to connect with the switch
2017.04.08 08:00:19 ERROR [epl] Failed to connect to switch: 192.169.6.2
2017.04.08 08:00:21 ERROR [epl] Failed to unconfigure BGP neighbor or failed to connect to
 switch or fabric information not provided.
2017.04.08 08:00:21 ERROR [epl] Failed to unconfigure BGP neighbor or failed to connect to
 switch or fabric information not provided.
2017.04.08 08:00:21 INFO [epl] Received response: configure terminal
Interface Ethernet1/1
no ip address
switchport
end
Enter configuration commands, one per line. End with CNTL/Z.
(config)# Interface Ethernet1/1
(config-if)# no ip address
(config-if)# switchport
(config-if)# end
# from 192.169.6.45
2017.04.08 08:00:21 ERROR [epl] Failed to disable EndPoint locator:
java.lang.NullPointerException
2017.04.08 08:00:21 INFO [epl] EPL disabled succesfully
```

In this example, the LAN credentials set in DCNM for accessing the switch are incorrect. There may be other reasons for which enablement of the EPL feature may fail. In all scenarios, an appropriate error message is displayed. You can fetch additional context information from `epl.log`.

After the EPL is enabled successfully, all the debug, error, and info logs associated with endpoint information are stored in `bgp.log`. Depending on the scale of the network and the number of endpoint events, the file size will increase. Therefore, there is a restriction on the maximum number and size of `bgp.log`. Up to 10 such files will be stored with each file size of maximum of 10MB.

The EPL relies on BGP updates to get endpoint information. In order for this to work, the switch loopback or VTEP interface IP addresses must be discovered on the DCNM for all switches that have endpoints. To validate, navigate to the Cisco DCNM **Web UI > Dashboard > Switch > Interfaces** tab, and verify if the IP address and the prefix associated with the corresponding Layer-3 interfaces (typically loopbacks) are displayed correctly.

In a Cisco DCNM Cluster deployment, if EPL cannot establish BGP peering and the active DCNM is able to ping the loopback IP address of the spine, while the EPL container cannot, it implies that the eth2 port group for Cisco DCNM and its computes does not have Promiscuous mode set to **Accept**. After changing this setting, the container can ping the spine and EPL will establish BGP.

In a large-scale setup, it may take more than 30 seconds (default timer set in Cisco DCNM) to get this information from the switch. If this occurs, the `ssh.read-wait-timeout` property (in the **Administration > DCNM Server > Server Properties**) must be changed from 30000 (default) to 60000 or a higher value.

Streaming Telemetry for LAN Deployments

In today's data center environments, granular visibility and tracking of network events has become critical. The traditional polling-based methods that pull the network state in predefined intervals need a fork-lift upgrade. More advanced streaming approaches are required that provide network event visibility in closer to real time through a push method. Streaming telemetry not only allows data to be pushed out at a much finer granularity with a lower cadence (shorter interval) but it also enables event-based notifications. While getting relevant data in a timely fashion is highly desirable, the data needs to be analyzed and converted into actionable insights.

As a first step toward LAN analytics, DCNM 11.0(1) enables subscriptions for environmental metrics through streaming telemetry for consumption and analysis. The environmental metrics that are streamed include CPU, Memory, Power, Temperature, and Fan Speed; all these are enabled with a single click. DCNM allows you to configure the streaming interval for these metrics. The default streaming interval for CPU, Memory is set to 30 seconds, and those for Power, Temperature, and Fan Speed is set to 300 seconds (5 minutes).

The per-metric real-time streaming dashboards allow filtering on a per fabric and per switch level including a per-switch drill-down where applicable. Streaming telemetry is currently supported on the Nexus 9000 platforms.

Guidelines and Recommendations

- In a cluster mode, a minimum of three compute nodes have to be up for LAN Telemetry to start properly. However, LAN Telemetry functions properly if any one of the three compute nodes is intermittently down.
- If two compute nodes go down, both nodes have to be restored for Zoo Keeper and Kafka Connect to bootstrap correctly and resume data transmission.
- We recommend using the LAN Telemetry feature for up to 30 switches.
- The LAN Telemetry feature is not supported on switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images.

Pre-Requisites for Enabling the LAN Telemetry Feature

- The Cisco Nexus 9000 switches and Cisco DCNM need to be time synchronized (NTP is recommended).
- Minimum software version on the Nexus 9000 switches must be 7.0(3)I6(1) or higher.

- In the LAN Classic mode, you need to manually enable the following configurations on all the switches before enabling telemetry:
    - **feature nxapi**
    - **nxapi http port 80**

Note If the preceding configurations are unavailable on the switches, the telemetry health on Cisco DCNM does not show the configurations and the connection status for the telemetry-enabled switches. The preceding commands can be manually defined in a new template, and then pushed to all the switches in the fabric from Cisco DCNM. Use an unused port (for example, port 80) configure nxapi.

Enabling the Streaming Telemetry Feature

Procedure

Step 1 Choose **Control** > **LAN Telemetry** > **Configure**. Select the fabric for which LAN Telemetry has to be enabled. Then press the **Enable** button.

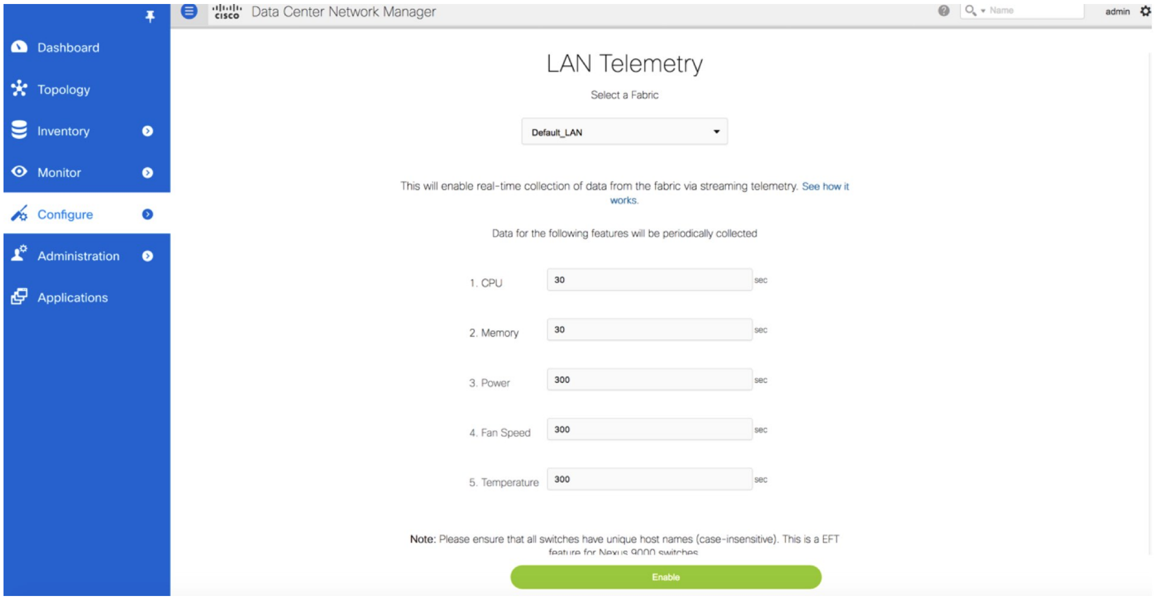

LAN Telemetry

Select a Fabric

Default_LAN

This will enable real-time collection of data from the fabric via streaming telemetry. See how it works.

Data for the following features will be periodically collected

1. CPU 30 sec

2. Memory 30 sec

3. Power 300 sec

4. Fan Speed 300 sec

5. Temperature 300 sec

Note: Please ensure that all switches have unique host names (case-insensitive). This is a EFT feature for Nexus 9000 switches

A warning message appears to indicate that the Cisco DCNM and switches need to be time-synchronized before this feature is enabled. Recall, that this is a prerequisite for this feature. If the prerequisite is met, acknowledge by clicking **Yes**.

Note When Telemetry is enabled, the NTP configuration is done on the switches for LAN Classic deployment, wherein the NTP server address is set to DCNM's out-of-band interface's IPv4 address. In case of HA setups, the NTP server address is set to the VIP address of the out-of-band interface.

Ensure that the NTP configurations are not removed/modified from the switches.

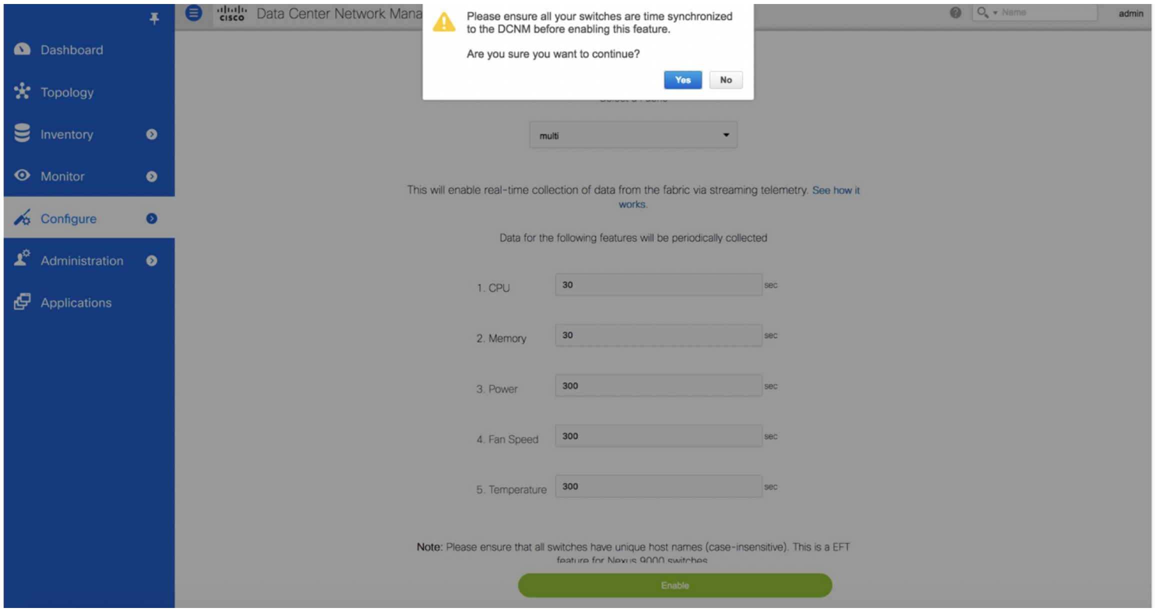

Please ensure all your switches are time synchronized to the DCNM before enabling this feature.

Are you sure you want to continue?

multi

This will enable real-time collection of data from the fabric via streaming telemetry. See how it works.

Data for the following features will be periodically collected

1. CPU 30 sec
2. Memory 30 sec
3. Power 300 sec
4. Fan Speed 300 sec
5. Temperature 300 sec

Note: Please ensure that all switches have unique host names (case-insensitive). This is a EFT feature for Nexus 9000 switches

Step 2 Once this feature is enabled, a message appears indicating the initialization process has begun, which takes a couple of minutes. This time is needed for the streaming configuration to be pushed to the switches. The initial data to be streamed out from the switches, which are consumed by DCNM, and depicted on the LAN telemetry dashboard.

Once the LAN telemetry preview feature is enabled, DCNM updates the switch telemetry configuration for the environmental metrics. Every switch that does not conform to the telemetry requirements (must be Cisco Nexus 9000) is excluded from the configuration update. The status of the switch configuration can be monitored by choosing **Control > LAN Telemetry > Health**.

Once the jobs are successfully executed, the required telemetry configuration has been applied to the switches and the streaming data appears once received and processed.

LAN Telemetry Health

The LAN Telemetry Health window provides a detailed break-down of how much data is streamed out by each switch per feature for the last 24 hours. This window shows the status of the configuration for every switch, apart from showing the statistics of the received data for every metric from every switch.

When the connection status is shown as **Disconnected** the port configuration may not be accepted by the switch correctly. On the switch image 7.0(3)I7(6), if a switch already had **nxapi** configuration, and later it was managed by DCNM and telemetry was enabled on that fabric, DCNM pushes **http port 80** configuration so that it could query some NXAPI commands such as **show telemetry transport** and **show telemetry data collector details**, to monitor telemetry connection statistics. In this case, the switch does not update **http port**

80 in its configuration even though the command was executed correctly. In such a scenario, issue the following commands on the switch:

```
switch# configure
switch(config)# no feature nxapi
switch(config)# feature nxapi
switch(config)# http port 80
```

Note You cannot configure ICAM on the Cisco Nexus 9000 Series Switches Release 7.0(3)I7(6), and therefore, the telemetry will fail until the switch issue is resolved.

To view the LAN Telemetry Health, perform the following steps:

Procedure

Step 1 Choose **Control > LAN Telemetry > Health**.

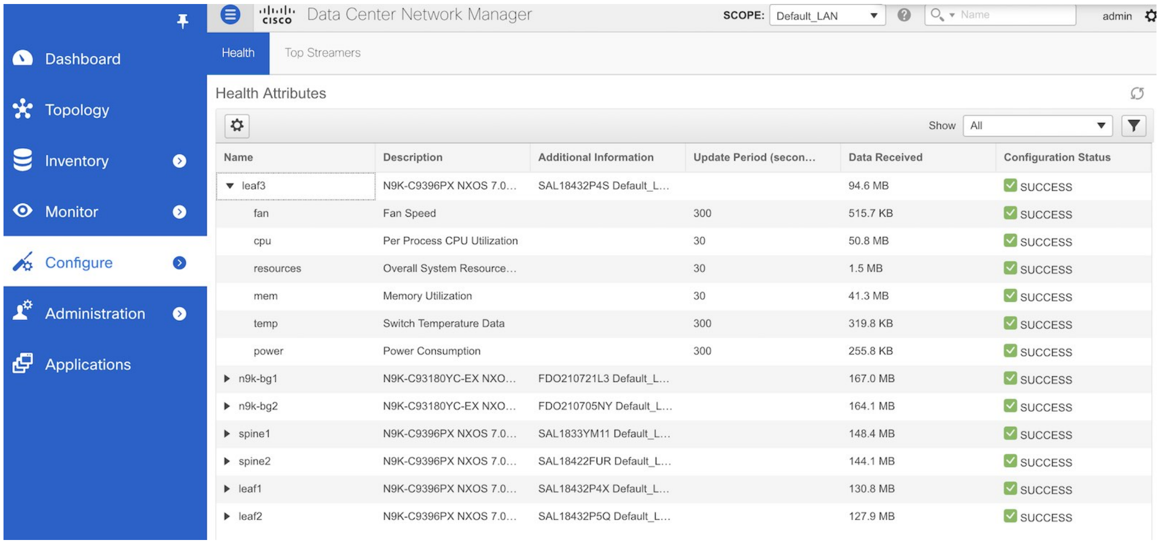

Health Attributes

Show All

Name	Description	Additional Information	Update Period (secon...	Data Received	Configuration Status
▼ leaf3	N9K-C9396PX NXOS 7.0...	SAL18432P4S Default_L...		94.6 MB	✓ SUCCESS
fan	Fan Speed		300	515.7 KB	✓ SUCCESS
cpu	Per Process CPU Utilization		30	50.8 MB	✓ SUCCESS
resources	Overall System Resource...		30	1.5 MB	✓ SUCCESS
mem	Memory Utilization		30	41.3 MB	✓ SUCCESS
temp	Switch Temperature Data		300	319.8 KB	✓ SUCCESS
power	Power Consumption		300	255.8 KB	✓ SUCCESS
▶ n9k-bg1	N9K-C93180YC-EX NXO...	FDO210721L3 Default_L...		167.0 MB	✓ SUCCESS
▶ n9k-bg2	N9K-C93180YC-EX NXO...	FDO210705NY Default_L...		164.1 MB	✓ SUCCESS
▶ spine1	N9K-C9396PX NXOS 7.0...	SAL1833YM11 Default_L...		148.4 MB	✓ SUCCESS
▶ spine2	N9K-C9396PX NXOS 7.0...	SAL18422FUR Default_L...		144.1 MB	✓ SUCCESS
▶ leaf1	N9K-C9396PX NXOS 7.0...	SAL18432P4X Default_L...		130.8 MB	✓ SUCCESS
▶ leaf2	N9K-C9396PX NXOS 7.0...	SAL18432P5Q Default_L...		127.9 MB	✓ SUCCESS

When Telemetry is enabled or disabled, there is a chance that enabling or disabling can fail in some or all the switches. When that happens, a pop-up similar to the following screen appears.

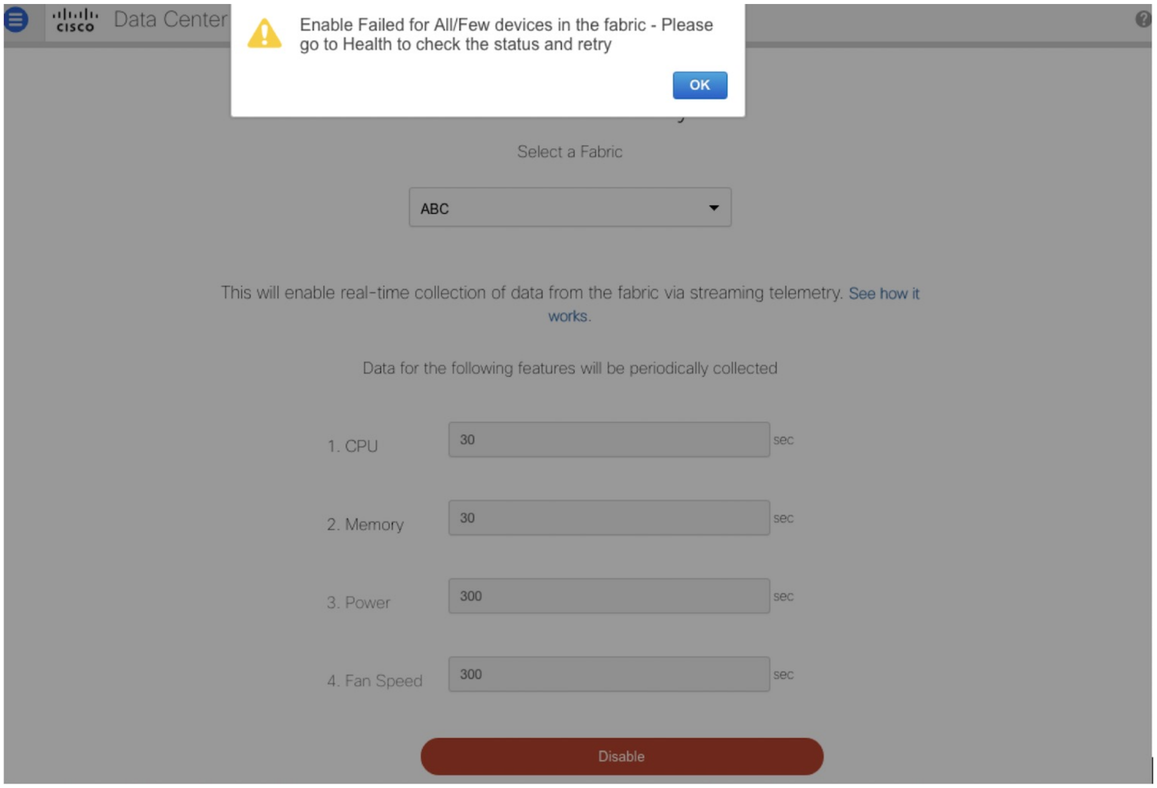

Enable Failed for All/Few devices in the fabric - Please go to Health to check the status and retry

OK

Select a Fabric

ABC

This will enable real-time collection of data from the fabric via streaming telemetry. See how it works.

Data for the following features will be periodically collected

1. CPU 30 sec

2. Memory 30 sec

3. Power 300 sec

4. Fan Speed 300 sec

There are two possible options:

1. You can go to the Health page, and retry the configuration for those switches that failed. When a configuration cannot be applied or removed on any switch, **Configuration Status** in the health page, appears as FAILED. Upon clicking the ‘FAILED’ link, a pop-up would show the reason for the failure. After you correct the failure, the configuration can be retried by clicking on the retry button appearing next to the Configuration Status for every switch. The screen-shot for that is also shown below.

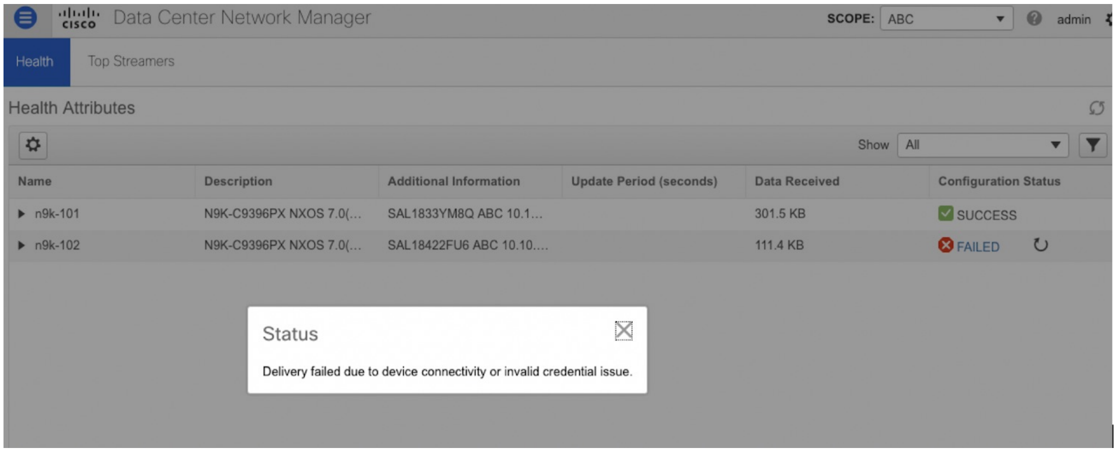

Name	Description	Additional Information	Update Period (seconds)	Data Received	Configuration Status
n9k-101	N9K-C9396PX NXOS 7.0(...	SAL1833YM8Q ABC 10.1...		301.5 KB	✓ SUCCESS
n9k-102	N9K-C9396PX NXOS 7.0(...	SAL18422FU6 ABC 10.10....		111.4 KB	☒ FAILED

Status

Delivery failed due to device connectivity or invalid credential issue.

2. You can stay in the main **Telemetry > Configure** page. It would display a dialogue box with the failed message. Then you can reverse the configuration for the successfully configured switches. In other words:

   - When “Enable” fails for some or all switches, the screen has a Red button with “disable” option. This means, for those switches, wherein enabling Telemetry was successful, you can disable Telemetry on those switches. If "Enable" failed on all switches, you will still see the Red button with "disable" option. Clicking on "disable" is a no-op. In both the cases, you will see the green button with the "enable" option in a few seconds after disabling is completed. This removes the "retry" option from the health page since you want to "disable" telemetry and there is nothing to retry.

3. Similarly, when "Disable" fails for some or all switches, the screen has a Green button with "Enable" option. This means, for those switches, wherein disabling Telemetry was successful, you can Enable Telemetry on those switches. If "Disable" failed on all switches, you will still see the Green button with "Enable" option. Clicking on "Enable" is a no-op. In both the cases, you will see the Red button with the "Disable" option in a few seconds after Enabling is completed. Doing this, removes the "retry" option from the health page since you want to "enable" telemetry and there's nothing to retry.

Step 2 Click the **Top Streamers** tab to view the graphs that depicts the top five streaming switches and has a drill-down capability for a feature-wise break-down.

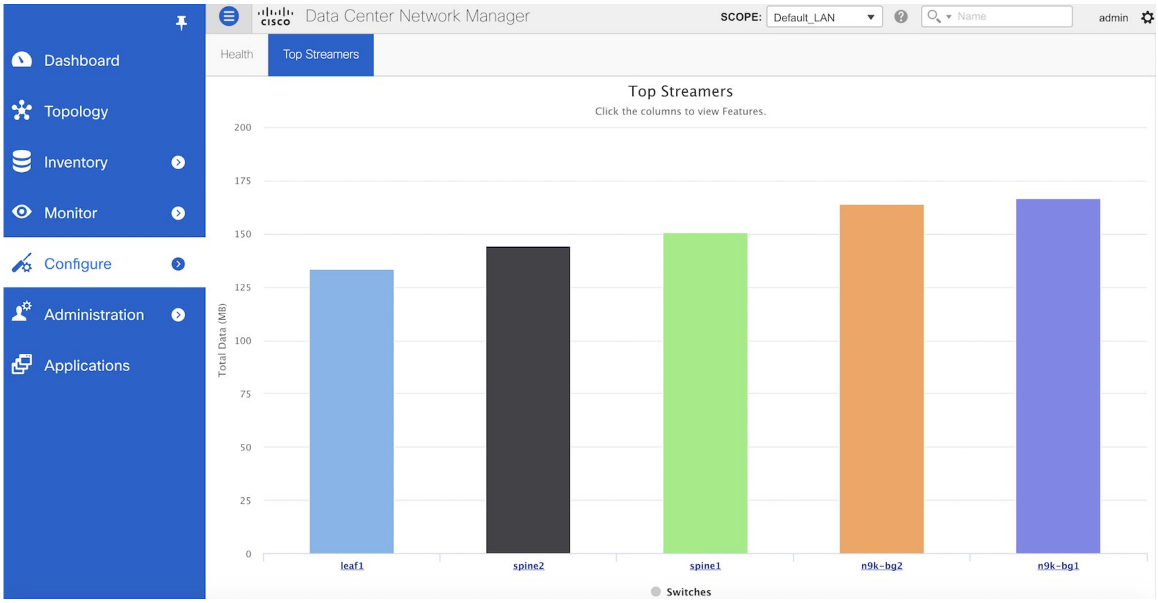

Telemetry Streaming Interface

Telemetry data, by default is streamed through the management interface of the switches to the Cisco DCNM. This is the Out-of-Band network. This is a global configuration for all fabrics or switch-groups in DCNM. The switches can also stream the Telemetry data through their front panel ports to DCNM assuming there's connectivity from the switches to the DCNM. This is the In-band network. To use the in-band network, do the following:

Procedure

Step 1 Disable Telemetry on all the Enabled fabrics.

Step 2 Go to the Health window and change the settings by clicking on the gear icon on the Heath window. In the Telemetry Settings window that comes up, select **`In-Band`** from the Interface drop-down list. The VRF option is set to default. Click Submit.

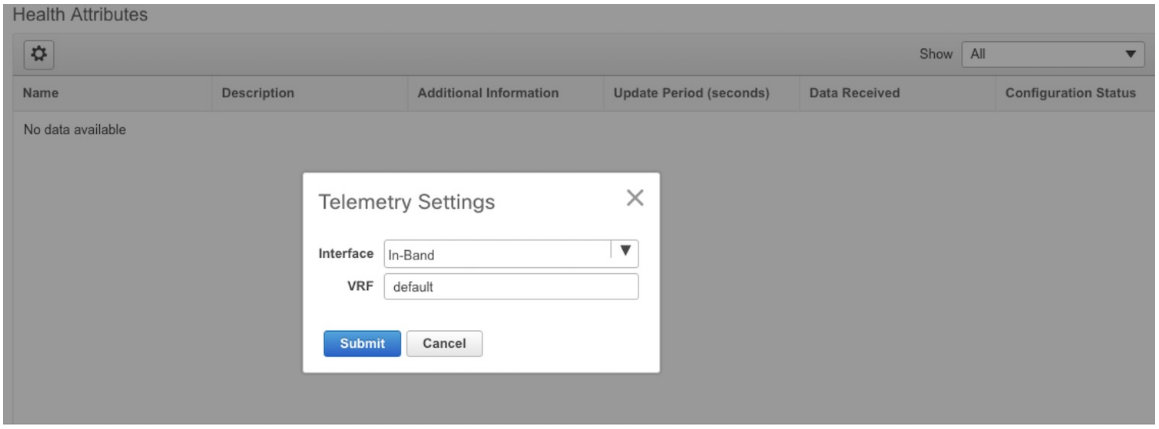

Health Attributes

Show: All

Name	Description	Additional Information	Update Period (seconds)	Data Received	Configuration Status
No data available					

Telemetry Settings

Interface: In-Band

VRF: default

The VRF option is set to default. You can also change this to any other VRF assuming that the switches are configured with the VRF and there is availability to the in-band interface of the DCNM through that VRF.

Note If Telemetry is already enabled for some fabrics, you should first disable Telemetry on all the enabled fabrics and only then modify the Telemetry network setting. After modifying the Telemetry network settings, you can enable Telemetry on the fabrics. Now, Telemetry data start coming through the in-band interface.