Control

Note

Click 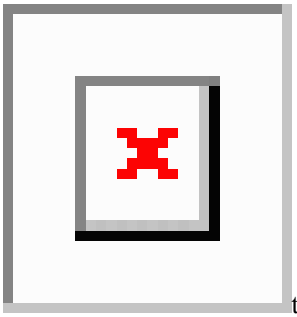 to view the contents of this chapter.

This chapter contains the following topics:

Fabrics

The following terms are referred to in the document:

- Greenfield Deployments: Applicable for provisioning new VXLAN EVPN fabrics.
- Brownfield Deployments: Applicable for existing VXLAN EVPN fabrics:
  - Migrate NFM-Managed VXLAN EVPN Fabrics to DCNM.
- Upgrades: Applicable for VXLAN EVPN fabrics created with previous DCNM versions
  - Upgrade for VXLAN fabrics built with DCNM 11.0(1) to DCNM 11.2(1).
  - Upgrade for VXLAN fabrics built with DCNM 11.1(1) to DCNM 11.2(1).

This section contains the following topics:

VXLAN BGP EVPN Fabrics Provisioning

In DCNM 11.0(1), fabric creation is enhanced to provision VXLAN BGP EVPN underlay network parameters to the fabric switches. The concept of Multi-Site Domain (MSD) fabrics was introduced.

In the DCNM 11.1(1) and 11.2(1) releases, further enhancements are made. For the LAN Fabric deployment type, fabric template support is introduced for Cisco Nexus 3000 Series switches, in addition to the existing support for Cisco Nexus 9000 Series switches.

Support of simplified CLIs for VXLAN EVPN fabrics is not supported in either greenfield or brownfield deployments.

The DCNM GUI functions for creating, deploying, and migrating VXLAN fabrics are as follows

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

Create, edit, and delete a fabric:

- Create new VXLAN, MSD and external VXLAN fabrics.
- View the VXLAN and MSD fabric topologies, including connections between fabrics.
- Update fabric settings.
- Save and deploy updated changes.
- Delete a fabric (if devices are removed).

Fabric Membership changes

- Transition existing VXLAN fabric management to DCNM (through the Preserve Config = Yes option).
- Deploy new fabrics or add new devices to an existing fabric (through the bootstrap or Preserve Config = No options).
- Move fabrics into or out of an MSD.

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

- Add switch instances to the fabric.
- Provision start-up configurations and an IP address to a new switch through POAP configuration.
- Update switch policies, save and deploy updated changes.
- Create intra-fabric and inter-fabric links (also called Inter-Fabric Connections [IFCs]).

Transitioning VXLAN fabric management to DCNM

In DCNM 11.1(1) release, transitioning existing VXLAN fabric management to DCNM is introduced.

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 subinterface.

- 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.

Note vPC support is added for BGWs in the DCNM 11.1(1) release.

Control > Networks and **Control > VRFs** menu options (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.

This chapter mostly covers standalone fabric-related configurations. MSD fabric documentation is available in a separate chapter. The deployment of networks and VRFs is covered under the Creating and Deploying Networks and VRFs section. Step by step configuration:

Guidelines for VXLAN BGP EVPN Fabrics Provisioning

- For any switch to be successfully imported into DCNM, the user defined on the switch via local or remote AAA, and used for import into DCNM should have the following permissions:
  - SSH access to the switch
  - Ability to perform SNMPv3 queries
  - Ability to run **show** commands
  - Ability to execute the **guestshell** commands, which are prefixed by **run guestshell** for the DCNM tracker
- When an invalid command is deployed by DCNM to a device, for example, a command with an invalid key chain due to an invalid entry in the fabric settings, an error is generated displaying this issue. This error is not cleared after correcting the invalid fabric entry. You need to manually cleanup or delete the invalid commands to clear the error.

  Note that the fabric errors related to the command execution are automatically cleared only when the same failed command succeeds in the subsequent deployment.
- When LAN credentials are not set for a device, DCNM moves this device to the maintenance mode. However, DCNM also displays a pop-up message saying that this device is not set to the maintenance mode. Ignore this message because the switch will be in the maintenance mode as seen in the **Topology** view.
- Persistent configuration diff is seen for the command line: **system nve infra-vlan** *int* **force** . The persistent diff occurs if you have deployed this command via the freeform configuration to the switch. Although

the switch requires the **force** keyword during deployment, the running configuration that is obtained from the switch in DCNM does not display the **force** keyword. Therefore, the **system nve infra-vlan** *int* **force** command always shows up as a diff.

The intent in DCNM contains the line:

```
system nve infra-vlan int force
```

The running config contains the line:

```
system nve infra-vlan int
```

Note that the switch does not display the **force** keyword as being applied. However, the **force** keyword is required by the switch to be deployed.

As a workaround to fix the persistent diff, edit the freeform config to remove the **force** keyword after the first deployment such that it is **system nve infra-vlan** *int* **force**.

The **force** keyword is required for the initial deploy and must be removed after a successful deploy. You can confirm the diff by using the **Side-by-side Comparison** tab in the **Config Preview** window.

The persistent diff is also seen after a write erase and reload of a switch. Update the intent on DCNM to include the **force** keyword, and then you need to remove the **force** keyword after the first deployment.

- The **Save & Deploy** button triggers the intent regeneration for the entire fabric as well as a configuration compliance check for all the switches within the fabric. This button is required but not limited to the following cases:
  - A switch or a link is added, or any change in the topology
  - A change in the fabric settings that must be shared across the fabric
  - A switch is removed or deleted
  - A new vPC pairing or unpairing is done
  - A change in the role for a device

  When you click **Save & Deploy**, the changes in the fabric are evaluated, and the configuration for the entire fabric is generated. You can preview the generated configuration, and then deploy it at a fabric level. Therefore, **Save & Deploy** can take more time depending on the size of the fabric.

  When you right-click on a switch icon, you can use the **Deploy Config** option to deploy per switch configurations. This option is a local operation for a switch, that is, the expected configuration or intent for a switch is evaluated against it's current running configuration, and a config compliance check is performed for the switch to get the **In-Sync** or **Out-of-Sync** status. If the switch is out of sync, the user is provided with a preview of all the configurations running in that particular switch that vary from the intent defined by the user for that respective switch.

  Note that the fabric builder does not re-evaluate the topology or generate any dependent configuration for that switch or any other devices that are part of the fabric.

- When the switch contains the **hardware access-list tcam region arp-ether 256** command, which is deprecated without the **double-wide** keyword, the below warning is displayed:

  WARNING: Configuring the arp-ether region without "double-wide" is deprecated and can result in silent non-vxlan packet drops. Use the "double-wide" keyword when carving TCAM space for the arp-ether region.

Since the original **hardware access-list tcam region arp-ether 256** command does not match the policies in DCNM, this config is captured in the **switch_freeform** policy. After the **hardware access-list tcam region arp-ether 256 double-wide** command is pushed to the switch, the original **tcam** command that does not contain the **double-wide** keyword is removed.

You must manually remove the **hardware access-list tcam region arp-ether 256** command from the **switch_freeform** policy. Otherwise, config compliance shows a persistent diff.

Here is an example of the **hardware access-list** command on the switch:

```
switch(config)# show run | inc arp-ether
switch(config)# hardware access-list tcam region arp-ether 256
Warning: Please save config and reload the system for the configuration to take effect
switch(config)# show run | inc arp-ether
hardware access-list tcam region arp-ether 256
switch(config)#
switch(config)# hardware access-list tcam region arp-ether 256 double-wide
Warning: Please save config and reload the system for the configuration to take effect
switch(config)# show run | inc arp-ether
hardware access-list tcam region arp-ether 256 double-wide
```

You can see that the original **tcam** command is overwritten.

Creating a New VXLAN BGP EVPN Fabric

This procedure shows how to create a new VXLAN BGP EVPN fabric.

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

   The **Fabric Builder** screen 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** screen, wherein a rectangular box represents each fabric.

   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** screen appears.

   The fields are explained:

   Fabric Name - Enter the name of the fabric.

   Fabric Template - From the drop-down menu, choose the **Easy_Fabric_11_1** fabric template. The fabric settings for creating a standalone fabric comes up.

   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 Multi-Site Domain for VXLAN BGP EVPN Fabrics topic before member fabric creation.

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

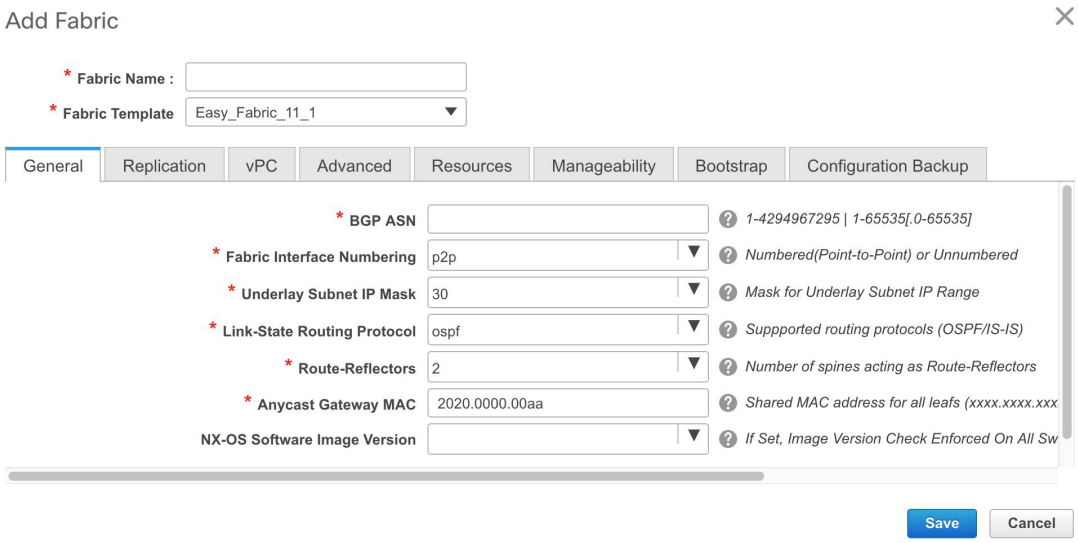

BGP ASN: Enter the BGP AS number the fabric is associated with.

Fabric Interface Numbering : Specifies whether you want to use point-to-point (**p2p**) or unnumbered networks.

Underlay Subnet IP Mask - Specifies the subnet mask for the fabric interface IP addresses.

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

Route-Reflectors – The number of spine switches that are used as route reflectors for transporting BGP traffic. Choose 2 or 4 from the drop down box. The default value is 2.

To deploy spine devices as RRs, DCNM sorts the spine devices based on their serial numbers, and designates two or four spine devices as RRs. If you add more spine devices, existing RR configuration will not change.

Increasing the count - You can increase the route reflectors from two to four at any point in time. Configurations are automatically generated on the other 2 spine devices designated as RRs.

Decreasing the count - When you reduce four route reflectors to two, you must remove the unneeded route reflector devices from the fabric. Follow these steps to reduce the count from 4 to 2.

a. Change the value in the drop-down box to 2.

b. Identify the spine switches designated as route reflectors.

An instance of the **rr_state** policy is applied on the spine switch if it is a route reflector. To find out if the policy is applied on the switch, right-click the switch, and choose **View/edit policies**. In the View/Edit Policies screen, search **rr_state** in the **Template** field. It is displayed on the screen.

c. Delete the unneeded spine devices from the fabric (right-click the spine switch icon and choose **Discovery** > **Remove from fabric**).

If you delete existing RR devices, the next available spine switch is selected as the replacement RR.

d. Click Save and Deploy at the top right part of the fabric topology screen.

You can preselect RRs and RPs before performing the first **Save & Deploy** operation. For more information, see *Preselecting Switches as Route-Reflectors and Rendezvous-Points.*.

Anycast Gateway MAC : Specifies the 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 should 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. Till all devices run the specified image, the deployment process will be 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 **Replication** tab. Most of the fields are auto generated. You can update the fields if needed.

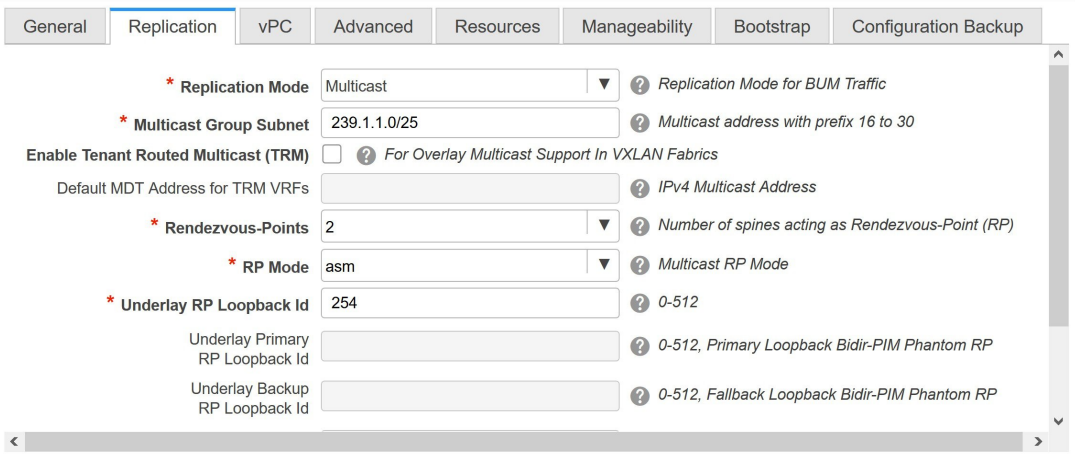

Replication Mode : The mode of replication that is used in the fabric, Ingress Replication, or Multicast.

When you choose Ingress replication, the multicast replication fields get disabled.

You can change the fabric setting from one mode to the other, if no overlay profile exists for the fabric.

Multicast Group Subnet : IP address prefix used for multicast communication. An unique IP address is allocated from this group for each overlay network.

In the DCNM 11.0(1) release, the replication mode change is not allowed if a policy template instance is created for the current mode. For example, if a multicast related policy is created and deployed, you cannot change the mode to Ingress.

Enable Tenant Routed Multicast (TRM) – Select the checkbox to enable Tenant Routed Multicast (TRM) as the fabric overlay multicast protocol.

Default MDT Address for TRM VRFs: The multicast address for Tenant Routed Multicast traffic is populated. By default, this address is from the IP prefix specified in the **Multicast Group Subnet** field. When you update either field, ensure that the TRM address is chosen from the IP prefix specified in **Multicast Group Subnet**.

Rendezvous-Points - Enter the number of spine switches acting as rendezvous points.

RP mode – Choose from the two supported multicast modes of replication, ASM (for Any-Source Multicast [ASM]) or BiDir (for Bidirectional PIM [BIDIR-PIM]).

When you choose ASM, the BiDir related fields are not enabled. When you choose BiDir, the BiDir related fields are enabled.

Note BIDIR-PIM is supported on Cisco's Cloud Scale Family platforms 9300-EX and 9300-FX/FX2, and software release 9.2(1) onwards.

When you create a new VRF for the fabric overlay, this address is populated in the **Underlay Multicast Address** field, in the **Advanced** tab.

Underlay RP Loopback ID – The loopback ID used for the rendezvous point (RP), for multicast protocol peering purposes in the fabric underlay.

The next two fields are enabled if you choose BIDIR-PIM as the multicast mode of replication.

Underlay Primary RP Loopback ID – The primary loopback ID used for the phantom RP, for multicast protocol peering purposes in the fabric underlay.

Underlay Backup RP Loopback ID – The secondary loopback ID used for the phantom RP, for multicast protocol peering purposes in the fabric underlay.

Underlay Second Backup RP Loopback Id and **Underlay Third Backup RP Loopback Id**: Used for the second and third fallback Bidir-PIM Phantom RP.

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

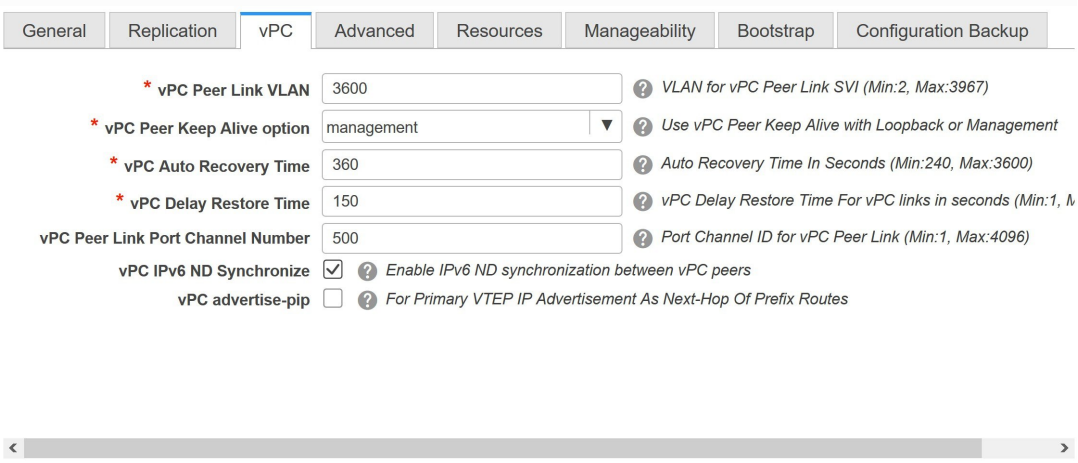

vPC Peer Link VLAN – VLAN used for the vPC peer link SVI.

vPC Peer Keep Alive option – Choose the management or loopback option. If you want to use IP addresses assigned to the management port and the management VRF, choose management. If you use IP addresses assigned to loopback interfaces (and a non-management VRF), choose loopback.

If you use IPv6 addresses, you must use loopback IDs.

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.

vPC Peer Link Port Channel Number - Specifies the Port Channel ID for a vPC Peer Link. By default, the value in this field is 500.

vPC IPv6 ND Synchronize – Enables IPv6 Neighbor Discovery synchronization between vPC switches. The check box is enabled by default. Clear the check box to disable the function.

vPC advertise-pip - Select the check box to enable the Advertise PIP feature.

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

VRF Template and **VRF Extension Template**: Specifies the VRF template for creating VRFs, and the VRF extension template for enabling VRF extension to other fabrics.

Network Template and **Network Extension Template**: Specifies the network template for creating networks, and the network extension template for extending 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.

Underlay Routing Loopback Id - The loopback interface ID is populated as 0 since loopback0 is usually used for fabric underlay IGP peering purposes.

Underlay VTEP Loopback Id - The loopback interface ID is populated as 1 since loopback1 is usually used for the VTEP peering purposes.

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

OSPF Area ID – The OSPF area ID, if OSPF is used as the IGP within the fabric.

Note The OSPF or IS-IS authentication fields are enabled based on your selection in the **Link-State Routing Protocol** field in the **General** tab.

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 - The Key ID is populated.

OSPF Authentication Key - The OSPF authentication key must be the 3DES key from the switch.

Note Plain text passwords are not supported. Login to the switch, retrieve the encrypted key and enter it in this field. Refer the Retrieving the Authentication Key section for details.

Enable ISIS Authentication - Select the check box to enable IS-IS authentication. Deselect the check box to disable it. If you enable this field, the IS-IS authentication fields are enabled.

ISIS Authentication Keychain Name - Enter the Keychain name, such as CiscoisisAuth.

ISIS Authentication Key ID - The Key ID is populated.

ISIS Authentication Key - Enter the Cisco Type 7 encrypted key.

Note Plain text passwords are not supported. Login to the switch, retrieve the encrypted key and enter it in this field. Refer the Retrieving the Authentication Key section for details.

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 for existing switches.

This is enabled by default. Clear the check box to disable VXLAN OAM function.

If you want to enable the VXLAN OAM function on specific switches and disable on other switches in the fabric, you can use freeform configurations to enable OAM and disable OAM in the fabric settings.

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

Enable Tenant DHCP – Select the checkbox to enable the tenant DHCP support.

Note Ensure that **Enable Tenant DHCP** is enabled before enabling DHCP related parameters in the overlay profiles.

Enable BFD – Select the checkbox to enable **feature bfd** on all switches in the fabric.

Note Additional BFD related configurations must be added by using the appropriate freeform config fields.

The BFD feature is disabled by default.

Greenfield Cleanup Option – Enable the switch cleanup option for greenfield switches without a switch reload. This option is typically recommended only for the data center environments with the Cisco Nexus 9000v Switches.

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 Encryption Type and BGP Authentication Key fields are enabled.

Note If you enable BGP authentication using this field, leave the iBGP Peer-Template Config field blank to avoid duplicate configuration.

BGP Authentication Key Encryption Type – Choose the 3 for 3DES encryption type, or 7 for Cisco encryption type.

BGP Authentication Key - Enter the encrypted key based on the encryption type.

Note Plain text passwords are not supported. Login to the switch, retrieve the encrypted key and enter it in the BGP Authentication Key field. Refer the Retrieving the Authentication Key section for details.

iBGP Peer-Template Config – Add iBGP peer template configurations on the leaf switches to establish an iBGP session between the leaf switch and route reflector.

If you use BGP templates, add the authentication configuration within the template and clear the Enable BGP Authentication check box to avoid duplicate configuration.

In the sample configuration, the 3DES password is displayed after password 3.

```
router bgp 65000
    password 3 sd8478fswerdfw3434fsw4f4w34sdsd8478fswerdfw3434fsw4f4w
```

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*, *Border Spine*, and *Border Gateway Spine* roles.

7. Click the **Resources** tab.

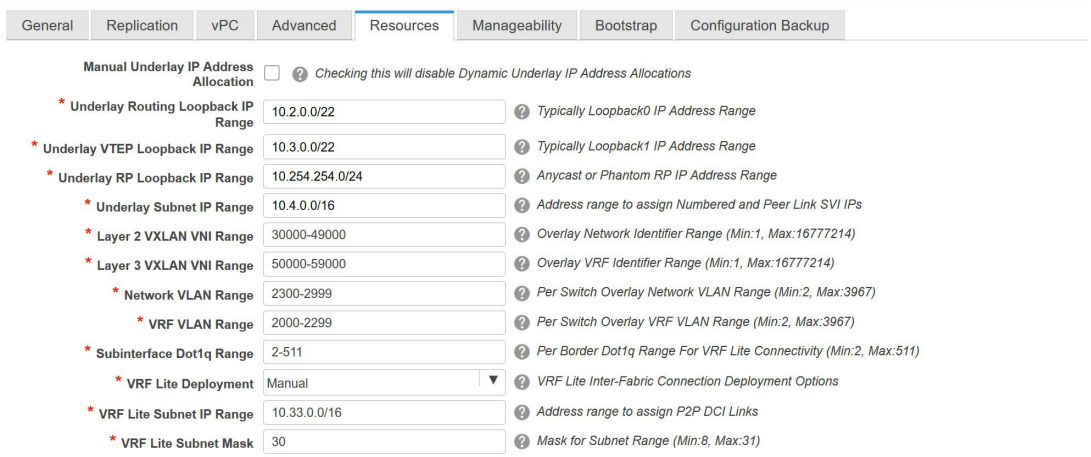

Manual Underlay IP Address Allocation – *Do not* select this check box if you are transitioning your VXLAN fabric management to DCNM.

- By default, DCNM allocates the underlay IP address resources (for loopbacks, fabric interfaces, etc) dynamically from the defined pools. If you select the check box, the allocation scheme switches to static, and some of the dynamic IP address range fields are disabled.
- For static allocation, the underlay IP address resources must be populated into the Resource Manager (RM) using REST APIs.

  Refer the Cisco DCNM REST API Reference Guide, Release 11.2(1) for more details. The REST APIs must be invoked after the switches are added to the fabric, and before you use the Save & Deploy option.
- The Underlay RP Loopback IP Range field stays enabled if BIDIR-PIM function is chosen for multicast replication.
- Changing from static to dynamic allocation keeps the current IP resource usage intact. Only future IP address allocation requests are taken from dynamic pools.

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 RP Loopback IP Range - Specifies the anycast or phantom RP IP address range.

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.

VRF Lite Deployment - Specify the VRF Lite method for extending inter fabric connections.

If you select Manual, the VRF Lite subnet details are required so that the resource manager can reserve the address space.

If you select Back2BackOnly, ToExternalOnly, or Both, then the VRF Lite subnet fields are enabled.

VRF Lite Subnet IP Range and **VRF Lite Subnet Mask** – These fields are populated with the DCI subnet details. Update the fields as needed.

The values shown in your screen 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:

Note When you update a range of values, ensure that it does not overlap with other ranges. You should 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.

a. Update the L2 range and click **Save**.

b. Click the **Edit Fabric** option again, update the L3 range and click **Save**.

8. Click the **Manageability** tab.

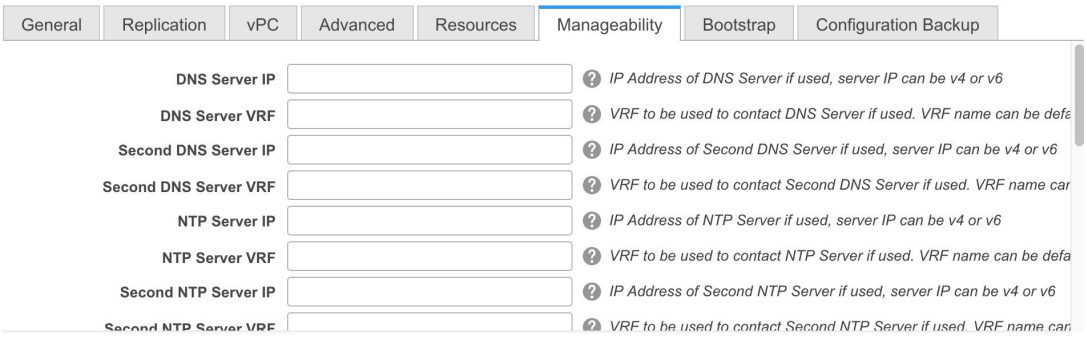

The fields in this tab are:

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.

Syslog Server IP – IP address of the syslog server, if used.

Syslog Server Severity – Severity level of the syslog server. To specify a higher severity, enter a higher number.

Syslog Server VRF – The default or management VRF that the syslog server IP address is assigned to.

Second Syslog Server IP – IP address of the second syslog server, if used.

Second Syslog Server Severity – Severity level of the second syslog server. To specify a higher severity, enter a higher number.

Second Syslog Server VRF – The default or management VRF that the second syslog server's IP address is assigned to.

9. Click the **Bootstrap** tab.

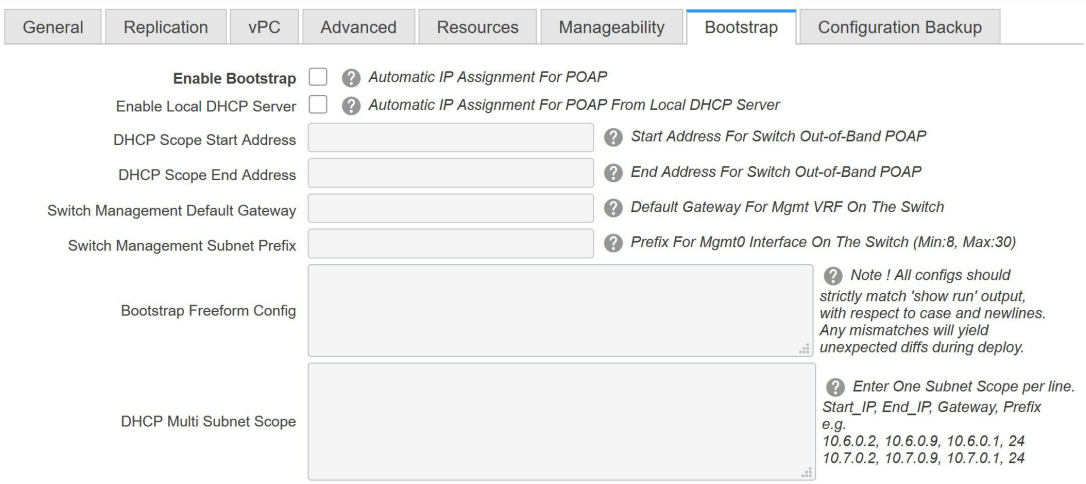

Enable Bootstrap - Select this check box to enable the bootstrap feature.

After you enable bootstrap, you can enable the DHCP server for automatic IP address assignment using one of the following methods:

- External DHCP Server: Enter information about the external DHCP server in the **Switch Management Default Gateway** and **Switch Management Subnet Prefix** fields.
- Local DHCP Server: Enable the **Local DHCP Server** checkbox and enter details for the remaining mandatory fields.

Enable Local DHCP Server - Select this check box to initiate enabling of automatic IP address assignment through the local DHCP server. When you select this check box, the **DHCP Scope Start Address** and **DHCP Scope End Address** fields become editable.

If you do not select this check box, DCNM uses the remote or external DHCP server for automatic IP address assignment.

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 10.0.1.1 and subnet mask 24, ensure that the DHCP scope is within the specified subnet, between 10.0.1.2 and 10.0.1.254.

Bootstrap Freeform Config - (Optional) Enter additional commands as needed. For example, if you are using AAA or remote authentication related configurations, you need to add these configurations in this field to save the intent. After the devices boot up, they contain the intent defined in the **Bootstrap Freeform Config** field.

Copy-paste the running-config to a **freeform config** field with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. For more information, see Resolving Freeform Config Errors in Switches, on page 204.

DHCP Multi Subnet Scope - Specifies the field to enter one subnet scope per line. This field is editable after you check the **Enable Local DHCP Server** check box.

The format of the scope should be defined as:

DHCP Scope Start Address, DHCP Scope End Address, Switch Management Default Gateway, Switch Management Subnet Prefix

For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24

10. Click the **Configuration Backup** tab. The fields on this tab are:

Hourly Fabric Backup: Select the check box to enable an hourly backup of fabric configurations and the intent. The backup process is initiated only when you click **Save and Deploy**, and the subsequent configuration compliance activity is successfully completed.

You can enable an hourly backup for fresh fabric configurations and the intent as well. If there is a configuration push in the previous hour, DCNM takes a backup.

Intent refers to configurations that are saved in DCNM but yet to be provisioned on the switches.

Scheduled Fabric Backup: Check the check box to enable a daily backup. This backup tracks changes in running configurations on the fabric devices that are not tracked by configuration compliance.

Scheduled Time: Specify the scheduled backup time in a 24-hour format. This field is enabled if you check the **Scheduled Fabric Backup** check box.

Select both the check boxes to enable both back up processes. If you update settings, execute the **Save & Deploy** option on the fabric topology screen (click within the fabric box to access the fabric topology screen).

The backup configuration files are stored in the following path in DCNM:
`/usr/local/cisco/dcm/dcnm/data/archive`

The number of archived files that can be retained is set in the **# Number of archived files per device to be retained:** field in the **Server Properties** window.

Note Hourly and scheduled backup processes happen only during the next periodic configuration compliance activity, and there can be a delay of up to an hour. To trigger an immediate backup, do the following:

a. Choose **Control > Fabric Builder**. The Fabric Builder screen comes up.

b. Click within the specific fabric box. The fabric topology screen comes up.

c. From the **Actions** pane at the left part of the screen, click **Re-Sync Fabric**.

You can also initiate the fabric backup in the fabric topology window. Click **Backup Now** in the **Actions** pane.

11. 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.

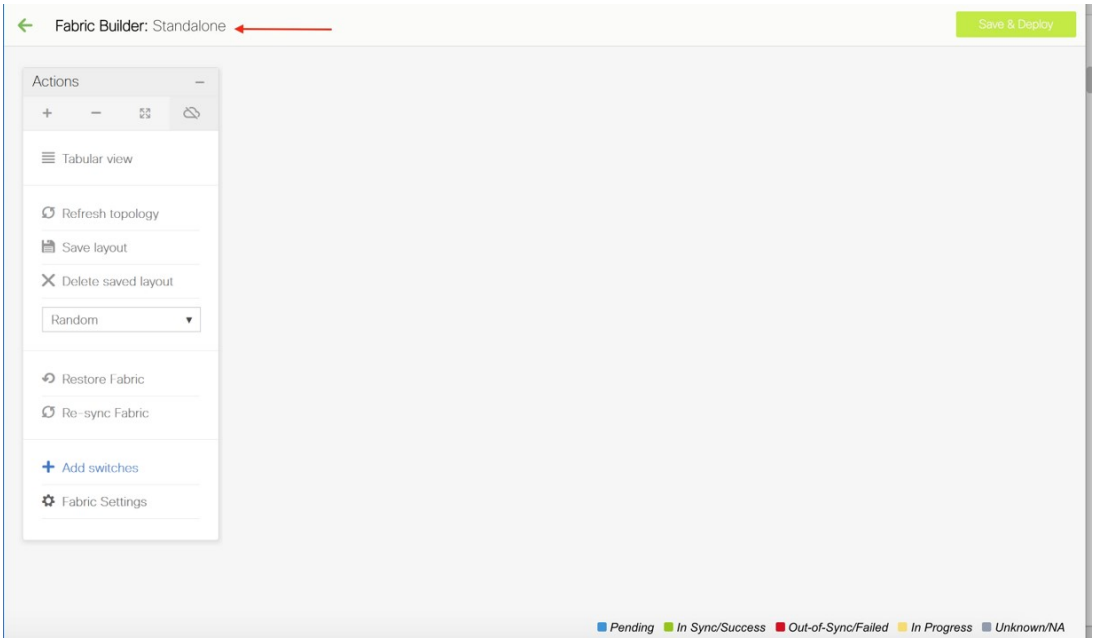

(At the same time, the newly created fabric instance appears on the **Fabric Builder** screen. To go to the **Fabric Builder** screen, click the left arrow (←) button above the **Actions** pane [to the left of the screen]).

The **Actions** pane 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 options are explained:

- **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.
- **Save Layout** – Saves a custom view of the topology. You can create a specific view in the topology and save it for ease of use.
- **Delete saved layout** – Deletes the custom view of the topology
- **Topology views** - 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 window. 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.
- **Restore Fabric** – Allows you to restore the fabric to a prior DCNM configuration state (one month back, two months back, and so on). For more information, see the *Restore Fabric* section.

- **Resync Fabric** - Use this option to resynchronize DCNM state when there is a large scale out-of-band change, or if configuration changes do not register in the DCNM properly. The resync operation does a full CC run for the fabric switches and recollects “show run” and “show run all” commands from the switches. When you initiate the re-sync process, a progress message is displayed on the window. During the re-sync, the running configuration is taken from the switches. Then, the Out-of-Sync/In-Sync status for the switch is recalculated based on the intent or expected configuration defined in DCNM versus the current running configuration that was taken from the switches.
- **Add Switches** – Allows you to add switch instances to the fabric.
- **Fabric Settings** – Allows you to view or edit fabric settings.
- **Cloud** icon - Click the **Cloud** icon to display (or not display) an **Undiscovered** cloud.

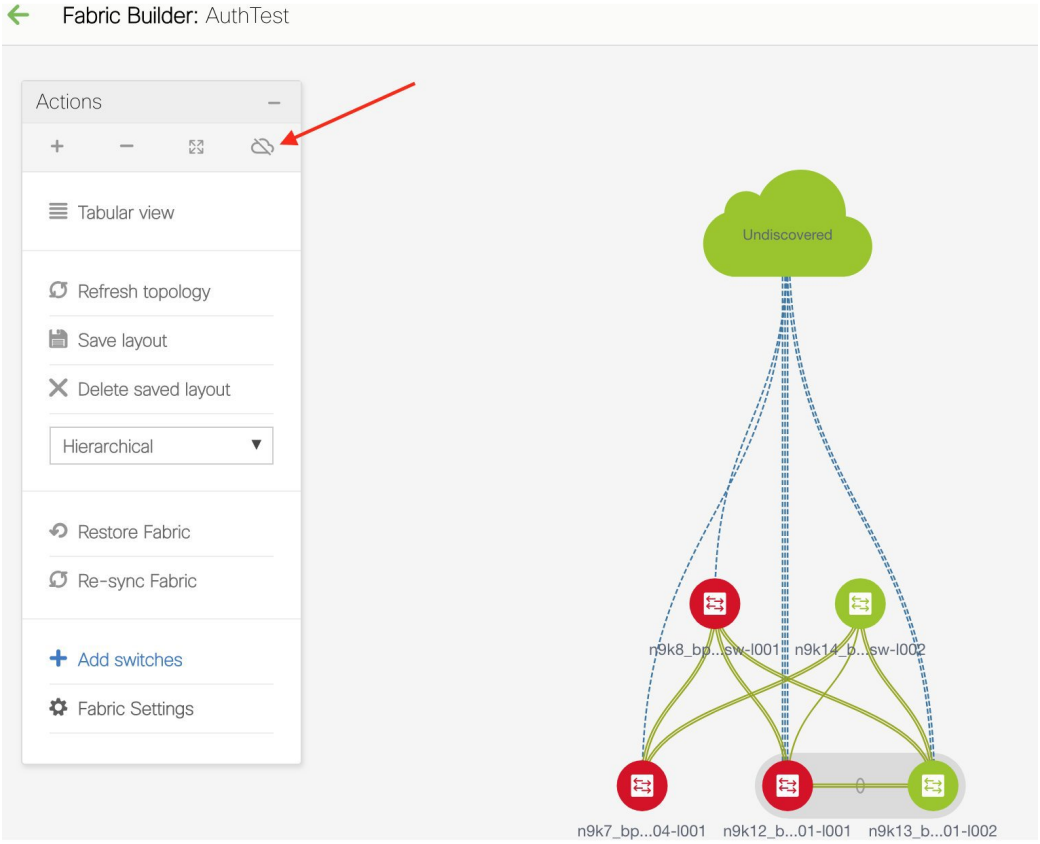

When you click the icon, the Undiscovered cloud and its links to the selected fabric topology are not displayed.

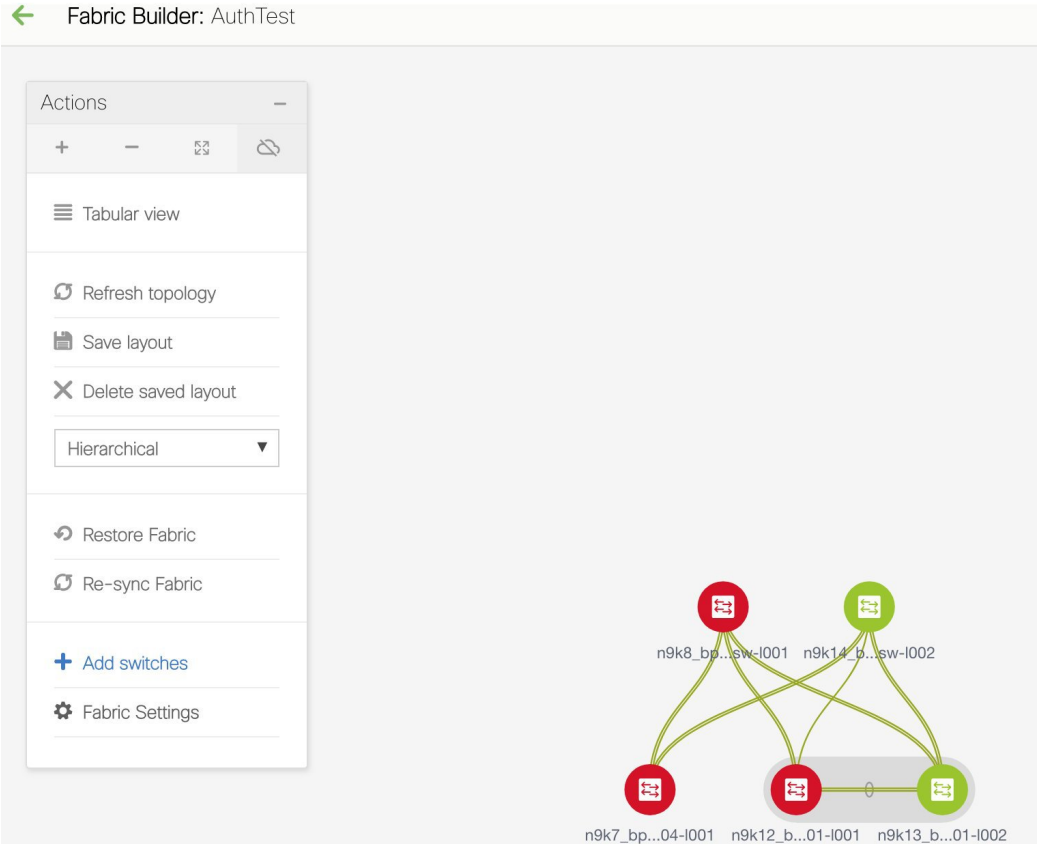

Click the **Cloud** icon again to display the **Undiscovered** cloud.

SCOPE - You can toggle between fabrics by using the SCOPE drop-down box at the top right. The current fabric is highlighted. An MSD and its member fabrics are distinctly displayed, wherein the member fabrics are indented, under the MSD fabric.

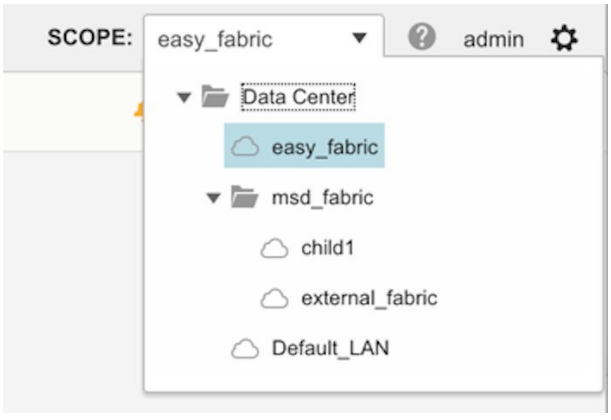

Adding Switches to a Fabric

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 created in DCNM. 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.

Additionally, you can pre-provision switches. For more information, see Pre-provisioning a Device , on page 33.

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.

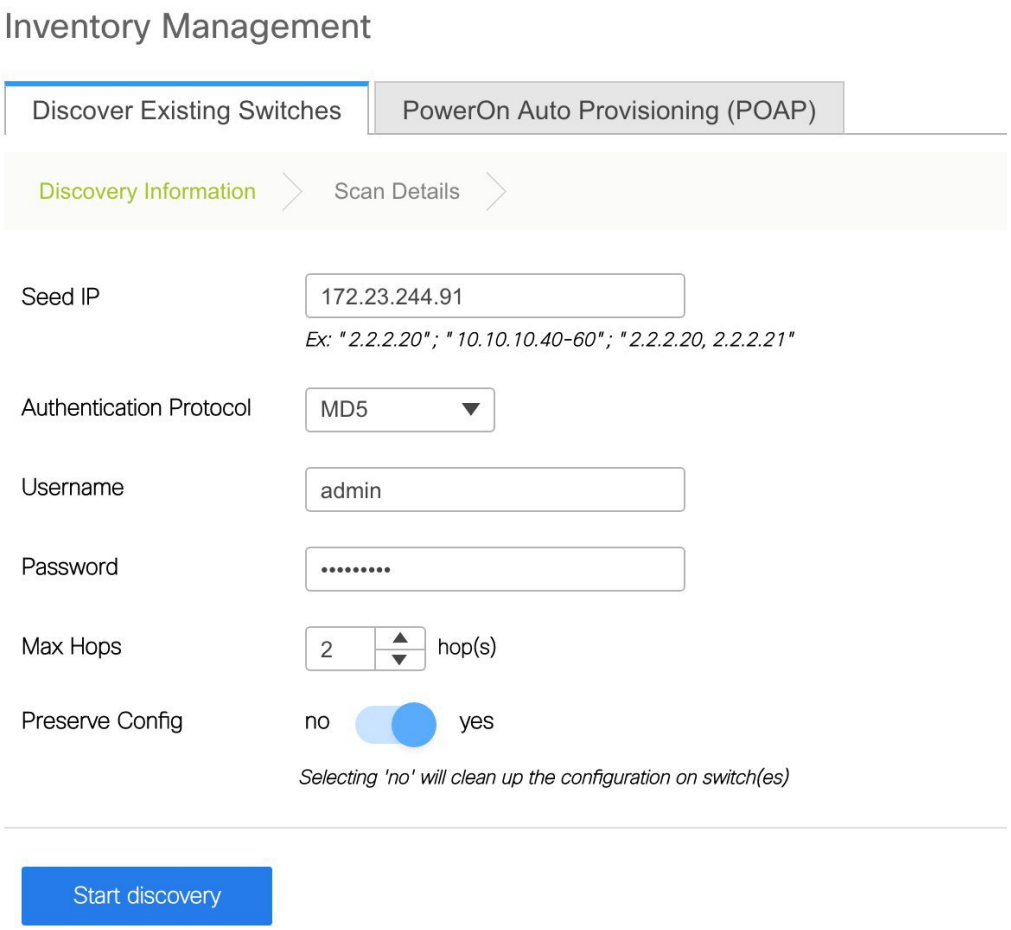

2. Click **Start discovery**. The **Scan Details** window comes up shortly. Since the **Max Hops** field was populated with 2, the switch with the specified IP address (leaf-91) and switches two hops from it are populated in the **Scan Details** window.

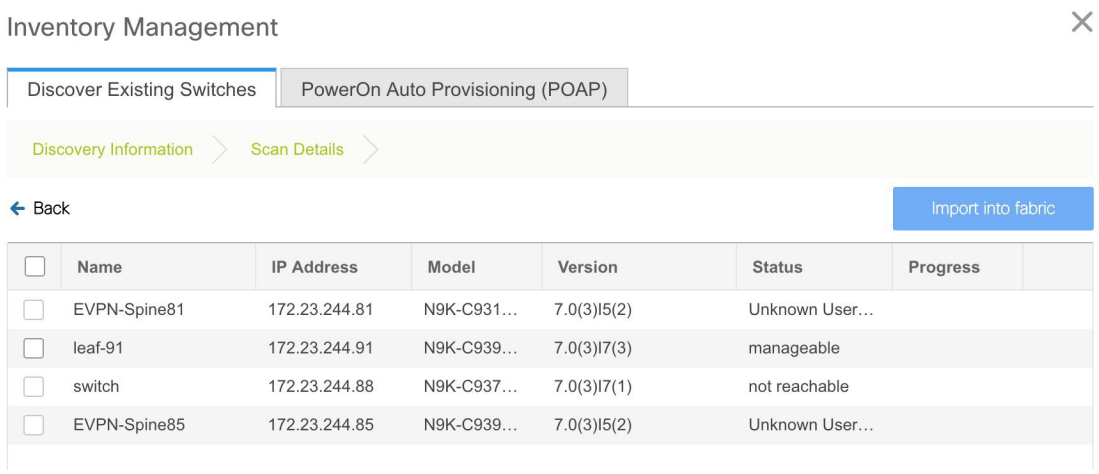

3. Check the check box next to the concerned switch and click **Import into fabric**.

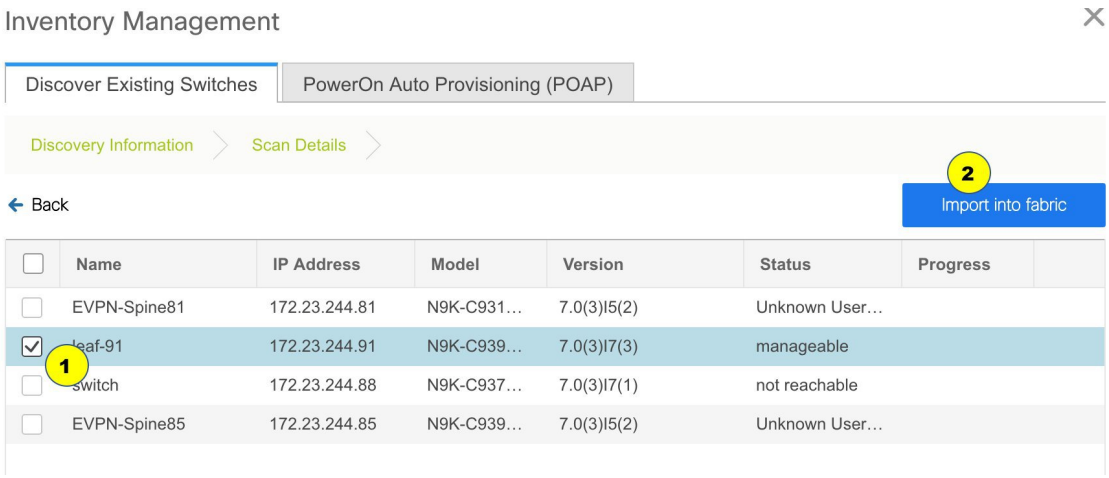

Though this example describes the discovery of one switch, it is a best practice to discover multiple switches at once. The switches must be properly cabled and connected to the DCNM server and the switch status must be manageable.

The switch discovery process is initiated. The **Progress** column displays progress for all the selected switches. It displays **done** for each switch on completion.

Note You must not close the screen (and try to add switches again) until all selected switches are imported or an error message comes up.

If an error message comes up, close the screen. The fabric topology screen comes up. The error messages are displayed at the top right part of the screen. Resolve the errors wherever applicable and initiate the import process again by clicking **Add Switches** in the Actions panel.

After DCNM discovers all the switches, and the Progress column displays **done** for all switches, close the screen. The *Standalone* fabric topology screen comes up again. The switch icons of the added switches are displayed in it.

Note You will encounter the following errors during switch discovery sometimes.

Discovery error - The switch discovery process might fail for a few switches, and the Discovery Error message displayed. However, such switches are displayed in the fabric topology. You must remove such switches from the fabric (right-click the switch icon and click **Discovery** > **Remove** from fabric), and import them again.

Device connectivity issue: Before proceeding further, wait for ten minutes for the switch-internal processes to complete. Else, you might encounter a device connectivity failure message at a later stage.

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

   When all switches are added and roles assigned to them, the fabric topology contains the switches and connections between them.

5. After discovering the switches, assign the fabric role to each switch. Since each switch is assigned the leaf role by default, assign other roles as needed. Right click the switch, and use the **Set role** option to set the appropriate role.

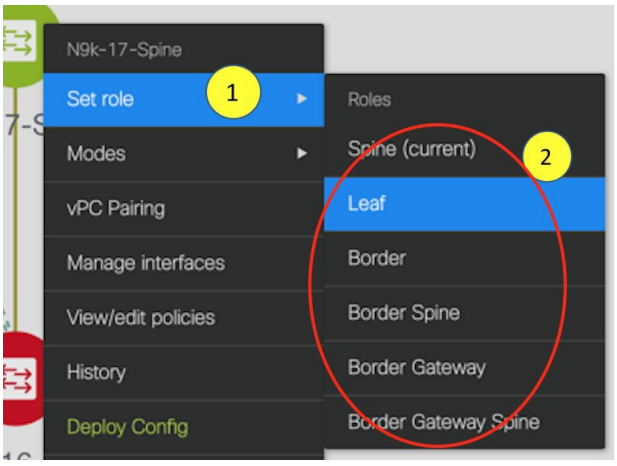

Note

- Starting from DCNM 11.1(1), switch roles can be changed if there are no overlays on the switches, but only as per the list of allowed switch role changes given at Switch Operations, on page 109.
- After you upgrade to Cisco DCNM Release 11.1(1) with an existing fabric with the Easy_Fabric template, you cannot set the Border Spine or Border Gateway Spine roles to switches, because these roles are not supported with the Easy_Fabric template. You need to use the **Easy_fabric_11_1** template to set these roles for switches in a fabric.

If you choose the Hierarchical layout for display (in the Actions panel), 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 border switches 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* or *Border Gateway Spine*. To connect fabrics using the VRF Lite feature, you must change the role of the border leaf switch to *Border* or *Border Spine*. If you want to deploy VRF Lite and EVPN Multi-Site features in a fabric, you must set the device role to *Border Gateway* or *Border Gateway Spine* 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 DCNM 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.

Note vPC support is added for BGWs in the DCNM 11.1(1) release.

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.

When you enable or disable a vPC setup or the advertise-pip option, or update Multi-Site configuration, you should use the **Save & Deploy** operation. At the end of the operation, an error prompts you to configure the **shutdown** or **no shutdown** command on the nve interface. A sample error screenshot when you enable a vPC setup:

To resolve, go to the **Control > Interfaces** screen and deploy the **No Shutdown** or **Shutdown** configuration on the nve interface (**nve1** in the screenshot).

Click **Save & Deploy** in the Fabric Builder topology screen again to complete the task.

If the non-overlay SVIs are captured in the DCNM intent while the switch is in the standalone mode, and then the switch becomes a part of a vPC pair, the switch generates the following configuration:

```
no ip redirects
no ipv6 redirects
```

To avoid a diff from the configuration compliance in DCNM, you must update the intent with the same config.

When a new vPC pair is created and deployed successfully using Cisco DCNM, one of the peers might be out-of-sync for the **no ip redirects** CLI even if the command exists on the switch. This out-of-sync is due to a delay on the switch to display the CLI in the running configuration, which causes a diff in the configuration compliance. Re-sync the switches in the **Config Deployment** window to resolve the diff.

6. Click **Save & Deploy** at the top right part of the screen.

   The template and interface configurations form the underlay network configuration on the switches. Also, freeform CLIs that were entered as part of fabric settings (leaf and spine switch freeform configurations entered in the Advanced tab) are deployed. For more details on freeform configurations, refer Enabling Freeform Configurations on Fabric Switches .

   Configuration Compliance: If the provisioned configurations and switch configurations do not match, the **Status** column displays out-of-sync. For example, if you enable a function on the switch manually through a CLI, then it results in a configuration mismatch.

   To ensure configurations provisioned from DCNM to the fabric are accurate or to detect any deviations (such as out-of-band changes), DCNM's Configuration Compliance engine reports and provides necessary remediation configurations.

   When you click **Save & Deploy**, the **Config Deployment** window appears.

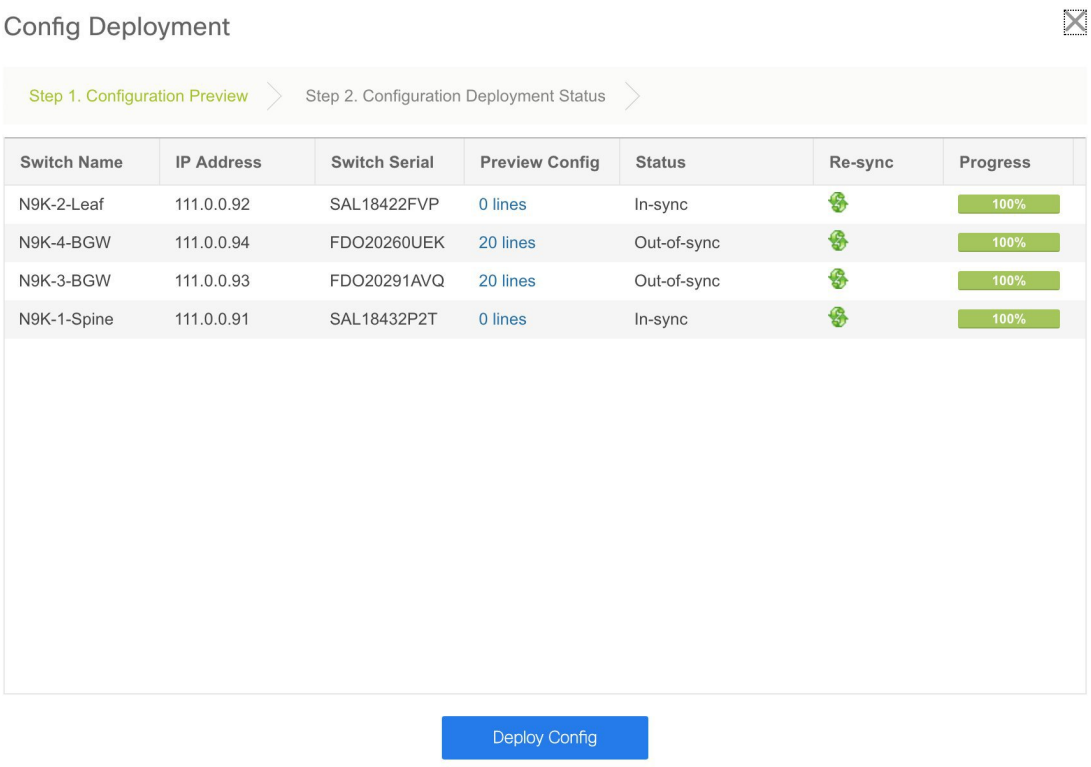

   If the status is out-of-sync, it suggests that there is inconsistency between the DCNM and configuration on the device.

   The Re-sync button is displayed for each switch in the Re-sync column. Use this option to resynchronize DCNM state when there is a large scale out-of-band change, or if configuration changes do not register in the DCNM properly. The re-sync operation does a full CC run for the switch and recollects "show run" and "show run all" commands from the switch. When you initiate the re-sync process, a progress message

is displayed on the screen. During the re-sync, the running configuration is taken from the switch. The Out-of-Sync/In-Sync status for the switch is recalculated based on the intent defined in DCNM.

Click the Preview Config column entry (updated with a specific number of lines). The Config Preview screen comes up.

The Pending Config tab displays the pending configurations for successful deployment.

The **Side-by-side Comparison** tab displays the current configurations and expected configurations together.

Note that multi-line banner configuration support is available in Cisco DCNM Release 11.1(1).

In DCNM 11.0, Configuration Compliance only supports single-line banner motd configuration. In DCNM 11.1, multi-line banner motd configuration is supported. Multi-line banner motd configuration can be configured in DCNM with freeform configuration policy, either per switch using **switch_freeform**, or per fabric using leaf/spine freeform configuration. Note that after the multi-line banner motd is configured, deploy the policy by executing the **Save & Deploy** option in the (top right part of the) fabric topology screen. Else, the policy may not be deployed properly on the switch. The **banner** policy is only to configure single-line banner configuration. Also, you can only create one banner related freeform configuration/policy. Multiple policies for configuring banner motd is not supported.

7. 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 successful configuration provisioning (when all switches display a progress of 100%), close the screen.

   The fabric topology is displayed. The switch icons turn green to indicate successful configuration.

   If a switch icon is in red color, it indicates that the switch and DCNM configurations are not in sync. When deployment is pending on a switch, the switch is displayed in blue color.

Note If there are any warning or errors in the CLI execution, a notification will appear in the **Fabric builder** window. Warnings or errors that are auto-resolvable have the **Resolve** option.

You can right click the switch icon and update switch related settings.

SCOPE: You can toggle between fabrics by using the **SCOPE** drop-down list at the top right part of the screen. By default, the current fabric is highlighted. An MSD and its member fabrics are distinctly displayed, wherein the member fabrics are indented under the MSD fabric.

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 Enabling Freeform Configurations on Fabric Switches 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 in the external fabric 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.

   Choose **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 a standalone fabric (Click **Control** > **Fabric Builder** and click a standalone fabric). The fabric topology is displayed.

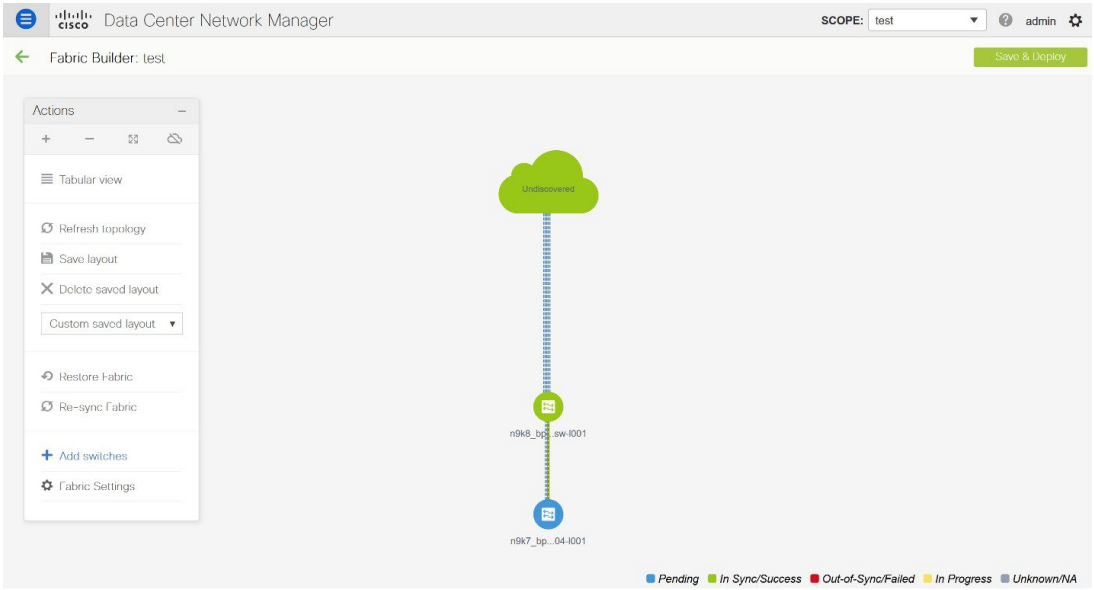

Note If you want to POAP with DCHP, make sure that DHCP is enabled on the fabric settings. Click **Fabric Settings** and edit the DHCP information in the **Bootstrap** tab.

5. Go to the fabric topology window and click the **Add switches** option from the **Actions** panel. The Inventory Management window 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 window. Also, an option to add the IP address, hostname, and password are made available. If the switch information is not retrieved, refresh the window.

Note

- Before initiating POAP, make sure that password for the device should contain characters from at least three of the following classes: lower case letters, upper case letters, digits, and special characters.

  If a switch password is changed, then the **nfm_switch_user** PTI has to be updated with encrypted password, that is, copy and paste from the switch. This PTI update is apart from the device and LAN credentials update. The device-config is updated immediately if you click **Save & Deploy** in **Fabric Builder**.

- At the top left part of the window, *export* and *import* options are provided to export and import the .csv file that contains the switch information. You can pre-provision devices using the *import* option as well.

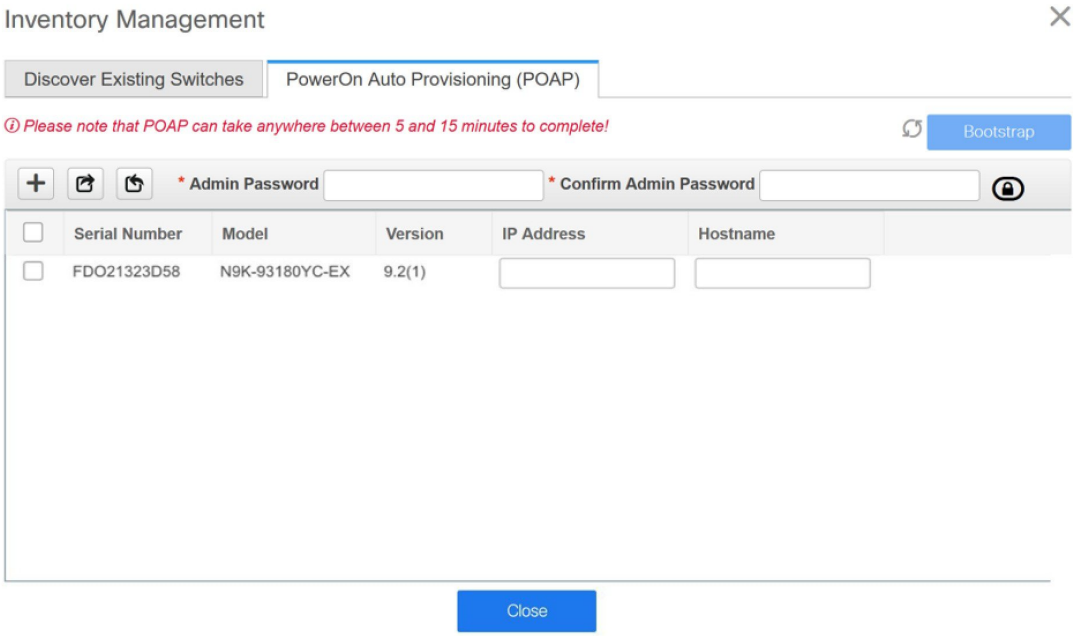

Select the checkbox next to the switch and add switch credentials: IP address and host name.

Beginning with Release 11.2(1), you can provision devices in advance. To pre-provision devices, refer to Pre-provisioning a Device , on page 33.

7. In the **Admin Password** and **Confirm Admin Password** fields, enter and confirm the admin password.

   This admin password is applicable for all the switches displayed in the POAP window.

Note If you do not want to use admin credentials to discover switches, you can instead use the AAA authentication, that is, RADIUS or TACACS credentials for discovery only.

8. (Optional) Use discovery credentials for discovering switches.
   a. Click the **Add Discovery Credentials** icon to enter the discovery credentials for switches.

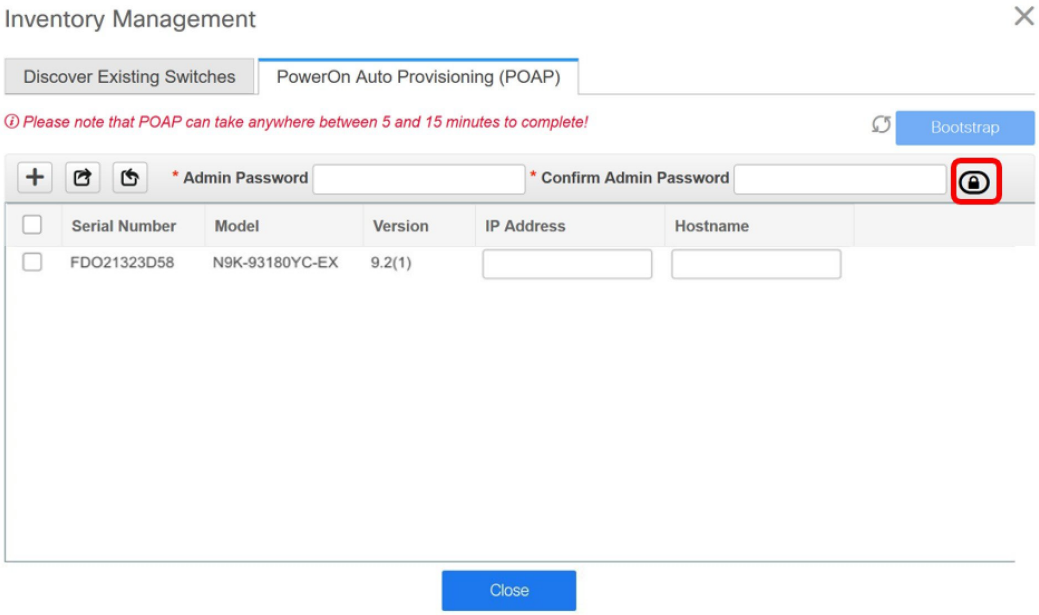

b. In the **Discovery Credentials** window, enter the discovery credentials such as discovery username and password.

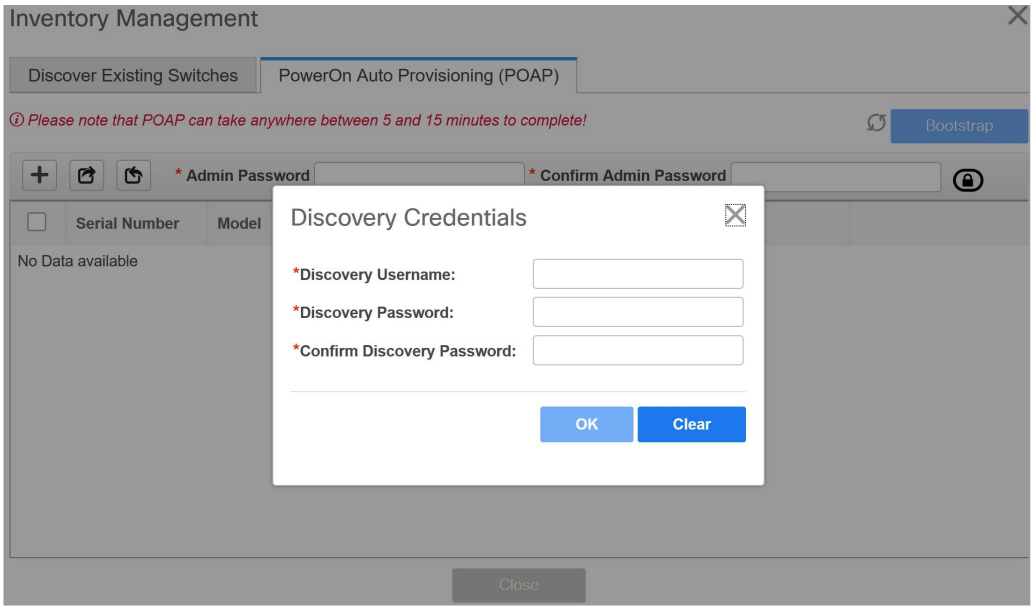

Click **OK** to save the discovery credentials.

If the discovery credentials are not provided, DCNM uses the admin user and password to discover switches.

Note

- The discovery credentials that can be used are AAA authentication based credentials, that is, RADIUS or TACACS.
- The discovery credential is not converted as commands in the device configuration. This credential is mainly used to specify the remote user (or other than the admin user) to discover the switches. If you want to add the commands as part of the device configuration, add them in the **Bootstrap Freeform Config** field under the **Bootstrap** tab in the fabric settings. Also, you can add the respective policy from **View/Edit Policies** window.

9. Click **Bootstrap** at the top right part of the screen.

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

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

11. 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.

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

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

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

15. 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.

16. 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.

When you enable or disable a vPC setup or the advertise-pip option, or update Multi-Site configuration, you should use the **Save & Deploy** operation. At the end of the operation, an error prompts you to configure the **shutdown** or **no shutdown command** on the nve interface. A sample error screenshot when you enable a vPC setup:

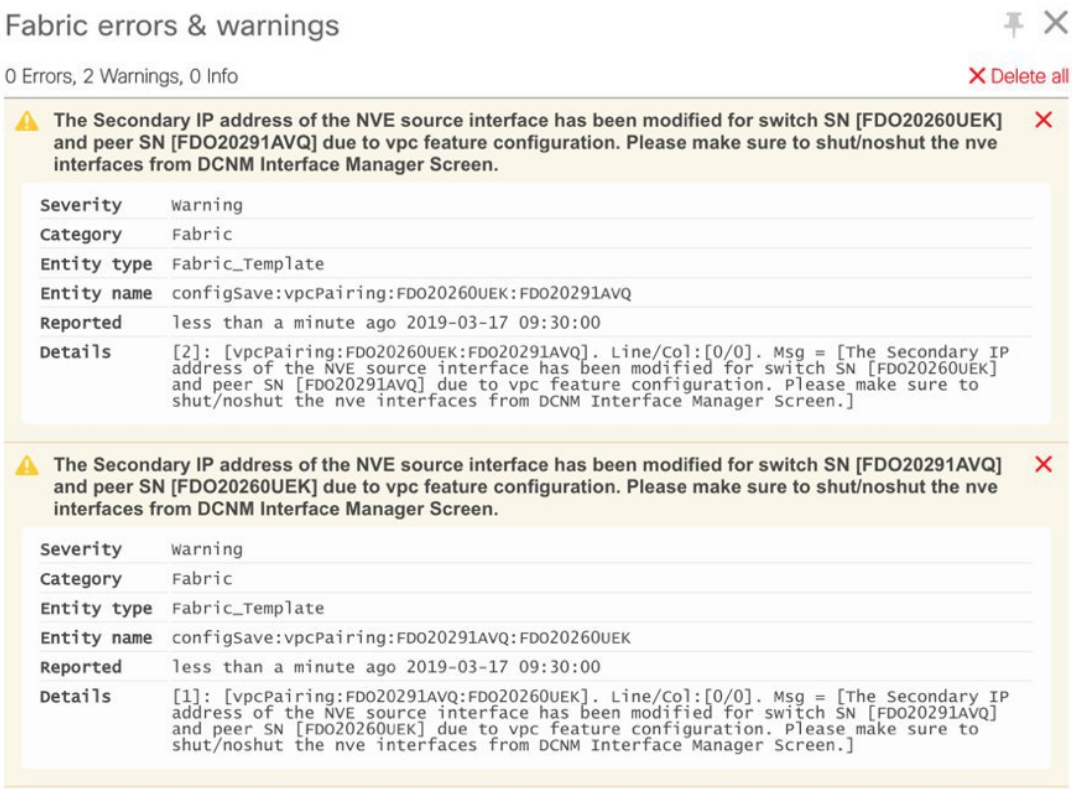

To resolve, go to the **Control > Interfaces** screen and deploy the **No Shutdown** or **Shutdown** configuration on the nve interface.

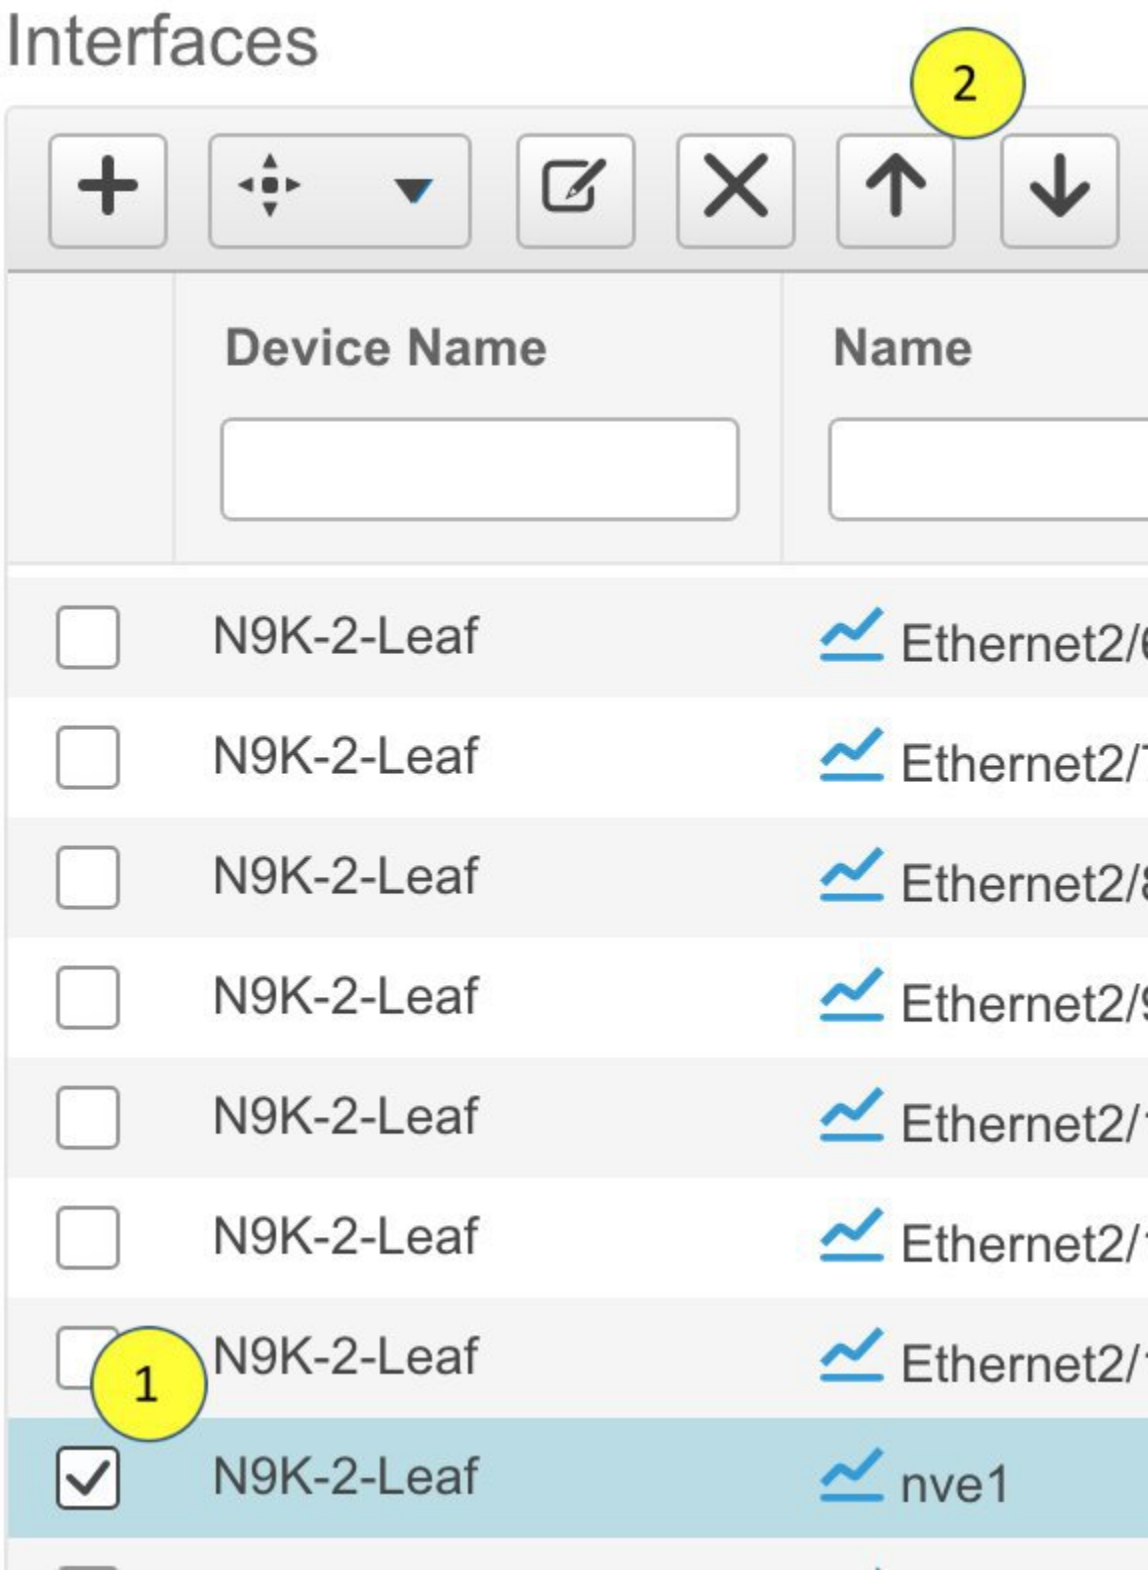

Click **Save & Deploy** in the Fabric Builder topology screen again to complete the task.

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.
- You might encounter an issue with module discovery after bootstrap. In such cases, the discovery happens after a delay. If not, go through the discovery process again.

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**.
  - Starting from DCNM 11.1(1), switch roles can be changed if there are no overlays on the switches, but only as per the list of allowed switch role changes given at Switch Operations, on page 109.
  - After you upgrade to Cisco DCNM Release 11.1(1) with an existing fabric with the Easy_Fabric template, you cannot set the Border Spine or Border Gateway Spine roles to switches, because these roles are not supported with the **Easy_Fabric** template. You need to use the **Easy_fabric_11_1** template to set these roles for switches in a fabric.

- **Modes** - Maintenance and Active/Operational modes.
- **vPC Pairing** - Select a switch for vPC and then select its peer.

  You can create a virtual link for a vPC pair or change the existing physical link to a virtual link for a vPC pair.
- **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 configuration partially 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. [Refer Interfaces].
- Create overlay networks and VRFs and deploy them on the switches. [Refer Creating and Deploying Networks and VRFs].

Pre-provisioning a Device

In DCNM 11.2, you can provision devices in advance.

Note Ensure that you enter DHCP details in the Bootstrap tab in the fabric settings.

- The pre-provisioned devices support the following configurations in DCNM:
  - Base management
  - vPC Pairing
  - Intra-Fabric links
  - Interface breakout configuration
- The pre-provisioned devices do not support the following configurations in DCNM:
  - Inter-Fabric links
  - Host ports
  - vPCs to the access switches or hosts
  - FEX
  - Overlay network configurations
- When a device is being pre-provisioned has breakout links, you need to specify the corresponding breakout command along with the switch's model and gateway in the **Data** field in the **Add a new device to pre-provisioning** window in order to generate the breakout PTI.

  Note the following guidelines:
  - Multiple breakout commands can be separated by a semicolon (;).
  - The definitions of the fields in the data JSON object are as follows:
    - **modulesModel**: (Mandatory) Specifies the switch module’s model information.
    - **gateway**: (Mandatory) Specifies the default gateway for the management VRF on the switch. This field is required to create the intent to pre-provision devices. You need to enter the gateway even if it is in the same subnet as DCNM to create the intent as part of pre-provisioning a device.
    - **breakout**: (Optional) Specifies the breakout command provided in the switch.
    - **portMode**: (Optional) Specifies the port mode of the breakout interface.

  The examples of the values in the **Data** field are as follows:
  - {"modulesModel": ["N9K-C93180LC-EX"], "gateway": "10.1.1.1/24"}
  - {"modulesModel": ["N9K-C93180LC-EX"],"breakout": "interface breakout module 1 port 1 map 10g-4x", "portMode": "hardware profile portmode 4x100G+28x40G", "gateway": "172.22.31.1/24" }

- {"modulesModel": ["N9K-X9736C-EX", "N9K-X9732C-FX", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-SUP-B+", "N9K-SC-A", "N9K-SC-A"], "gateway": "172.22.31.1/24"}
- {"breakout":"interface breakout module 1 port 50 map 10g-4x" , "gateway": "172.16.1.1/24", "modulesModel": ["N9K-C93180YC-EX "]}
- {"modulesModel": ["N9K-X9732C-EX", "N9K-X9732C-EX", "N9K-C9504-FM-E", "N9K-C9504-FM-E", "N9K-SUP-B", "N9K-SC-A", "N9K-SC-A"], "gateway": "172.29.171.1/24", "breakout":"interface breakout module 1 port 1,11,19 map 10g-4x; interface breakout module 1 port 7 map 25g-4x"}
- {"modulesModel": ["N9K-C93180LC-EX"], "gateway": "10.1.1.1/24","breakout": "interface breakout module 1 port 1-4 map 10g-4x","portMode": "hardware profile portmode 48x25G + 2x100G + 4x40G"}

Procedure

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

The **Fabric Builder** screen is displayed.

Step 2 Click within the fabric box.

Step 3 From the Actions panel, click the **Add switches** option.

The **Inventory Management** screen is displayed.

Step 4 Click the **POAP** tab.

Step 5 In the **POAP** tab, do the following:

a. Click + from the top left part of the screen.

The Add a new device screen comes up.

b. Fill up the device details as shown in the screenshot.

c. Click **Save**.

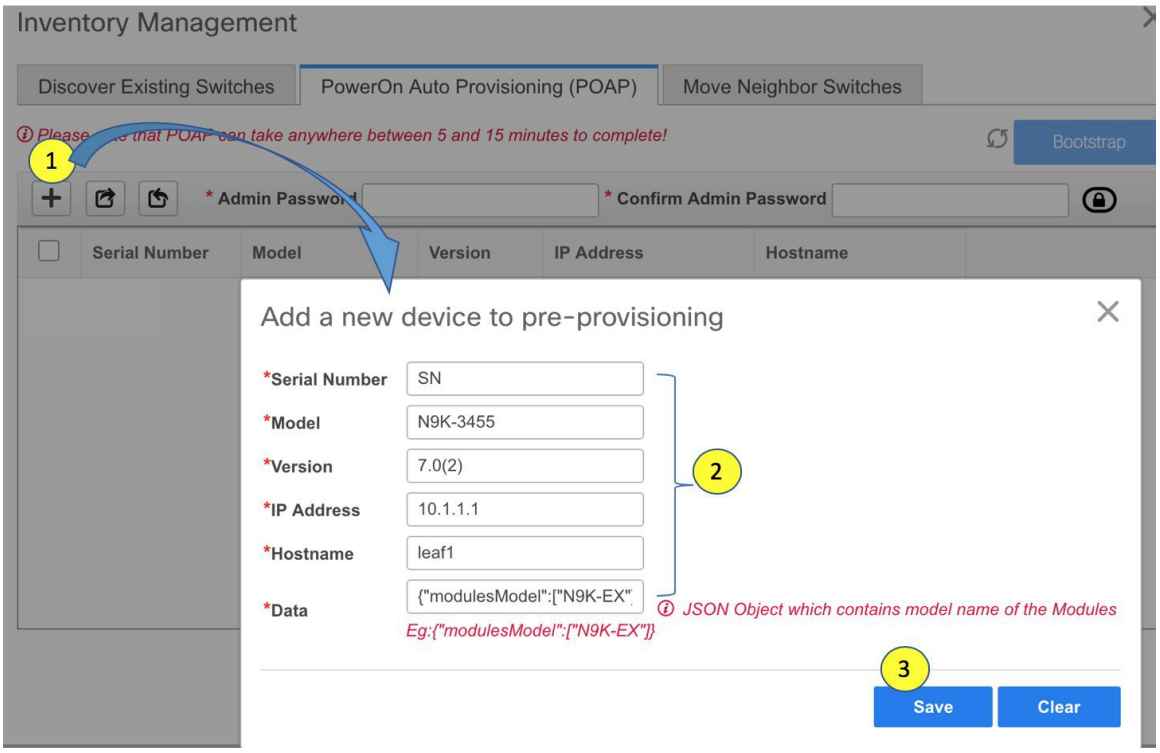

Serial Number: The serial number for the new device. This number can be a dummy serial number if the device serial number is not available.

For information about the **Data** field, see the examples provided in guidelines.

The device details appear in the POAP screen. You can add more devices for pre-provisioning.

At the top left part of the window, **Export** and **Import** icons are provided to export and import the .csv file that contains the switch information.

Using the **Import** option, you can pre-provision multiple devices.

Add new devices' information in the .csv file with all the mandatory fields (SerialNumber, Model, version, IpAddress, Hostname and Data fields [JSON Object]).

The Data column consists of the model name of the module to identify the hardware type from the fabric template. A .csv file screenshot:

	A	B	C	D	E	F	G
1	#SerialNumber(Eg:FD01344GH5)	#Model(Eg:N9k-C9236C)	#Version(Eg:7.0(3)12(3))	#IPAddress of the device	#HostName	#Data(JSON Field contains model name of the modules	
2	Serial Number	Model	Version	IP Address	Hostname	Data	
3	FDO21331SND	N9K-93180YC-EX	7.0(3)I5(2)	1.1.1.1	leaf1	{"modulesModel":["N9K-93180YC-EX"]}	
4	FDO21351N3X	N9K-C9236C	7.0(3)I4(1)	11.1.1.1	spine1	{"modulesModel":["N9K-C9236C"]}	
5	FDO21491A5K	N9K-C93240YC-FX2	7.0(3)I7(3)	12.1.1.1	leaf2	{"modulesModel":["N9K-C93240YC-FX2"]}	
6							

Step 6 Enter the administration password in the **Admin Password** and **Confirm Admin Password** fields.

Step 7 Select the device(s) and click **Bootstrap** at the top right part of the screen.

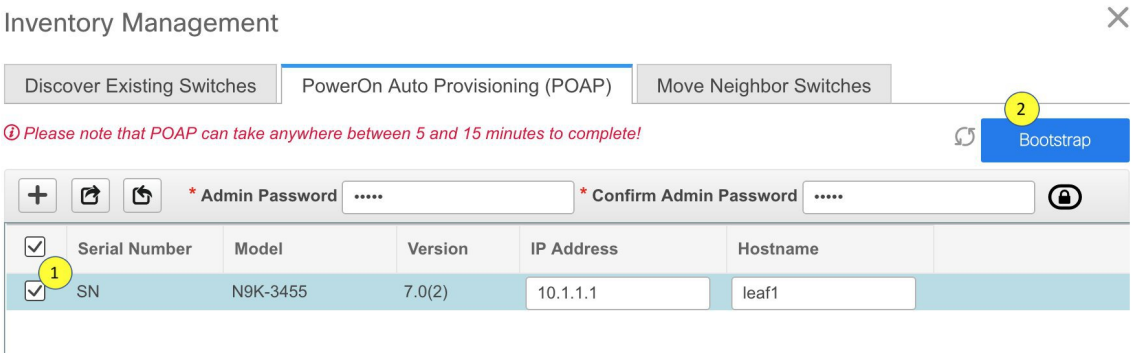

The leaf1 device appears in the fabric topology.

From the **Actions** panel, click **Tabular View**. You cannot deploy the fabric till the status of all the pre-provisioned switch(es) are displayed as **ok** under the **Discovery Status** column.

When you connect leaf1 to the fabric, the switch is provisioned with the IP address 10.1.1.1.

Step 8 Navigate to **Fabric Builder** and set roles for the device.

Create intra-link policy using one of the templates:

- **int_pre_provision_intra_fabric_link** to automatically generate intra fabric interface configuration with DCNM allocated IP addresses
- **int_intra_fabric_unnum_link_11_1** if you are using unnumbered links
- **int_intra_fabric_num_link_11_1** if you want to manually assign IP addresses to intra-links

Click **Save & Deploy**.

Configuration for the switches are captured in corresponding PTIs and can be seen in the **View/Edit Policies** window.

Step 9 To bring in the physical device, you can follow the manual RMA or POAP RMA procedure.

For more information, see Return Material Authorization (RMA), on page 135.

If you use the POAP RMA procedure, ignore the error message of failing to put the device into maintenance mode due to no connectivity since it is expected to have no connectivity to a non-existing device.

You need to click **Save & Deploy** in the fabric after the switch(es) are online to provision the host ports. This action must be performed before overlays are provisioned for the host port attachment.

Changing the TCAM Configuration on a Device

If you are onboarding the Cisco Nexus 9300 Series switches and Cisco Nexus 9500 Series switches with X9500 line cards using the bootstrap feature with POAP, DCNM pushes the following policies depending on the switch models:

- Cisco Nexus 9300 Series Switches: **tcam_pre_config_9300** and **tcam_pre_config_vxlan**
- Cisco Nexus 9500 Series Switches: **tcam_pre_config_9500** and **tcam_pre_config_vxlan**

Perform the following steps to change the TCAM carving of a device in DCNM.

1. Choose **Control > Fabrics > Fabric Builder**.
2. Click the fabric containing the specified switches that have been onboarded using the bootstrap feature.
3. Click **Tabular View** under the **Actions** menu in the **Fabric Builder** window.
4. Select all the specified switches and click the **View/Edit Policies** icon.
5. Search for **tcam_pre_config** policies.
6. If the TCAM config is incorrect or not applicable, select all these policies and click the Delete icon to delete policies.
7. Add one or multiple tcam_config policies and provide the correct TCAM configuration. For more information about how to add a policy, see *Adding PTIs for Multiple Switches*.
8. Reload the respective switches.

If the switch is used as a leaf, border leaf, border gateway leaf, border spine, or border gateway spine, add the **tcam_config** policy with the following command and deploy.

```
hardware access-list tcam region racl 1024
```

This config is required on the switches so that the NGOAM and VXLAN Suppress ARP features are functional.

Make sure that the priority of this tcam_config policy is higher than the **tcam_pre_config_vxlan** policy so that the config policy with **racl 1024** is configured before the **tcam_pre_config_vxlan** policy.

Note The **tcam_pre_config_vxlan** policy contains the config: **hardware access-list tcam region arp-ether 256 double-wide**.

Adding a vPC L3 Peer Keep-Alive Link

This procedure shows how to add a vPC L3 peer keep-alive link.

Note

- vPC L3 Peer Keep-Alive link is not supported with fabric vPC peering.
- In Brownfield migration, You need to manually create a vPC pairing when the L3 keep alive is configured on the switches. Otherwise, the vPC configuration is automatically picked up from the switches.

Procedure

Step 1 From DCNM, navigate to **Control > Template Library**.

Step 2 Search for the **vpc_serial_simulated** policy, select it, and click the **Edit** icon.

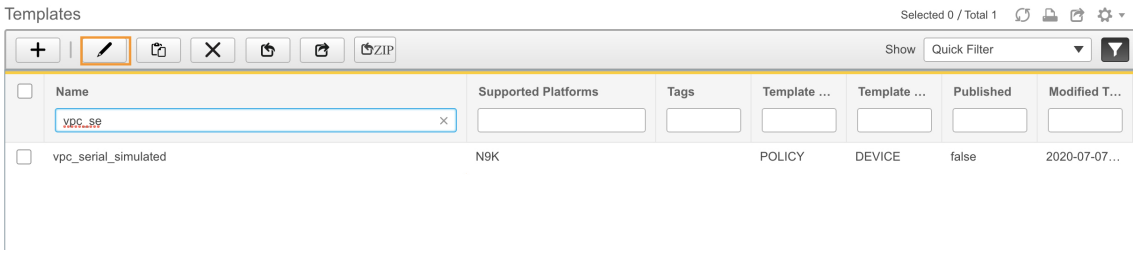

Step 3 Edit the template properties and set the **Template Sub Type** to **Device** so that this policy appears in **View/Edit Policies**.

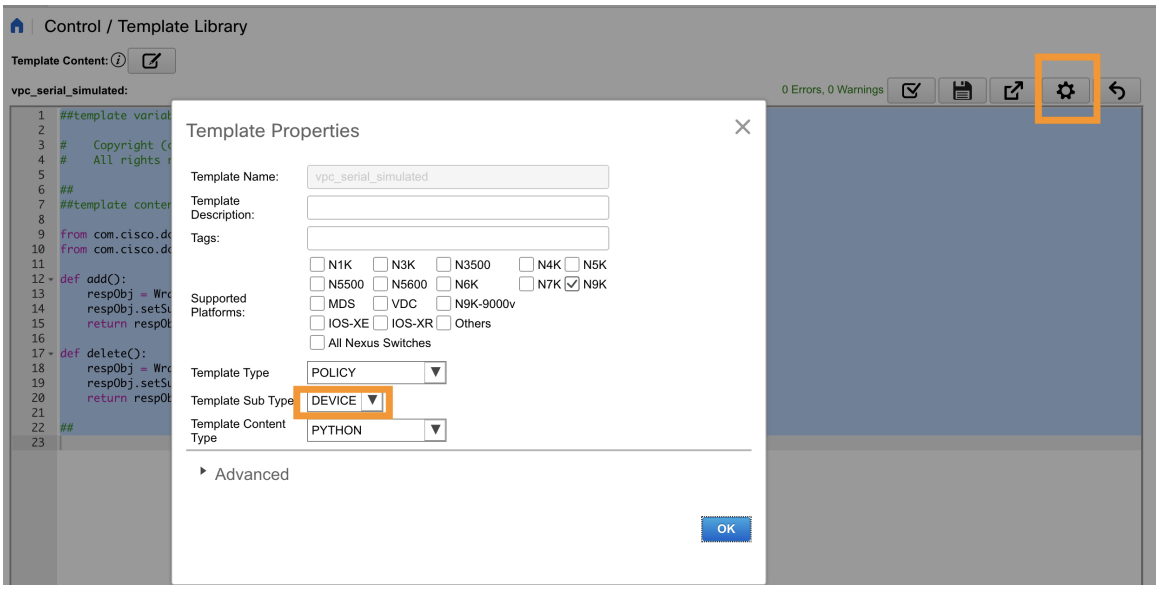

Step 4 Navigate to the **Fabric Builder** window and click on the fabric containing the vPC pair switches.

Step 5 Click **Tabular View** and select the vPC pair switches, and then click **View/Edit Policies**.

You can also right-click the switches individually in the topology and select **View/Edit Policies**.

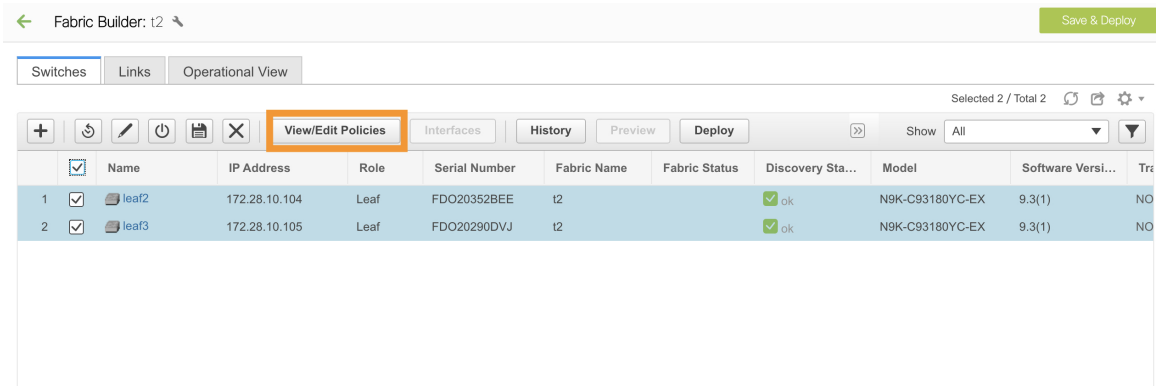

Step 6 Click + to add policies.

Step 7 From the **Policy** drop-down list, select **vpc_serial_simulated policy** and add priority. Click **Save**.

Note that if both switches are selected, then this policy will be created on both vPC pair switches.

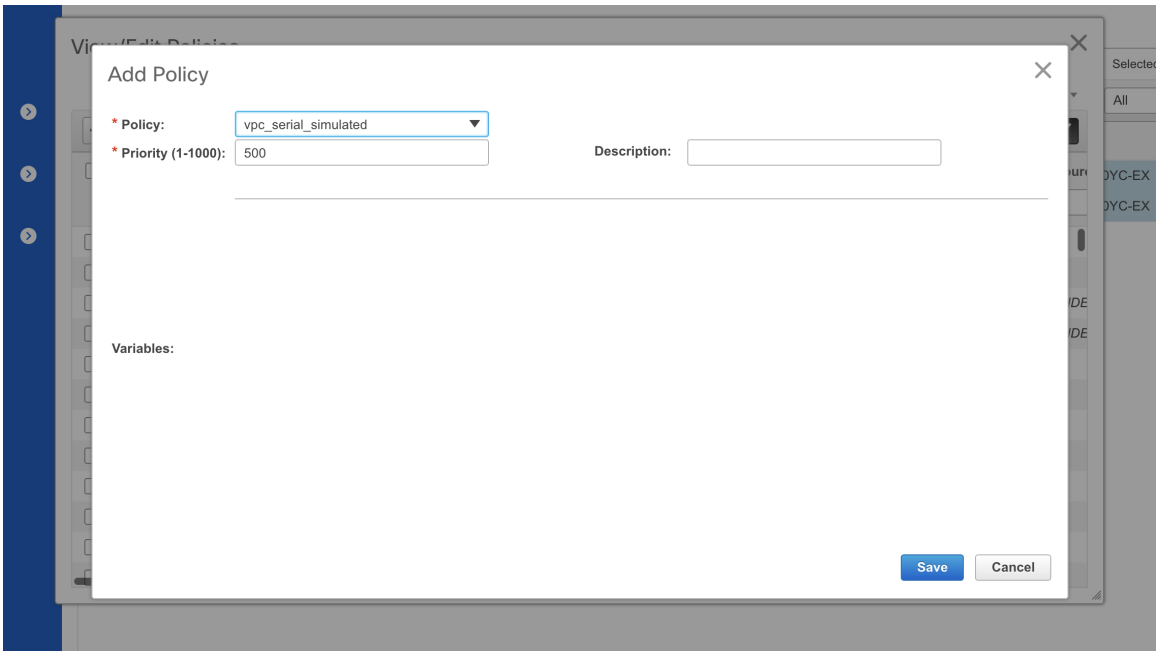

Step 8 Navigate back to **Tabular View** and click the **Links** tab.

Step 9 Select the link between vPC pair, which has to be a vPC peer keep alive and click **Edit**.

Step 10 From the **Link Template** drop-down list, select **int_intra_vpc_peer_keep_alive_link_11_1**.

Enter values for the remaining fields. Make sure to leave the field empty for the default VRF and click **Save**.

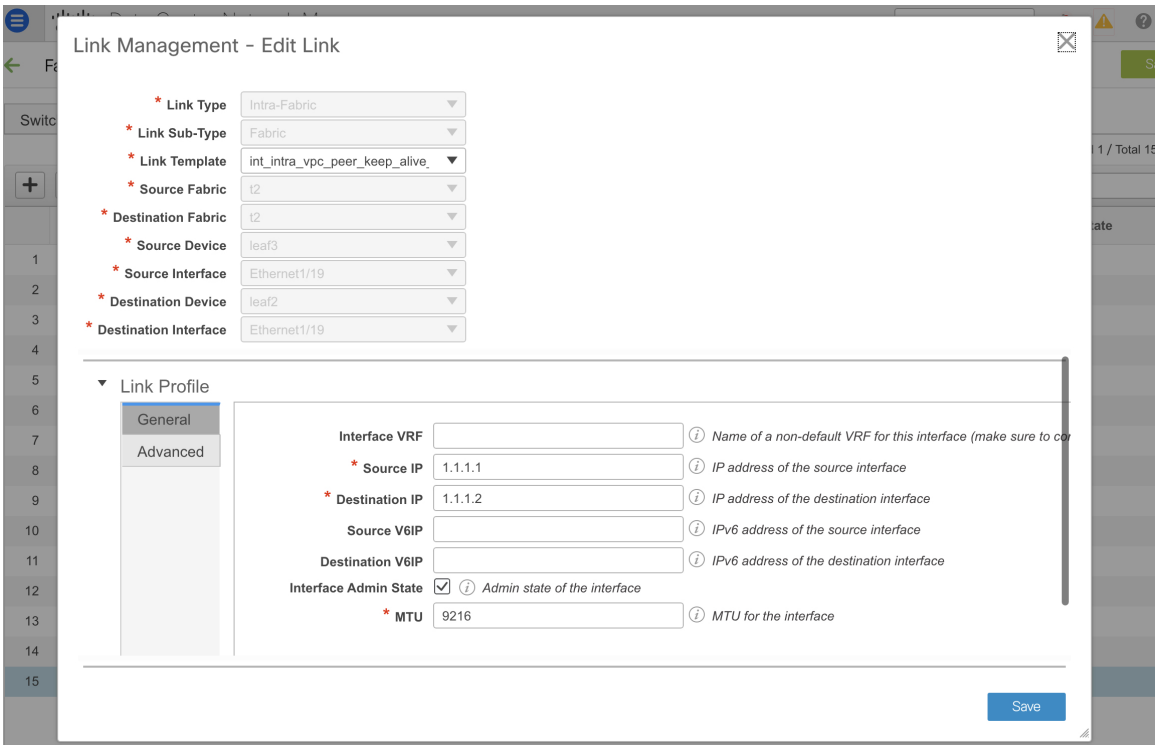

Step 11 Click **Save & Deploy**, and click **Preview Config** for one of the switches.

```
vpc domain 1
  ip arp synchronize
  peer-gateway
  peer-switch
  delay restore 150
  peer-keepalive destination 1.1.1.1 source 1.1.1.2 vrf default
  auto-recovery reload-delay 360
  ipv6 nd synchronize
interface port-channel500
```

If VRF is non-default, use **switch_freeform** to create the respective VRF.

Navigate to the topology and click the vPC pair switch to see the details.

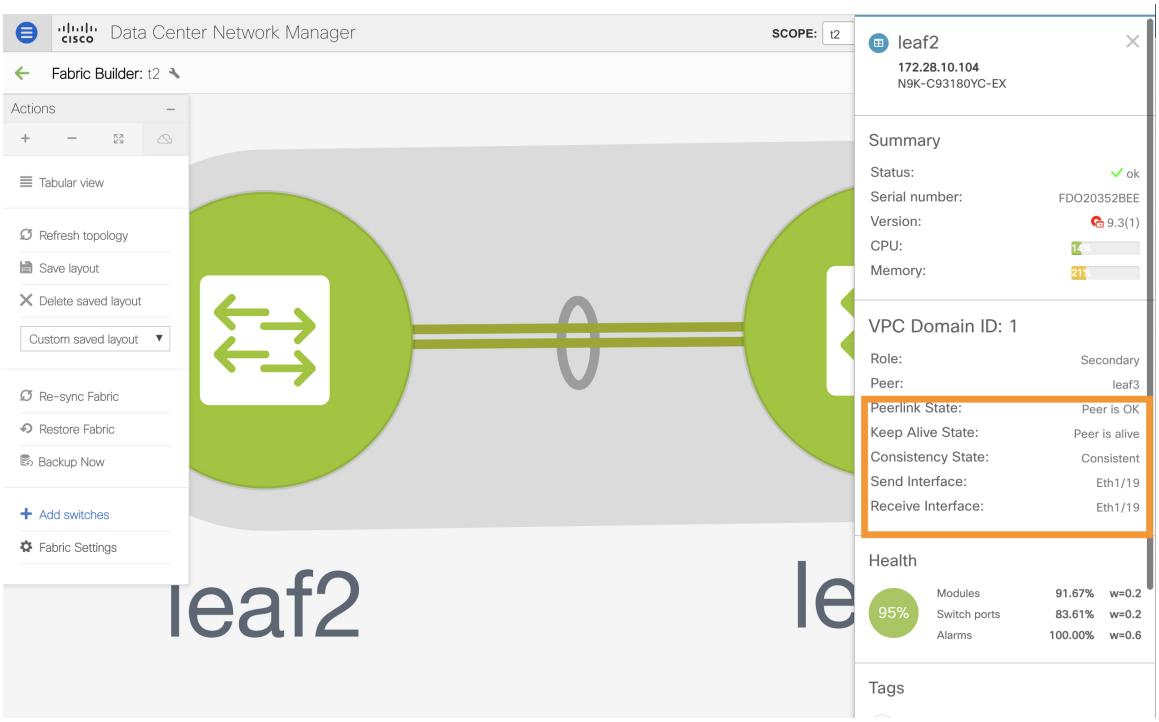

Brownfield Deployment-Transitioning VXLAN Fabric Management to DCNM

DCNM supports Brownfield deployments, wherein you transition your VXLAN BGP EVPN fabric management to DCNM. The transition involves migrating existing network configurations to DCNM. For information, see *Managing a Brownfield VXLAN BGP EVPN Fabric*.

Creating a New Fabric for EBGP-Based Underlay

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

   The **Fabric Builder** screen 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** screen, wherein a rectangular box represents each fabric.

   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).

   The technology is for a fabric with eBGP Routed Fabric or eBGP VXLAN EVPN Fabric. The mode of replication is only applicable for the eBGP VXLAN EVPN fabric, and not eBGP Routed fabric.

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

   The fields are explained:

   Fabric Name - Enter the name of the fabric.

   Fabric Template - From the drop-down menu, choose the **Easy_Fabric_eBGP** fabric template. The fabric settings for creating a standalone routed fabric comes up.

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

   BGP ASN for Spines: Enter the BGP AS number of the fabric's spine switches.

   BGP AS Mode: Choose **Multi-AS** or **Dual-AS**.

   In a **Multi-AS** fabric, the spine switches have a unique BGP AS number and each leaf switch has a unique AS number. If two leaf switches form a vPC switch pair, then they have the same AS number.

   In a **Dual-AS** fabric, the spine switches have a unique BGP AS number and the leaf switches have a unique AS number.

   The fabric is identified by the spine switch AS number.

   Underlay Subnet IP Mask - Specifies the subnet mask for the fabric interface IP addresses.

   Routing Loopback Id - The loopback interface ID is populated as 0 by default. It is used as the BGP router ID.

   Static Underlay IP Address Allocation – Check the check box to enable static IP address allocation for the fabric underlay.

   a. By default, DCNM allocates the underlay IP address resources (for loopbacks, fabric interfaces, etc) dynamically from the defined pools. If you select the check box, the allocation scheme switches to static, and some of the dynamic IP address range fields are disabled.

   b. For static allocation, the underlay IP address resources must be populated into the Resource Manager (RM) using REST APIs.

      See *Cisco DCNM REST API Guide* for more details. The REST APIs must be invoked after the switches are added to the fabric, and before you use the Save & Deploy option.

   c. Changing from static to dynamic allocation keeps the current IP resource usage intact. Only future IP address allocation requests are taken from dynamic pools.

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

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

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

NX-OS Software Image Version: Select an image from the drop-down 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 should 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. Till all devices run the specified image, the deployment process will be 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 **EVPN**. Most of the fields in this tab are auto-populated. The fields are:

Enable EVPN VXLAN Overlay: Enables the VXLAN overlay provisioning for the fabric.

You can convert a routed fabric to a VXLAN enabled fabric by selecting this option. When the fabric is VXLAN enabled, you can create and deploy overlay networks or VRFs. The procedure for creating and deploying networks or VRFs is the same as in Easy_Fabric_11_1. For more information, see *Creating and Deploying Networks and VRFs* in the Control chapter in *Cisco DCNM LAN Fabric Configuration Guide*.

Routed Fabric: You must disable the Enable EVPN VXLAN Overlay field for Routed fabric (an IP fabric with no VXLAN encapsulation) creation.

Whether you create an eBGP Routed or eBGP VXLAN fabric, the fabric uses eBGP as the control plane to build intra-fabric connectivity. Links between spine and leaf switches are autoconfigured with point-to-point (p2p) numbered IP addresses with eBGP peering built on top.

With an eBGP Routed Fabric, the VXLAN overlay fabric options such as creating networks/VRFs are disabled.

Note The rest of the fields in the EVPN tab section are only applicable if you enable the EVPN VXLAN Overlay.

Anycast Gateway MAC: Anycast gateway MAC address for the leaf switches.

VTEP Loopback Id: The loopback interface ID is populated as 1 since loopback1 is usually used for the VTEP peering purposes.

Enable VXLAN OAM: Enables the VXLAM OAM function for existing switches. This is enabled by default. Clear the check box to disable VXLAN OAM function.

If you want to enable the VXLAN OAM function on specific switches and disable on other switches in the fabric, you can use freeform configurations to enable OAM and disable OAM in the fabric settings.

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

Enable Tenant DHCP: Enables tenant DHCP support.

vPC advertise-pip: Check the check box to enable the Advertise PIP feature.

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

Multicast Group Subnet: IP address prefix used for multicast communication. A unique IP address is allocated from this group for each overlay network.

Enable Tenant Routed Multicast: Check the check box to enable Tenant Routed Multicast (TRM) as the fabric overlay multicast protocol.

Rendezvous-Points: Enter the number of spine switches acting as rendezvous points.

Replication mode: Choose from the two supported multicast modes of replication, ASM (for Any-Source Multicast [ASM]) or BiDir (for Bidirectional PIM [BIDIR-PIM]). When you choose ASM, the BiDir related fields are not enabled. When you choose BiDir, the BiDir related fields are enabled.

Note BIDIR-PIM is supported on Cisco's Cloud Scale Family platforms 9300-EX and 9300-FX/FX2, and software release 9.2(1) onwards.

Multicast address for TRM: The multicast address for Tenant Routed Multicast traffic is populated. By default, this address is from the IP prefix specified in the Multicast Group Subnet field. When you create a new VRF for the fabric overlay, this address is populated in the Underlay Multicast Address field, in the Advanced tab.

Underlay RP Loopback ID: The loopback ID used for the rendezvous point (RP), for multicast protocol peering purposes in the fabric underlay. The default is 254.

The following fields are enabled if you choose **bidir**. Depending on the RP count, either 2 or 4 phantom RP loopback ID fields are enabled.

- **Underlay Primary RP Loopback ID**: The primary loopback ID used for the phantom RP, for multicast protocol peering purposes in the fabric underlay.
- **Underlay Backup RP Loopback ID**: The secondary (or backup) loopback ID used for the phantom RP, for multicast protocol peering purposes in the fabric underlay.

The following Loopback ID options are applicable only when the RP count is 4.

- **Underlay Second Backup RP Loopback ID**: The second backup loopback ID used for the phantom RP, for multicast protocol peering purposes in the fabric underlay.
- **Underlay Third Backup RP Loopback ID**: The third backup loopback ID used for the phantom RP, for multicast protocol peering purposes in the fabric underlay.

VRF Template and VRF Extension Template: Specifies the VRF template for creating VRFs, and the VRF extension template for enabling VRF extension to other fabrics.

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

Underlay VTEP Loopback IP Range: Specifies the loopback IP address range for VTEPs.

Underlay RP Loopback IP Range: Specifies the anycast or phantom RP IP address range.

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.

VRF Lite Deployment: Specifies the VRF Lite method for extending inter fabric connections. Only the 'Manual' option is supported.

5. Click **vPC**. The fields in the tab are:

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

   vPC Peer Keep Alive option: Choose the management or loopback option. If you want to use IP addresses assigned to the management port and the management VRF, choose management. If you use IP addresses assigned to loopback interfaces (and a non-management VRF), choose loopback. If you use IPv6 addresses, you must use loopback IDs.

   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.

   vPC Peer Link Port Channel Number - Specifies the Port Channel ID for a vPC Peer Link. By default, the value in this field is 500.

   vPC IPv6 ND Synchronize: Enables IPv6 Neighbour Discovery synchronization between vPC switches. The check box is enabled by default. Clear the check box to disable the function.

6. Click the **Advanced** tab. The fields in the tab are:

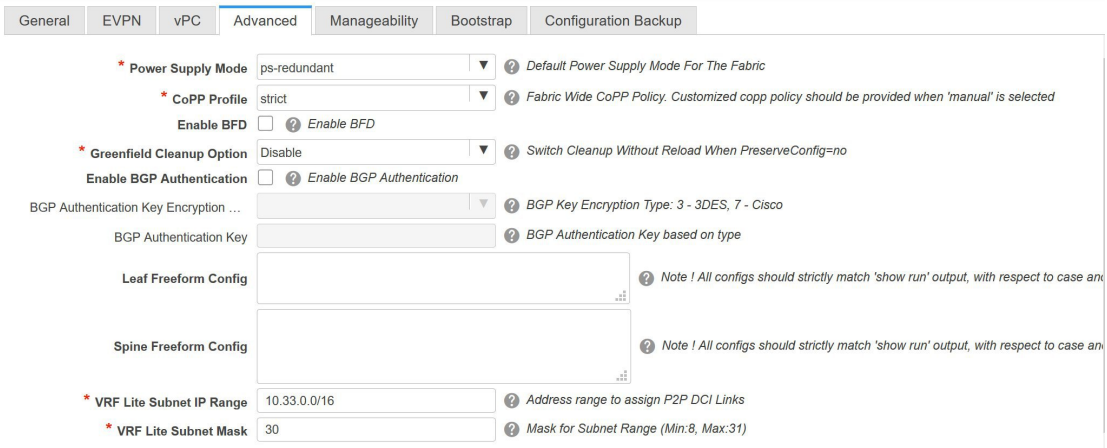

Power Supply Mode: Choose the appropriate power supply mode.

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

Enable BFD – Select the checkbox to enable **feature bfd** on all switches in the fabric.

Note Additional BFD related configurations must be added by using the appropriate freeform config fields.

The BFD feature is disabled by default.

Greenfield Cleanup Option: Enable the switch cleanup option for greenfield switches without a switch reload. This option is typically recommended only for the data center environments with the Cisco Nexus 9000v Switches.

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 Encryption Type and BGP Authentication Key fields are enabled.

BGP Authentication Key Encryption Type: Choose the 3 for 3DES encryption type, or 7 for Cisco encryption type.

BGP Authentication Key: Enter the encrypted key based on the encryption type.

Note Plain text passwords are not supported. Login to the switch, retrieve the encrypted key and enter it in the BGP Authentication Key field. Refer the Retrieving the Authentication Key section for details.

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, Border Spine, and Border Gateway Spine roles.

VRF Lite Subnet IP Range and **VRF Lite Subnet Mask** – These fields are populated with the DCI subnet details. Update the fields as needed.

7. Click the **Manageability** tab.

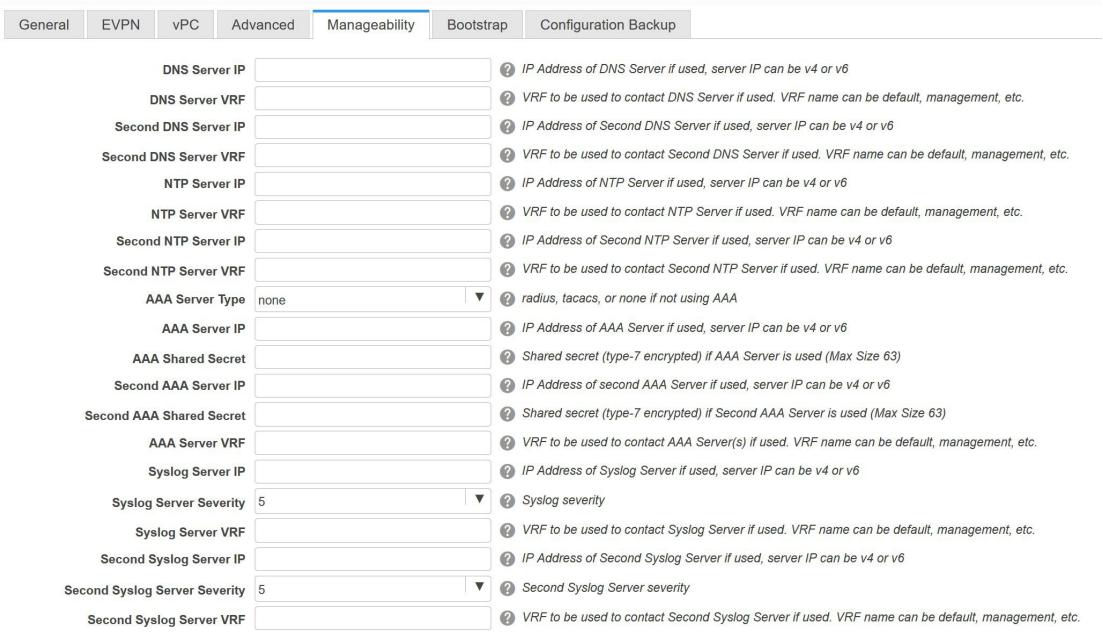

The fields in this tab are:

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.

Syslog Server IP – IP address of the syslog server, if used.

Syslog Server Severity – Severity level of the syslog server. To specify a higher severity, enter a higher number.

Syslog Server VRF – The default or management VRF that the syslog server IP address is assigned to.

Second Syslog Server IP – IP address of the second syslog server, if used.

Second Syslog Server Severity – Severity level of the second syslog server. To specify a higher severity, enter a higher number.

Second Syslog Server VRF – The default or management VRF that the second syslog server's IP address is assigned to.

8. Click the **Bootstrap** tab.

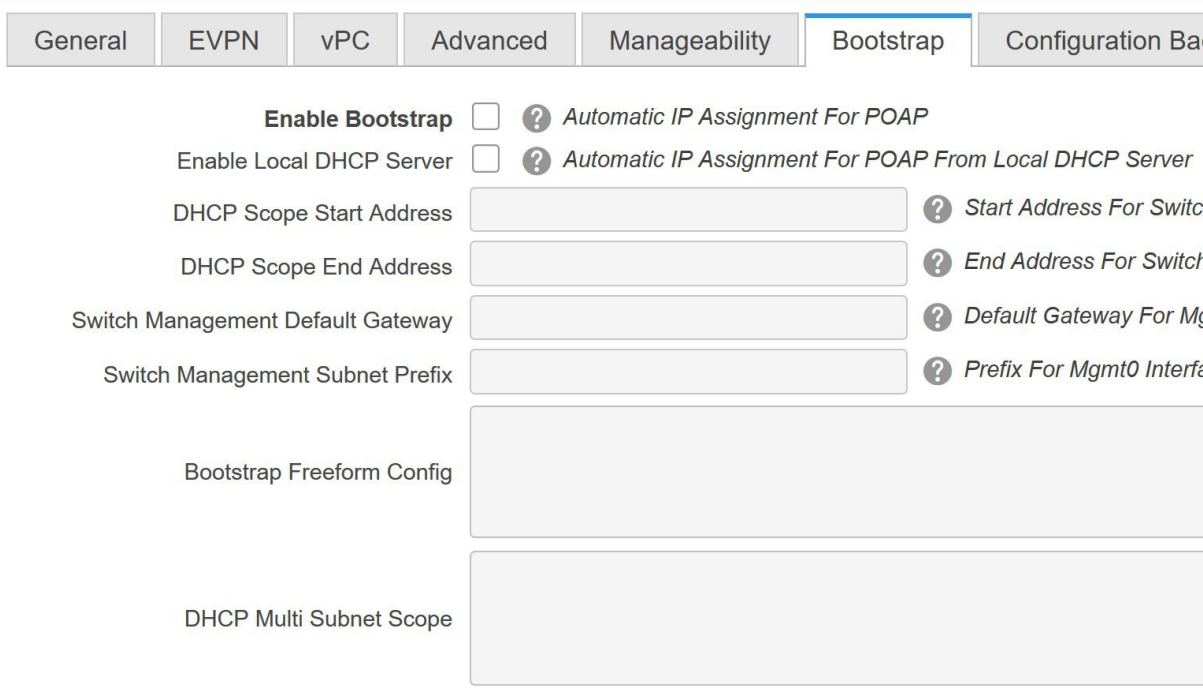

Enable Bootstrap - Select this check box to enable the bootstrap feature.

After you enable bootstrap, you can enable the DHCP server for automatic IP address assignment using one of the following methods:

- External DHCP Server: Enter information about the external DHCP server in the **Switch Management Default Gateway** and **Switch Management Subnet Prefix** fields.
- Local DHCP Server: Enable the **Local DHCP Server** checkbox and enter details for the remaining mandatory fields.

Enable Local DHCP Server - Select this check box to initiate enabling of automatic IP address assignment through the local DHCP server. When you select this check box, the **DHCP Scope Start Address** and **DHCP Scope End Address** fields become editable.

If you do not select this check box, DCNM uses the remote or external DHCP server for automatic IP address assignment.

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 10.0.1.1 and subnet mask 24, ensure that the DHCP scope is within the specified subnet, between 10.0.1.2 and 10.0.1.254..

Bootstrap Freeform Config - (Optional) Enter additional commands as needed. For example, if you are using AAA or remote authentication related configurations, you need to add these configurations in this field to save the intent. After the devices boot up, they contain the intent defined in the **Bootstrap Freeform Config** field.

Copy-paste the running-config to a **freeform config** field with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. For more information, see *Resolving Freeform Config Errors in Switches* in *Enabling Freeform Configurations on Fabric Switches*.

DHCP Multi Subnet Scope - Specifies the field to enter one subnet scope per line. This field is editable after you check the **Enable Local DHCP Server** check box.

The format of the scope should be defined as:

DHCP Scope Start Address, DHCP Scope End Address, Switch Management Default Gateway, Switch Management Subnet Prefix

For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24

9. Click the **Configuration Backup** tab. The fields on this tab are:

Hourly Fabric Backup: Select the check box to enable an hourly backup of fabric configurations and the intent. The backup process is initiated only when you click **Save and Deploy**, and the subsequent configuration compliance activity is successfully completed.

You can enable an hourly backup for fresh fabric configurations and the intent as well. If there is a configuration push in the previous hour, DCNM takes a backup.

Intent refers to configurations that are saved in DCNM but yet to be provisioned on the switches.

Scheduled Fabric Backup: Check the check box to enable a daily backup. This backup tracks changes in running configurations on the fabric devices that are not tracked by configuration compliance.

Scheduled Time: Specify the scheduled backup time in a 24-hour format. This field is enabled if you check the **Scheduled Fabric Backup** check box.

Select both the check boxes to enable both back up processes. If you update settings, execute the **Save & Deploy** option on the fabric topology screen (click within the fabric box to access the fabric topology screen).

Note Hourly and scheduled backup processes happen only during the next periodic configuration compliance activity, and there can be a delay of up to an hour. To trigger an immediate backup, do the following:

a. Choose **Control** > **Fabric Builder**. The Fabric Builder screen comes up.

b. Click within the specific fabric box. The fabric topology screen comes up.

c. From the **Actions** panel at the left part of the screen, click **Re-Sync Fabric**.

You can also initiate the fabric backup in the fabric topology window. Click **Backup Now** in the **Actions** pane.

Click **Save** after filling and updating relevant information.

VXLAN Fabric With eBGP Underlay – Pointers

- Deploy the leaf overlay and underlay policies on all leaf switches at once, since they have a common AS number.
- Brownfield migration is not supported for eBGP fabric.
- You cannot change the leaf switch AS number after it is created and the Save & Deploy operation is executed. You need to delete the **leaf_bgp_asn** policy and execute the Save & Deploy operation to remove BGP configuration related to this AS first. Then, you can add the leaf_bgp_asn policy with the new AS number.
- If you want to switch between Multi-AS and Dual-AS modes, remove all manually added BGP policies (including leaf_bgp_asn on the leaf switch and the ebgp overlay policies), and execute the **Save & Deploy** operation before the mode change.
- You cannot change or delete the leaf switch leaf_bgp_asn policy if there are ebgp overlay policies present on the device. You need to delete the ebgp overlay policy first, and then delete the leaf_bgp_asn policy.
- The supported roles are leaf, spine, and border leaf.
- On the border device, VRF-Lite is supported with manual mode. There is no Multi-Site support for external connectivity.
- TRM is supported.
- You must apply policies on the leaf and spine switches for a functional fabric.
- For a VXLAN enabled fabric, you can create and deploy overlay networks and VRFs the same way as in Easy Fabric. For more information, see *Creating and Deploying Networks and VRFs* in the Control chapter in *Cisco DCNM LAN Fabric Configuration Guide*.

Applying Policies On A Fabric With An eBGP Underlay

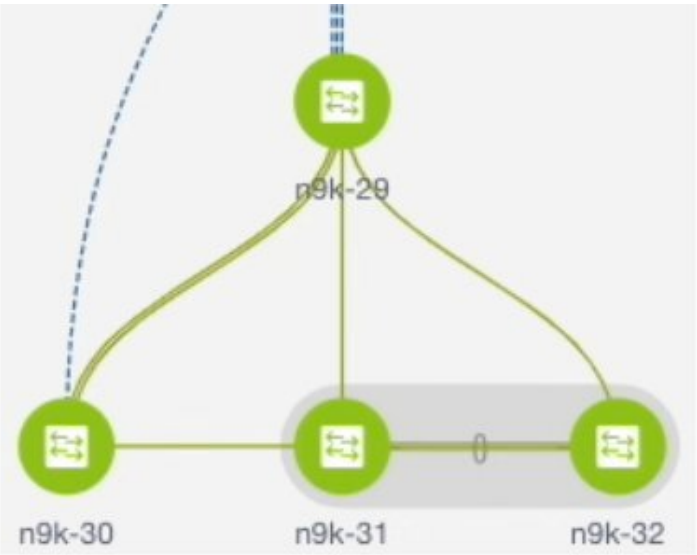

The topology shows a VXLAN fabric enabled with eBGP for the underlay. In DCNM, a fabric with the Easy_Fabric_eBGP template is created. One spine switch (n9k-29) and three leaf switches (n9k-30, and vPC switch pair n9k-31 and n9k-32) are imported to it.

This topic covers the following:

- **Creating a Multi-AS mode fabric**: This section mainly covers Multi-AS mode fabric creation. In a Multi-AS mode fabric, spine switches have a common BGP AS number and each leaf switch has a unique BGP AS number. Use the same steps for Dual-AS to Multi-AS mode fabric conversion.
- **Creating a Dual-AS mode fabric**: Alternate steps are mentioned for Dual-AS mode fabric creation. Use the same steps for Multi-AS to a Dual-AS mode fabric conversion.

  In a Dual-AS fabric, all spine switches have a common BGP AS number and all leaf switches have a common BGP AS number (differing from the spine switches' BGP AS number). You must deploy policies as explained in the next section.

Deploying Fabric Underlay eBGP Policies

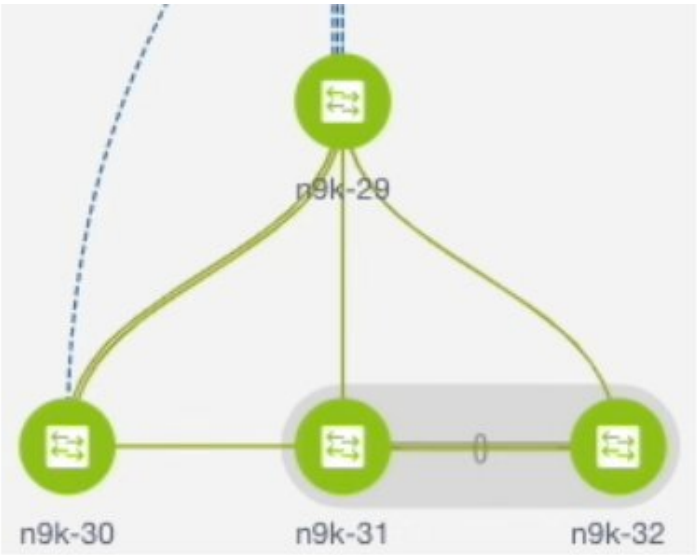

The topology shows a VXLAN fabric enabled with eBGP for the underlay. In DCNM, a fabric with the **Easy_Fabric_eBGP** template is created. One spine switch (n9k-29) and three leaf switches (n9k-30, and vPC switch pair n9k-31 and n9k-32) are imported to it.

The two different types of fabrics are:

- **Creating a Multi-AS mode fabric**: In a Multi-AS mode fabric, spine switches have a common BGP AS number and each leaf switch has a unique BGP AS number. Use the same steps for Dual-AS to Multi-AS mode fabric conversion.
- **Creating a Dual-AS mode fabric**: Alternate steps are mentioned for Dual-AS mode fabric creation. Use the same steps for Multi-AS to a Dual-AS mode fabric conversion.

  In a Dual-AS fabric, all spine switches have a common BGP AS number and all leaf switches have a common BGP AS number (differing from the spine switches' BGP AS number). You must deploy policies as explained in the next section.

To deploy fabric underlay eBGP policy, you must manually add the **leaf_bgp_asn** policy on each leaf switch to specify the BGP AS number used on the switch. Implementing the **Save & Deploy** operation afterward will generate eBGP peering over the physical interface between the leaf and spine switches to exchange underlay reachability information.

1. Click **Tabular View** at the left part of the screen. The **Switches | Links** screen comes up.
2. Select the leaf switch (n9k-30 check box for example) and click **View/Edit Policies**. The View/Edit Policies screen comes up.

Note When you create an eBGP fabric in the Dual-AS mode (or change from the Multi-AS mode to Dual-AS mode), select all leaf switches since they have a common BGP AS number.

3. Click **Add**. The **Add Policy** screen comes up.
4. From the Policy drop down box, select **leaf_bgp_asn** and enter the BGP AS number in the **BGP AS #** field.
5. Click **Save**.
6. Repeat the procedure for the vPC switches. For a vPC switch pair, select both switches and apply the **leaf_bgp_asn** policy.

Note This step is not needed if you create a fabric in the Dual-AS mode (or converting to the Dual-AS mode), and you have assigned a BGP AS number to all of them, as explained in the earlier steps.

7. Close the **View/Edit Policies** window.
8. In the topology screen, click **Save & Deploy** at the top right part of the screen.
9. Deploy configurations as per the **Config Deployment** wizard.

Deploying Fabric Overlay eBGP Policies

You must manually add the eBGP overlay policy for overlay peering. DCNM provides the eBGP leaf and spine overlay peering policy templates that you can manually add to the leaf and spine switches to form the EVPN overlay peering.

Deploying Spine Switch Overlay Policies

Add the ebgp_overlay_spine_all_neighbor policy on the spine switch n9k-29. This policy can be deployed on all spine switches at once, since they share the same field values.

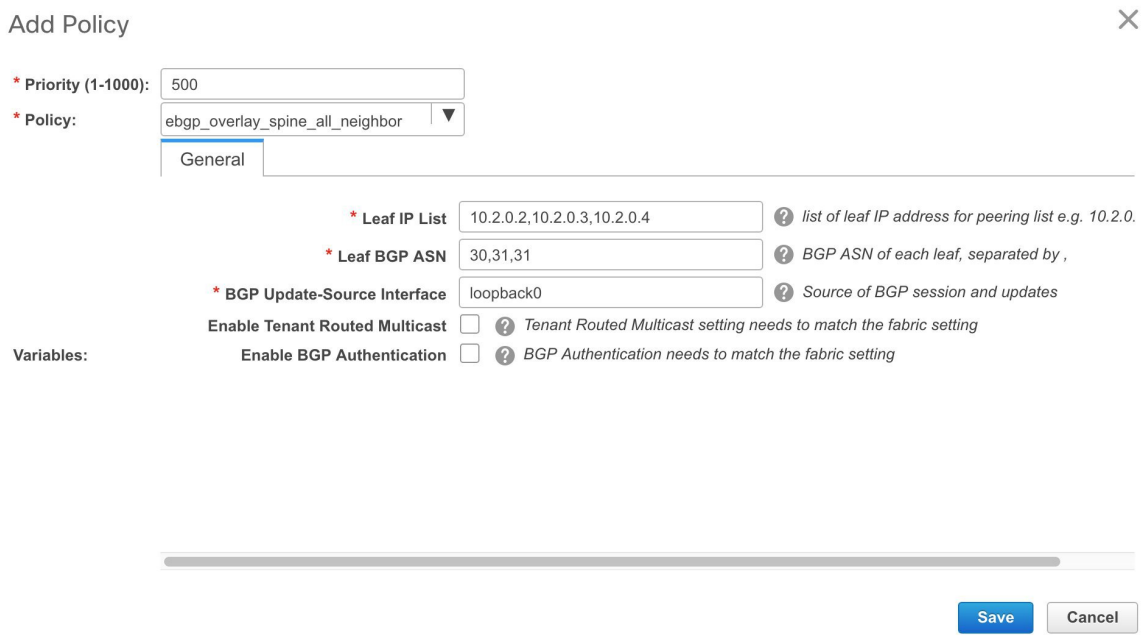

The fields on the screen are:

Leaf IP List - IP addresses of the connected leaf switch routing loopback interfaces.

10.2.0.2 is the loopback 0 peering IP address of leaf switch n9k-30. 10.2.0.3 and 10.2.0.4 are the IP addresses of the vPC switch pair n9k-31 and n9k-32.

Leaf BGP ASN – The BGP AS numbers of the leaf switches. Note that the AS number of vPC switches is the same, 31.

Note When you create fabric in the Dual-AS mode, (or convert to Dual-AS mode), you must update this field with the common BGP AS number all the leaf switches belong to.

BGP Update-Source Interface – This is the source interface of the BGP update. You can use loopback0 in this field, that is, the loopback interface for underlay routing.

Enable Tenant Routed Multicast – Select the checkbox to enable TRM for handling overlay multicast traffic. TRM enabling must match the fabric setting.

Enable BGP Authentication – Select the checkbox to enable BGP authentication.

The BGP authentication must match the fabric setting. Refer the Retrieving the Authentication Key section to know more about BGP authentication.

Deploying Leaf Switch Overlay Policies

Add the **ebgp_overlay_leaf_all_neighbor** policy on all the leaf switches, to establish eBGP overlay peering towards the spine switch. This policy can be deployed on all leaf switches at once, since they share the same field values.

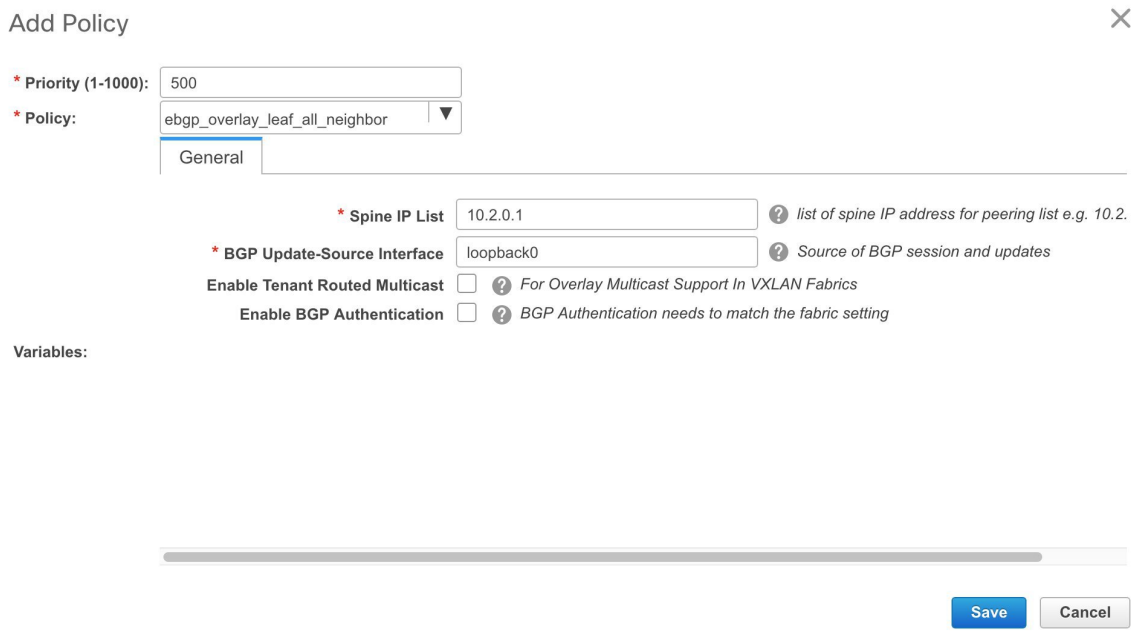

The fields on the screen are:

Spine IP List – IP addresses of the spine switch routing loopback interfaces.

10.2.0.1 is the loopback 0 peering IP address of spine switch n9k-29.

BGP Update-Source Interface – This is the source interface of the BGP update. You can use loopback0 in this field, that is, the loopback interface for underlay routing.

Enable Tenant Routed Multicast – Select the checkbox to enable TRM for handling overlay multicast traffic. TRM enabling must match the fabric setting.

Enable BGP Authentication – Select the checkbox to enable BGP authentication.

The BGP authentication must match the fabric setting. Refer the Retrieving the Authentication Key section to know more about BGP authentication.

Click **Save & Deploy** at the top right part of the screen, and deploy configurations as per the Config Deployment wizard. Or, use the **View/Edit Policy** option to select the policy and click **Push Config** to deploy the configuration.

Guidelines

- Deploy the leaf overlay and underlay policies on all leaf switches at once, since they have a common AS number.
- Brownfield migration is not supported for eBGP fabric.
- You cannot change the leaf switch AS number after it is created and the Save & Deploy operation is executed. You need to delete the **leaf_bgp_asn** policy and execute the Save & Deploy operation to remove BGP configuration related to this AS first. Then, you can add the leaf_bgp_asn policy with the new AS number.
- If you want to switch between Multi-AS and Dual-AS modes, remove all manually added BGP policies (including leaf_bgp_asn on the leaf switch and the ebgp overlay policies), and execute the **Save & Deploy** operation before the mode change.
- You cannot change or delete the leaf switch leaf_bgp_asn policy if there are ebgp overlay policies present on the device. You need to delete the ebgp overlay policy first, and then delete the leaf_bgp_asn policy.

Creating an External Fabric

In DCNM 11.1(1) release, you can add switches to the external fabric. Generic pointers:

- An external fabric is a monitor-only or managed 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.
- If an external fabric is an MSD fabric member, then the MSD topology screen displays the external fabric with its devices, along with the member fabrics and their devices.

  When viewed from an external fabric topology screen, any connections to non-DCNM managed switches are represented by a cloud icon labeled as **Undiscovered**.

- You can set up a Multi-Site or a VRF-lite IFC by manually configuring the links for the border devices in the VXLAN fabric or by using an automatic Deploy Border Gateway Method or VRF Lite IFC Deploy Method. If you are configuring the links manually for the border devices, we recommend using the Core Router role to set up a Multi-Site eBGP underlay from a Border Gateway device to a Core Router and the Edge Router role to set up a VRF-lite Inter-Fabric Connection (IFC) from a Border device to an Edge device.
- You can connect a Cisco data center to a public cloud using Cisco CSR 1000v. See the *Connecting Cisco Data Center and a Public Cloud* chapter for a use case.
- For the Cisco Network Insights for Resources (NIR) Release 2.1 and later, and flow telemetry, **feature lldp** command is one of the required configuration.

  Cisco DCNM pushes **feature lldp** on the switches only for the Easy Fabric deployments, that is, for the eBGP routed fabric or VXLAN EVPN fabric.

  Therefore, NIR users need to enable **feature lldp** on all the switches in the following scenarios:

  - External fabric in Monitored or Managed Mode

Creating External Fabric from Fabric Builder

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.

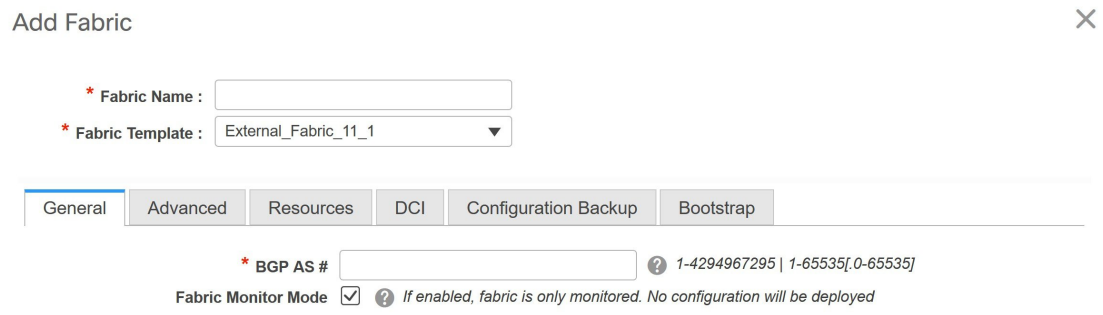

3. Fill up the General, Advanced, Resources, and DCI tabs as shown below.

   General tab

   BGP AS # - Enter the BGP AS number.

   Fabric Monitor Mode – Clear the checkbox if you want DCNM to manage the fabric. Keep the checkbox selected to enable a monitor only external fabric.

   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.

When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on its switches. If you click **Save & Deploy** in the fabric topology screen, it displays an error message.

However, the following settings (available when you right-click the switch icon) are allowed:

Advanced tab

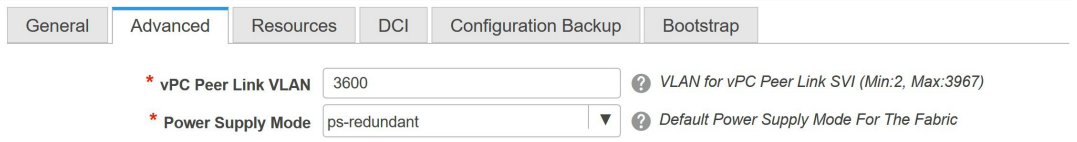

vPC Peer Link VLAN - The vPC peer link VLAN ID is autopopulated. Update the field to reflect the correct value.

Power Supply Mode - Choose the appropriate power supply mode.

Enable NX-API - Specifies enabling of NX-API on HTTPS.

Enable NX-API on HTTP - Specifies enabling of NX-API on HTTP. Enable this check box and the **Enable NX-API** check box to use HTTP.

Resources tab

Subinterface Dot1q Range - The subinterface 802.1Q range and the underlay routing loopback IP address range are autopopulated.

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

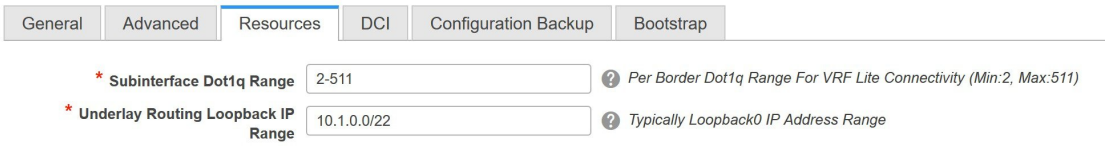

DCI tab – The DCI subnet IP prefix and subnet mask information are populated.

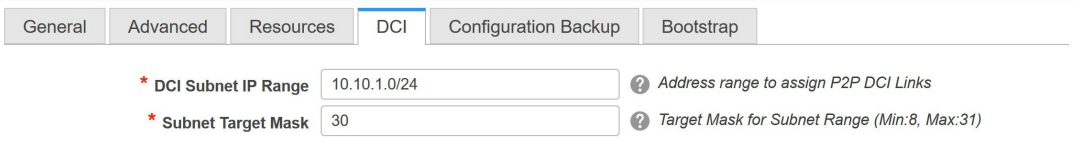

4. Click the **Configuration Backup** tab.

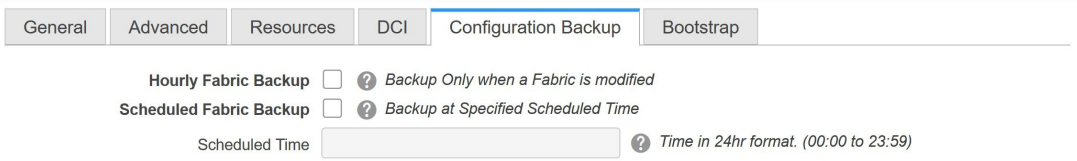

The fields on this tab are:

Hourly Fabric Backup: Select the check box to enable an hourly backup of fabric configurations and the intent. The backup process is initiated only when you click **Save and Deploy**, and the subsequent configuration compliance activity is successfully completed.

You can enable an hourly backup for fresh fabric configurations and the intent as well. If there is a configuration push in the previous hour, DCNM takes a backup. In case of the external fabric, the entire configuration on the switch is not converted to intent on DCNM as compared to the VXLAN fabric. Therefore, for the external fabric, both intent and running configuration are backed up.

Intent refers to configurations that are saved in DCNM but yet to be provisioned on the switches.

Scheduled Fabric Backup: Check the check box to enable a daily backup. This backup tracks changes in running configurations on the fabric devices that are not tracked by configuration compliance.

Scheduled Time: Specify the scheduled backup time in a 24-hour format. This field is enabled if you check the **Scheduled Fabric Backup** check box.

Select both the check boxes to enable both back up processes. If you update settings, execute the **Save & Deploy** option on the fabric topology screen (click within the fabric box to access the fabric topology screen).

You can also initiate the fabric backup in the fabric topology window. Click **Backup Now** in the **Actions** pane.

Pointers for hourly and scheduled backup:

- If you update a field in the Configuration Backup Tab, execute the Save & Deploy option on the fabric topology screen (click within the fabric box in the Fabric Builder screen to go to the fabric topology screen).
- The backups contain running configuration and intent pushed by DCNM. Configuration compliance forces the running config to be the same as the DCNM config. Note that for the external fabric, only some configurations are part of intent and the remaining configurations are not tracked by DCNM. Therefore, as part of backup, both DCNM intent and running config from switch are captured.
- The backups happen only during the next periodic configuration compliance activity, and there can be a delay of up to an hour.
- If you encounter an error during a device backup in a fabric, the backup for the entire fabric fails.

5. Click **Save**.

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

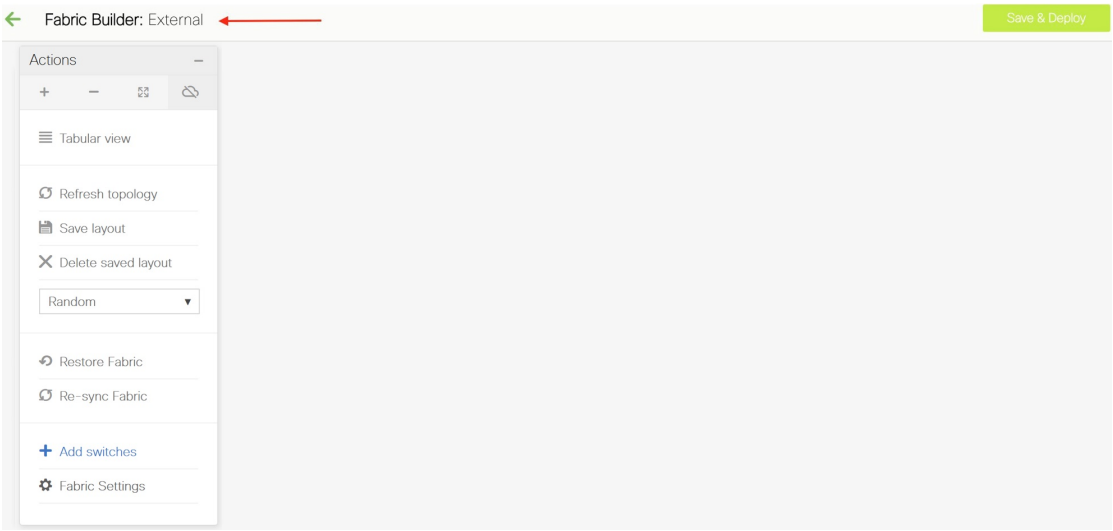

After creating the external fabric, add switches to it.

Add Switches to the External Fabric

1. Click Add switches. The Inventory Management screen comes up.

   You can also add switches by clicking Tabular View > Switches > + .

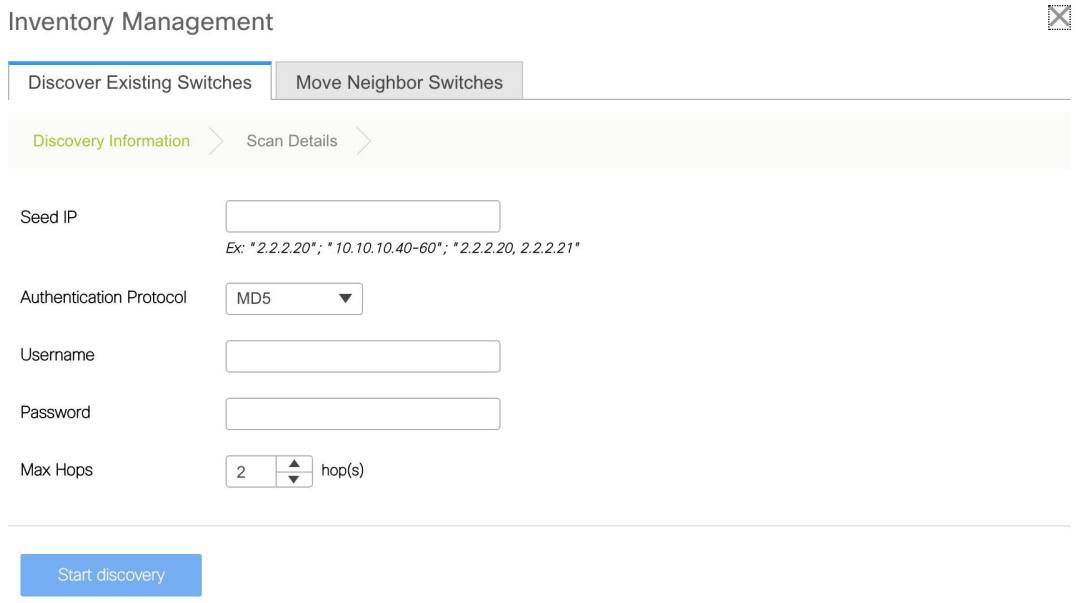

2. Enter the IP address (Seed IP) of the switch.
3. Enter the administrator username and password of the switch.
4. Click Start discovery at the bottom part of the screen. The Scan Details section comes up shortly. Since the Max Hops field was populated with 2, the switch with the specified IP address and switches two hops from it are populated.
5. Select the check boxes next to the concerned switches and click Import into fabric.

   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 fabric screen comes up again. The switch icons are seen at the centre of the fabric screen.
6. Click Refresh topology to view the latest topology view.
7. *External Fabric Switch Settings* - The settings for external fabric switches vary from the VXLAN fabric switch settings. Right-click on the switch icon and set or update switch options.

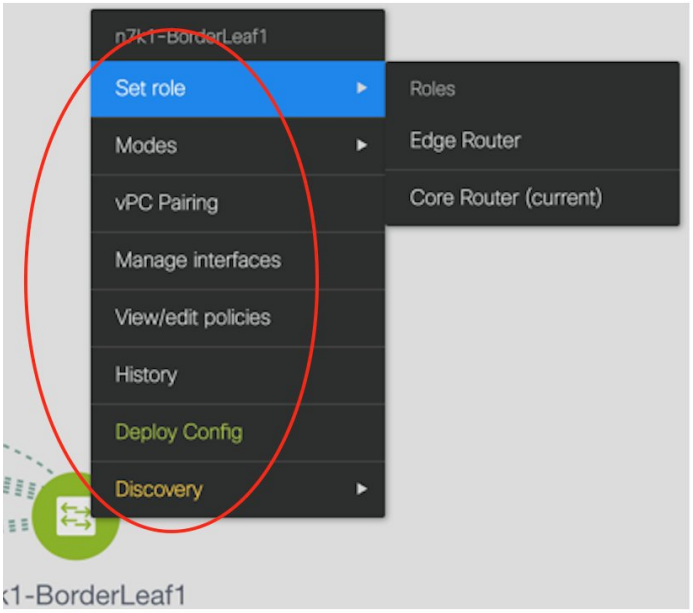

The options are:

Set Role – By default, no role is assigned to an external fabric switch. The allowed roles are Edge Router and Core Router. Assign the Core Router role for a Multi-Site Inter-Fabric Connection (IFC) and the Edge Router role for a VRF Lite IFC between the external fabric and VXLAN fabric border devices.

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

Modes – Active/Operational mode.

vPC Pairing – Select a switch for vPC and then select its peer.

Manage Interfaces – Deploy configurations on the switch interfaces.

Straight-through FEX, Active/Active FEX, and breakout of interfaces are not supported for external fabric switch interfaces.

View/edit Policies – Add, update, and delete policies on the switch. The policies you add to a switch are template instances of the templates available in the template library. After creating policies, deploy them on the switch using the Deploy option available in the View/edit Policies screen.

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.

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

   When you click Save & Deploy, the Configuration Deployment screen comes up.

9. Click Deploy Config at the bottom part of the screen to initiate pending configuration onto the switch.
10. Close the screen after deployment is complete.

Note If a switch in an external fabric does not accept default credentials, you should perform one of the following actions:

- Remove the switch in the external fabric from inventory, and then rediscover.
- LAN discovery uses both SNMP and SSH, so both passwords need to be the same. You need to change the SSH password to match the SNMP password on the switch. If SNMP authentication fails, discovery is stopped with authentication error. If SNMP authentication passes but SSH authentication fails, DCNM discovery continues, but the switch status shows a warning for the SSH error.

Move an External Fabric Under an MSD Fabric

You should go to the MSD fabric page to associate an external fabric as its member.

1. Click Control > Fabric Builder to go to the Fabric Builder screen.
2. Click within the MSD-Parent-Fabric box to go to its topology screen.
3. In the topology screen, go to the Actions panel and click Move Fabrics.

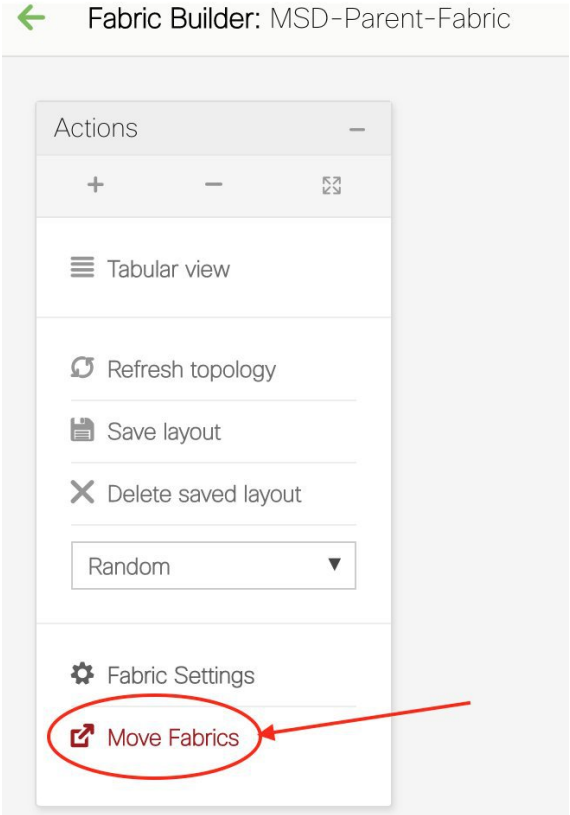

The Move Fabric screen comes up. It contains a list of fabrics. The external fabric is displayed as a standalone fabric.

4. Select the radio button next to the external fabric and click Add.

   Now, in the Scope drop-down box at the top right, you can see that the external fabric appears under the MSD fabric.

5. Click ← at the top left part of the screen to go to the Fabric Builder screen. In the MSD fabric box's Member Fabrics field, the external fabric is displayed.

External Fabric Depiction in an MSD Fabric Topology

The MSD topology screen displays MSD member fabrics and external fabrics together. The external fabric External65000 is displayed as part of the MSD topology.

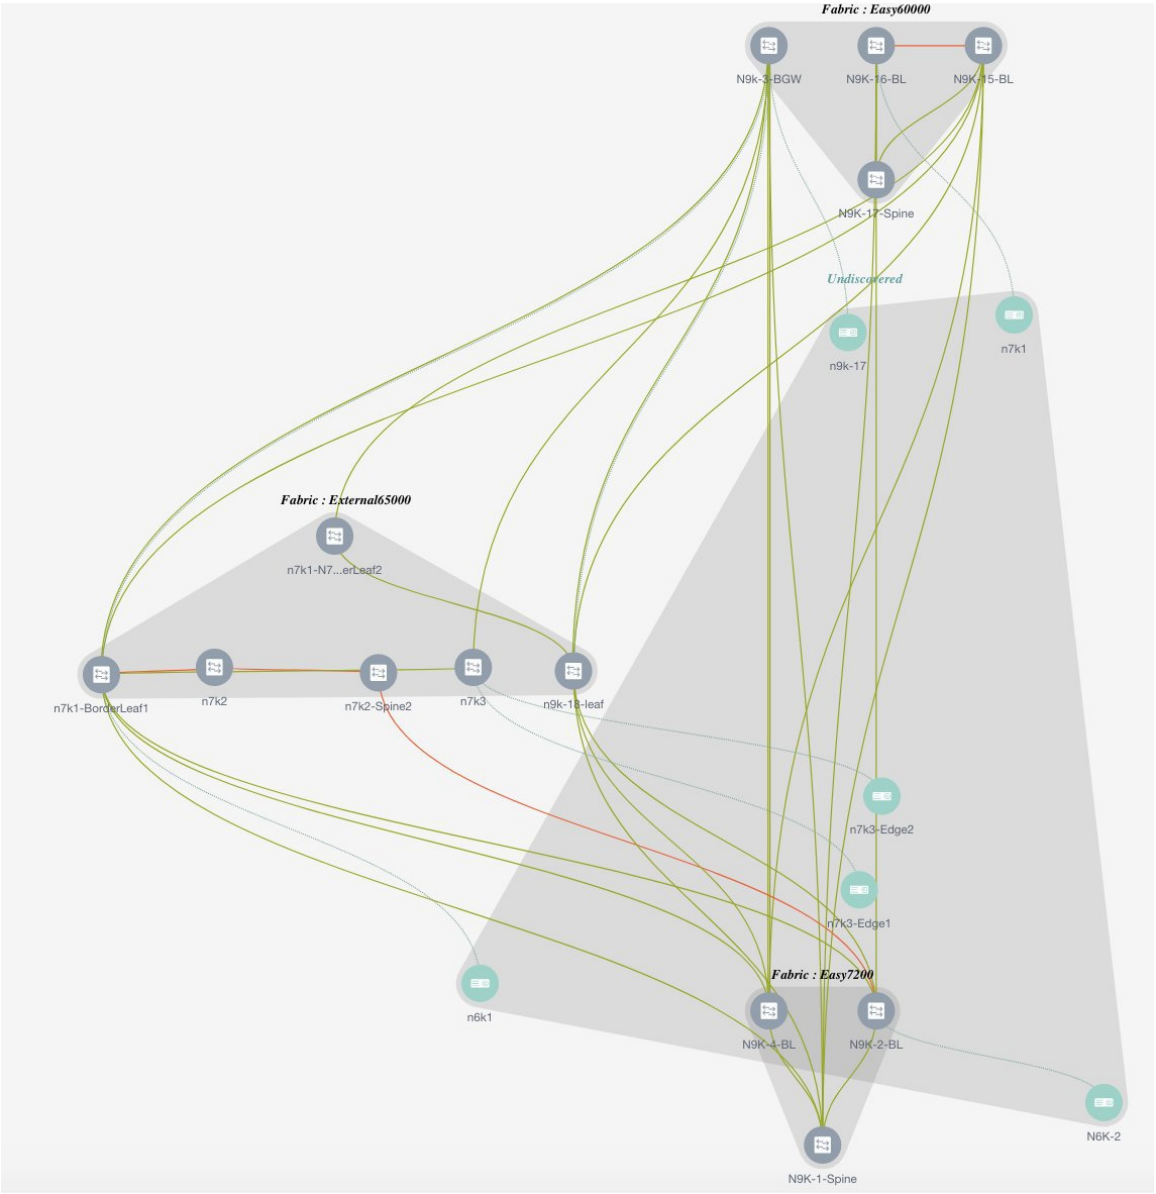

Note When you deploy networks or VRFs for the VXLAN fabric, the deployment page (MSD topology view) shows the VXLAN and external fabrics that are connected to each other.

External Fabric Switch Operations

In the external fabric topology screen, click Tabular view option in the Actions panel, at the left part of the screen. The Switches | Links screen comes up.

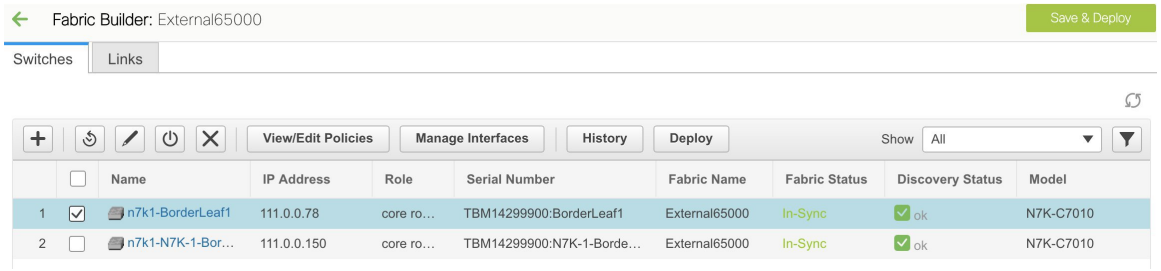

The Switches tab is for managing switch operations and the Links tab is for viewing fabric links. Each row represents a switch in the external fabric, and displays switch details, including its serial number.

The buttons at the top of the table are explained, from left to right direction. Some options are also available when you right-click the switch icon. However, the Switches tab enables you to provision configurations on multiple switches (for adding and deploying policies, and so on) simultaneously.

- Add switches to the fabric. This option is also available in the topology page (Add switches option in Actions panel).
- Initiate the switch discovery process by DCNM afresh.
- Update device credentials such as authentication protocol, username, and password.
- Reload the switch.
- Remove the switch from the fabric.
- View/edit Policies – Add, update, and delete a policy on multiple switches simultaneously. The policies are template instances of templates in the template library. After creating a policy, deploy it on the switches using the Deploy option available in the View/edit Policies screen.

  Note If you select multiple switches and deploy a policy instance, then it will be deployed on all the selected switches.

- Manage Interfaces – Deploy configurations on the switch interfaces.
- History – View deployment history on the selected switch.
- Deploy – Deploy switch configurations.

External Fabric Links

You can only view and delete external fabric links. You cannot create links or edit them.

To delete a link in the external fabric, do the following:

1. Go to the topology screen and click the Tabular view option in the Actions panel, at the left part of the screen.

   The Switches | Links screen comes up.

2. Choose one or more checkboxes and click the Delete icon at the top left.

   The links are deleted.

Move Neighbor Switch to External Fabric

1. Click Add switches. The Inventory Management screen comes up.
2. Click Move Neighbor Switches tab.
3. Select the switch and click Move Neighbor at the top right part of the screen.

   To delete a neighbor, select a switch and click Delete Neighbor at the top right.

Discovering New Switches

To discover new switches, perform the following steps:

Procedure

Step 1 Power on the new switch in the external fabric after ensuring that it is cabled to the DCNM server.

Boot the Cisco NX-OS and setup switch credentials.

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

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

Step 3 On the DCNM UI, choose **Control > Fabric Builder**.

The **Fabric Builder** screen is displayed. It contains a list of fabrics wherein a rectangular box represents each fabric.

Step 4 Click **Edit Fabric** icon at the top right part of the fabric box.

The **Edit Fabric** screen is displayed.

Step 5 Click the **Bootstrap** tab and update the DHCP information.

Step 6 Click **Save** at the bottom right part of the Edit Fabric screen to save the settings.

Step 7 In the Fabric Builder screen, click within the fabric box.

The fabric topology screen appears.

Step 8 In the fabric topology screen, from the Actions panel at the left part of the screen, click **Add switches**.

The Inventory Management screen comes up.

Step 9 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 management IP address, hostname, and password are made available. If the switch information is not retrieved, refresh the screen using the Refresh icon at the top right part of 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. You can pre-provision a device using the import option too.

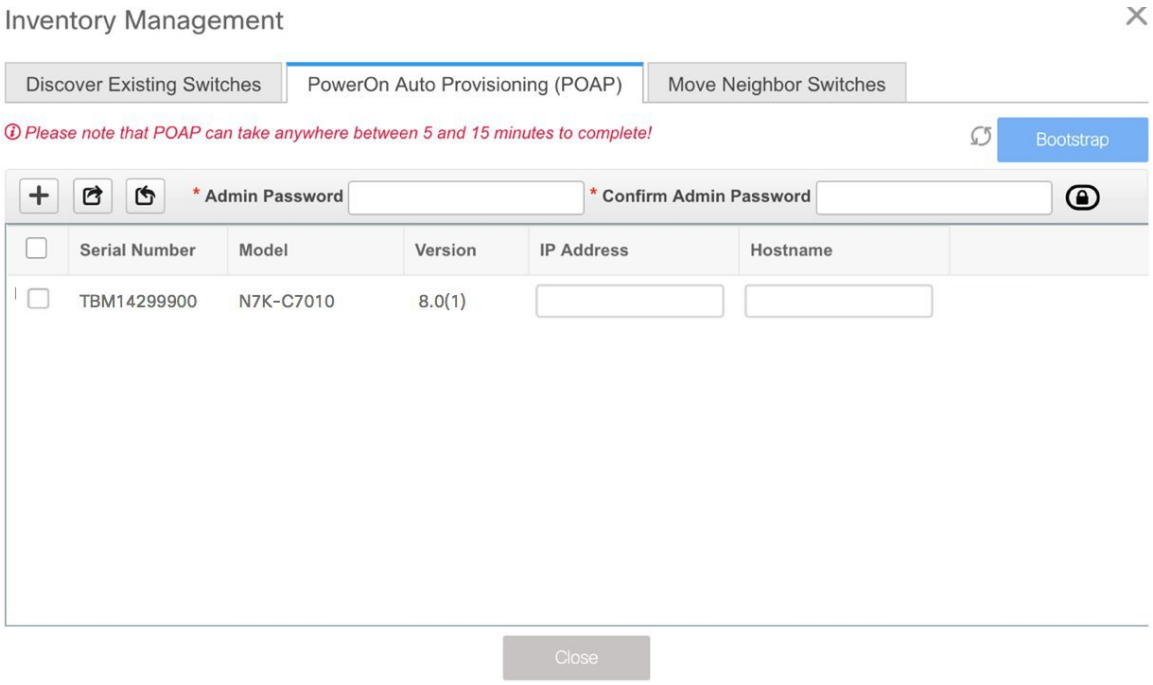

Select the checkbox next to the switch and add switch credentials: IP address and host name.

Beginning with Release 11.2(1), you can provision devices in advance. To pre-provision devices, refer to Pre-provisioning a Device , on page 33.

Step 10 In the **Admin Password** and **Confirm Admin Password** fields, enter and confirm the admin password.

This admin password is applicable for all the switches displayed in the POAP window.

Note If you do not want to use admin credentials to discover switches, you can instead use the AAA authentication, that is, RADIUS or TACACS credentials for discovery only.

Step 11 (Optional) Use discovery credentials for discovering switches.

a) Click the **Add Discovery Credentials** icon to enter the discovery credentials for switches.

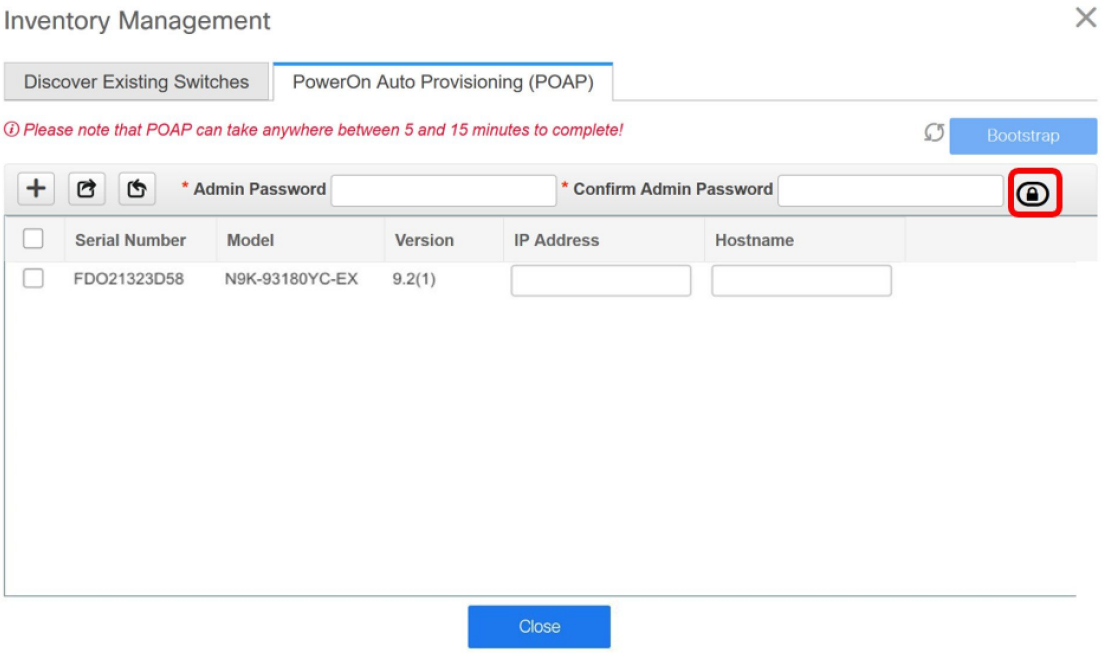

b) In the **Discovery Credentials** window, enter the discovery credentials such as discovery username and password.

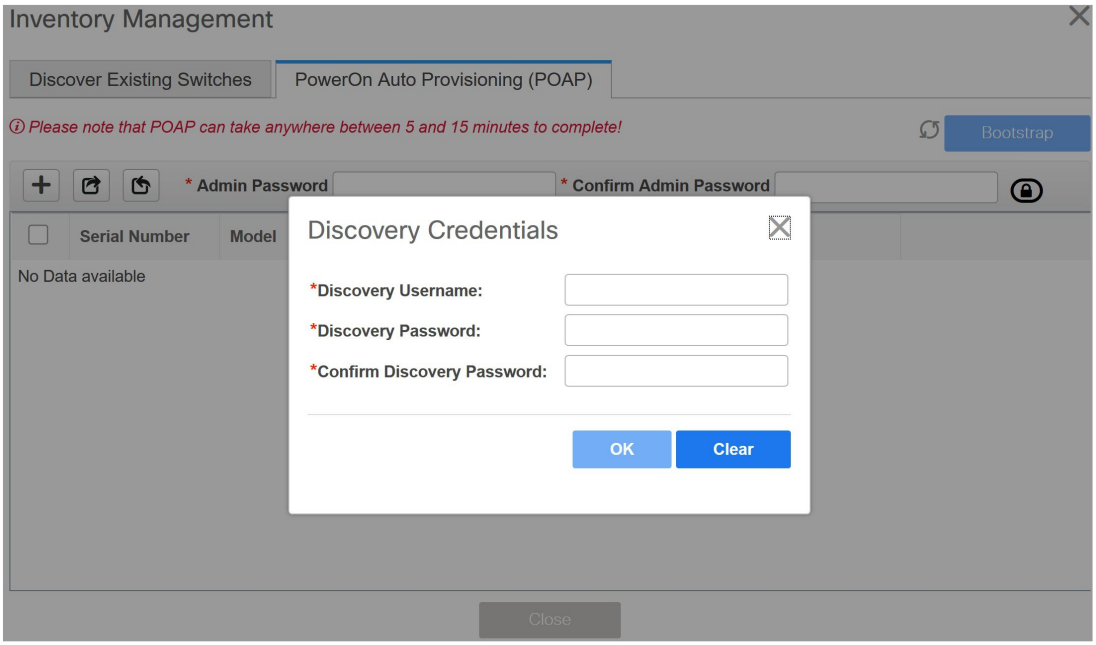

Click **OK** to save the discovery credentials.

If the discovery credentials are not provided, DCNM uses the admin user and password to discover switches.

Note
- The discovery credentials that can be used are AAA authentication based credentials, that is, RADIUS or TACACS.
- The discovery credential is not converted as commands in the device configuration. This credential is mainly used to specify the remote user (or other than the admin user) to discover the switches. If you want to add the commands as part of the device configuration, add them in the **Bootstrap Freeform Config** field under the **Bootstrap** tab in the fabric settings. Also, you can add the respective policy from **View/Edit Policies** window.

Step 12 Click **Bootstrap** at the top right part of the screen.

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

Step 13 After the bootstrapping is complete, close the **Inventory Management** screen to go to the fabric topology screen.

Step 14 In the fabric topology screen, from the **Actions** panel at the left part of the screen, click **Refresh Topology**.

After the added switch completes POAP, the fabric builder topology screen displays the added switch with some physical connections.

Step 15 Monitor and check the switch for POAP completion.

Step 16 Click **Save & Deploy** at the top right part of the fabric builder topology screen to deploy pending configurations (such as template and interface configurations) onto the switches.

Note
- If there is a sync issue between the switch and DCNM, the switch icon is displayed in red color, indicating that the fabric is Out-Of-Sync. For any changes on the fabric that results in the out-of-sync, you must deploy the changes. The process is the same as explained in the Discovering Existing Switches section.
- The discovery credential is not converted as commands in the device configuration. This credential is mainly used to specify the remote user (or other than the admin user) to discover the switches. If you want to add the commands as part of the device configuration, add them in the **Bootstrap Freeform Config** field under the **Bootstrap** tab in the fabric settings. Also, you can add the respective policy from **View/Edit Policies** window.

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.

Step 17 After the pending configurations are deployed, the **Progress** column displays 100% for all switches.

Step 18 Click **Close** to return to the fabric builder topology.

Step 19 Click **Refresh Topology** to view the update.

All switches must be in green color indicating that they are functional.

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.

Step 20 Right-click and select History to view the deployed configurations.

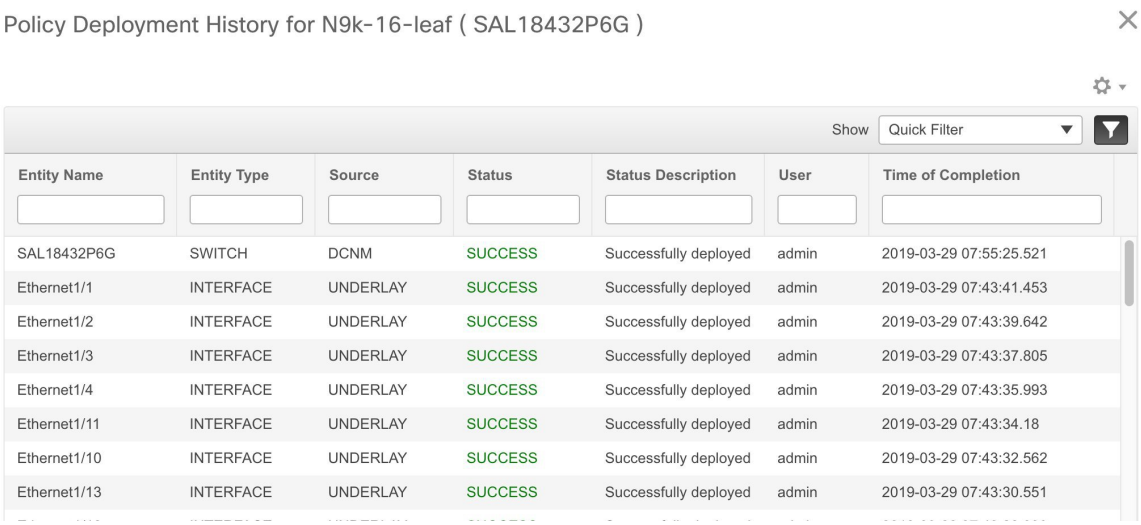

Click the **Success** link in the **Status** column for more details. An example:

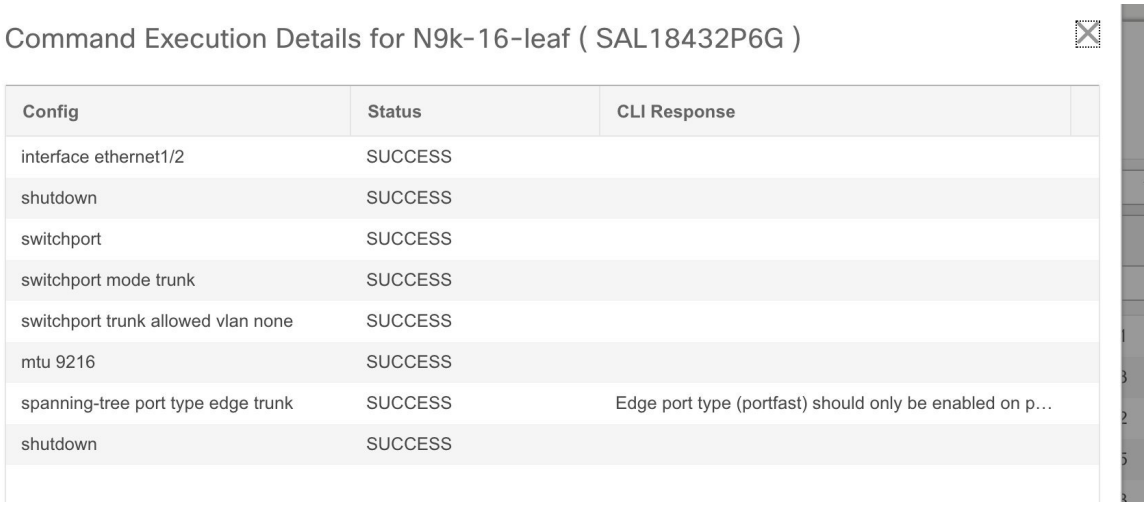

Step 21 On the DCNM UI, the discovered switches can be seen in the fabric topology.

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

  Support for breakout interfaces is available for 9000 Series switches.
- 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.

Pre-provisioning a Device

In DCNM 11.2, you can provision devices in advance.

Note Ensure that you enter DHCP details in the Bootstrap tab in the fabric settings.

- The pre-provisioned devices support the following configurations in DCNM:
  - Base management
  - vPC Pairing
  - Intra-Fabric links
  - Interface breakout configuration
- The pre-provisioned devices do not support the following configurations in DCNM:
  - Inter-Fabric links
  - Host ports
  - vPCs to the access switches or hosts
  - FEX
  - Overlay network configurations
- When a device is being pre-provisioned has breakout links, you need to specify the corresponding breakout command along with the switch's model and gateway in the **Data** field in the **Add a new device to pre-provisioning** window in order to generate the breakout PTI.

  Note the following guidelines:
  - Multiple breakout commands can be separated by a semicolon (;).
  - The definitions of the fields in the data JSON object are as follows:
    - **modulesModel**: (Mandatory) Specifies the switch module’s model information.
    - **gateway**: (Mandatory) Specifies the default gateway for the management VRF on the switch. This field is required to create the intent to pre-provision devices. You need to enter the gateway even if it is in the same subnet as DCNM to create the intent as part of pre-provisioning a device.
    - **breakout**: (Optional) Specifies the breakout command provided in the switch.
    - **portMode**: (Optional) Specifies the port mode of the breakout interface.

  The examples of the values in the **Data** field are as follows:
  - {"modulesModel": ["N9K-C93180LC-EX"], "gateway": "10.1.1.1/24"}
  - {"modulesModel": ["N9K-C93180LC-EX"],"breakout": "interface breakout module 1 port 1 map 10g-4x", "portMode": "hardware profile portmode 4x100G+28x40G", "gateway": "172.22.31.1/24" }

- {"modulesModel": ["N9K-X9736C-EX", "N9K-X9732C-FX", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-SUP-B+", "N9K-SC-A", "N9K-SC-A"], "gateway": "172.22.31.1/24"}
- {"breakout":"interface breakout module 1 port 50 map 10g-4x" , "gateway": "172.16.1.1/24", "modulesModel": ["N9K-C93180YC-EX "]}
- {"modulesModel": ["N9K-X9732C-EX", "N9K-X9732C-EX", "N9K-C9504-FM-E", "N9K-C9504-FM-E", "N9K-SUP-B", "N9K-SC-A", "N9K-SC-A"], "gateway": "172.29.171.1/24", "breakout":"interface breakout module 1 port 1,11,19 map 10g-4x; interface breakout module 1 port 7 map 25g-4x"}
- {"modulesModel": ["N9K-C93180LC-EX"], "gateway": "10.1.1.1/24","breakout": "interface breakout module 1 port 1-4 map 10g-4x","portMode": "hardware profile portmode 48x25G + 2x100G + 4x40G"}

Procedure

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

The **Fabric Builder** screen is displayed.

Step 2 Click within the fabric box.

Step 3 From the Actions panel, click the **Add switches** option.

The **Inventory Management** screen is displayed.

Step 4 Click the **POAP** tab.

Step 5 In the **POAP** tab, do the following:

a. Click + from the top left part of the screen.

The Add a new device screen comes up.

b. Fill up the device details as shown in the screenshot.

c. Click **Save**.

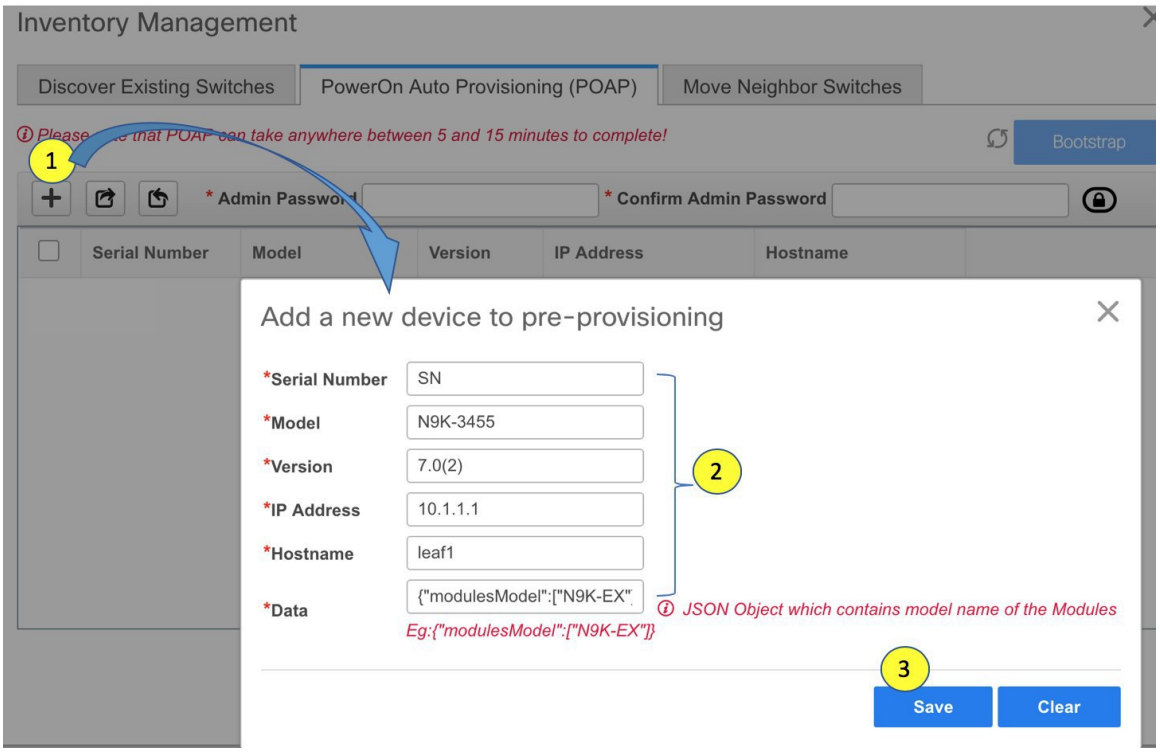

Serial Number: The serial number for the new device. This number can be a dummy serial number if the device serial number is not available.

For information about the **Data** field, see the examples provided in guidelines.

The device details appear in the POAP screen. You can add more devices for pre-provisioning.

At the top left part of the window, **Export** and **Import** icons are provided to export and import the .csv file that contains the switch information.

Using the **Import** option, you can pre-provision multiple devices.

Add new devices' information in the .csv file with all the mandatory fields (SerialNumber, Model, version, IpAddress, Hostname and Data fields [JSON Object]).

The Data column consists of the model name of the module to identify the hardware type from the fabric template. A .csv file screenshot:

	A	B	C	D	E	F	G
1	#SerialNumber(Eg:FD01344GH5)	#Model(Eg:N9k-C9236C)	#Version(Eg:7.0(3)12(3))	#IPAddress of the device	#HostName	#Data(JSON Field contains model name of the modules	
2	Serial Number	Model	Version	IP Address	Hostname	Data	
3	FDO21331SND	N9K-93180YC-EX	7.0(3)I5(2)	1.1.1.1	leaf1	{"modulesModel":["N9K-93180YC-EX"]}	
4	FDO21351N3X	N9K-C9236C	7.0(3)I4(1)	11.1.1.1	spine1	{"modulesModel":["N9K-C9236C"]}	
5	FDO21491A5K	N9K-C93240YC-FX2	7.0(3)I7(3)	12.1.1.1	leaf2	{"modulesModel":["N9K-C93240YC-FX2"]}	
6							

Step 6 Enter the administration password in the **Admin Password** and **Confirm Admin Password** fields.

Step 7 Select the device(s) and click **Bootstrap** at the top right part of the screen.

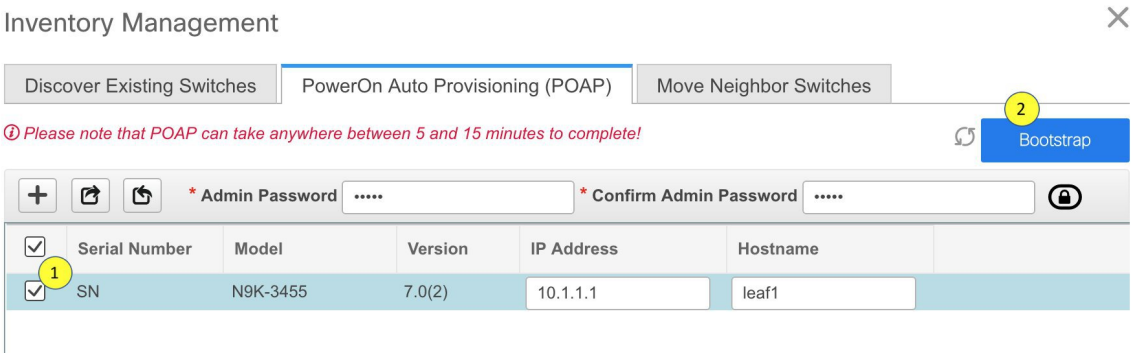

The leaf1 device appears in the fabric topology.

From the **Actions** panel, click **Tabular View**. You cannot deploy the fabric till the status of all the pre-provisioned switch(es) are displayed as **ok** under the **Discovery Status** column.

When you connect leaf1 to the fabric, the switch is provisioned with the IP address 10.1.1.1.

Step 8 Navigate to **Fabric Builder** and set roles for the device.

Create intra-link policy using one of the templates:

- **int_pre_provision_intra_fabric_link** to automatically generate intra fabric interface configuration with DCNM allocated IP addresses
- **int_intra_fabric_unnum_link_11_1** if you are using unnumbered links
- **int_intra_fabric_num_link_11_1** if you want to manually assign IP addresses to intra-links

Click **Save & Deploy**.

Configuration for the switches are captured in corresponding PTIs and can be seen in the **View/Edit Policies** window.

Step 9 To bring in the physical device, you can follow the manual RMA or POAP RMA procedure.

For more information, see Return Material Authorization (RMA), on page 135.

If you use the POAP RMA procedure, ignore the error message of failing to put the device into maintenance mode due to no connectivity since it is expected to have no connectivity to a non-existing device.

You need to click **Save & Deploy** in the fabric after the switch(es) are online to provision the host ports. This action must be performed before overlays are provisioned for the host port attachment.

Creating a vPC Setup in the External Fabric

You can create a vPC setup for a pair of switches in the external fabric. Ensure that the switches are of the same role and connected to each other.

Procedure

Step 1 Right-click one of the two designated **vPC switches** and choose **vPC Pairing**.

The **Select vPC peer** dialog box comes up. It contains a list of potential peer switches. Ensure that the **Recommended** column for the vPC peer switch is updated as **true**.

Step 2 Click the radio button next to the vPC peer switch and choose **vpc_pair** from the **vPC Pair Template** drop-down list. Only templates with the **VPC_PAIR** template sub type are listed here.

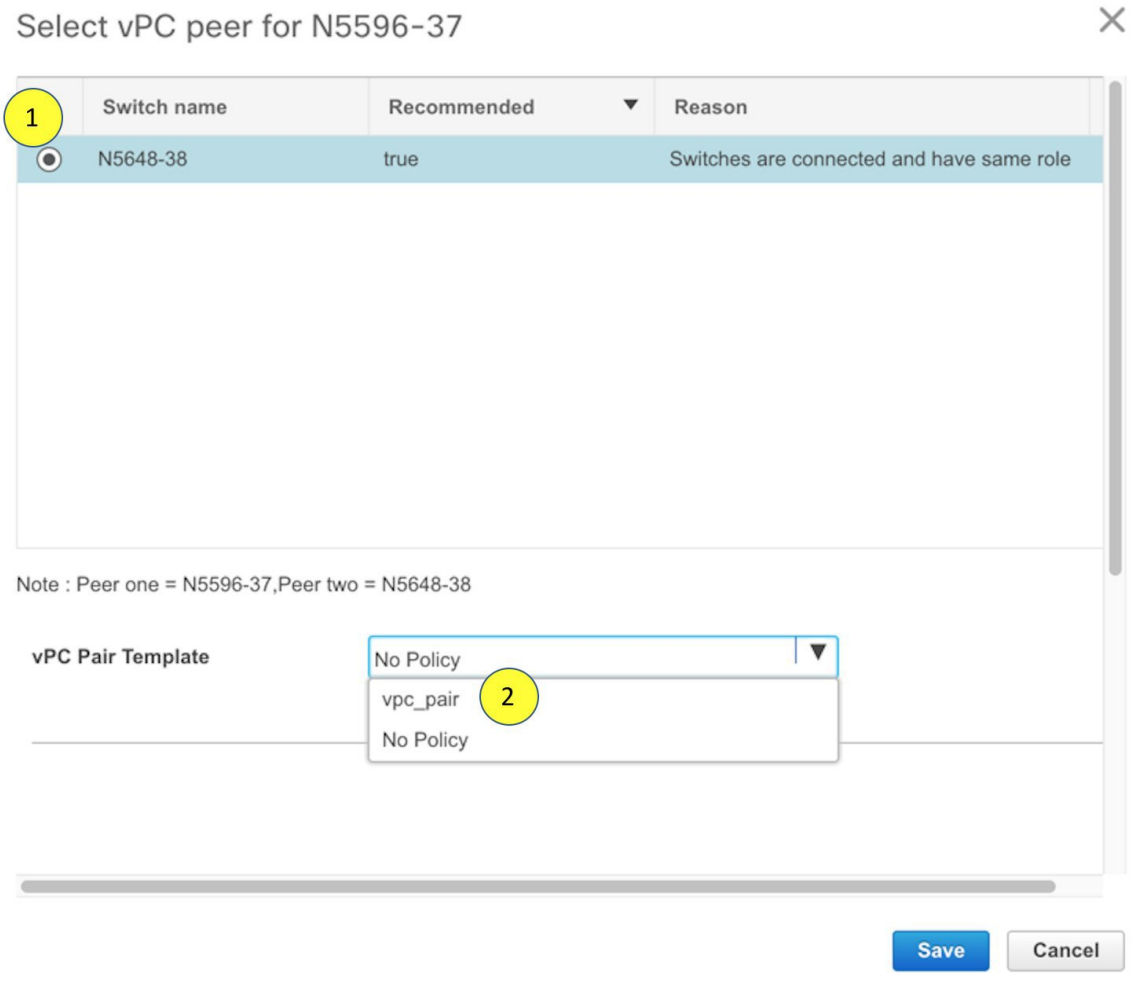

The **vPC Domain** and **vPC Peerlink** tabs appear. You must fill up the fields in the tabs to create the vPC setup. The description for each field is displayed at the extreme right.

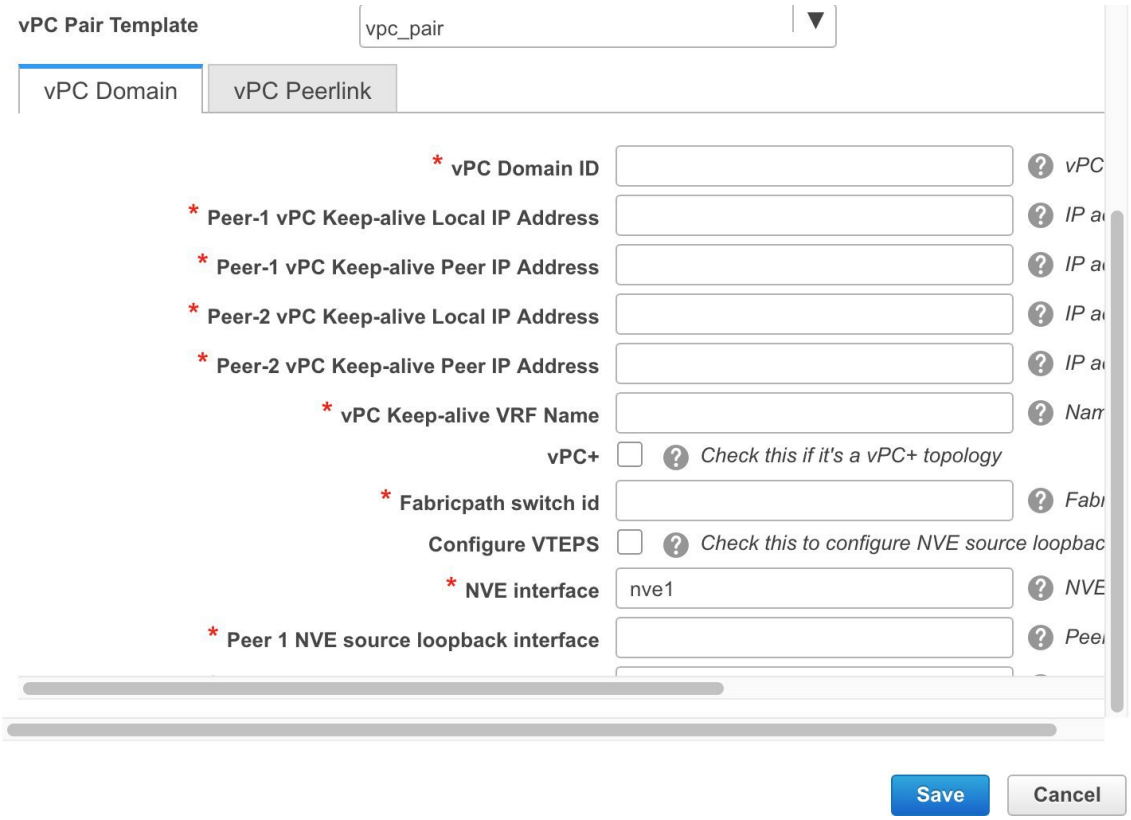

vPC Domain tab: Enter the vPC domain details.

vPC+: If the switch is part of a FabricPath vPC + setup, enable this check box and enter the **FabricPath switch ID** field.

Configure VTEPs: Check this check box to enter the source loopback IP addresses for the two vPC peer VTEPs and the loopback interface secondary IP address for NVE configuration.

NVE interface: Enter the NVE interface. vPC pairing will configure only the source loopback interface. Use the freeform interface manager for additional configuration.

NVE loopback configuration: Enter the IP address with the mask. vPC pairing will only configure primary and secondary IP address for loopback interface. Use the freeform interface manager for additional configuration.

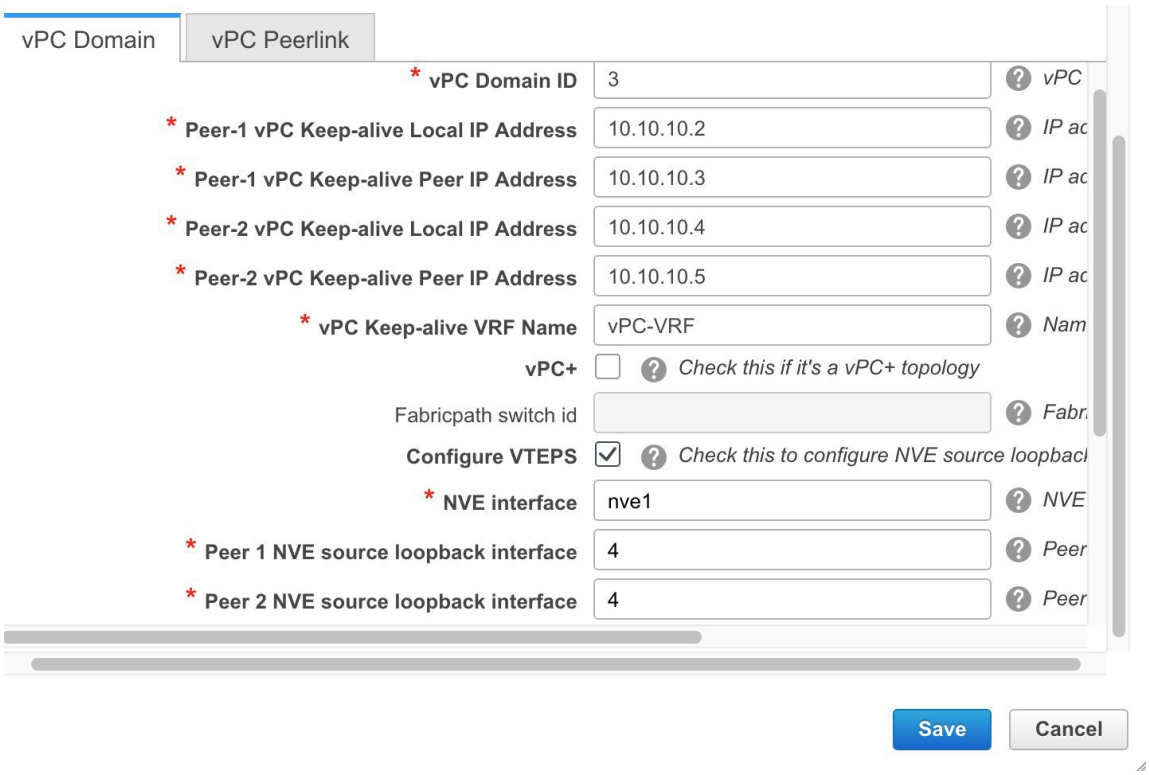

vPC Peerlink tab: Enter the vPC peer-link details.

Switch Port Mode: Choose **trunk** or **access** or **fabricpath**.

If you select trunk, then corresponding fields (**Trunk Allowed VLANs** and **Native VLAN**) are enabled. If you select **access**, then the **Access VLAN** field is enabled. If you select **fabricpath**, then the trunk and access port related fields are disabled.

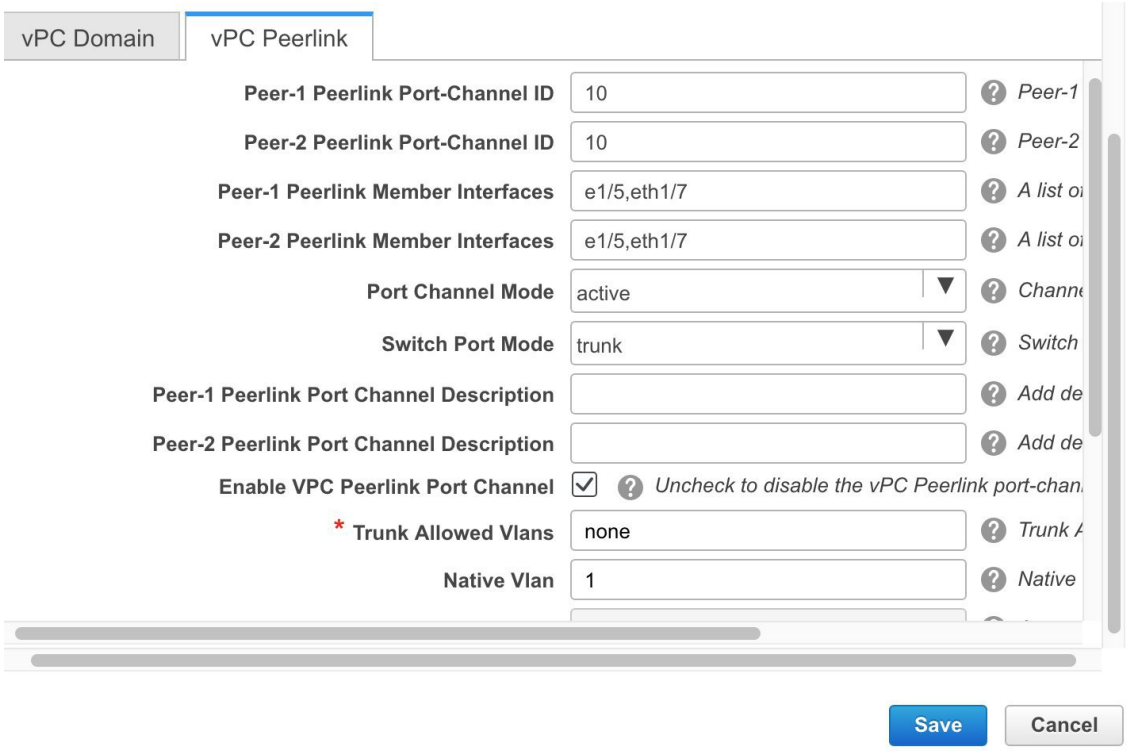

Step 3 Click **Save**.

The **fabric topology** window appears. The **vPC setup** is created.

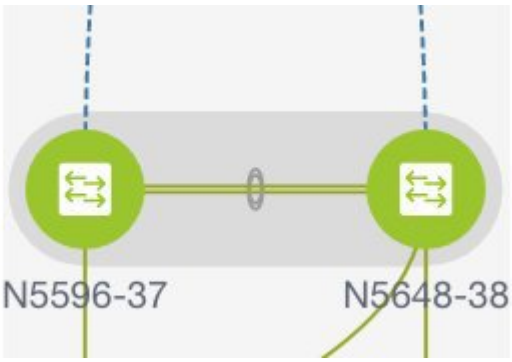

To update vPC setup details, do the following:

a. Right-click a vPC switch and choose vPC Pairing.

The **vPC peer** dialog box comes up.

b. Update the field(s) as needed.

When you update a field, the **Unpair** icon changes to **Save**.

c. Click **Save** to complete the update.

Undeploying a vPC Setup in the External Fabric

Procedure

Step 1 Right-click a **vPC** switch and choose **vPC Pairing**.

The vPC peer screen comes up.

Step 2 Click **Unpair** at the bottom right part of the screen.

The vPC pair is deleted and the fabric topology window appears.

Step 3 Click **Save & Deploy**.

The **Config Deployment** dialog box appears.

Step 4 (Optional) Click the value under the **Preview Config** column.

View the pending configuration in the **Config Preview** dialog box. The following configuration details are deleted on the switch when you unpair: vPC feature, vPC domain, vPC peerlink, vPC peerlink member ports, loopback secondary IPs, and host vPCs. However, the host vPCs and port channels are not removed. Delete these port channels from the **Interfaces** window if required.

Note Resync the fabric if it is out of sync.

When you unpair, only PTIs are deleted for following features, but the configuration is not cleared on the switch during **Save & Deploy**: NVE configuration, LACP feature, fabricpath feature, nv overlay feature, loopback primary ID. In case of host vPCs, port channels and their member ports are not cleared. You can delete these port channels from the **Interfaces** window if required. You can continue using these features on the switch even after unpairing.

If you are migrating from fabricpath to VXLAN, you need to clear the configuration on the device before deploying the VXLAN configuration.

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.

In DCNM 11.1(1) release, in addition to member fabrics, the topology view for the MSD fabric is introduced. This view displays all member fabrics, and how they are connected to each other, in one view.

Also, a deployment view is introduced for the MSD fabric. You can deploy overlay networks (and VRFs) on member fabrics from a single topology deployment screen, instead of visiting each member fabric deployment screen separately and deploying.

Note

- vPC support is added for BGWs in the DCNM 11.1(1) release.
- The MSD feature is unsupported on the switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images.
- 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 only be edited or updated in the fabric context from the VRFs and Networks window. The appropriate fabric should be selected in the **SCOPE** drop-down list to edit the fabric instance values. Some of the examples of fabric instance variables are BGP ASN, Multicast group per network or VRF, etc. For information about editing multicast group address, see Editing Networks in the Member Fabric, on page 100.

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:

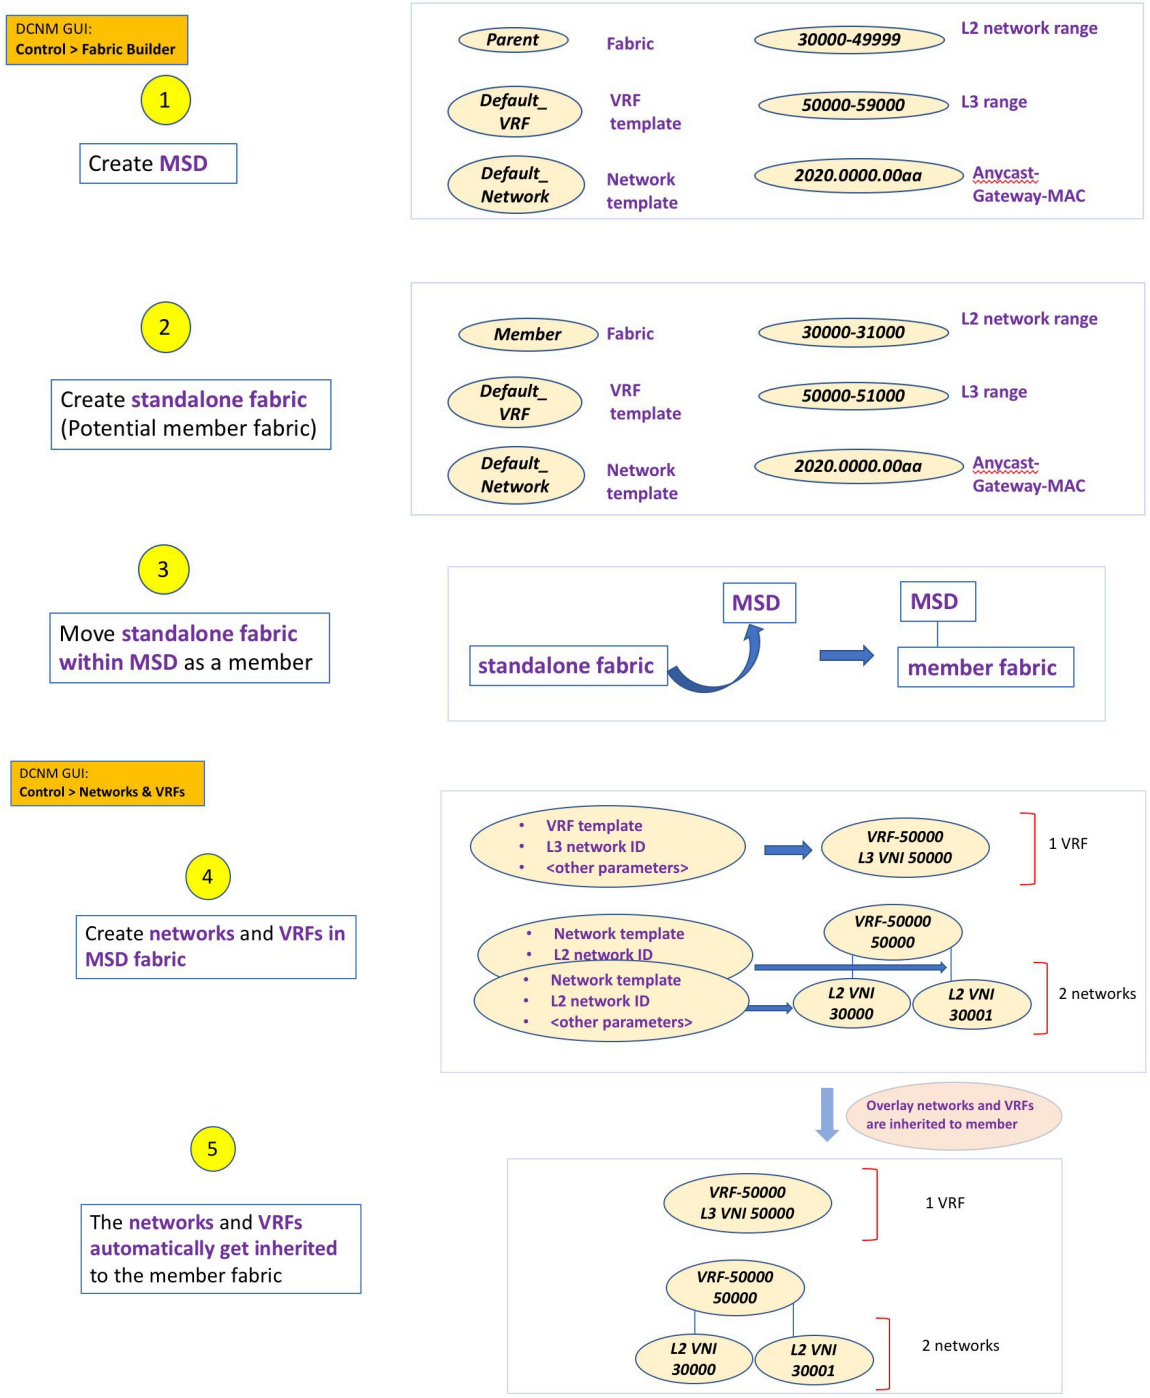

The sample flow explained the inheritance from the MSD to one member. An MSD has multiple sites (and hence, multiple member fabrics under an MSD). A sample flow from an MSD to multiple members:

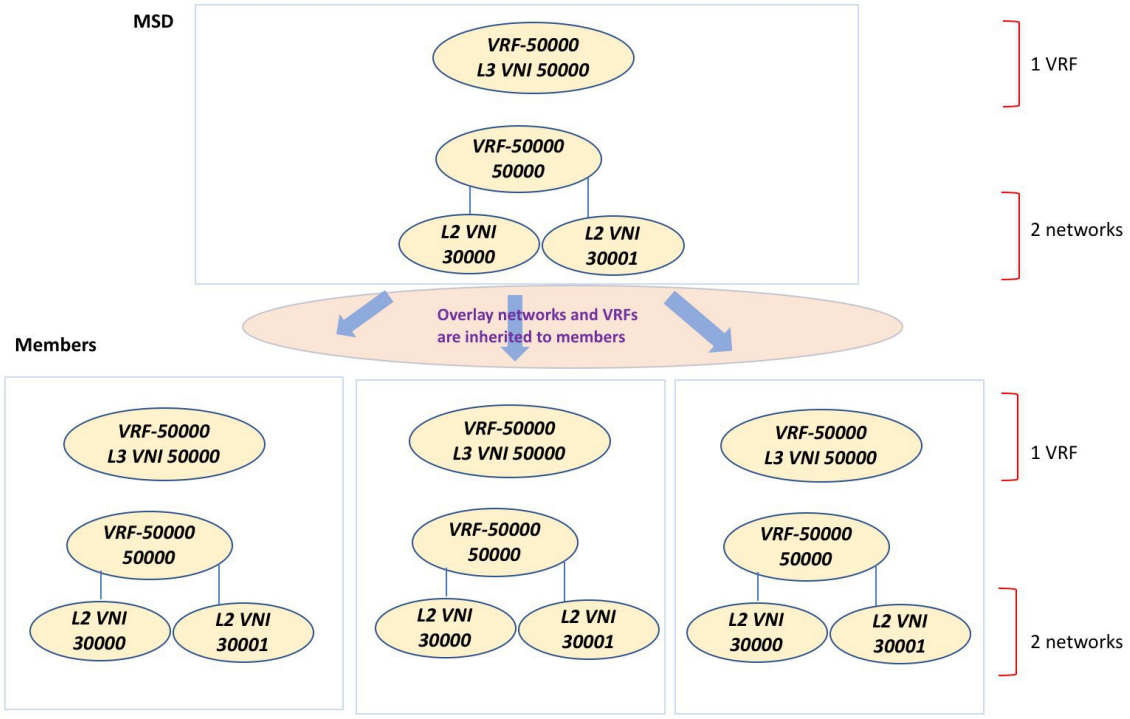

In this example, VRF-50000 (and L3 network with ID 50000), and L2 networks with IDs 30000 and 30001 are created in one go. Networks and VRFs are deployed on the member fabric switches, one after another, as depicted in the image.

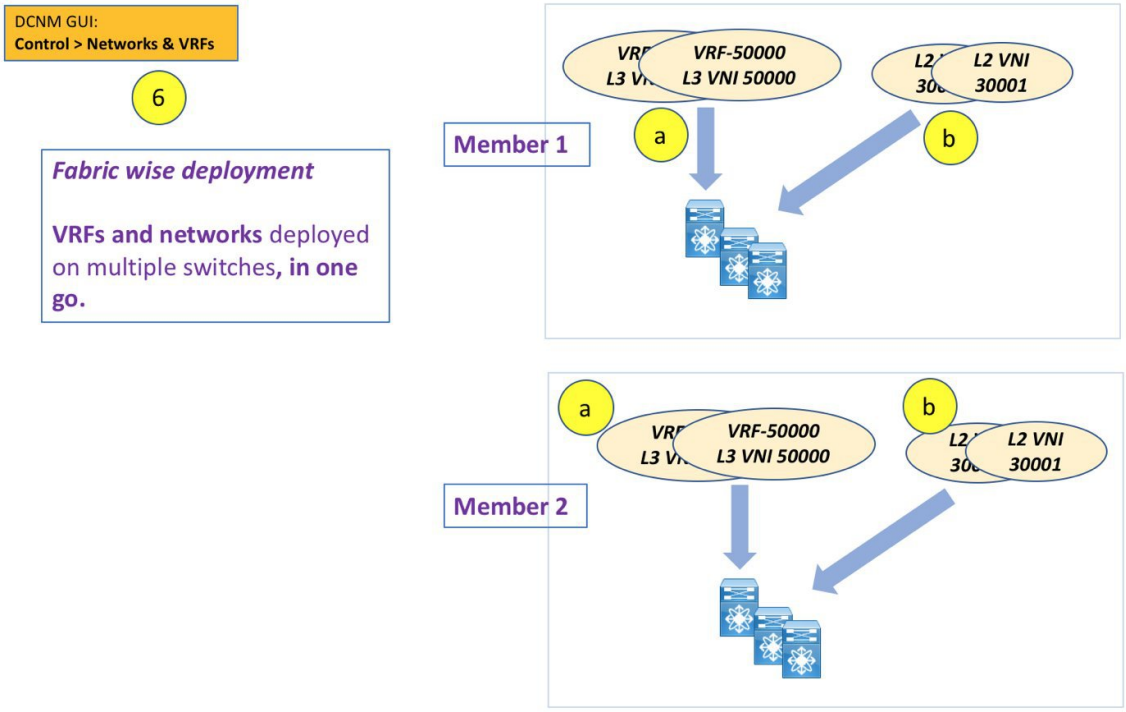

In DCNM 11.1(1), you can provision overlay networks through a single MSD deployment screen.

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 from the MSD and member fabric topology views.
- Other scenarios for fabric movement:
  - Standalone fabric with existing networks and VRFs to an MSD fabric.
  - Member fabric from one MSD to another.

Creating an MSD Fabric and Associating 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.

Creating an MSD Fabric

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

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

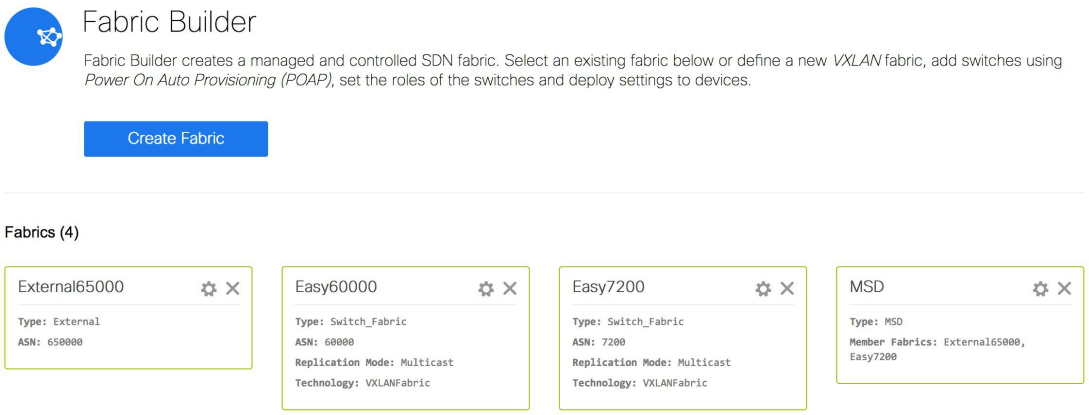

   A standalone or member fabric contains *Switch_Fabric* in the **Type** field, its AS number in the **ASN** field and mode of replication, *Multicast* or *Ingress Replication*, in the **Replication Mode** field. Since no device or network traffic is associated with an MSD fabric as it is a container, 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.

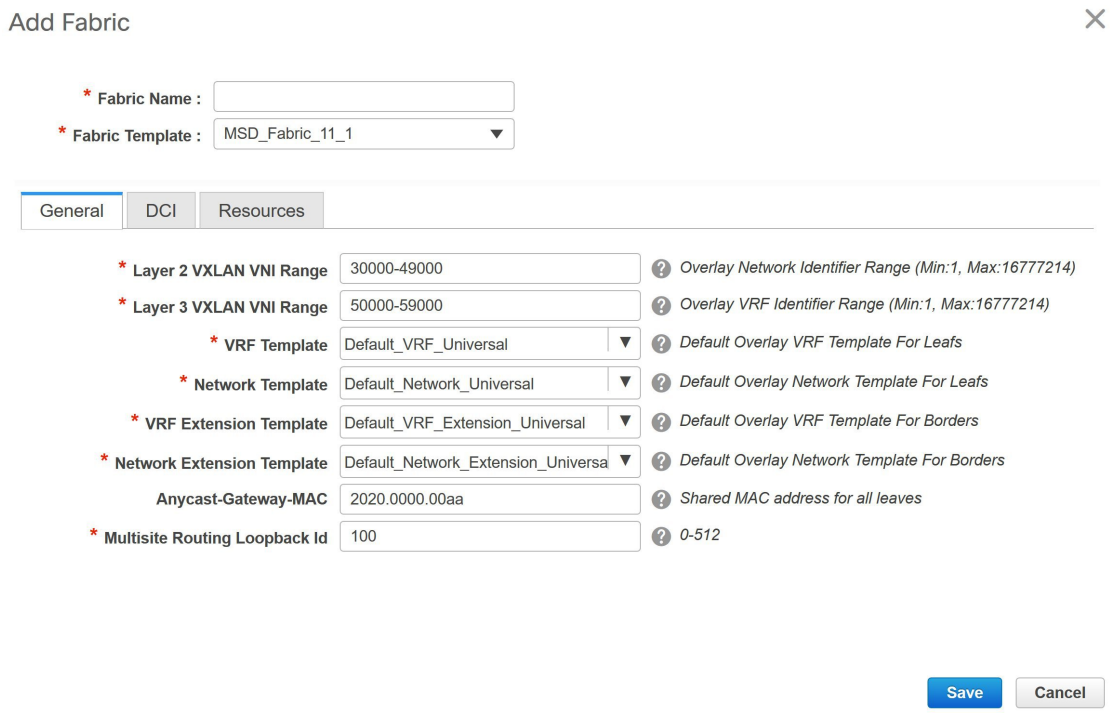

The fields in the screen are explained:

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.

Layer 2 VXLAN VNI Range - Layer 2 VXLAN segment identifier range.

Layer 3 VXLAN VNI Range - Layer 3 VXLAN segment identifier range.

VRF Template - Default VRF template.

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.

Multisite Routing Loopback Id – The multicast routing loopback ID is populated in this field.

3. Click the **DCI** tab.

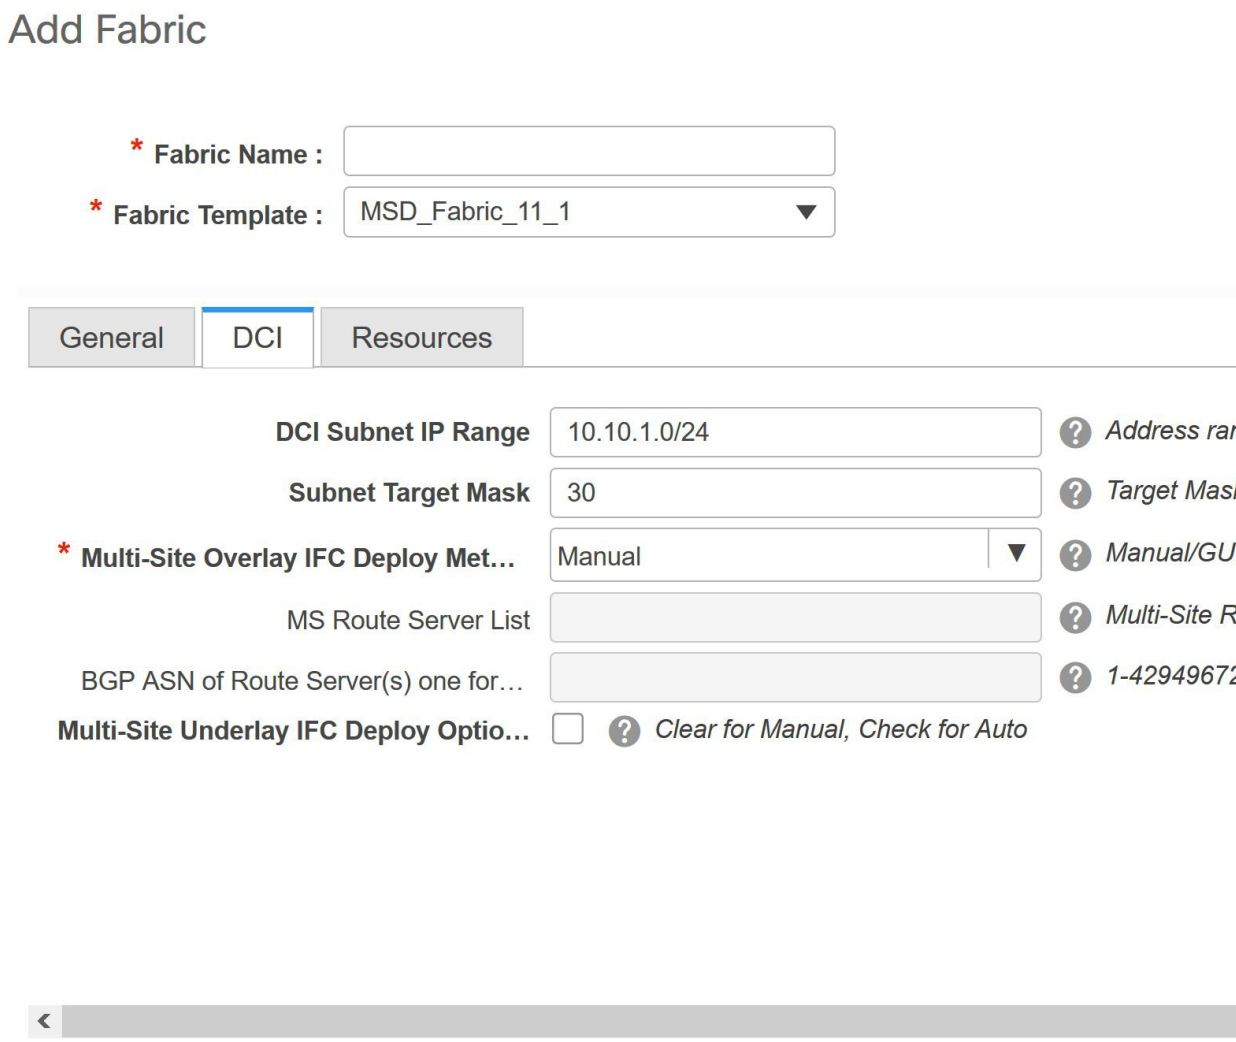

The fields are:

DCI Subnet IP Range and **Subnet Target Mask** – Specify the Data Center Interconnect (DCI) subnet IP address and mask.

Multi-Site Overlay IFC Deploy Method – Choose how you will connect the data centers through the BGW, manually, in a back-to-back fashion or through a route server.

If you choose to connect them through a route server, you should enter the route server details.

MS Route Server List – Specify the IP addresses of the route server. If you specify more than one, separate the IP addresses by a comma.

BGP ASN of Route Server(s) one for each route server – Specify the BGP AS Number of the router server. If you specify more than one route server, separate the AS Numbers by a comma.

Multi-Site Underlay IFC Deploy Options - Check the check box to enable auto configuration. Uncheck the check box for manual configuration.

4. Click the **Resources** tab.

MultiSite Routing Loopback IP Range – Specify the Multi-Site loopback IP address range used for the EVPN Multi-Site function.

A unique loopback IP address is assigned from this range to each member fabric because each member site must have a Loopback 100 IP address assigned for overlay network reachability. The per-fabric loopback IP address is assigned on all the BGWs in a specific member fabric.

5. 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.

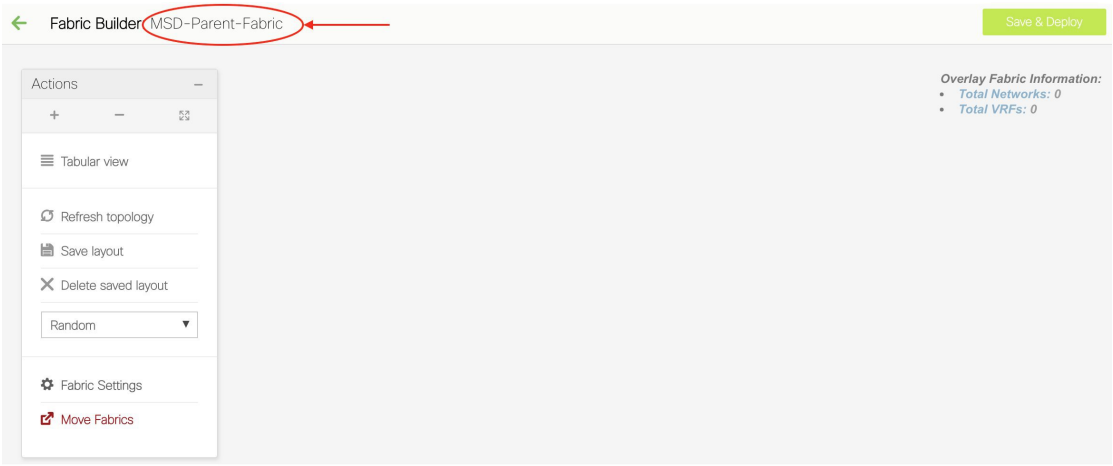

Since the MSD fabric is a container, you cannot add a switch to it. The **Add Switches** button that is available in the **Actions** panel for member and standalone fabrics is not available for the MSD fabric.

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.

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.

Creating and Moving 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. However, some pointers about a standalone (potential member) fabric:

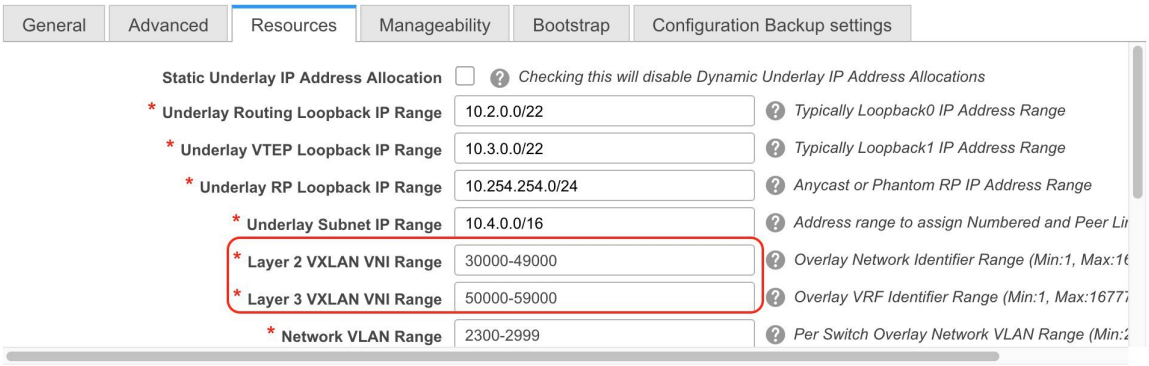

The values that are displayed in the screen are automatically generated. The VXLAN VNI ID ranges (in the L2 Segment ID Range and L3 Partition ID Range fields) allocated for new network and VRF creation are values from the MSD fabric segment ID range. If you want to update the VXLAN VNI ranges or the 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**.

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:

- 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.
- 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 fabric, *Member1*.

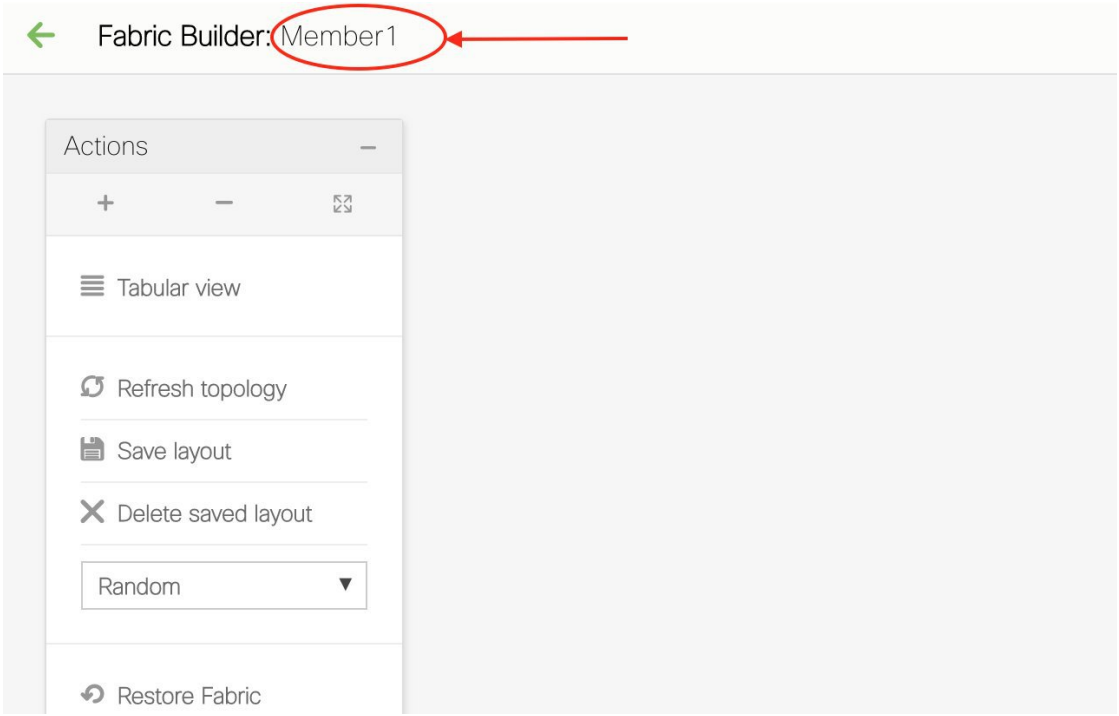

Simultaneously, the Fabric Builder page also displays the newly created fabric, Member1.

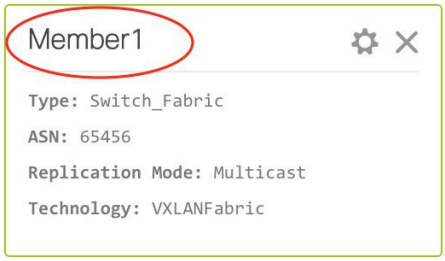

Moving the Member1 Fabric Under MSD-Parent-Fabric

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**.

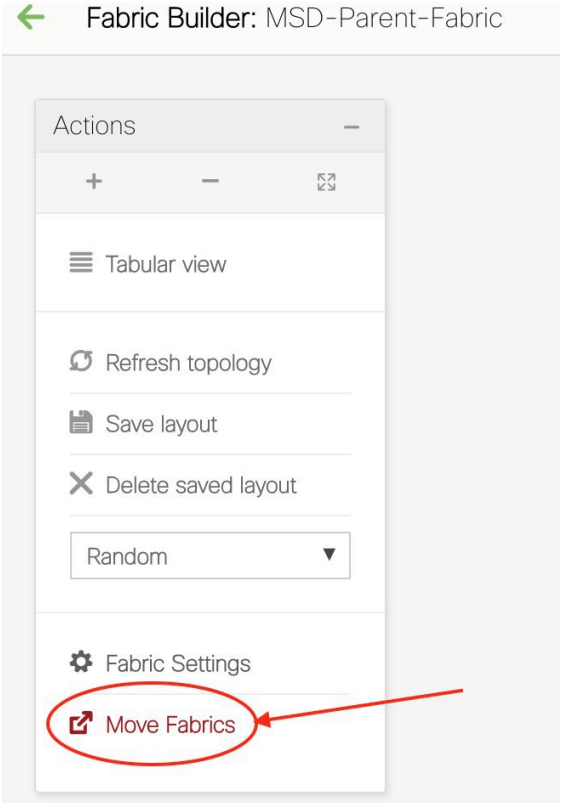

The Move Fabric screen comes up. It contains a list of fabrics.

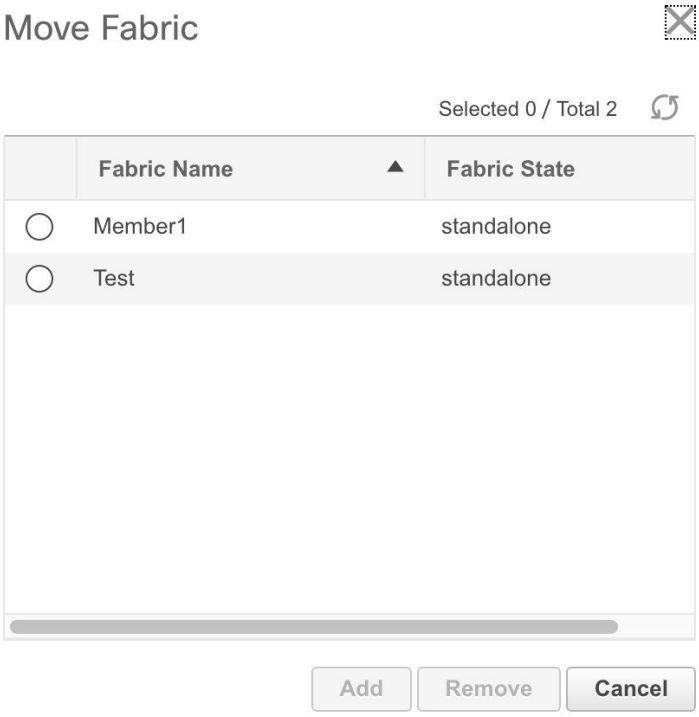

Member fabrics of other MSD container fabrics are not displayed here.

The *Member1* fabric is still a standalone fabric. 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.

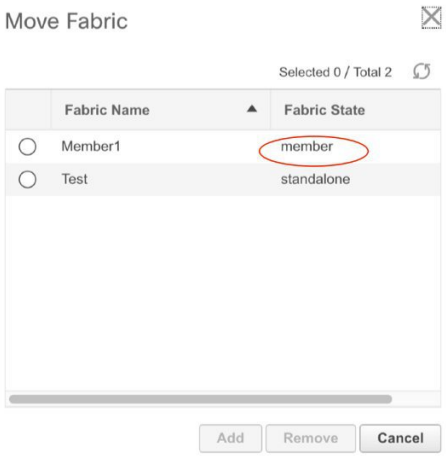

5. Close this screen.
6. Click ← above the Actions panel to go to the Fabric Builder page.

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

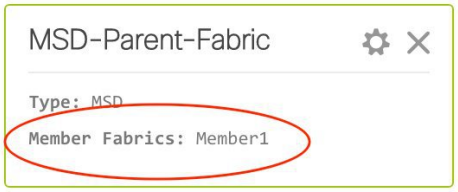

MSD Fabric Topology View Pointers

- **MSD fabric topology view** - Member fabrics and their switches are displayed. A boundary defines each member fabric. All fabric devices of the fabric are confined to the boundary.

  All links are displayed, including intra-fabric links and Multi-Site (underlay and overlay), and VRF Lite links to remote fabrics.

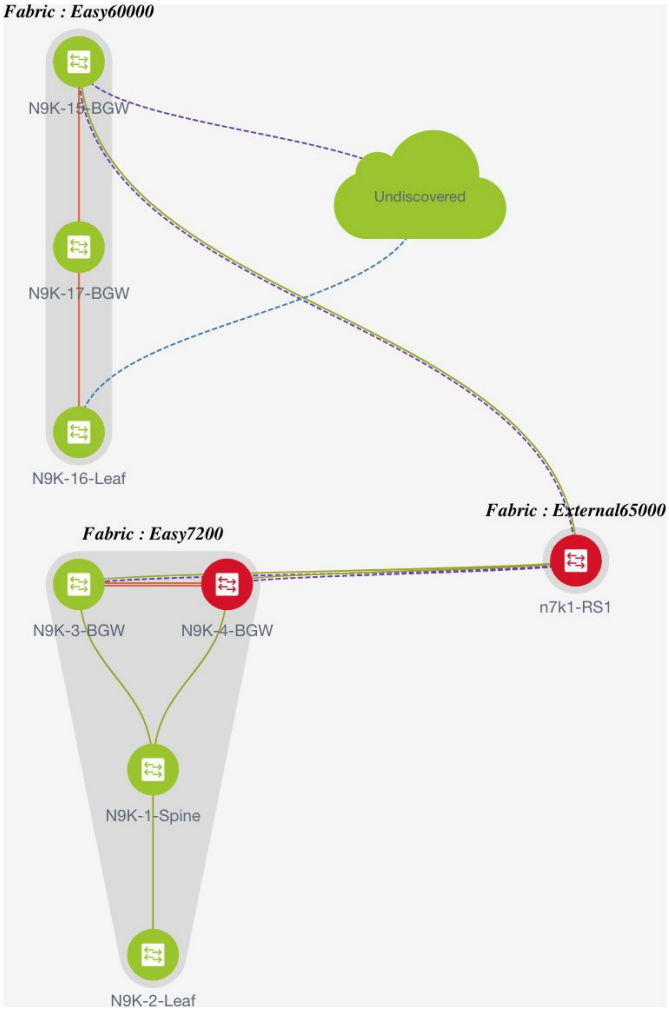

- **Member fabric topology view** - A member fabric and its switches are displayed. In addition, the connected external fabric is displayed.

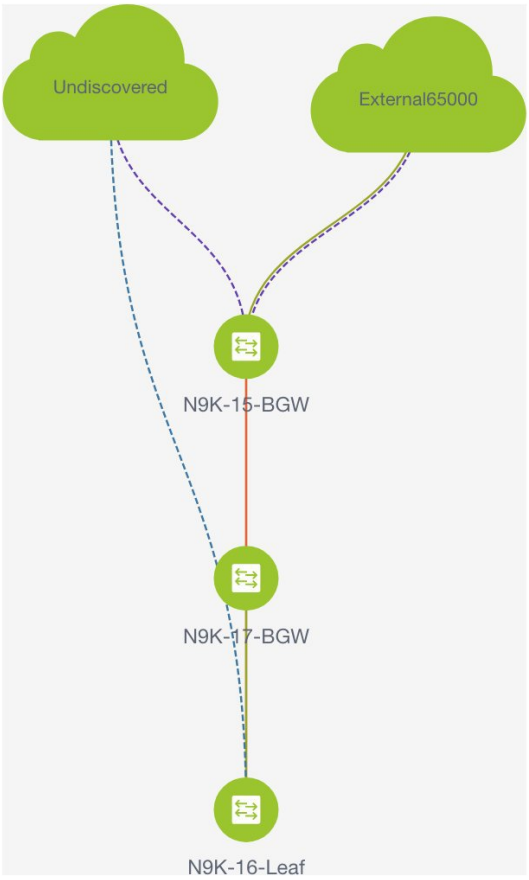

- A boundary defines a standalone VXLAN fabric, and each member fabric in an MSD fabric. A fabric's devices are confined to the fabric boundary. You can move a switch icon by dragging it. For a better user experience, in addition to switches, DNCM 11.2(1) release allows you to move an entire fabric. To move a fabric, place the cursor within the fabric boundary (but not on a switch icon), and drag it in the desired direction.

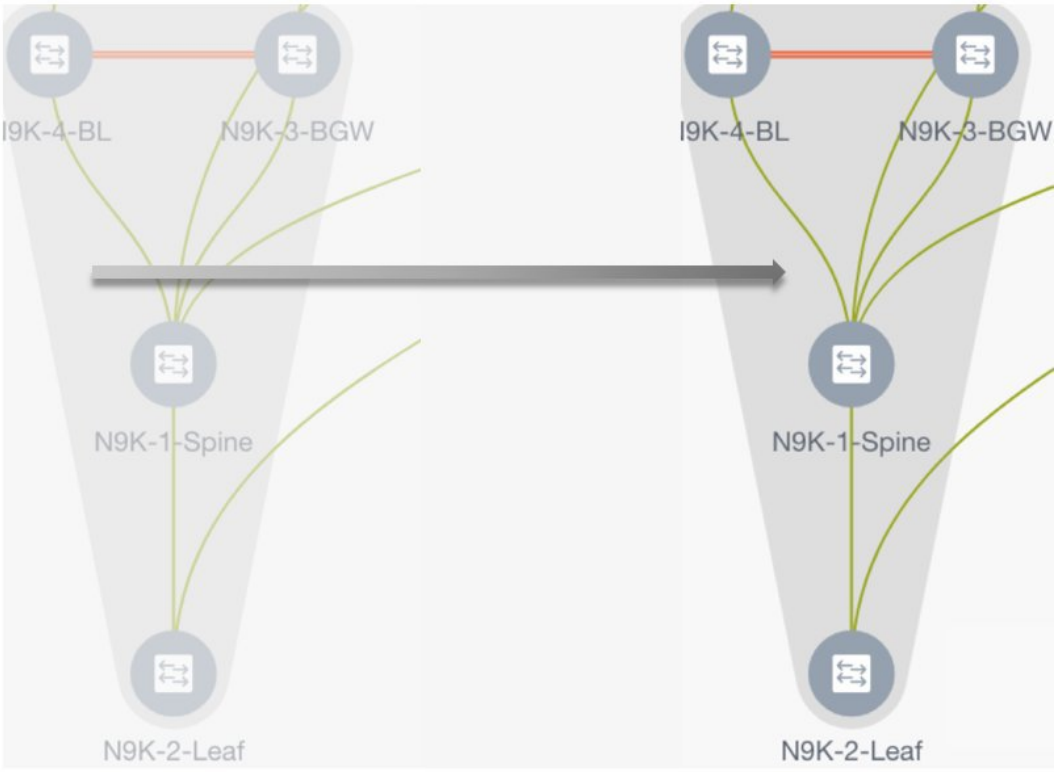

Adding and Editing Links

To add a link, right-click anywhere in the topology and use the **Add Link** option. To edit a link, right-click on the link and use the **Edit Link** option.

Alternatively, you can use the **Tabular view** option in the **Actions** panel.

To know how to add links between border switches of different fabrics (inter-fabric links) or between switches in the same fabric (intra-fabric links), refer the **Fabric Links** topic.

Creating and Deploying Networks and VRFs 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.

In DCNM 11.1(1) release, a deployment view is introduced for the MSD, in addition to the per-fabric deployment view. In this view, you can view and provision overlay networks for all member fabrics within the MSD, at once. However, you still have to apply and save network and VRF configurations 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.

In DCNM 11.1(1) release, you can deploy 30000 and 30001 on the border devices of all member fabrics through a single (MSD fabric) deployment screen. Prior to this, you had to access the first member fabric deployment screen, deploy 30000 and 300001 on the fabric's border devices, and then access the second member fabric deployment screen and deploy again.

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.

Creating Networks in the MSD Fabric

1. Click **Control** > **Networks** (under **Fabrics** submenu).

   The Networks screen comes up.

2. Choose the correct fabric from SCOPE. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

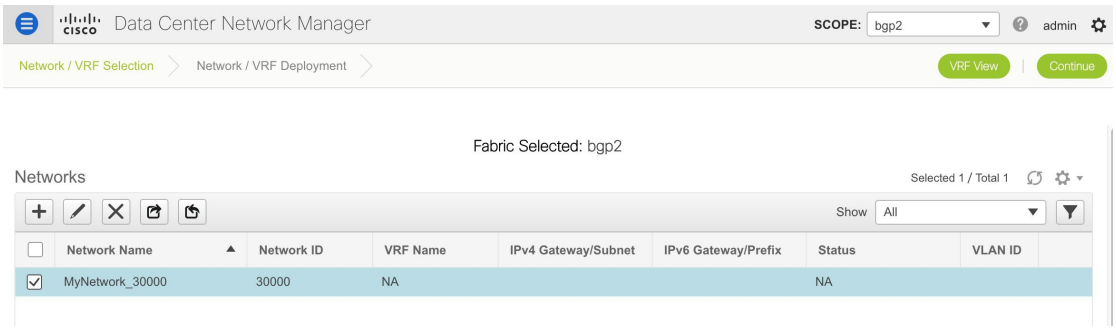

3. Select *MSD-Parent-Fabric* from the list and click **Continue** at the top right part of the screen.

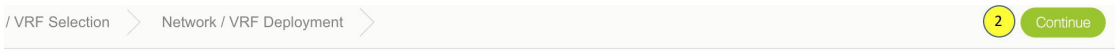

The Networks page comes up. This lists the list of networks created for the MSD fabric. Initially, this screen has no entries.

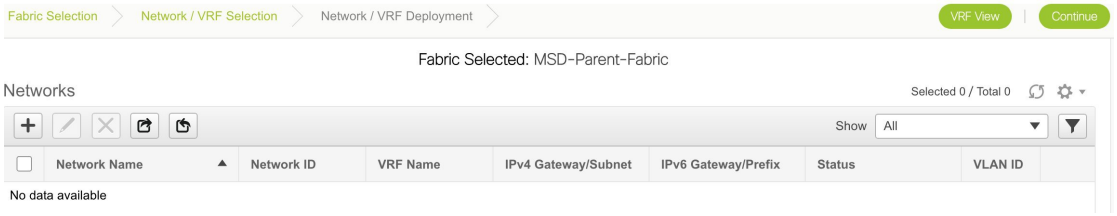

4. Click the + button at the top left part of the screen (under **Networks**) to add networks to the MSD fabric. The Create Network screen comes up. Most of the fields are autopopulated.

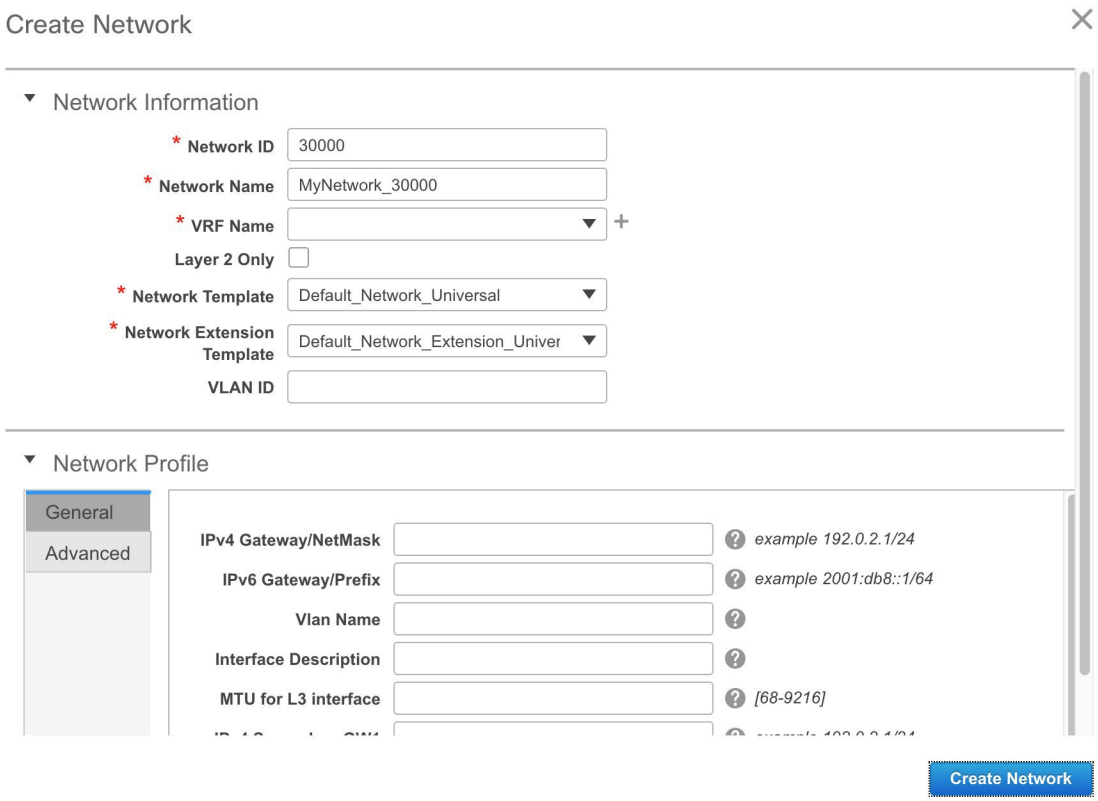

The fields in this screen are:

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 is 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.

VLAN Name - Enter the VLAN name.

If the VLAN is mapped to more than one subnet, enter the anycast gateway IP addresses for those subnets.

Interface Description - Specifies the description for the interface.

MTU for the L3 interface - Enter the MTU for Layer 3 interfaces.

IPv4 Secondary GW1 - Enter the gateway IP address for the additional subnet.

IPv4 Secondary GW2 - Enter the gateway IP address for the additional subnet.

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

- ARP Suppression
- DHCPv4 Server 1 and DHCPv4 Server 2 - Enter the DHCP relay IP address of the first and second DHCP servers.
- DHCPv4 Server VRF - Enter the DHCP server VRF ID.
- Loopback ID for DHCP Relay interface - Enter the loopback ID of the DHCP relay interface.
- Routing Tag – The routing tag is autopopulated. This tag is associated with each gateway IP address prefix.
- TRM enable – Select the checkbox to enable TRM.
- L2 VNI Route-Target Both Enable - Select the check box to enable automatic importing and exporting of route targets for all L2 virtual networks.
- Enable L3 Gateway on Border - Select the checkbox to enable the Layer 3 gateway on the border device.

  A sample of the Create Network screen:

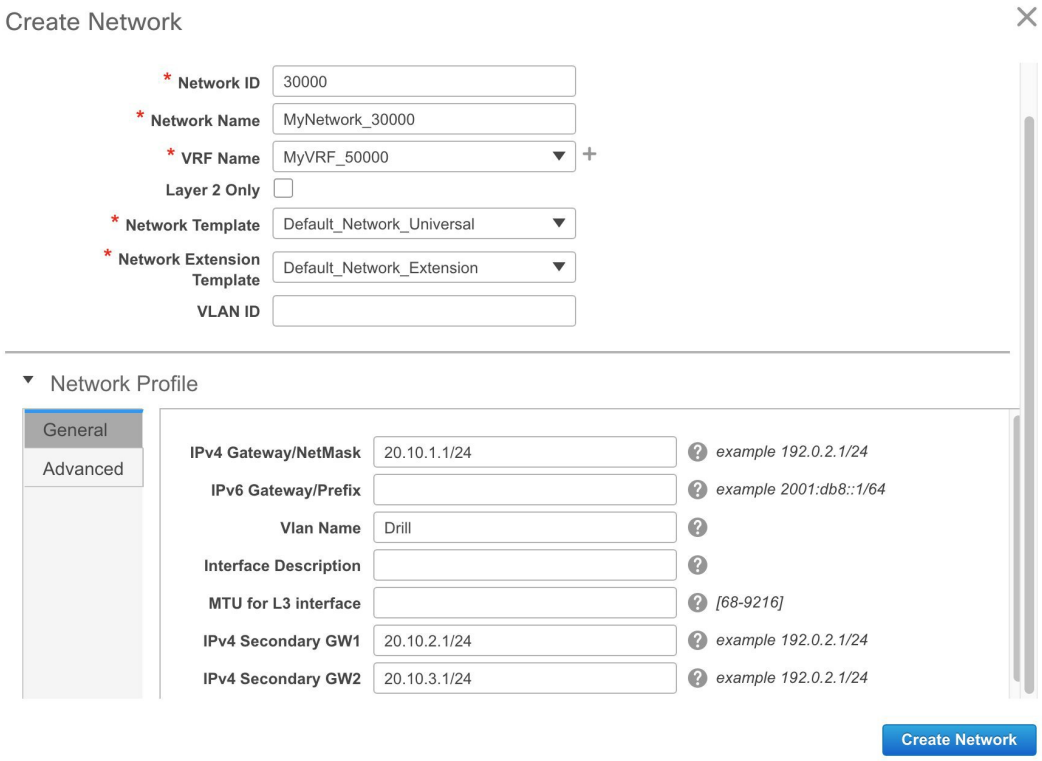

Advanced tab:

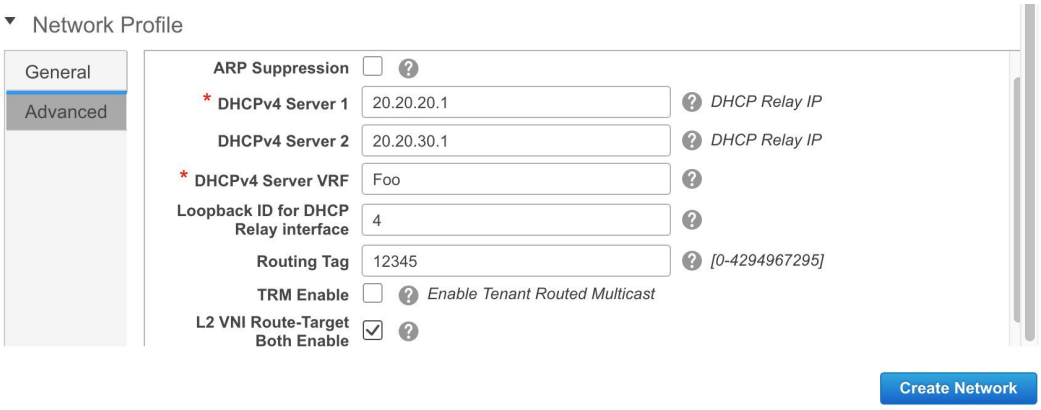

5. Click **Create Network**. A message appears at the bottom right part of the screen indicating that the network is created. The new network (*MyNetwork_30000*) appears on the Networks page that comes up.

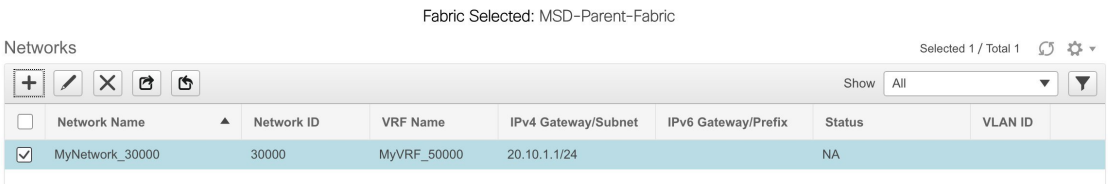

Editing Networks in the MSD Fabric

1. In the Networks screen of the MSD fabric, select the network you want to edit and click the Edit icon at the top left part of the screen.

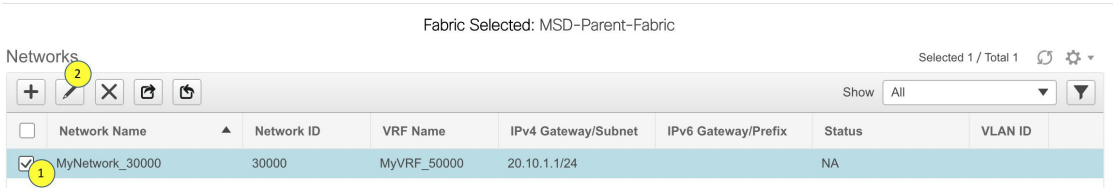

The Edit Network screen comes up.

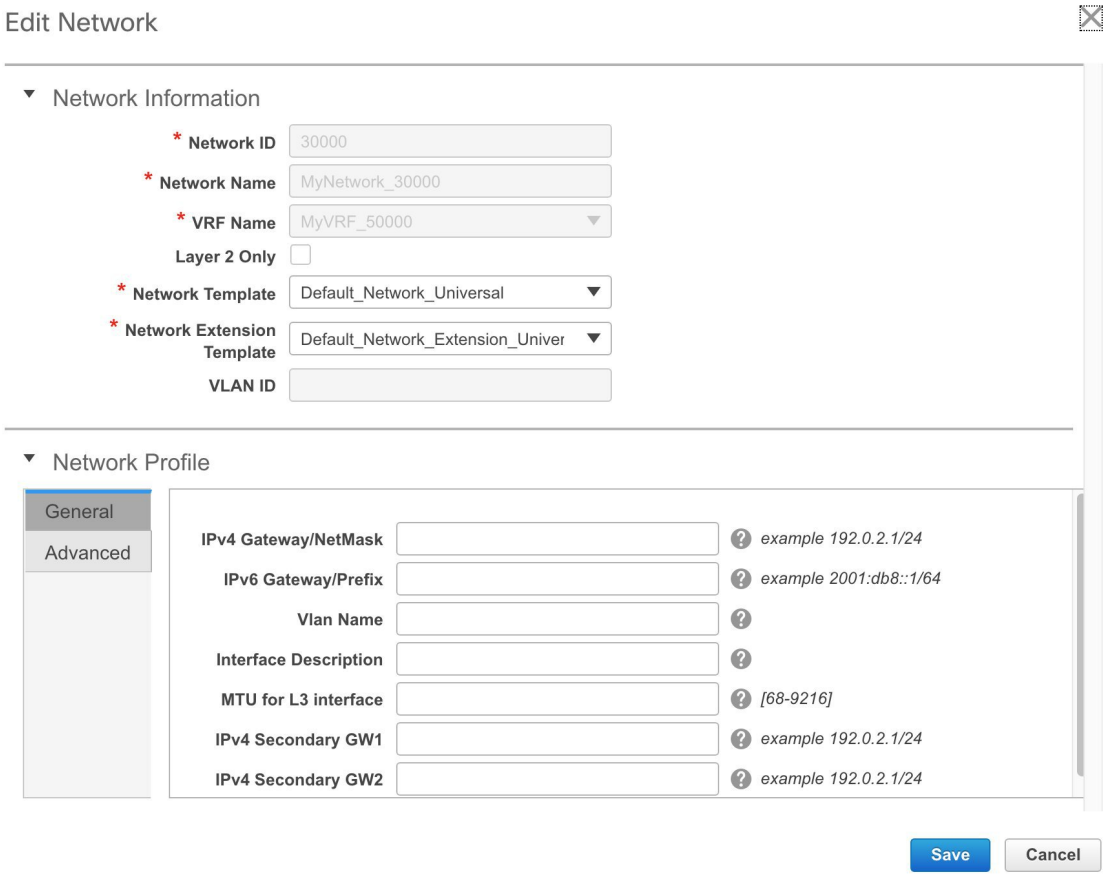

You can edit the **Network Profile** part (**General** and **Advanced** tabs) of the MSD fabric network.

2. Click **Save** at the bottom right part of the screen to save the updates.

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.

1. Click **Control > Networks** (under **Fabrics** submenu).

   The Networks screen comes up.

2. Choose the correct fabric from SCOPE. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

Editing Networks in the Member Fabric

An MSD can contain multiple fabrics. These fabrics forward BUM traffic via Multicast or Ingress replication. Even if all the fabrics use multicast for BUM traffic, the multicast groups within these fabrics need not be the same.

When you create a network in MSD, it is inherited by all the member fabrics. However, the multicast group address is a fabric instance variable. To edit the multicast group address, you need to navigate to the member fabric and edit the network. For more information about the **Multicast Group Address** field, see *Creating Networks for the Standalone Fabric*.

1. Select the network and click the **Edit** option at the top left part of the window. The **Edit Network** window comes up.
2. Update the multicast group address in one of the following ways:
   - Under **Network Profile**, click the **Generate Multicast IP** button to generate a new multicast group address for the selected network, and click **Save**.
   - Click the **Advanced** tab in the **Network Profile** section, update the multicast group address, and click **Save**.

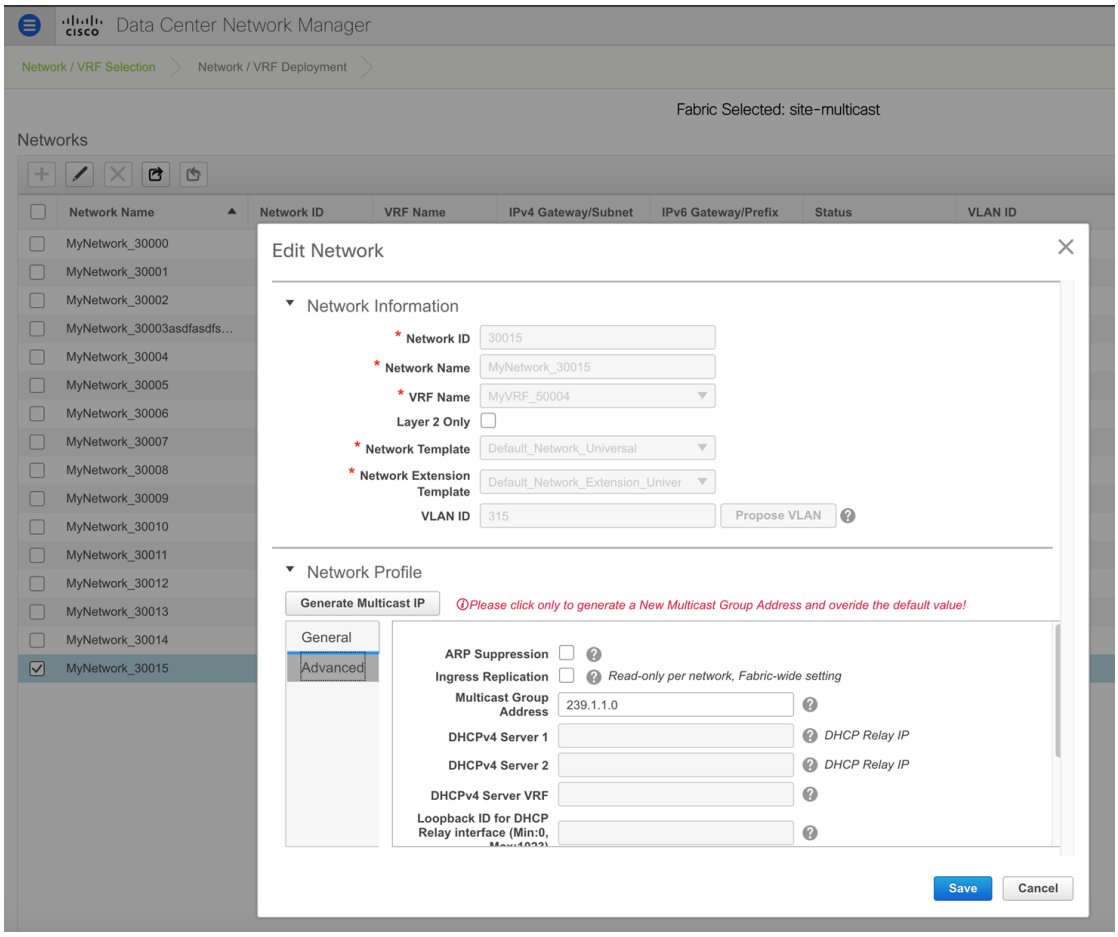

Note The **Generate Multicast IP** option is only available for member fabric networks and not MSD networks.

Deleting Networks in the MSD and Member Fabrics

You can only delete networks from the MSD fabric, and not member fabrics. To delete networks and corresponding VRFs in the MSD fabric, follow this order:

1. Undeploy the networks on the respective fabric devices before deletion.
2. Delete the networks from the MSD fabric. To delete networks, use the delete (**X**) option at the top left part of the Networks screen. You can delete multiple networks at once.

Note When you delete networks from the MSD fabric, the networks are automatically removed from the member fabrics too.

3. Undeploy the VRFs on the respective fabric devices before deletion.

4. Delete the VRFs from the MSD fabric by using the delete (**X**) option at the top left part of the screen. You can delete multiple VRF instances at once.

Creating 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.

   a. Choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

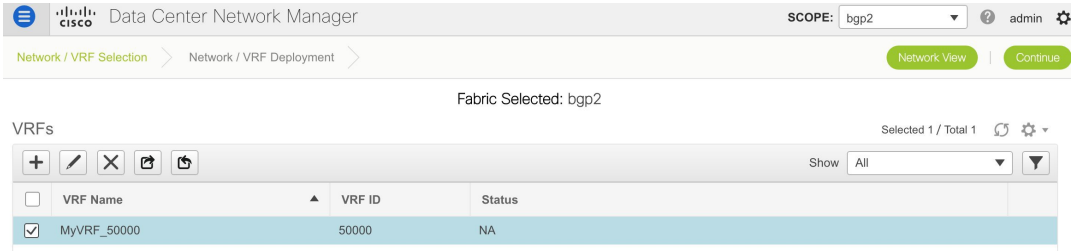

   b. Choose the MSD fabric (*MSD-Parent-Fabric*) from the drop-down box and click **Continue**. The Networks page comes up.

   c. 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.

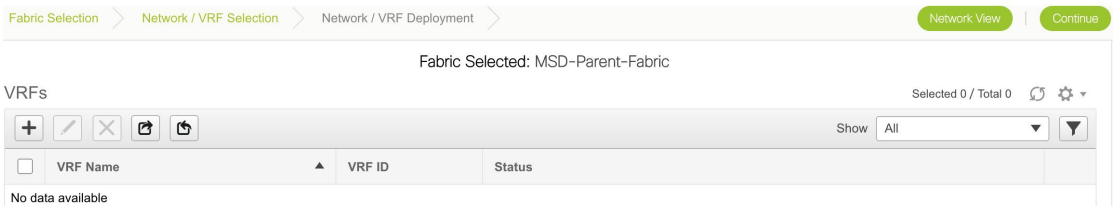

2. Click the + button at the top left part of the screen to add VRFs to the MSD fabric. The Create VRF screen comes up. Most of the fields are autopopulated.

   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. **General** tab – Enter the VLAN ID of the VLAN associated with the VRF, the corresponding Layer 3 virtual interface, and the VRF ID.

4. **Advanced** tab

Routing Tag – If a VLAN is associated with multiple subnets, then this tag is associated with the IP prefix of each subnet. Note that this routing tag is associated with overlay network creation too.

Redistribute Direct Route Map – Specifies the route map name for redistribution of routes in the VRF.

Max BGP Paths and **Max iBGP Paths** – Specifies the maximum BGP and iBGP paths.

TRM Enable – Select the checkbox to enable TRM.

If you enable TRM, then the RP address, the RP loopback ID and the underlay multicast address must be entered.

Is RP external - Select the checkbox if a fabric-external device is designated as RP.

RP Address and **RP Loopback ID** – Specifies the loopback ID and IP address of the RP.

Underlay Multicast Address – Specifies the multicast address associated with the VRF. The multicast address is used for transporting multicast traffic in the fabric underlay.

Overlay Multicast Groups – Specifies the multicast address for the VRF, used in the fabric overlay.

Enable IPv6 link-local Option - Select the check box to enable the IPv6 link-local option under the VRF SVI. If this check box is unchecked, IPv6 forward is enabled.

Advertise Host Routes - Select the checkbox to control advertisement of /32 and /128 routes to Edge Routers.

Advertise Default Route - Select the checkbox to control advertisement of default routes within the fabric.

A sample screenshot:

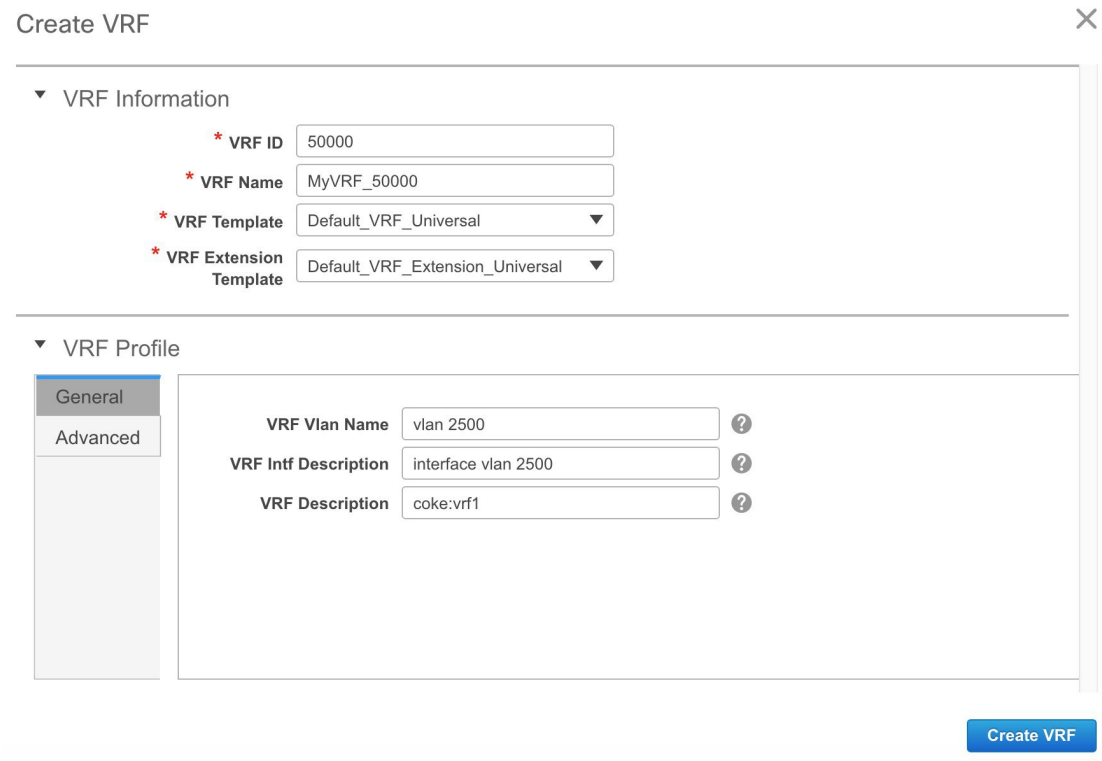

Advanced tab:

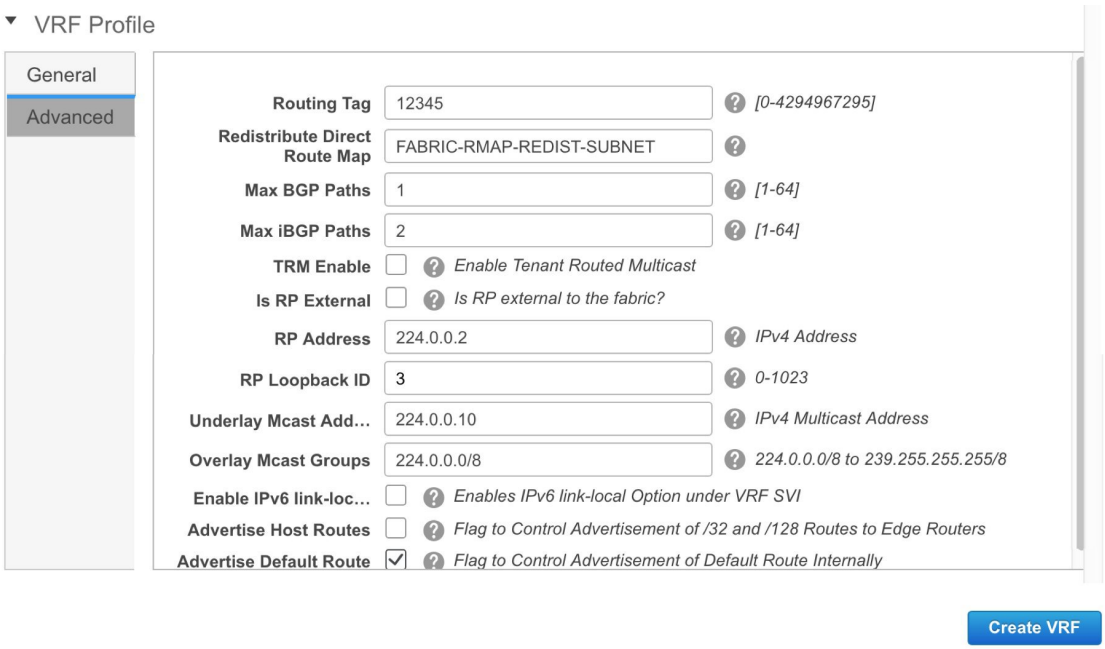

5. Click **Create VRF**.

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

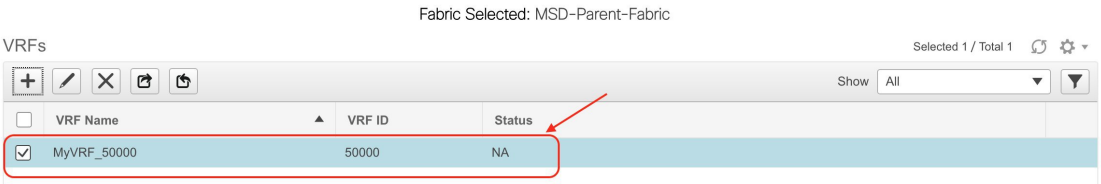

Editing VRFs in the MSD Fabric

1. In the VRFs screen of the MSD fabric, select the VRF you want to edit and click the Edit icon at the top left part of the screen.

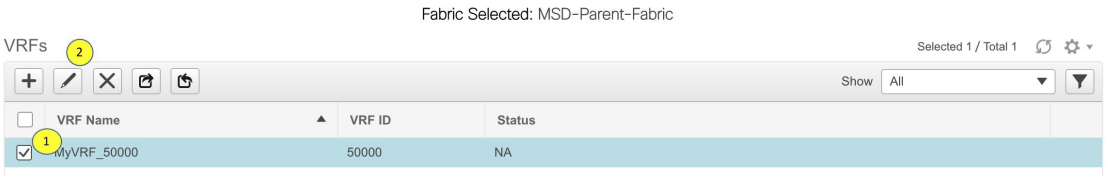

   The Edit VRF screen comes up.

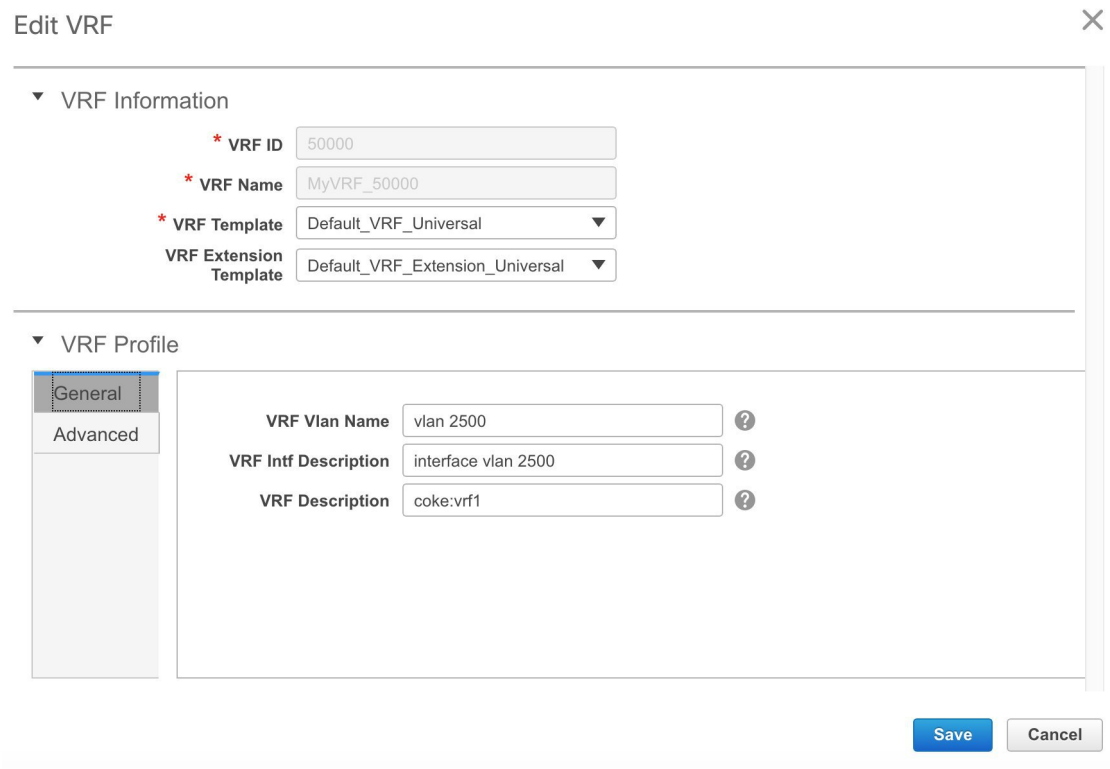

You can edit the **VRF Profile** part (**General** and **Advanced** tabs).

2. Click **Save** at the bottom right part of the screen to save the updates.

VRF Inheritance from MSD-Parent-Fabric to Member1

MSD-Parent-Fabric contains one member fabric, *Member1*. Do the following to access the member fabric page.

1. Choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

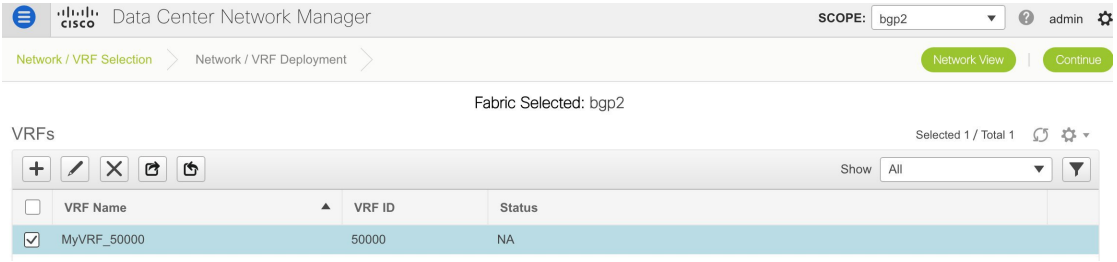

2. Click the **VRF View** button. On the VRFs page, you can see that the VRF created for the MSD is inherited to its member.

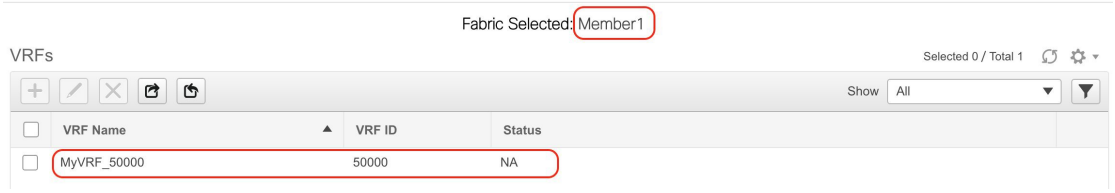

Deleting VRFs in the MSD and Member Fabrics

You can only delete networks from the MSD fabric, and not member fabrics. To delete networks and corresponding VRFs in the MSD fabric, follow this order:

1. Undeploy the networks on the respective fabric devices before deletion.
2. Delete the networks from the MSD fabric.
3. Undeploy the VRFs on the respective fabric devices before deletion.
4. Delete the VRFs from the MSD fabric by using the delete (**X**) option at the top left part of the screen. You can delete multiple VRF instances at once.

Note When you delete VRFs from the MSD fabric, they are automatically removed from the member fabrics too.

Editing VRFs in the Member Fabric

You cannot edit VRF parameters at the member fabric level. Update VRF settings in the MSD fabric. All member fabrics are automatically updated.

Deleting VRFs in the Member Fabric

You cannot delete VRFs at the member fabric level. Delete VRFs in the MSD fabric. The deleted VRFs are automatically removed from all member fabrics.

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 Creating and Deploying Networks and VRFs .

Removing a Fabric From an MSD

To remove a fabric from an MSD fabric, perform the following steps:

Before you begin

Make sure that there are no VRFs deployed on the border switches in the fabric that you want to remove. For more information, see Deployment and Undeployment of Networks and VRFs in Member Fabrics, on page 106.

Note Before removing a fabric from MSD, you need to manually remove overlay and underlay IFCs even with the auto deployment field enabled.

Procedure

Step 1 From the **Fabric Builder** window, click an MSD fabric.

Step 2 Click **Move Fabric** in the **Actions** menu.

Step 3 In the **Move Fabric** window, select the respective radio button of the fabric that you want to remove and click **Remove**.

In the fabric removal notification window, click **Close**.

Step 4 Click **Save & Deploy** for the MSD in the **Fabric Builder** window.

Step 5 Click **Deploy Config** in the **Config Deployment** window.

Click **Close**.

Step 6 Navigate to the fabric that you removed from MSD and click **Save & Deploy**.

Step 7 Click **Deploy Config** in the **Config Deployment** window.

Click **Close**.

Moving 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.

SSH Key RSA Handling

Bootstrap scenario

If the switch has the **ssh key rsa** command with the key-length variable value other than 1024 in the running configuration, the **ssh key rsa** *key-length* **force** command needs to be added to the bootstrap freeform configuration with the required value (any value other than 1024) during bootstrap.

Greenfield and Brownfield scenarios

Use the **ssh key rsa** *key-length* **force** command to change the key-length variable to a value other than 1024.

However, on Cisco Nexus 9000 Releases 9.3(1) and 9.3(2), the **ssh key rsa** *key-length* **force** command fails while the device is booting up during the ASCII replay process. For more information, refer CSCvs40704.

The configurations are considered to be in-sync when both the intent and switch running configurations have the same command. For example, the status is considered to be in-sync when the **ssh key rsa 2048** command is present in both in the intent and the running configuration. However, consider a scenario in which the **ssh key rsa 2040** command was pushed to the switch as an Out-Of-Band change. While the intent has a key-length value of 2048, the device has a key-length value of 2040. In such instances, the switch will be marked as out-of-sync.

The diff shown in the Pending Config tab (in both Strict Config-Compliance and non-Strict Config-Compliance mode) cannot be deployed onto the switch from DCNM as the **feature ssh** command has to be used to disable the SSH feature before making any change to the **ssh key rsa** command. This would lead to a dropped connection to DCNM. In such a scenario, the diff can be resolved by modifying the intent such that there is no diff.

With Strict Config-Compliance mode:

- Delete the Policy Template Instance (PTI) that has the **ssh key rsa 2048 force** command by clicking **View/Edit Policies** in the **Tabular View** of the **Fabric Builder** window.

- Create a new PTI with the **ssh key rsa 2040 force** command by clicking **View/Edit Policies**.

Without Strict Config-Compliance mode:

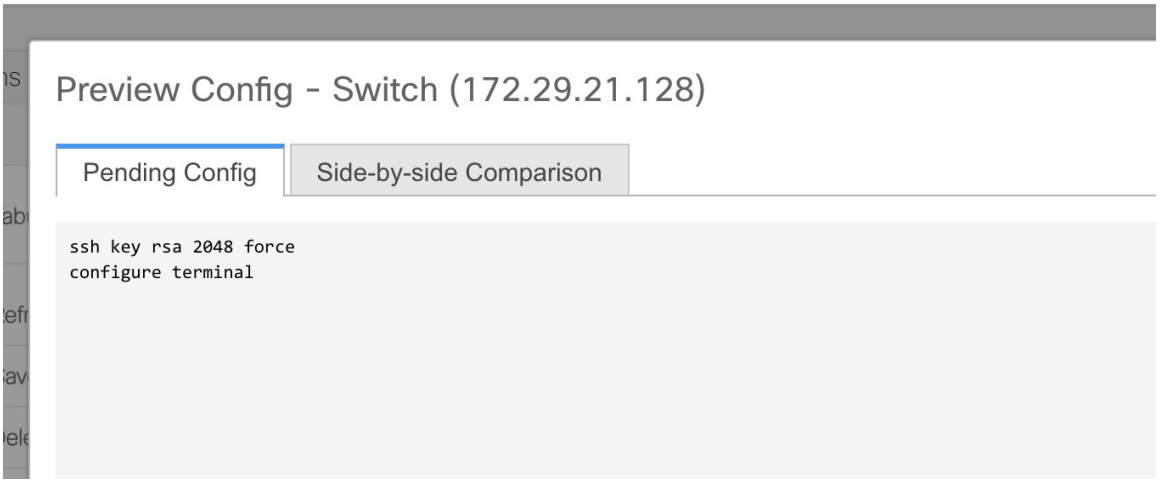

- Delete the PTI with the **ssh key rsa 2048 force** command in the intent by clicking **View/Edit Policies** in the **Tabular View** of the **Fabric Builder** window.

- Create a switch_freeform PTI with the **ssh key rsa 2040 force** command in the intent to match the Out-Of-Band change from the device.

Switch Operations

To view various options, right-click on switch:

- **Set Role** - Assign a role to the switch. You can assign any one of the following roles to a switch:

- Spine
- Leaf (Default role)
- Border
- Border Spine
- Border Gateway
- Access
- Aggregation
- Edge Router
- Core Router
- Super Spine
- Border Super Spine
- Border Gateway Spine
- ToR

Note You can change the switch role only before executing **Save & Deploy**.

From DCNM 11.1(1) release, you can shift the switch role from existing to required role if there are no overlays on the switches. Click **Save and Deploy** to generate the updated configuration. The following shifts are allowed for the switch role:

- Leaf to Border
- Border to Leaf
- Leaf to Border Gateway
- Border Gateway to Leaf
- Border to Border Gateway
- Border Gateway to Border
- Spine to Border Spine
- Border Spine to Spine
- Spine to Border Gateway Spine
- Border Gateway Spine to Spine
- Border Spine to Border Gateway Spine
- Border Gateway Spine to Border Spine

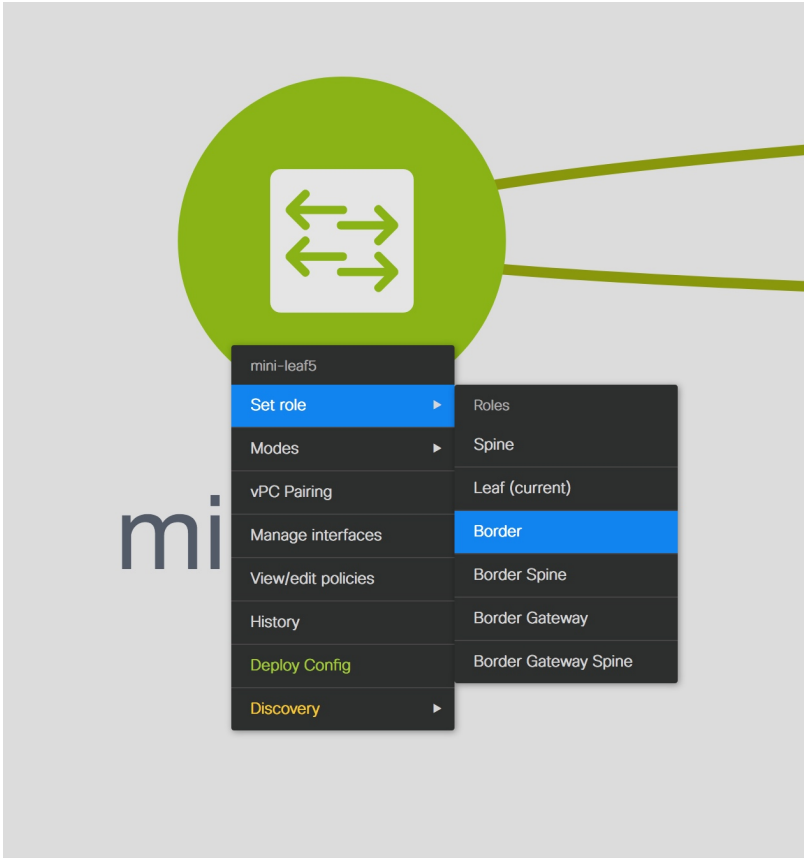

You cannot change the switch role from any Leaf role to any Spine role and from any Spine role to any Leaf role.

In case the switch role is not changed according to the allowed switch role changes mentioned above, the following error is displayed after you click **Save and Deploy**:

```
Switch[<serial-number>]: Role change from <switch-role> to <switch-role> is not permitted.
```

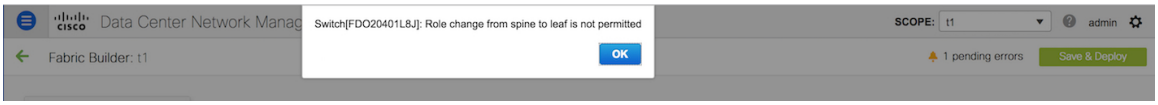

You can then change the switch role to the role that was set earlier, or set a new role, and configure the fabric.

If you have not created any policy template instances before clicking **Save and Deploy**, and there are no overlays, you can change the role of a switch to any other required role.

If you change the switch role of a vPC switch that is part of a vPC pair, the following error appears when you click **Save and Deploy**:

```
Switches role should be the same for VPC pairing. peer1 <serial-number>: [<switch-role>],
peer2 <serial-number>: [<switch-role>]
```

To prevent this scenario, change the switch roles of both the switches in the vPC pair to the same role.

Fabric Multi Switch Operations

In the fabric topology screen, click Tabular view option in the Actions panel, at the left part of the screen. The Switches | Links screen comes up.

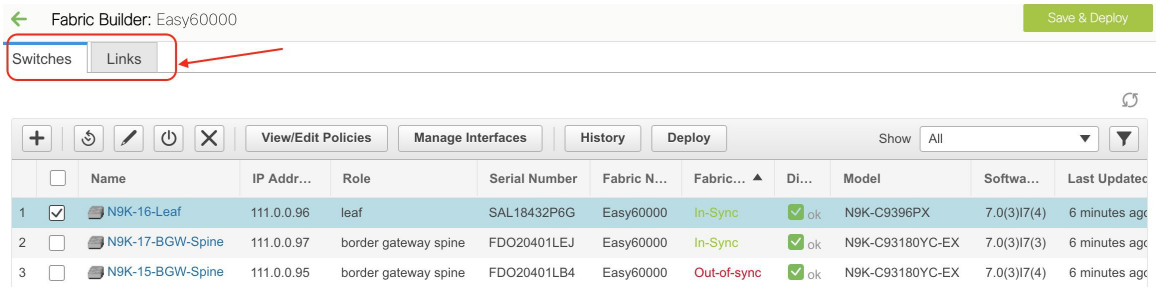

The Switches tab is for managing switch operations and the Links tab is for adding and updating fabric links. Each row represents a switch in the fabric, and displays switch details, including its serial number.

The buttons at the top of the table are explained, from left to right direction. Some options are also available when you right-click the switch icon. However, the Switches tab enables you to provision configurations on multiple switches (for example, adding and deploying policies) simultaneously.

- Add switches to the fabric. This option is also available in the topology page (Add switches option in Actions panel).
- Initiate the switch discovery process by DCNM afresh.
- Update device credentials such as authentication protocol, username and password.
- Reload the switch.
- View/Edit Policies: Add, update and delete a policy. The policies are template instances of templates in the template library. After creating a policy, you should deploy it on the switches using the Deploy option available in the View/edit Policies screen. You can select more than one policy and view them.

  Note If you select multiple switches and deploy a policy instance, then it will be deployed on all the selected switches.

- Manage Interfaces: Deploy configurations on the switch interfaces.
- **History** - View per switch deployment history.
- Deploy: Deploy switch configurations.

Fabric Links

You can add links between border switches of different fabrics (inter-fabric links) or between switches in the same fabric (intra-fabric links). You can only create an inter-fabric connection (IFC) for a switch that is managed by DCNM.

There are scenarios where you might want to define links between switches before connecting them physically. The links could be inter-fabric or intra-fabric links. Doing so, you can express and represent your intent to

add links. The links with intent are displayed in a different colour till they are actually converted to functional links. Once you physically connect the links, they are displayed as connected.

Management links might show up in the fabric topology as red colored links. To remove such links, right-click the link and click **Delete Link**.

From Cisco DCNM Release 11.1(1), the Border Spine and Border Gateway Spine roles are added to switch roles for border switches.

You can create links between existing and pre-provisioned devices as well by selecting the pre-provisioned device as the destination device.

Creating Intra-Fabric Links

1. Click Control > Fabric Builder to go to the Fabric Builder screen.
2. Click within the rectangular box that represents the fabric. The fabric topology screen comes up.
3. Click Tabular view in the Actions panel that is displayed at the left part of the screen.

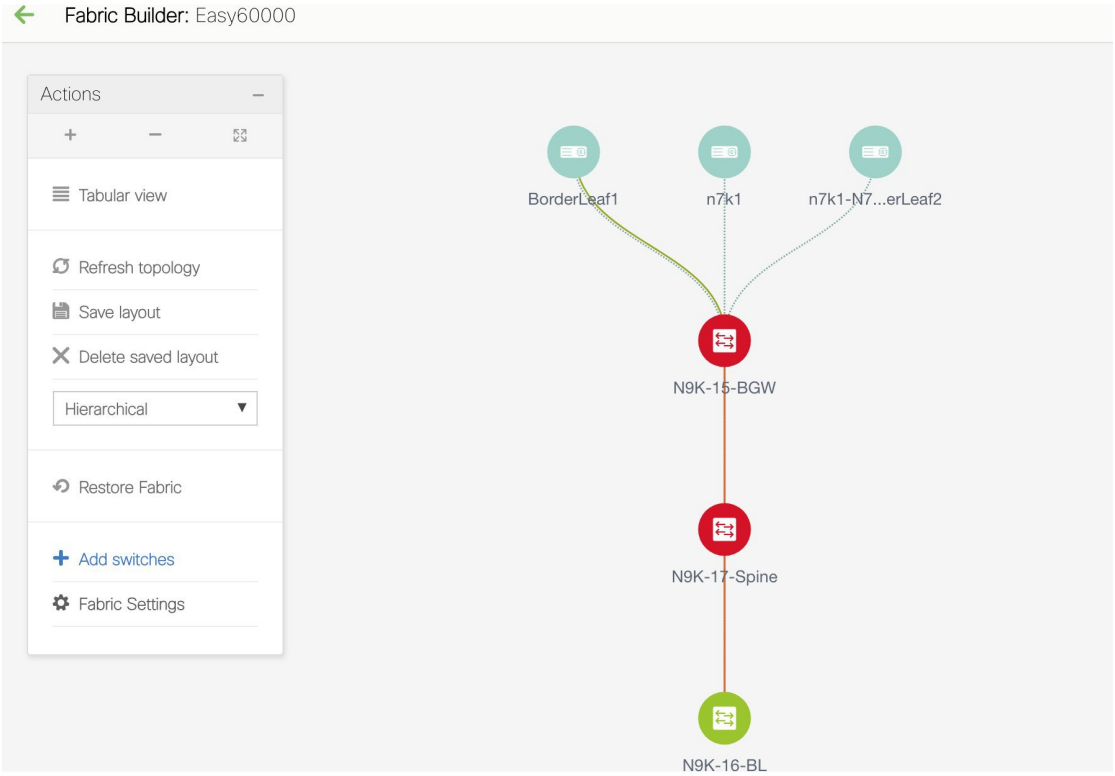

A screen with the tabs Switches and Links appears. They list the fabric switches and links in a table.

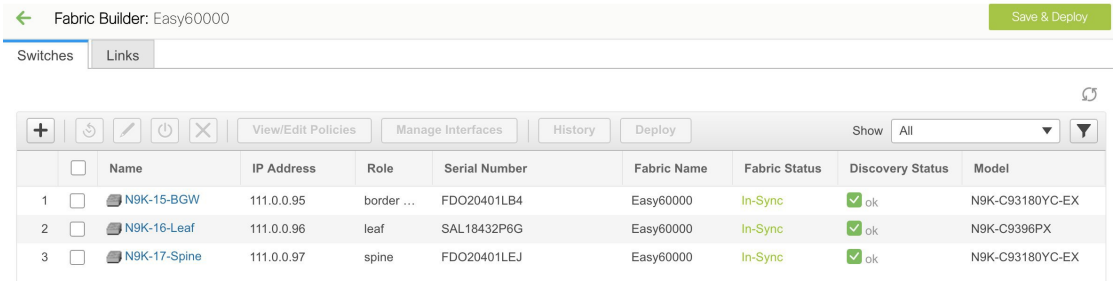

4. Click the Links tab. You can see a list of links.

   The list is empty when you are yet to create a link.

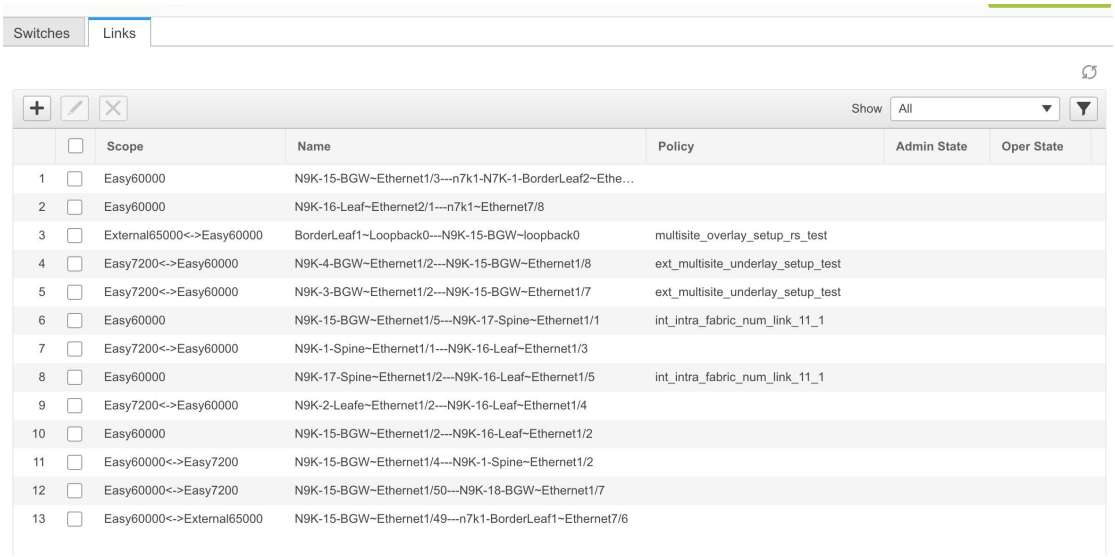

5. Click the Add (+) button at the top left part of the screen to add a link.

   The Add Link screen comes up. By default, the Intra-Fabric option is chosen as the link type.

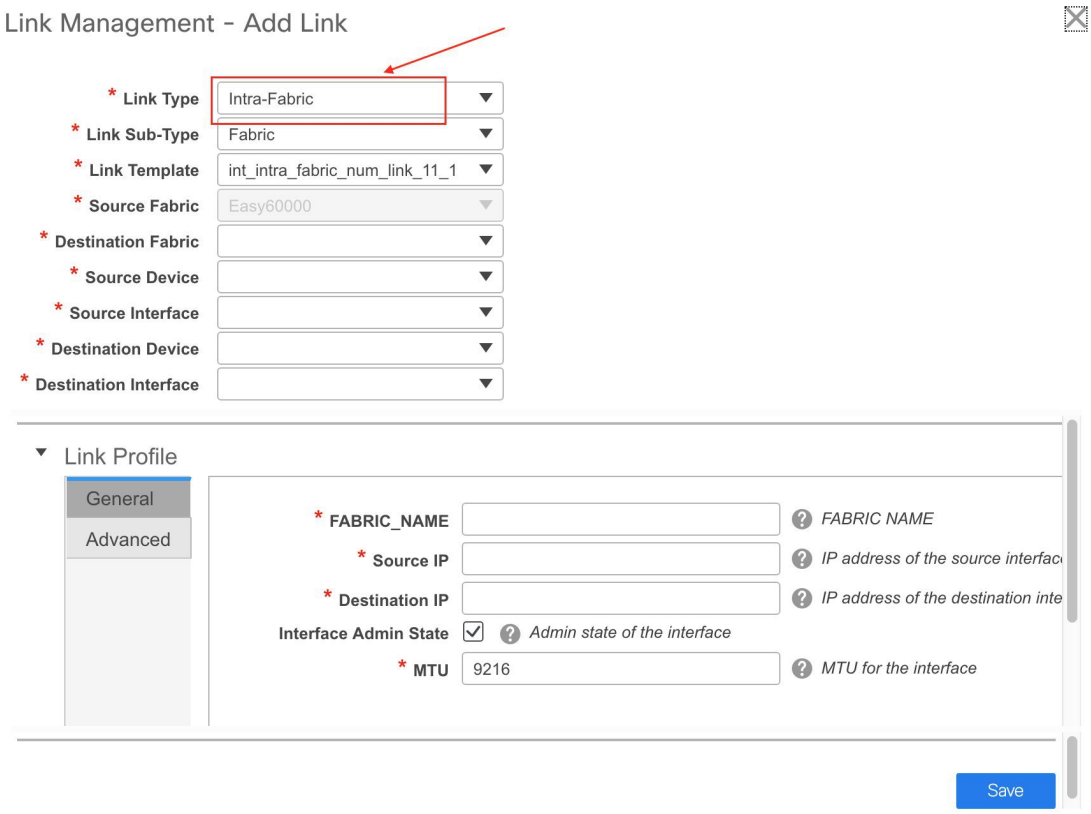

The fields are:

Link Type – Choose Intra-Fabric to create a link between two switches in a fabric.

Link Sub-Type – This field populates Fabric indicating that this is a link within the fabric.

Link Template: You can choose any of the following link templates.

- int_intra_fabric_num_link_11_1: If the link is between two ethernet interfaces assigned with IP addresses, choose int_intra_fabric_num_link_11_1.
- int_intra_fabric_unnum_link_11_1: If the link is between two IP unnumbered interfaces, choose int_intra_fabric_unnum_link_11_1.
- int_intra_vpc_peer_keep_alive_link_11_1: If the link is a vPC peer keep-alive link, choose int_intra_vpc_peer_keep_alive_link_11_1.
- int_pre_provision_intra_fabric_link: If the link is between two pre-provisioned devices, choose int_pre_provision_intra_fabric_link. After you click **Save & Deploy**, an IP address is picked from the underlay subnet IP pool.

Correspondingly, the Link Profile section fields is updated.

Source Fabric – The fabric name populates this field since the source fabric is known.

Destination Fabric – Choose the destination fabric. For an intra-fabric link, source and destination fabrics are the same.

Source Device and Source Interface – Choose the source device and interface.

Destination Device and Destination Interface – Choose the destination device and interface.

Note Select the pre-provisioned device as the destination device if you are creating a link between an existing device and a pre-provisioned device.

General tab in the Link Profile section

Interface VRF – Name of a non-default VRF for this interface.

Source IP and Destination IP – Specify the source and destination IP addresses of the source and destination interfaces, respectively.

Note The Source IP and Destination IP fields do not appear if you choose **int_pre_provision_intra_fabric_link** template.

Interface Admin State – Check or uncheck the check box to enable or disable the admin sate of the interface.

MTU – Specify the maximum transmission unit (MTU) through the two interfaces.

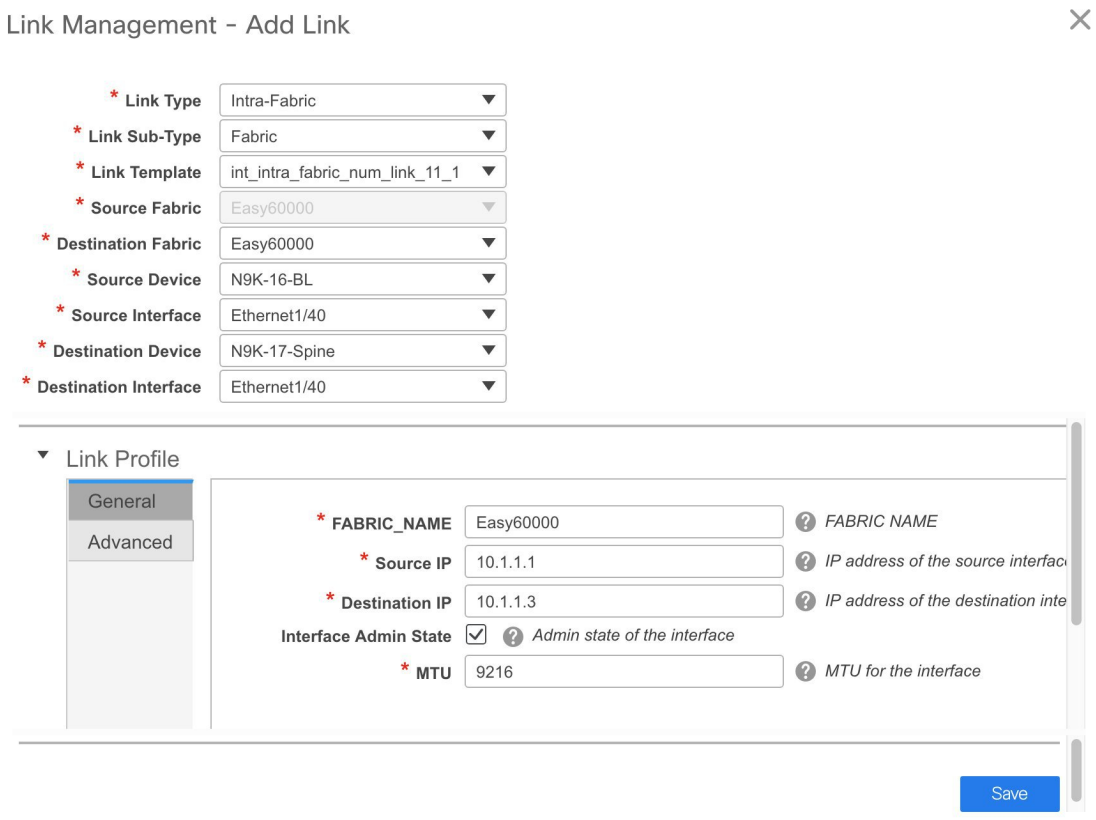

Advanced tab.

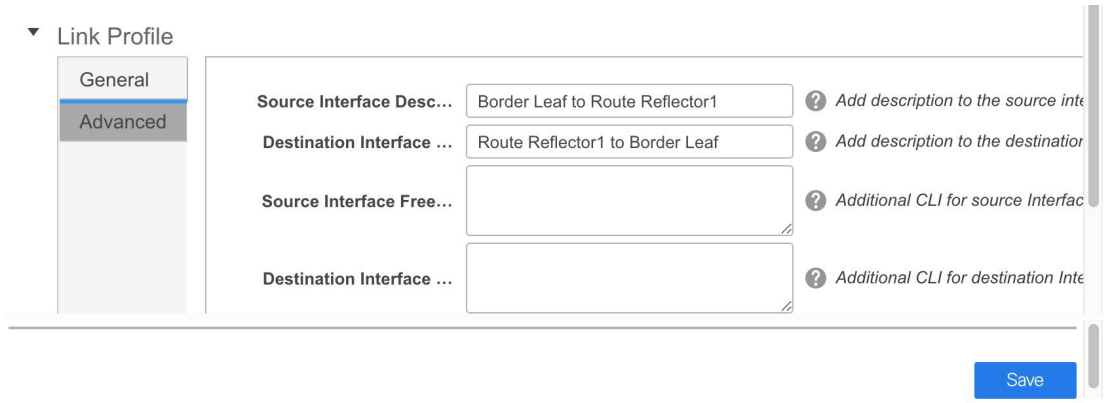

Source Interface Description and Destination Interface Description – Describe the links for later use. For example, if the link is between a leaf switch and a route reflector device, you can enter the information in these fields (Link from leaf switch to RR 1 and Link from RR 1 to leaf switch). This description will converted into a config, but will not be pushed into the switch. After **Save & Deploy**, it will reflect in the running configuration.

Source Interface Freeform CLIs and Destination Interface Freeform CLIs: Enter the freeform configurations specific to the source and destination interfaces. You should add the configurations as displayed in the running configuration of the switch, without indentation. For more information, refer Enabling Freeform Configurations on Fabric Switches .

6. Click Save at the bottom right part of the screen.

   The new link appears in the Links tab.

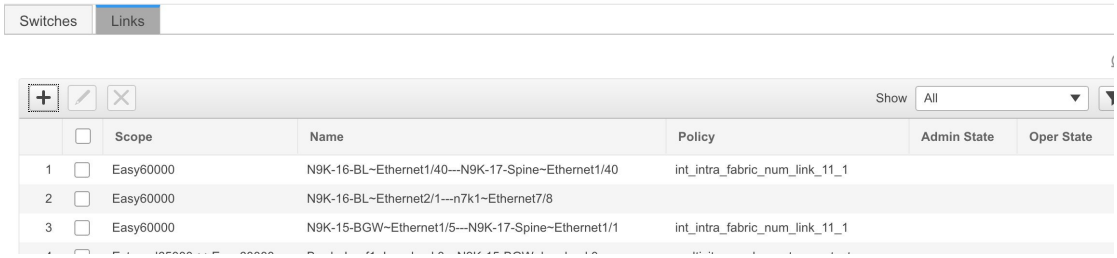

7. Click **Save & Deploy** to deploy the link configurations on the switches.

   The Config Deployment screen comes up. It displays the configuration status on the switches. You can also view the pending configurations by clicking the respective link in the Preview Config column. When you click a link in the Preview Config column, the Config Preview window comes up. It lists the pending configurations on the switch. The Side-by-side Comparison tab displays the running configuration and expected configuration side-by-side.

8. Close the preview screen and click Deploy Config. The pending configurations are deployed.

9. After ensuring that the progress is 100% in all the rows, click Close at the bottom part of the screen. The Links screen comes up again.

   Click <- at the top left part of the screen to go to the fabric topology. In the fabric topology, you can see that the link between the two devices is displayed.

Creating Inter-Fabric Links

1. Click the Links tab in the Switches | Links page. The list of previously created links are displayed. The list contains intra-fabric links (between switches in a fabric), and inter-fabric links (between BGWs or border leaf/spine switches of different fabrics).

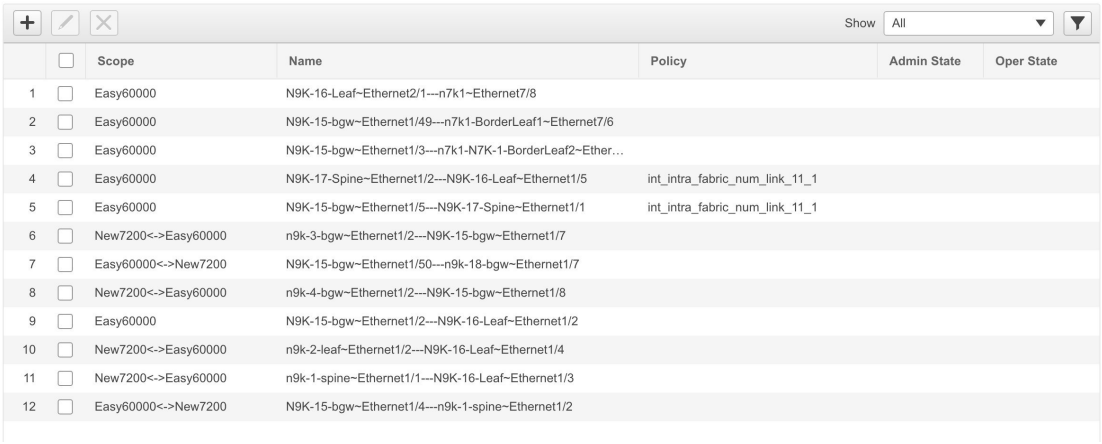

Show All

	Scope	Name	Policy	Admin State	Oper State
1	Easy60000	N9K-16-Leaf~Ethernet2/1---n7k1~Ethernet7/8			
2	Easy60000	N9K-15-bgw~Ethernet1/49---n7k1-BorderLeaf1~Ethernet7/6			
3	Easy60000	N9K-15-bgw~Ethernet1/3---n7k1-N7K-1-BorderLeaf2~Ether…			
4	Easy60000	N9K-17-Spine~Ethernet1/2---N9K-16-Leaf~Ethernet1/5	int_intra_fabric_num_link_11_1		
5	Easy60000	N9K-15-bgw~Ethernet1/5---N9K-17-Spine~Ethernet1/1	int_intra_fabric_num_link_11_1		
6	New7200<->Easy60000	n9k-3-bgw~Ethernet1/2---N9K-15-bgw~Ethernet1/7			
7	Easy60000<->New7200	N9K-15-bgw~Ethernet1/50---n9k-18-bgw~Ethernet1/7			
8	New7200<->Easy60000	n9k-4-bgw~Ethernet1/2---N9K-15-bgw~Ethernet1/8			
9	Easy60000	N9K-15-bgw~Ethernet1/2---N9K-16-Leaf~Ethernet1/2			
10	New7200<->Easy60000	n9k-2-leaf~Ethernet1/2---N9K-16-Leaf~Ethernet1/4			
11	New7200<->Easy60000	n9k-1-spine~Ethernet1/1---N9K-16-Leaf~Ethernet1/3			
12	Easy60000<->New7200	N9K-15-bgw~Ethernet1/4---n9k-1-spine~Ethernet1/2			

2. Click the Add (+) button at the top left part of the screen to add a link. The Add Link screen comes up.

   By default, the Intra-Fabric option is chosen as the link type.

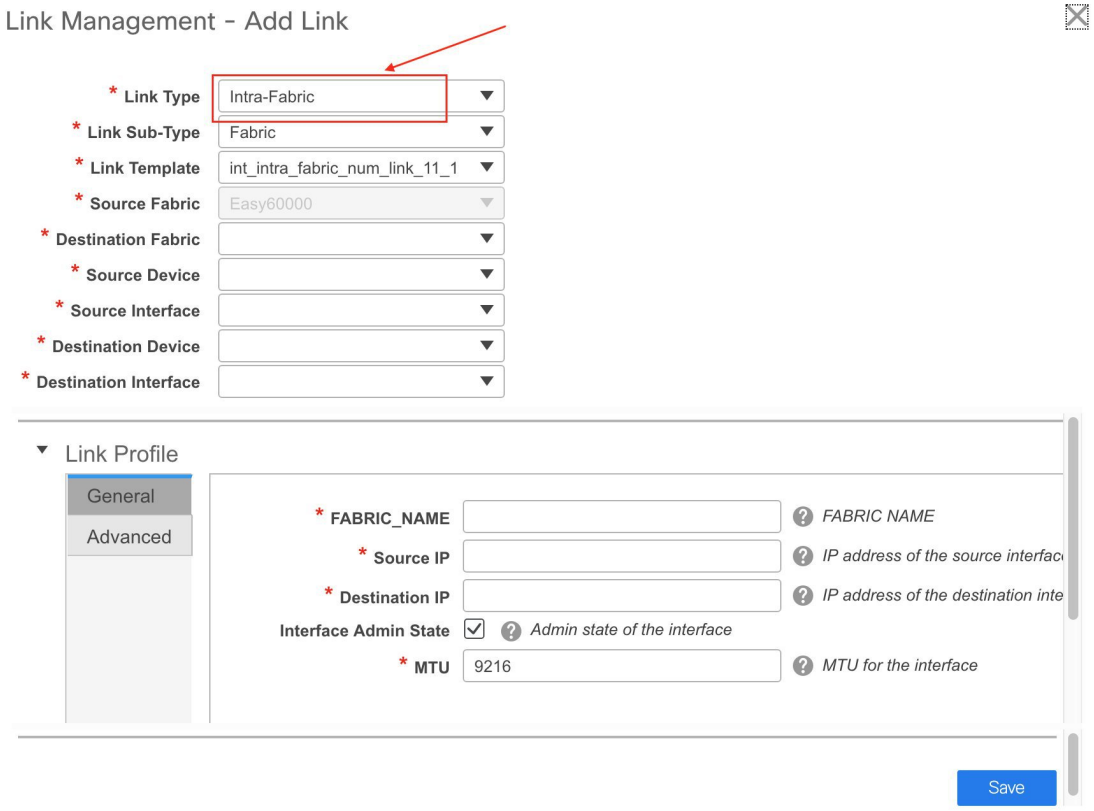

3. From the Link Type drop-down box, choose Inter-Fabric since you are creating an IFC. The screen changes correspondingly.

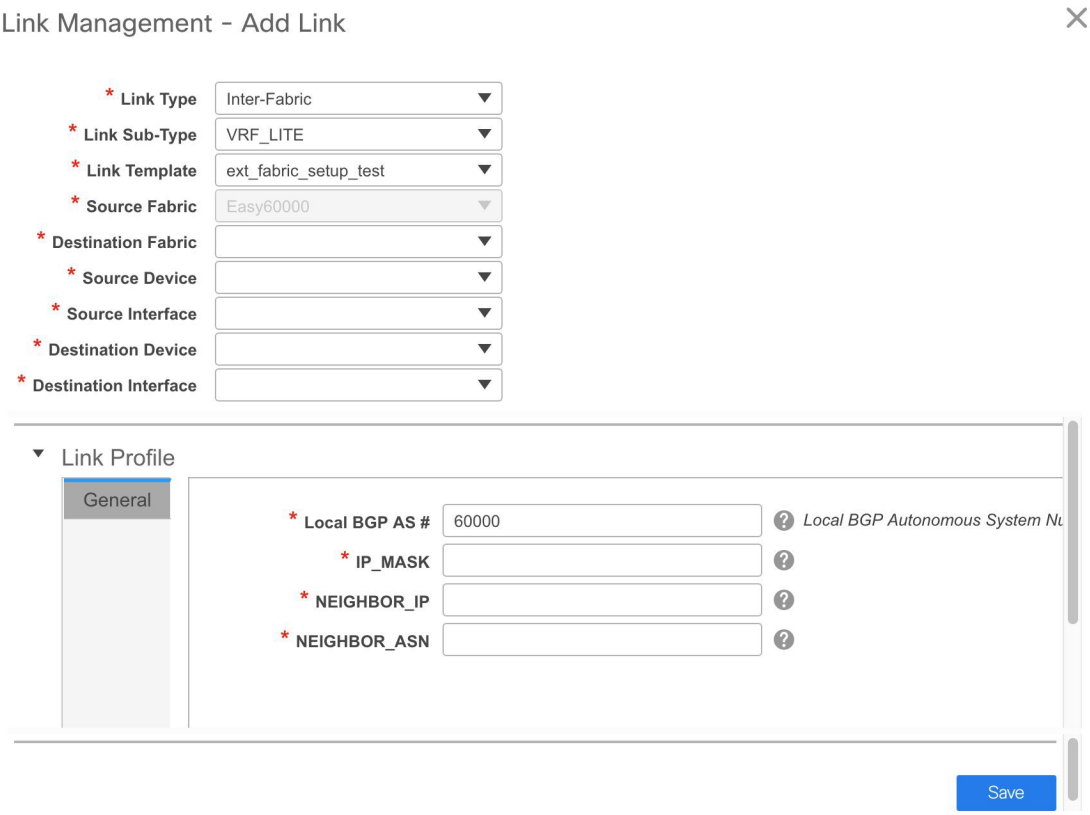

The fields for inter-fabric link creation are explained:

Link Type – Choose Inter-Fabric to create an inter-fabric connection between two fabrics, via their border switches.

Link Sub-Type – This field populates the IFC type. Choose **VRF_LITE**, **MULTISITE_UNDERLAY**, or **MULTISITE_OVERLAY** from the drop-down list.

The Multi-Site options are explained in the Multi-Site use case.

Link Template: The link template is populated.

The templates are autopopulated with corresponding pre-packaged default templates that are based on your selection.

Note You can add, edit, or delete user-defined templates. See *Template Library* section in the Control chapter for more details.

Source Fabric - This field is prepopulated with the source fabric name.

Destination Fabric - Choose the destination fabric from this drop-down box.

Source Device and Source Interface - Choose the source device and Ethernet interface that connects to the destination device.

Destination Device and Destination Interface—Choose the destination device and Ethernet interface that connects to the source device.

Based on the selection of the source device and source interface, the destination information is autopopulated based on Cisco Discovery Protocol information, if available. There is an extra validation performed to ensure that the destination external device is indeed part of the destination fabric.

General tab in the Link Profile section.

Local BGP AS# - In this field, the AS number of the source fabric is autopopulated.

IP_MASK—Fill up this field with the IP address of the source interface that connects to the destination device.

NEIGHBOR_IP—Fill up this field with the IP address of the destination interface.

NEIGHBOR_ASN—In this field, the AS number of the destination device is autopopulated.

After filling up the Add Link screen, it looks like this:

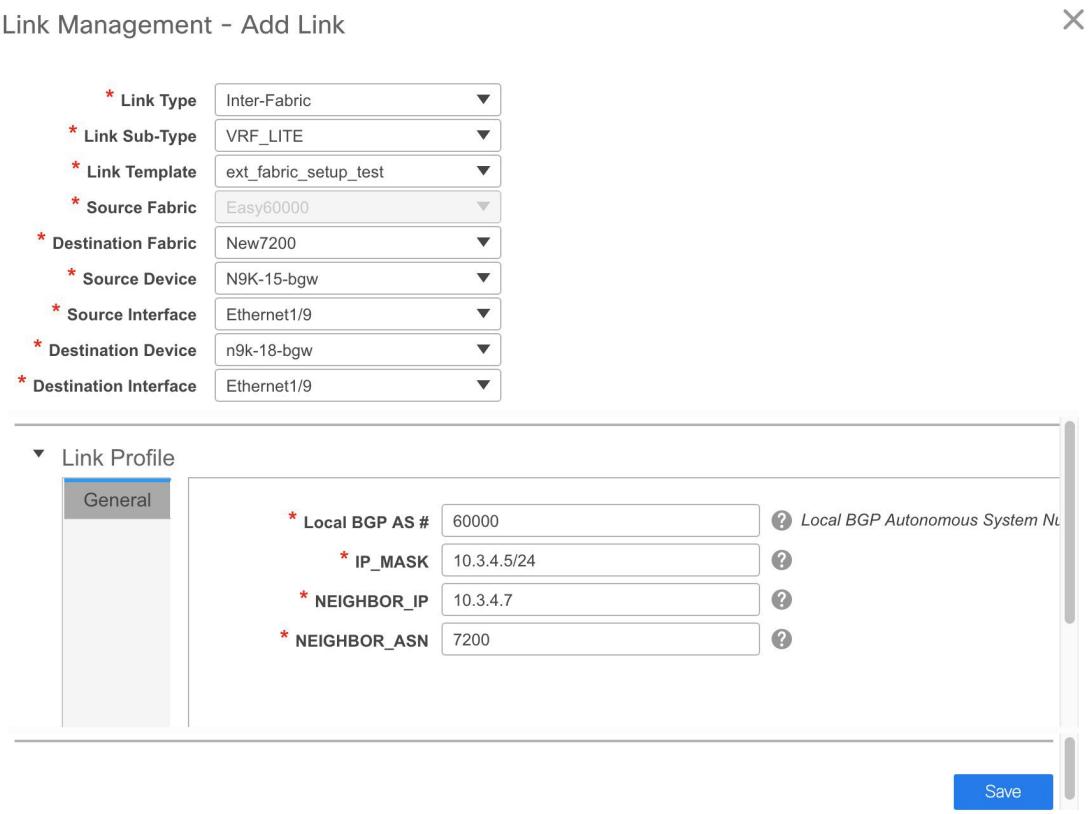

4. Click Save at the bottom right part of the screen.

   The Switches|Links screen comes up again. You can see that the IFC is created and displayed in the list of links.

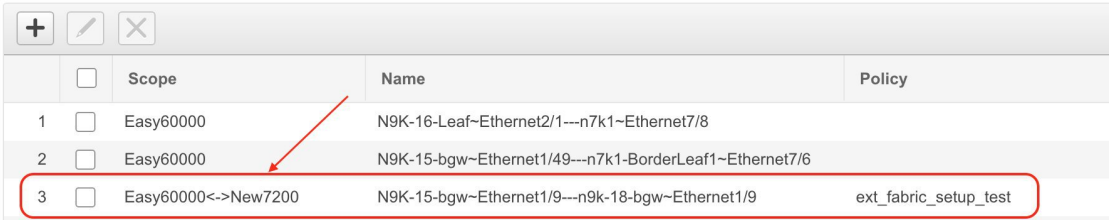

5. Click on Save & Deploy to deploy the link configurations on the switches.

   The Config Deployment screen comes up. It displays the configuration status on the switches. You can also view the pending configurations by clicking the respective link in the Preview Config column. When you click a link in the Preview Config column, the Config Preview window comes up. It lists the pending configurations on the switch. The Side-by-side Comparison tab displays the running configuration and expected configuration side-by-side.

6. Close the preview screen and click Deploy Config. The pending configurations are deployed.

7. After ensuring that the progress is 100% in all the rows, click Close at the bottom part of the screen. The Links screen comes up again.

8. Click <- at the top left part of the screen to go to the fabric topology. In the fabric topology, you can see that the link between the two devices is displayed.

   If the two fabrics are member fabric of an MSD, then you can see the link in the MSD topology too.

When you enable the VRF Lite function using the ToExternalOnly method or Multisite function via MSD fabric, IFCs are automatically created between the (VXLAN fabric) border/BGW device and connected (external fabric) edge router/core device. When you remove the ER/core/border/BGW device, the corresponding IFCs (link PTIs) to/from that switch are deleted on DCNM. Subsequently, DCNM removes the corresponding IFC configurations, if any, from the remaining devices on the next Save & Deploy operation. Also, if you want to remove a device that has an IFCs and overlay extensions over those IFCs, you should undeploy all overlay extensions corresponding to those IFCs for switch delete to be possible.

To undeploy VRF extensions, click Control > Networks & VRFs, select the VXLAN fabric and the extended VRFs, and undeploy the VRFs in the VRF deployment screen.

To delete the IFCs, click Control > Fabric Builder, go to the fabric topology screen, click Tabular view, and delete the IFCs from the Links tab.

Ensure that the fabric switch names are unique. If you deploy VRF extensions on switches with the same name, it leads to erroneous configuration.

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. Refer Interfaces.
- Create overlay networks and VRFs and deploy them on the switches. Refer Creating and Deploying Networks and VRFs .

Exporting Links

1. Choose **Control** > **Fabric Builder**, and select a fabric.

The fabric topology window appears.

2. Click **Tabular view** in the **Actions** panel.

   A window with the **Switches** and **Links** tabs appears.

3. Click the **Links** tab.

   You can see a list of links. The list is empty when you are yet to create a link.

4. Click the **Export Links** icon to export the links in a CSV file.

   The following details of links are exported: link template, source fabric, destination fabric, source device, destination device, source switch name, destination switch name, source interface, destination interface, and nvPairs. The nvPairs field consits JSON object.

Importing Links

You can import a CSV file containing details of links to add new links to the fabric. The CSV file should have the following details of links: link template, source fabric, destination fabric, source device, destination device, source switch name, destination switch name, source interface, destination interface, and nvPairs.

Note

- You cannot update existing links.
- The **Import Links** icon is disabled for external fabric.

1. Choose **Control** > **Fabric Builder**, and select a fabric.

   The fabric topology window appears.

2. Click **Tabular view** in the **Actions** panel.

   A window with the **Switches** and **Links** tabs appears.

3. Click the **Links** tab.

   You can see a list of links. The list is empty when you are yet to create a link.

4. Click the **Import Links** icon.

   The file server directory opens.

5. Browse the directory and select the CSV file that you want to import.

6. Click **Open**.

   A confirmation screen appears.

7. Click **Yes** to import the selected file.

Viewing Details of Fabric Links

You can view information about a fabric link, like IP subnet between links to deploy underlay, MTU, speed mismatch, and so on, in the topology view of a fabric builder. To view the details of a link from the Cisco DCNM Web client, perform the following steps:

Procedure

Step 1 Choose **Control** > **Fabrics** > **Fabric Builder** and select a fabric.

The topology view of the fabric appears.

Step 2 Double-click any of the links.

The details window appears. You can view the devices that are connected using this link, summary, and the data traffic.

Step 3 Click **Show more details**.

A comparison table of the two devices connected by the link appears. It includes the following parameters of the devices: device name, name, admin status, operation status, reason, policies, overlay network, status, PC, vPC ID, speed, MTU, mode, VLANs, IP or prefix, VRF, neighbor, and description.

Note
- You can view the traffic details of a fabric link by clicking the device name with hyperlink. Alternatively, you can view these traffic details in the details window. See *Viewing the Traffic Details of the Fabric Links* section for more information.
- You can view the expected configuration of a fabric link by clicking the policy with the hyperlink.

Step 4 Click the **Back** icon to go back to the details window.

Note You can click the **Close** icon to exit the details window.

Viewing the Traffic Details of Fabric Links

In the details window of a fabric link, you can choose how you want to view the traffic details. You can view the traffic details based on the time duration, format, and export this information.

You can view the data traffic of a link for the following durations from the duration drop-down list:

- 24 Hours
- Week
- Month
- Year

Show: Click **Show**, and choose **Chart**, **Table**, or **Chart and Table** from the drop-down list to see how you want to view the traffic details. Enlarge your browser window to view the details in **Chart and Table** format.

If you choose **Chart**, hover over the traffic chart to view the Rx and Tx values, along the Y axis, for the corresponding time, along X axis. You can change the time duration values of the X axis by moving the sliders in the time range selector. You can choose the Y-axis values by checking or unchecking the Rx and Tx check boxes.

Note If you select **Week**, **Month**, or **Year** as the time duration, you can also view the Peak Rx and Peak Tx values along the Y axis.

Select **Table** to view the traffic information in tabular format.

Chart Type and Chart Options: Choose **Area Chart** or **Line Chart** from the **Chart Type** drop-down list.

You can choose the following chart options:

- **Show Fill Patterns**
- **Show Datamarkers**
- **Y Axis Log Scale**

Actions: Export or print the traffic information by choosing the appropriate options from the **Actions** drop-down list.

vPC Fabric Peering

You can create a virtual peer link for two switches or change the existing physical peer link to a virtual peer link. Only greenfield deployments support vPC fabric peering in Cisco DCNM, Release 11.2(1). This feature is applicable for **Easy_Fabric_11_1** and **Easy_Fabric_eBGP** fabric templates.

Guidelines and Limitations

The following are the guidelines and limitations for vPC fabric pairing.

- vPC fabric peering is supported from Cisco DCNM Release 11.2(1) and Cisco NX-OS Release 9.2(3).
- Only Cisco Nexus N9K-C9332C Switch, Cisco Nexus N9K-C9364C Switch, Cisco Nexus N9K-C9348GC-FXP Switch as also the Cisco Nexus 9000 Series Switches that ends with FX, FX2, and FX2-Z support vPC fabric peering.
- If you use other Cisco Nexus 9000 Series Switches, a warning will appear during **Save & Deploy**. A warning appears in this case because these switches will be supported in future releases.
- If you try pairing switches that do not support vPC fabric peering, using the **Use Virtual Peerlink** option, a warning will appear when you deploy the fabric.
- You can convert a physical peer link to a virtual peer link and vice-versa with or without overlays.
- Switches with border gateway leaf roles do not support vPC fabric peering.
- vPC fabric peering is not supported for Cisco Nexus 9000 Series Modular Chassis and FEXs. An error appears during **Save & Deploy** if you try to pair any of these.
- Only greenfield deployments support vPC fabric peering.
- However, you can import switches that are connected using physical peer links and convert the physical peer links to virtual peer links after **Save & Deploy**. To update a TCAM region during the feature configuration, use the **hardware access-list tcam ingress-flow redirect** *512* command in the configuration terminal.

Fields and Description

To view the vPC pairing window of a switch, from the fabric topology window, right-click the switch and choose **vPC Pairing**. The vPC pairing window for a switch has the following fields:

Field	Description
Use Virtual Peerlink	Allows you to enable or disable the virtual peer linking between switches.
Switch name	Specifies all the peer switches in a fabric. **Note** When you have not paired any peer switches, you can see all the switches in a fabric. After you pair a peer switch, you can see only the peer switch in the vPC pairing window.
Recommended	Specifies if the peer switch can be paired with the selected switch. Valid values are **true** and **false**. Recommended peer switches will be set to **true**.
Reason	Specifies why the vPC pairing between the selected switch and the peer switches is possible or not possible.
Serial Number	Specifies the serial number of the peer switches.

You can perform the following with the **vPC Pairing** option:

Creating a Virtual Peer Link

To create a virtual peer link from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control** > **Fabrics**.

The **Fabric Builder** window appears.

Step 2 Choose a fabric with the **Easy_Fabric_11_1** or **Easy_Fabric_eBGP** fabric templates.

The fabric topology window appears.

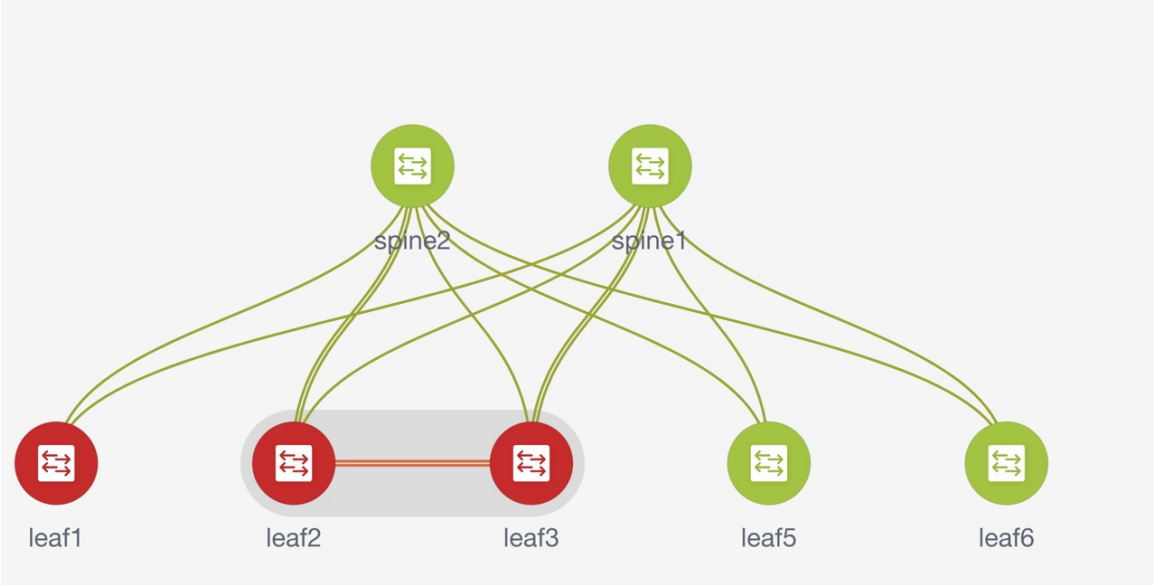

Step 3 Right-click a switch and choose **vPC Pairing** from the drop-down list.

The window to choose the peer appears.

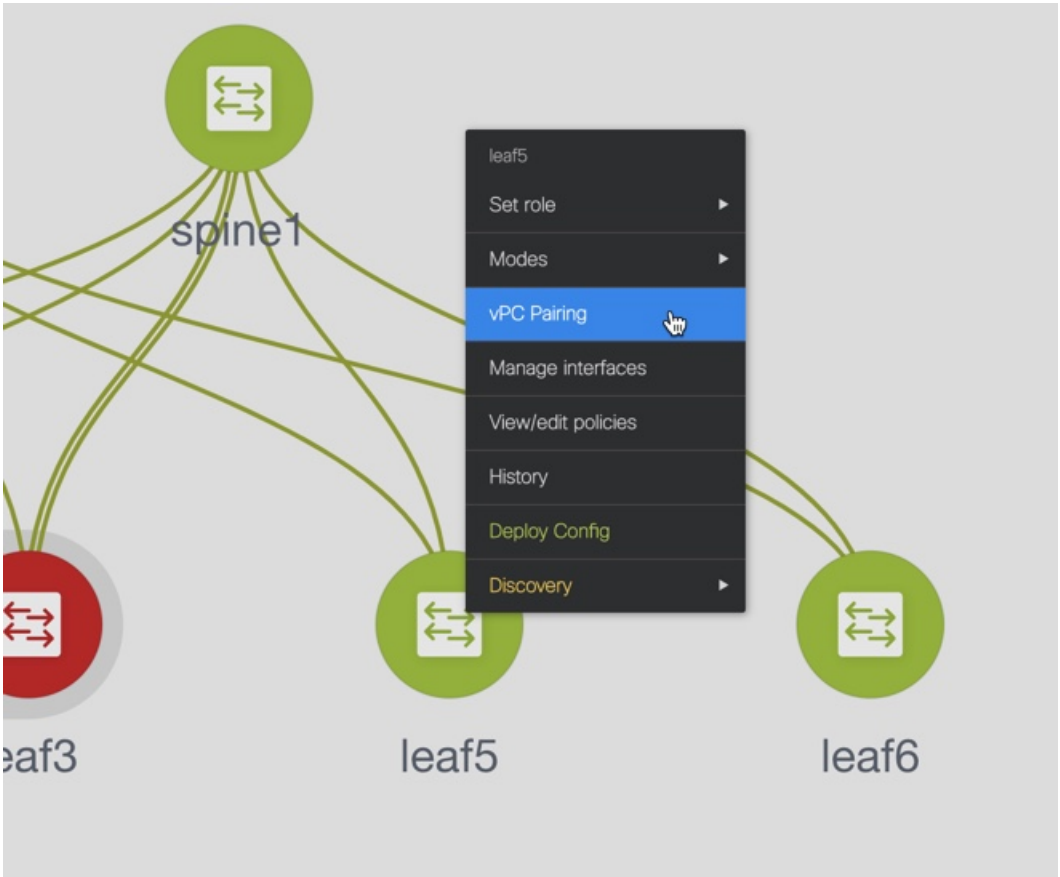

Note You will get the following error when you choose a switch with the border gateway leaf role.

```
<switch-name> has a Network/VRF attached. Please detach the Network/VRF before vPC
Pairing/Unpairing
```

Step 4 Check the **Use Virtual Peerlink** check box.

Step 5 Choose a peer switch and check the **Recommended** column to see if pairing is possible.

If the value is **true**, pairing is possible. You can pair switches even if the recommendation is **false**. However, you will get a warning or error during **Save & Deploy**.

Step 6 Click **Save**.

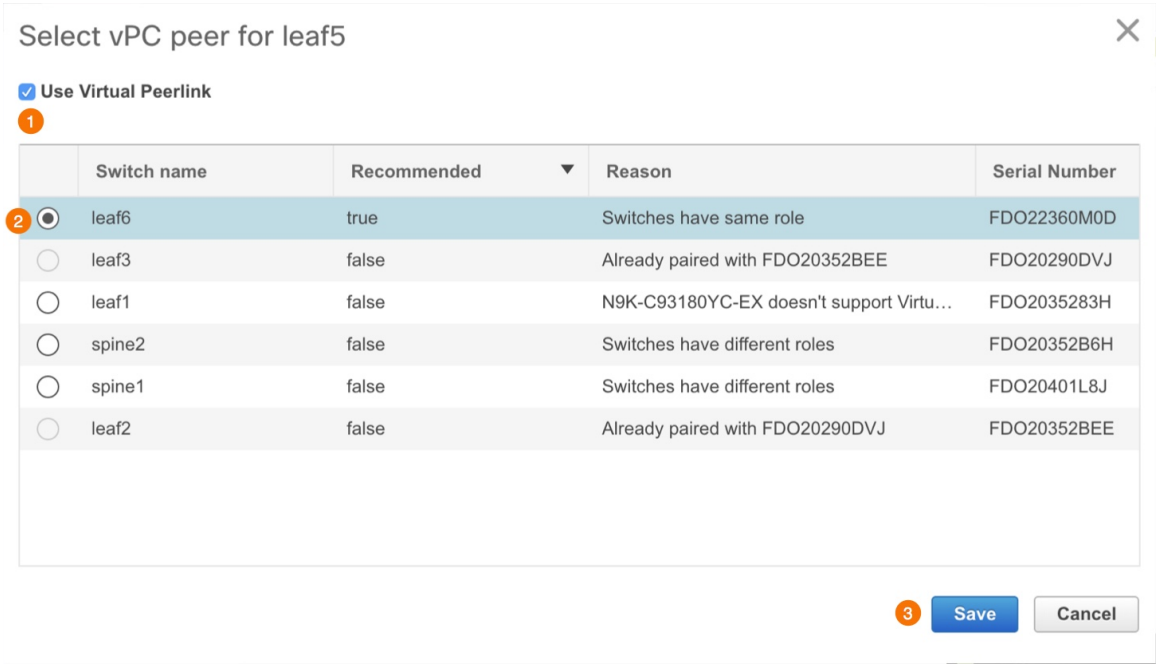

Step 7 In the **Fabric Topology** window, click **Save & Deploy**.

The **Config Deployment** window appears.

Step 8 Click the field against the switch in the **Preview Config** column.

The **Config Preview** window appears for the switch.

Step 9 View the vPC link details in the pending configuration and the side-by-side configuration.

Step 10 Close the window.

Step 11 Click the pending errors icon next to the **Save & Deploy** icon to view errors and warnings, if any.

If you see any warnings that are related to TCAM, click the **Resolve** icon. A confirmation dialog box about reloading switches appears. Click **OK**. You can also reload the switches from **Tabular view** in the fabric topology window.

The switches that are connected through vPC fabric peering, are enclosed in a gray cloud.

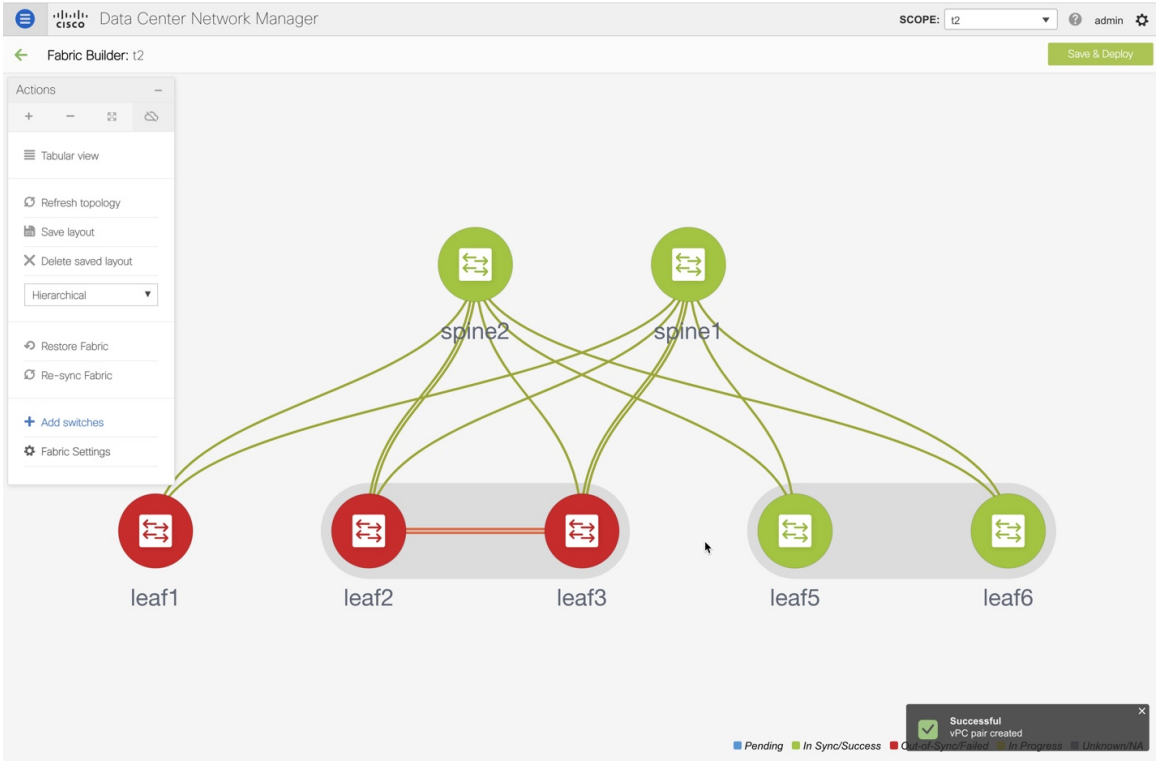

Converting a Physical Peer Link to a Virtual Peer Link

To convert a physical peer link to a virtual peer link from the Cisco DCNM Web UI, perform the following steps:

Before you begin

- Plan the conversion from physical peer link to virtual peer link during the maintenance window of switches.
- Ensure the switches support vPC fabric peering. Only the following switches support vPC fabric peering:
  - Cisco Nexus N9K-C9332C Switch, Cisco Nexus N9K-C9364C Switch, and Cisco Nexus N9K-C9348GC-FXP Switch
  - Cisco Nexus 9000 Series Switches that ends with FX, FX2, and FX2-Z

Procedure

Step 1 Choose **Control > Fabrics**.

The **Fabric Builder** window appears.

Step 2 Choose a fabric with the **Easy_Fabric_11_1** or **Easy_Fabric_eBGP** fabric templates.

Step 3 Right-click the switch that is connected using the physical peer link and choose **vPC Pairing** from the drop-down list.

The window to choose the peer appears.

Note You will get the following error when you choose a switch with the border gateway leaf role.

```
<switch-name> has a Network/VRF attached. Please detach the Network/VRF before vPC
Pairing/Unpairing
```

Step 4 Check the **Recommended** column to see if pairing is possible.

If the value is **true**, pairing is possible. You can pair switches even if the recommendation is **false**. However, you will get a warning or error during **Save & Deploy**.

Step 5 Check the **Use Virtual Peerlink** check box.

The **Unpair** icon changes to **Save**.

Step 6 Click **Save**.

Note After you click **Save**, the physical vPC peer link is automatically deleted between the switches even without deployment.

Step 7 In the **Fabric Topology** window, click **Save & Deploy**.

The **Config Deployment** window appears.

Step 8 Click the field against the switch in the **Preview Config** column.

The **Config Preview** window appears for the switch.

Step 9 View the vPC link details in the pending configuration and the side-by-side configuration.

Step 10 Close the window.

Step 11 Click the pending errors icon next to the **Save & Deploy** icon to view errors and warnings, if any.

If you see any warnings that are related to TCAM, click the **Resolve** icon. A confirmation dialog box about reloading switches appears. Click **OK**. You can also reload the switches from **Tabular view** in the fabric topology window.

The physical peer link between the peer switches turns red. Delete this link. The switches are connected only through a virtual peer link and are enclosed in a gray cloud.

Converting a Virtual Peer Link to a Physical Peer Link

To convert a virtual peer link to a physical peer link from the Cisco DCNM Web UI, perform the following steps:

Before you begin

Connect the switches using a physical peer link before disabling the vPC fabric peering.

Procedure

Step 1 Choose **Control** > **Fabrics**.

The **Fabric Builder** window appears.

Step 2 Choose a fabric with the **Easy_Fabric_11_1** or **Easy_Fabric_eBGP** fabric templates.

Step 3 Right-click the switch that is connected through a virtual peer link and choose **vPC Pairing** from the drop-down list.

The window to choose the peer appears.

Step 4 Uncheck the **Use Virtual Peerlink** check box.

The **Unpair** icon changes to **Save**.

Step 5 Click **Save**.

Step 6 In the **Fabric Topology** window, click **Save & Deploy**.

The **Config Deployment** window appears.

Step 7 Click the field against the switch in the **Preview Config** column.

The **Config Preview** window appears for the switch.

Step 8 View the vPC peer link details in the pending configuration and the side-by-side configuration.

Step 9 Close the window.

Step 10 Click the pending errors icon next to the **Save & Deploy** icon to view errors and warnings, if any.

If you see any warnings that are related to TCAM, click the **Resolve** icon. The confirmation dialog box about reloading switches appears. Click **OK**. You can also reload the switches from **Tabular view** in the fabric topology window.

The virtual peer link, represented by a gray cloud, disappears and the peer switches are connected through a physical peer link.

Viewing and Editing Policies

Cisco DCNM provides the ability to group a set of switches, and allows you to push a set of underlay configurations to the group. This release enables you to create a policy template, and apply it to multiple selected switches.

To view, add, deploy, or edit a policy, perform the following steps:

Procedure

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

Step 2 Select any available fabric, and then click **Tabular view**.

Step 3 Select multiple switches in switches tab, and click **View/Edit Policies**.

Viewing Policies

Procedure

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

Step 2 Select any available fabric, and then click **Tabular view**.

Step 3 Select multiple switches in the switches tab and click **View/Edit Policies**.

Policies are listed in view or edit policies table for multiple switches.

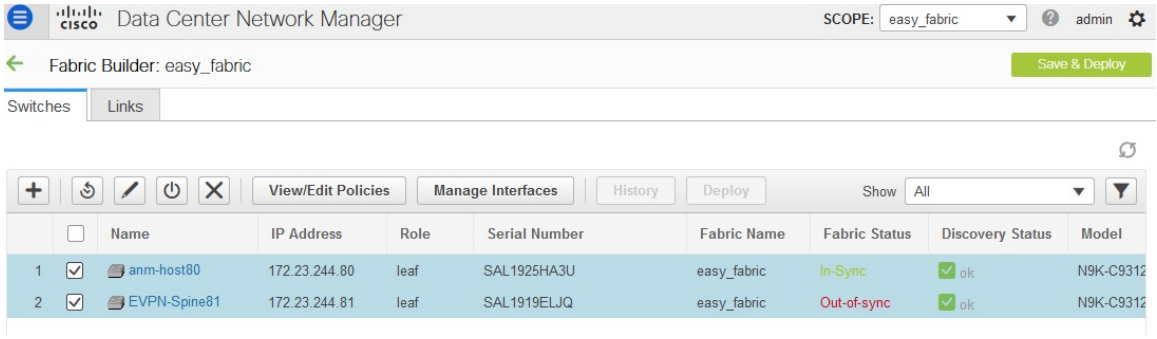

Step 4 Select a policy and click the **View** button to view its configs.

Adding a Policy

Procedure

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

Step 2 Select any available fabric, and then click Tabular view.

Step 3 Select a single or multiple switches in the **Switches** tab, and click the **View/Edit Policies** button.

Step 4 Click the **Add** icon.

Step 5 Select a policy template and enter the mandatory parameters data and click **Save**. PTI is added per each device based on n-number of devices selection.

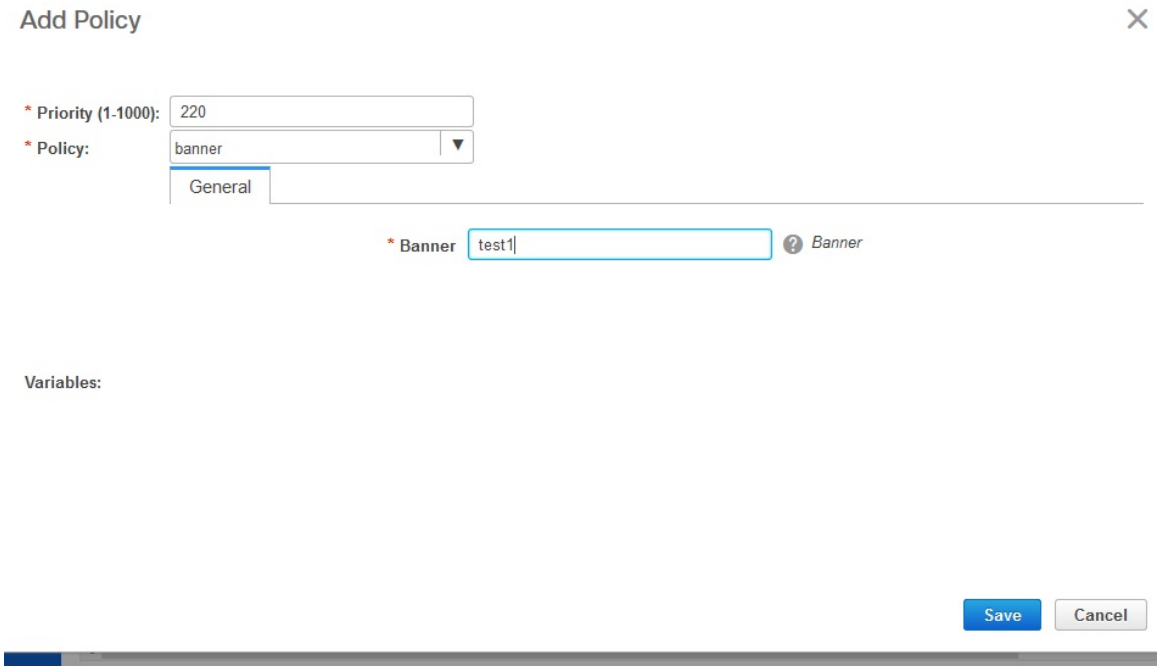

Policy: Select a policy from this drop-down list.

Priority: Specify a priority for the policy. The applicable values are from 1 to 1000. The default value is 500. The lower number in the **Priority** field means that there is a higher priority for the generated configuration and POAP startup-configuration. For example, features are 50, route-maps are 100, and vpc-domain is 200.

Deploying Policies

Procedure

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

Step 2 Select any available fabric, and then click Tabular view.

Step 3 Select multiple switches in the switches tab, and click the **View/Edit Policies** button.

Step 4 Select multiple polices, and then click **Push Config**. The selected PTI's configs are pushed to the group of switches.

Editing a Policy

Note Multiple policy editing is not supported.

Procedure

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

Step 2 Select any available fabric, and then click **Tabular view**.

Step 3 Select multiple switches in the switches tab, and click the **View/Edit Policies** button.

Step 4 Select a PTI, click **Edit** to modify the required data, and then click **Save** to save the PTI.

Step 5 Select a PTI, click **Edit** to modify the required data, and then click **Push Config** to push the policy config to the device.

Note
- A warning appears if you push config for a Python policy.
- A warning appears if you edit, delete, or push config a mark-deleted policy. A mark-deleted policy is set to **true** under the **Mark Deleted** column. The switch freeform child policies of **Mark Deleted** policies appears in the **View/Edit Policies** dialog box. You can edit only **Python** switch_freeform policies. You cannot edit **Template_CLI** switch_freeform_config policies.

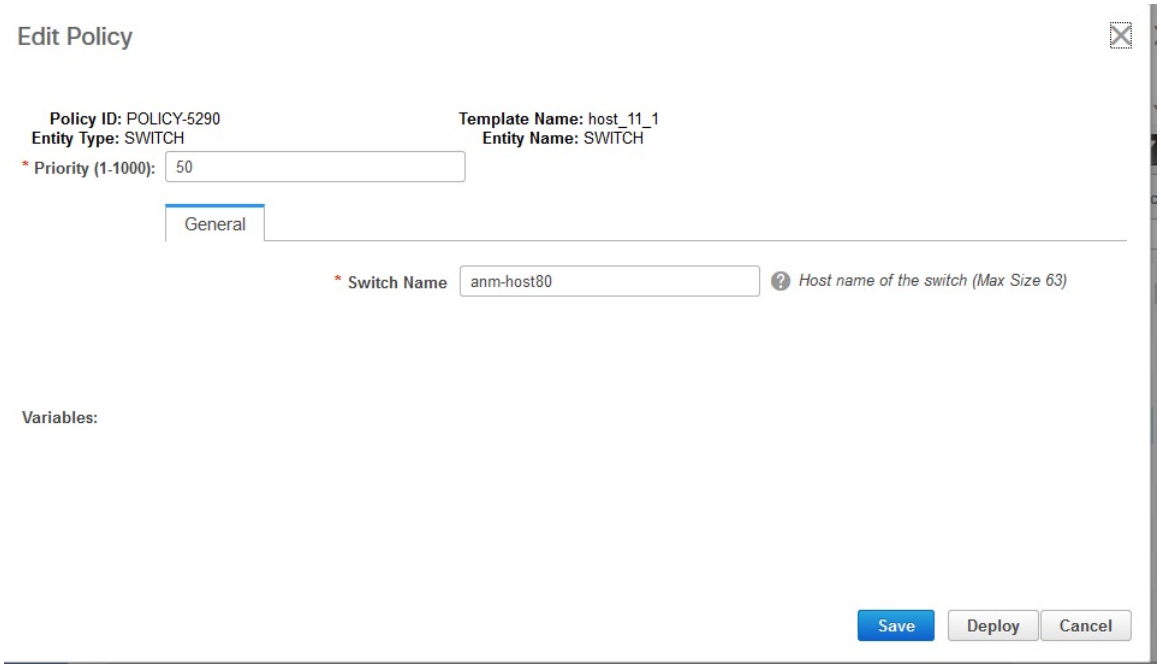

Current Switch Configuration

Procedure

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

Step 2 Select any available fabric, and then click **Tabular** view.

Step 3 Select multiple switches in the switches tab, and click **View/Edit Policies**.

Step 4 Click **Current Switch Config**.

The current switch configuration appears in the **Running Config** dialog box.

Note The running configuration will not appear for the Cisco CSR 1000v when you click **Current Switch Config** if the user role cannot access the enable prompt by default.

Retrieving the Authentication Key

Retrieving the 3DES Encrypted OSPF Authentication Key

1. SSH into the switch.
2. On an unused switch interface, enable the following:

```
config terminal
    feature ospf
    interface Ethernet1/1
       no switchport
       ip ospf message-digest-key 127 md5 ospfAuth
```

In the example, **ospfAuth** is the unencrypted password.

Note This Step 2 is needed when you want to configure a new key.

3. Enter the **show run interface Ethernet1/1** command to retrieve the password.

```
Switch # show run interface Ethernet1/1
   interface Ethernet1/1
     no switchport
     ip ospf message-digest key 127 md5 3 sd8478f4fsw4f4w34sd8478fsdfw
     no shutdown
```

The sequence of characters after **md5 3** is the encrypted password.

4. Update the encrypted password into the **OSPF Authentication Key** field.

Retrieving the Encrypted IS-IS Authentication Key

To get the key, you must have access to the switch.

1. SSH into the switch.
2. Create a temporary keychain.

```
config terminal
   key chain isis
   key 127
   key-string isisAuth
```

In the example, **isisAuth** is the plaintext password. This will get converted to a Cisco type 7 password after the CLI is accepted.

3. Enter the **show run | section "key chain"** command to retrieve the password.

```
key chain isis
   key 127
       key-string 7 071b245f5a
```

   The sequence of characters after key-string 7 is the encrypted password. Save it.

4. Update the encrypted password into the ISIS Authentication Key field.
5. Remove any unwanted configuration made in Step 2.

Retrieving the 3DES Encrypted BGP Authentication Key

1. SSH into the switch and enable BGP configuration for a non-existent neighbor.

Note Non-existent neighbor configuration is a temporary BGP neighbor configuration for retrieving the password.

```
router bgp
   neighbor 10.2.0.2 remote-as 65000
   password bgpAuth
```

   In the example, **bgpAuth** is the unencrypted password.

2. Enter the show run bgp command to retrieve the password. A sample output:

```
neighbor 10.2.0.2
   remote-as 65000
   password 3 sd8478fswerdfw3434fsw4f4w34sdsd8478fswerdfw3434fsw4f4w3
```

   The sequence of characters after password 3 is the encrypted password.

3. Update the encrypted password into the **BGP Authentication Key** field.
4. Remove the BGP neighbor configuration.

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**.

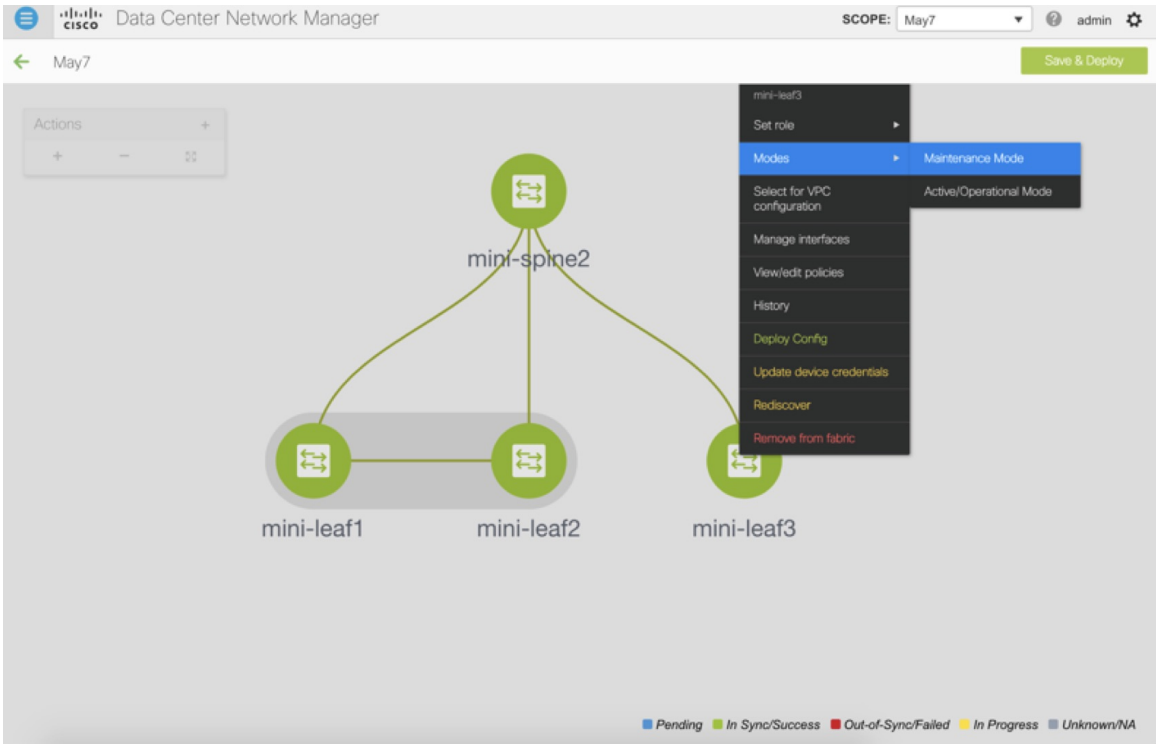

Step 4 Physically replace the device in the network. Physical connections should be made in the same place on the replacement switch as they existed on the original switch.

Step 5 Provision RMA flow and select the replacement device.

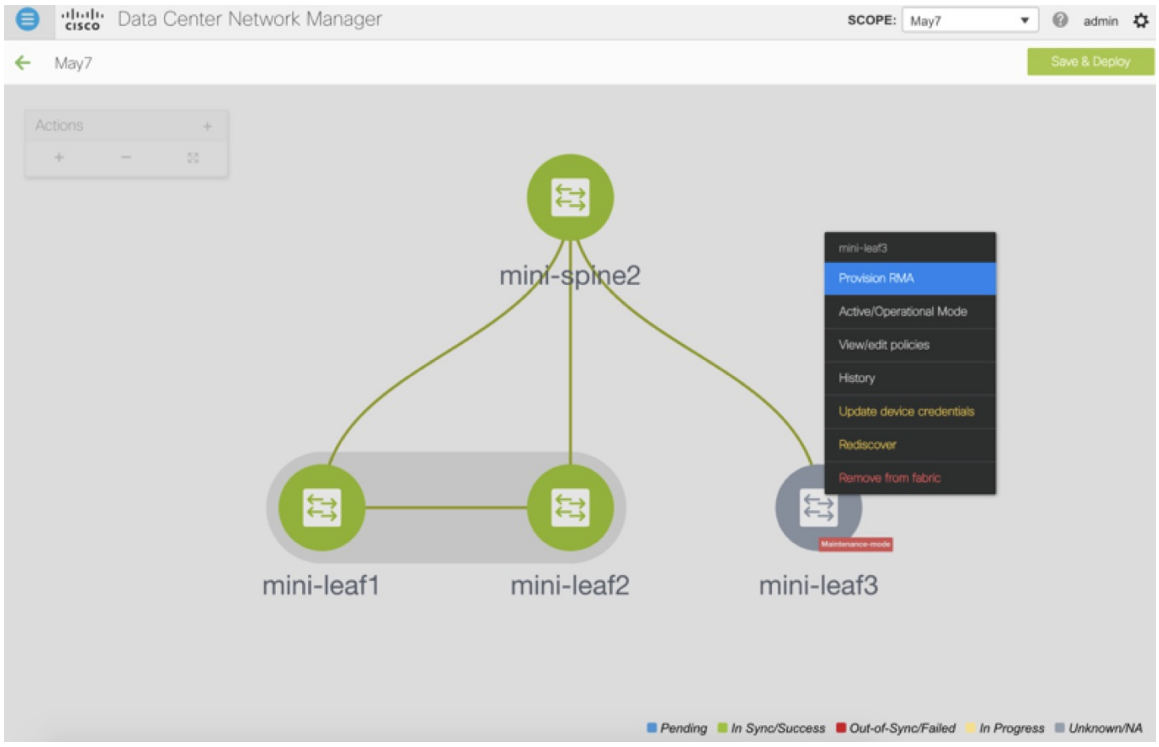

Step 6 The Provision RMA UI will show the replacement device 5-10 minutes after it is powered on.

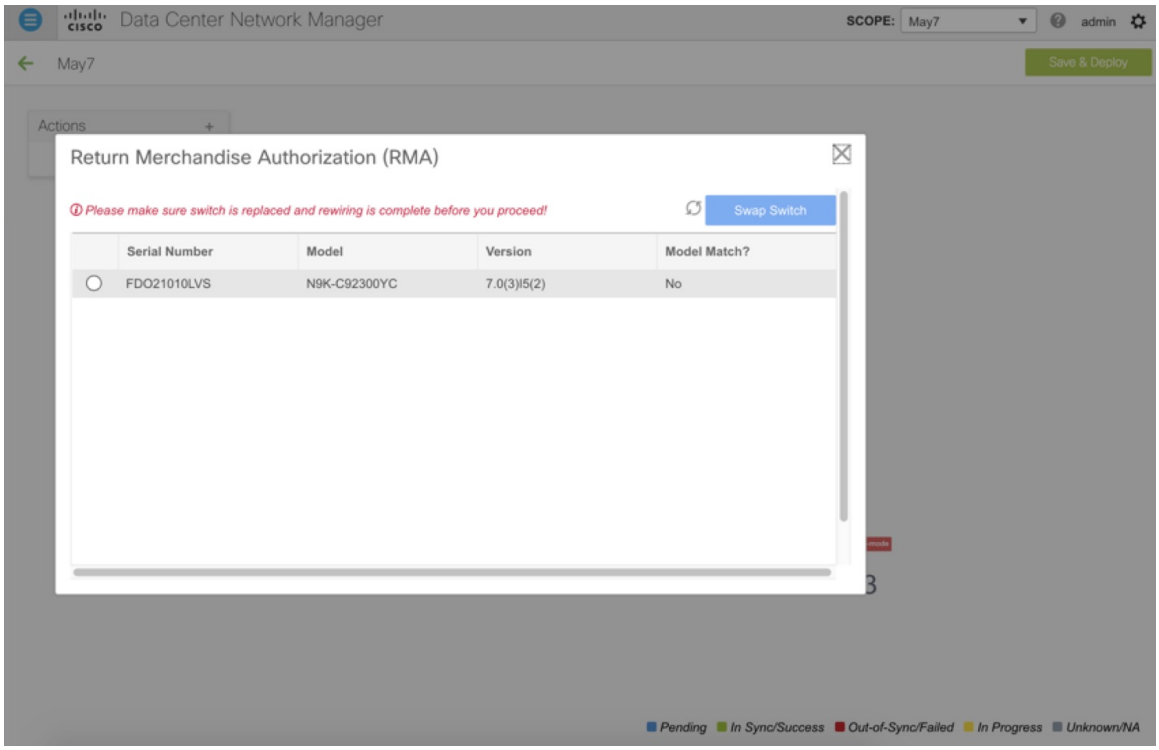

Step 7 Select the correct replacement device and click **Swap Switch**. This begins POAP with the full “expected” configuration for that device. Total POAP time is generally around 10-15 minutes.

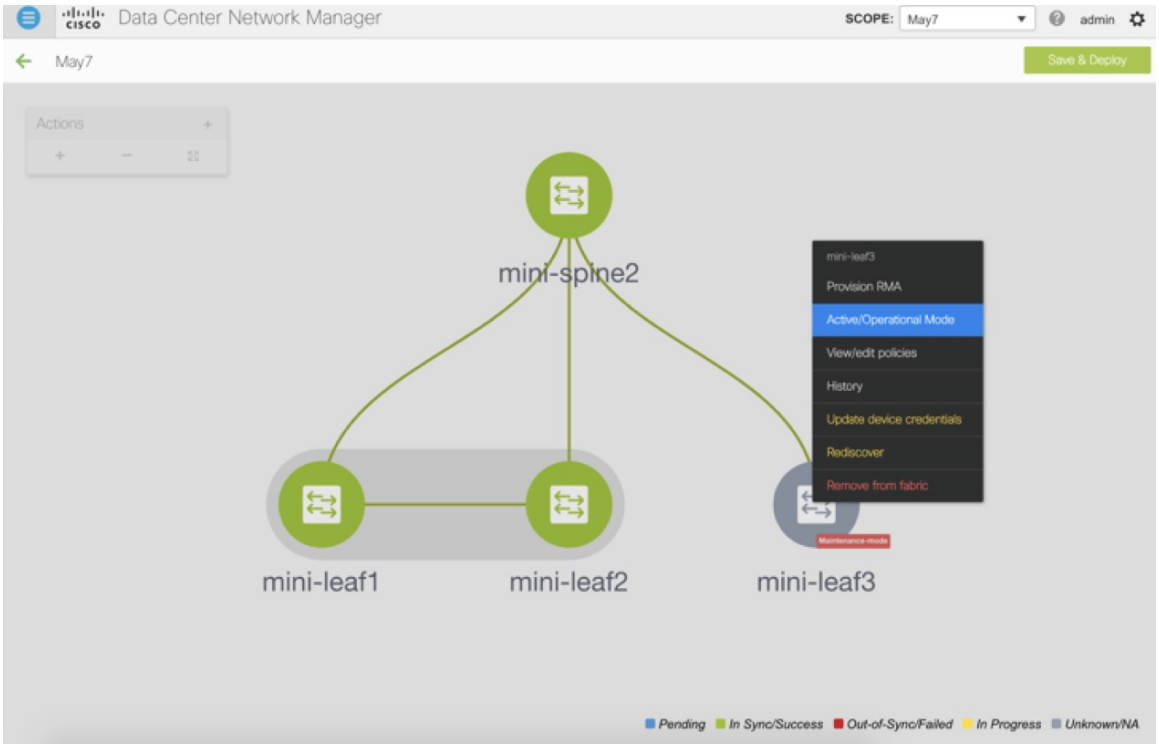

Manual RMA Flow

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).

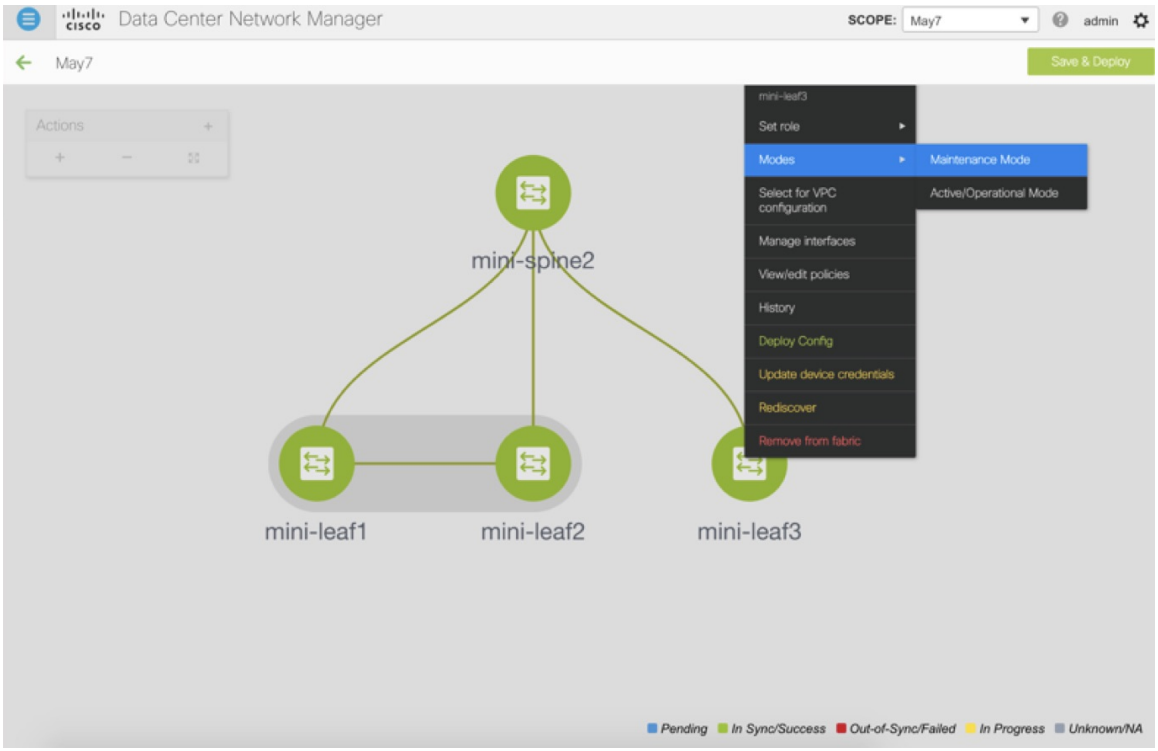

Step 2 Physically replace the device in the network.

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**.

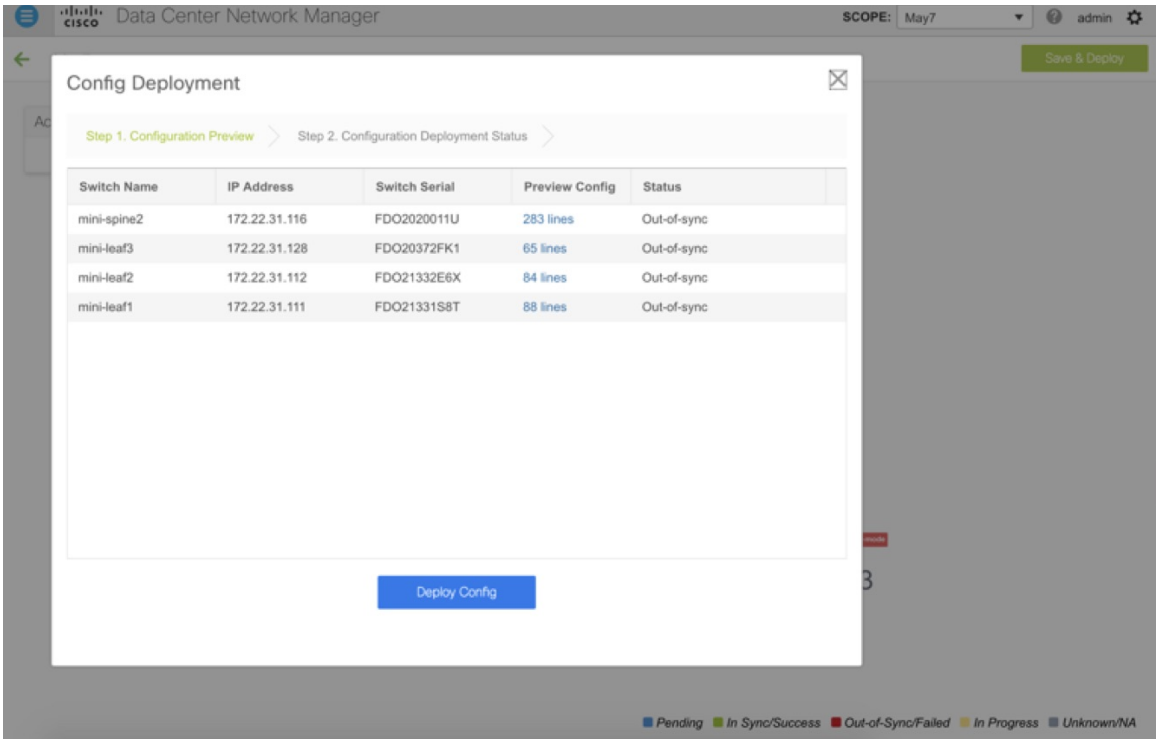

Step 6 Depending on the configuration, if breakout ports or FEX ports are in use, you have to deploy again to completely restore the configuration.

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

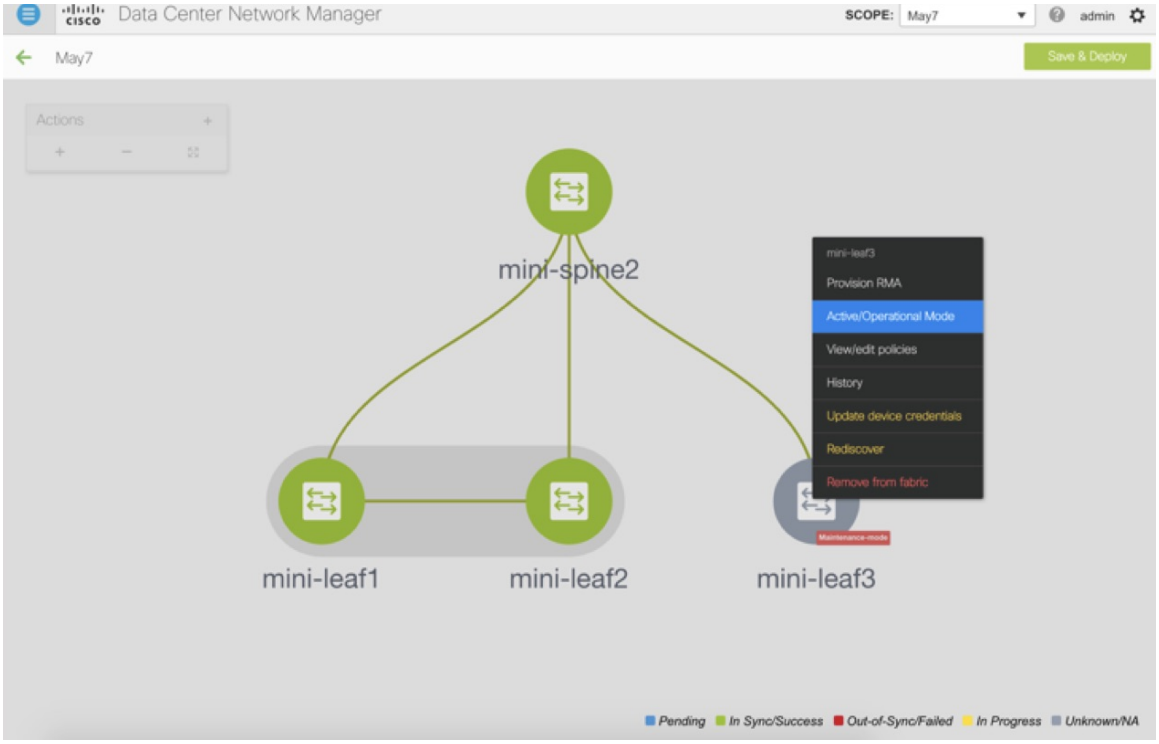

RMA for User with Local Authentication

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.

  Note
  - The following features are unsupported for the brownfield migration of switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images:
    - FEX on switches other than Cisco Nexus 9300 Series switches and Cisco Nexus 9500 Series switches with X9500 line cards
    - AA-FEX

    For information about the platform support for FEX, refer to your platform and NX-OS documentation to check the feature compatibility.
  - To edit interfaces associated with fabric links such as intra-fabric links and inter-fabric links, see Editing Interfaces Associated with Links, on page 146.

- Create tunnel interfaces for Cisco Cloud Services Router 1000v Series (Cisco CSR 1000v Series).
- 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 **Switch** dashboard.
- Click the graph icon in the Name column to view the interface performance chart for the last 24 hours. However, note that performance data for VLAN interfaces that are associated with overlay networks is not displayed in this chart.

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.

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

If you perform admin operations (shutdown/no shutdown) on SVI, which is part of a config profile, successive **Save & Deploy** operations generate **no interface vlan** command.

For SVI with no policy, on performing admin operation, that is, shutdown/no shutdown command pushed from **Interface Manager**, **int_vlan_admin_state** policy is associated with the SVI.

For example, create and deploy the SVI from **switch_freeform**.

```
interface vlan1234
  description test
  no shutdown
  no ip redirects
  no ipv6 redirects
```

If you shutdown the SVI from interface manager, the **int_vlan_admin_state** policy is associated with the SVI.

Pending diff is shown as:

```
interface Vlan1234
  shutdown
  no ip redirects
  no ipv6 redirects
  description test
  no shutdown
```

Remove the **no shutdown** CLI from the free-form config.

If the user has performed admin operation on SVI, device will have interface in running cofig. Therefore, post network detach **interface vlan** will be still present and interface will be discovered. You need to manually delete the interface from **Interface Manager**.

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. If you want to view interfaces for a specific fabric, select the fabric window from the list.

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

The **Add Interface** window appears.

Step 3 In the **Type** drop-down list, choose the type of the interface.

Valid values are Port Channel, virtual Port Channel (vPC), Straight-through (ST) FEX, Active-Active (AA) FEX, Loopback, Subinterface, Tunnel. The respective interface ID field (Port-channel ID, vPC ID, Loopback ID,Subinterface ID, or Tunnel ID) is displayed when you select an interface Type. 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 won't 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 isn't valid.
- To add vPC hosts, you must designate vPC switches in the fabric topology (through the Fabric Builder) and deploy vPC and peer-link configurations using the **Save and Deploy** option. Once the vPC pair configurations are deployed, it appears in the Select a vPC pair drop-down box.

  You can create a vPC using the **int_vpc_trunk_host_11_1** policy.
- 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.

Devices are listed based on the fabric and interface type. External fabric devices aren't listed for ST FEX and AA FEX. In the case of vPC or Active to Active FEX, select the vPC switch pair.

Step 5 Enter the ID value in the respective interface ID field (**Port-channel ID**, **vPC ID**, **Loopback ID** and **Subinterface ID**) that is displayed, based on the selected interface.

You can override this value. The new value is used only if it's 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 shouldn't 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 only modify the policy attribute after the first save. If you try to use an ID that is already used, you encounter the *Resource could not be allocated* error.

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 unbreakout an interface by using the **breakout** option at the top left.

Editing Interfaces

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

Note 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.

During overlay network deployment on switches, the network can be associated with trunk interfaces. The trunk interface to network association is reflected in the **Interfaces** screen. You can update such interfaces.

For interface policies that are not created from the **Control > Interfaces** screen, you can edit some configurations but not change the policy itself. The policy and fields that cannot be edited are grayed out.

The following are some examples of policies that cannot be edited:

- Loopback interface policies - The int_fabric_loopback_11_1 policy is used to create a loopback interface. You can edit the loopback IP address and description but not the int_fabric_loopback_11_1 policy instance.
- Fabric underlay network interface policies (int_fabric_num_11_1, for example) and fabric overlay network interface (NVE) policies.
- Policies associated with port channels and member ports of port channels, including the port channels and member ports associated with a vPC.
- SVIs created during network and VRF creation. The associated VLANs appear in the interfaces list.

Editing Interfaces Associated with Links

There are two types of links, namely intra-fabric links and inter-fabric links. As the name implies, intra-fabric links are set up between devices within the same Easy fabric and are typically used for spine-leaf connectivity. Inter-fabric links are set up between the Easy fabric, and typically other external or Easy fabrics. They are used for external WAN and/or DCI connectivity. A policy is associated with each link that effectively states the configuration that is applied to both ends of the link. In other words, the link policy becomes the parent of the individual child interface policies that are associated with the two interfaces that form the link. In this scenario, you must edit the link policy to edit the interface policy fields such as description, IP address, and any per interface freeform config. The following procedure shows how to edit the interfaces associated with links:

Procedure

Step 1 Choose **Control > Fabric Builder**, and select the fabric containing the link.

Step 2 Click **Tabular view** in the **Actions** panel.

A window with the **Switches** and **Links** tabs appears.

Step 3 Click the **Links** tab.

Step 4 Select the link that you want to edit and click the **Update Link** icon.

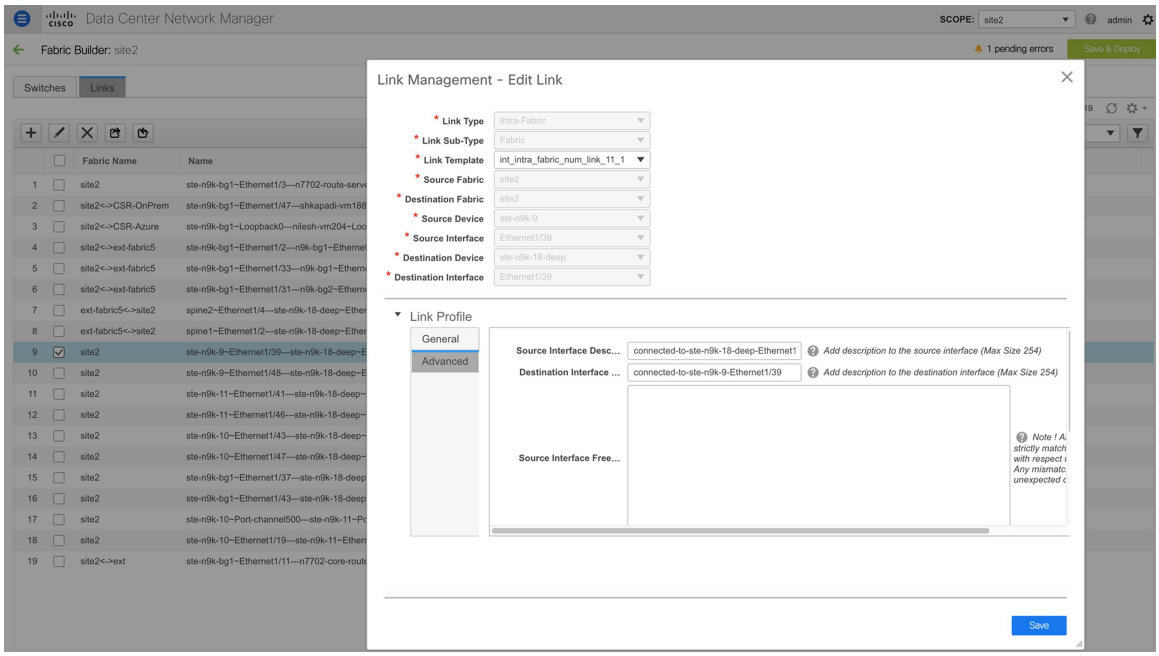

Update the link based on your requirements and click **Save**.

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**.

You cannot delete logical interfaces created in the fabric underlay.

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 for interface or interface sub types like port channel or vPC, 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

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

Step 2 Choose an interface you want to deploy.

Note You can select multiple interfaces and deploy pending configurations.

Step 3 Click **Deploy** to deploy or redeploy configurations that are saved for an interface.

Creating External Fabric Interfaces

You can add and edit port channel, vPC, subinterface, and loopback interfaces for external fabric devices. You cannot add Straight-through FEX and Active-Active FEX functions.

The Breakout port function is only supported for Cisco Nexus 9000 and 3000 series switches in the external fabric.

When you add an interface to an external fabric device, the Resource Manager is not in sync with the device. So, ensure that the value populated in the ID field (Port-channel ID, vPC ID, Loopback ID, etc) is not previously configured on the switch.

If you want to configure a portchannel in the external fabric, you should add and deploy the **feature_lacp** policy on the switches where the portchannel will be configured.

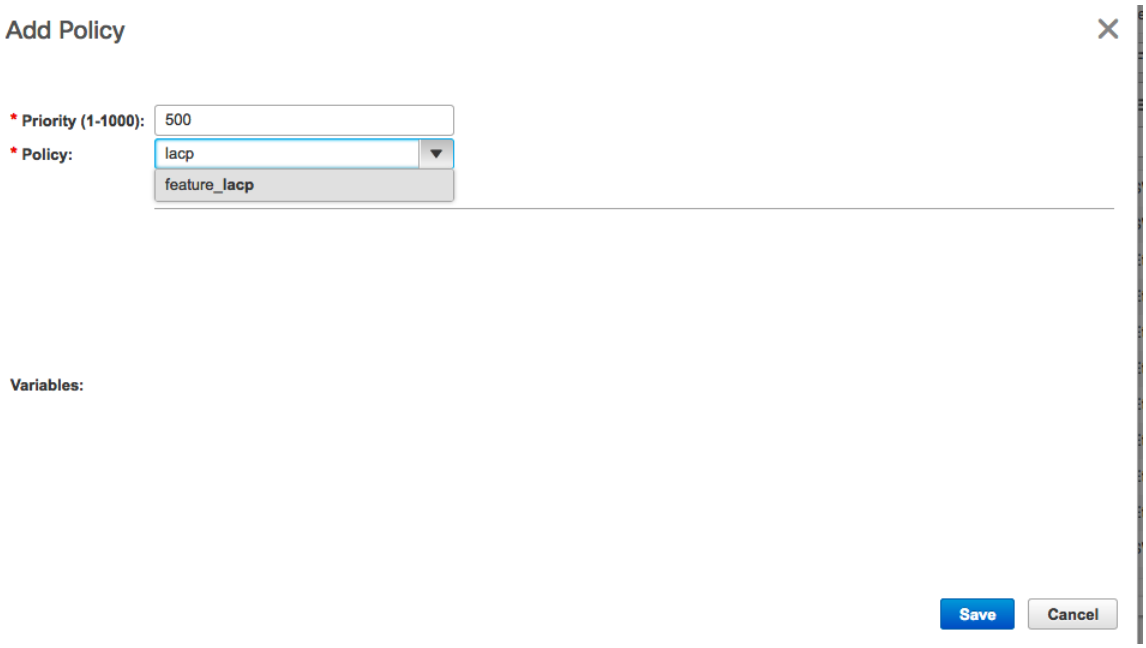

When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on its switches. If you click **Save & Deploy** in the fabric topology screen, it displays an error message. However, the following settings (available when you right-click the switch icon) are allowed:

vPC pairing - You can designate a vPC switch pair, but it is only for reference.

View/edit policy - You can add a policy but you cannot deploy it on the switch.

Manage interfaces – You can only create intent for adding interfaces. If you try to deploy, edit, or delete interfaces, it results in an error message.

Creating and Deploying Networks and VRFs

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.

You can navigate to the networks and VRFs window through any of the following options:

- From the home page: Click the **Networks & VRFs** button in the Cisco DCNM Web UI landing page.
- From the Control menu: From the home page of the Cisco DCNM Web UI, choose **Control > Fabrics > Networks** to navigate to the **Networks** window. Choose **Control > Fabrics > VRFs** to navigate to the **VRFs** window.

You can toggle between the network view and VRF view in both the windows by clicking the **VRF View** or **Network View** button. When you are in the networks or VRFs window, ensure you choose the appropriate fabric from the **Scope** drop-down list before you create any networks or VRFs.

Viewing Networks and VRFs for a Fabric

- Click **Control** > **Networks** from the main menu.

  The **Networks** screen comes up. The SCOPE drop down box (at the top right part of the screen) lists all fabrics managed by the DCNM instance, in alphabetical order. You can choose the correct fabric from SCOPE. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

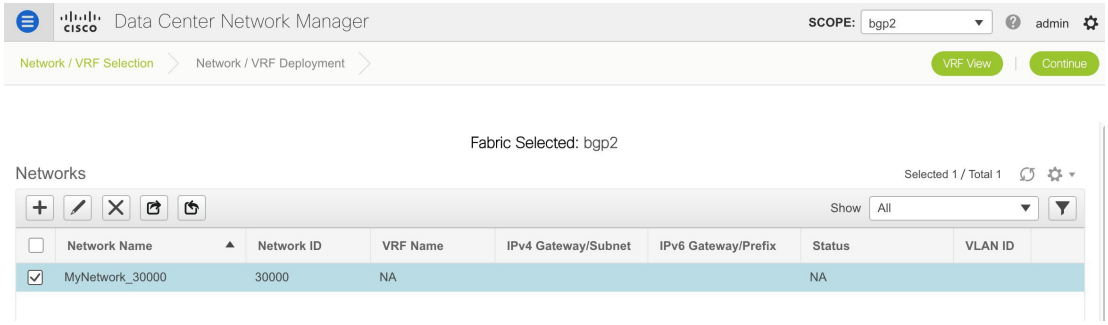

- Click **Control** > **VRFs** from the main menu.

  The VRFs screen comes up. The SCOPE drop down box (at the top right part of the screen) lists all fabrics managed by the DCNM instance, in alphabetical order. You can choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

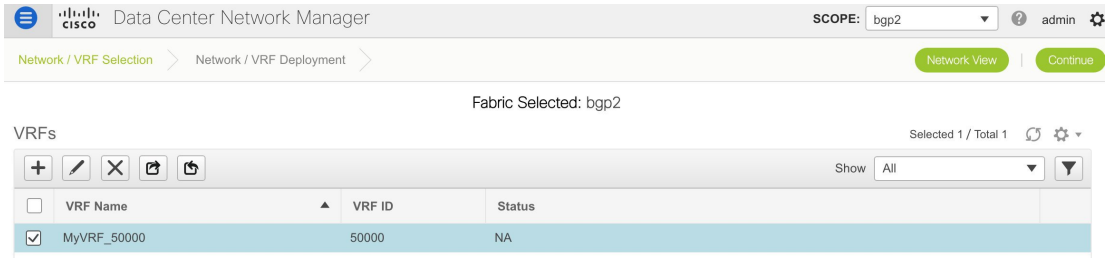

Note The **Networks** or **VRFs** windows are applicable only for the Easy or MSD fabrics.

Creating Networks for the Standalone Fabric

1. Click **Control** > **Networks** (under **Fabrics** submenu).

   The Networks screen comes up.

2. Choose the correct fabric from SCOPE. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

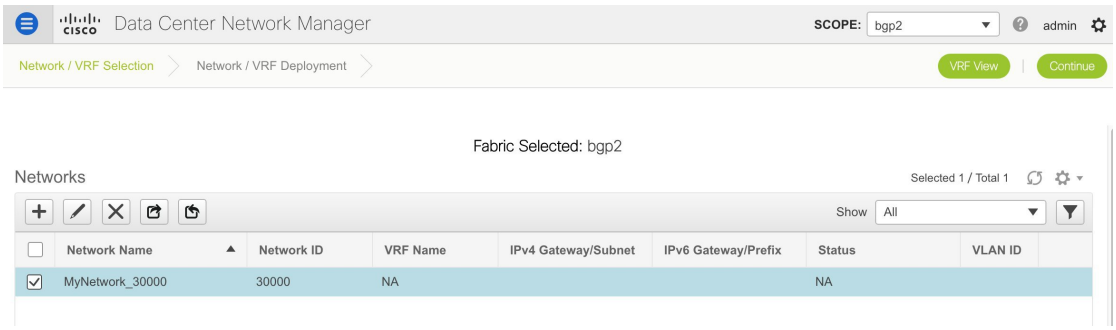

3. Click the + button at the top left part of the screen (under **Networks**) to add networks to the fabric. The Create Network screen comes up. Most of the fields are autopopulated.

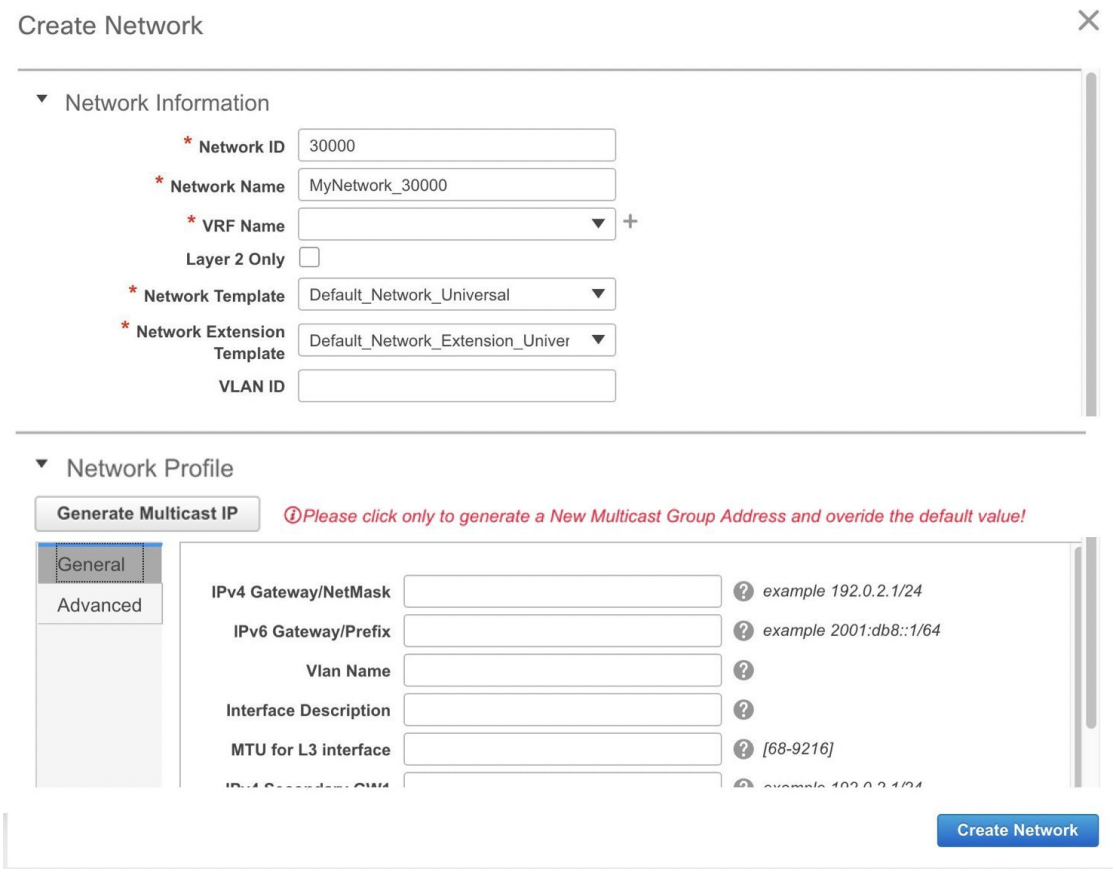

The fields in this screen are:

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: A universal template is autopopulated. This is only applicable for leaf switches.

Network Extension Template: A universal extension template is autopopulated. This allows you to extend this network to another fabric. 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.

Note If the same IP address is configured in the IPv4 Gateway and IPv4 Secondary GW1 or GW2 fields of the network template, DCNM does not show an error, and you will be able to save this configuration.

However, after the network configuration is pushed to the switch, it would result in a failure as the configuration is not allowed by the switch.

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.

VLAN Name - Enter the VLAN name.

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

MTU for the L3 interface - Enter the MTU for Layer 3 interfaces.

IPv4 Secondary GW1 - Enter the gateway IP address for the additional subnet.

IPv4 Secondary GW2 - Enter the gateway IP address for the additional subnet.

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

ARP Suppression – Select the checkbox to enable the ARP Suppression function.

Ingress Replication - The checkbox is selected if the replication mode is Ingress replication.

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

Multicast Group Address- The multicast IP address for the network is autopopulated.

Multicast group address is a per fabric instance variable. The number of underlay multicast groups supported is only 128. If all networks are deployed on all switches, you need not use a different multicast group per L2 VNI or a network. Therefore, multicast group for all networks in a fabric remains same. If a new multicast group address is required, you can generate it by clicking the **Generate Multicast IP** button.

DHCPv4 Server 1 - Enter the DHCP relay IP address of the first DHCP server.

DHCPv4 Server 2 - Enter the DHCP relay IP address of the next DHCP server.

DHCPv4 Server VRF- Enter the DHCP server VRF ID.

Routing Tag – The routing tag is autopopulated. This tag is associated with each gateway IP address prefix.

TRM enable – Select the checkbox to enable TRM.

L2 VNI Route-Target Both Enable - Select the check box to enable automatic importing and exporting of route targets for all L2 virtual networks.

Enable L3 Gateway on Border - Select the checkbox to enable a Layer 3 gateway on the border switches.

A sample of the Create Network screen is given below.

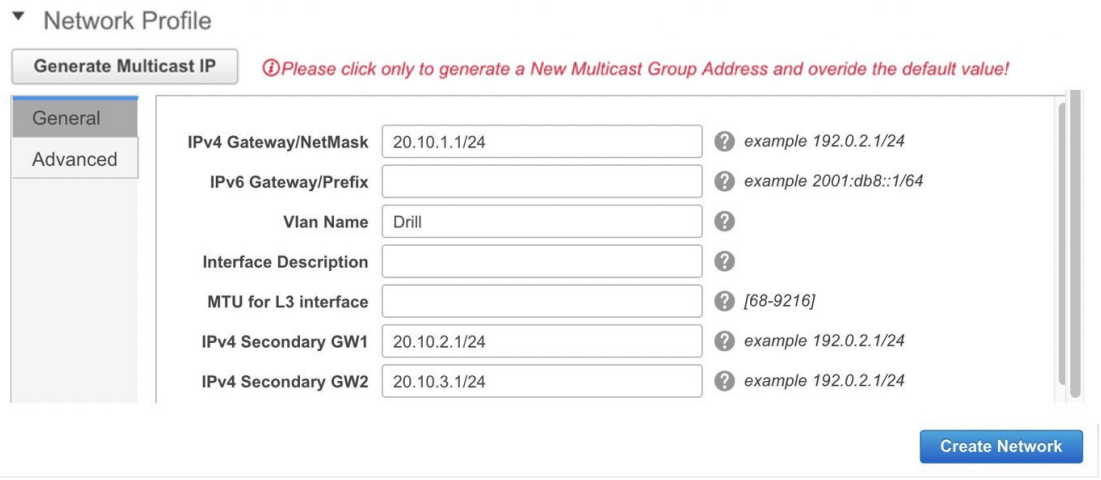

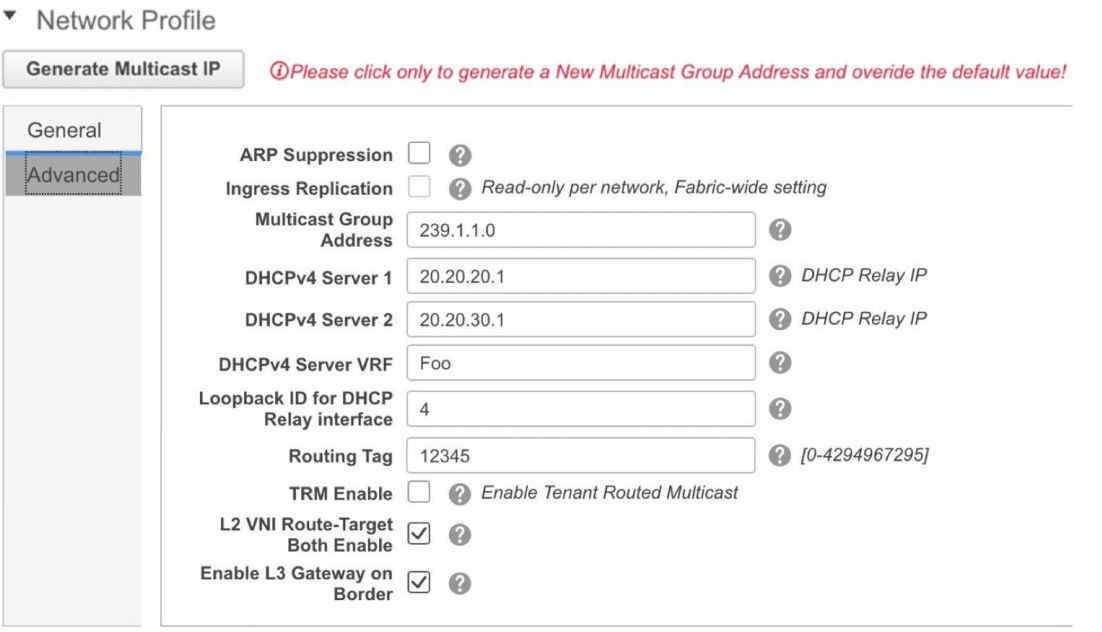

4. Click **Create Network**. A message appears at the bottom right part of the screen indicating that the network is created.

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

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.

Export and Import Network Information

You can export network information to a .CSV file. The exported file contains information pertaining to each network, including the fabric it belongs to, the associated VRF, the network templates used to create the network, and all other configuration details that you saved during network creation.

In the Networks screen, click the Export icon to export network information as a .CSV file.

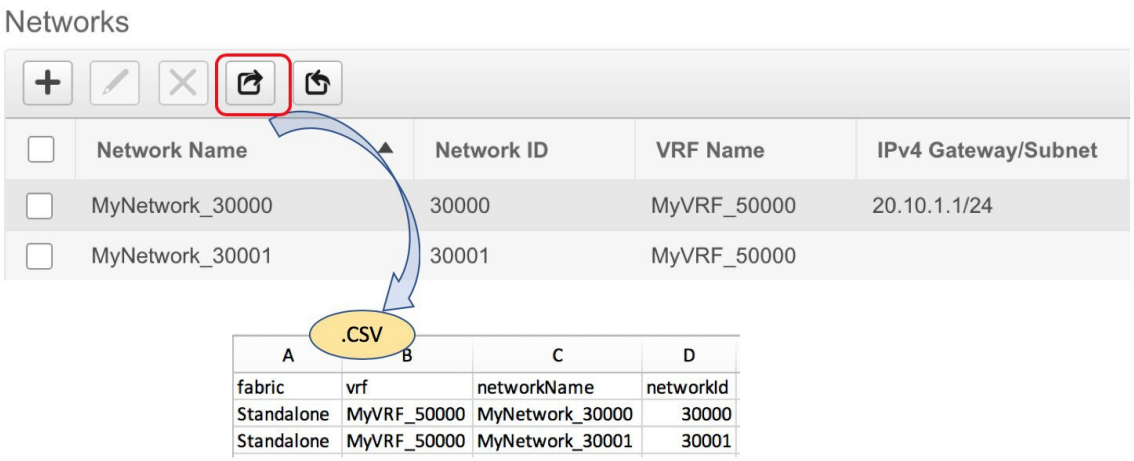

You can use the exported .CSV file for reference or use it as a template for creating new networks. To import networks, do the following:

1. Update new records in the .CSV file. Ensure that the networkTemplateConfig field contains the JSON Object. A message at the bottom right part of the screen displays errors and success messages. This screenshot depicts two new networks being imported.

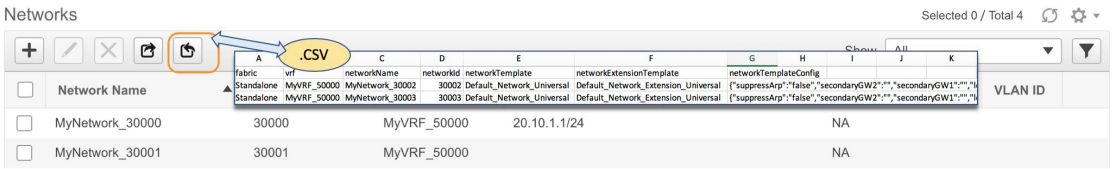

2. In the Networks screen, click the Import icon and import the .CSV file into DCNM.

You can see that the imported networks are displayed in the Networks screen.

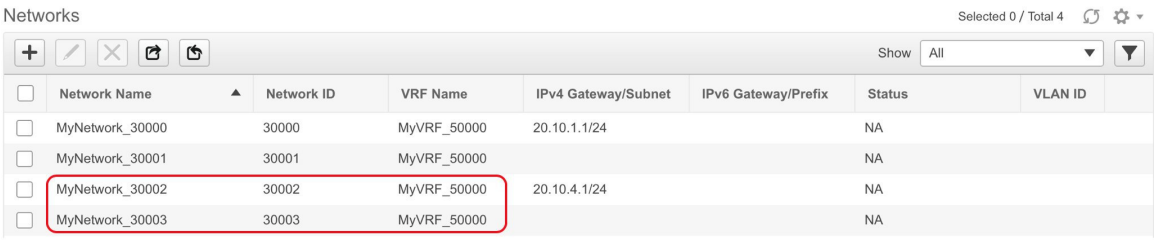

Editing Networks for the Standalone Fabric

To edit networks for standalone fabrics from Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Click **Control** > **Networks**.
The **Networks** window appears.

Step 2 Choose a fabric from the **SCOPE** drop-down list.
The **Networks** window refreshes and lists the networks in the fabric.

Step 3 Choose a network.

Step 4 Click the **Edit** icon.
The **Edit Network** window appears.

Step 5 Update the fields in the **General** and **Advanced** tabs of the **Network Profile** area as needed.

Step 6 Click **Save** at the bottom right part of the window to save the updates.

Creating VRFs for the Standalone Fabric

1. Click **Control** > **VRFs** (under **Fabrics** submenu).

   The VRFs screen comes up.

2. Choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

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

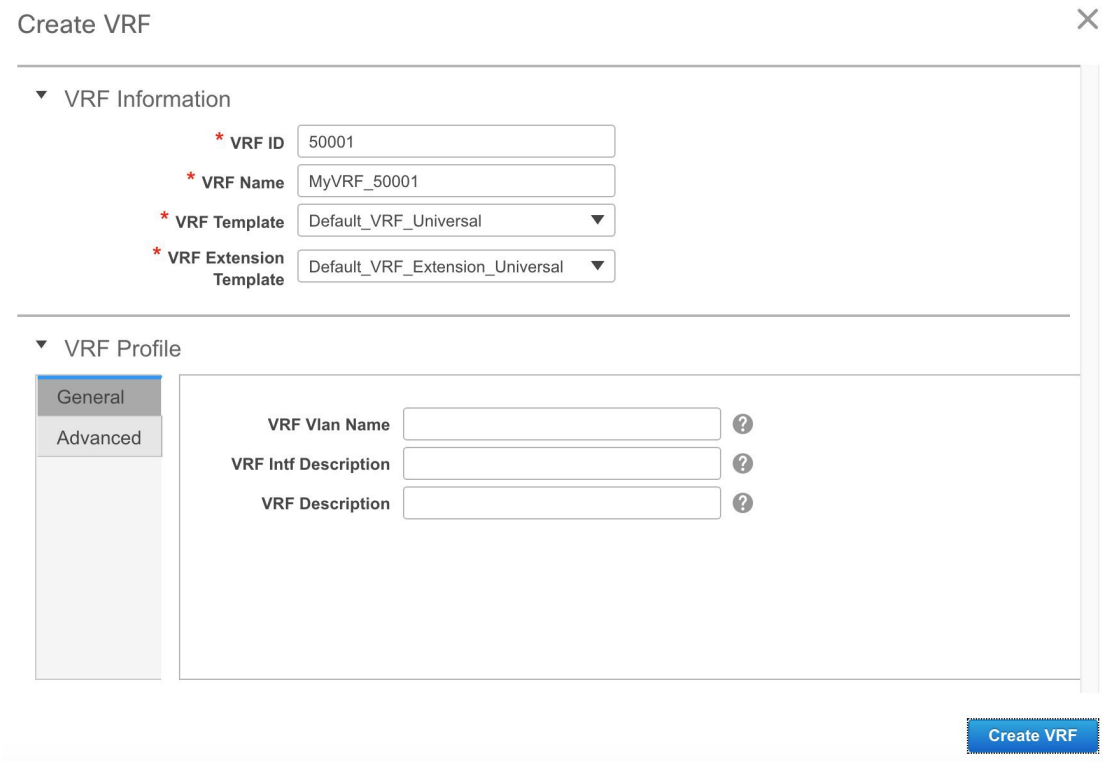

The fields in this screen are:

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 devices.

Fill the fields in the **VRF Profile** section.

General tab – Enter the VLAN ID of the VLAN associated with the VRF, the corresponding Layer 3 virtual interface, and the VRF ID.

Advanced tab – The fields in the tab are autopopulated.

Routing Tag – If a VLAN is associated with multiple subnets, then this tag is associated with the IP prefix of each subnet. Note that this routing tag is associated with overlay network creation too.

Redistribute Direct Route Map – Specifies the route map name for redistribution of routes in the VRF.

Max BGP Paths and **Max iBGP Paths** – Specifies the maximum BGP and iBGP paths.

TRM Enable – Select the checkbox to enable TRM.

If you enable TRM, then the RP address, the RP loopback ID and the underlay multicast address must be entered.

Is RP External – Enable this checkbox if the RP is external to the fabric.

RP Address and **RP Loopback ID** – Specifies the loopback ID and IP address of the RP.

Underlay Multicast Address – Specifies the multicast address associated with the VRF. The multicast address is used for transporting multicast traffic in the fabric underlay.

Note The multicast address in the **Multicast address for TRM** field in the fabric settings screen is populated in this field.

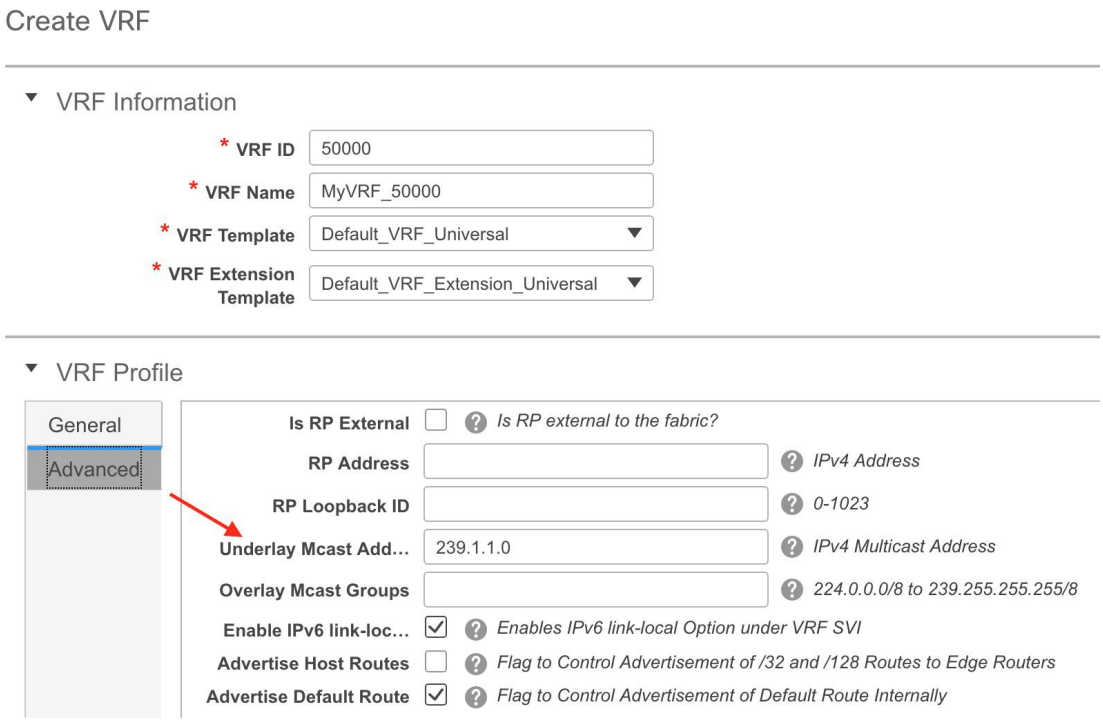

Overlay Multicast Groups – Specifies the multicast address for the VRF, used in the fabric overlay.

Enable IPv6 link-local Option - Select the check box to enable the IPv6 link-local option under the VRF SVI. If this check box is unchecked, IPv6 forward is enabled.

Advertise Host Routes – Enable the checkbox to control advertisement of /32 and /128 routes to Edge Routers.

Advertise Default Route – Enable the checkbox to control advertisement of default routes internally.

To allow inter-subnet communication between end hosts in different VXLAN fabrics, where the subnets are present in both fabrics, you must disable the **Advertise Default Route** feature (clear the **Advertise Default Route** checkbox) for the associated VRF. This will result in /32 routes for hosts being seen in both fabrics. For example, Host1 (VNI 30000, VRF 50001) in Fabric1 can send traffic to Host2 (VNI 30001, VRF 50001) in Fabric2 only if the host route is present in both fabrics. When a subnet is present in only one fabric then default route is sufficient for inter-subnet communication.

Sample screenshots of the Create VRF screen:

Advanced tab:

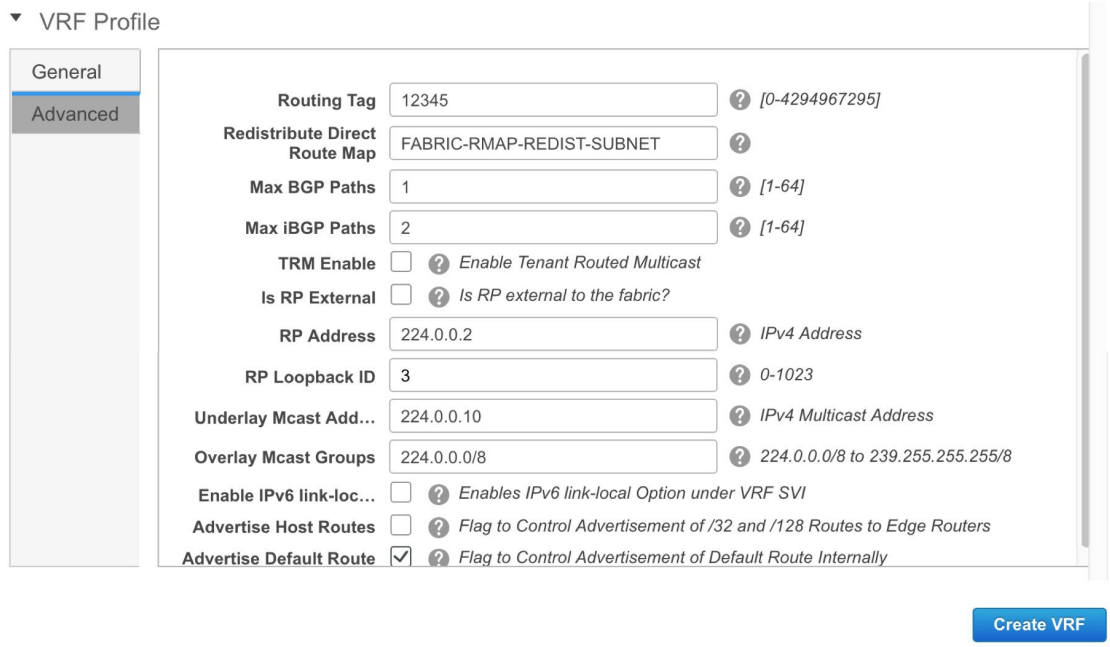

4. Click **Create VRF**.

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

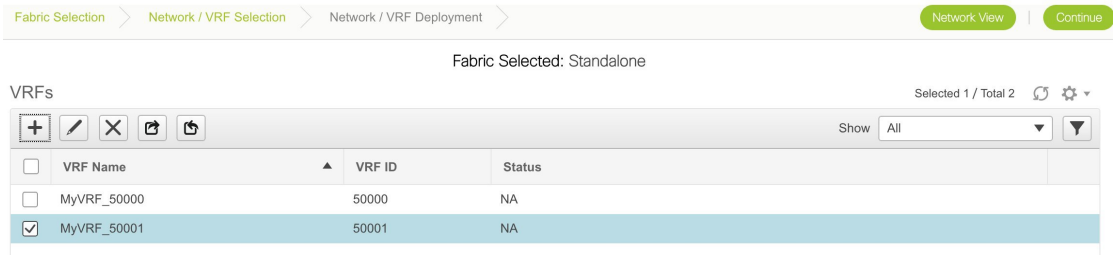

Export and Import VRF Information

You can export VRF information to a .CSV file. The exported file contains information pertaining to each VRF, including the fabric it belongs to, the templates used to create the VRF, and all other configuration details that you saved during VRF creation.

In the VRFs screen, click the Export icon to export VRF information as a .CSV file.

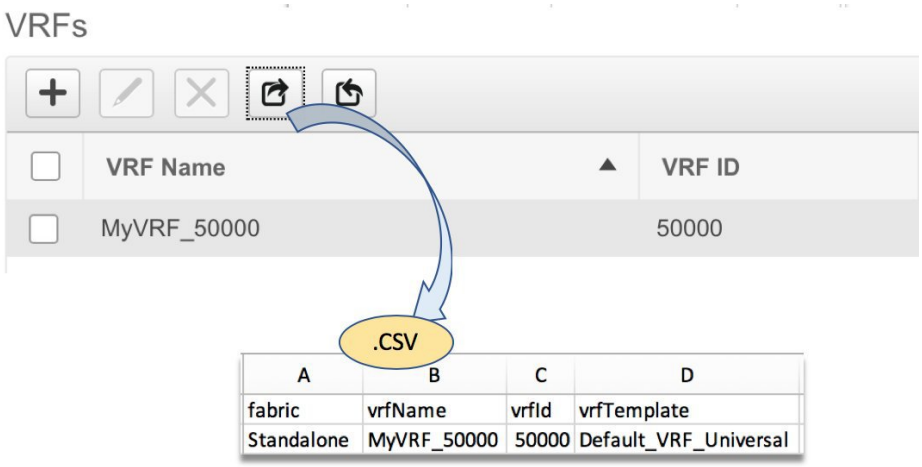

You can use the exported .CSV file for reference or use it as a template for creating new VRFs. To import VRFs, do the following:

1. Update new records in the .CSV file. Ensure that the **vrfTemplateConfig** field contains the JSON Object.
2. In the VRFs screen, click **Import** icon and import the .CSV file into DCNM.

   A message at the bottom right part of the screen displays errors and success messages. This screenshot depicts a new VRF being imported.

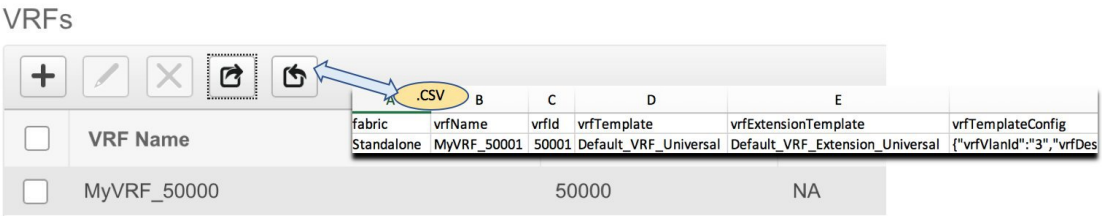

You can see that the imported VRF is displayed in the VRFs screen.

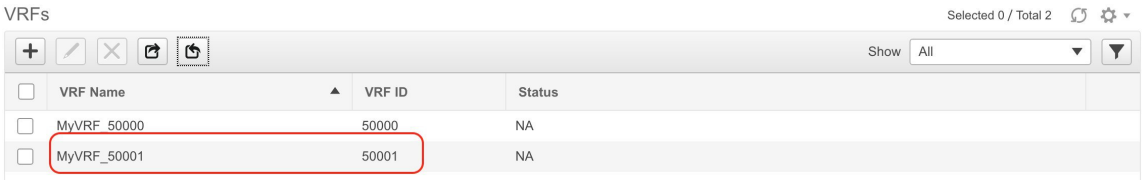

Editing VRFs for the Standalone Fabric

1. Choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

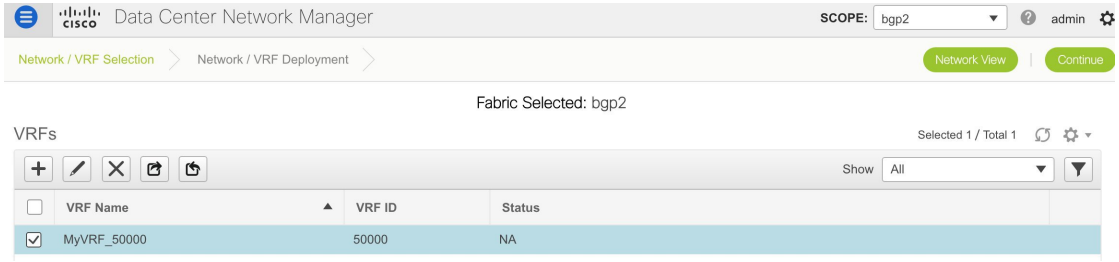

2. From the **Select a Fabric** drop-down list, select the fabric *Standalone*, and click **Continue** on the top right part of the screen. The Networks page is displayed.
3. Click the **VRF View** at the top right part of the screen. The VRFs page appears.

4. Select the **VRF** and click the **Edit** option at the top left part of the screen. The **Edit VRF** screen comes up.

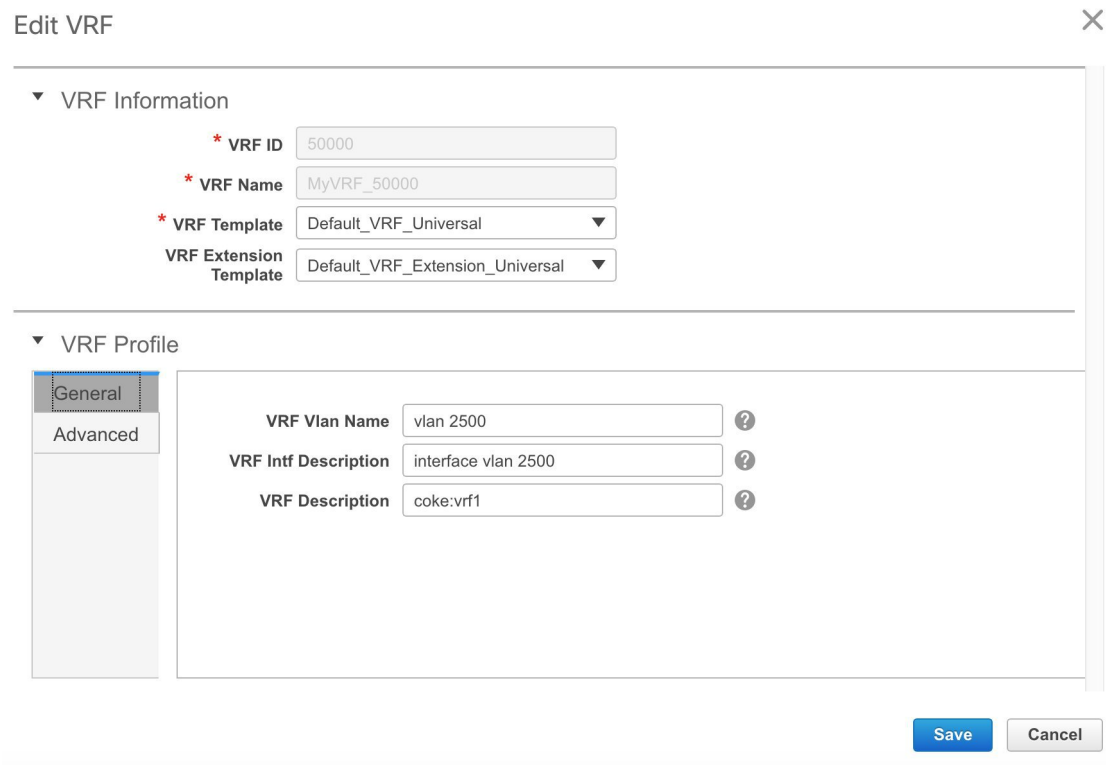

5. Update the fields in the **General** and **Advanced** tabs of the **VRF Profile** section as needed.
6. Click **Save** at the bottom right part of the screen to save the updates.

Deploying Networks for the Standalone and MSD Fabrics

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

1. Click **Control** > **Networks** (under **Fabrics** submenu).

   The Networks screen comes up.

2. Choose the correct fabric from SCOPE. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

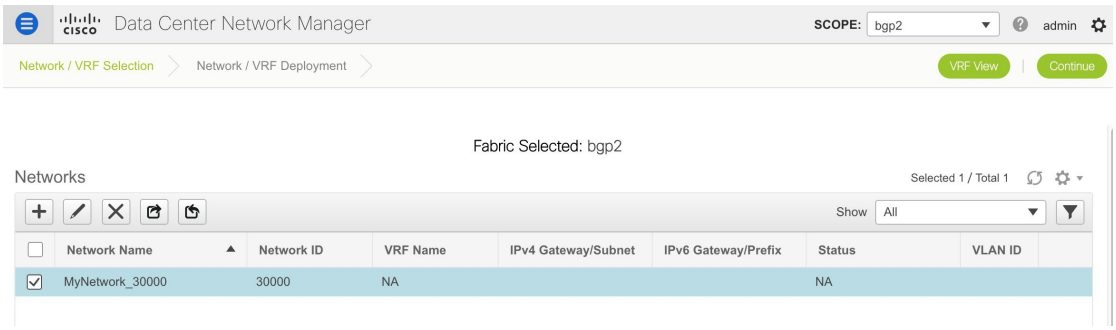

3. Select networks that you want to deploy. In this case, select the check boxes next to both the networks and click **Continue** at the top right part of the screen.

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).

Note In an MSD fabric, all member fabrics are visible from this screen.

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.

Undiscovered cloud display – To display (or not display) an **Undiscovered** cloud in this screen, click the cloud icon in the vertical panel, at the top-right part of the screen. When you click the icon, the **Undiscovered** cloud and its links to the fabric topology are not displayed. Click the icon again to display the **Undiscovered** cloud.

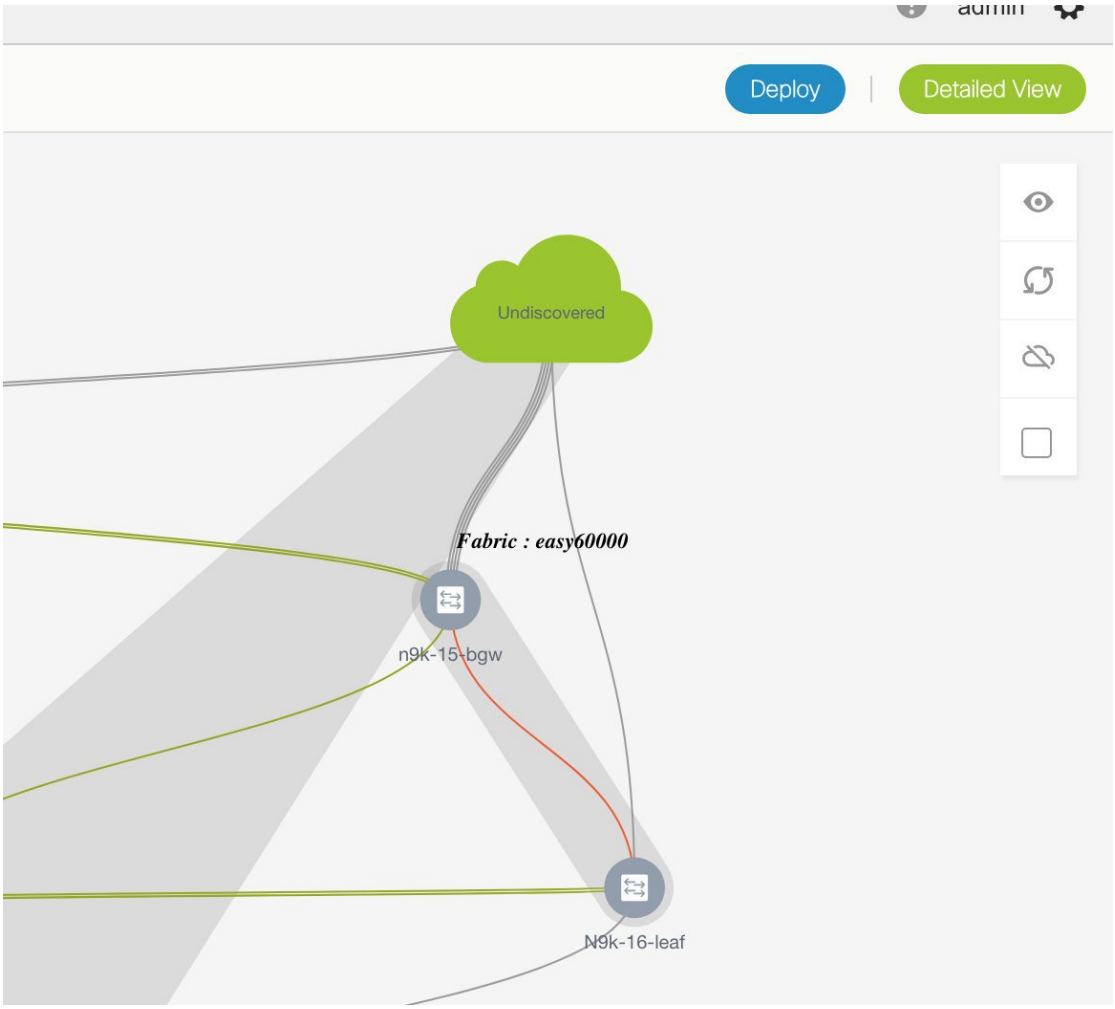

You can move the topology around the screen by clicking the left mouse button on the screen and moving it in the direction you desire. You can enlarge or shrink the switch icons proportionately by moving the cursor roller. You can also use corresponding alternatives on the touchpad.

4. Click **...** in the **Interfaces** column.

   The **Interfaces** box opens up. It lists interfaces or port channels. You can select interfaces/port channels to associate them with the selected network. For each interface, port type and description, channel number and connected neighbor interface details are displayed.

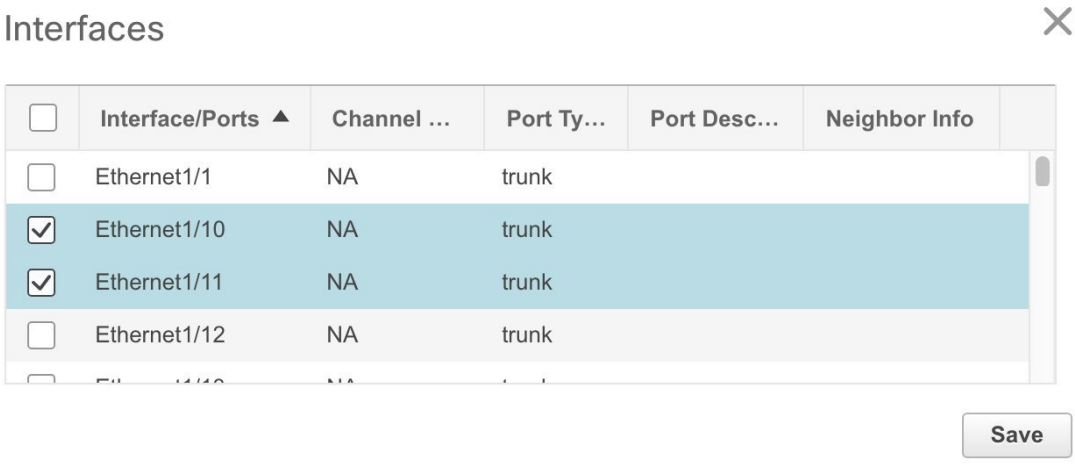

5. Double-click a switch to deploy the networks on it. For deployment of networks on multiple switches, click Multi-Select from the panel at the top right part of the screen (the topology freezes to a static state), and drag the cursor across the switches.

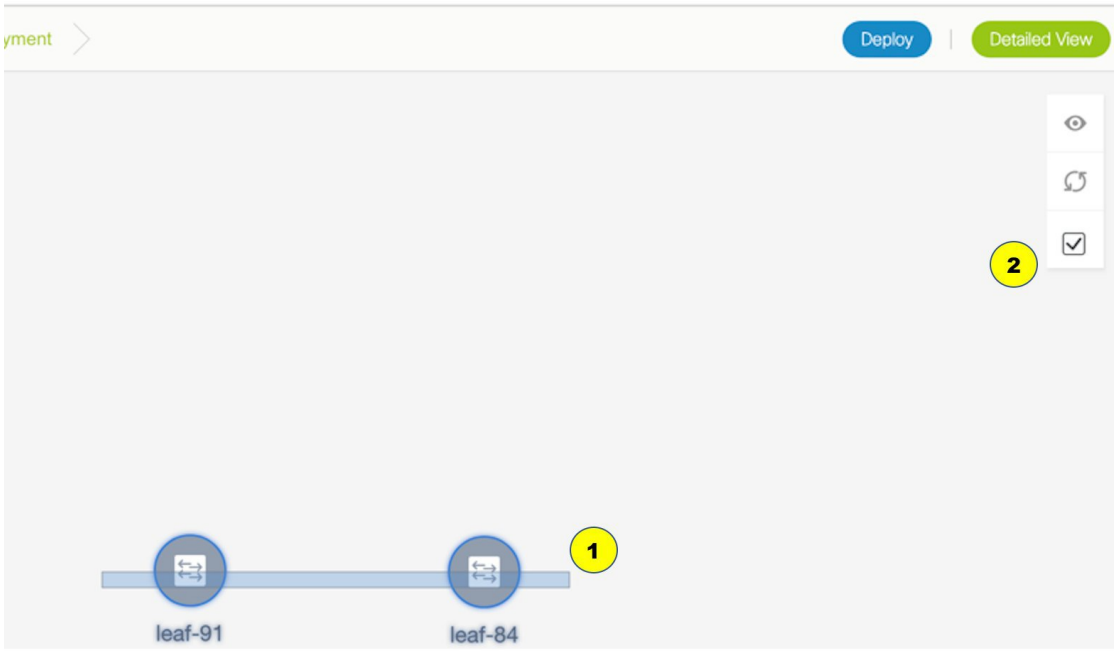

Immediately the Network Attachment dialog box appears.

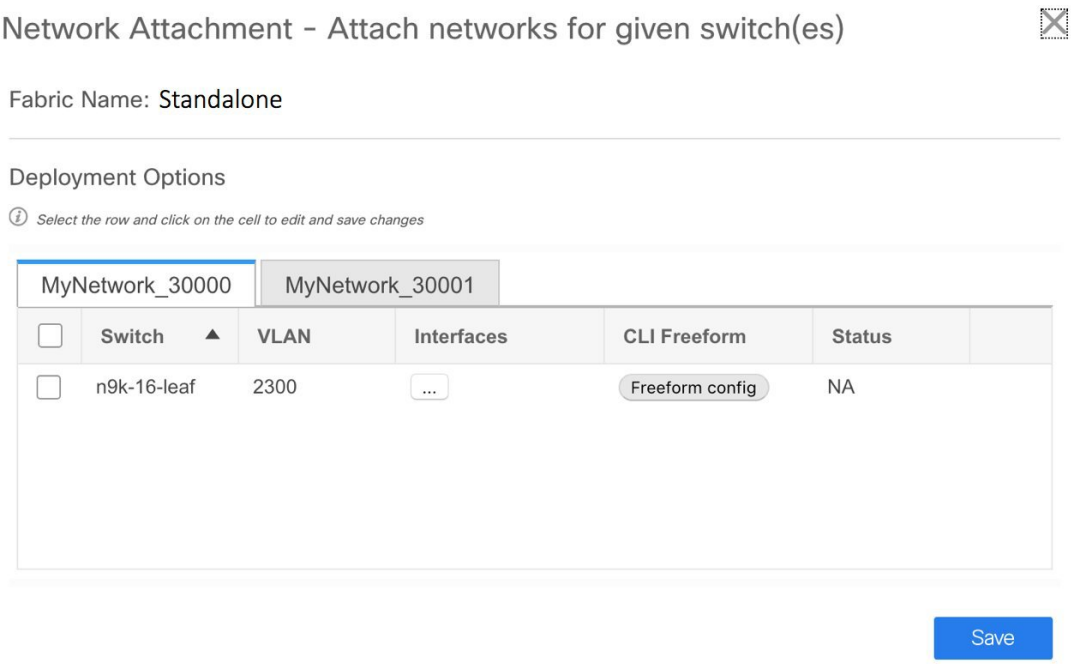

A tab represents each network (the first network is displayed by default) that is being deployed. In each network tab, the switches are displayed. Each row represents a switch.

Click the check box next to the **Switch** column to select all switches. The network is ready to be provisioned on the switches.

VLAN - Update the VLAN ID if needed.

When you update a VLAN ID and complete the network deployment process, the old VLAN is not automatically removed. To complete the process, you should go to the fabric topology screen (click **Control > Fabric Builder** and click within the corresponding fabric box to go to the screen) and use the Save and Deploy option.

When updating the VLAN ID for a given network, the original VLAN ID is not automatically removed from the attached trunk interface. In order to remove the old or original VLAN ID, you must perform **Save and Deploy** + **Config Deploy** operation from within the fabric in Fabric Builder. For this, go to the fabric topology screen (click **Control > Fabric Builder** and click within the corresponding fabric box to go to the screen) and execute the **Save and Deploy** operation. Verify that config compliance is removing the expected config, then execute **Deploy Config** operation to remove the configs.

Interfaces – Click … in the column to add interfaces associated with the selected network.

VLAN to trunk port mapping – The selected trunk ports include the VLAN as an allowed VLAN on the port.

VLAN to vPC domain mapping - If you want to associate the VLAN to port channels of a vPC domain, add the port channels from the list of interfaces. The vPC port channels include the VLAN as an allowed VLAN.

Freeform configurations – Click Freeform config to enable additional configurations on the switch. After the configurations are saved, the Freeform config button gets highlighted.

6. Select the other network tab and make the same selections.

7. 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 window appears 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.

8. 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.

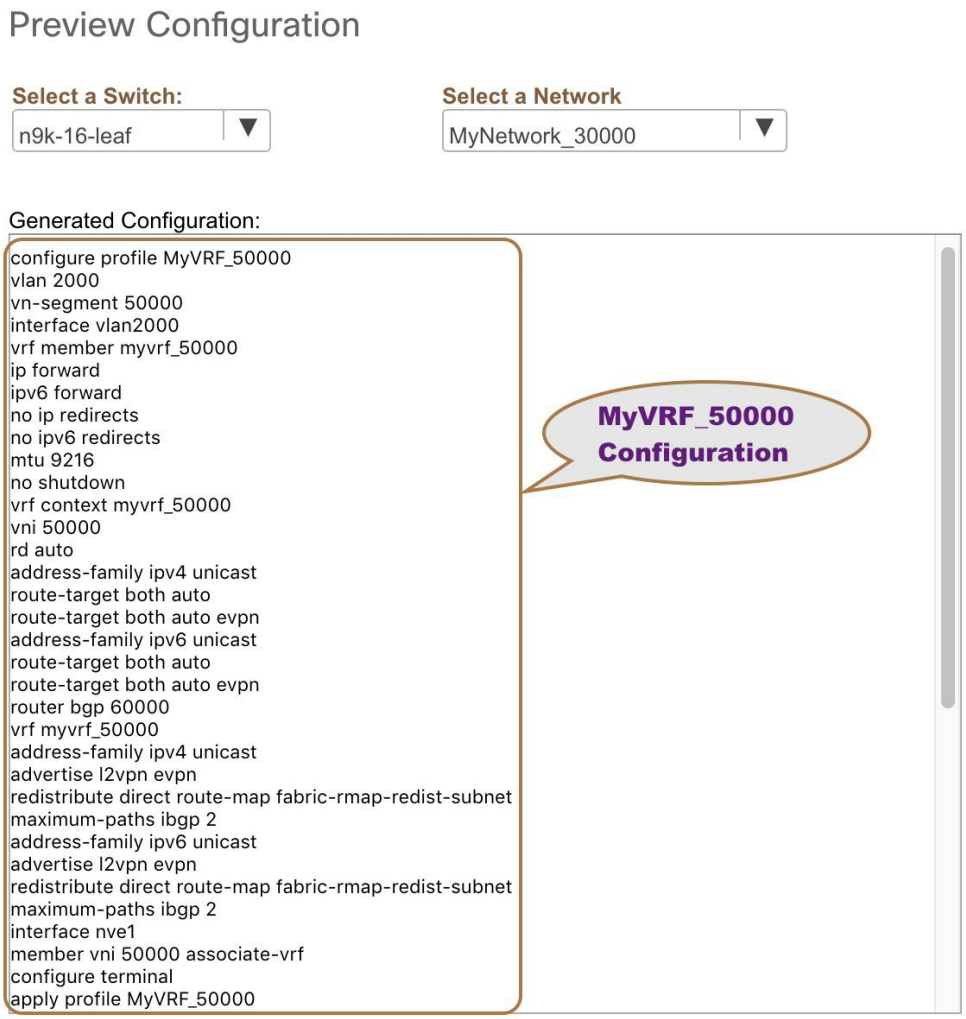

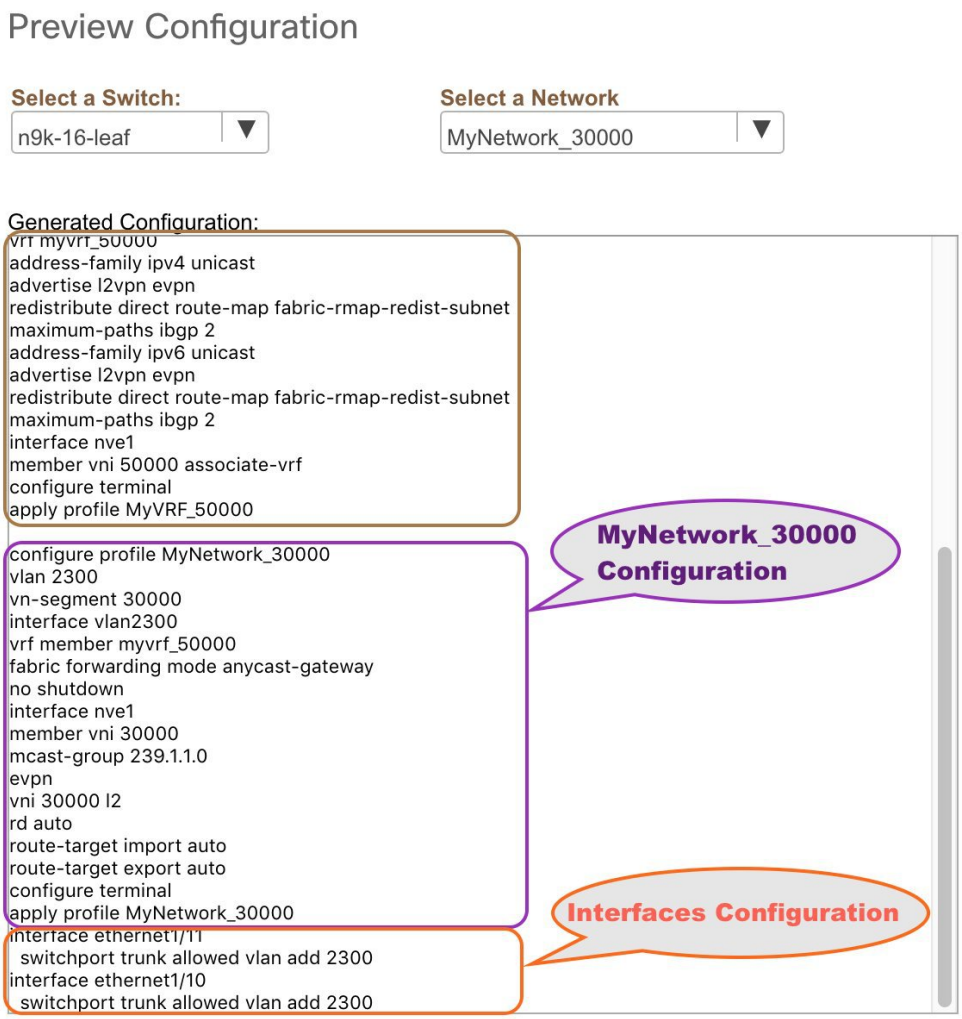

On the preview screen, you can select from the **Select a switch** and **Select a network** drop-down boxes at the top of the screen to view other network configurations.

After checking the configurations, close the screen. The Topology screen appears again.

9. 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 In case you click **Deploy** and there is no configuration diff that has to be deployed, a pop-up window comes up stating **No switches PENDING for deployment**.

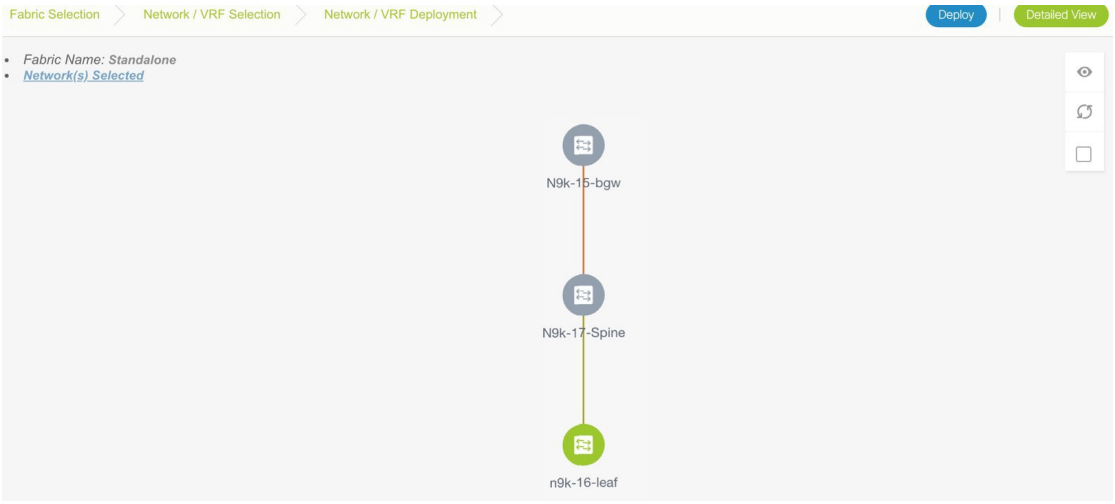

Note When you select multiple networks on the *Topology View* screen and proceed to the deployment screen, the switch color reflects the status of the first network in the selected list of networks. In this example, the switch color turns green when *MyNetwork_30000* is provisioned on the switch.

Go to the Networks page to view the individual status for all networks.

Network Deployment for an MSD Fabric

Consider a scenario wherein you are deploying the same networks on different member fabric border devices. You can choose one fabric, deploy networks on its border devices, and then choose the second fabric and deploy networks.

Alternatively, you can choose the MSD fabric, and deploy the networks from a single topology view of all member fabric border devices.

This is a topology view of an MSD fabric wherein the two member fabrics topologies and their connections are depicted. You can deploy networks on the BGWs of the fabrics at once.

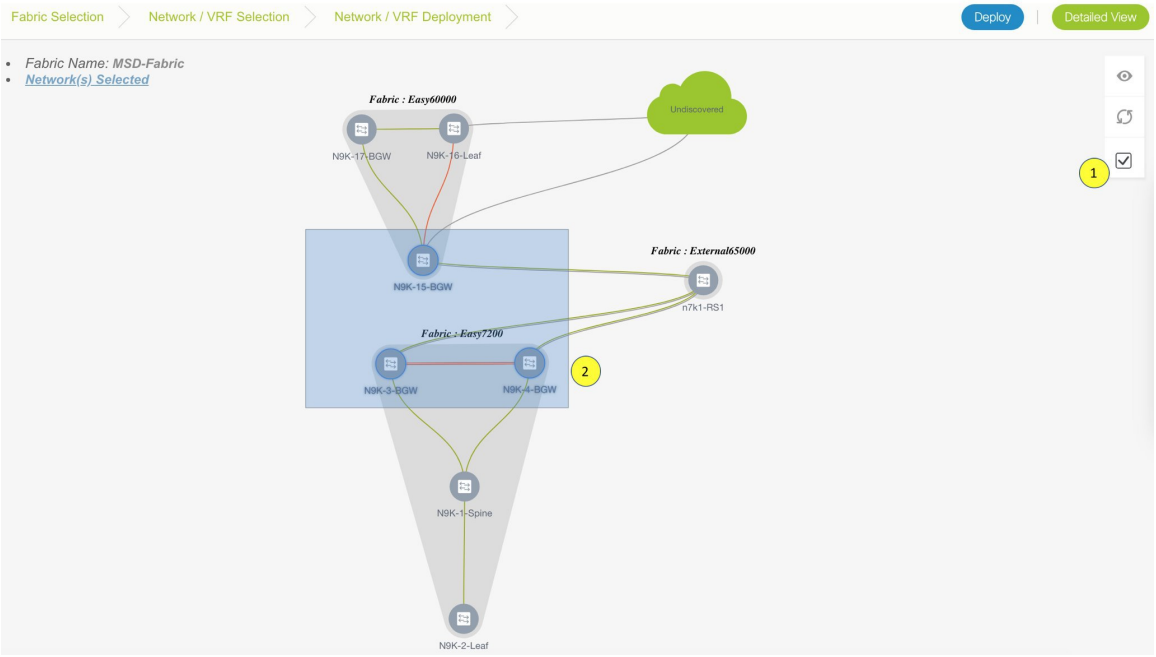

Detailed View

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 window appears. This lists the networks in a tabular view.

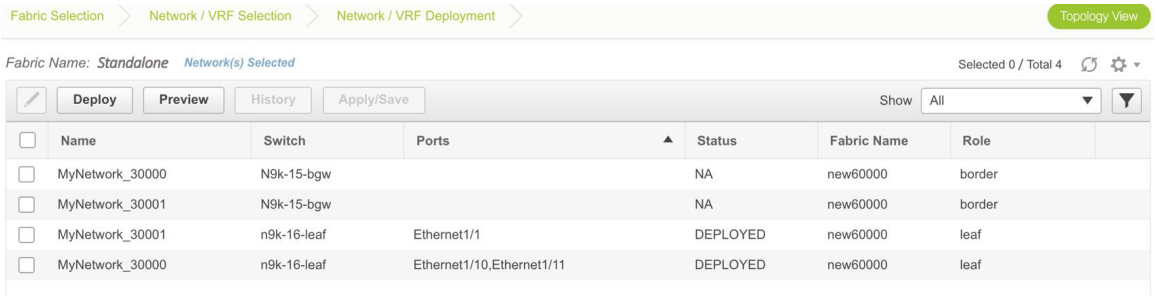

The options:

Edit - Select a network and click the Edit icon at the top left part of the screen.

Note If you select one network/switch entry and click on Edit, the Network Attach dialog box appears. To maintain consistency across the Topology View and Detailed View screens, the Network Attach screen displays all networks, and not just the selected network/switch.

Preview – Click Preview to preview configurations before deployment. You can only preview pending configurations, and not uninitiated or deployed configurations.

Deploy – Click Deploy to provision networks onto the switches.

History – Select a row and click History to view the configuration instances and status. Network and VRF-wise configurations are displayed. Click in the Status column of any instance for more details.

The fields in the table contain the configuration instance in each row, the associated switch and fabric names, the switch role, trunk ports (if any), and the deployment status.

Apply/Save – Selecting a network and clicking Apply/Save will select a switch for the network to be deployed on.

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.

Deploying VRFs for the Standalone and MSD Fabrics

1. Choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

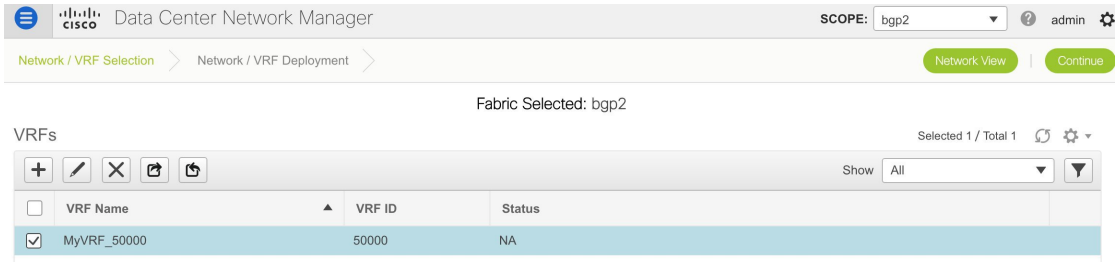

2. Select check boxes next to the VRFs that you want to deploy and click Continue at the top right part of the screen.

   The VRF Deployment screen appears. On this page, you can see the topology of the Standalone fabric. The following example shows you how to deploy the VRFs MyVRF_50000 and MyVRF_50001 on the leaf switch. You can deploy VRFs simultaneously on multiple switches but of the same role (Leaf, Border Gateway, and so on).

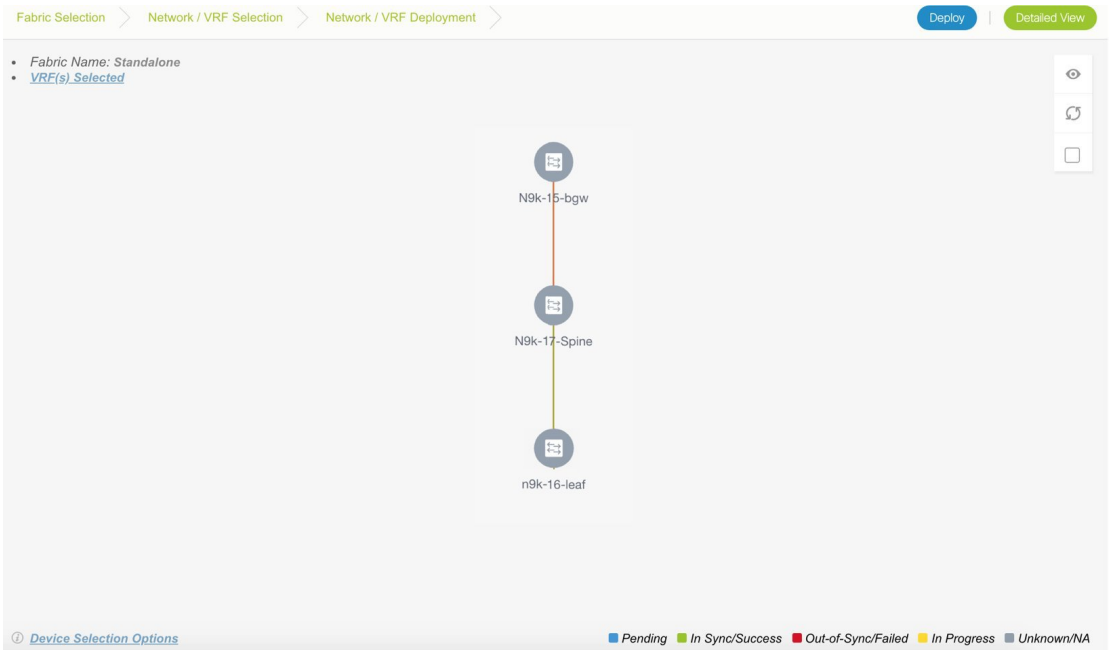

At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly. Blue for *Pending* state, yellow for *In Progress* state when the provisioning is in progress, red for failure state, green when successfully 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 VRFs are yet to be deployed on any switch in this fabric.

Undiscovered cloud display – To display (or not display) an **Undiscovered** cloud in this screen, click the cloud icon in the vertical panel, at the top-right part of the screen. When you click the icon, the **Undiscovered** cloud and its links to the fabric topology are not displayed. Click the icon again to display the **Undiscovered** cloud.

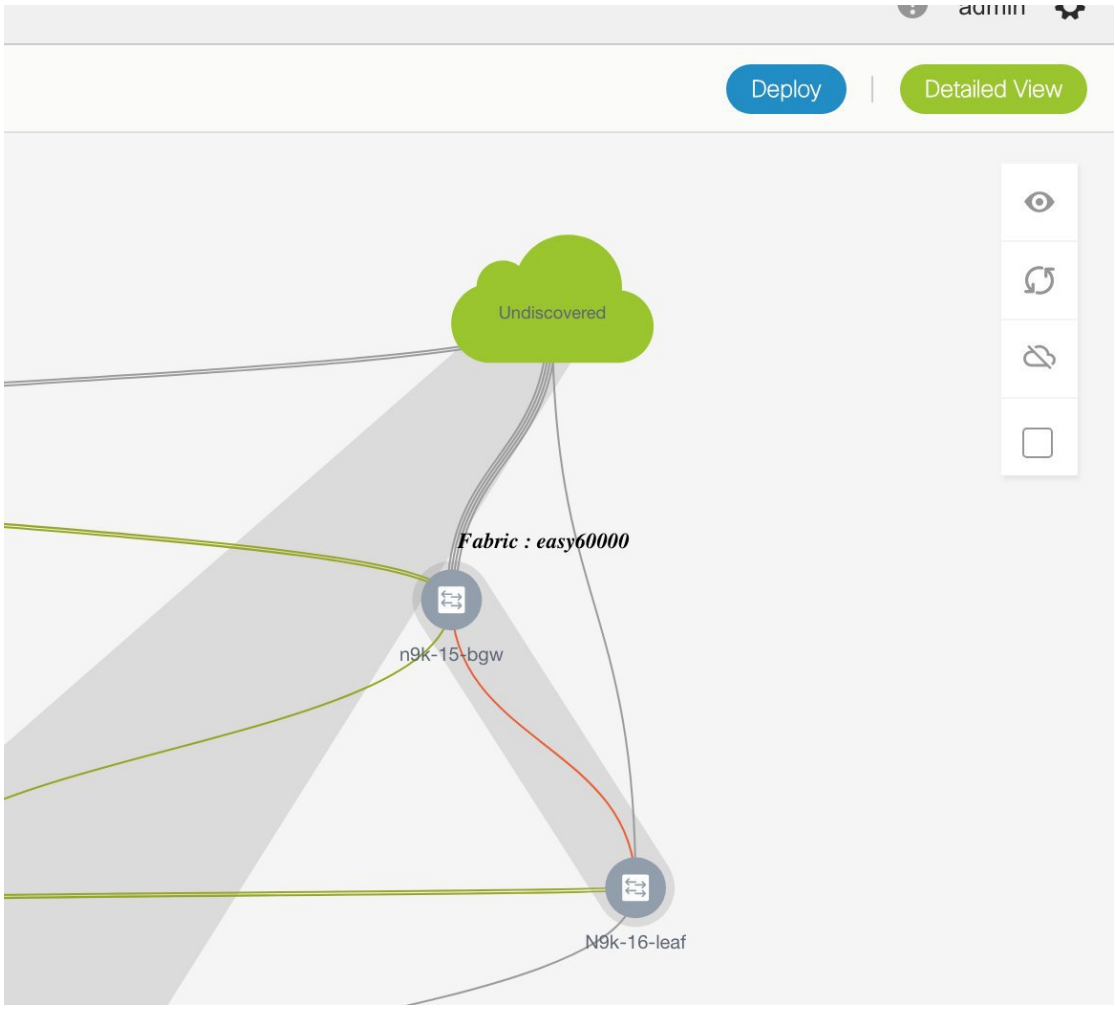

You can move the topology around the screen by clicking the left mouse button on the screen and moving it in the direction you desire. You can enlarge or shrink the switch icons proportionately by moving the cursor roller. You can also use corresponding alternatives on the touchpad.

3. Double-click a switch to deploy VRFs on it. The VRF Attachment screen comes up.

Note For deployment of VRFs on multiple switches, click the Multi-Select option from the panel at the top right part of the screen (This freezes the topology to a static state), and drag the cursor across the switches.

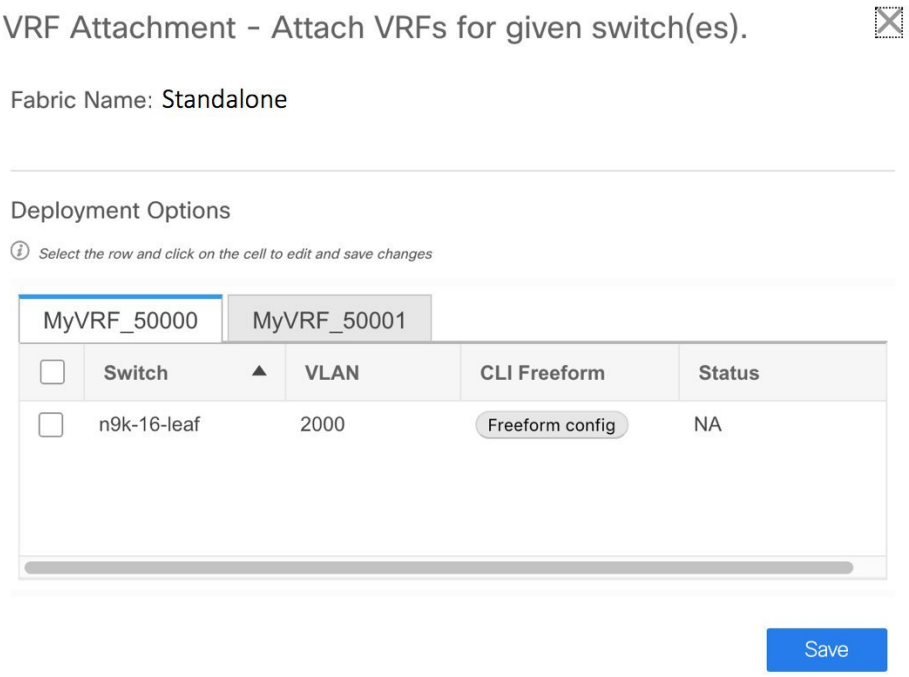

A tab represents each VRF that is being deployed (the first selected VRF is displayed by default). In each VRF tab, the selected switches are displayed. Each row represents a switch.

VLAN ID - Click within the VLAN column to update the VRF VLAN ID, if needed.

Freeform configurations – Click Freeform config to enable additional configurations on the switch. After you save freeform configurations, the Freeform config button gets highlighted.

Click the checkbox next to the Switch column to select all switches. VRF MyVRF_50000 is ready to be provisioned on the switch

4. Select the other VRF tab and make the same selections.
5. 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).

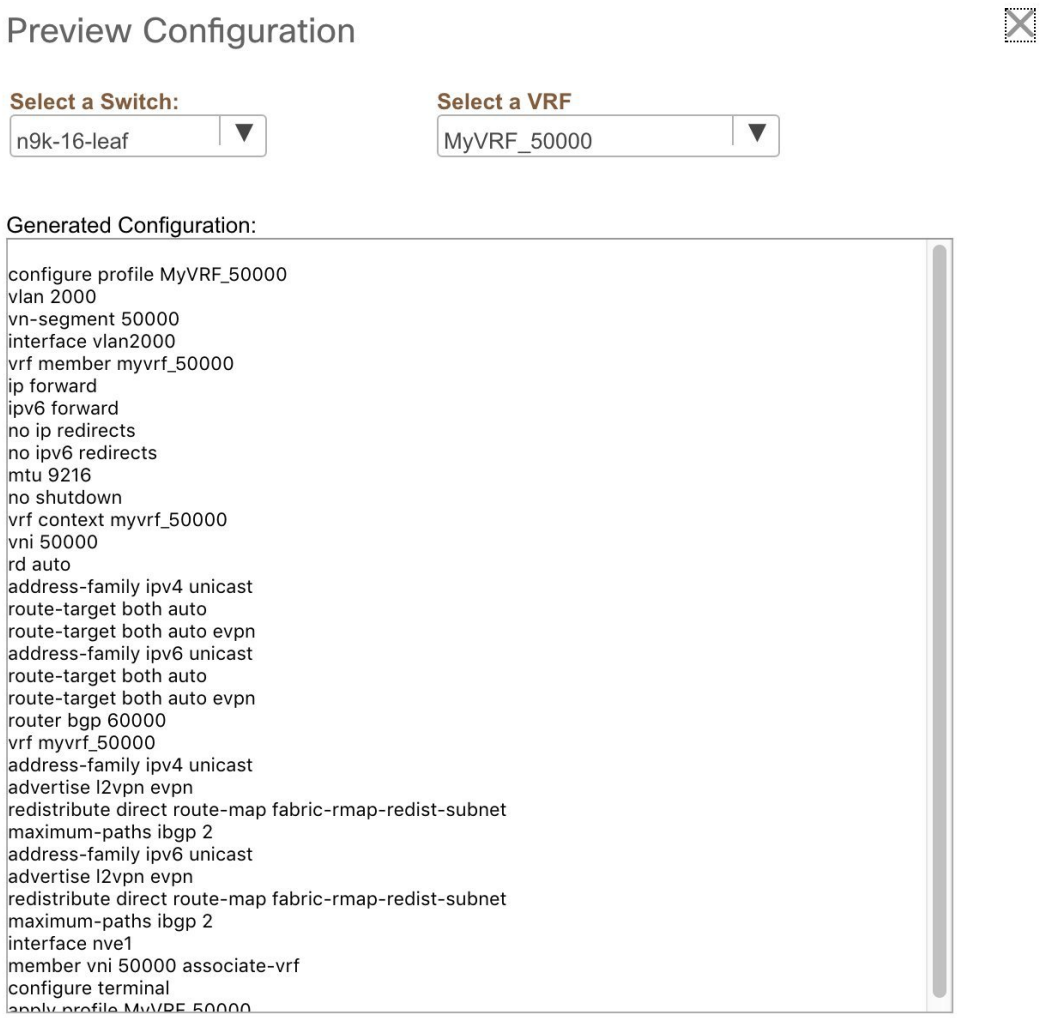

After checking the configurations, close the screen. The *Topology View* screen appears.

6. 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.

Note In case you click **Deploy** and there is no configuration diff that has to be deployed, a pop-up window comes up stating **No switches PENDING for deployment**.

VRFs Deployment for an MSD Fabric

Consider a scenario wherein you are deploying the same VRFs on different member fabric border devices. You can choose one fabric, deploy VRFs on its border devices, and then choose the second fabric and deploy the VRFs.

Alternatively, you can choose the MSD fabric, and deploy the VRFs from a single topology view of all member fabric border devices at once.

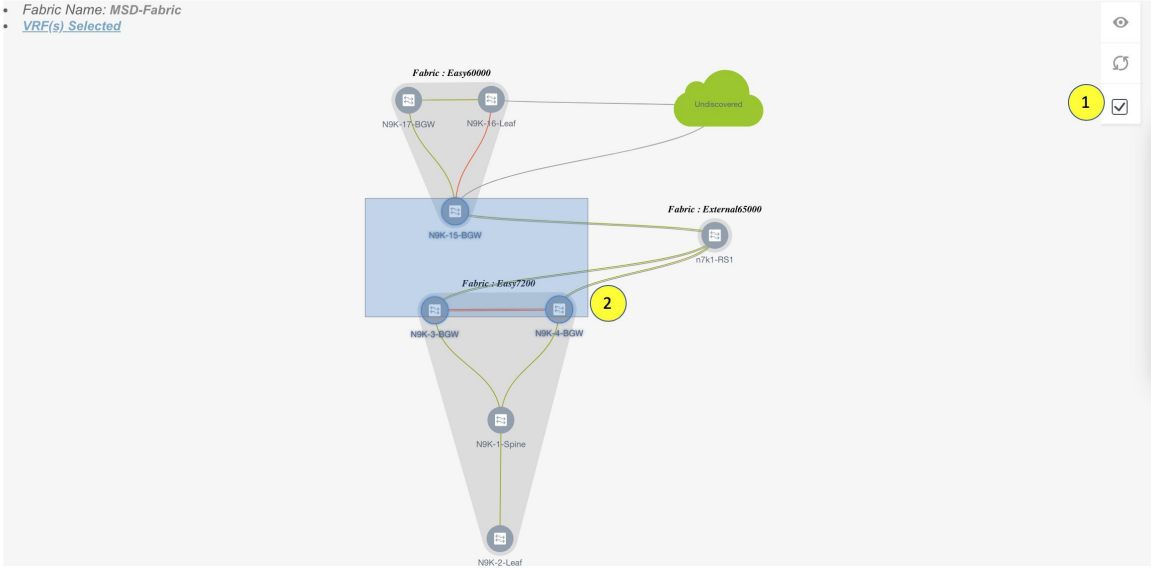

Detailed View

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. This lists the VRFs in a tabular view.

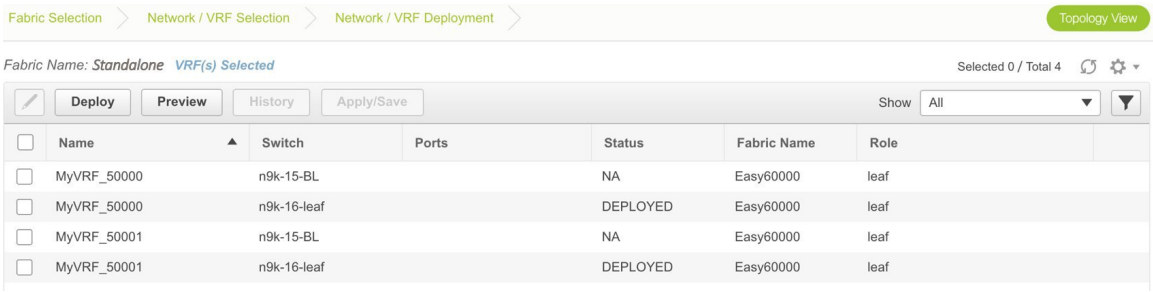

The options:

Edit - Select a VRF and click the Edit icon at the top left part of the screen.

Note If you select one VRF/switch entry, the VRF Attach screen comes up. To maintain consistency across the Topology View and Detailed View screens, the VRF Attach screen displays all VRFs, and not just the selected VRF/switch entry.

Preview – Click Preview to preview configurations before deployment. You can only preview pending configurations, and not uninitiated or deployed configurations.

Deploy – Click Deploy to provision VRFs onto the switches.

History – Select a row and click History to view the configuration instances and status. Network and VRF-wise configurations are displayed. Click in the Status column of any instance for more details.

The fields in the table contain the configuration instance in each row, the associated switch and fabric names, the switch role, and the deployment status.

Apply/Save – Selecting a VRF and clicking Apply/Save will select a switch for the VRF to be deployed on.

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 for the Standalone Fabric

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

1. Click **Control > Networks** (under **Fabrics** submenu).

   The Networks screen comes up.

2. Choose the correct fabric from SCOPE. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

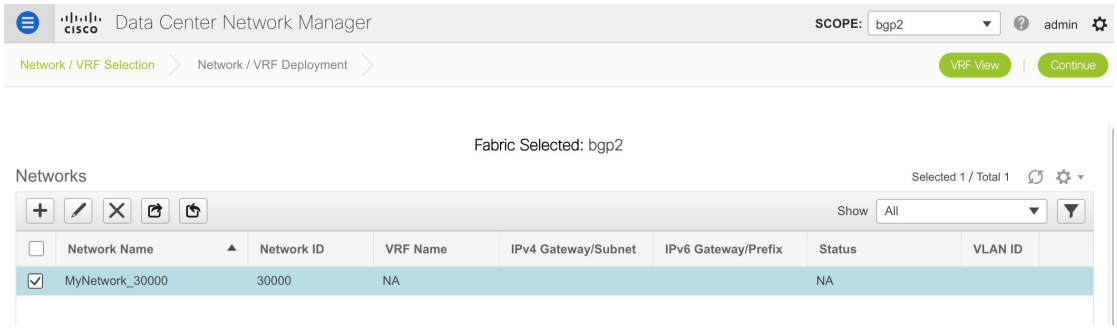

3. Select the networks that you want to undeploy and click Continue. The topology view comes up.
4. Select the Multi-Select button (if you are undeploying the networks from multiple switches), and drag the cursor across switches with the same role. The Network Attachment screen comes up.

   (For a single switch, double-click the switch and the Network Attachment screen comes up).

   (For a single switch, double-click the switch and the Switches Deploy screen comes up).

5. In the Network Attachment screen, the Status column for the deployed networks is displayed as DEPLOYED. Clear 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 are undeployed.

Undeploying VRFs for the Standalone Fabric

You can undeploy VRFs from the deployment screen. The DCNM screen flow for undeployment is similar to the deployment process flow.

1. Choose **Control > Fabrics > VRFs**.
2. Choose the correct fabric from **SCOPE**. When you select a fabric, the **VRFs** screen refreshes and lists networks of the selected fabric.
3. Select the VRFs that you want to undeploy and click **Continue**. The *Topology View* page comes up.
4. Select the Multi-Select option (if you are undeploying the VRFs from multiple switches), and drag the cursor across switches with the same role. The VRF Attachment screen comes up.

   (For a single switch, double-click the switch and the VRF Attachment screen comes up).
5. In the Switches Deploy screen, the **Status** column for the deployed VRFs is displayed as DEPLOYED. Clear the check boxes next to the switches, as needed. Ensure that you repeat this on all tabs since each tab represents a VRF.
6. 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.

7. 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.
8. Go to the VRFs page to verify if the networks are undeployed.

Deleting Networks and VRFs

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.

Configuring Multiple VLAN IDs to a Single VNI

The following procedure shows how to tag multiple VLAN IDs to a single VNI in DCNM.

Procedure

Step 1 Navigate to **Control** > **Networks**.

Step 2 Select the fabric from the **Scope** drop-down list and then select the network. Click **Continue**.

Step 3 Check the **Multi-Select** check box and drag the cursor over the switches that needs to be updated with VLAN IDs.

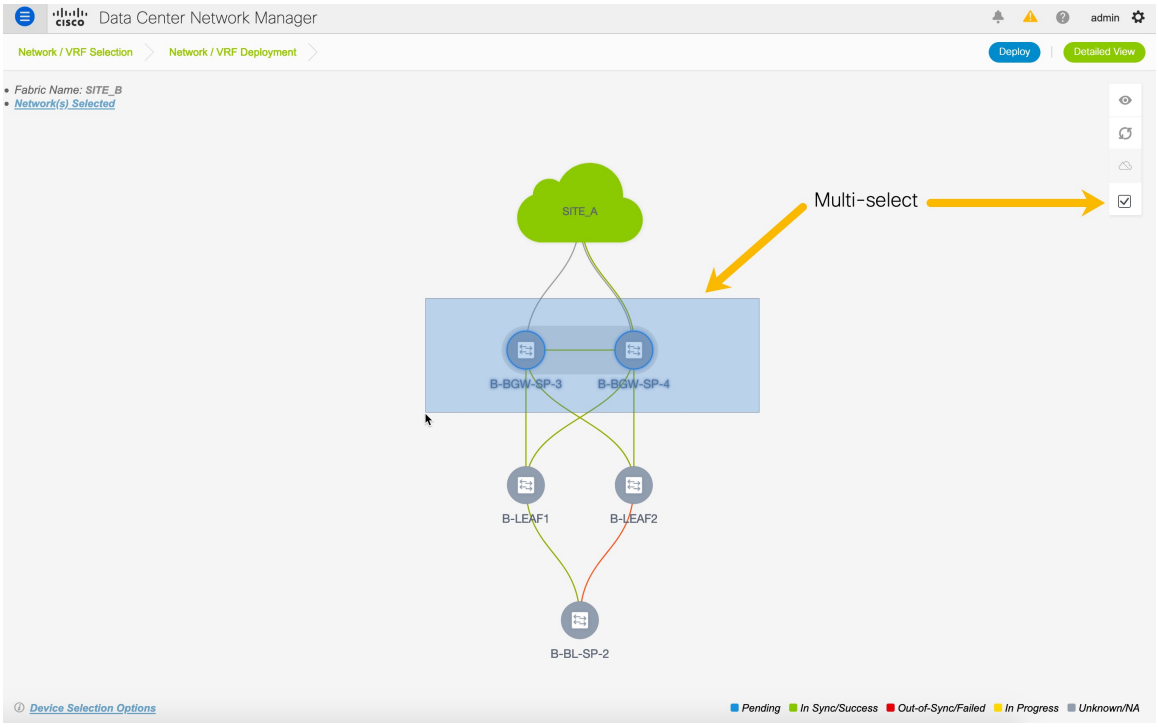

Step 4 In the **Network Attachment** window, edit the VLAN ID for the switches and click **Save**.

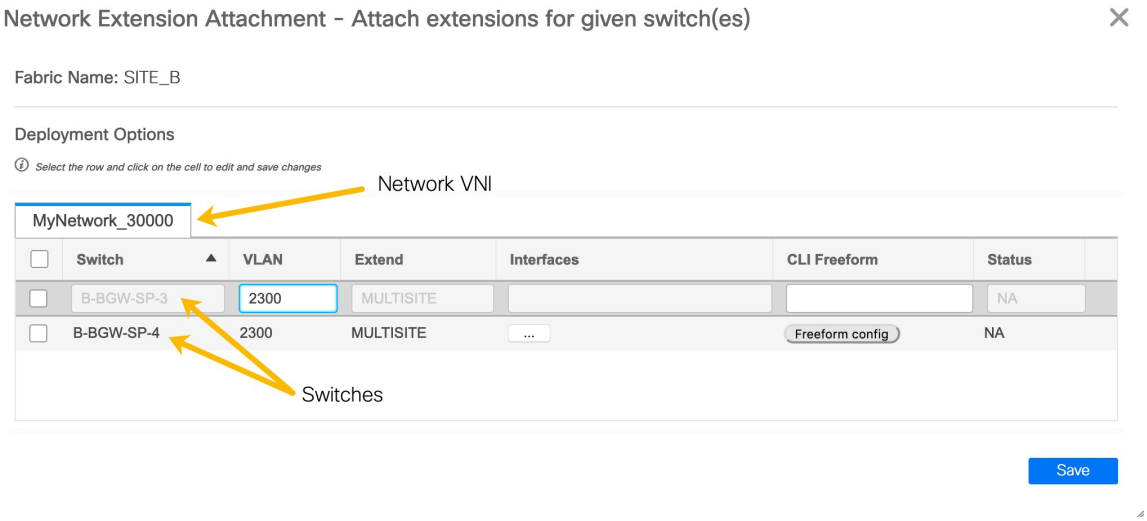

Step 5 Click **Deploy** to deploy the configuration.

Fabric Backup and Restore

This section describes the fabric backup and restore in Cisco DCNM.

Backing Up Fabrics

You can back up all fabric configurations and intents automatically or manually. You can save configurations in DCNM, which are the intents. The intent may or may not be pushed on to the switches.

DCNM doesn't back up the following fabrics:

- External fabrics in monitor-only mode: Backing up of external fabrics in monitor-only mode isn't supported because you can't restore any configurations or intent.
- Parent MSD fabrics in releases earlier than Cisco DCNM, Release 11.4(1): You can only back up the configurations and intent of member fabrics in an MSD fabric individually.

Note From Cisco DCNM, Release 11.4(1), you can take backups of MSD fabrics. When you initiate a backup from the parent fabric, the backup process is applicable for the member fabrics as well. However, DCNM stores all the backed-up information of the member fabrics and the MSD fabric together in a single directory.

The backup does not capture the intent related to IFC. When you're backing up an external fabric, the checkpoints are copied from the switches to DCNM. The backup configuration files are stored in the following path in DCNM: `/usr/local/cisco/dcm/dcnm/data/archive`

By default, DCNM archives only 50 backups, and removes the older backups.

You can set the number of backup files to be archived in the **Server Properties** window. Search for the **# Number of archived files per fabric to be retained:** section in the **Server Properties** window. Enter a value in the **archived.versions.limit** field.

Backing Up Fabrics Automatically

You can enable an automatic hourly backup or scheduled backup for fabric configurations and intents. There are two types of automatic backup.

The backup has the information related to intent and fabric configurations in addition to associated state of the resource manager in terms of used resources on fabrics. DCNM backs up only when there's a configuration push. DCNM triggers the automatic backup only if you didn't trigger any manual backup after the last configuration push.

There are two types of automatic backup.

- **Hourly Fabric Backup**: You can enable an hourly backup.
- **Scheduled Fabric Backup**: You can schedule a fabric backup for regular intervals.

Note In external fabrics, DCNM backs up the changes in the running configurations as well. The configuration push happens after a deploy. If you didn't deploy the changes, you can't back up them in an intent.

Hourly and scheduled backup processes happen only during the next periodic configuration compliance activity, and there can be a delay of up to an hour.

Hourly and Scheduled Backup of Fabrics

To enable automatic backup of fabric configurations and intents from the Cisco DCNM Web client, perform the following steps:

Procedure

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

The **Fabric Builder** window appears.

Step 2 Click the **Edit Fabric** icon for the fabric you want to backup.

Step 3 Click the **Configuration Backup** tab.

Step 4 Choose the nature of backup by checking the appropriate check box.

The valid options are **Hourly Fabric Backup** and **Scheduled Fabric Backup**. If you want to enable both the backups, check the **Hourly Fabric Backup** check box and the **Scheduled Fabric Backup** check box.

Note If you check the **Scheduled Fabric Backup** check box, specify the scheduled backup time in the **Scheduled Time** field. Enter the value in `HH:MM` format.

Step 5 Click **Save**.

Step 6 Click the fabric and go to the fabric topology window.

Step 7 Click **Save & Deploy**.

Backing Up Fabrics Manually

You can enable a manual backup for fabric configurations and intents. Regardless of the settings you choose under the **Configuration Backup** tab in the **Edit Fabric** dialog box, you can initiate a backup using this option.

To initiate a manual backup of fabric configurations and intents from the Cisco DCNM Web UI, perform the following steps:

Procedure

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

The **Fabric Builder** window appears.

Step 2 Click the fabric for which you want to backup immediately.

The fabric topology window appears.

Step 3 Click **Backup Now** in the **Actions** pane.

The **Backup Now** dialog appears.

Step 4 Enter a tag name in the **Tag** field.

Step 5 Click **OK**.

A confirmation message appears that the backup is triggered successfully.

Note The confirmation message only states that the backup is triggered and not if the backup is successful.

Step 6 (Optional) Click **Restore Fabric** from the **Actions** pane to confirm if the manual backup is successful or not.

When you hover over the backup, the name has the tag you mentioned in *Step 4* confirming that it's a manual backup.

Restoring Fabrics

This section describes the fabric restoring for different types of fabrics. Cisco DCNM supports configuration restore at fabric level. Take a backup of the configuration to restore it.

Restoring Easy Fabrics

To restore an easy fabric in Cisco DCNM, perform the following steps from the Cisco DCNM Web UI:

Procedure

Step 1 Choose **Control > Fabrics > Fabric Builder** and select a fabric.

Step 2 Select **Restore Fabric** from the **Actions** menu.

The **Restore Fabric** window appears.

Step 3 Choose the time for which you want to restore the configuration.

Valid values are **1m**, **3m**, **6m**, **YTD**, **1y**, and **All**. You can zoom into the graph. By default the backup information for **1m**, which is one month, appears. You can also select a custom date range. The backup information includes the following information:

- Backup date
- Total number of devices
- Number of devices in sync
- Number of devices out of sync

Step 4 Click **View Backup Summary** to see the selected backup information of the devices in sync.

The switch name, switch serial number, IP address, status, and the configuration details of the devices appear.

Note If you add or remove devices from the fabric, the backup isn't valid. You can restore only the valid backups.

Step 5 Click **Get Config** to preview the configuration details.

Config Preview window appears, which has two tabs.

- **Backup Config**: This tab displays the backup configuration for the selected device.
- **Current Config**: This tab displays the current configuration for the selected device.

Step 6 Go back to **View Backup Summary** window.

Step 7 Click **Restore Intent** to proceed with the restoring.

The **Restore Status** window appears. You can view the status of the following:

- **Validating Backup**
- **Restoring fabric intent**
- **Restoring underlay intent**
- **Restoring interface intent**
- **Restoring overlay intent**

The valid values for the status of any action are **In Progress**, **Pending**, or **Failed**.

Note If the status of **Validating Backup** is **Failed**, other restoring actions won't be listed in this window.

Step 8 Click **Next** after the intent is restored.

The **Configuration Preview** window appears. You can view the following details in this window:

- Switch name
- IP address
- Switch serial number
- Preview configuration
- Status
- Progress

Step 9 Click **Deploy** to deploy the restored configuration.

The **Configuration Deployment Status** window appears. You can view the details of the switch name, IP address, status, status description, and the progress.

Step 10 Click **Close** after the restoring process is complete.

Restoring External Fabrics

When you restore an external fabric, the backed-up checkpoint is copied from DCNM to switches. To restore an external fabric in Cisco DCNM, perform the following steps from the Cisco DCNM Web UI:

Procedure

Step 1 Choose **Control** > **Fabrics** > **Fabric Builder** and select a fabric.

Step 2 Select **Restore Fabric** from the Actions menu.

The **Restore Fabric** window appears.

Step 3 Select the time for which you want to restore the configuration.

Valid values are **1m**, **3m**, **6m**, **YTD**, **1y**, and **All**. You can zoom into the graph. By default the backup information for **1m**, which is one month, appears.

When you select a backup version, the vertical bar representing it turns grey, and corresponding information is displayed at the bottom part of the screen. It includes the following information:

- Backup date
- DCNM Version
- Total number of devices
- Number of devices in sync
- Number of devices out of sync

You can select a custom date range either by rearranging the date slide below the vertical bars, or using the **From** and **To** boxes at the top right part of the screen.

Step 4 Click **View Backup Summary** to see the selected backup information of the devices in sync.

The switch name, switch serial number, IP address, status, Restore Supported (indicating whether the device supports checkpoint rollback or not), the configuration details of the devices, and the VRF appear.

Note For information about the support for the checkpoint rollback feature in platforms, refer to the respective platform documentation.

By default, the management VRF is displayed in the VRF column because it is used for the copy operation during the restore process. If you want to use a different VRF for the copy operation, update the VRF column. To update the same VRF for all devices, use the Apply for all devices option at the bottom-left part of the screen. A sample screenshot:

Note If you added or removed devices to the fabric, you can't restore a fabric from the present day to a past date.

Step 5 Click **Get Config** to preview device configuration details.

The **Config Preview** window appears, which has three tabs.

- **Backup Config**: This tab displays the backup configuration for the selected device.
- **Current Config**: This tab displays the current running configuration of the selected device.
- **Side-by-side Comparison**: This tab displays current running configuration on the switch, and the backup configuration (or expected configuration).

Step 6 Go back to the **View Backup Summary** window.

Step 7 Click **Restore Intent** to proceed with the restoring.

The **Restore Status** window appears. You can view the status of the following:

- **Validating Backup**
- **Restoring fabric intent**
- **Restoring underlay intent**
- **Restoring interface intent**
- **Restoring overlay intent**
- **Intent Regeneration**

The valid values for the status of any action are **In Progress**, **Pending**, **Completed**, or **Failed**.

Note If the status of **Validating Backup** is **Failed**, other restoring actions won't be listed in this window.

Step 8 Click **Close** after the restore process is complete.

Deleting a VXLAN BGP EVPN 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.

Post DCNM 11.2(1) Upgrade for VXLAN BGP EVPN, External, and MSD Fabrics

Note the following guidelines after you upgrade to the DCNM Release 11.2(1):

- As part of the upgrade from an earlier DCNM release, the fabric and associated templates are carried over to the DCNM Release 11.2(1).

- You can use the old fabric template, but you will not be able to use the features introduced in the DCNM Release 11.2(1). Perform the following steps to use the new features:

  1. Edit the settings of each fabric by updating the old fabric template with the equivalent new fabric template and clicking **Save**.

     The following table shows the old and new fabric template names in DCNM.

Old Template Fabric Name	New Template Fabric Name
Easy_Fabric	Easy_Fabric_11_1
External_Fabric	External_Fabric_11_1
MSD_Fabric	MSD_Fabric_11_1

  2. Under the **Advanced** tab, the **iBGP Peer-Template Config** field displays the value **null**. If this value is displayed, remove it and click **Save**.

     Note You can skip this step if you are upgrading from the DCNM Release 11.1(1) to the DCNM Release 11.2(1).

  3. Navigate to each fabric in the Topology view, and click **Save & Deploy** to deploy any changes.

     If you encounter any new or unexpected pending configurations after you click **Save & Deploy**, refer Configuration Compliance in DCNM, on page 188.

     Caution Some configuration changes can be expected as part of this step. Therefore, perform it only during a scheduled maintenance window.

- After a multi-level upgrade from Cisco DCNM 10.4(2) or 11.0(1), you can change the VRF templates to **Default_VRF_Universal** or **Default_VRF_Extension_Universal** to enable **ipv6 address use-link-local-only**.

Changing ISIS Configuration from Level 1 to Level 2

This procedure shows how to change ISIS configuration on switches from Level 1 to Level 2 in a VXLAN fabric deployment.

1. Choose **Control > Fabrics > Fabric Builder**.
2. Click a fabric in the **Fabric Builder** window.
3. Click **Tabular view** under **Actions** menu.
4. Search for all the **base_isis** policies in the **Template** search field.
5. Select all the **base_isis** policies and click the **Delete** icon to delete policies
6. Click **Save & Deploy**.

After all the **base_isis** policies are deleted, DCNM considers the migrated brownfield fabric as a greenfield fabric and creates the **base_isis_level2** policies on the switches.

Configuration Compliance in DCNM

The entire intent or expected configuration defined for a given switch is stored in DCNM. When you want to push this configuration down to one or more switches, the configuration compliance (CC) module is triggered. CC takes the current intent, the current running configuration, and then comes up with the set of configurations that are required to go from the current running configuration to the current expected config so that everything will be In-Sync.

When performing a software or firmware upgrade on the switches, the current running configuration on the switches is not changed. Post upgrade, if CC finds that the current running configuration does not have the current expected configuration or intent, it reports an Out-of-Sync status. There is no auto deployment of any configurations. You can preview the diffs that will get deployed to get one or more devices back In-Sync.

With CC, the sync is always from the DCNM to the switches. There is no reverse sync. So, if you make a change out-of-band on the switches that conflicts with the defined intent in DCNM, CC captures this diff, and indicates that the device is Out-of-Sync. The pending diffs will undo the configs done out-of-band to bring back the device In-Sync. Note that such conflicts due to out-of-band changes are captured by the periodic CC run that occurs every 60 mins by default, or when you click the RESYNC option either on a per fabric or per switch basis. Note that you can also capture the out-of-band changes for the entire switch by using the CC REST API. For more information, see *Cisco DCNM REST API Guide, Release 11.2(1)*.

From Cisco DCNM Release 11.2(1), to improve ease of use and readability of deployed configurations, CC in DCNM has been enhanced with the following:

- All displayed configurations in DCNM are easily readable and understandable.
- Repeated configuration snippets are not displayed.
- Pending configurations precisely show only the diff configuration.
- Side-by-side diffs has greater readability, integrated search or copy, and diff summary functions.

All freeform configurations have to strictly match the **show running configuration** output on the switch and any deviations from the configuration will show up as a diff during **Save & Deploy**. You need to adhere to the leading space indentations.

You can typically enter configuration snippets in DCNM using the following methods:

- User-defined profile and templates
- Switch, interface, overlay, and vPC freeform configurations
- Network and VRF per switch freeform configurations
- Fabric settings for Leaf, Spine, or iBGP configurations

Caution The configuration format should be identical to the **show running configuration** of the corresponding switch. Otherwise, any missing or incorrect leading spaces in the configuration can cause unexpected deployment errors and unpredictable pending configurations. If any unexpected diffs or deployment errors are displayed, check the user-provided or custom configuration snippets for incorrect values.

If DCNM displays the "Out-of-Sync" status due to unexpected pending configurations, and this configuration is either unable to be deployed or stays consistent even after a deployment, perform the following steps to recover:

1. Check the lines of config highlighted under the **Pending Config** tab in the **Config Preview** window.
2. Check the same lines in the corresponding **Side-by-side Comparison** tab. This tab shows whether the diff exists in "intent", or "show run", or in both with different leading spaces. Leading spaces are highlighted in the **Side-by-side Comparison** tab.
3. If the pending configurations or switch with an out-of-sync status is due to any identifiable configuration with mismatched leading spaces in "intent" and "running configuration", this indicates that the intent has incorrect spacing and needs to be edited.
4. To edit incorrect spacing on any custom or user-defined policies, navigate to the switch and edit the corresponding policy:
   a. If the source of the policy is **UNDERLAY**, you will need to edit this from the Fabric settings screen and save the updated configuration.
   b. If the source is blank, it can be edited from the **View/Edit policies** window for that switch.
   c. If the source of the policy is **OVERLAY**, but it is derived from a switch freeform configuration. In this case, navigate to the appropriate **OVERLAY** switch freeform configuration and update it.
   d. If the source of the policy is **OVERLAY** or a custom template, perform the following steps:
      1. Navigate to **Administration > DCNM Server > Server Properties**, set the **template.in_use.check** property to **false**. This allows the profiles or templates to be editable.
      2. Edit the specific profile or template from the **Control > Template Library** edit window, and save the updated profile template with the right spacing.
      3. Click **Save & Deploy** to recompute the diffs for the impacted switches.
      4. After the configurations are updated, set the **template.in_use.check** property to **true**, as it slows down the performance of the DCNM system, specifically for **Save & Deploy** operations.

To confirm that the diffs have been resolved, click **Save & Deploy** after updating the policy to validate the changes.

Note DCNM checks only leading spaces, as it implies hierarchy of the command, especially in case of multi-command sequences. DCNM does not check any trailing spaces in command sequences.

Example 1: Configuration Compliance in Switch Freeform Policy

Let us consider an example with an incorrect spacing in the Switch Freeform Config field.

The switch freeform policy is created as shown:

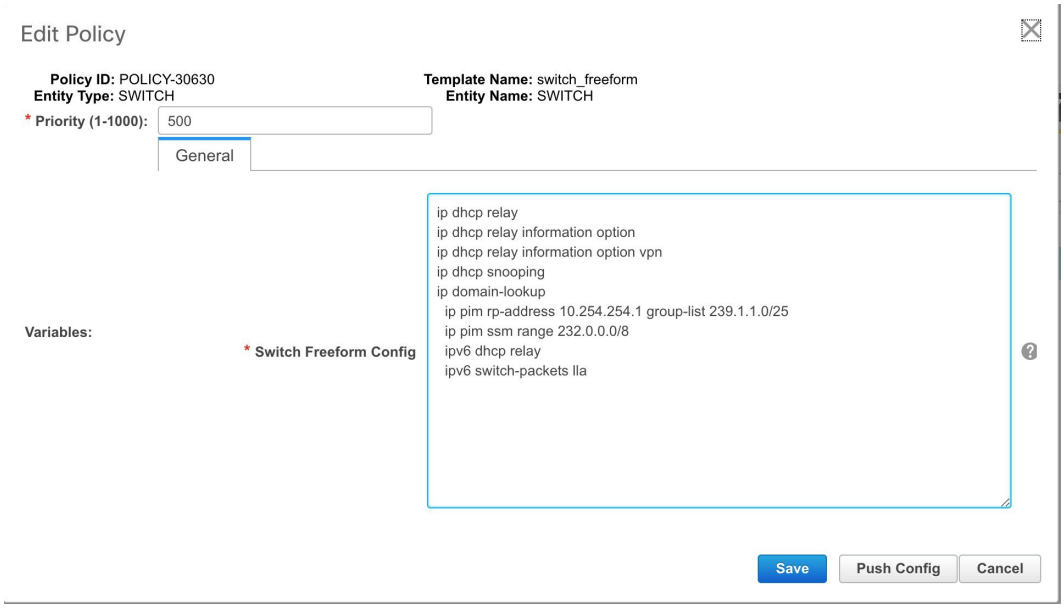

After deploying this policy successfully to the switch, DCNM persistently reports the following diffs:

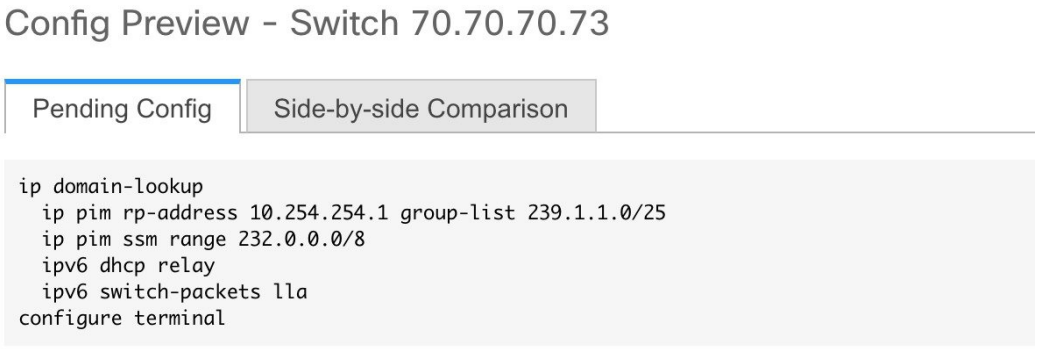

After clicking the **Side-by-side Comparison** tab, you can see the cause of the diff. As seen below, the **ip pim rp-address** line has 2 leading spaces, while the running configuration has 0 leading spaces.

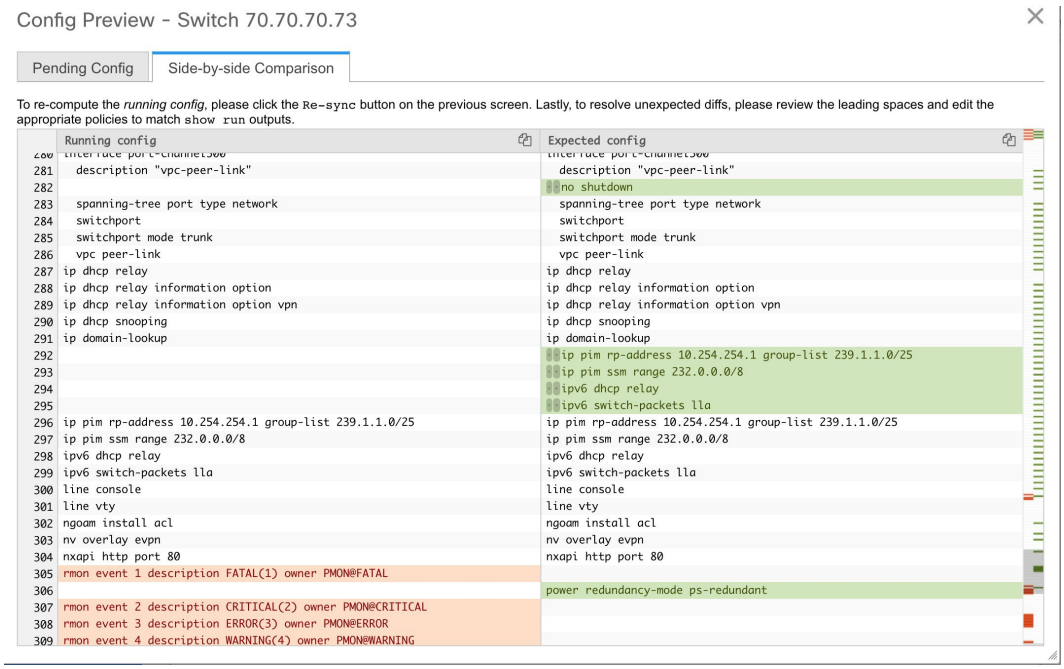

To resolve this diff, edit the corresponding Switch Freeform policy so that the spacing is correct.

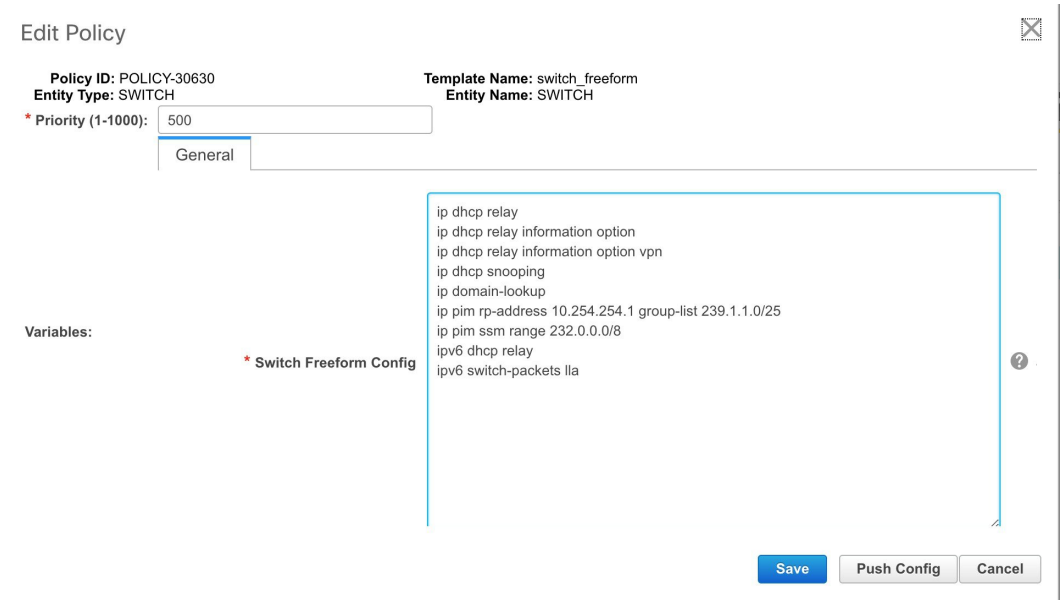

After you save, you can use the **Push Config** or **Save & Deploy** option to re-compute diffs.

As shown below, the diffs are now resolved. The **Side-by-side Comparison** tab confirms that the leading spaces are updated.

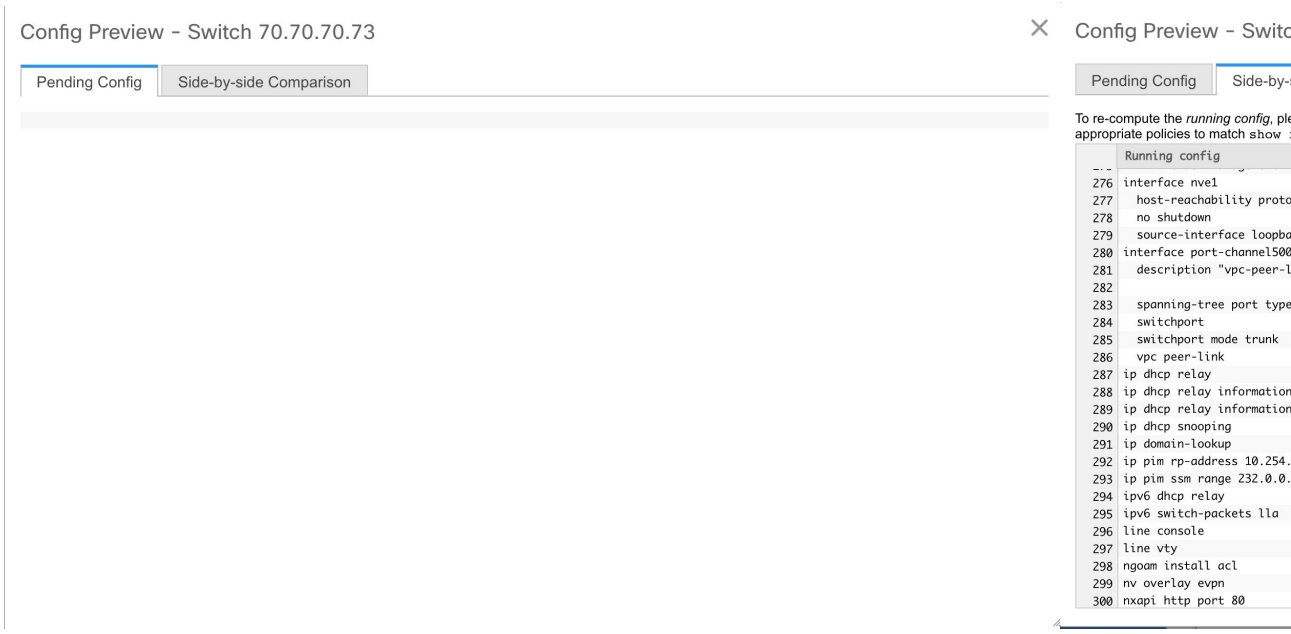

Example 2: Resolving a Leading Space Error in Overlay Configurations

Let us consider an example with a leading space error that is displayed in the **Pending Config** tab.

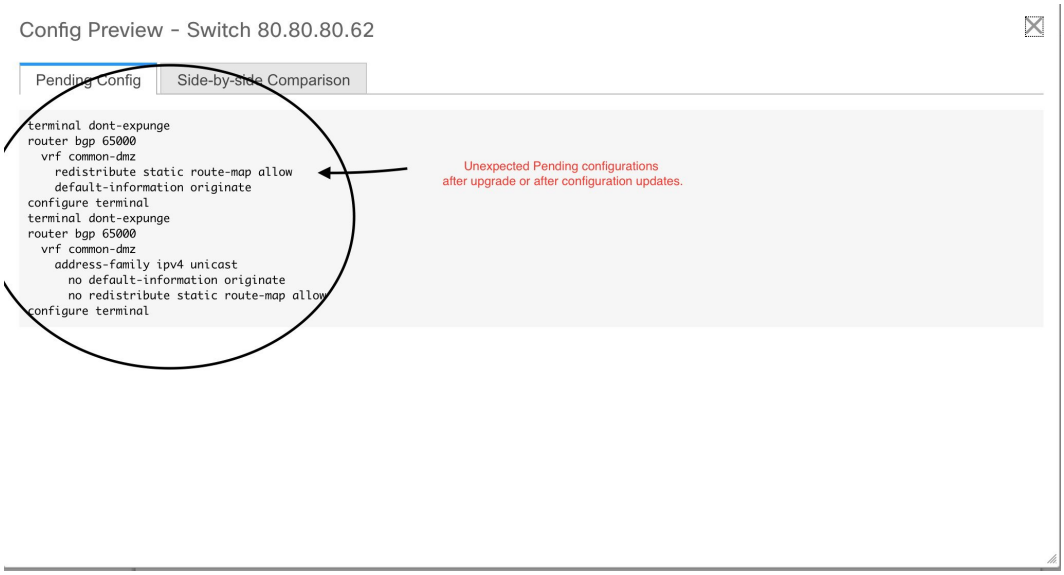

In the **Side-by-side Comparison** tab, search for diffs line by line to understand context of the deployed configuration.

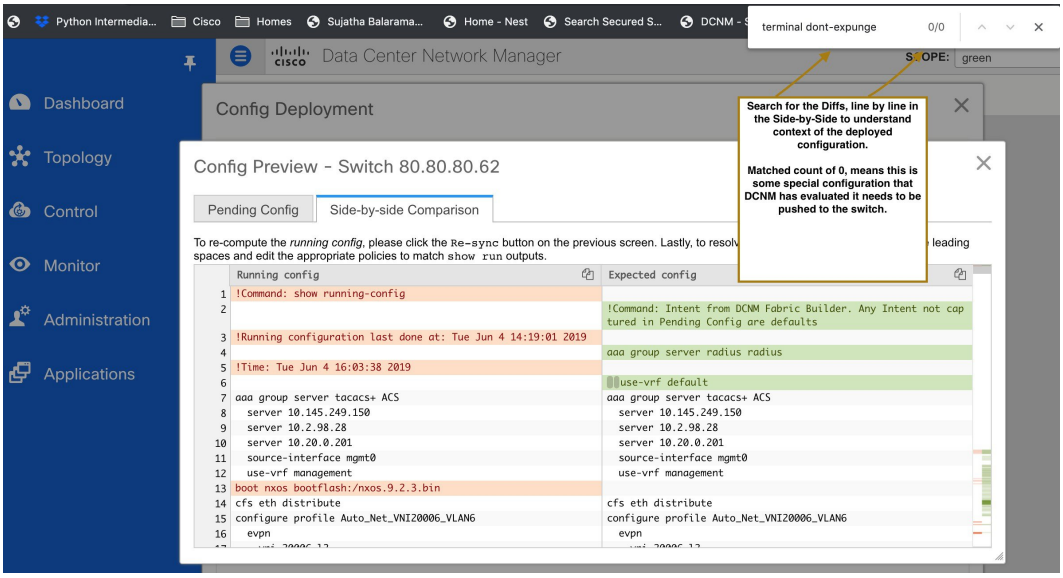

A matched count of 0 means that it is a special configuration that DCNM has evaluated to push it to the switch.

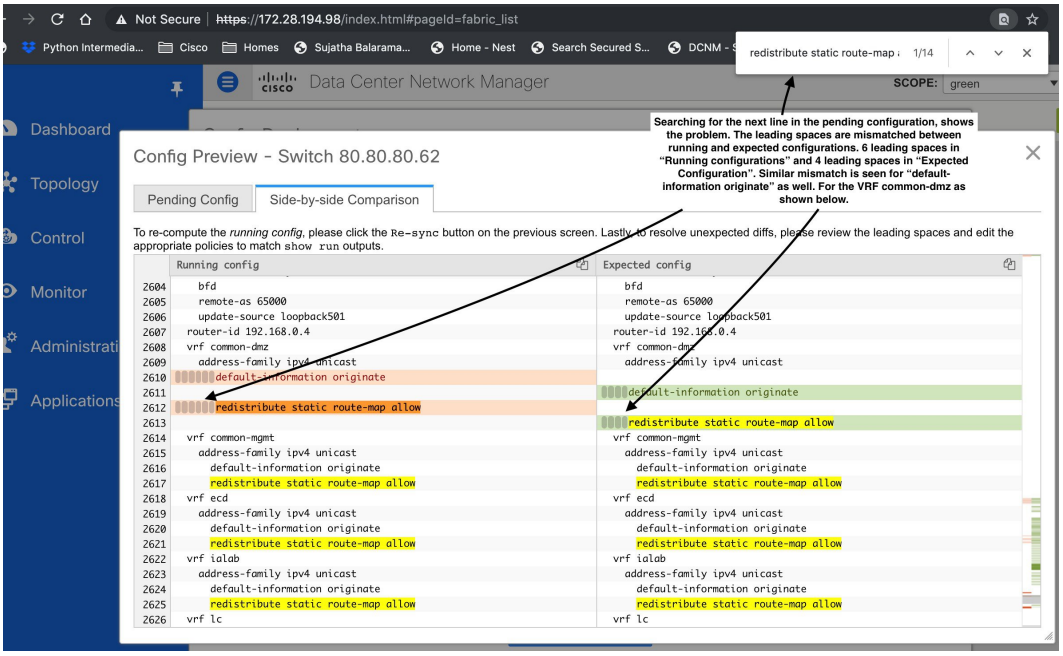

You can see that the leading spaces are mismatched between running and expected configurations.

Navigate to the respective freeform configs and correct the leading spaces, and save the updated configuration.

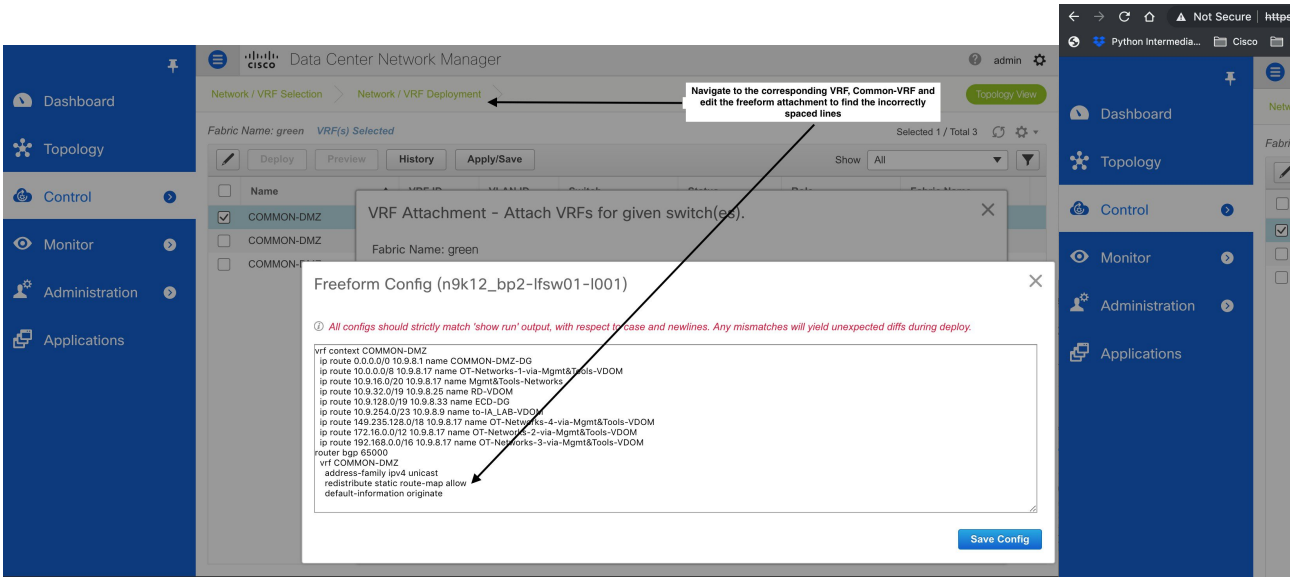

Navigate to the **Fabric Builder** window for the fabric and click **Save & Deploy**.

In the **Config Deployment** window, you can see that all the devices are in-sync.

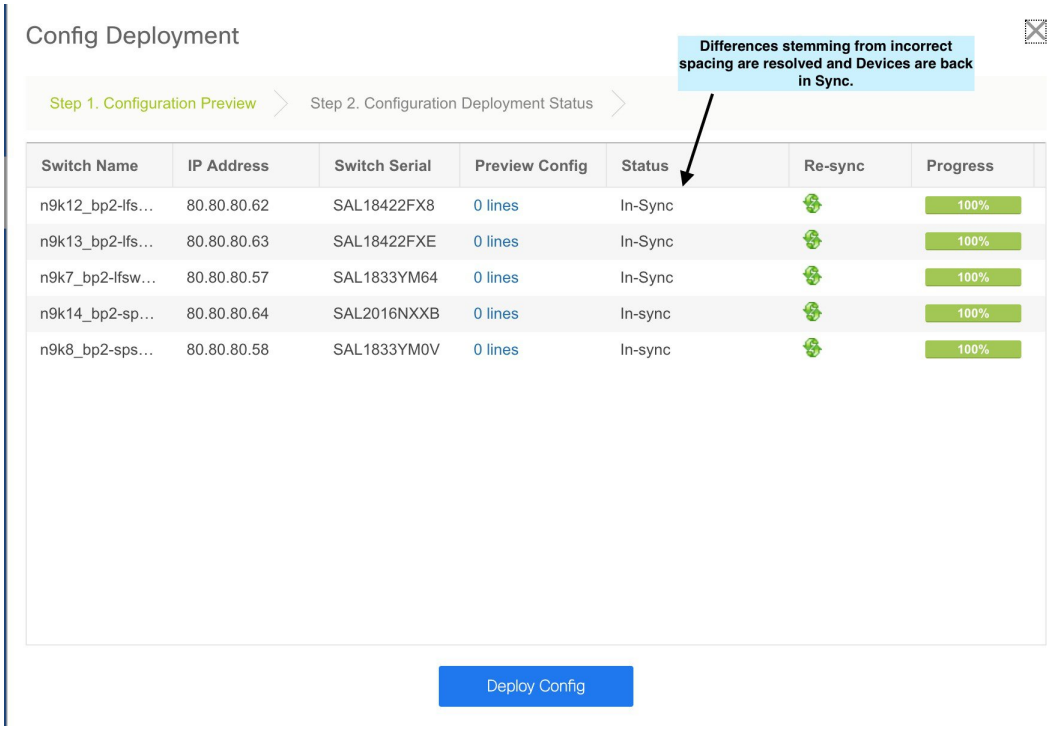

Configuration Compliance in External Fabrics

With external fabrics, any Nexus switch can be imported into the fabric, and there is no restriction on the type of deployment. It can be LAN Classic, VXLAN, FabricPath, vPC, HSRP, etc. When switches are imported into an external fabric, the configuration on the switches is retained so that it is non-disruptive. Only basic policies such as the switch username and mgmt0 interface are created after a switch import.

In the external fabric, for any intent that is defined in the DCNM, configuration compliance (CC) ensures that this intent is present on the corresponding switch. If this intent is not present on the switch, CC reports an Out-of-Sync status. Additionally, there will be a Pending Config generated to push this intent to the switch to change the status to In-Sync. Any additional configuration that is on the switch but not in intent defined in DCNM, will be ignored by CC, as long as there is no conflict with anything in the intent.

When there is user-defined intent added on DCNM and the switch has additional configuration under the same top-level command, as mentioned earlier, CC will only ensure that the intent defined in DCNM is present on the switch. When this user defined intent on DCNM is deleted as a whole with the intention of removing it from the switch and the corresponding configuration exists on the switch, CC will report an Out-of-Sync status for the switch and will generate **Pending Config** to remove the config from the switch. This **Pending Config** includes the removal of the top-level command. This action leads to removal of the other out-of-band configurations made on the switch under this top-level command as well. If you choose to override this behavior, the recommendation is that, you create a freeform policy and add the relevant top-level command to the freeform policy.

Let us see this behavior with an example.

1. A **switch_freeform** policy defined by the user in DCNM and deployed to the switch.

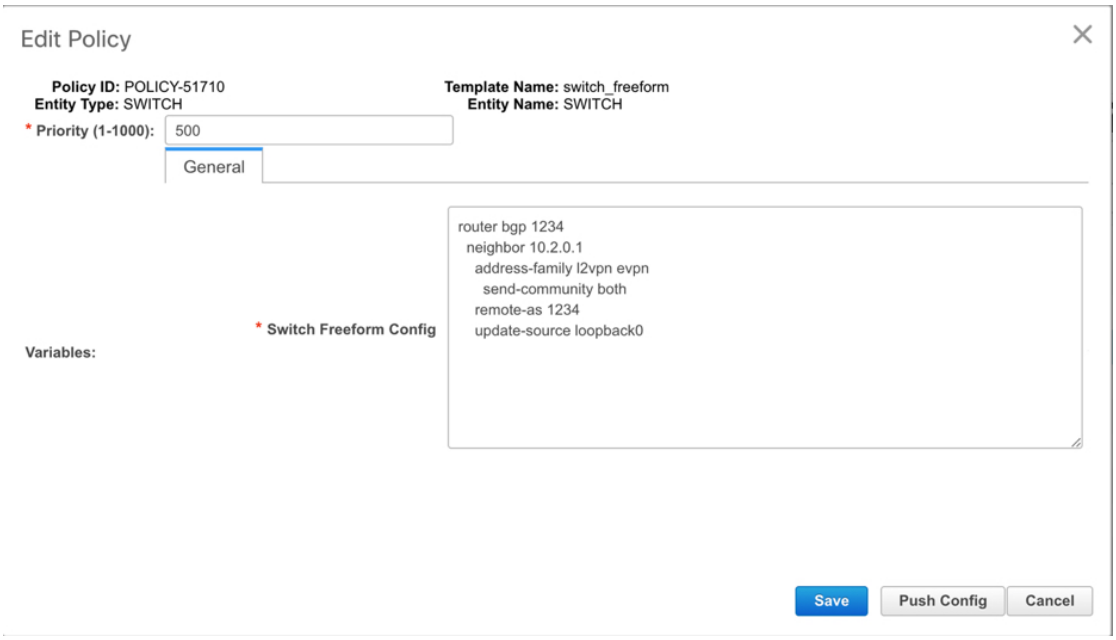

2. Additional configuration exists under **router bgp** in **Running config** that does not exist in user-defined DCNM intent **Expected config**. Note that there is no **Pending Config** to remove the additional config that exists on the switch without a user defined intent on DCNM.

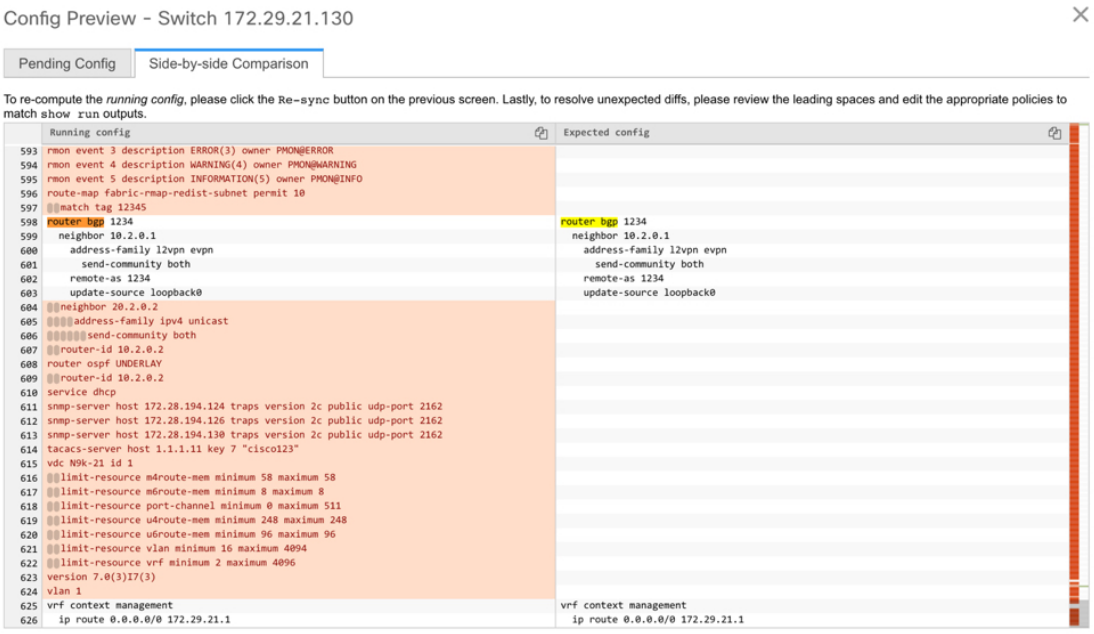

3. The **Pending Config** and the **Side-by-side Comparison** when the intent that was pushed earlier via DCNM is deleted from DCNM by deleting the switch_freeform policy that was created in the Step 1.

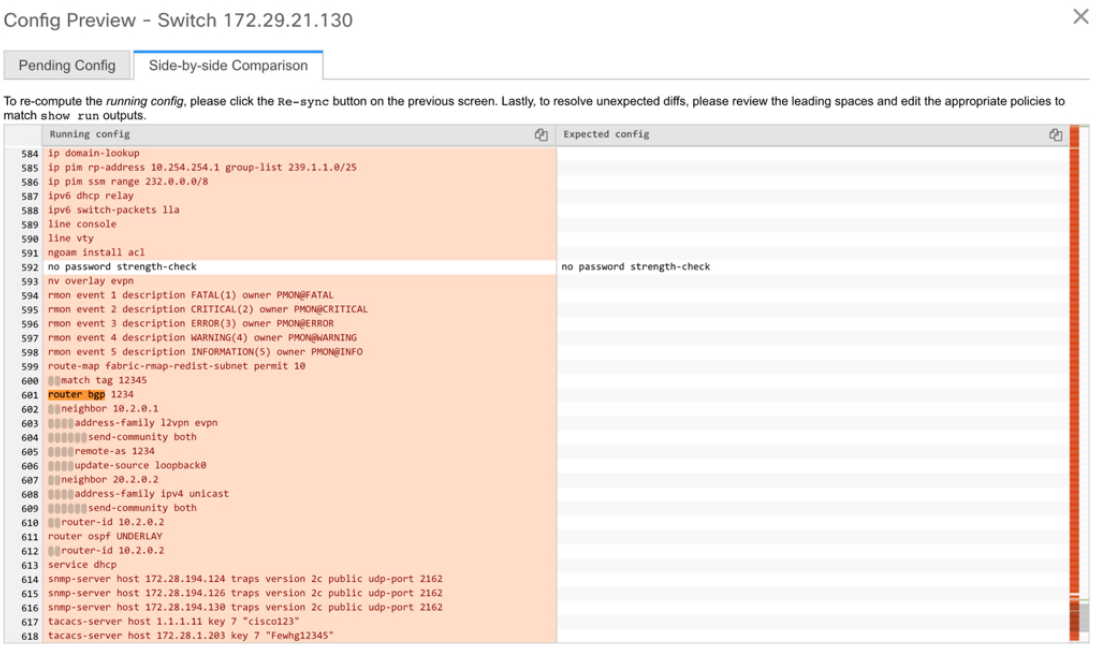

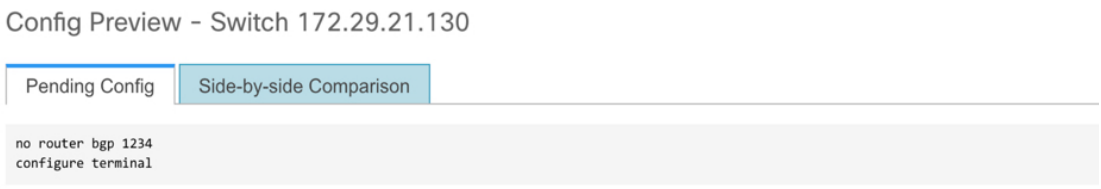

4. A **switch_freeform** policy with the top-level **router bgp** command needs to be created. This enables CC to generate the configuration needed to remove only the desired sub-config which was pushed from DCNM earlier.

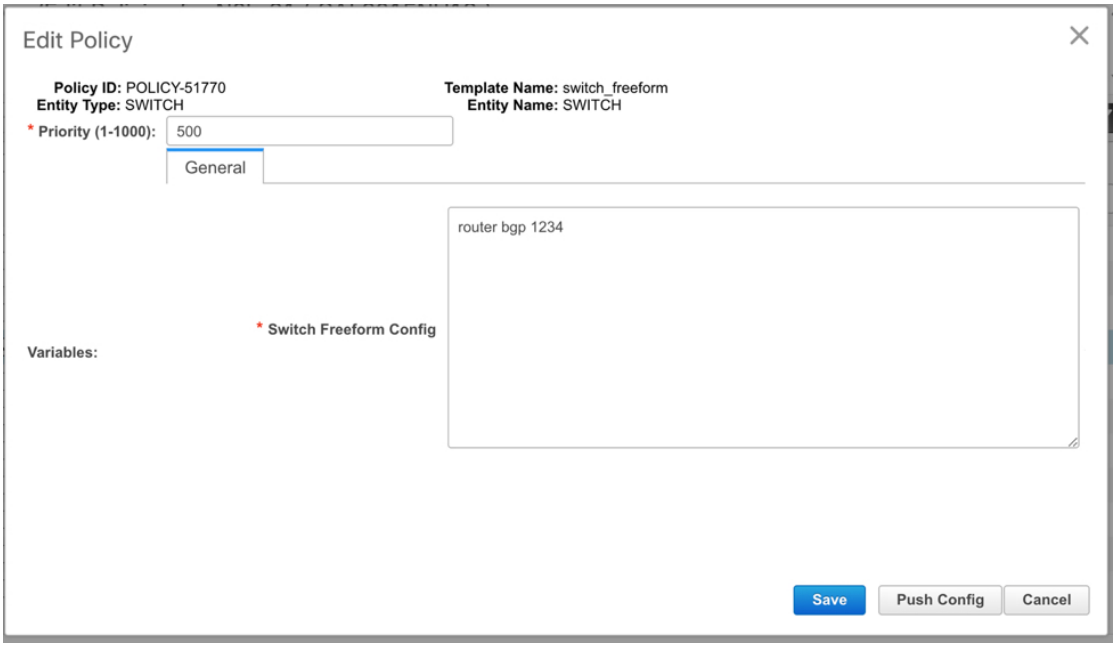

5. The removed configuration is only the subset of the configuration that was pushed earlier from DCNM.

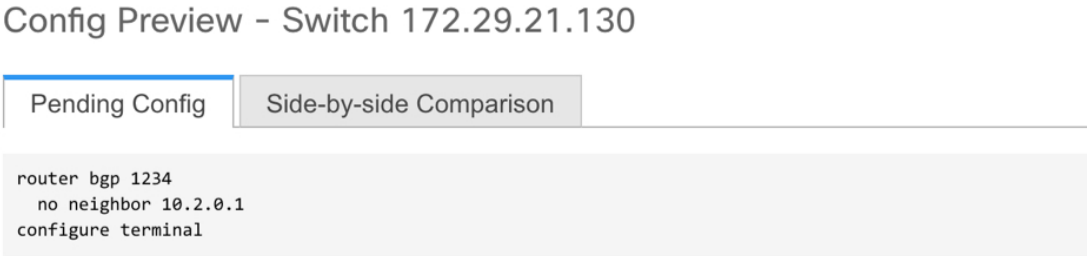

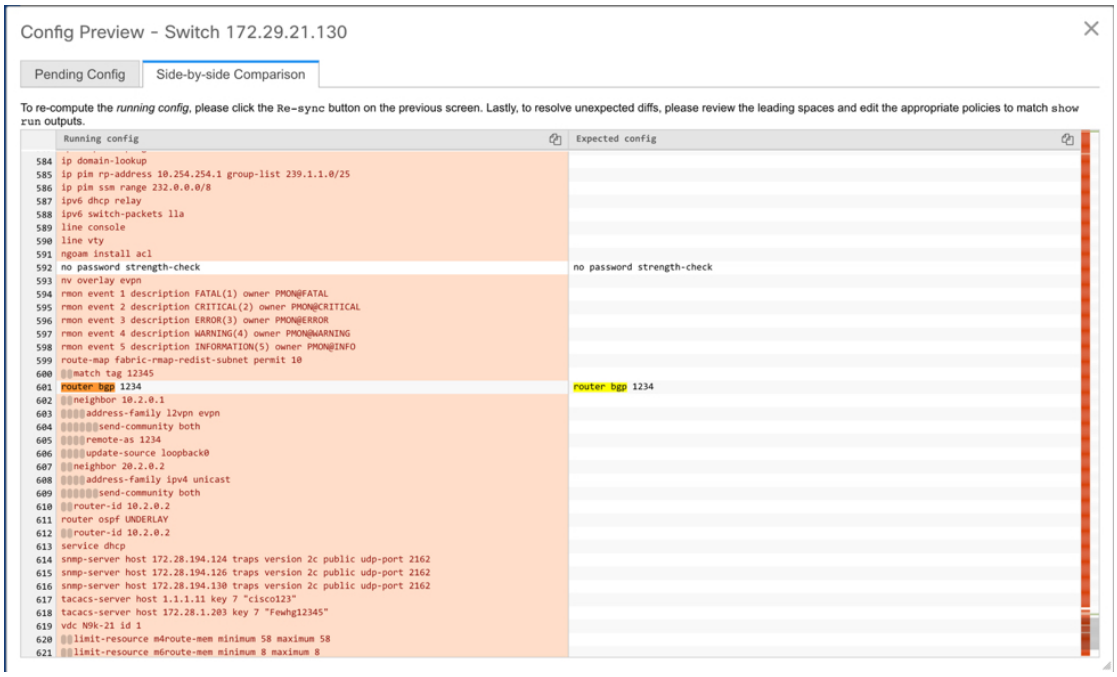

For interfaces on the switch in the external fabric, DCNM either manages the entire interface or does not manage it at all. CC checks interfaces in the following ways:

- For any interface, if there is a policy defined and associated with it, then this interface is considered as managed. All configurations associated with this interface must be defined in the associated interface policy. This is applicable for both logical and physical interfaces. Otherwise, CC removes any out-of-band updates made to the interface to change the status to **In-Sync**.
- Interfaces created out-of-band (applies for logical interfaces such as port-channels, sub interfaces, SVIs, loopbacks, etc.), will be discovered by DCNM as part of the regular discovery process. However, since there is no intent for these interfaces, CC will not report an **Out-of-Sync** status for these interfaces.
- For any interface, there can always be a monitor policy associated with it in DCNM. In this case, CC will ignore the interface's configuration when it reports the **In-Sync** or **Out-of-Sync** config compliance status.

Special Configuration CLIs Ignored for Configuration Compliance

The following configuration CLIs are ignored during configuration compliance checks:

- Any CLI having 'username' along with 'password'
- Any CLI that starts with 'snmp-server user'

Any CLIs that match the above will not show up in pending diffs and clicking **Save & Deploy** in the **Fabric Builder** window will not push such configurations to the switch. These CLIs will not show up in the **Side-by-side Comparison** window also.

To deploy such configuration CLIs, perform the following procedure:

1. Select **Control>Template Library**, and click **+** to create a new custom template with the required configuration CLIs and the following attributes:
   - Template Type: Policy
   - Template Sub Type: Device
   - Template Content Type: TEMPLATE_CLI
2. Select **Control>Fabric Builder**, click **Tabular View**, and select a switch in the **Name** column or select **Control>Fabric Builder** and right-click on the device.
3. Click **View/Edit Policies**.
4. Select the custom template created in **Step 1** and click **Push Config** to deploy the configuration to the switch(es).

Resolving Diffs for Case Insensitive Commands

By default, all diffs generated in DCNM while comparing intent, also known as Expected Configuration, and Running Configuration, are case sensitive. However, the switch has many commands that are case insensitive, and therefore it may not be appropriate to flag these commands as differences. These outlier cases are captured in the **compliance_case_insensitive_clis.txt** text file.

There could be additional commands not included in the existing **compliance_case_insensitive_clis.txt** file that should be treated as case insensitive. If the pending configuration is due to the differences of cases between the Expected Configuration in DCNM and the Running Configuration, you can configure DCNM to ignore these case differences as follows:

1. Modify the following file on the DCNM file system:

   `/usr/local/cisco/dcm/dcnm/model-config/compliance_case_insensitive_clis.txt`

   The sample entries in **compliance_case_insensitive_clis.txt** file are displayed as:

```
[root@dcnm98 model-config]# pwd
/usr/local/cisco/dcm/dcnm/model-config
[root@dcnm98 model-config]# cat compliance_case_insensitive_clis.txt
"^(no |)interface\s+Port(.)"
"^(no |)interface\s+Loo(.)"
"^(no |)interface\s+Eth(.)"
"^update-source\s+Loo(.)"
"^vrf\s+"
"^hardware profile portmode\s+"
"^(.*)route-map\s+(.)"
"^(.*)neighbor-policy(.)"
"(no |)encapsulation\s+(.)"
"(.*)alert-group\s+(.)"
"^streetaddress\s+(.)"
"^transport email\s+(.)"
"(no |)action\s+(.)"
"(no|)\s+\d*\s+remark.*"
[root@dcnm98 model-config]#
```

   If newer patterns are detected during deployment, and they are triggering pending configurations, you can add these patterns to this file. The patterns need to be valid regex patterns.

This enables DCNM to treat the documented configuration patterns as case insensitive while performing comparisons.

2. Run the following command for each fabric to restart the config compliance container:

```
# docker exec -it `docker ps | grep compliance | grep <fabric name> | awk '{print $1}'`
 /usr/bin/pkill python
```

3. Click **Save & Deploy** for fabrics to see the updated comparison outputs.

Enabling 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, border leaf, and border gateway leaf switches in the fabric, at once.
   - On all spine, super spine, border spine, border super spine, border gateway spine and border switches, at once.
2. On a specific switch at the global level.
3. On a specific switch on a per Network or per VRF level.

Leaf switches are identified by the roles Leaf, Border, and Border Gateway. The spine switches are identified by the roles Spine, Border Spine, Border Gateway Spine, Super Spine, Border Super Spine, and Border Gateway Super 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 this as a reference for a new fabric.

Deploying 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 **Edit Fabric** 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, border leaf, and border gateway leaf switches in the fabric.

   Spine Freeform Config - In this field, add configurations for all Spine, Border Spine, Border Gateway Spine, Super Spine, Border Super Spine, and Border Gateway Super 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 204.

4. Click **Save**. The fabric topology screen comes up.
5. Click **Save & Deploy** at the top right part of the screen to save and deploy configurations.

   Configuration Compliance functionality will ensure that the 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, there could be a mismatch between the intended and switch configuration. To resolve the issue, add a **switch_freeform** policy to the affected switch (as explained in the *Deploy Freeform CLIs on a Specific Switch* section). For example, consider the following persistent pending configurations:

```
line vty
logout-warning 0
```

After adding the above configurations in a policy and saving the updates, click **Save and Deploy** in the topology screen to complete the deployment process.

To bring the switch back in-sync, you can add the above configuration in a **switch_freeform** policy saved and deployed onto the switch.

Deploying 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.
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**.
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 204.
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 Policy Configuration:

- You can create multiple instances of the policy.
- For a vPC switch pair, create consistent **switch_freeform** policies on both the vPC switches.
- When you edit a **switch_freeform** policy and deploy it onto the switch, you can see the changes being made (in the **Side-by-side** tab of the Preview option).

Freeform CLI Configuration Examples

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
```

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 intended and running configuration. 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** policy, update the policy with sequence numbers *as shown in the running configuration of the switch*.

After the policy is updated and saved, right click the device and select the per switch **Deploy Config** option to deploy the configuration. Alternatively, use the **Save and Deploy** option in the fabric topology screen

(within Fabric Builder) so that the fabric triggers Configuration Compliance and resolves the configuration mismatch.

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**, **DevicePair**, **Fabric**, and **Link**.
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 are managed with device, device interface, or fabric. Valid values are ID type, subnet, or IP addresses.
Allocated To	Specifies the entity name for which the resource is allocated.
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 **Control** > **Management** > **Virtual Machine Manager**.

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 **Control** > **Management** > **Virtual Machine Manager**.

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 **Control** > **Management** > **Virtual Machine Manager**.

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 **Control** > **Management** > **Virtual Machine Manager**.

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.

Note Notifications appear next to **Import Template Zip File** if there are issues while loading templates after restarting the server. Click the notifications to see the errors in the **Issues in loading Template** window. Templates with errors are not listed in the **Templates** window. To import these templates, correct the errors, and import them.

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.
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.

Table 3: Advanced Template Properties

Field	Description
Implements	Displays the abstract template to be implemented.
Dependencies	Specifies the specific feature of a switch.
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
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, IOS-XE, IOS-XR, Others, 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_VPC • INTERFACE_ETHERNET • INTERFACE_BD • INTERFACE_PORT_CHANNEL • INTERFACE_FC • INTERFACE_MGMT • INTERFACE_LOOPBACK • INTERFACE_NVE • INTERFACE_VFC • INTERFACE_SAN_PORT_CHANNEL • DEVICE • FEX • INTRA_FABRIC_LINK • INTER_FABRIC_LINK	

Property Name	Description	Valid Values	Optional?
		• INTERFACE • SHOW • VLAN • INTERFACE_VLAN • INTERFACE_VPC • INTERFACE_ETHERNET • INTERFACE_BD • INTERFACE_PORTCHANNEL • INTERFACE_FC • INTERFACE_MGMT • INTERFACE_LOOPBACK • INTERFACE_NVE • INTERFACE_VFC • INTERFACE_SAN_PORTCHANNEL • DEVICE • FEX • INTRA_FABRIC_LINK • INTER_FABRIC_LINK • 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 • INTRA_FABRIC_LINK • INTER_FABRIC_LINK • 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

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

Variable Type	Valid Value	Iterative?

ipAddressList	You can have a list of IPv4, IPv6, or a combination of both types of addresses.
	
	`Example 1: 172.22.31.97,`
	`172.22.31.99,`
	`172.22.31.105,`
	`172.22.31.109`
	`Example 2:`
	`2001:0db8:85a3:0000:0000:8a2e:0370:7334,`
	
	`2001:0db8:85a3:0000:0000:8a2e:0370:7335,`
	
	`2001:0db8:85a3:1230:0000:8a2f:0370:7334`
	`Example 3: 172.22.31.97,`
	`172.22.31.99,`
	
	`2001:0db8:85a3:0000:0000:8a2e:0370:7334,`
	
	`172.22.31.254`
		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
ISISNetAddress	`Example:` `49.0001.00a0.c96b.c490.00`	No
long	`Example: 100`	No
macAddress	14 or 17 character length MAC address format	No

string	Free text, for example, used for the description of a variable
	
	`Example:`
	`string scheduledTime`
	`{`
	
	`regularExpr=^([01]\d|2[0-3]):([0-5]\d)$;`
	`}`
		No

Variable Type	Valid Value	Iterative?
string[]	Example: {a,b,c,str1,str2}	Yes

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 []>];
	
	struct interface_detail {
	string inf_name;
	string inf_description;
	ipAddress inf_host;
	enum duplex {
	validValues = auto, full, half;
	};
	}myInterface, myInterfaceArray[];
		No
		Note If the struct variable is declared as an array, the variable is iterative.

wwn (Available only in Cisco DCNM Web Client)	Example: 20:01:00:08:02:11:05:03	No

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										

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
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			
ipAddress	IP address in IPv4 or IPv6 format	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
ipAddrList	You can have a list of IPv4, IPv6, or a combination of both types of addresses. Example 1: 172.23.1.97, 172.23.1.98, 172.23.1.105, 172.23.1.109 Example 2: [redacted] Example 3: 172.23.1.97, 172.23.1.98, [redacted], 172.23.1.254 **Note** Separate the addresses in the list using commas and not hyphens.	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	IPv4 or IPv6 Address (does not require prefix/subnet)												
ipv4Address	IPv4 address	Yes											
ipv4Prefix	IPv4 Address with Subnet	Yes											
ipv6Address	IPv6 address	Yes											
ipv6Prefix	IPv6 Address with prefix	Yes											
ipv6Subnet	IPv6 Address with Subnet	Yes											
ISISNetAddress	Example: 49.0000.0000.0000.0000												
long	Example: 100	Yes			Yes	Yes							
macAddress	MAC address												

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
string	literal string Example for string Regular expression: string [illegible]	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>; ….. } [structvar1] [, <structvar2>] [, <structvar3> []>];												
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
AutoPopulate	Text	Copies values from one field to another
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"	

Annotation Key	Valid Values	Description
IsFEXID	“true” or “false”	
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”	
IsShow	“true” or “false”	Displays or hides a field on the window
IsSiteId	“true” or “false”	
IsSourceDevice	“true” or “false”	
IsSourceFabric	“true” or “false”	
IsSourceInterface	“true” or “false”	

Annotation Key	Valid Values	Description
IsSourceSwitchName	“true” or “false”	
IsSwitchName	“true” or “false”	
IsRMID	“true” or “false”	
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
PeerOneFEXID	“true” or “false”	
PeerTwoFEXID	“true” or “false”	
PeerOnePCID	“true” or “false”	
PeerTwoPCID	“true” or “false”	
PrimaryAssociation		
ReadOnly	“true” or “false”	Makes the field read-only
ReadOnlyOnEdit	“true” or “false”	
SecondaryAssociation	Text	
Section		
UsePool	“true” or “false”	
UseDNSReverseLookup		
Username	Text	Displays the username field on the window
Warning	Text	Provides text to override the Description annotation

Example: AutoPopulate Annotation

```
##template variables
string BGP_AS;
 @(AutoPopulate="BGP_AS")
```

```
  string SITE_ID;
##
```

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
```

Example: Warning Annotation

```
##template variables
@(Warning="This is a warning msg")
  string SITE_ID;
##
```

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]**

Template Content Editor

The template content editor has the following features:

- Syntax highlighting: The editor highlights the syntax, like different types of statements, keywords, and so on, for Python scripting.
- Autocompletion: The editor suggests the template datatypes, annotations, or metaproperties when you start typing.
- Go to line: You can navigate to the exact line in the template content editor instead of scrolling. Press **Command-L** in Mac or **Ctrl-L** in Windows, and enter the line number to which you want to navigate to in the pop-up window.

  If you enter a value greater than the number of lines in the editor, you will be navigated to the last line in the editor window.

- Template search and replace: Press **Command-F** in Mac or **Ctrl-F** in Windows, enter the search term in the **Search for** field, and select the type of search in the search window. You can perform the following searches in the editor:
  - **RegExp Search**: You can perform the regular expression search in the editor.
  - **CaseSensitive Search**: You can perform a case-sensitive search in the editor.
  - **Whole Word Search**: You can perform a whole word search to find the exact words in the editor. For example, a regular search for the word "play" returns results where it is part of words like "display," but the whole word search returns results only when there is an exact match for the word "play".
  - **Search In Selection**: You can perform a search in the selected content. Select the content to which you want to limit the search and enter the search term.

  Choose the + icon in the search window to use the replace option. Enter the replacing word in the **Replace with** field. You can replace the selected word once by selecting **Replace**. To replace all the occurrences of the selected word, select **All**.
- Code folding: You can expand or group code blocks in the editor by clicking the arrow next to their line numbers.
- Other features: The editor automatically indents the code, the closing braces, and highlights the matching parenthesis.

Template Editor Settings

You can edit the following features of a template editor by clicking **Template Editor Settings**.

- **Theme**: Select the required theme for the editor from the drop-down list.
- **KeyBinding**: Select the editor mode from the **KeyBinding** drop-down list to customize the editor. **Vim** and **Ace** modes are supported. The default is **Ace**.
- **Font Size**: Select the required font size for the editor.

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.

The Template Properties window appears.

Step 3 Specify a template name, description, tags, and supported platforms for the new template.

Step 4 Specify a **Template Type** for the template.

Step 5 Select a **Template Sub Type** and **Template Content Type** for the template.

Step 6 Click the **Advanced** tab to edit other properties like **Implements**, **Dependencies**, **Published**, and **Imports**. Select **Published** to make the template read-only. You cannot edit a published template.

Step 7 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 8 Click **OK** to save the template properties, or click the cancel icon at the top-right corner of the window to revert the changes.

Note You can edit the template properties by clicking **Template Property**.

Step 9 Click **Template Content** to edit the template syntax. For information about the structure of the Configuration Template, see the *Template Structure* section.

Step 10 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** by clicking the error and warnings field.

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. Click the line number under the Start Line column to locate the error in the template content. You will get an error if you validate a template that does not have a template name.

Step 11 Click **Save** to save the template.

Step 12 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 232.

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

Upgrading your devices to the latest software version manually might take a long time and prone to error, which requires a separate maintenance window. To ensure rapid and reliable software upgrades, image management automates the steps associated with upgrade planning, scheduling, downloading, and monitoring. Image management is supported only for Cisco Nexus switches.

Note Before you upgrade, ensure that the POAP boot mode is disabled for Cisco Nexus 9000 Series switches and Cisco Nexus 3000 Series switches. To disable POAP, run the `no boot poap enable` command on the switch console. You can however, enable it after the upgrade.

The **Image Management** menu includes the following submenu:

This feature allows you to upload or delete images that are used during POAP and switch upgrade. To view the window from the Cisco DCNM Web UI homepage, choose .

You can view the following details in the window.

Deleting an Image

To delete an image from the repository from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose .

The window appears.

Step 2 Choose an existing image from the list and click the **Delete Image** icon.

A confirmation window appears.

Step 3 Click **Yes** to delete the image.

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 or image upgrade.

Your user role should be **network-admin** to upload an image. You can't perform this operation with the **network-stager** user role.

Procedure

Step 1 Choose .

The window appears.

Step 2 Click **Image Upload**.

The **Select File to Upload** dialog box appears.

Step 3 Click **Choose file** to choose a file from the local repository of your device.

Step 4 Choose the file and click **Upload**.

Step 5 Click **OK**.

The upload takes some time depending on the file size and network bandwidth.

Note You can upload images for all Cisco Nexus Series Switches.

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.

Field	Description
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 **Note** If the job fails on a single or multiple devices, the status field shows COMPLETED WITH EXCEPTION indicating a failure.
Created Time	Specifies the time when the task was created.
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.
Completed Time	Specifies the time when the task was completed.
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:

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 appears.

Step 3 Click **Settings**. Expand the **Columns** menu and choose the details you want to view.

You can view the following information in this window:

- Location of the kickstart and system images
- Compatibility check status
- Installation status
- Descriptions
- Logs

Step 4 Select the device.

The detailed status of the task appears. For the completed tasks, the response from the device appears.

If the upgrade task is in progress, a live log of the installation process appears.

Note
- This table autorefreshes every 30 secs for jobs in progress, when you're 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.

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**.

Step 2 Choose **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 3 Select the check box to the left of the switch name.

You can select more than one switch and move the switches to the right column.

Step 4 Click **Add** or **Remove** icons to include the appropriate switches for upgrade.

The selected switches appear in a column on the right.

Step 5 Click **Next**.

The **Specify Software Images** window appears. This tab displays the switches that you selected in the previous screen. You can choose the images for upgrade as well.

- 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 the **Image Upload** window.
- The **Path** field is disabled, and the default image path is used.

Step 6 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's an issue in viewing the **Software Image Browser** dialog box, reduce the font size of your browser and retry.

Step 7 Click **Select Image** in the **System Image** column.

The **Software Image Browser** dialog box appears.

Step 8 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 Default_SCP_Repository 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.

Note Only files with BIN extension are listed if you select **File Server**. To view other files, choose **Administration > DCNM Server > Server Properties**, set **FILE_SELECTION_FILTER** to **false**, and restart the server. It is set to **true** by default.

Note Only image files present in the **Image Upload** window can be selected. You can't select images present in any other paths.

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.

Note Only files with BIN extension are listed if you select **Switch File System**. To view other files, choose **Administration > DCNM Server > Server Properties**, set **FILE_SELECTION_FILTER** to **false**, and restart the server. It is set to **true** by default.

b) Click **OK** to choose the kickstart image or **Cancel** to revert to the **Specify Software Images** dialog box.

Step 9 The **Vrf** column indicates the name of the virtual routing and forwarding (VRF).

Step 10 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's 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 11 **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 12 Drag and drop the switches to reorder the upgrade task sequence.

Step 13 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 14 Select **Select Parallel Line Card upgrade** to upgrade all the line cards at the same time.

Upgrading a parallel line card isn't applicable for Cisco MDS devices.

Step 15 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**, **Upgrade Option 2** is disabled.

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 16 Click **Next**.

If you didn't 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 17 Click **Finish Installation Later** to perform the upgrade later.

Step 18 Click **Next**.

Step 19 Check the check box to save the running configuration to the startup configuration before upgrading the device.

Step 20 You can schedule the upgrade process to occur immediately or later.

a. Select **Deploy Now** to upgrade the device immediately.

b. 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 21 You can choose the execution mode based on the devices and the line cards you have chosen to upgrade.

a. Select **Sequential** to upgrade the devices in the order you chose them.

b. Select **Concurrent** to upgrade all the devices at the same time.

Step 22 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. DCNM discovers 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.

a. Select **Deploy Now** to upgrade the device immediately.

b. 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.

a. Select **Sequential** to upgrade the devices in the order in which they were chosen.

b. Select **Concurrent** to upgrade the devices at the same time.

Step 7 Click **Finish** to complete the upgrade process.

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
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 (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.

Information about the Endpoint Locator is displayed on a single landing page or dashboard . The dashboard displays an almost real-time view of data (refreshed every 30 seconds) pertaining to all the active endpoints on a single pane. The data that is displayed on this landing page depends on the scope selected by you from the **SCOPE** drop-down list.

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 (IPv4 and\or IPv6) and MAC address. 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 or Route Servers
- 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 iBGP and eBGP based VXLAN EVPN fabrics. From Release 11.2(1), the fabrics may be created as Easy Fabrics or External Fabrics. EPL can be enabled with an option to automatically configure the spine or RRs with the appropriate BGP configuration (new in DCNM 11.2).
- 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.
- 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 three interfaces:

- eth0 interface for external access
- eth1 interface for fabric management (Out-of-band or OOB)
- eth2 interface for in-band network connectivity

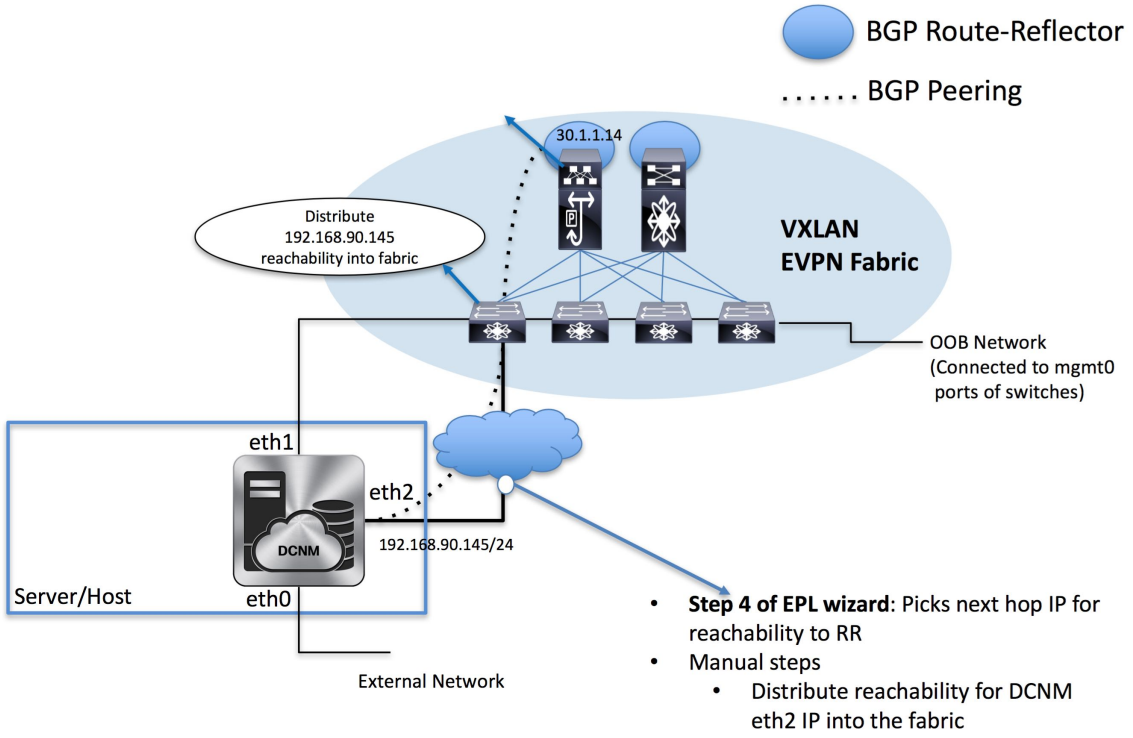

The eth1 interface provides reachability to the devices via the mgmt0 interface either Layer-2 or Layer-3 adjacent. This allows DCNM to manage and monitor these devices including POAP. EPL requires BGP peering between the DCNM and the Route-Reflector. Since the BGP process on Nexus devices typically runs on the default VRF, in-band IP connectivity from the DCNM to the fabric is required. For this purpose, the eth2 interface can be configured using the **appmgr setup inband** command. Optionally, you can configure the eth2 interface during the Cisco DCNM installation.

If you need to modify the already configured in-band network (eth2 interface), execute the **ifconfig eth2 0.0.0.0** command and run the **appmgr setup inband** command again. Refer Editing Network Properties Post DCNM Installation to run the **appmgr setup inband** command.

Note The setup of eth2 interface on the DCNM is a prerequisite of any application that requires the in-band connectivity to the devices within fabric. This includes EPL and Network Insights Resources (NIR).

Note For configuring EPL in standalone mode, you must add a single neighbor to EPL. DCNM eth2 IP address is EPL IP.

On the fabric side, for a standalone DCNM deployment, if the DCNM eth2 port is directly connected to one of the front-end interfaces on a leaf, then that interface can be configured using the **epl_routed_intf** template.

An example scenario of how this can be done when IS-IS or OSPF is employed as the IGP in the fabric, is depicted below:

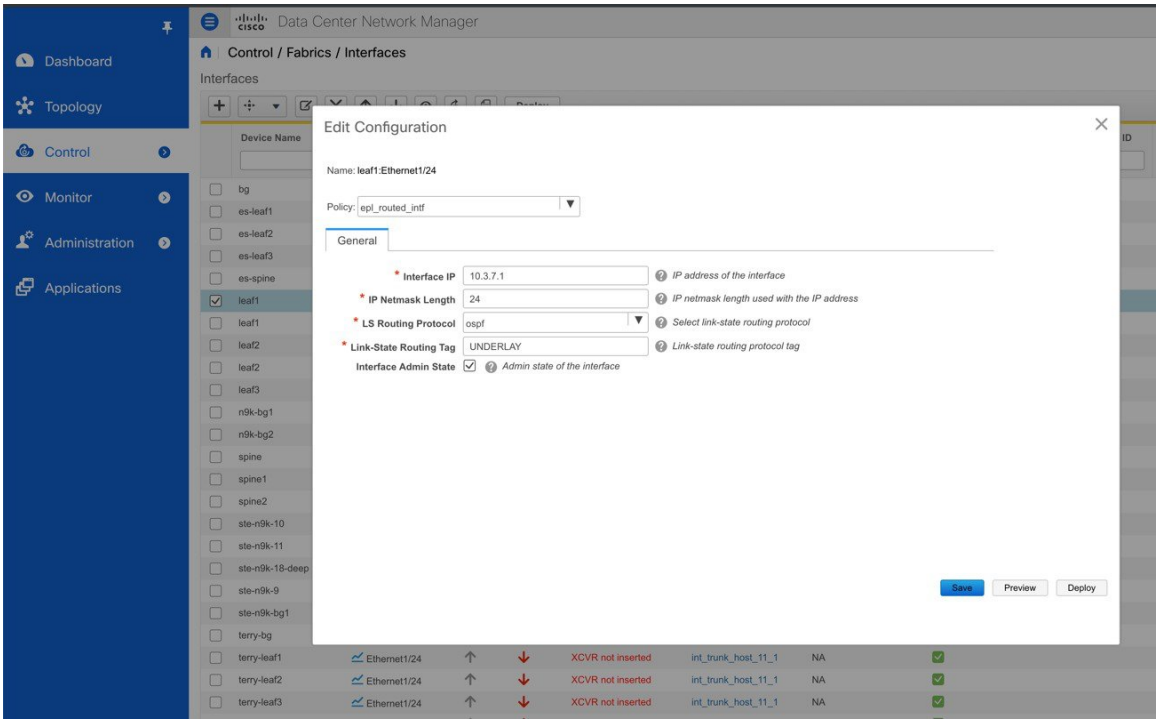

However, for redundancy purposes, it is always advisable to have the server on which the DCNM is installed to be dual-homed or dual-attached. With the OVA DCNM deployment, the server can be connected to the switches via a port-channel. This provides link-level redundancy. To also have node-level redundancy on the network side, the server may be attached to a vPC pair of Leaf switches. In this scenario, the switches must be configured such that the HSRP VIP serves as the default gateway of the eth2 interface on the DCNM. The following image depicts an example scenario configuration:

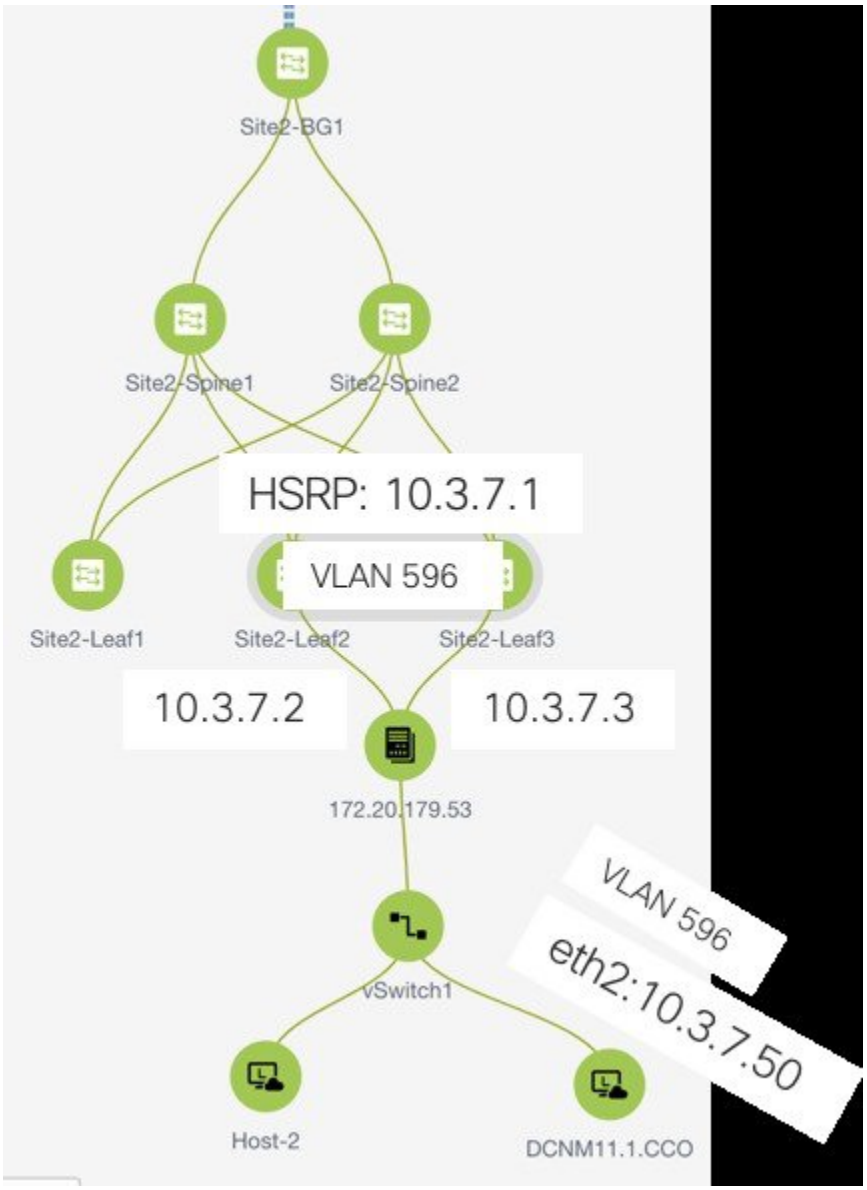

In this example, the server with the DCNM VM is dual-attached to a vPC pair of switches that are named Site2-Leaf2 and Site2-Leaf3 respectively. VLAN 596 associated with the IP subnet 10.3.7.0/24 is employed for in-band connectivity. You can configure the vPC host port toward the server using the **interface vpc trunk host** policy as shown is the following image:

For the HSRP configuration on Site2-Leaf2, the **switch_freeform** policy may be employed as shown in the following image:

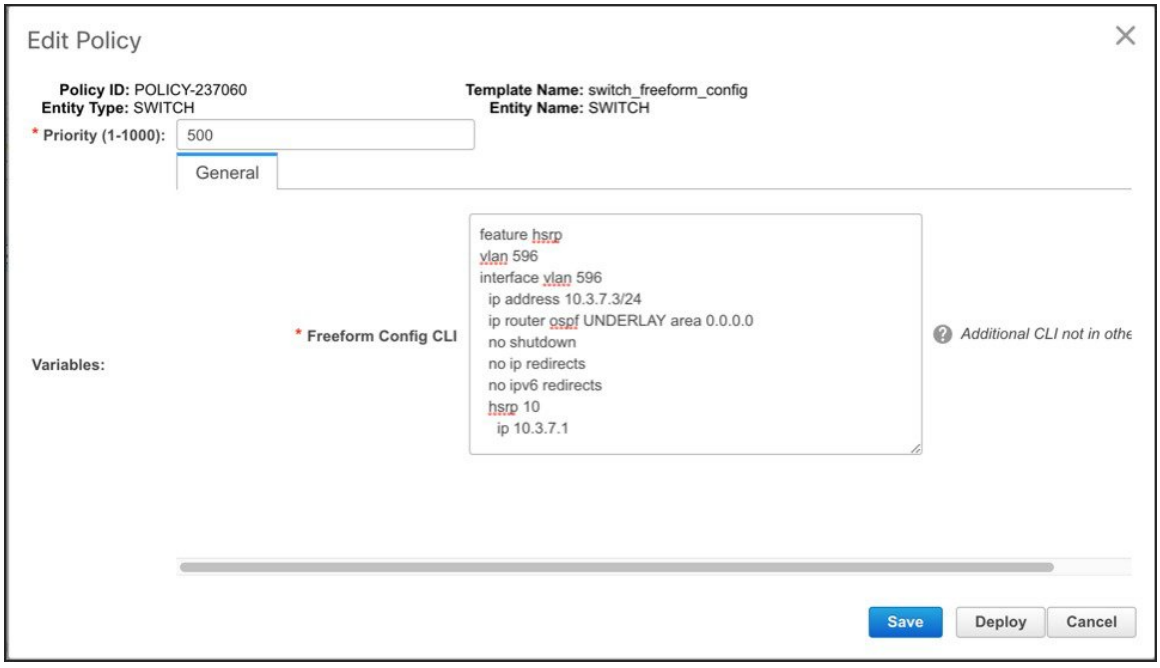

You can deploy a similar configuration on Site2-Leaf3 while using IP address 10.3.7.2/24 for SVI 596. This establishes an in-band connectivity from the DCNM to the fabrics over the eth2 interface with the default gateway set to 10.3.7.1.

After you establish the in-band connectivity between the physical or virtual DCNM and the fabric, you can establish BGP peering. There is a simple wizard for enabling Endpoint Locator.

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 on the spines/RRs via the eth2 gateway.

Note Cisco DCNM queries the BGP RR to glean information for establishment of the peering, like ASN, RR, IP, and so on.

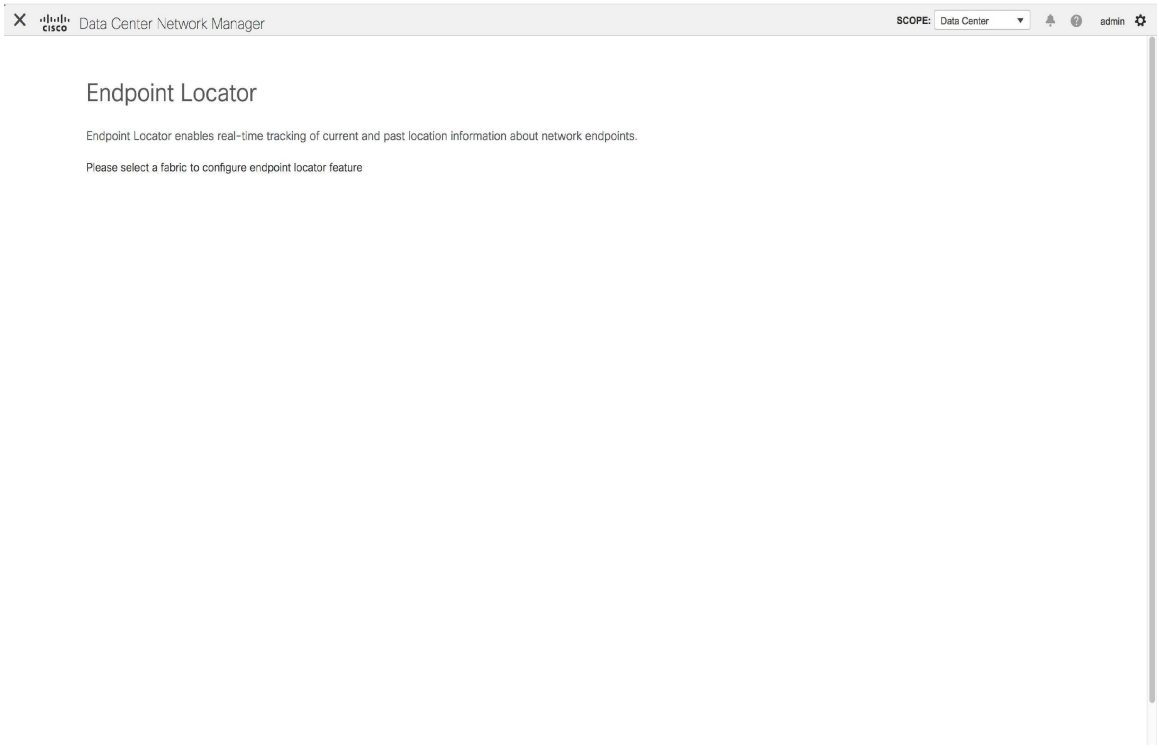

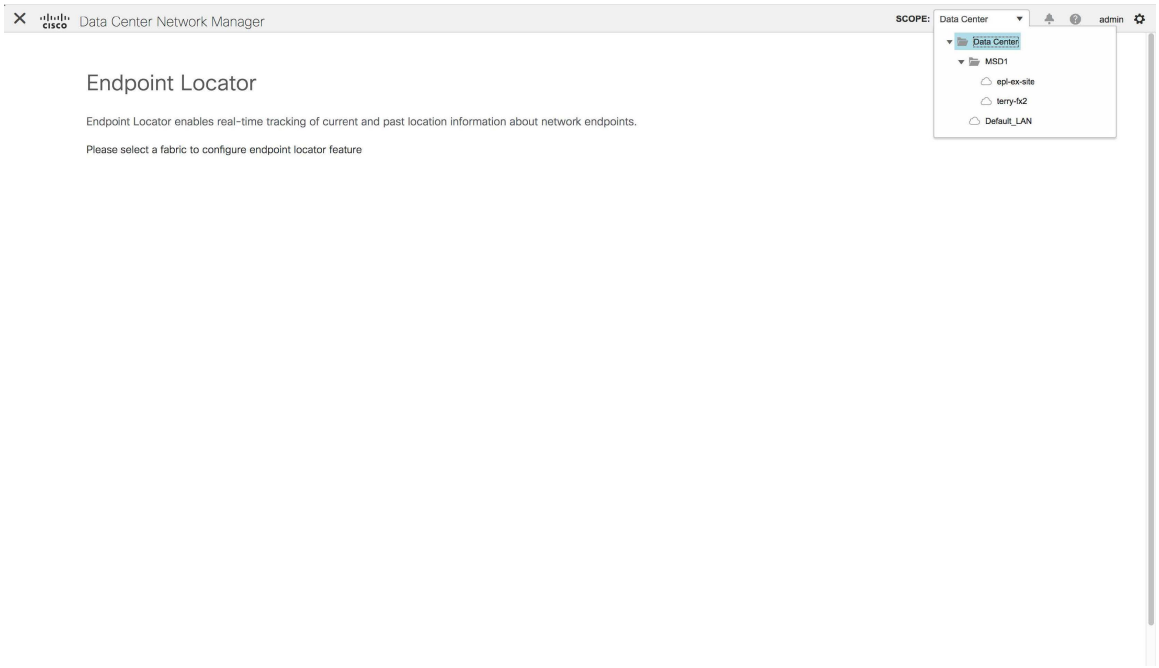

Fore more information about the EPL dashboard, refer Monitoring Endpoint Locator.

To configure Endpoint Locator from Cisco Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Endpoint Locator > Configure**.

The **Endpoint Locator** window appears, with a **See how it works** help link.

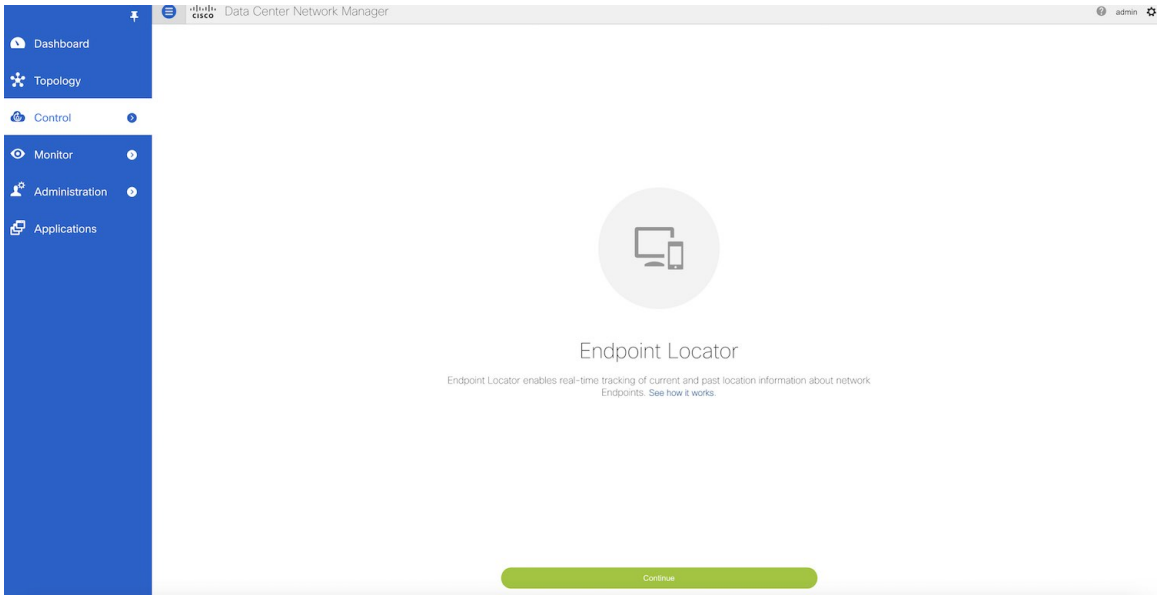

Step 2 Click **Continue**.

Step 3 Select the appropriate fabric on which the endpoint locator feature should be enabled to track endpoint activity.

You can enable EPL for one fabric. It can be DFA or EVPN.

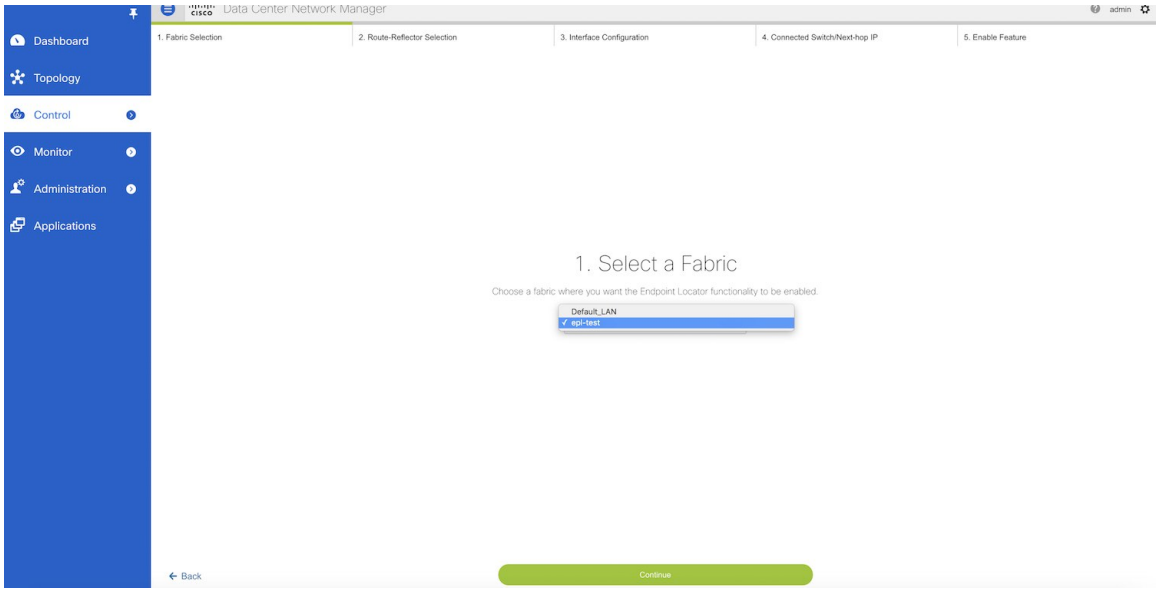

Step 4 Select the switches on the fabric hosting the RRs. Cisco DCNM will peer with the RRs.

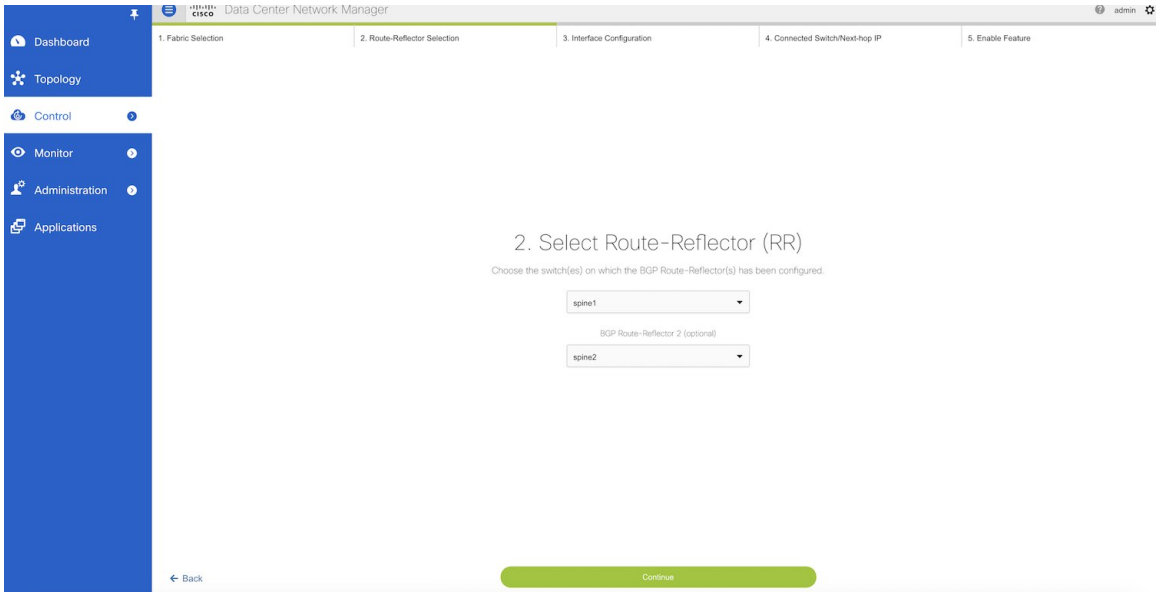

Step 5 Check the Cisco DCNM eth2 configuration for IP reachability to the RR.

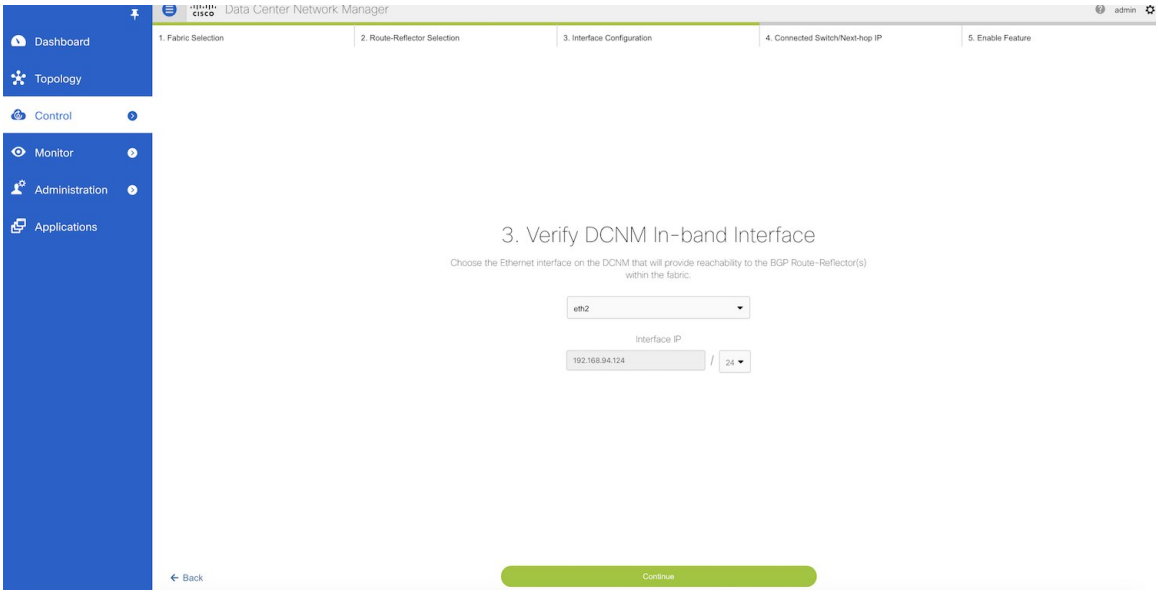

Step 6 By default, the “Configure my fabric” option is selected. This knob controls whether BGP configuration will be pushed to the selected spines/RRs as part of the enablement of the EPL feature. If the spine/RR needs to be configured manually with a custom policy for the EPL BGP neighborship, then this option should be unchecked. Check Next-hop IP and ensure the eth2 gateway IP is correct. If there is an error go to command line and reconfigure the eth2 port using the **appmgr setup inband** command. To flush the prior eth2 configuration, perform **ifconfig eth2 0.0.0.0** before proceeding with the **appmgr setup inband** command.

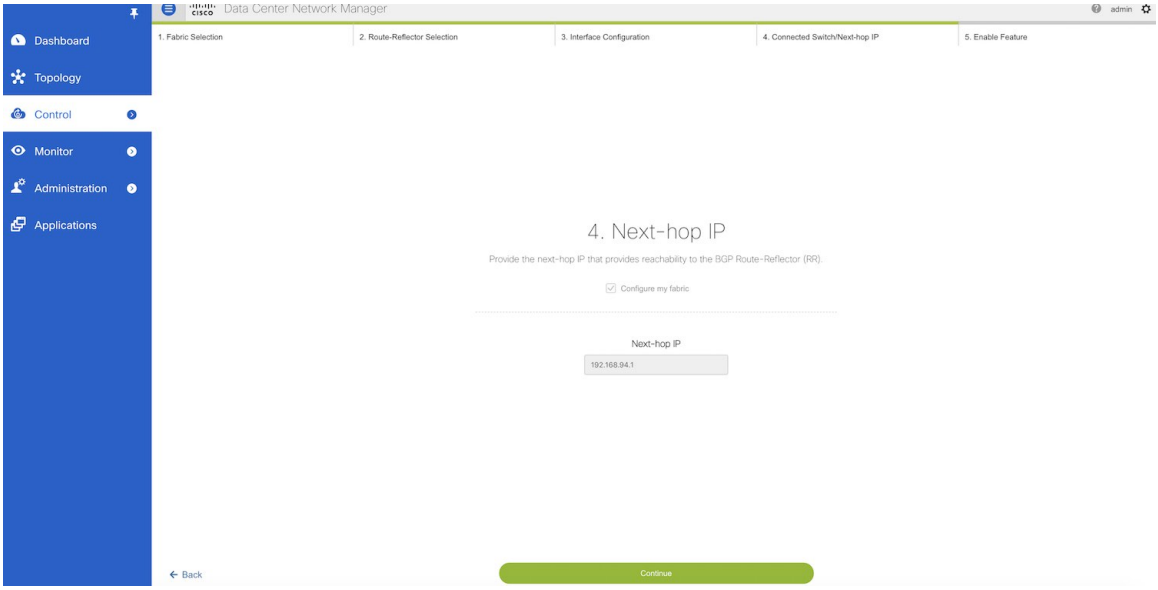

Step 7 The last step provides a summary of the information entered in the previous steps. The wizard view allows navigation to any particular step whereby one can make necessary changes/edits. In this step, one must specify whether additional information such as PORT, VLAN, VRF etc. is required when enabling the EPL feature. If the **No** option is selected, then this information will not be collected and reported by EPL.

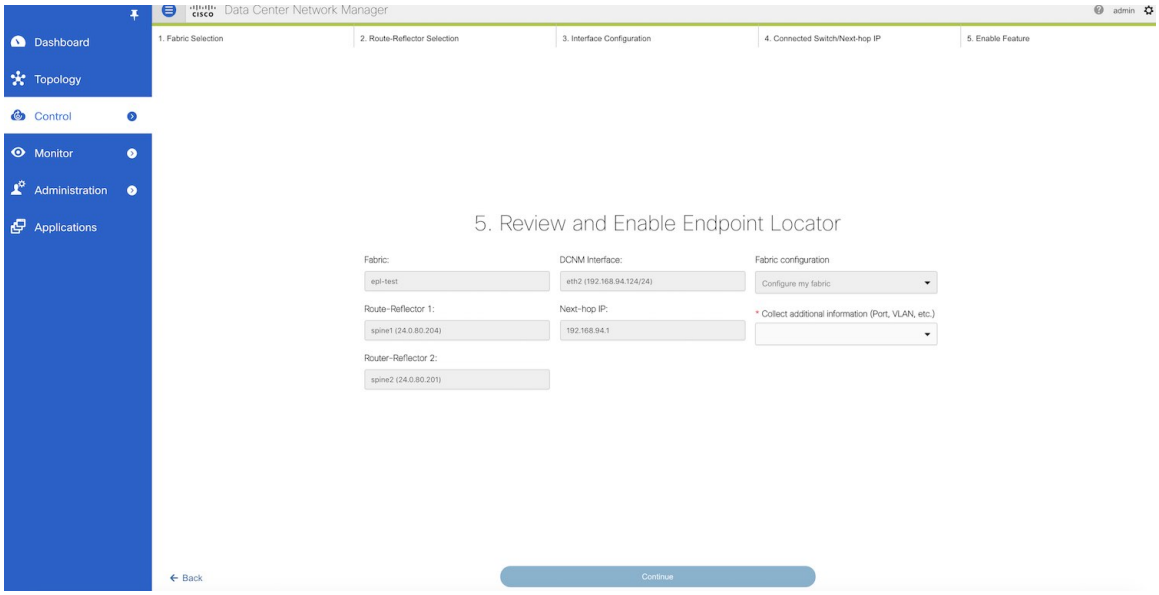

However, if the **Yes** option is selected in the drop-down, a warning pop-up appears that feature NX-API must be supported and enabled on the switches, ToRs, and leafs to gather this information. Otherwise, you cannot fetch or report this additional information.

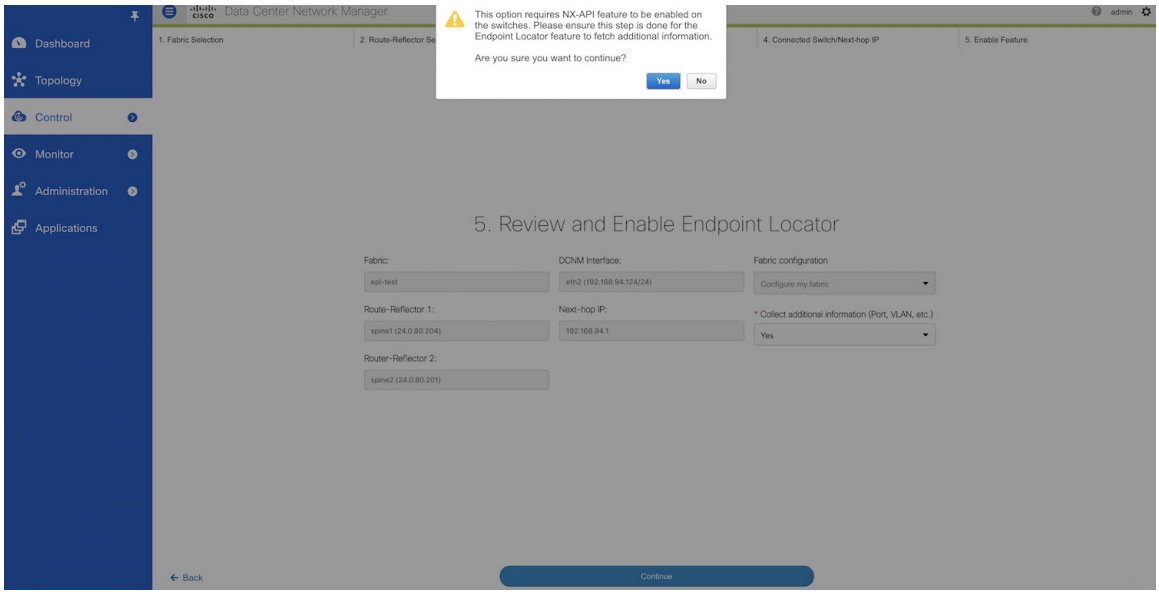

Step 8 Once the appropriate selections are made and various inputs have been reviewed, click **Continue** to enable EPL. A progress bar will appear indicating the status of the EPL feature enable process.

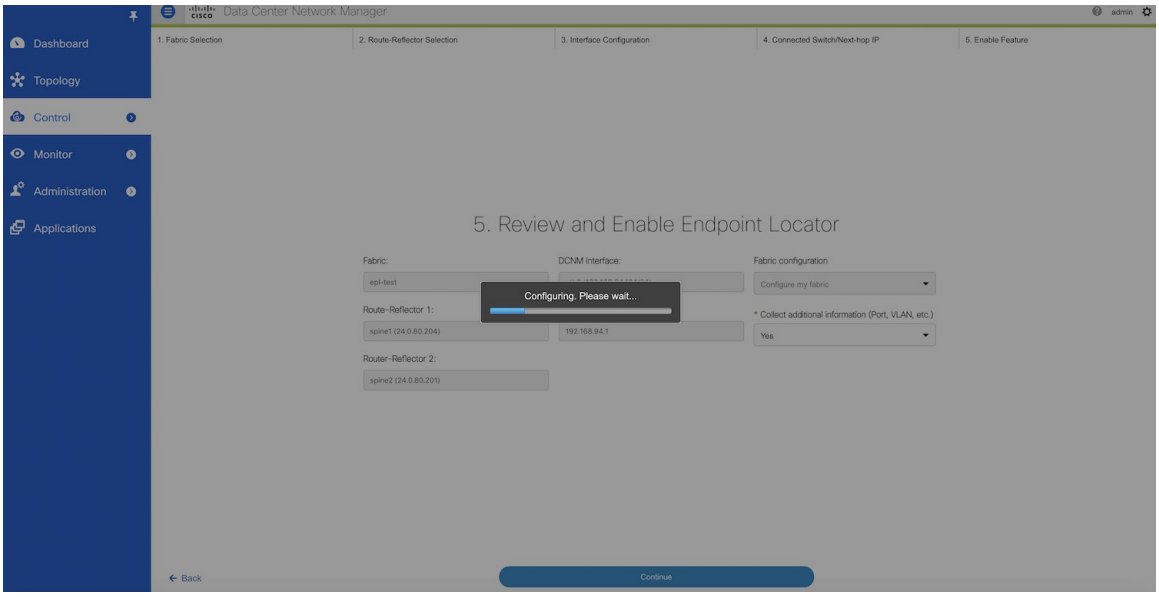

If there are any errors while you enable EPL, the enable process aborts and the appropriate error message is displayed. Otherwise, EPL is successfully enabled and on clicking **OK**, the screen is automatically redirected to the EPL dashboard.

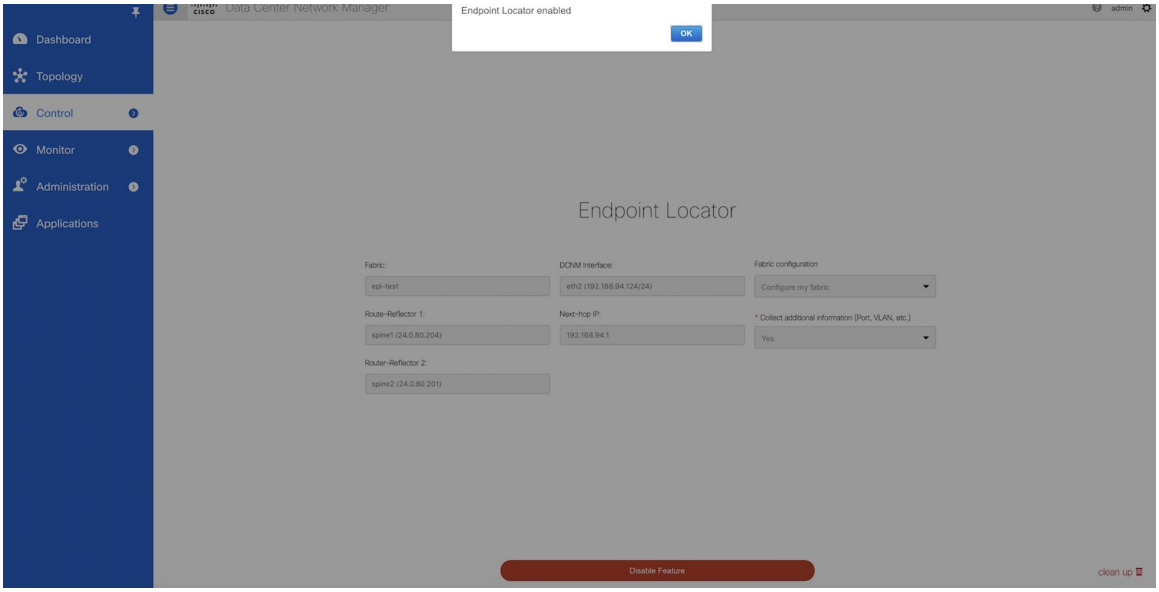

When the Endpoint Locator feature is enabled, there are a number of steps that occur in the background. DCNM contacts the selected RRs and determines the ASN. It also determines the interface IP that is bound to the BGP process. Also, appropriate BGP neighbor statements are added on the RRs or spines in case of eBGP underlay, to get them ready to accept the BGP connection that will be initiated from the DCNM. The neighbor address is the same as that of the eth2 interface shown in step 2. For the native HA DCNM deployment, both the primary and secondary DCNM eth2 interface IPs will be added as BGP neighbors but only one of them will be active at any given time. 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.

What to do next

For more information about monitoring EPL, see Monitoring Endpoint Locator.

Flushing the Endpoint Database

After you enable the Endpoint Locator feature, you can clean up or flush all the Endpoint information. This allows starting from a clean-slate with respect to ensuring no stale information about any endpoint is present in the database. After the database is clean, the BGP client re-populates all the endpoint information learnt from the BGP RR.

To flush all the Endpoint Locator information from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control** > **Endpoint Locator** > **Configure**, and click **clean up** link.

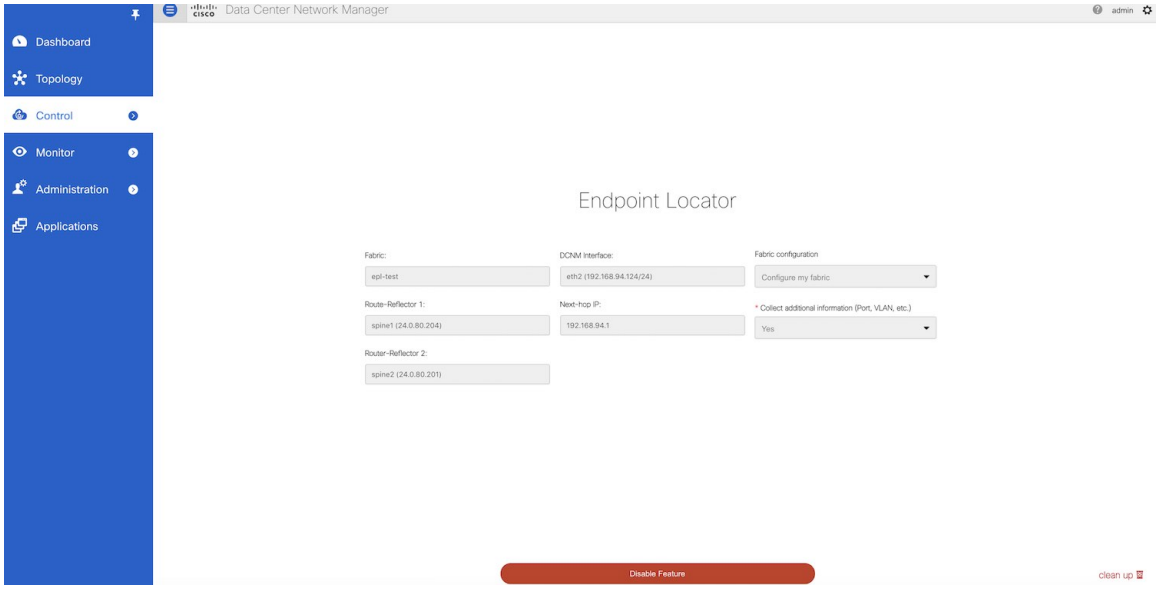

A warning is displayed with a message indicating that all the endpoint information that is stored in the database will be flushed.

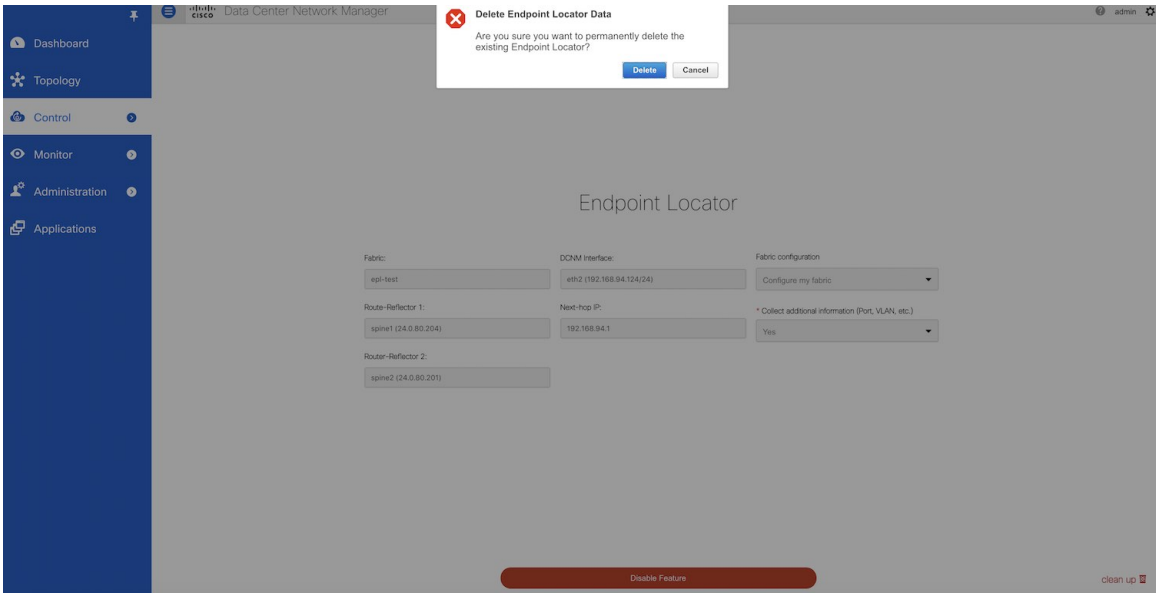

Step 2 Click **Delete** to continue or **Cancel** to abort.

Configuring Endpoint Locator in DCNM Cluster Mode

Note For configuring EPL in cluster mode, you must add a single neighbor to EPL. DCNM EPL container Inband IP address is EPL IP.

With the DCNM cluster mode deployment, in addition to the DCNM nodes, an additional 3 compute nodes are present in the deployment. For information about deploying applications in cluster mode, see *Cisco DCNM in Clustered Mode*.

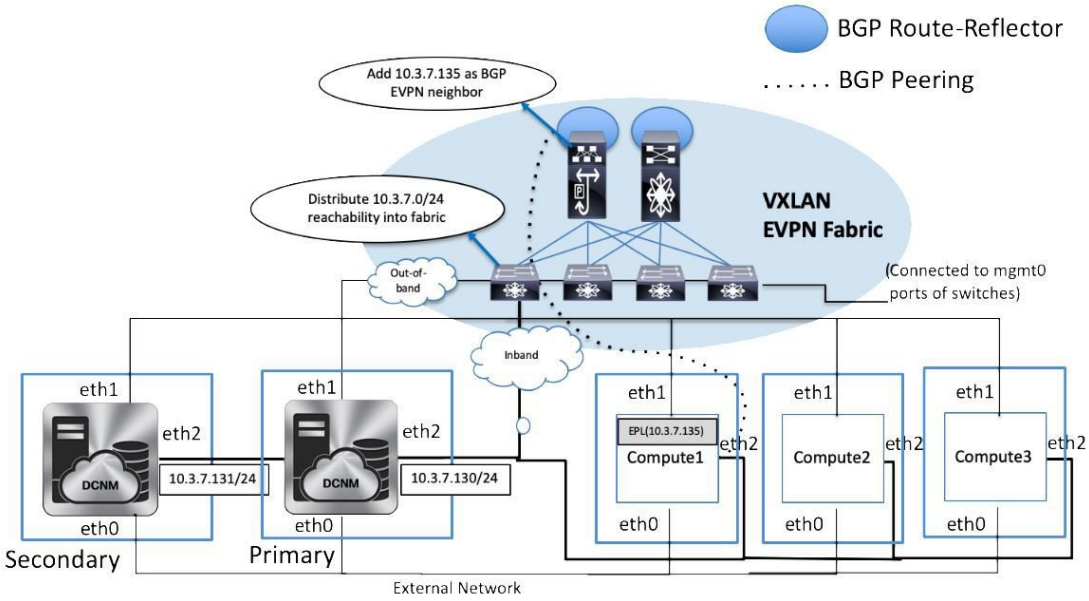

In DCNM Cluster mode, all applications including EPL run on the compute nodes. The DCNM application framework takes care of the complete life cycle management of all applications that run on the compute nodes. The EPL instance runs as a container that has its own IP address allocated out of the inband pool assigned to the compute nodes. This IP address will be in the same IP subnet as the one allocated to the eth2 or inband interface. Using this IP address, the EPL instance forms a BGP peering with the spines/RRs when the EPL feature is enabled. If a compute node hosting the EPL instance will go down, the EPL instance will be automatically respawned on one of the remaining 2 compute nodes. All IP addresses and other properties associated with the EPL instance are retained.

The Layer-2 adjacency requirement of the compute nodes dictates that the compute node eth2 interfaces should be part of the same IP subnet as the DCNM nodes. Again, in this case, connecting the compute nodes to the same vPC pair of switches is the recommended deployment option. Note that for cluster mode DCNM OVA setups, ensure that promiscuous mode is enabled in the port group corresponding to eth2 interface in order to establish inband connectivity as depicted below:

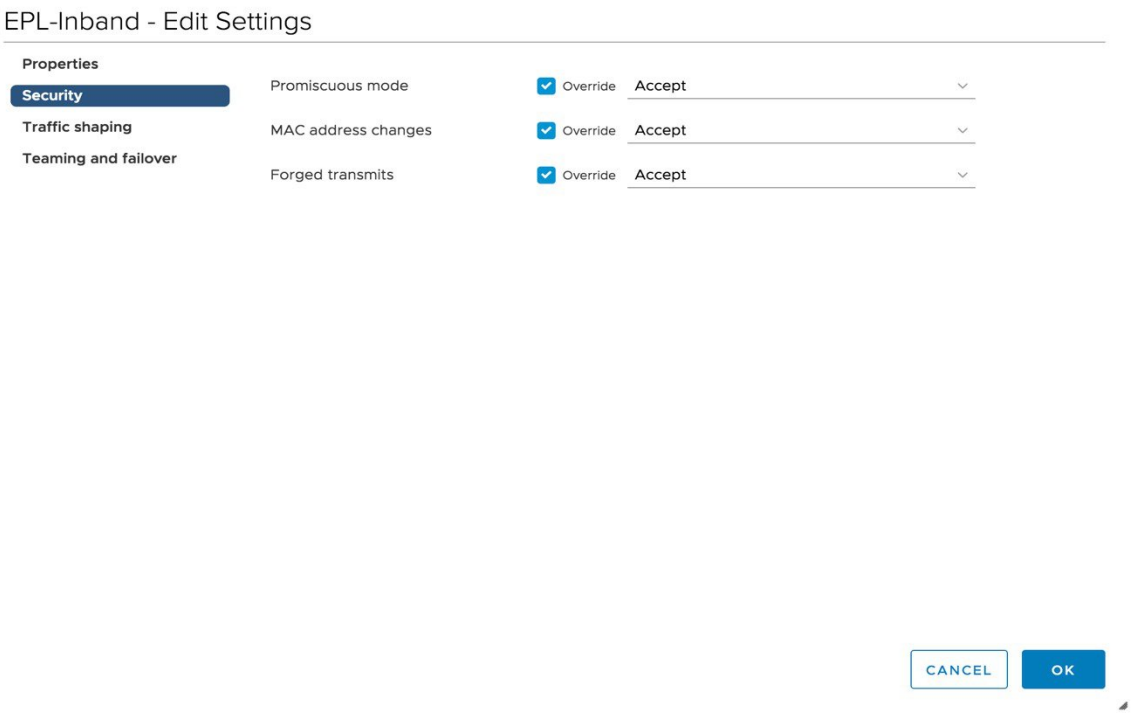

The enablement of the EPL feature for DCNM cluster mode is identical to that in the non-cluster mode. The main difference is that on the spine/RRs, only a single BGP neighborship is required that points to the IP address allocated to the EPL instance. Recall that for the DCNM native HA deployment in the non-cluster mode, all spines/RRs always had 2 configured BGP neighbors, one pointing to the DCNM primary eth2 interface and other one pointing to the DCNM secondary eth2 interface. However, only one neighbor would be active at any given time.

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 Monitor Mode** 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.

Configuring Endpoint Locator for eBGP EVPN Fabrics

From Cisco DCNM Release 11.2(1), you can enable EPL for VXLAN EVPN fabrics, where eBGP is employed as the underlay routing protocol. Note that with an eBGP EVPN fabric deployment, there is no traditional RR similar to iBGP. The reachability of the in-band subnet must be advertised to the spines that behave as Route Servers. To configure EPL for eBGP EVPN fabrics from the Cisco DCNM Web UI, perform the following steps:

Procedure

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

Select the fabric to configure eBGP on or create eBGP fabric with the **Easy_Fabric_eBGP** template.

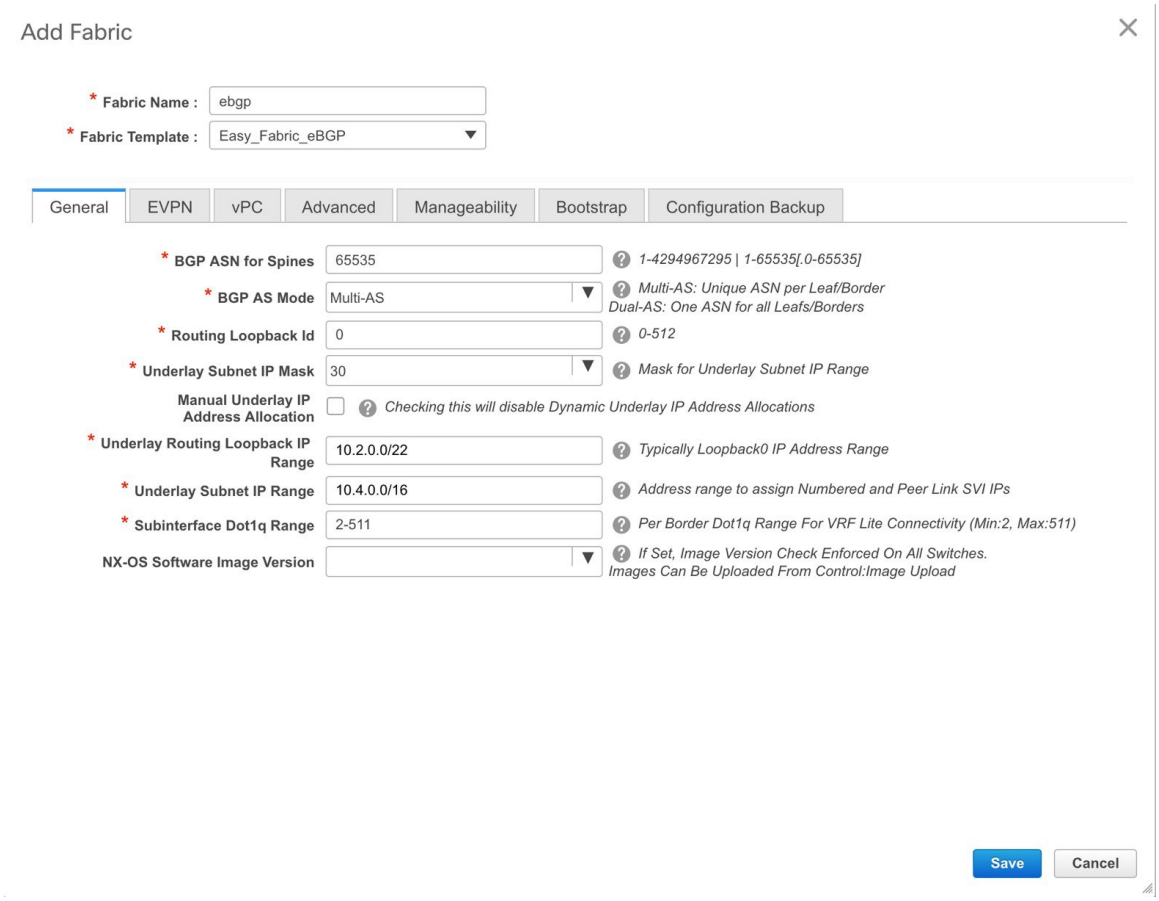

Step 2 Use the **leaf_bgp_asn** policy to configure unique ASNs on all leaves.

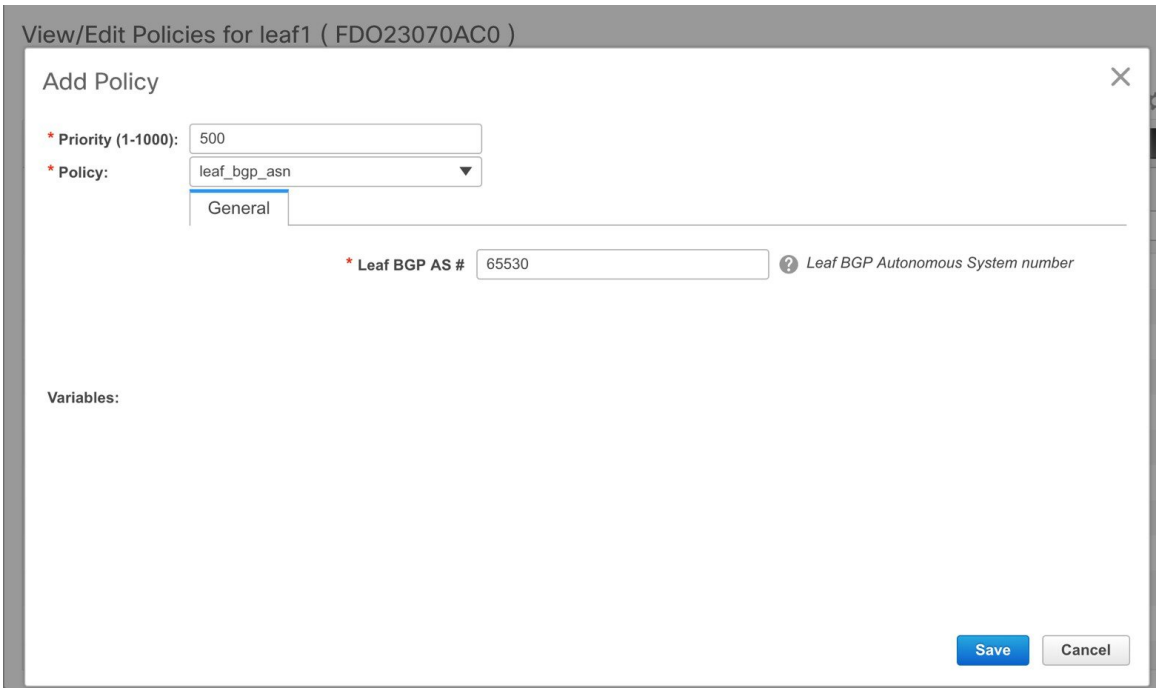

Step 3 Add the **ebgp_overlay_leaf_all_neighbor** policy to each leaf.

Fill **Spine IP List** with the spines' BGP interface IP addresses, typically the loopback0 IP addresses.

Fill **BGP Update-Source Interface** with the leaf's BGP interface, typically loopback0.

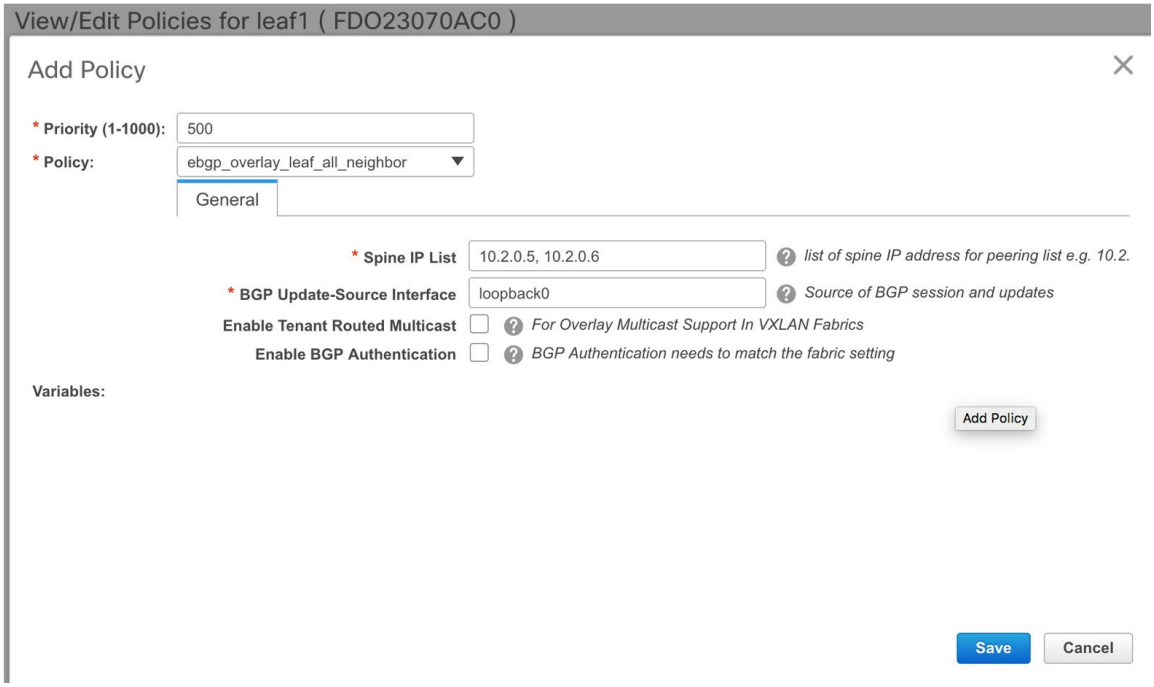

Step 4 Add the **ebgp_overlay_spine_all_neighbor** policy to each spine.

Fill **Leaf IP List** with the leaves' BGP interface IPs, typically the loopback0 IPs.

Fill **Leaf BGP ASN** with the leaves' ASNs in the same order as in **Leaf IP List**.

Fill **BGP Update-Source Interface** with the spine's BGP interface, typically loopback0.

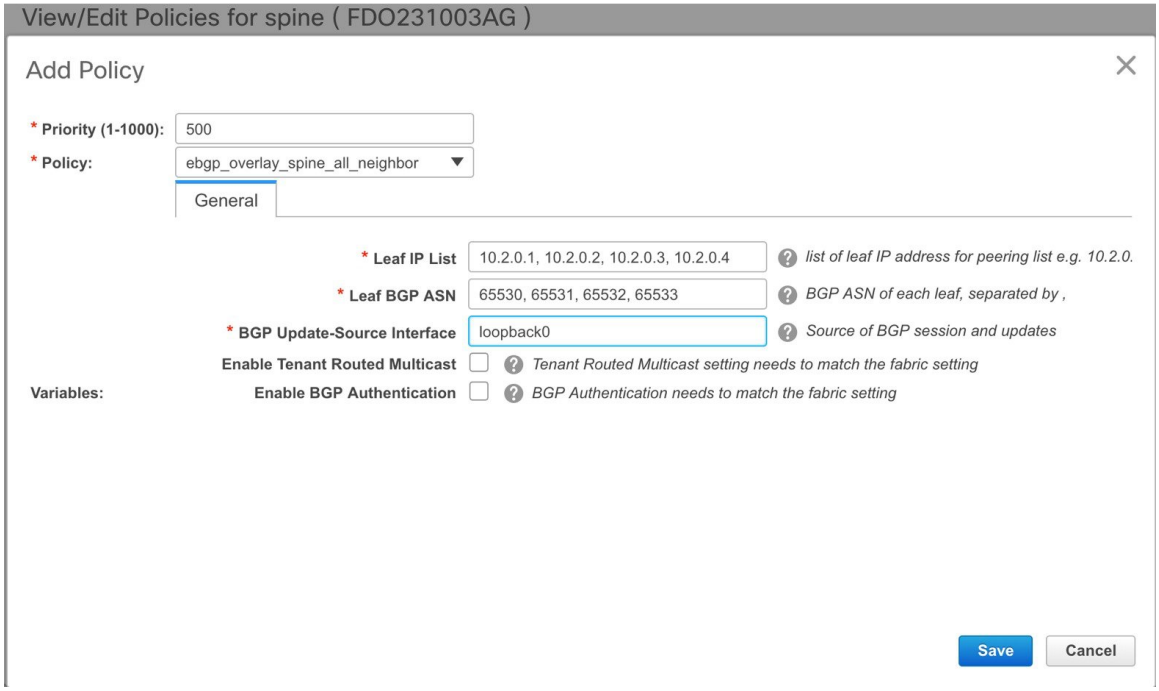

After the in-band connectivity is established, the enablement of the EPL feature remains identical to what is listed so far. EPL becomes a iBGP neighbor to the Route Servers running on the spines.

EPL Connectivity Options

Sample topologies for the varios EPL connectivity options are as given below.

Cisco DCNM supports the following web browsers:

DCNM Cluster Mode: Physical Server to VM Mapping

We recommend a minimum of 3 physical servers, or a maximum of 5 physical servers in which each DCNM and compute is located on an individual physical server.

Figure 1: A minimum of 3 physical servers

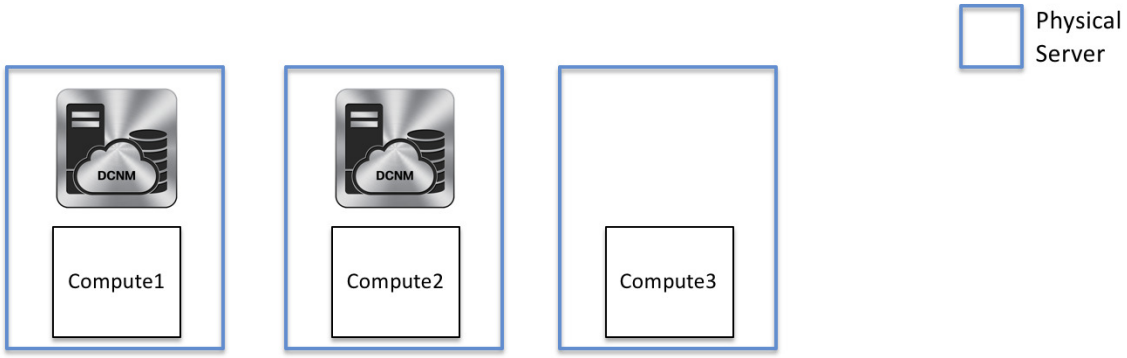

Figure 2: A maximum of 5 physical servers

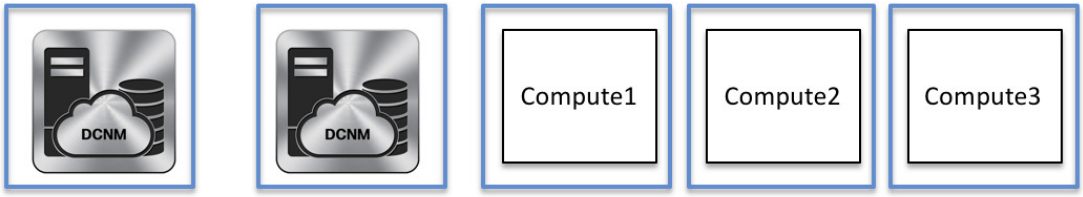

DCNM/Compute VM Physical Connectivity

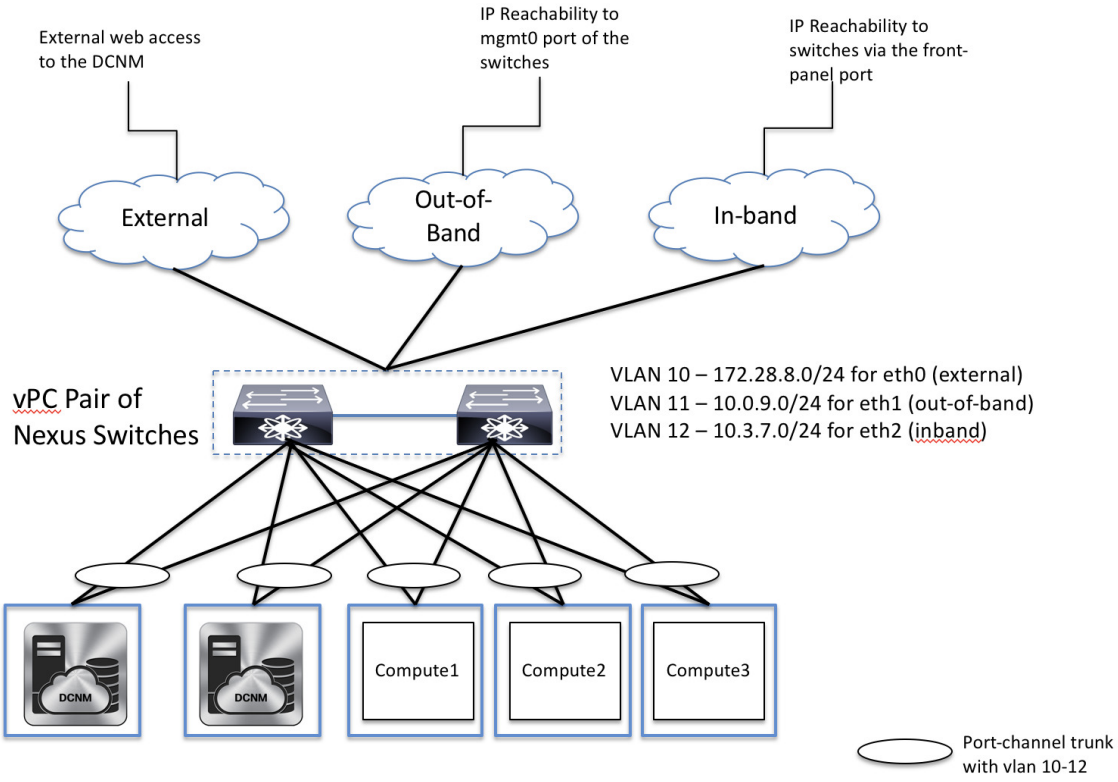

DCNM Cluster Mode

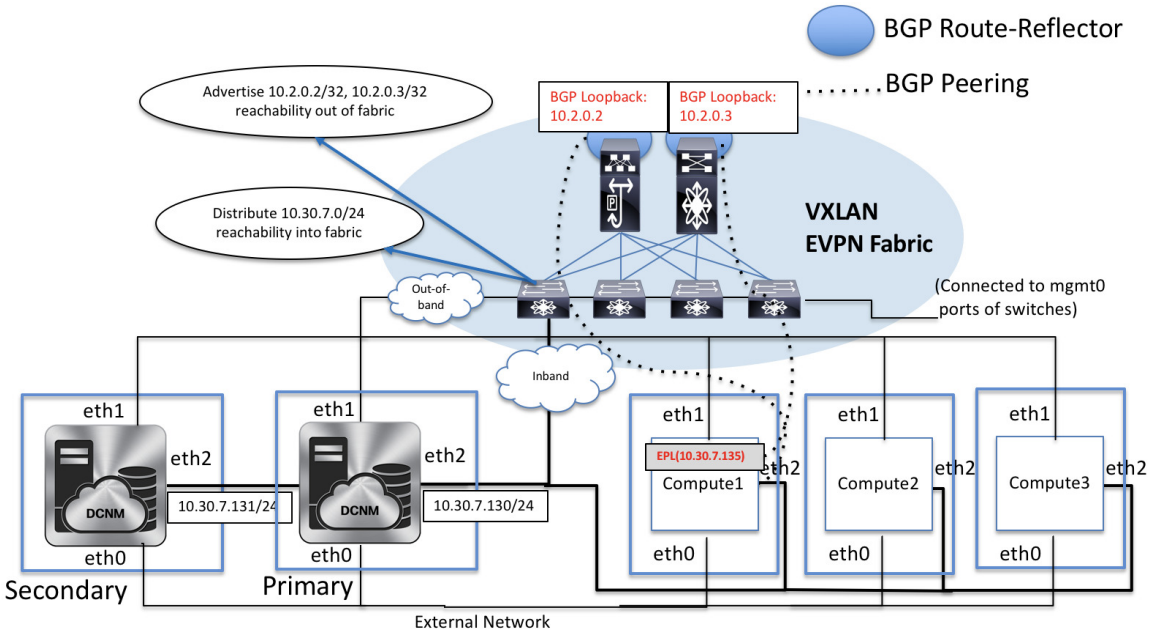

DCNM Multi-Fabric Connectivity

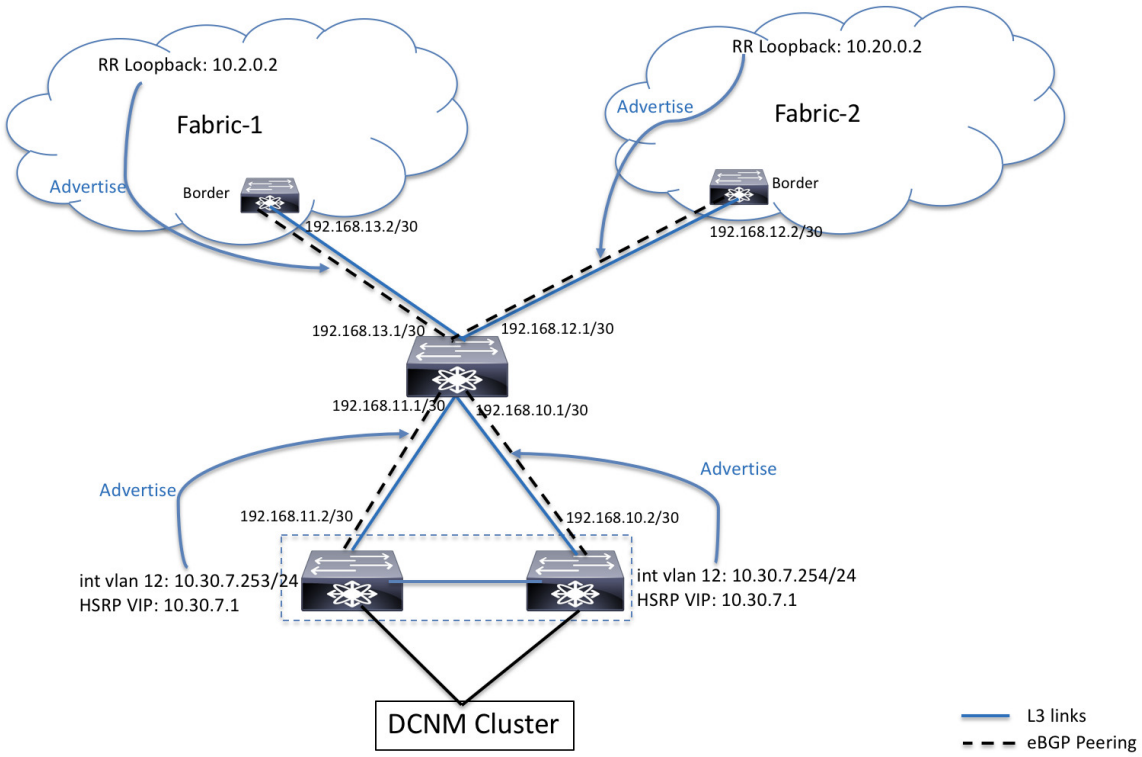

EPL Connectivity for Native HA

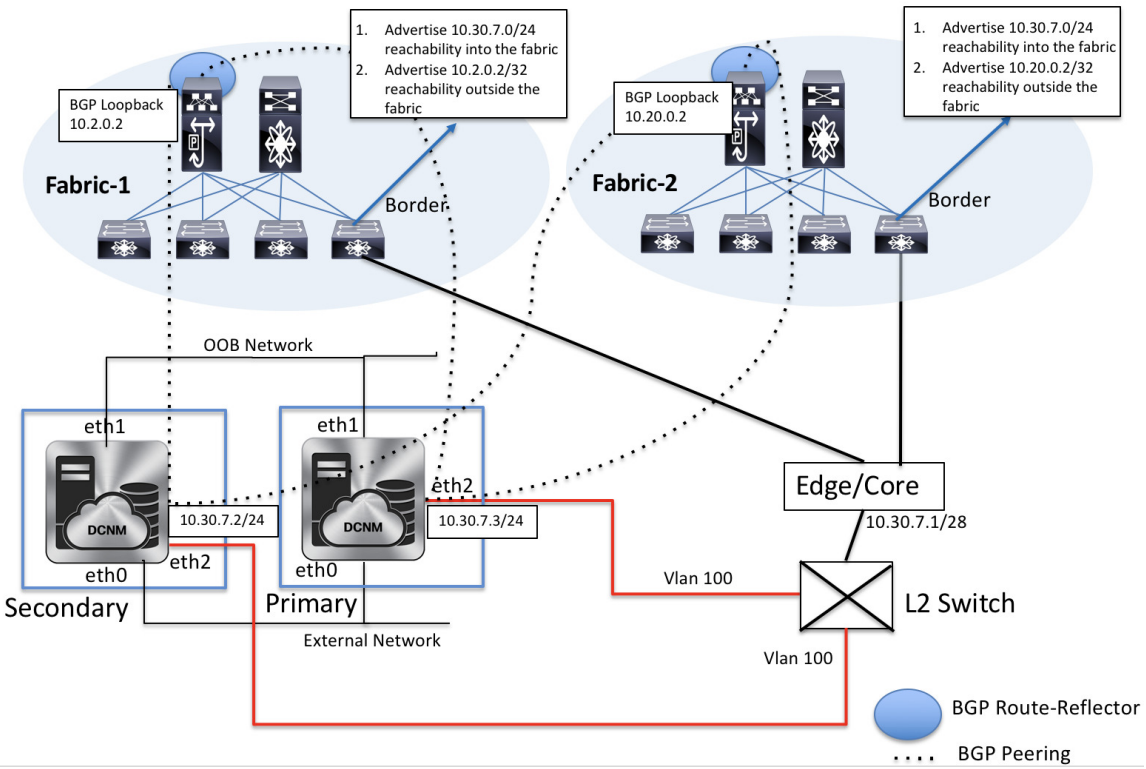

Disabling Endpoint Locator

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 /var/afw/applogs/ under the directory for the associated fabric. For example, if EPL is enabled for the **test** fabric, the logs will be in /var/afw/applogs/epl_cisco_test_afw_log/epl/ starting with filename afw_bgp.log.1. 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 `afw_bgp.log`. Up to 10 such files will be stored with each file size of maximum of 10MB.

Note EPL creates a symlink in this directory inside the docker container, hence it appears broken when accessed natively.

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.

Monitoring Endpoint Locator

Information about the Endpoint Locator is displayed on a single landing page or dashboard. The dashboard displays an almost real-time view of data (refreshed every 30 seconds) pertaining to all the active endpoints on a single pane. The data that is displayed on this dashboard depends on the scope selected by you from the **SCOPE** drop-down list. The DCNM scope hierarchy starts with the fabrics. Fabrics can be grouped into a Multi-Site Domain (MSD). A group of MSDs constitute a Data Center. The data that is displayed on the Endpoint Locator dashboard is aggregated based on the selected scope. From this dashboard, you can access Endpoint History, Endpoint Search, and Endpoint Life.

Endpoint Locator Dashboard

To explore endpoint locator details from the Cisco DCNM Web UI, choose **Monitor > Endpoint Locator > Explore**. The **Endpoint Locator** dashboard is displayed.

Note Due to an increase in scale from Cisco DCNM Release 11.3(1), the system may take some time to collect endpoint data and display it on the dashboard. Also, on bulk addition or removal of endpoints, the endpoint information displayed on the EPL dashboard takes a few minutes to refresh and display the latest endpoint data.

You can also filter and view the endpoint locator details for a specific **Switch, VRF, Network**, and **Type**, by using the respective drop-down lists. Starting from Cisco DCNM Release 11.3(1), you can select MAC type of endpoints as a filter attribute. By default, the selected option is **All** for these fields.

You can reset the filters to the default options by clicking the **Reset Filters** icon.

The 'top pane' of the window displays the number of active endpoints, active VRFs, active networks, dual attached endpoints, single attached endpoints and dual stacked endpoints, for the selected scope. Support for displaying the number of dual attached endpoints, single attached endpoints and dual stacked endpoints has been added from Cisco DCNM Release 11.3(1). A dual attached endpoint is an endpoint that is behind at least two switches. A dual stacked endpoint is an endpoint that has at least one IPv4 address and one IPv6 address.

Historical analysis of data is performed and a statement mentioning if any deviation has occurred or not over the previous day is displayed at the bottom of each tile.

Click any tile in the top pane of the EPL dashboard to go to the Endpoint History window.

The 'middle pane' of the window displays the following information:

- **Top 10 Networks by Endpoints** - A pie chart is displayed depicting the top ten networks that have the most number of endpoints. Hover over the pie chart to display more information. Click on the required section to view the number of IPv4, IPv6, and MAC addresses.
- **Top 10 Switches by Endpoints** - A pie chart is displayed depicting the top ten switches that are connected to the most number of endpoints. Hover over the pie chart to display more information. Click on the required section to view the number of IPv4, IPv6, and MAC addresses.
- **Top Switches by Networks** - Bar graphs are displayed depicting the number of switches that are associated with a particular network. For example, if a vPC pair of switches is associated with a network, the number of switches associated with the network is 2.

The 'bottom pane' of the window displays the list of active endpoints.

Click the search icon in the **Endpoint Identifier** column to search for specific IP addresses.

In certain scenarios, the datapoint database may go out-of-sync and information, such as the number of endpoints, may not be displayed correctly due to network issues such as -

- Endpoint moves under the same switch between ports and the port information needs some time to be updated.
- An orphan endpoint is attached to the second VPC switch and is no longer an orphan endpoint.
- NX-API not enabled initially and then enabled at a later point in time.
- NX-API failing initially due to misconfiguration.
- Change in Route Reflector (RR).
- Management IPs of the switches are updated.

In such cases, clicking the **Resync** icon leads to the dashboard syncing to the data currently in the RR. However, historical data is preserved. We recommend not clicking **Resync** multiple times as this is a compute-intense activity.

Click the **Pause** icon to temporarily stop the near real-time collection and display of data.

Consider a scenario in which EPL is first enabled and the **Process MAC-Only Advertisements** checkbox is selected. Then, EPL is disabled and enabled again without selecting the **Process MAC-Only Advertisements** checkbox. As the cache data in elasticsearch is not deleted on disabling of EPL, the MAC endpoint information is still displayed in the EPL dashboard. The same behavior is observed when a Route-Reflector is disconnected. Depending on the scale, the endpoints are deleted from the EPL dashboard after some time. In certain cases, it may take up to 30 minutes to remove the older MAC-only endpoints. However, to display the latest endpoint data, you can click the **Resync** icon at the top right of the EPL dashboard.

Endpoint History

Click any tile in the top pane of the EPL dashboard to go to the **Endpoint History** window. A graph depicting the number of active endpoints, VRFs and networks, dual attached endpoints and dual stacked MAC endpoints at various points in time is displayed. The graphs that are displayed here depict all the endpoints and not only the endpoints that are present in the selected fabric. Endpoint history information is available for the last 180 days amounting to a maximum of 100 GB storage space.

Endpoint Snapshots

Starting from Cisco DCNM Release 11.3(1), you can compare endpoint data at two specific points in time. To display the **Endpoint Snapshot** window, click the **Endpoint Snapshot** icon at the top right of the **Active Endpoints** graph in the **Endpoint History** window.

By default, endpoint snapshot comparison data for the previous hour is displayed.

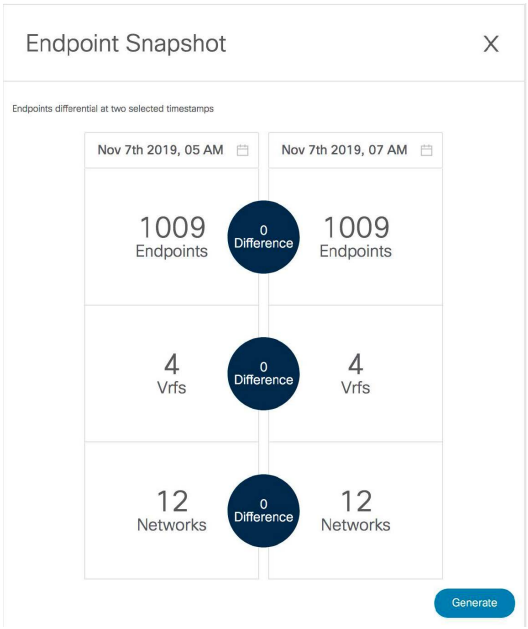

To compare endpoint snapshots at specific points in time, select two points in time, say T1 and T2, and click **Generate**.

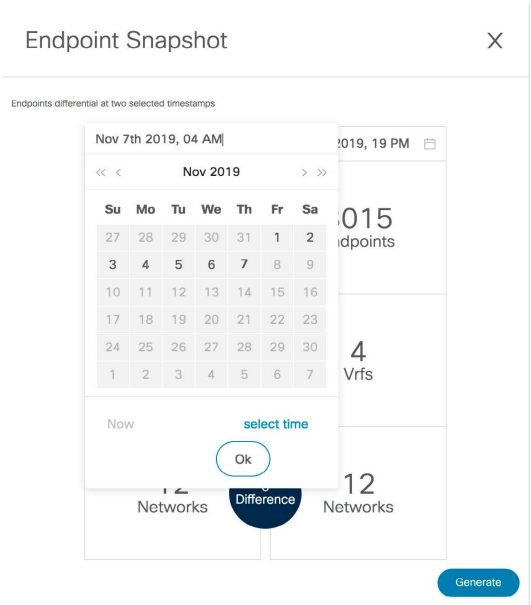

A comparison of the endpoints, VRFs, and networks at the selected points in time are displayed. Click each tile to download more information about the endpoints, VRFs, or networks. Click the **Difference** icon to

download details about the differences in data for the specified time interval. Snapshots are stored for a maximum of three months and then discarded.

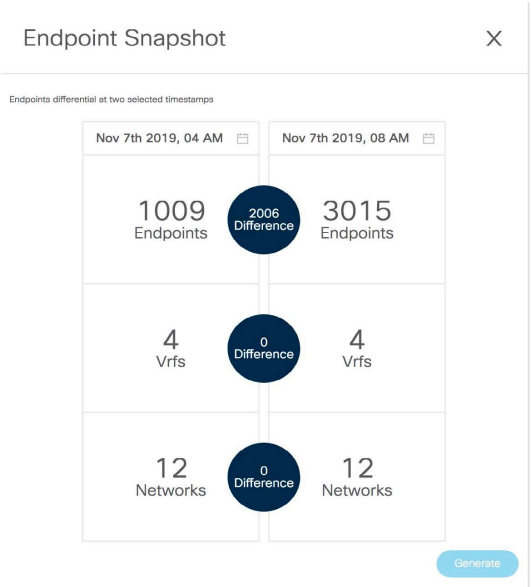

Endpoint Search

Click the **Endpoint Search** icon at the top right of the Endpoint Locator landing page to view a real-time plot displaying endpoint events for the period specified in a date range.

Endpoint Life

Click the **Endpoint Life** icon at the top right of the Endpoint Locator landing page to display a time line of a particular endpoint in its entire existence within the fabric.

Specify the IP or MAC address of an endpoint and the VXLAN Network Identifier (VNI) to display the list of switches that an endpoint was present under, including the associated start and end dates. Click **Submit**.

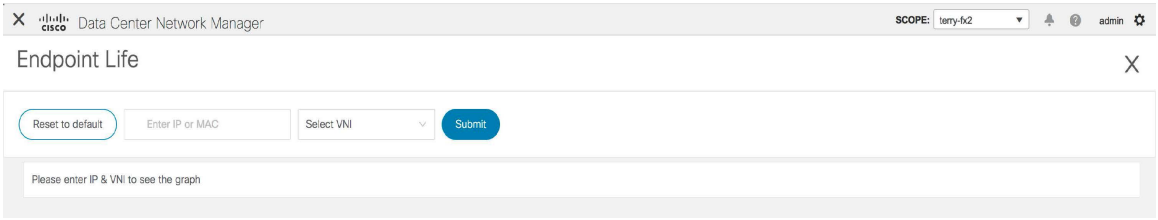

The window that is displayed is essentially the endpoint life of a specific endpoint. The bar that is orange in color represents the active endpoint on that switch. If the endpoint is viewed as active by the network, it will have a band here. If an endpoint is dual-homed, then there will be two horizontal bands reporting the endpoint existence, one band for each switch (typically the vPC pair of switches). In case the endpoints are deleted or moved, you can also see the historical endpoint deletions and moves on this window.

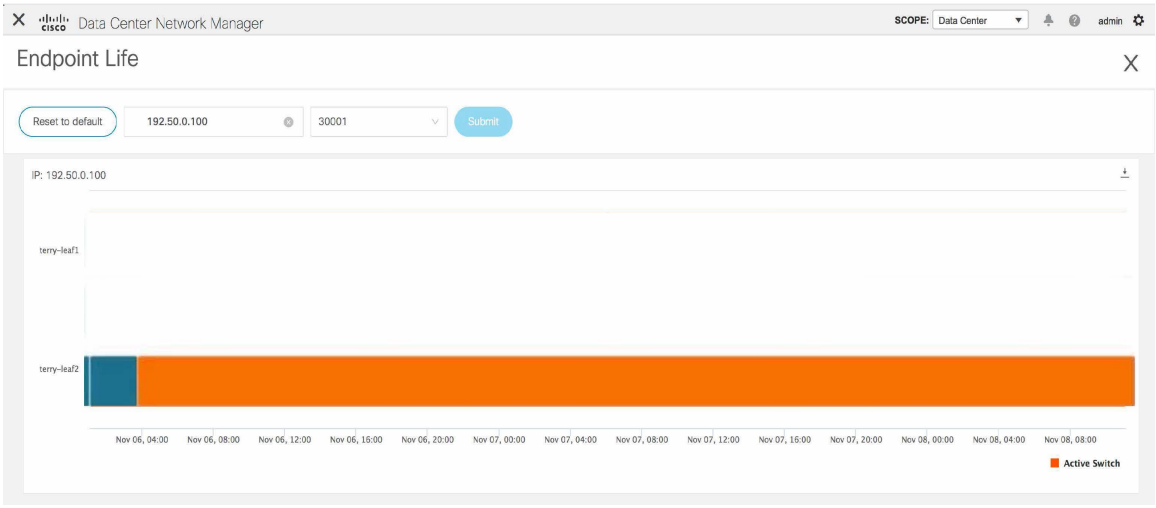

LAN Telemetry Health

Starting from DCNM 11.2(1), Streaming LAN Telemetry preview feature in DCNM is obsolete and is replaced by Network Insights Resources (NIR) application. NIR can be deployed using Cisco DCNM Applications Framework on **Web UI > Applications**. After the NIR is enabled on a fabric, you can monitor the status on the window in the Cisco DCNM Web UI.

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.

LAN Telemetry has the following topics:

Health

Cisco DCNM allows you to monitor the health attributes for each fabric. The attributes are displayed for a particular fabric or all fabrics based on the selected **SCOPE**. **Default_LAN** displays all fabrics.

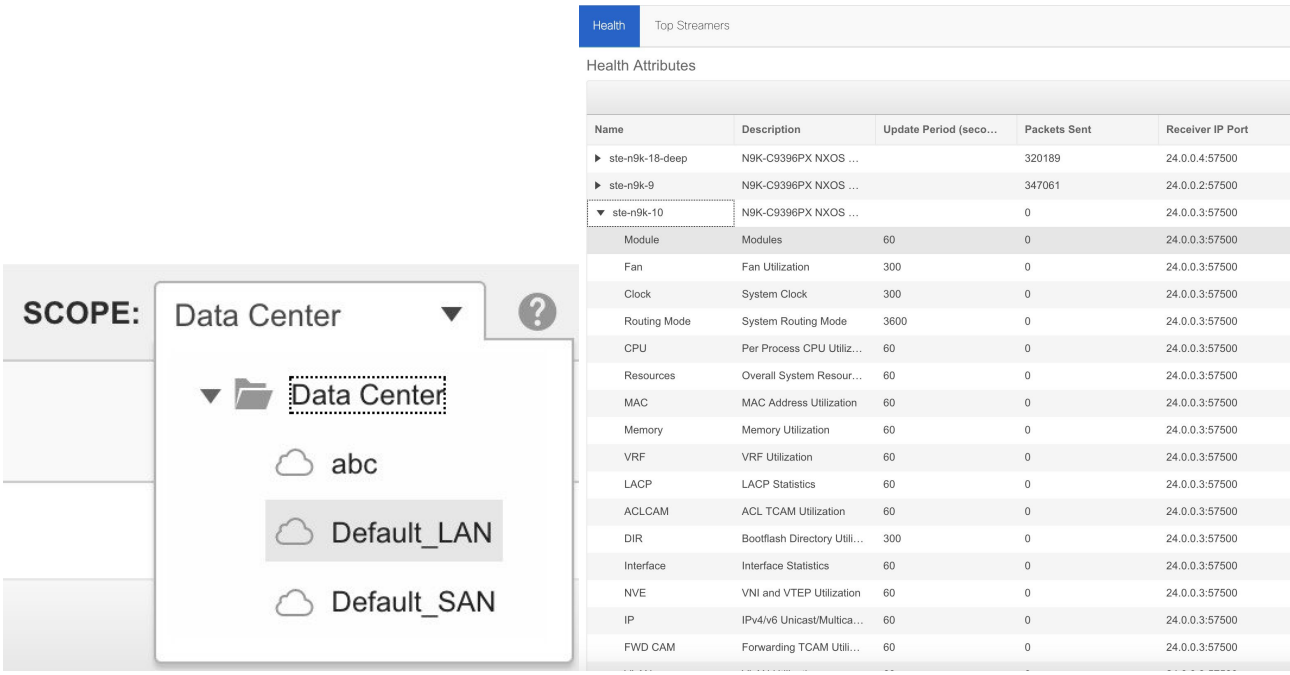

The following table describes the columns in the **LAN Telemetry > Health** tab.

Table 4: Fields and Description on Health tab

Field	Description
Name	Displays the switch name at the top. The drop-down list displays all the metric names such as CPU, Memory, and so on, that are configured on the switch by the telemetry manager.
Description	At the switch level, it displays the switch model and switch image version. At the metric level, it displays the metric description.
Update Period (seconds)	At the switch level, nothing is displayed. At the metric level, it displays the **metric collection interval in seconds**. Ex: 60 means that the switch shall stream that metric every 60 seconds.
Packets Sent	At the switch level, nothing is displayed. At the metric level, it displays the number of metric samples is collected till time.
Receiver IP Port	Displays the IP address of the UTR micro-service and the port that runs as a part of the Network Insights application, which collects the streamed telemetry data from the switches.

Field	Description
Configuration Status	

Field	Description
	Displays the telemetry configuration status on the switches. The following statuses are displayed: • **MONITOR** implies that the switch in the fabric was configured as "Monitored" in the NIR app. In this case, it is the user's responsibility to configure these switches with the required telemetry configurations and stream it to the right destination UTR IP address and Port displayed in the "Receiver IP Port" column. • **PROCESSING**: This means that the swtich belonging to the fabric was configured as "Managed" in the NIR app. In this case, the telemetry manager will configure the switches and when configuration is in progress, it is displayed as "PROCESSING". • **SUCCESS**: This means that the switches were successfully configured. • **FAILED**: This means that the switches could not be configured successfully. It could be a partial failure i.e. some of the metrics could not be configured or a full failure i.e. none of configurations went through successfully. Click on the FAILED link to the failure reason. 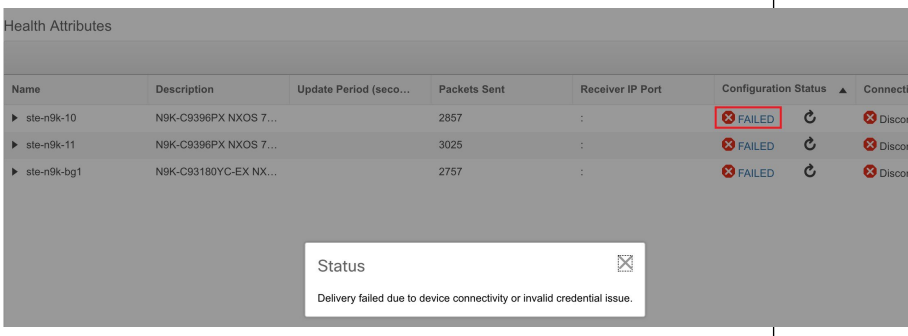 Click on the **Name** drop-down list to check which metrics failed. 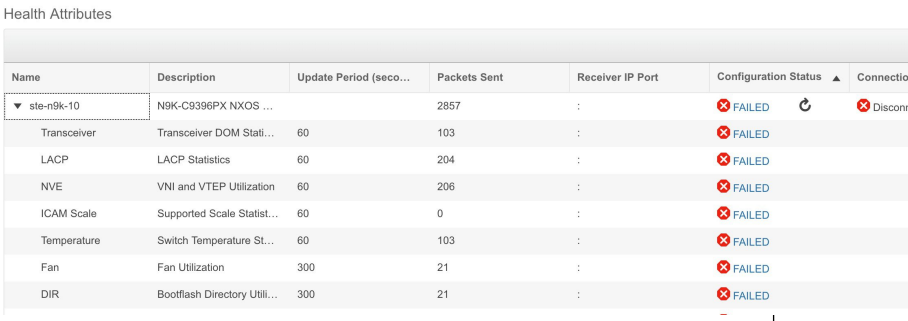 **Note** FWD TCAM and ACLCAM metric configurations are not supported on Cisco Nexus C9504, C9508, C9516 series platforms. Check the release notes for any limitations or caveats that could cause a failure. • **Failure Retry**: Click on the retry button to the right of FAILED to reconfigure the switches again.

Field	Description
	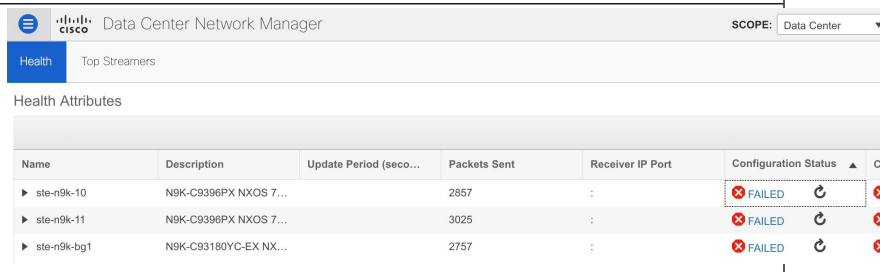

Connection Status	This indicates the status of the connection used to transport telemetry data between the switch and the UTR micro-service running in the NIR app. The value can either be **Connected** or **Disconnected**.
	
	When the connection status is **Disconnected**, and if **Configuration Status** shows either **MONITOR** or **SUCCESS**, login to the switch and check the **nxapi** configuration. When a Cisco DCNM managed fabric is enabled with telemetry, the telemetry manager pushes the **http port 80** configuration. If the switch does not have **http port 80** configuration, run the following commands on the switch:
	
	`switch# configure`
	`switch(config)# no feature nxapi`
	`switch(config)# feature nxapi`
	`switch(config)# http port 80`

Additional Information	Displays the switch serial number, fabric name, and the switch management IP address.

Note The **Health** table data gets refreshed every 2 minutes automatically. It can be manually refreshed by clicking the refresh button on the top right corner.

Failure Troubleshooting

To get more details on the failures, choose **Control > Fabric Builder**. Navigate to **Jobs** tab. Click on the fabric that you want to debug for failures.

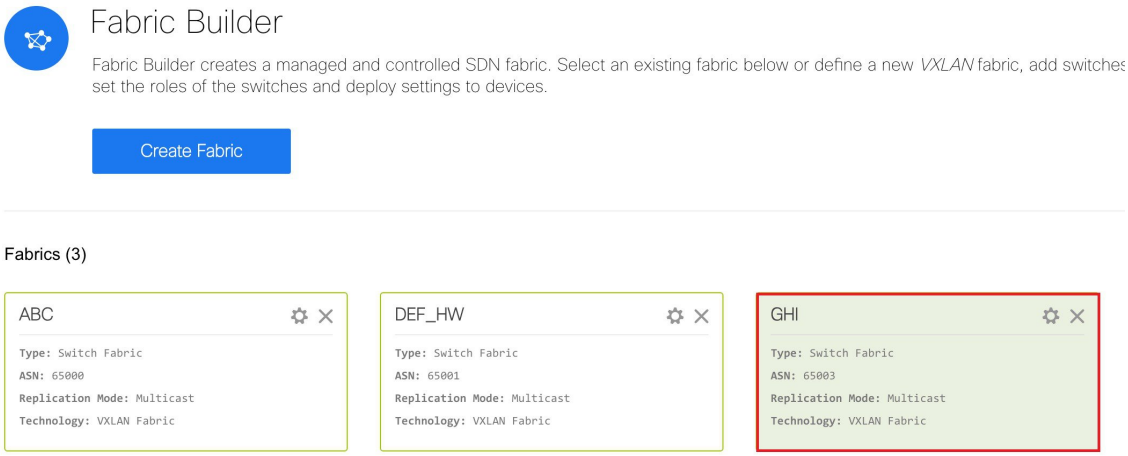

Click on the Tabular View.

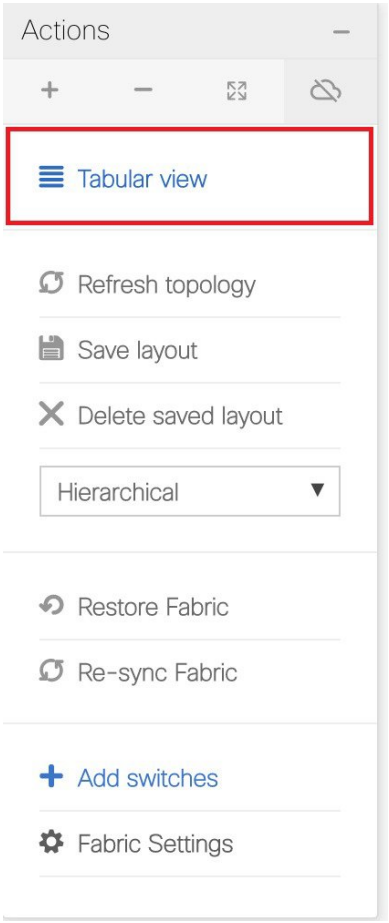

Select a switch. Click **History**.

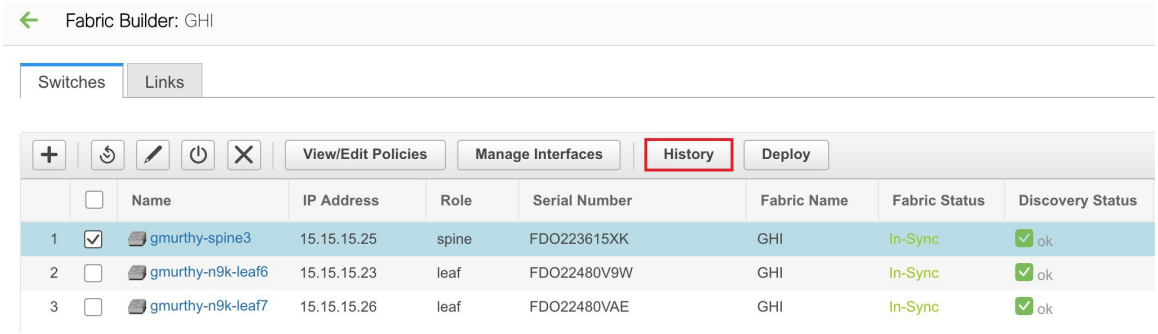

This will list all the Policies that were deployed on that switch and their Status. Check the Status Description column to see the failure reason.

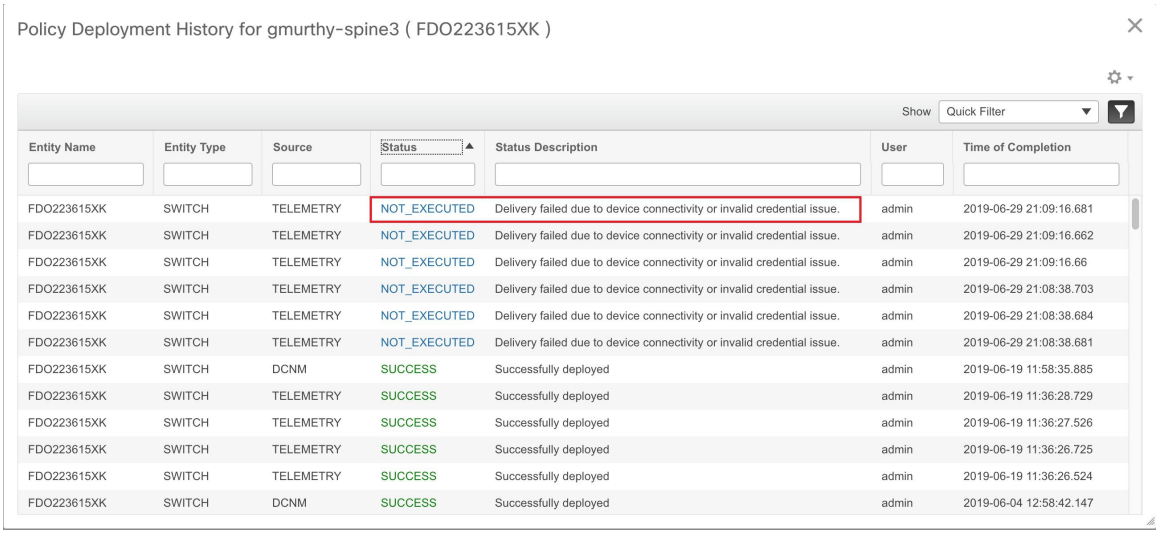

Click on the NOT_EXECUTED in the Status column to the check the commands that were not executed or failed.

Command Execution Details for gmurthy-spine3 (FDO223615XK)

Config	Status	CLI Response
telemetry	NOT_EXECUTED	
destination-profile	NOT_EXECUTED	
use-vrf pepsi	NOT_EXECUTED	
destination-group 500	NOT_EXECUTED	
ip address 17.17.17.242 port 57500 p…	NOT_EXECUTED	
sensor-group 500	NOT_EXECUTED	
data-source NX-API	NOT_EXECUTED	
path "show interface" depth unbounded	NOT_EXECUTED	
path "show lacp counters detail" dept…	NOT_EXECUTED	
path "show lacp interface" depth unb…	NOT_EXECUTED	
path "show interface transceiver detai…	NOT_EXECUTED	
path "show mac address-table count"…	NOT_EXECUTED	
path "show nve vrf" depth unbounded	NOT_EXECUTED	
path "show ip route summary vrf all" …	NOT_EXECUTED	
path "show ipv6 route summary vrf all…	NOT_EXECUTED	
path "show ip mroute summary vrf all…	NOT_EXECUTED	
path "show ipv6 mroute summary vrf …	NOT_EXECUTED	
path "show lldp traffic interface all" de…	NOT_EXECUTED	
path "show lldp all" depth unbounded	NOT_EXECUTED	

Top Streamers

Choose **LAN Telemetry > Health > Top Streamers** to view the graphs that depicts the top five streaming switches.

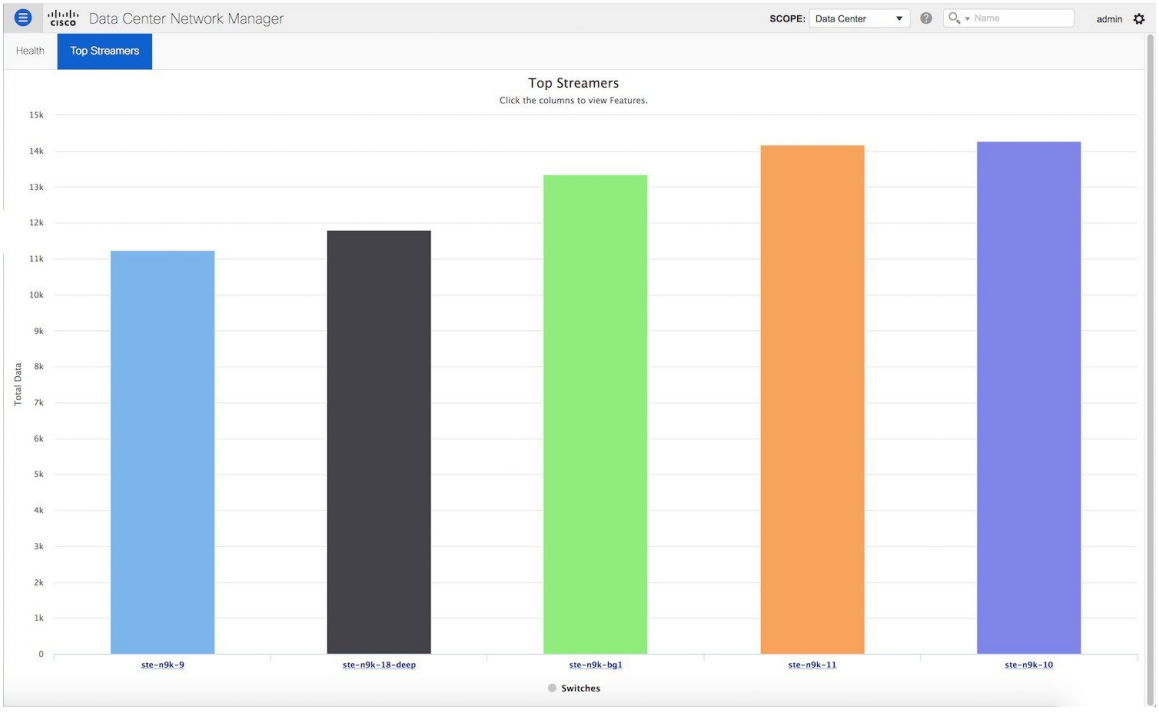

Click on the switch level bar chart to visualize a feature-wise break-down.

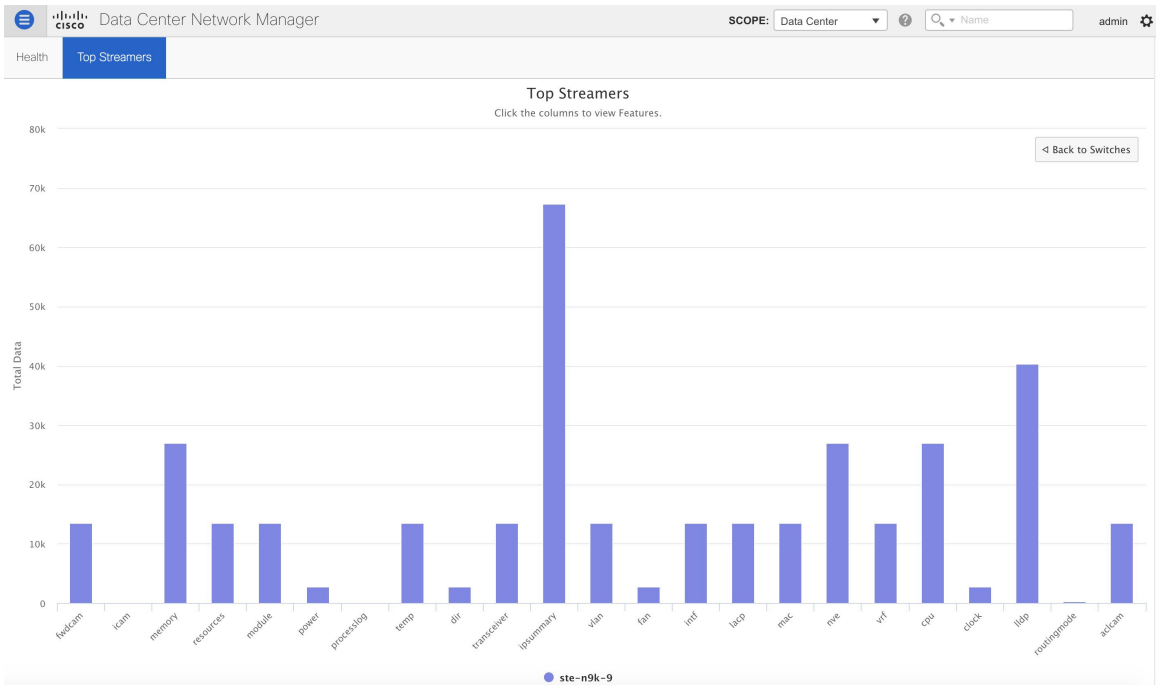
Data Center Network Manager
SCOPE: Data Center
admin
Health
Top Streamers
Top Streamers
Click the columns to view Features.
Back to Switches
Total Data
80k
70k
60k
50k
40k
30k
20k
10k
0
fwdcam
icam
memory
resources
module
power
processlog
temp
dir
transceiver
ipsummary
vlan
fan
intf
lacp
mac
nve
vrf
cpu
clock
lldp
routingmode
aclcam
ste-n9k-9