



# Control

---

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.
  - Migrate VXLAN fabrics that are built with DCNM 10.4.2 using POAP templates for Underlay provisioning and Top-down Overlay provisioning, to DCNM 11.1.
  - Migrate VXLAN fabrics that are built with DCNM 11.0 or DCNM 11.1.

This chapter contains the following topics:

- [Fabrics, on page 1](#)
- [Management, on page 138](#)
- [Template Library, on page 140](#)
- [Image Management, on page 167](#)
- [Endpoint Locator, on page 175](#)
- [Streaming Telemetry for LAN Deployments, on page 188](#)

## Fabrics

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) release, 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 & VRFs** menu option (under the **Fabrics** sub menu).

Overlay network provisioning.

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

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:

## 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:

Add Fabric ✕

\* Fabric Name :

\* Fabric Template Easy\_Fabric\_11\_1 ▼

General | Replication | vPC | Advanced | Resources | Manageability | Bootstrap | Configuration Backup

\* BGP ASN  ? 1-4294967295 | 1-65535[0-65535]

\* Fabric Interface Numbering p2p ▼ ? Numbered(Point-to-Point) or Unnumbered

\* Underlay Subnet IP Mask 30 ▼ ? Mask for Underlay Subnet IP Range

\* Link-State Routing Protocol ospf ▼ ? Supported routing protocols (OSPF/IS-IS)

\* Route-Reflectors 2 ▼ ? Number of spines acting as Route-Reflectors

\* Anycast Gateway MAC 2020.0000.00aa ▼ ? Shared MAC address for all leaves (xxxx.xxxx.xxx)

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

**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.

In the DCNM 11.1(1) release, you can change the fabric setting from one mode to the other, if no overlay profile exists for the fabric.

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.

**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.

**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.

- 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 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.

- 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.

**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 VXLAN 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.

---

**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.

**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.

**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.

**Static 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 11.1(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 DHCP** - Click this check box to initiate enabling of automatic IP address assignment through DHCP. When you click the check box, the other fields become editable. They are:

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

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

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

*DHCP scope and management default gateway IP address specification* - If you specify the management default gateway IP address 10.0.1.1 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 137](#).

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/dcm/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.

General Replication vPC Advanced Resources Manageability Bootstrap Configuration Backup

Hourly Fabric Backup  ? Backup Only when a Modified Fabric is In-Sync

Scheduled Fabric Backup  ? Backup at Specified Scheduled Time

\* Scheduled Time  ? Time in 24hr format. (00:00 to 23:59)

Save Cancel

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.

**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

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

Click the **Add Switches** option from the **Actions** panel to add switches to the fabric 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.

### 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.

## Inventory Management

Discover Existing Switches

PowerOn Auto Provisioning (POAP)

Discovery Information > Scan Details >

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

Authentication Protocol

Username

Password

Max Hops    
hop(s)

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

Start discovery

- 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.

Discover Existing Switches

PowerOn Auto Provisioning (POAP)

Discovery Information > Scan Details >

← Back
Import into fabric

<input type="checkbox"/>	Name	IP Address	Model	Version	Status	Progress
<input type="checkbox"/>	EVPN-Spine81	172.23.244.81	N9K-C931...	7.0(3)I5(2)	Unknown User...	
<input type="checkbox"/>	leaf-91	172.23.244.91	N9K-C939...	7.0(3)I7(3)	manageable	
<input type="checkbox"/>	switch	172.23.244.88	N9K-C937...	7.0(3)I7(1)	not reachable	
<input type="checkbox"/>	EVPN-Spine85	172.23.244.85	N9K-C939...	7.0(3)I5(2)	Unknown User...	

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

## Inventory Management



Discover Existing Switches | PowerOn Auto Provisioning (POAP)

Discovery Information > Scan Details >

← Back 2 Import into fabric

<input type="checkbox"/>	Name	IP Address	Model	Version	Status	Progress
<input type="checkbox"/>	EVPN-Spine81	172.23.244.81	N9K-C931...	7.0(3)I5(2)	Unknown User...	
<input checked="" type="checkbox"/>	Leaf-91	172.23.244.91	N9K-C939...	7.0(3)I7(3)	manageable	
<input type="checkbox"/>	Switch	172.23.244.88	N9K-C937...	7.0(3)I7(1)	not reachable	
<input type="checkbox"/>	EVPN-Spine85	172.23.244.85	N9K-C939...	7.0(3)I5(2)	Unknown User...	

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.

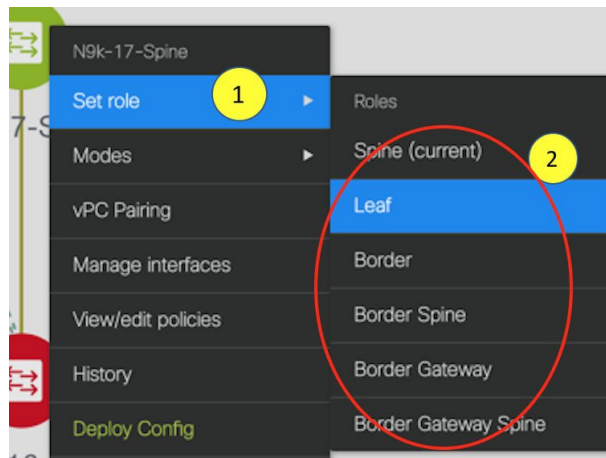
Device connectivity issue: Before proceeding further, wait for ten minutes for the switch-internal processes to complete. Else, you might encounter a device connectivity failure message at a later stage.

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.



- 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 24](#).
- 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.

**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.

## Config Deployment



Step 1. Configuration Preview &gt;

Step 2. Configuration Deployment Status &gt;

Switch Name	IP Address	Switch Serial	Preview Config	Status	Re-sync	Progress
N9K-2-Leaf	111.0.0.92	SAL18422FVP	0 lines	In-sync		100%
N9K-4-BGW	111.0.0.94	FDO20260UEK	20 lines	Out-of-sync		100%
N9K-3-BGW	111.0.0.93	FDO20291AVQ	20 lines	Out-of-sync		100%
N9K-1-Spine	111.0.0.91	SAL18432P2T	0 lines	In-sync		100%

Deploy Config

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. 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 switch.

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



## Config Preview - Switch 111.0.0.94

Pending Config	Expected Config	Current Config	Side-by-side Comparison
<pre> interface ethernet1/21 description connected-to-N9K-3-BGW-Ethernet1/21 no switchport medium p2p ip address 10.23.0.1/30 ip router ospf UNDERLAY area 0.0.0.0 ip ospf network point-to-point ip pim sparse-mode mtu 9216 no shutdown interface ethernet1/22 description connected-to-N9K-3-BGW-Ethernet1/22 no switchport medium p2p ip address 10.23.0.5/30 ip router ospf UNDERLAY area 0.0.0.0 ip ospf network point-to-point ip pim sparse-mode mtu 9216 no shutdown           </pre>			

The Pending Config tab displays the pending configurations for successful deployment. The Expected Config and Current Config tabs display the expected and current configurations on the switch.

The **Side-by-side Comparison** tab displays the current configurations and expected configurations together. Common configurations appear next to each other and are not highlighted. In the **Expected config** column within this tab, the additional configurations are highlighted in green. In the **Running config** column within this tab, the additional configurations of the running config are highlighted in a distinct color.

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

Config Preview - Switch 111.0.0.94			
Pending Config	Expected Config	Current Config	Side-by-side Comparison
110	vrf context management	vrf context management	
111	ip route 0.0.0.0/0 111.0.0.251	ip route 0.0.0.0/0 111.0.0.251	
112	nxapi http port 80	nxapi http port 80	
113	interface vlan1		
114	interface nve1	interface nve1	
115	no shutdown	no shutdown	
116	host-reachability protocol bgp	host-reachability protocol bgp	
117	source-interface loopback1	source-interface loopback1	
118	multisite border-gateway interface loopback100	multisite border-gateway interface loopback100	
119		multisite border-gateway interface loopback100	
120	interface ethernet1/1	interface ethernet1/1	

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\_config**, 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.

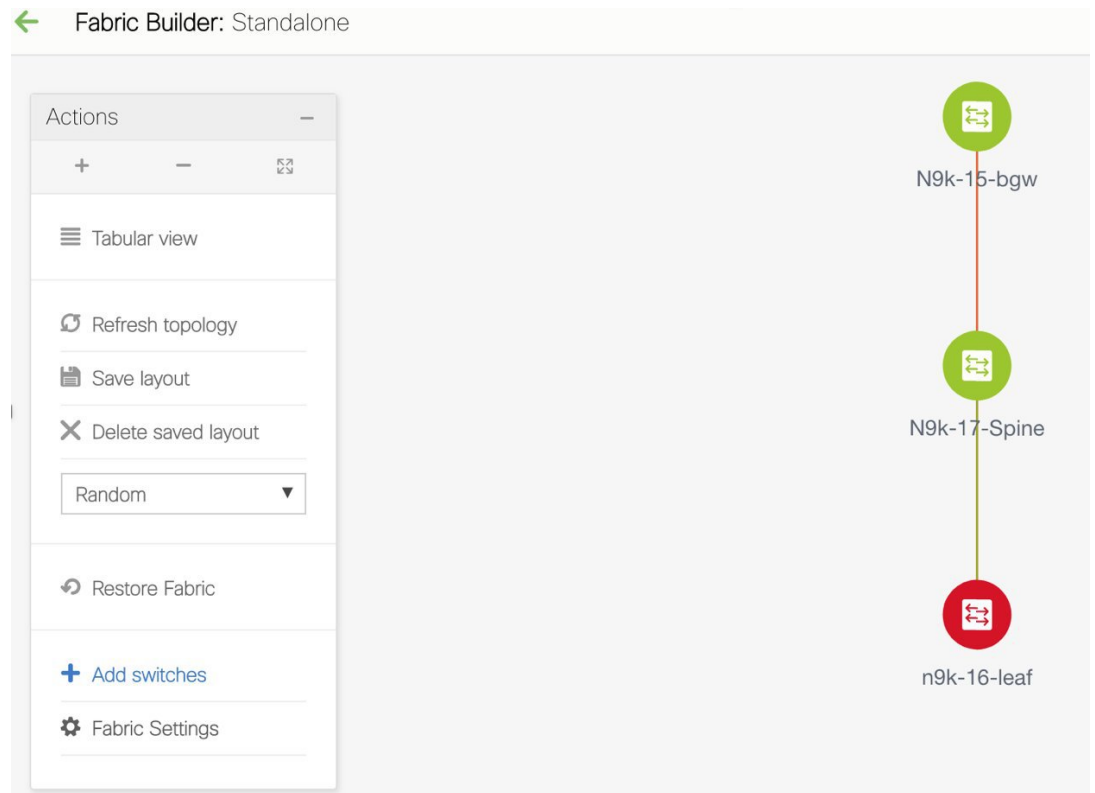
The Configuration Compliance function and principles are applicable for discovering existing and new switches. New switch discovery in DCNM (through a simplified POAP process) is explained next.

### 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.

- 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.
- 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 DHCP, make sure that DHCP is enabled on the fabric settings. Click **Fabric Settings** and edit the DHCP information in the **Bootstrap** tab.

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

Inventory Management

Discover Existing Switches PowerOn Auto Provisioning (POAP)

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

Bootstrap

\* Password  \* Confirm Password

<input type="checkbox"/>	Serial Number	Model	Version	IP Address	Hostname
<input type="checkbox"/>	FDO21323D58	N9K-93180YC-EX	9.2(1)	<input type="text"/>	<input type="text"/>

Close

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

7. 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.
8. Click **Refresh Topology** to get updated information. The added switch goes through the POAP cycle. Monitor and check the switch for POAP completion.
9. 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.

---

10. After the pending configurations are deployed, the **Progress** column displays 100% for all switches.
11. Click **Close** to return to the fabric builder topology.
12. Click **Refresh Topology** to view the update. All switches must be in green color indicating that they are functional.
13. 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.
14. 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.

# Interfaces

<div style="display: flex; justify-content: space-between; align-items: center;"> <span>+</span> <span>↻</span> <span>✎</span> <span>✕</span> <span>↑</span> <span>↓</span> </div>		
	Device Name	Name
<input type="checkbox"/>	N9K-2-Leaf	Ethernet2/6
<input type="checkbox"/>	N9K-2-Leaf	Ethernet2/7
<input type="checkbox"/>	N9K-2-Leaf	Ethernet2/8
<input type="checkbox"/>	N9K-2-Leaf	Ethernet2/9
<input type="checkbox"/>	N9K-2-Leaf	Ethernet2/10
<input type="checkbox"/>	N9K-2-Leaf	Ethernet2/11
<input type="checkbox"/>	N9K-2-Leaf	Ethernet2/12
<input checked="" type="checkbox"/>	N9K-2-Leaf	nve1

2

1

**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 24](#).
- 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.
- **Manage Interfaces** - Deploy configurations on the switch interfaces.
- **View/Edit Policies** - See switch policies and edit them as required.
- **History** - View per switch deployment history.
- **Deploy Config** - Deploy per switch configurations.
- **Discovery** - You can use this option to update the credentials of the switch, reload the switch, rediscover the switch, and remove the switch from the fabric.

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

- Provision interface configurations such as vPCs, loopback interface, and subinterface configurations. [Refer [Interfaces](#)].
- Create overlay networks and VRFs and deploy them on the switches. [Refer [Creating and Deploying Networks and VRFs](#)].

## Switch Operations

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

- **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
  - Border Gateway Spine



---

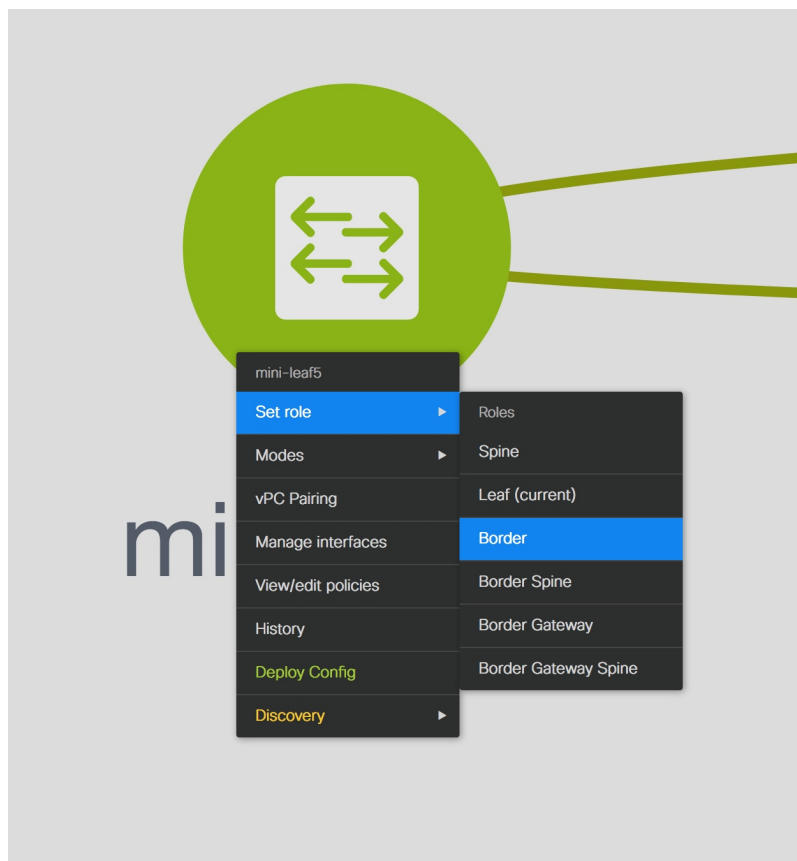
**Note**

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

Starting from DCNM 11.1(1), you can change switch roles if there are no overlays on the switches. The updated configuration is then generated after you click **Save and Deploy**. The following switch role changes are allowed:

- 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.




---

**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.

## Changing Fabric Interface Numbering

This procedure shows how to change the **Fabric Interface Numbering** setting of an existing fabric to **unnumbered**.

### Procedure

---

- Step 1** Select an existing fabric from the **Fabric Builder** window.
- Step 2** Click **Tabular view** in the **Actions** menu.
- Step 3** Click the **Links** tab.
- Step 4** Select the link connecting a Spine and a Leaf, and click the **Update Link** icon.
- Step 5** In the **Link Template** field, select **int\_intra\_fabric\_unnum\_link\_11\_1**.
- Step 6** Click **Save** and close the **Link Management - Edit Link** window.
- Step 7** Repeat this procedure for the all the links connecting a Spine and a Leaf.
- Step 8** Navigate back to the fabric, and click the **Fabric Settings** in the **Actions** menu.
- Step 9** Under the **General** tab, select **unnumbered** from the **Fabric Interface Numbering** drop-down list.

- Step 10** Click **Save** and close the window.
- Step 11** Click **Save & Deploy** to deploy the updated configuration.

## 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.

The screenshot shows the Cisco Data Center Network Manager interface. The breadcrumb navigation is 'Fabric Builder: easy\_fabric'. The 'Switches' tab is active, and the 'View/Edit Policies' button is selected. The table below lists the switches and their associated policies.

	<input type="checkbox"/>	Name	IP Address	Role	Serial Number	Fabric Name	Fabric Status	Discovery Status	Model
1	<input checked="" type="checkbox"/>	anm-host80	172.23.244.80	leaf	SAL1925HA3U	easy_fabric	In-Sync	<input checked="" type="checkbox"/> ok	N9K-C9312
2	<input checked="" type="checkbox"/>	EVPN-Spine81	172.23.244.81	leaf	SAL1919ELJQ	easy_fabric	Out-of-sync	<input checked="" type="checkbox"/> ok	N9K-C9312

## View/Edit Policies

Selected 0 / Total 1139

View All Deploy Show Quick Filter

<input type="checkbox"/>	Template	Priority	Fabric Name	Serial Number	Editable	Entity Type	Entity Name	Source
<input type="checkbox"/>	no_shut_interface	400	easy_fabric	SAL1925HA3U	false	INTERFACE	nve1	nve1
<input type="checkbox"/>	mgmt_interface_11_1	900	easy_fabric	SAL1925HA3U	false	INTERFACE	mgmt0	mgmt0
<input type="checkbox"/>	int_mgmt_11_1	900	easy_fabric	SAL1925HA3U	false	INTERFACE	mgmt0	mgmt0
<input type="checkbox"/>	no_shut_interface	910	easy_fabric	SAL1925HA3U	false	INTERFACE	mgmt0	mgmt0
<input type="checkbox"/>	int_eth	910	easy_fabric	SAL1925HA3U	false	INTERFACE	mgmt0	mgmt0
<input type="checkbox"/>	cdp_disable_interface...	910	easy_fabric	SAL1925HA3U	false	INTERFACE	mgmt0	mgmt0
<input type="checkbox"/>	copp_policy	50	easy_fabric	SAL1919ELJQ	false	SWITCH	SWITCH	UNDE
<input type="checkbox"/>	feature_pim	50	easy_fabric	SAL1919ELJQ	false	SWITCH	SWITCH	UNDE
<input type="checkbox"/>	base_feature_leaf_upg	50	easy_fabric	SAL1919ELJQ	false	SWITCH	SWITCH	UNDE
<input type="checkbox"/>	feature_ospf	50	easy_fabric	SAL1919ELJQ	false	SWITCH	SWITCH	UNDE

**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.

### Add Policy ✕

\* Priority (1-1000):

\* Policy:  ▼

General

\* Banner  ? Banner

Variables:

### View/Edit Policies ✕

Selected 0 / Total 2 ↻ ⚙

+ ✎ ✕ View View All Deploy
Show  ▼ 🔍

☐	Template	Priority	Fabric Name	Serial Number	Editable ▼	Entity Type	Entity Name	Source
<input type="checkbox"/>	<input type="text"/>	<input type="text" value="220"/> ✕	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
<input type="checkbox"/>	banner	220	easy_fabric	SAL1925HA3U	true	SWITCH	SWITCH	
<input type="checkbox"/>	banner	220	easy_fabric	SAL1919ELJQ	true	SWITCH	SWITCH	

**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 **Deploy**. The selected PTI's configs are pushed to the group of switches.

View/Edit Policies ✕

Selected 4 / Total 1141  

<input type="checkbox"/>	Template	Priority	Fabric Name	Serial Number	Editable	Entity Type	Entity Name	Source
<input type="checkbox"/>	nfm_switch_user	100	easy_fabric	SAL1925HA3U	true	SWITCH	SWITCH	
<input type="checkbox"/>	nfm_switch_user	100	easy_fabric	SAL1925HA3U	true	SWITCH	SWITCH	
<input type="checkbox"/>	nfm_switch_snmp_user	150	easy_fabric	SAL1925HA3U	true	SWITCH	SWITCH	
<input type="checkbox"/>	nfm_switch_snmp_user	150	easy_fabric	SAL1925HA3U	<input type="checkbox"/>	SWITCH	SWITCH	
<input checked="" type="checkbox"/>	banner	220	easy_fabric	SAL1925HA3U	true	SWITCH	SWITCH	
<input checked="" type="checkbox"/>	v4_mgmt_default_gat...	910	easy_fabric	SAL1925HA3U	true	SWITCH	SWITCH	
<input checked="" type="checkbox"/>	switch_role_simulated	10	easy_fabric	SAL1919ELJQ	true	SWITCH	SWITCH	
<input checked="" type="checkbox"/>	nve_lb_id	10	easy_fabric	SAL1919ELJQ	true	SWITCH	SWITCH	
<input type="checkbox"/>	bgp_lb_id	10	easy_fabric	SAL1919ELJQ	true	SWITCH	SWITCH	
<input type="checkbox"/>	power_redundancy	50	easy_fabric	SAL1919ELJQ	true	SWITCH	SWITCH	
<input type="checkbox"/>	host_11_1	50	easy_fabric	SAL1919ELJQ	true	SWITCH	SWITCH	

### 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 **Deploy** to push the policy config to the device.

## Edit Policy



Policy ID: POLICY-5290  
Entity Type: SWITCH

Template Name: host\_11\_1  
Entity Name: SWITCH

\* Priority (1-1000):

General

\* Switch Name  Host name of the switch (Max Size 63)

Variables:

Save

Deploy

Cancel

### 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.

#### 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.

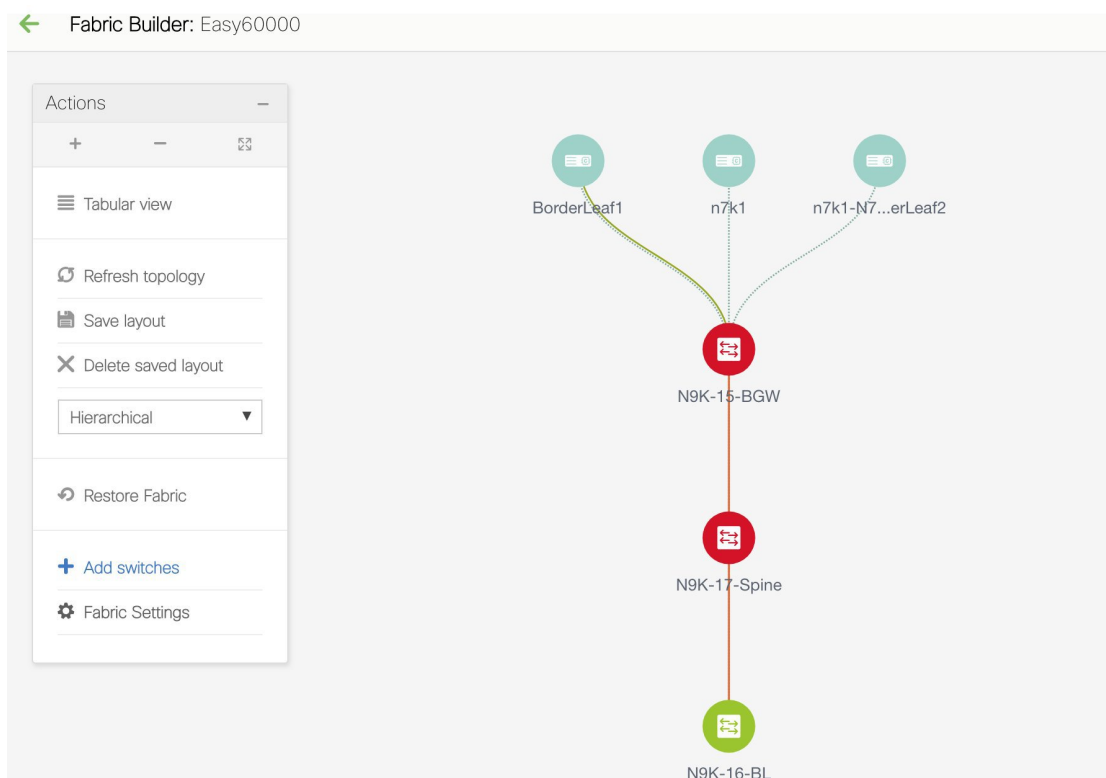
There are scenarios where you might want to define links between switches before connecting them physically. The links could be inter-fabric or intra-fabric links. Doing so, you can express and represent your intent to add links. The links with intent are displayed in a different colour till they are 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.

### 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.

	Name	IP Address	Role	Serial Number	Fabric Name	Fabric Status	Discovery Status	Model
1	N9K-15-BGW	111.0.0.95	border ...	FDO20401LB4	Easy60000	In-Sync	ok	N9K-C93180YC-EX
2	N9K-16-Leaf	111.0.0.96	leaf	SAL18432P6G	Easy60000	In-Sync	ok	N9K-C9396PX
3	N9K-17-Spine	111.0.0.97	spine	FDO20401LEJ	Easy60000	In-Sync	ok	N9K-C93180YC-EX

4. Click the Links tab. You can see a list of links.  
The list is empty when you are yet to create a link.



		Scope	Name	Policy	Admin State	Oper State
1	<input type="checkbox"/>	Easy60000	N9K-15-BGW-Ethernet1/3---n7k1-N7K-1-BorderLeaf2-Ethe...			
2	<input type="checkbox"/>	Easy60000	N9K-16-Leaf-Ethernet2/1---n7k1-Ethernet7/8			
3	<input type="checkbox"/>	External65000<->Easy60000	BorderLeaf1-Loopback0---N9K-15-BGW-loopback0	multisite_overlay_setup_rs_test		
4	<input type="checkbox"/>	Easy7200<->Easy60000	N9K-4-BGW-Ethernet1/2---N9K-15-BGW-Ethernet1/8	ext_multisite_underlay_setup_test		
5	<input type="checkbox"/>	Easy7200<->Easy60000	N9K-3-BGW-Ethernet1/2---N9K-15-BGW-Ethernet1/7	ext_multisite_underlay_setup_test		
6	<input type="checkbox"/>	Easy60000	N9K-15-BGW-Ethernet1/5---N9K-17-Spine-Ethernet1/1	int_intra_fabric_num_link_11_1		
7	<input type="checkbox"/>	Easy7200<->Easy60000	N9K-1-Spine-Ethernet1/1---N9K-16-Leaf-Ethernet1/3			
8	<input type="checkbox"/>	Easy60000	N9K-17-Spine-Ethernet1/2---N9K-16-Leaf-Ethernet1/5	int_intra_fabric_num_link_11_1		
9	<input type="checkbox"/>	Easy7200<->Easy60000	N9K-2-Leaf-Ethernet1/2---N9K-16-Leaf-Ethernet1/4			
10	<input type="checkbox"/>	Easy60000	N9K-15-BGW-Ethernet1/2---N9K-16-Leaf-Ethernet1/2			
11	<input type="checkbox"/>	Easy60000<->Easy7200	N9K-15-BGW-Ethernet1/4---N9K-1-Spine-Ethernet1/2			
12	<input type="checkbox"/>	Easy60000<->Easy7200	N9K-15-BGW-Ethernet1/50---N9K-18-BGW-Ethernet1/7			
13	<input type="checkbox"/>	Easy60000<->External65000	N9K-15-BGW-Ethernet1/49---n7k1-BorderLeaf1-Ethernet7/6			

- 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.

#### Link Management - Add Link

\* Link Type 
  
 \* Link Sub-Type 
  
 \* Link Template 
  
 \* Source Fabric 
  
 \* Destination Fabric 
  
 \* Source Device 
  
 \* Source Interface 
  
 \* Destination Device 
  
 \* Destination Interface

▼ Link Profile

General

\* FABRIC\_NAME  ? FABRIC NAME

\* Source IP  ? IP address of the source interface

\* Destination IP  ? IP address of the destination interface

Interface Admin State  ? Admin state of the interface

\* MTU  ? MTU for the interface

Save

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`.
- 

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.

**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 template.

---

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

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

## Link Management - Add Link



* Link Type	Intra-Fabric	▼
* Link Sub-Type	Fabric	▼
* Link Template	int_intra_fabric_num_link_11_1	▼
* Source Fabric	Easy60000	▼
* Destination Fabric	Easy60000	▼
* Source Device	N9K-16-BL	▼
* Source Interface	Ethernet1/40	▼
* Destination Device	N9K-17-Spine	▼
* Destination Interface	Ethernet1/40	▼

▼ Link Profile

General

Advanced

\* FABRIC\_NAME Easy60000 ? FABRIC NAME

\* Source IP 10.1.1.1 ? IP address of the source interface

\* Destination IP 10.1.1.3 ? IP address of the destination interface

Interface Admin State  ? Admin state of the interface

\* MTU 9216 ? MTU for the interface

Save

## Advanced tab.

▼ Link Profile

General

Advanced

Source Interface Desc... Border Leaf to Route Reflector1 ? Add description to the source interface

Destination Interface ... Route Reflector1 to Border Leaf ? Add description to the destination interface

Source Interface Freeform CLI... ? Additional CLI for source interface

Destination Interface ... ? Additional CLI for destination interface

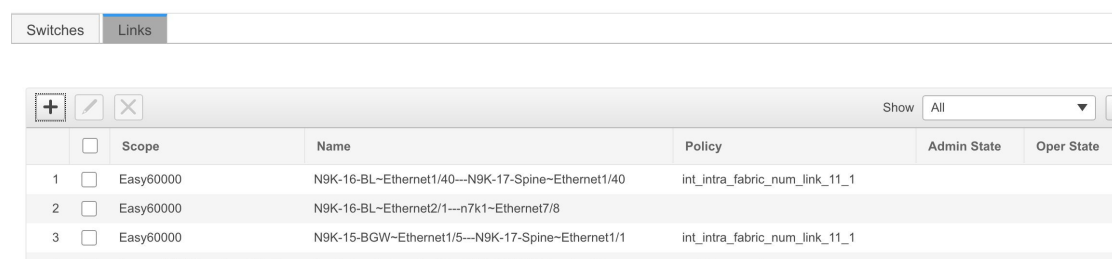
Save

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 be 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](#).

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

The new link appears in the Links tab.



	<input type="checkbox"/>	Scope	Name	Policy	Admin State	Oper State
1	<input type="checkbox"/>	Easy60000	N9K-16-Leaf-Ethernet1/40---N9K-17-Spine-Ethernet1/40	int_intra_fabric_num_link_11_1		
2	<input type="checkbox"/>	Easy60000	N9K-16-Leaf-Ethernet2/1---n7k1-Ethernet7/8			
3	<input type="checkbox"/>	Easy60000	N9K-15-BGW-Ethernet1/5---N9K-17-Spine-Ethernet1/1	int_intra_fabric_num_link_11_1		

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

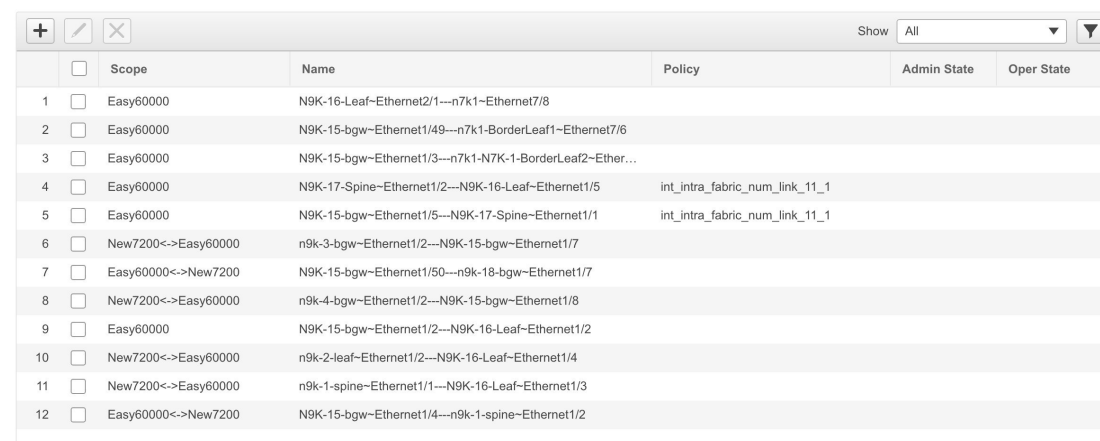
The Config Deployment screen comes up. It displays the configuration status on the switches. You can also view the pending configurations by clicking the respective link in the Preview Config column. When you click a link in the Preview Config column, the Config Preview window comes up. It lists the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

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

- 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).



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

- 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.

## Link Management - Add Link



\* Link Type  ▼

\* Link Sub-Type  ▼

\* Link Template  ▼

\* Source Fabric  ▼

\* Destination Fabric  ▼

\* Source Device  ▼

\* Source Interface  ▼

\* Destination Device  ▼

\* Destination Interface  ▼

---

▼ Link Profile

General

Advanced

\* FABRIC\_NAME  ? FABRIC NAME

\* Source IP  ? IP address of the source interface

\* Destination IP  ? IP address of the destination interface

Interface Admin State  ? Admin state of the interface

\* MTU  ? MTU for the interface

Save

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

## Link Management - Add Link



* Link Type	Inter-Fabric	▼
* Link Sub-Type	VRF_LITE	▼
* Link Template	ext_fabric_setup_test	▼
* Source Fabric	Easy60000	▼
* Destination Fabric		▼
* Source Device		▼
* Source Interface		▼
* Destination Device		▼
* Destination Interface		▼

▼ Link Profile

General

\* Local BGP AS #  ? Local BGP Autonomous System Nu

\* IP\_MASK  ?

\* NEIGHBOR\_IP  ?

\* NEIGHBOR\_ASN  ?

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:

Link Management - Add Link ✕

* Link Type	Inter-Fabric
* Link Sub-Type	VRF_LITE
* Link Template	ext_fabric_setup_test
* Source Fabric	Easy60000
* Destination Fabric	New7200
* Source Device	N9K-15-bgw
* Source Interface	Ethernet1/9
* Destination Device	n9k-18-bgw
* Destination Interface	Ethernet1/9

▼ Link Profile

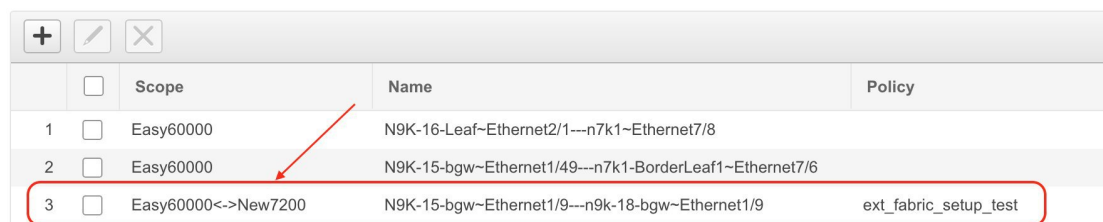
General

* Local BGP AS #	60000	? Local BGP Autonomous System Nu
* IP_MASK	10.3.4.5/24	?
* NEIGHBOR_IP	10.3.4.7	?
* NEIGHBOR_ASN	7200	?

[Save](#)

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.



	<input type="checkbox"/>	Scope	Name	Policy
1	<input type="checkbox"/>	Easy60000	N9K-16-Leaf~Ethernet2/1---n7k1~Ethernet7/8	
2	<input type="checkbox"/>	Easy60000	N9K-15-bgw~Ethernet1/49---n7k1~BorderLeaf1~Ethernet7/6	
3	<input type="checkbox"/>	Easy60000<->New7200	N9K-15-bgw~Ethernet1/9---n9k-18-bgw~Ethernet1/9	ext_fabric_setup_test

- 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.

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

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](#).

## Restore Fabric

Cisco DCNM supports configuration restore at fabric level. Take a backup of the configuration to restore it.



## Procedure

---

- Step 1** Choose **Control > Fabrics > Fabric Builder** and select a fabric.
- Step 2** Select **Restore Fabric** from the **Actions** menu.  
**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 **1m**, which is one month, backup information will be displayed. You can also select a custom date range. The backup information includes the backup date, total number of devices, number of devices in sync, and the 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** The backup is not valid if devices are added or removed from the fabric.
- 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.  
**Restore Status** window appears. You can view the status of **Validating Backup**, **Restoring fabric intent**, **Restoring underlay intent**, **Restoring interface intent**, and **Restoring overlay intent**. The valid values for the status of any action will be **In Progress**, **Pending**, or **Failed**.  
**Note** If the status of **Validating Backup** is **Failed**, other restoring actions will not be listed in this window.
- Step 8** Click **Next** after the intent is restored.  
**Configuration Preview** window appears. You can view the details of the switch name, IP address, switch serial number, preview configuration, status, and the progress in this window.
- Step 9** Click **Deploy** to deploy the restored configuration.  
**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.
- 

## 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.

## 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

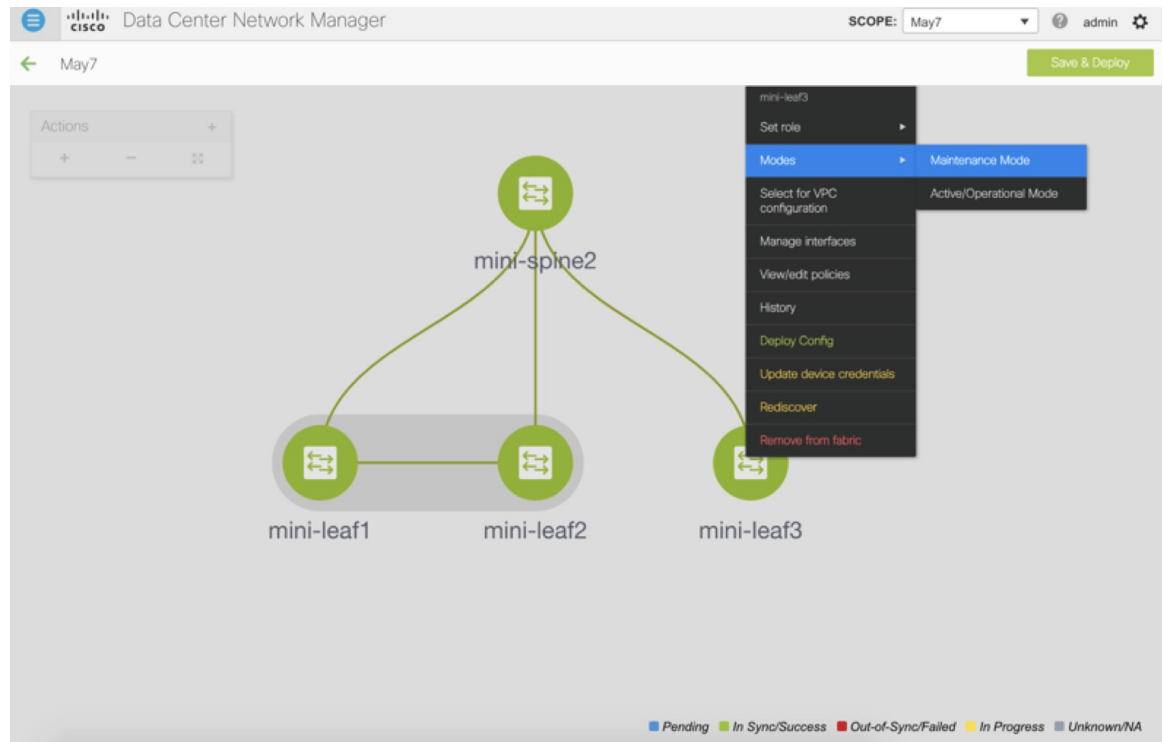
- The switch must be replaced with a switch of the same model (ASIC type) and physical port configuration. If not, the old switch must be removed and a new switch (replacement) added as a new switch into the fabric.

### POAP RMA Flow


#### Procedure

---

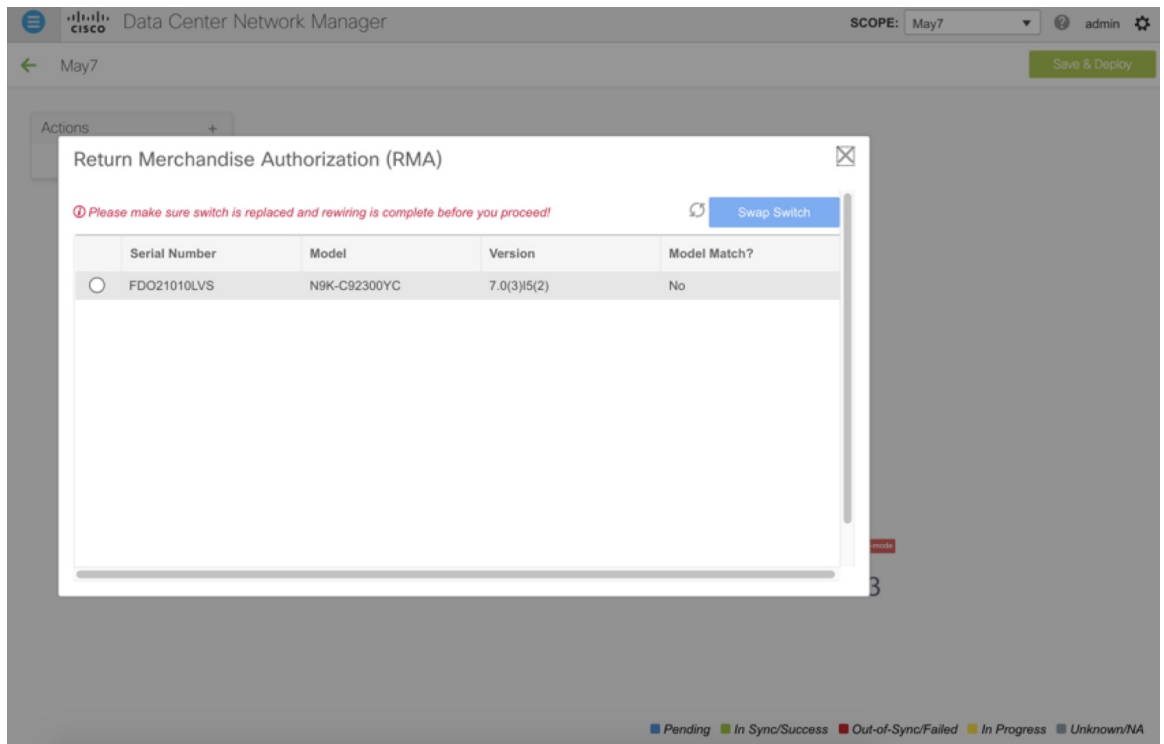
- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Click the Fabric where you want to perform RMA.
- Step 3** Move the device into maintenance mode. To move a device into maintenance mode, right-click on the device, and then choose **Modes > Maintenance Mode**.



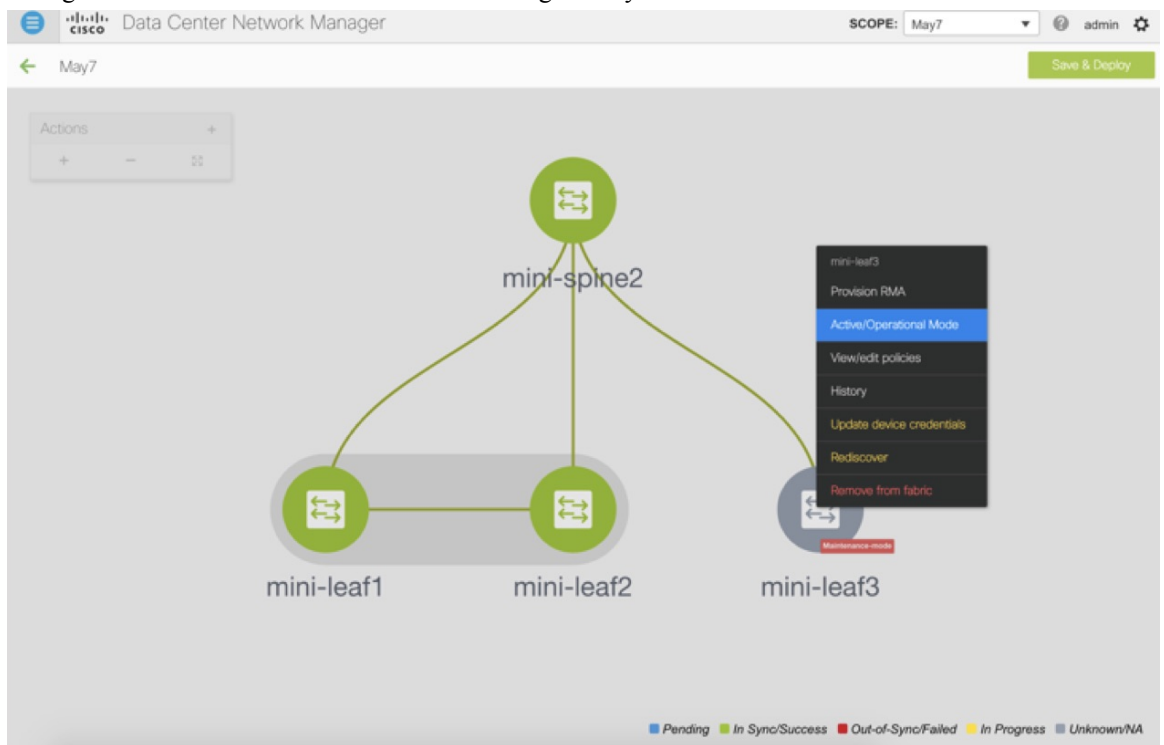
- Step 4** Physically replace the device in the network. Physical connections should be made in the same place on the replacement switch as they existed on the original switch.
- Step 5** Provision RMA flow and select the replacement device.



- 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.

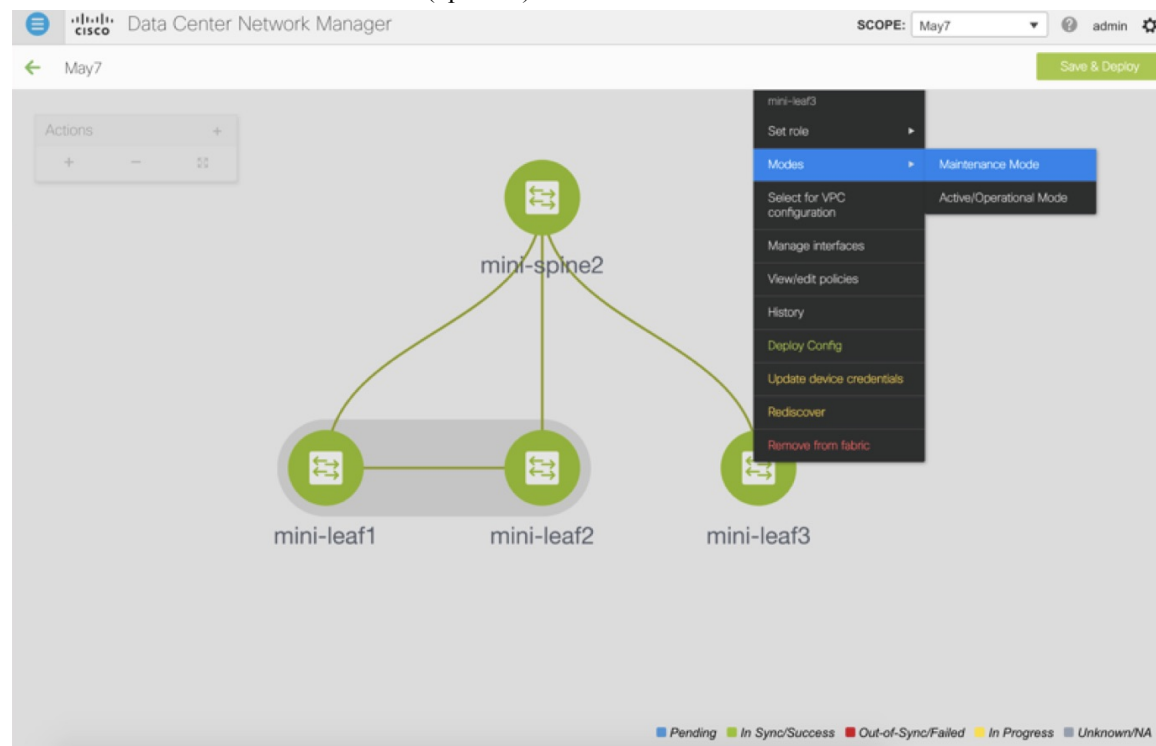


## Manual RMA Flow

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).

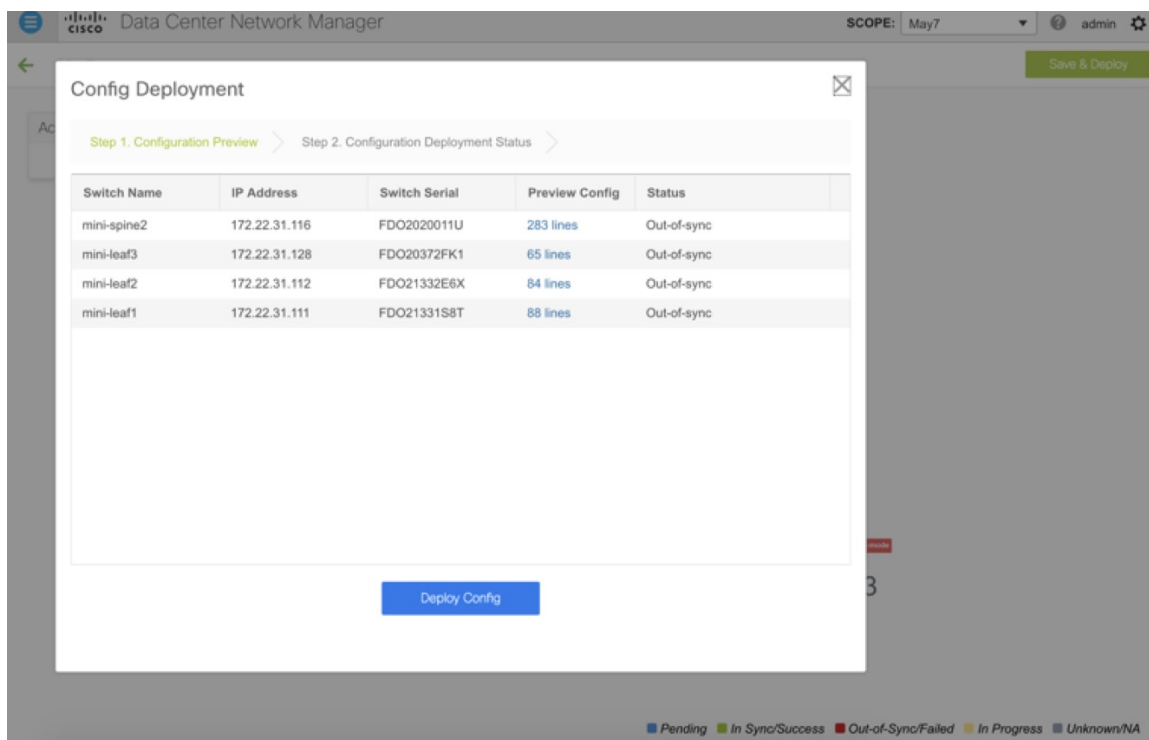


**Step 2** Physically replace the device in the network.

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

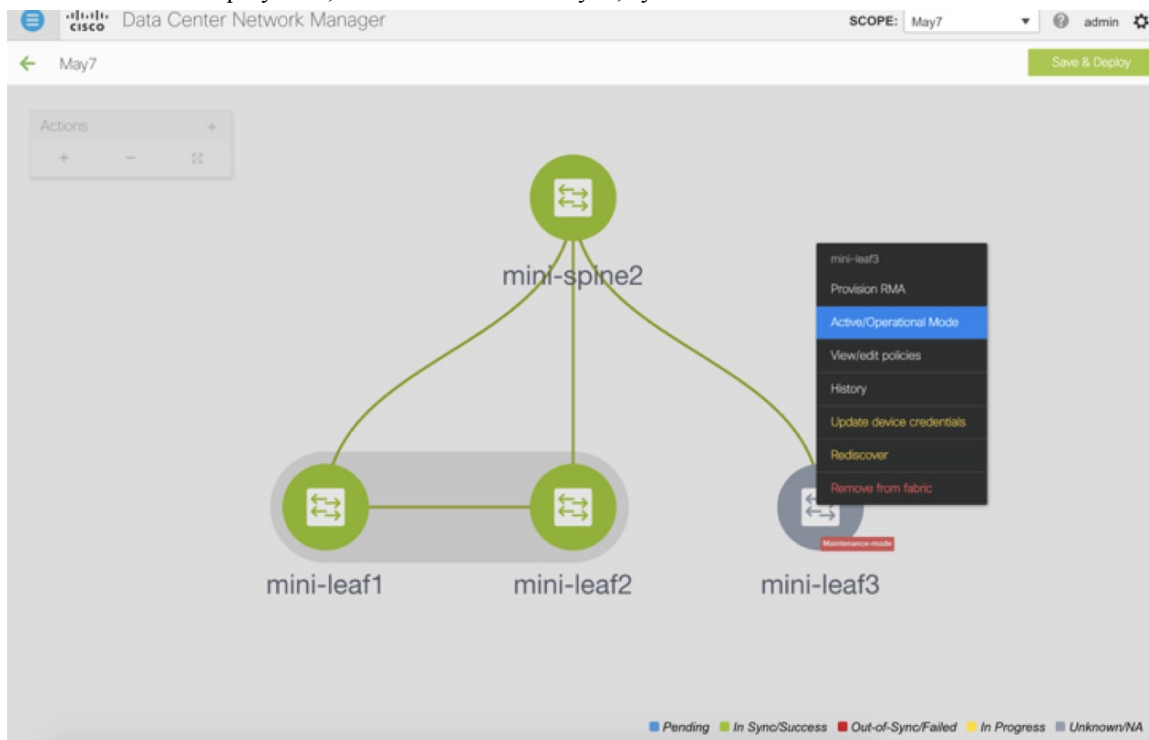
**Step 4** The Cisco DCNM rediscovers the new device (or you can manually choose **Discovery > Rediscover**).

**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.
- Create breakout and unbreakout ports.
- Shut down and bring up interfaces.
- Rediscover ports and view interface configuration history.
- Apply host policies on interfaces and vPCs. For example, int\_trunk\_host\_11\_1, int\_access\_host\_11\_1, and so on.
- View interface information such as its admin status, operation status, reason, policy, speed, MTU, mode, VLANs, IP/Prefix, VRF, port channel, and the neighbor of the interface.



### Note

- The **Neighbor** column provides details of connected switches that are discovered, intent links, and Virtual Machine Manager (VMM) connectivity. You can navigate to the **Switch** dashboard of the corresponding switch by clicking it. However, intent links and VMM links are not hyperlinked and you cannot navigate to the corresponding dashboard.
- 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.
Breakout, Unbreakout	Allows you to <i>breakout</i> an interface or unbreakout interfaces that are in <i>breakout</i> 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.



Field	Description
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.

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, and Subinterface. The respective interface ID field (Port-channel ID, vPC ID, Loopback ID, or Subinterface 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.

## Deleting Interfaces

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** to delete the interface.

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.

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.

## Creating Networks for the Standalone Fabric

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




### Select a Fabric

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

Standalone ▾

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



Fabric Selected: Standalone

Networks Selected 0 / Total 0  

<input type="checkbox"/>	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
No data available							

+ ✍ ✕ 🔄 📄
Show All ▾ 

- 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.

**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.

## Create Network



\* Network ID

\* Network Name

\* VRF Name  +

Layer 2 Only

\* Network Template  ▼

\* Network Extension Template  ▼

VLAN ID

▼ Network Profile

General

Advanced

IPv4 Gateway/NetMask  ? *example 192.0.2.1/24*

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

Vlan Name  ?

Interface Description  ?

MTU for L3 interface  ? *[68-9216]*

IPv4 Secondary GW1  ? *example 192.0.2.1/24*

IPv4 Secondary GW2  ? *example 192.0.2.1/24*

Create Network

▼ Network Profile

General

Advanced

ARP Suppression  ?

Ingress Replication  ? *Read-only per network, Fabric-wide setting*

Multicast Group Address  ?

DHCPv4 Server 1  ? *DHCP Relay IP*

DHCPv4 Server 2  ? *DHCP Relay IP*

DHCPv4 Server VRF  ?

Loopback ID for DHCP Relay interface  ?

Routing Tag  ? *[0-4294967295]*

TRM Enable  ? *Enable Tenant Routed Multicast*

L2 VNI Route-Target Both Enable  ?

Enable L3 Gateway on Border  ?

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

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

Fabric Selection > Network / VRF Selection > Network / VRF Deployment > VRF View | Continue

Fabric Selected: Standalone

Networks Selected 1 / Total 1

Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
<input checked="" type="checkbox"/> MyNetwork_30000	30000	MyVRF_50000	20.10.1.1/24		NA	

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.

Networks

Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet
<input type="checkbox"/> MyNetwork_30000	30000	MyVRF_50000	20.10.1.1/24
<input type="checkbox"/> MyNetwork_30001	30001	MyVRF_50000	

.CSV

A	B	C	D
fabric	vrf	networkName	networkId
Standalone	MyVRF_50000	MyNetwork_30000	30000
Standalone	MyVRF_50000	MyNetwork_30001	30001

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.

Networks Selected 0 / Total 4

Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	Status	VLAN ID
<input type="checkbox"/> MyNetwork_30000	30000	MyVRF_50000	20.10.1.1/24	NA	
<input type="checkbox"/> MyNetwork_30001	30001	MyVRF_50000		NA	

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.

Networks Selected 0 / Total 4

Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
MyNetwork_30000	30000	MyVRF_50000	20.10.1.1/24		NA	
MyNetwork_30001	30001	MyVRF_50000			NA	
MyNetwork_30002	30002	MyVRF_50000	20.10.4.1/24		NA	
MyNetwork_30003	30003	MyVRF_50000			NA	

## Editing Networks for the Standalone Fabric

1. Click **Control** > **Networks & VRFs** (under Fabrics submenu). The Networks & VRFs screen comes up.
2. Click **Continue**. The **Select a Fabric** screen is displayed.
3. From the **Select a Fabric** drop-down list, select the fabric *Standalone*, and click **Continue** on the top right part of the screen. The Networks page is displayed. This page lists the networks that are created for the fabric.
4. Select the network and click the **Edit** option at the top left part of the screen.

Fabric Selected: MSD-Parent-Fabric

Networks Selected 1 / Total 1

Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
<input checked="" type="checkbox"/> MyNetwork_30000	30000	MyVRF_50000	20.10.1.1/24		NA	

The **Edit Network** screen comes up.

Edit Network
✕

---

▼ Network Information

\* Network ID

\* Network Name

\* VRF Name

Layer 2 Only

\* Network Template

\* Network Extension Template

VLAN ID

---

▼ Network Profile

General

Advanced

IPv4 Gateway/NetMask  ? example 192.0.2.1/24

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

Vlan Name  ?

Interface Description  ?

MTU for L3 interface  ? [68-9216]


IPv4 Secondary GW1  ? example 192.0.2.1/24

IPv4 Secondary GW2  ? example 192.0.2.1/24

5. Update the fields in the **General** and **Advanced** tabs of the **Network Profile** section as needed.
6. Click **Save** at the bottom right part of the screen to save the updates.

## Creating VRFs for the Standalone Fabric

1. From the Networks page, click the **VRF View** button at the top right part of the screen to create VRFs. (If you have freshly logged in to DCNM, do the following:  
 Click **Control > Networks & VRFs**.  
 Click **Continue** in the LAN Fabric Provisioning page.  
 Choose the fabric (*Standalone*) from the drop-down list and click **Continue** to reach the Networks page.  
 Click **VRF View** at the top right part of the Networks page).  
 The VRFs page comes up. The page lists the list of VRFs created for the fabric. Initially, this page has no entries. One VRF is already created for this fabric. Let us create one more VRF.



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

### Create VRF ✕

▼ VRF Information

\* VRF ID

\* VRF Name

\* VRF Template

\* VRF Extension Template

---

▼ VRF Profile

General

Advanced

VRF Vlan Name

VRF Intf Description

VRF Description

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.

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

**Enable IPv6 link-local Option** – Enables the IPv6 link-local option under the VRF SVI.

**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:

## Create VRF



▼ VRF Information

\* VRF ID

\* VRF Name

\* VRF Template

\* VRF Extension Template

▼ VRF Profile

General

Advanced

VRF Vlan Name

VRF Intf Description

VRF Description

Create VRF

## Advanced tab:

▼ VRF Profile

General

Advanced

Routing Tag

Redistribute Direct Route Map

Max BGP Paths

Max iBGP Paths

TRM Enable

Is RP External

RP Address

RP Loopback ID

Underlay Mcast Add...

Overlay Mcast Groups

Enable IPv6 link-loc...

Advertise Host Routes


Advertise Default Route

Create VRF

3. Click **Create VRF**.

The *MyVRF\_50001* VRF is created and appears on the VRFs page.

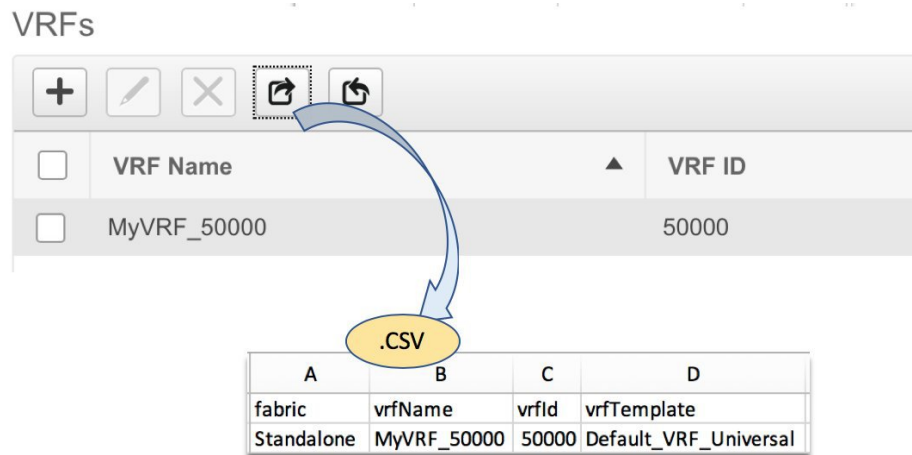




### Export and Import VRF Information

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.

VRFs Selected 0 / Total 2  

     Show  


<input type="checkbox"/>	VRF Name	VRF ID	Status
<input type="checkbox"/>	MyVRF_50000	50000	NA
<input type="checkbox"/>	MyVRF_50001	50001	NA

## 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. Click **Control > Networks & VRFs** (under Fabrics submenu). The Networks & VRFs screen comes up.
3. Click **Continue**. The **Select a Fabric** screen is displayed.
4. 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.
5. Click the **VRF View** at the top right part of the screen. The VRFs page appears.

Fabric Selected: New7200

VRFs Selected 0 / Total 2  

     Show  

<input type="checkbox"/>	VRF Name	VRF ID	Status
<input type="checkbox"/>	MyVRF_50000	50000	NA
<input type="checkbox"/>	MyVRF_50001	50001	NA

6. Select the **VRF** and click the **Edit** option at the top left part of the screen. The **Edit VRF** screen comes up.

Edit VRF
✕

---

▼ VRF Information

\* VRF ID

\* VRF Name

\* VRF Template

VRF Extension Template

---

▼ VRF Profile

General

Advanced

VRF Vlan Name  ?

VRF Intf Description  ?

VRF Description  ?

7. Update the fields in the **General** and **Advanced** tabs of the **VRF Profile** section as needed.
8. 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. Go to the Select a Fabric page.

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

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

OR

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

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

For an MSD fabric, you can either choose the MSD fabric or the member fabric. If you choose the MSD fabric, you can view all member fabrics in the same topology screen. So, you can provision networks from a single topology screen, one member fabric at a time.

The Networks page comes up.

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



**Note** You can edit or delete networks from this screen.

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).



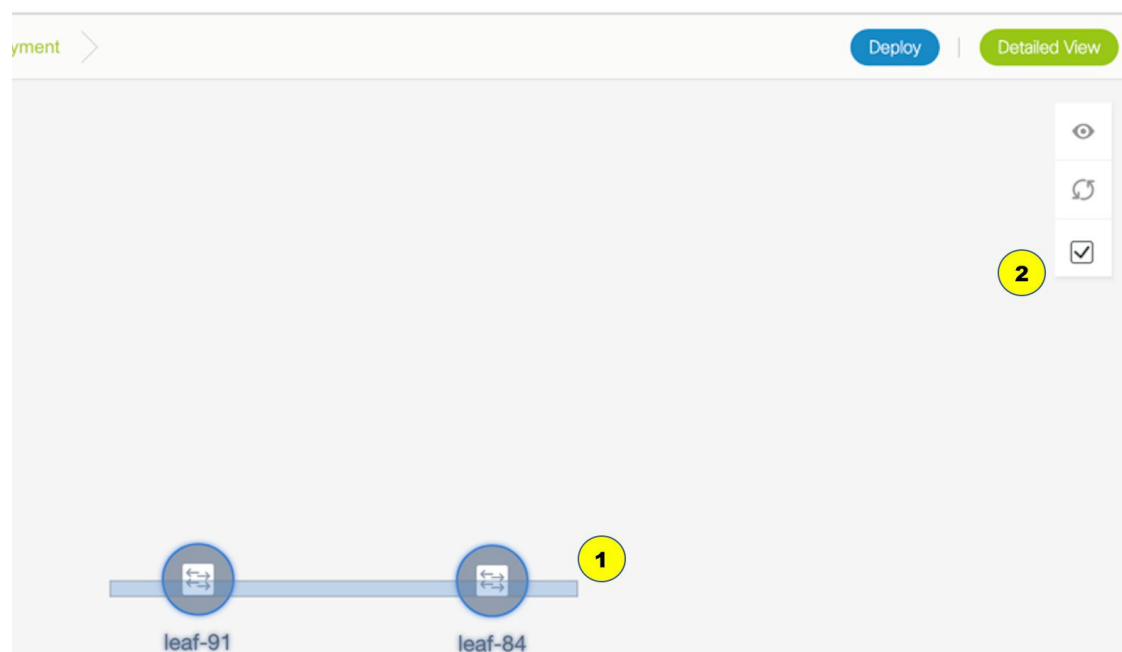
**Note** In an MSD fabric, all member fabrics are visible from this screen.

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.

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

4. Double-click a switch 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.

## Network Attachment - Attach networks for given switch(es) ✕

Fabric Name: Standalone

### Deployment Options

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

MyNetwork_30000		MyNetwork_30001				
<input type="checkbox"/>	Switch ▲	VLAN	Interfaces	CLI Freeform	Status	
<input type="checkbox"/>	n9k-16-leaf	2300	...	Freeform config	NA	

[Save](#)

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.

5. Select the other network tab and make the same selections.

- 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.

- 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.

## Preview Configuration

Select a Switch:

Select a Network

Generated Configuration:

```
configure profile MyVRF_50000
vlan 2000
vn-segment 50000
interface vlan2000
vrf member myvrf_50000
ip forward
ipv6 forward
no ip redirects
no ipv6 redirects
mtu 9216
no shutdown
vrf context myvrf_50000
vni 50000
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
address-family ipv6 unicast
route-target both auto
route-target both auto evpn
router bgp 60000
vrf myvrf_50000
address-family ipv4 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redist-subnet
maximum-paths ibgp 2
address-family ipv6 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redist-subnet
maximum-paths ibgp 2
interface nve1
member vni 50000 associate-vrf
configure terminal
apply profile MyVRF_50000
```

**MyVRF\_50000  
Configuration**

## Preview Configuration

Select a Switch:

n9k-16-leaf

Select a Network

MyNetwork\_30000

Generated Configuration:

```
vrf myvrf_50000
address-family ipv4 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redis-subnet
maximum-paths ibgp 2
address-family ipv6 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redis-subnet
maximum-paths ibgp 2
interface nve1
member vni 50000 associate-vrf
configure terminal
apply profile MyVRF_50000
```

```
configure profile MyNetwork_30000
vlan 2300
vn-segment 30000
interface vlan2300
vrf member myvrf_50000
fabric forwarding mode anycast-gateway
no shutdown
interface nve1
member vni 30000
mcast-group 239.1.1.0
evpn
vni 30000 l2
rd auto
route-target import auto
route-target export auto
configure terminal
apply profile MyNetwork_30000
```

```
interface ethernet1/11
switchport trunk allowed vlan add 2300
interface ethernet1/10
switchport trunk allowed vlan add 2300
```

**MyNetwork\_30000  
Configuration**

**Interfaces Configuration**

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.

- 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.

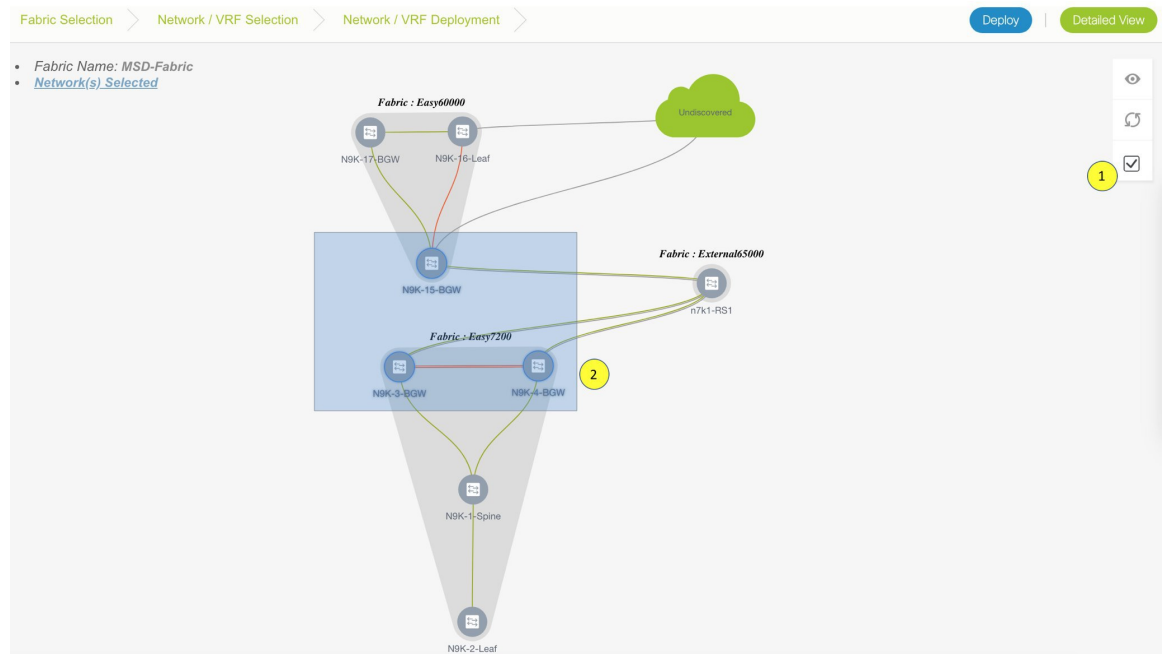
### Network Deployment for an MSD Fabric

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.

Name	Switch	Ports	Status	Fabric Name	Role
MyNetwork_30000	N9k-15-bgw		NA	new60000	border
MyNetwork_30001	N9k-15-bgw		NA	new60000	border
MyNetwork_30001	n9k-16-leaf	Ethernet1/1	DEPLOYED	new60000	leaf
MyNetwork_30000	n9k-16-leaf	Ethernet1/10,Ethernet1/11	DEPLOYED	new60000	leaf

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

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

Click **Control > Networks & VRFs**.

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

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

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

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

Fabric Selection > Network / VRF Selection > Network / VRF Deployment > Network View | Continue

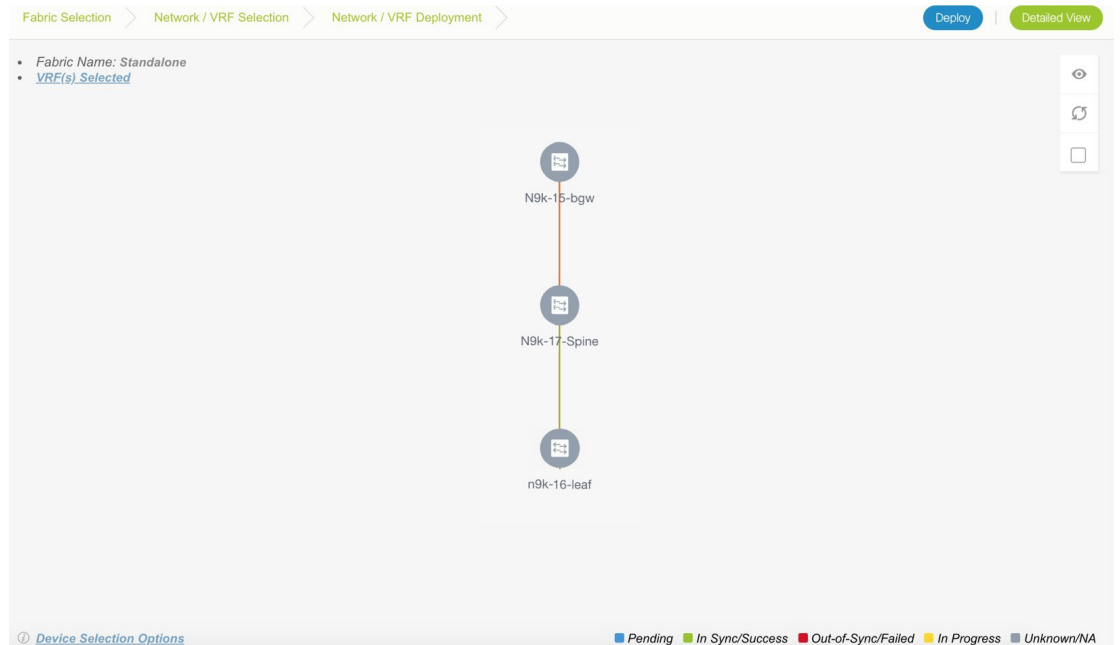
Fabric Selected: Standalone

VRFs Selected 0 / Total 2

<input type="checkbox"/>	VRF Name	VRF ID	Status
<input type="checkbox"/>	MyVRF_50000	50000	NA
<input type="checkbox"/>	MyVRF_50001	50001	NA

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.

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.

## VRF Attachment - Attach VRFs for given switch(es).



Fabric Name: Standalone

## Deployment Options

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

MyVRF_50000		MyVRF_50001			
<input type="checkbox"/>	Switch	▲	VLAN	CLI Freeform	Status
<input type="checkbox"/>	n9k-16-leaf		2000	Freeform config	NA

Save

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).

## Preview Configuration



Select a Switch:

n9k-16-leaf ▼

Select a VRF

MyVRF\_50000 ▼

Generated Configuration:

```
configure profile MyVRF_50000
vlan 2000
vn-segment 50000
interface vlan2000
vrf member myvrf_50000
ip forward
ipv6 forward
no ip redirects
no ipv6 redirects
mtu 9216
no shutdown
vrf context myvrf_50000
vni 50000
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
address-family ipv6 unicast
route-target both auto
route-target both auto evpn
router bgp 60000
vrf myvrf_50000
address-family ipv4 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redis-subnet
maximum-paths ibgp 2
address-family ipv6 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redis-subnet
maximum-paths ibgp 2
interface nve1
member vni 50000 associate-vrf
configure terminal
apply profile MyVRF_50000
```

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.

Fabric Selection > Network / VRF Selection > Network / VRF Deployment > Topology View

Fabric Name: Standalone VRF(s) Selected Selected 0 / Total 4

<input type="checkbox"/>	Name	Switch	Ports	Status	Fabric Name	Role
<input type="checkbox"/>	MyVRF_50000	n9k-15-BL		NA	Easy60000	leaf
<input type="checkbox"/>	MyVRF_50000	n9k-16-leaf		DEPLOYED	Easy60000	leaf
<input type="checkbox"/>	MyVRF_50001	n9k-15-BL		NA	Easy60000	leaf
<input type="checkbox"/>	MyVRF_50001	n9k-16-leaf		DEPLOYED	Easy60000	leaf

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. Choose **Control > Networks and VRFs**.
2. In the **Select a Fabric** page, click **Continue** (at the top right part of the screen). The Networks page comes up.
3. Select the networks that you want to undeploy and click Continue. The topology view 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 > Networks and VRFs**.

2. In the Select a Fabric page, click **Continue** (at the top right part of the screen). The Networks page comes up.
3. Click the **VRF View** button (at the top right part of the screen) to go to the VRFs screen.
4. Select the VRFs that you want to undeploy and click **Continue**. The *Topology View* page comes up.
5. 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).
6. 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.
7. Click **Save** (at the bottom right part of the screen) to initiate the undeployment of the VRFs. The topology view comes up again.




---

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

---

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

## 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.

### 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:

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.

BGP Send-Community-Both Option – Select the checkbox to send standard and extended BGP communities to BGP peers. If the checkbox is not selected, only the extended community is sent.

#### Advanced tab

General | **Advanced** | Resources | DCI

\* vPC Peer Link VLAN  ? VLAN for vPC Peer Link SVI

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

**Enable NX-API** - Specifies enabling of NX-API on HTTPS. This check box is checked by default.

**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. This check box is checked by default.

#### 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.

General | Advanced | **Resources** | DCI

\* Subinterface Dot1q Range  ? Per Border Dot1q Range For VRF Lite Connectivity

\* Underlay Routing Loopback IP Range  ? Typically Loopback0 IP Address Range

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

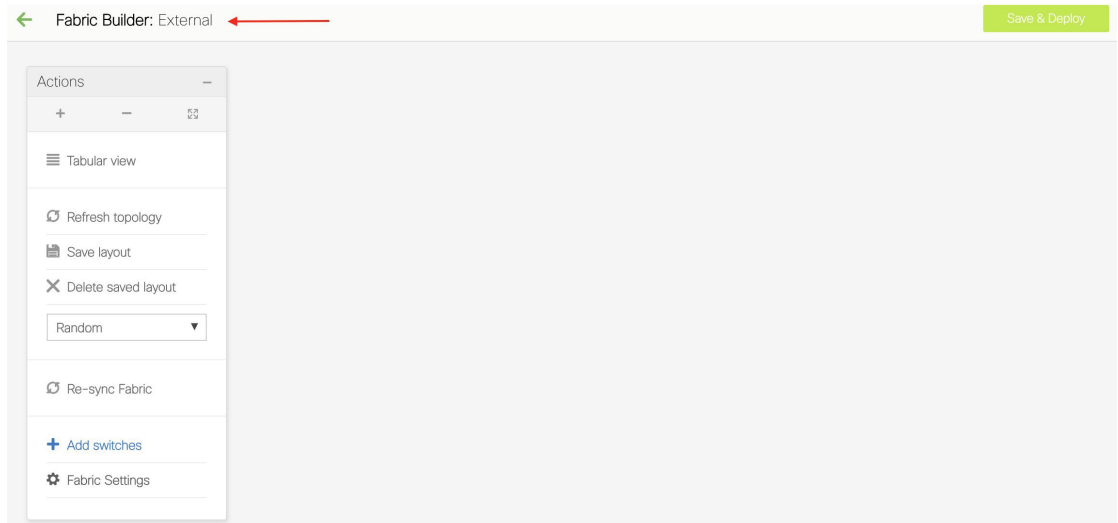
General | Advanced | Resources | **DCI**

\* DCI Subnet IP Range  ? Address range to assign P2P DCI Links

\* Subnet Target Mask  ? Target Mask for Subnet Range

#### 4. 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.

## 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>Fabric Builder**, click **Tabular View**, and select a switch in the **Name** column or select **Control>Fabric Builder** and right-click on the device.
2. Click **View/Edit Policies** and click on + to add a new policy. The **Add Policy** window comes up.
3. Add a PTI with the required configuration CLIs using the **switch\_freeform\_config** template and click **Save**.
4. Select the created policy and click **Deploy** to deploy the configuration to the switch(es).

## 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.

---



---

**Note** 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.

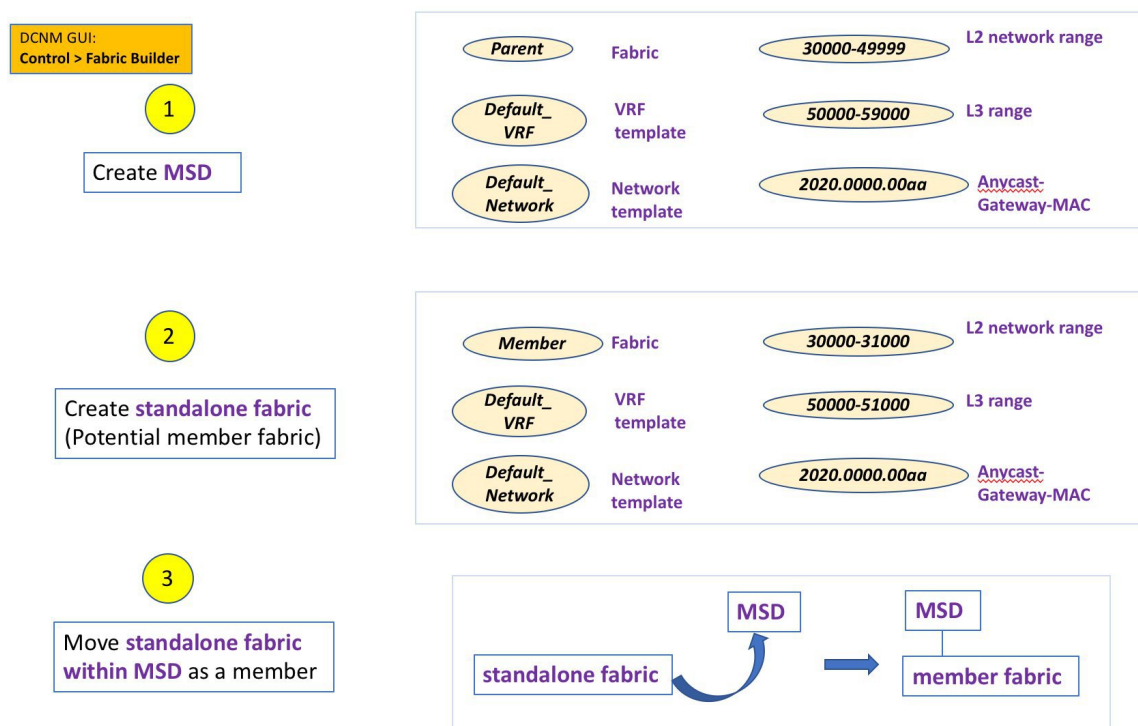
Fabric instance values can only be edited or updated in the fabric context from the VRFs and Networks window. The appropriate fabric should be selected in the **SCOPE** drop-down list to edit the fabric instance values. Some of the examples of fabric instance variables are BGP ASN, Multicast group per network or VRF, etc. For information about editing multicast group address, see [Editing Networks in the Member Fabric, on page 107](#).

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:






DCNM GUI:  
Control > Networks & VRFs

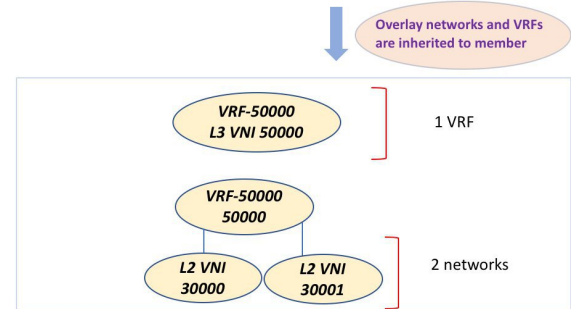
4

Create **networks** and **VRFs** in **MSD fabric**




5


The **networks** and **VRFs** automatically get inherited to the member fabric



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.

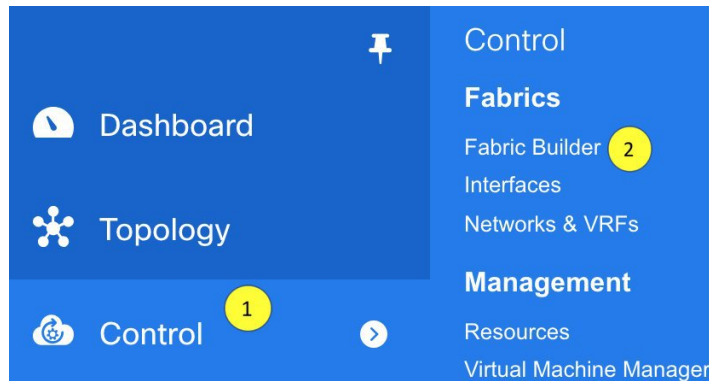
### Creating an MSD Fabric and Associating Member Fabrics to It

The process is explained in two steps:

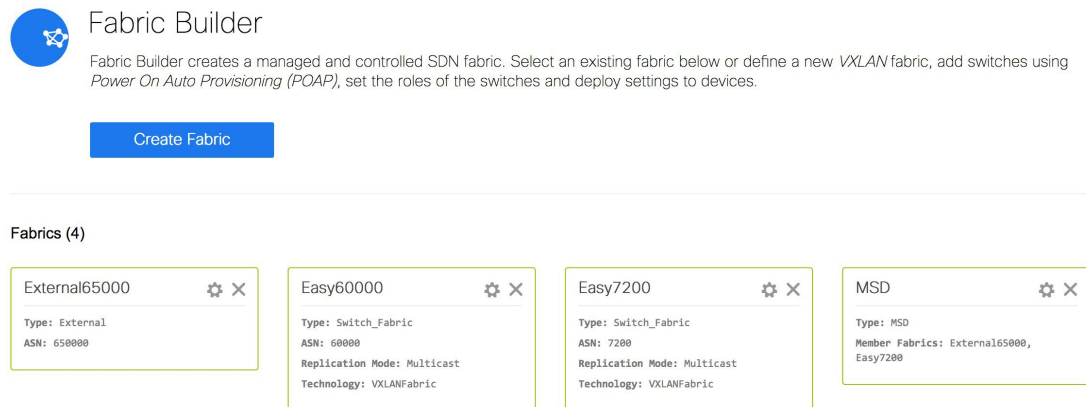
1. Create an MSD fabric.
2. Create a new standalone fabric and move it under the MSD fabric as a member fabric.

## 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.

Add Fabric ✕

\* Fabric Name :

\* Fabric Template

General DCI Resources

L2 Segment ID Range  ? L2 Segment ID Range

L3 Partition ID Range  ? L3 Partition ID Range

\* VRF Template  ? Default Overlay VRF Template For Leafs

\* Network Template  ? Default Overlay Network Template For Leafs

\* VRF Extension Template  ? Default Overlay VRF Template For Borders

\* Network Extension Template  ? Default Overlay Network Template For Borders

Anycast-Gateway-MAC  ? Shared MAC address for all leaves

\* Multisite Routing Loopback Id  ? 0-512

The fields in the screen are explained:

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

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

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

**VRF Template** - Default VRF template.

**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.

General DCI Resources

DCI Subnet IP Range  ? Address range to assign P2P DCI Links

Subnet Target Mask  ? Target Mask for Subnet Range

\* Deploy Border Gateway Method  ? Deploy Border Gateway Method

MS Route Server List  ? Multi-Site Router-Server peer list e.g. 128.89.0.1, 1

BGP ASN of Route Server(s) one for each route server  ? 1-4294967295 | 1-65535[.0-65535]

The fields are:

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

**Deploy Border Gateway 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.

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.

Fabrics (5)

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**.

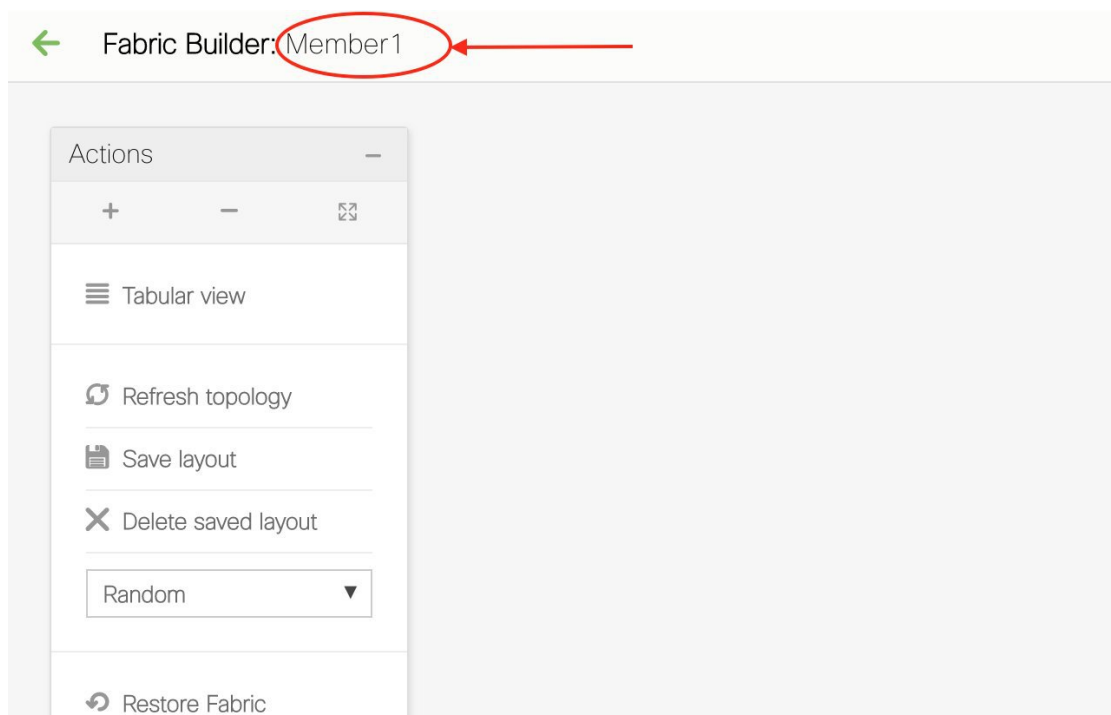
Ensure that the **Anycast Gateway MAC**, the **Network Template** and the **VRF Template** field values are the same as the MSD fabric. Else, member fabric movement to the MSD fail.

Other pointers:

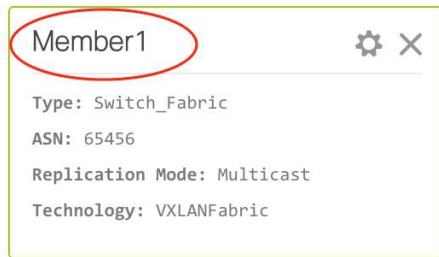
- Ensure that the Anycast Gateway MAC, the Network Template and the VRF Template field values are the same as the MSD fabric. Else, member fabric movement to the MSD fail.
- The member fabric should have a Site ID configured and the Site ID must be unique among the members.
- The BGP AS number should 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.



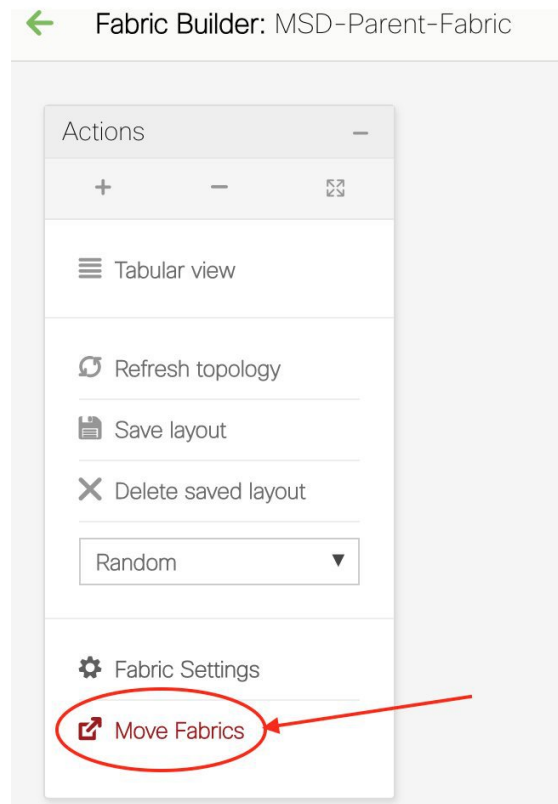
### Moving the Member1 Fabric Under MSD-Parent-Fabric

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.



## Move Fabric



Selected 0 / Total 2

	Fabric Name ▲	Fabric State
<input type="radio"/>	Member1	standalone
<input type="radio"/>	Test	standalone

Add

Remove

Cancel


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.

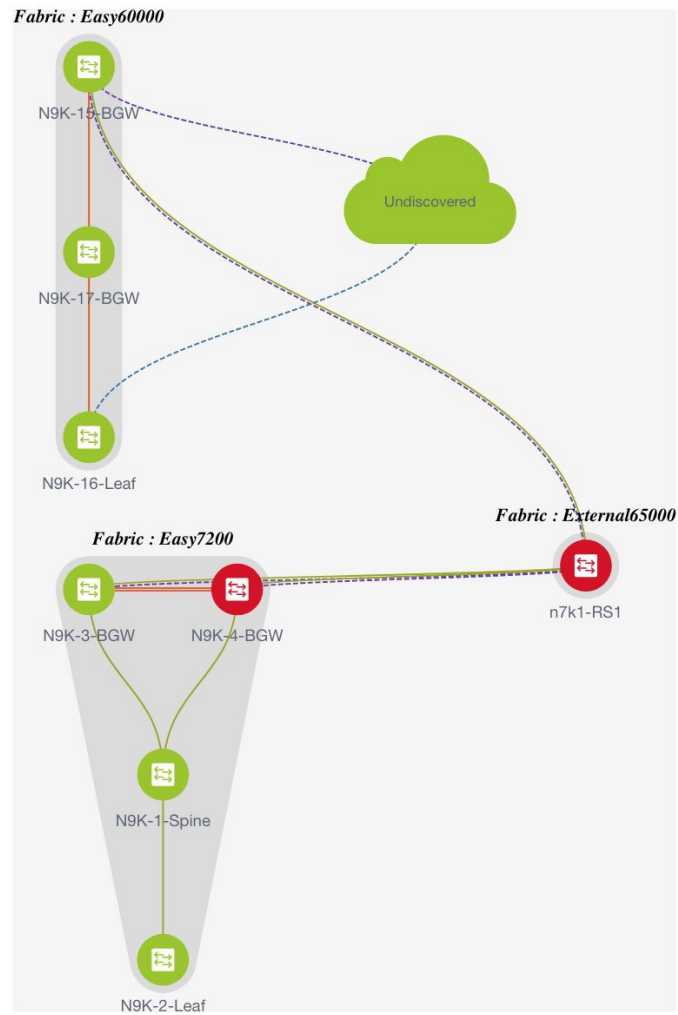
You can see that *Member1* is now added to MSD fabric and is displayed in the **Member Fabrics** field.



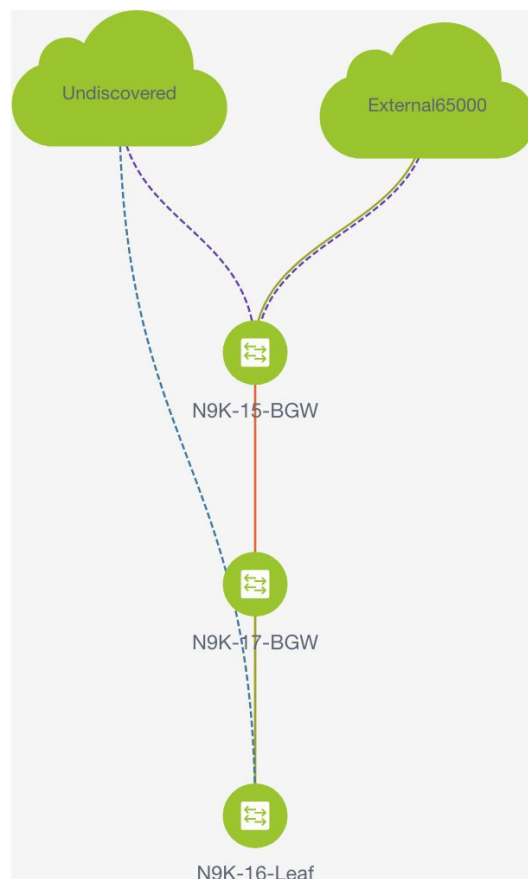
### MSD Fabric Topology View Pointers

- **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.



### 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 & VRFs** (under Fabrics submenu). The **Networks & VRFs** page comes up.
2. Click **Continue**. The Select a Fabric page comes up.



## Select a Fabric


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

You can click the **Select a Fabric** drop-down box to see the list of fabrics.


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

## Select a Fabric

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



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




## Select a Fabric

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



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.

Create Network
✕

▼ Network Information

\* Network ID

\* Network Name

\* VRF Name  +

Layer 2 Only

\* Network Template

\* Network Extension Template

VLAN ID

---

▼ Network Profile

General

Advanced

IPv4 Gateway/NetMask  ? example 192.0.2.1/24

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

Vlan Name  ?

Interface Description  ?

MTU for L3 interface  ? [68-9216]

Create Network

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:



## Create Network



\* Network ID

\* Network Name

\* VRF Name  +

Layer 2 Only

\* Network Template

\* Network Extension Template

VLAN ID

## Network Profile

General

Advanced

IPv4 Gateway/NetMask  ? *example 192.0.2.1/24*

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

Vlan Name  ?

Interface Description  ?

MTU for L3 interface  ? *[68-9216]*

IPv4 Secondary GW1  ? *example 192.0.2.1/24*

IPv4 Secondary GW2  ? *example 192.0.2.1/24*

**Create Network**

## Advanced tab:

## Network Profile

General

Advanced

ARP Suppression  ?

\* DHCPv4 Server 1  ? *DHCP Relay IP*

DHCPv4 Server 2  ? *DHCP Relay IP*

\* DHCPv4 Server VRF  ?

Loopback ID for DHCP Relay interface  ?

Routing Tag  ? *[0-4294967295]*

TRM Enable  ? *Enable Tenant Routed Multicast*

L2 VNI Route-Target Both Enable  ?

**Create Network**

5. Click **Create Network**. A message appears at the bottom right part of the screen indicating that the network is created. The new network (*MyNetwork\_30000*) appears on the Networks page that comes up.

Fabric Selected: MSD-Parent-Fabric

Networks Selected 1 / Total 1

	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
<input checked="" type="checkbox"/>	MyNetwork_30000	30000	MyVRF_50000	20.10.1.1/24		NA	

## 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.

Fabric Selected: MSD-Parent-Fabric

Networks Selected 1 / Total 1

Show All

Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
MyNetwork_30000	30000	MyVRF_50000	20.10.1.1/24		NA	

The Edit Network screen comes up.

### Edit Network

▼ Network Information

\* Network ID

\* Network Name

\* VRF Name

Layer 2 Only

\* Network Template

\* Network Extension Template

VLAN ID

▼ Network Profile

General

Advanced

IPv4 Gateway/NetMask  ? example 192.0.2.1/24

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

Vlan Name  ?

Interface Description  ?

MTU for L3 interface  ? [68-9216]

IPv4 Secondary GW1  ? example 192.0.2.1/24

IPv4 Secondary GW2  ? example 192.0.2.1/24

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

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

Networks

Selected 0 / Total 1

Show All

<input type="checkbox"/>	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
<input type="checkbox"/>	MyNetwork_30000	30000	MyVRF_50000	20.10.1.1/24		NA	

### 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*.

- Select the network and click the **Edit** option at the top left part of the window. The **Edit Network** window comes up.
- Click the **Advanced** tab in the **Network Profile** section, update the multicast group address, and click **Save**.

Edit Network
✕

---

▼ Network Information

\* Network ID

\* Network Name

\* VRF Name

Layer 2 Only

\* Network Template

\* Network Extension Template

VLAN ID

---

▼ Network Profile

General

1 Advanced

ARP Suppression  ?

Ingress Replication  ? *Read-only per network, Fabric-wide setting*

Multicast Group Address  2 ?

\* DHCPv4 Server 1  ? *DHCP Relay IP*

DHCPv4 Server 2  ? *DHCP Relay IP*

\* DHCPv4 Server VRF  ?

3



**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. Click **Control > Networks & VRFs**. The Networks & VRFs page comes up.
  - b. Click **Continue**. The Select a Fabric page comes up.
  - c. Choose the MSD fabric (*MSD-Parent-Fabric*) from the drop-down box and click **Continue**. The Networks page comes up.
  - d. 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 checkbox to enable the IPv6 link-local option.

**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:

Create VRF
✕

---

▼ VRF Information

\* VRF ID

\* VRF Name

\* VRF Template  ▼

\* VRF Extension Template  ▼

---

▼ VRF Profile

General

Advanced

VRF Vlan Name  ?

VRF Intf Description  ?

VRF Description  ?

**Advanced tab:**

### ▼ VRF Profile

General	Advanced
Routing Tag	12345 <small>[0-4294967295]</small>
Redistribute Direct Route Map	FABRIC-RMAP-REDIST-SUBNET
Max BGP Paths	1 <small>[1-64]</small>
Max iBGP Paths	2 <small>[1-64]</small>
TRM Enable	<input type="checkbox"/> <small>Enable Tenant Routed Multicast</small>
Is RP External	<input type="checkbox"/> <small>Is RP external to the fabric?</small>
RP Address	224.0.0.2 <small>IPv4 Address</small>
RP Loopback ID	3 <small>0-1023</small>
Underlay Mcast Add...	224.0.0.10 <small>IPv4 Multicast Address</small>
Overlay Mcast Groups	224.0.0.0/8 <small>224.0.0.0/8 to 239.255.255.255/8</small>
Enable IPv6 link-loc...	<input type="checkbox"/> <small>Enables IPv6 link-local Option under VRF SVI</small>
Advertise Host Routes	<input type="checkbox"/> <small>Flag to Control Advertisement of /32 and /128 Routes to Edge Routers</small>
Advertise Default Route	<input checked="" type="checkbox"/> <small>Flag to Control Advertisement of Default Route Internally</small>

Create VRF

### 5. Click **Create VRF**.

The *MyVRF\_50000* VRF is created and appears on the VRFs page.

Fabric Selected: MSD-Parent-Fabric

VRFs Selected 1 / Total 1

+	✎	✕	↺	↻	Show	All	▼	⌵
<input type="checkbox"/>	VRF Name	VRF ID	Status					
<input checked="" type="checkbox"/>	MyVRF_50000	50000	NA					

### Editing VRFs in the MSD Fabric

- 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.

Fabric Selected: MSD-Parent-Fabric

VRFs Selected 1 / Total 1

+	✎	✕	↺	↻	Show	All	▼	⌵
<input type="checkbox"/>	VRF Name	VRF ID	Status					
<input checked="" type="checkbox"/>	MyVRF_50000	50000	NA					

The Edit VRF screen comes up.

Edit VRF ✕

▼ VRF Information

\* VRF ID

\* VRF Name

\* VRF Template

VRF Extension Template

---

▼ VRF Profile

General  
 Advanced

VRF Vlan Name  ?

VRF Intf Description  ?

VRF Description  ?

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. From the main menu, click **Control > Networks & VRFs**. In the Networks & VRFs page, click **Continue**.
2. Choose *Member1* in the Select a Fabric drop-down box. and click **Continue**. The Networks page comes up.
3. Click the **VRF View** button. On the VRFs page, you can see that the VRF created for the MSD is inherited to its member.

Fabric Selected: Member1

VRFs Selected 0 / Total 1

	VRF Name	VRF ID	Status
<input type="checkbox"/>	MyVRF_50000	50000	NA



### 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](#) .

---

## 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.

## Brownfield Deployment-Transitioning VXLAN Fabric Management to DCNM

This document explains Brownfield deployments, wherein you transition your VXLAN BGP EVPN fabric management to DCNM. The transition involves migrating existing networks configurations to DCNM.

Typically, your fabric is created and managed through manual CLI configuration or custom automation scripts. Now, you want to start managing the fabric through DCNM. After the migration, the fabric underlay and overlay networks will be managed by DCNM.

The migration procedure only supports VXLAN BGP EVPN networks that use the best practices mentioned in the Prerequisites section.

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

For information about the MSD fabric migration, see *Migrating an MSD Fabric with Border Gateway Switches*.




---

**Note** The Brownfield deployment section is applicable for the **Easy\_Fabric\_11\_1** template.

---

### Prerequisites

- DCNM-supported NX-OS software versions. For details, refer *Cisco DCNM Release Notes, Release 11.1(1)*.
- Underlay routing protocol is OSPF or IS-IS.
- The supported underlay is based on the DCNM 10.2(1) POAP template's best practices for the VXLAN fabric (dcnm\_ip\_vxlan\_fabric\_templates.10.2.1.ST.1.zip) available on Cisco.com.
- The following fabric-wide loopback interface IDs must not overlap:
  - Routing loopback interface for IGP/BGP.
  - VTEP loopback ID
  - Underlay rendezvous point loopback ID if ASM is used for multicast replication.
- BGP configuration uses the 'router-id', which is the IP address of the routing loopback interface.
- If the iBGP peer template is configured, then it must be configured on the leaf switches and route reflectors. The template name that needs to be used between leaf and route reflector should be identical.

- The BGP route reflector and multicast rendezvous point (if applicable) functions are implemented on spine switches. Leaf switches do not support the functions.
- Install DCNM 11.1(1) release software. Refer the Installation Guide for more details. Log in to DCNM and set the default LAN Credentials when prompted.
- Familiarity with the DCNM 11.1(1) fabric management and monitoring features before initiating the migration process.
- Familiarity with VXLAN BGP EVPN fabric concepts and functioning of the fabric from the DCNM perspective.
- Fabric switch nodes are operationally stable and functional and all fabric links are up.
- vPC switches and the peer links are up before the migration. Ensure that no configuration updates are in progress or changes pending.
- Create an inventory list of the switches in the fabric with their IP addresses and credentials. DCNM uses this information to connect to the switches.
- Shut down any other controller software you are using presently so that no further configuration changes are made to the VXLAN fabric. Alternatively, disconnect the network interfaces from the controller software (if any) so that no changes are allowed on the switches.
- In the Cisco DCNM Release 11.1(1), a brownfield import captures all the overlay network or VRF configurations found on the switch in the respective overlay freeform config.

These freeform configs will have configs that are already part of the profiles and any extra configurations. This action creates a double intent scenario, that is, the configurations are captured twice in DCNM to avoid any network outages during conversion of regular CLI configuration on NX-OS devices to config-profile based templates for deployed networks.

Similarly, the double intent is created during Brownfield migration if the switches are running Cisco NX-OS Release 7.0(3)I7(6) or lower, and 9.2(3) or lower.

The following workarounds can be used to avoid issues with the double intent:

- Whenever the overlay parameters are updated, review the updated configurations present in the freeform configs such that they are consistent.
- We recommend that you contact Cisco Technical Assistance Center (TAC) to help you with removing the double intent via a script. The requirement is that all the switches in the fabric should be running the below versions:
  - Cisco NX-OS Release 7.0(3)I7(6) or higher
  - Cisco NX-OS Release 9.2(3) or higher
- All the overlay network and VRF profile parameters such as VLAN name and route map name should be consistent across all devices in the fabric for the brownfield migration to be successful.

### Guidelines and Limitations

- Fabric interfaces can be numbered or unnumbered.
- Various other interface types are supported.
- The following features are unsupported.

- eBGP underlay
  - BIDIR-PIM function
  - TRM
  - Border Spine or Border Gateway Spine
  - Layer 3 port channel
  - Configuration profiles present in the brownfield configurations (the expectation is that the overlays should be configured through regular CLIs).
- Take a backup of the switch configurations and save them before the migration.
  - No configuration changes (unless instructed to do so in this document) must be made to the switches until the migration is completed. Else, significant network issues can occur.
  - Migration to Cisco DCNM is only supported for Cisco Nexus 9000 switches.
  - Multi-line banner configuration on the switch is preserved in the switch\_freeform configuration, along with other configurations captured in the switch\_freeform configuration, if any.

### Procedure

Transitioning VXLAN fabric management to DCNM involves these steps.

1. Creating a new VXLAN BGP EVPN fabric in DCNM – This step creates a VXLAN fabric outline.
2. Initiating VXLAN fabric management transition to DCNM – This step adds switch instances to DCNM and initiates the transition.

### Creating a New VXLAN BGP EVPN Fabric

First, guidelines for updating the settings are noted. Then each VXLAN fabric settings tab is explained:

- Some values (BGP AS Number, OSPF, etc) are considered as reference points to your existing fabric, and the values you enter must match the existing fabric values.
- For some fields (such as IP address range, VXLAN ID range), the values that are auto-populated or entered in the settings are only used for future allocation. The existing fabric values are honored during migration.
- Some fields relate to new functions that may not exist in your existing fabric (such as advertise-pip). Enable or disable it as per your need.
- At a later point in time, after the fabric transition is complete, you can update settings if needed.

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. Click the **Replication** tab. Most of the fields are auto generated.

The screenshot shows the configuration interface for a fabric template. The 'Replication' tab is selected. The following fields are visible:

- Replication Mode:** Multicast (dropdown menu)
- Enable Tenant Routed Multicast:**  (checkbox)
- RP Mode:** asm (dropdown menu)
- Multicast Group Subnet:** 239.1.1.0/25 (text input)
- Rendezvous-Points:** 2 (dropdown menu)
- Underlay RP Loopback Id:** 254 (text input)
- Underlay Primary RP Loopback Id:** (grayed out text input)
- Underlay Backup RP Loopback Id:** (grayed out text input)

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

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

**Multicast Group Subnet** - The IP address prefix for multicast communication is used for post-migration allocation. The IP address prefix used in your existing fabric is honored during the transition.

A unique IP address is allocated from this group for each overlay network.

**Enable Tenant Routed Multicast** – *Do not* enable the check box. TRM is not supported for transitioning fabric management.

**Rendezvous-Points** - The RP count is only applicable post-migration. The existing RP configuration is honored when importing into the DCNM setup.

**RP mode** – Retain **asm** (for Any-Source Multicast [ASM]) mode. *Do not* change the selection to **bidir** since BIDIR-PIM is not supported for fabric migration.

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

**Underlay RP Loopback ID** – The loopback ID has to match your existing setup's loopback ID. This is the loopback ID used for the rendezvous point (RP), for multicast protocol peering purposes in the fabric underlay.

The other two fields are grayed out.

The next two fields are enabled if **Rendezvous-Points** is set to 4. However, the fabric can have only 2 RPs for the brownfield migration.

4. Click the **vPC** tab. Most of the fields are auto generated.

General	Replication	vPC	Advanced	Resources	Manageability	Bootstrap	Configuration Backup
			* vPC Peer Link VLAN	3600	<a href="#">?</a> VLAN for vPC Peer Link SVI		
			* vPC Peer Keep Alive option	management	<a href="#">?</a> Use vPC Peer Keep Alive with Loopback or Management		
			* vPC Auto Recovery Time	360	<a href="#">?</a> Auto Recovery Time In Seconds (Min:240, Max:3600)		
			* vPC Delay Restore Time	150	<a href="#">?</a> vPC Delay Restore Time For vPC links in seconds (Min:1, Max:3600)		
			vPC IPv6 ND Synchronize	<input checked="" type="checkbox"/>	<a href="#">?</a> Enable IPv6 ND synchronization between vPC peers		
			vPC advertise-pip	<input type="checkbox"/>	<a href="#">?</a> For Primary VTEP IP Advertisement As Next-Hop Of Prefix Routes		

**vPC Peer Link VLAN** - Enter the VLAN ID used for the vPC peer link SVI in the existing fabric.

**vPC Peer Keep Alive option** – Choose the management or loopback option, as used in the existing fabric. 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 only use IPv6 addresses on the management interface, you must use the loopback option.

During the transition, the switch configuration is not checked for the following fields in the vPC tab. The switch configurations will get updated if they are different.

**vPC Auto Recovery Time** - Specify the vPC auto recovery time-out period in seconds, as needed.

**vPC Delay Restore Time** - Specify the vPC delay restore period in seconds, as needed.

**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 as needed.

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

- Click the **Advanced** tab. Most of the fields are auto generated.

General	Replication	vPC	Advanced	Resources	Manageability	Bootstrap	Configuration Backup
			* VRF Template	Default_VRF_Universal	<a href="#">?</a> Default Overlay VRF Template For Leafs		
			* Network Template	Default_Network_Universal	<a href="#">?</a> Default Overlay Network Template For Leafs		
			* VRF Extension Template	Default_VRF_Extension_Universal	<a href="#">?</a> Default Overlay VRF Template For Borders		
			* Network Extension Template	Default_Network_Extension_Universa	<a href="#">?</a> Default Overlay Network Template For Borders		
			Site Id		<a href="#">?</a> For EVPN Multi-Site Support (Min:1, Max:16777216)		
			* Underlay Routing Loopback Id	0	<a href="#">?</a> 0-512		
			* Underlay VTEP Loopback Id	1	<a href="#">?</a> 0-512		
			* Link-State Routing Protocol Tag	UNDERLAY	<a href="#">?</a> Routing Process Tag		
			* OSPF Area Id	0.0.0.0	<a href="#">?</a> OSPF Area Id in decimal format or IP address format		
			* Power Supply Mode	ps-redundant	<a href="#">?</a> Default Power Supply Mode For The Fabric		
			* CoPP Profile	strict	<a href="#">?</a> Fabric Wide CoPP Policy		
			Enable VXLAN OAM	<input checked="" type="checkbox"/>	<a href="#">?</a> For Operations And Management Of VXLAN Fabrics		
			* Greenfield Cleanup Option	Disable	<a href="#">?</a> Switch Cleanup Without Reload When PreserveConfig=no		
			iBGP Peer-Template Config		<a href="#">?</a> Leaf to RR iBGP session establishment		
			Leaf Freeform Config		<a href="#">?</a> Additional CLIs For All Leafs As Captured From Show Running Configuration		
			Spine Freeform Config		<a href="#">?</a> Additional CLIs For All Spines As Captured From Show Running Configuration		

**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.

You must not change the templates when migrating. Only the Universal templates are supported for overlay migration.

**Site ID** - The ID for this fabric if you are moving this fabric within an MSD. You can update this field post-migration.

**Underlay Routing Loopback Id** - The loopback interface ID is populated as 0 since loopback0 is usually used for fabric underlay IGP peering purposes. This must match the existing configuration on the switches. This must be the same across all the switches.

**Underlay VTEP Loopback Id** - The loopback interface ID is populated as 1 since loopback1 is usually used for the VTEP peering purposes. This must match the existing configuration on the switches. This must be the same across all the switches where VTEPs are present.

**Link-State Routing Protocol Tag** - Enter the existing fabric's routing protocol tag in this field to define the type of network.

**OSPF Area ID** - The OSPF area ID of the existing fabric, if OSPF is used as the IGP within the fabric.

**Power Supply Mode** - Choose the appropriate power supply mode.

**CoPP Profile** - Choose the Control Plane Policing (CoPP) profile policy used in the existing fabric. By default, the strict option is populated.

**Enable VXLAN OAM** - Enables the VXLAN 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.

---

**Greenfield Cleanup Option** - Enable or disable the switch cleanup option for Greenfield switches. This is applicable post-migration when new switches are added.

**iBGP Peer-Template Config** - Add iBGP peer template configurations on the leaf switches and route reflectors to establish an iBGP session between the leaf switch and route reflector. Set this field based on switch configuration. If this field is blank, it implies that the iBGP peer template is not used. If the iBGP peer template is used, enter the peer template definition as defined on the switch. The peer template name on devices configured with BGP should be the same as defined here.

**Leaf Freeform Config** and **Spine Freeform Config** - You can enter these fields after fabric transitioning is complete, as needed.

6. Click the **Resources** tab.

General	Replication	vPC	Advanced	Resources	Manageability	Bootstrap	Configuration Backup
Static Underlay IP Address Allocation <input type="checkbox"/> ? <i>Checking this will disable Dynamic Underlay IP Address Allocations</i>							
* Underlay Routing Loopback IP Range	10.2.0.0/22			? Typically Loopback0 IP Address Range			
* Underlay VTEP Loopback IP Range	10.3.0.0/22			? Typically Loopback1 IP Address Range			
* Underlay RP Loopback IP Range	10.254.254.0/24			? Anycast or Phantom RP IP Address Range			
* Underlay Subnet IP Range	10.4.0.0/16			? Address range to assign Numbered and Peer Link SVI IP			
* Layer 2 VXLAN VNI Range	30000-49000			? Overlay Network Identifier Range (Min:1, Max:16777214)			
* Layer 3 VXLAN VNI Range	50000-59000			? Overlay VRF Identifier Range (Min:1, Max:16777214)			
* Network VLAN Range	2300-2999			? Per Switch Overlay Network VLAN Range (Min:2, Max:39)			
* VRF VLAN Range	2000-2299			? Per Switch Overlay VRF VLAN Range (Min:2, Max:3967)			
* Subinterface Dot1q Range	2-511			? Per Border Dot1q Range For VRF Lite Connectivity (Min:2)			
* VRF Lite Deployment	Manual			? VRF Lite Inter-Fabric Connection Deploy Options			
VRF Lite Subnet IP Range				? Address range to assign P2P DCI Links			
VRF Lite Subnet Mask				? Mask for Subnet Range			

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

Review the ranges and ensure they are consistent with the existing fabric. The migration will honor the existing resources as found on the fabric. The range settings apply to post migration allocation.

**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.

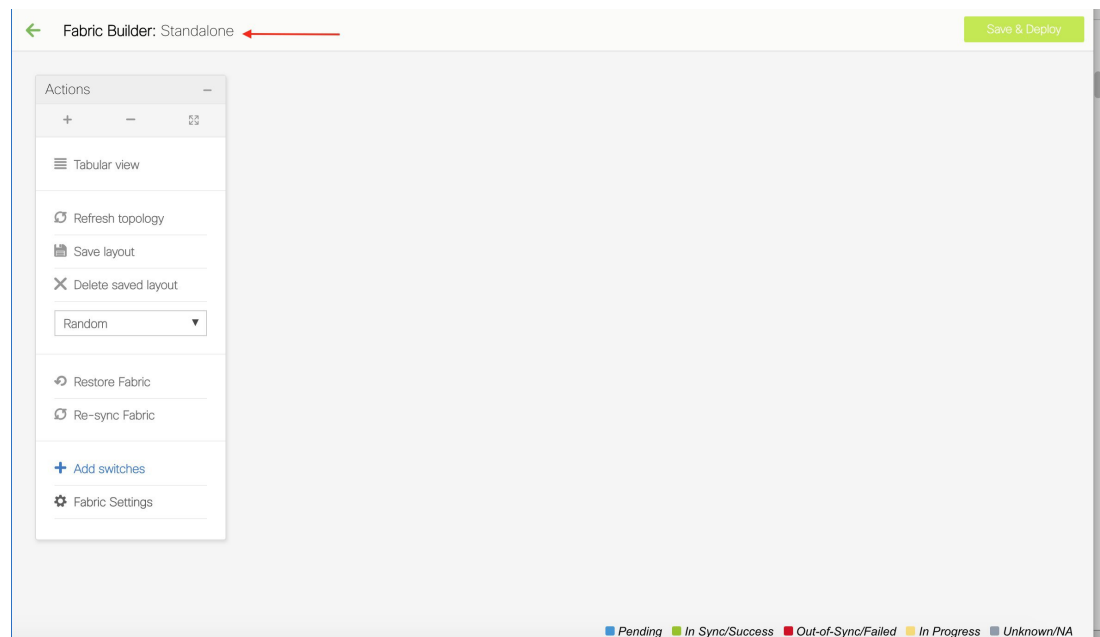
- a. Update the L2 range and click **Save**.
- b. Click the **Edit Fabric** option again, update the L3 range and click **Save**.

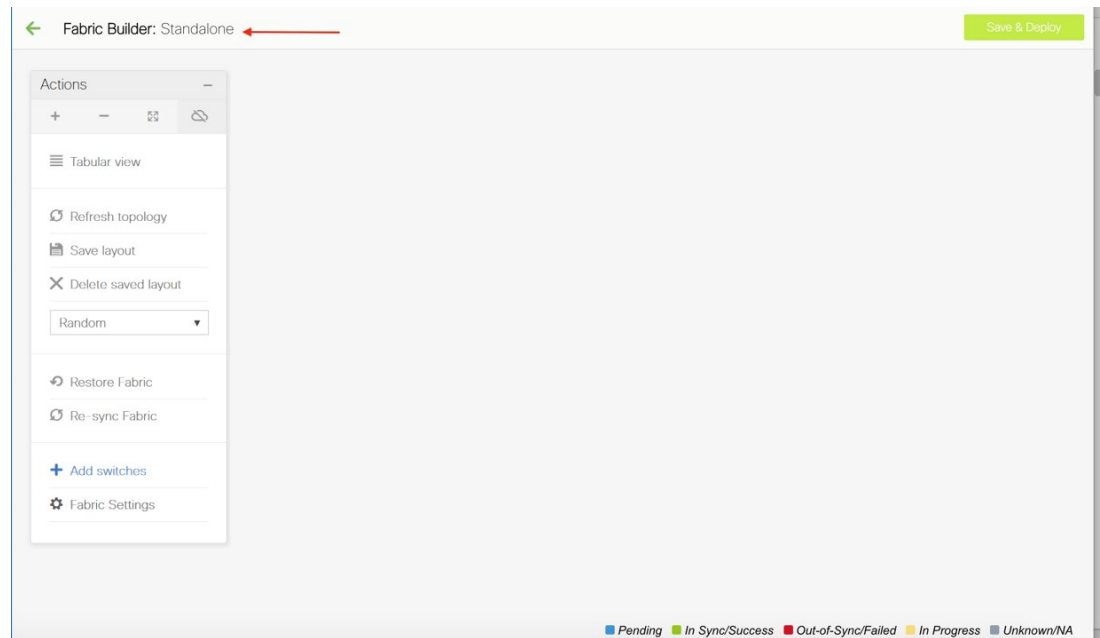
The remaining tabs do not require updates. However, their purpose is mentioned.

7. Click the **Manageability** tab - Leave the fields in this tab blank to retain existing DNS, NTP, AAA, and syslog configurations. Policies are created using the source "".

Post transition, for any new device added to the fabric, you must manually enter the configuration in the switch\_freeform policy configuration. If the tab has any field filled before or after migration, it will overwrite the corresponding feature configuration on the switch.

8. Click the **Bootstrap** tab. Update the fields in this tab post transition, when new switches are added to the fabric.
9. Click the **Configuration Backup** tab. Leave the fields in this tab blank. You can update post transition.
10. 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.





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

### Adding Switch Instances and Transitioning VXLAN Fabric Management to DCNM

1. In the fabric topology screen, click Add switches. The Inventory Management screen comes up. The Discover Existing Switches tab is displayed by default.

## Inventory Management

Discover Existing Switches PowerOn Auto Provisioning (POAP)

Discovery Information > Scan Details >

Seed IP   
*Ex: "2.2.2.20"; "10.10.10.40-60"; "2.2.2.20, 2.2.2.21"*

Authentication Protocol

Username

Password

Max Hops    hop(s)

Preserve Config  no  yes  
*Selecting 'no' will clean up the configuration on switch(es)*

---

The POAP tab is only used for adding new switches to the fabric. Use the tab only after migrating your existing fabric to DCNM.

2. Enter the IP address (Seed IP), administrator username and password (Username and Password fields) of the switch, and set the Max Hops count for the switch. Ensure that all fabric switches can be added to DCNM at once.

**Important** - Ensure that the Preserve Config field remains set to **yes**. Selecting 'no' can cause significant configuration loss and fabric disruption.

## Inventory Management

Discover Existing Switches | PowerOn Auto Provisioning (POAP)

Discovery Information > Scan Details >

Seed IP   
*Ex: "2.2.2.20"; "10.10.10.40-60"; "2.2.2.20, 2.2.2.21"*

Authentication Protocol

Username

Password

Max Hops  hop(s)

Preserve Config  no  yes  
*Selecting 'no' will clean up the configuration on switch(es)*

3. Click Start discovery, at the bottom part of the screen. The switch with the specified IP address and switches up to two hops away (depending on the setting of Max Hops) from it are populated in the Scan Details section.

## Inventory Management



Discover Existing Switches | PowerOn Auto Provisioning (POAP)

Discovery Information > Scan Details >

← Back Note: Preserve Config selection is 'yes'. Import into fabric

<input type="checkbox"/>	Name	IP Address	Model	Version	Status	Progress
<input type="checkbox"/>	80.80.80.1	80.80.80.1	cisco WS-...	12.2(55)SE5,	timeout	
<input type="checkbox"/>	n9k-12	80.80.80.62	N9K-C939...	7.0(3)I7(3)	manageable	
<input type="checkbox"/>	88.88.88.3	88.88.88.3	cisco WS-...	12.2(55)SE5,	not reachable	
<input type="checkbox"/>	n9k-7	80.80.80.57	N9K-C939...	7.0(3)I7(3)	manageable	
<input type="checkbox"/>	n9k-8-spine	80.80.80.58	N9K-C939...	7.0(3)I7(3)	manageable	
<input type="checkbox"/>	n9k-13	80.80.80.63	N9K-C939...	7.0(3)I7(3)	manageable	
<input type="checkbox"/>	host-26-vinci-mgmt...	0.0.0.0	VMware ESX	Releasebuild-799733	not reachable	
<input type="checkbox"/>	n9k-14-spine	80.80.80.64	N9K-C921...	7.0(3)I7(3)	manageable	
<input type="checkbox"/>	n9k-15-spine	80.80.80.65	N9K-C921...	7.0(3)I7(1)	manageable	

Close

- Select the check box next to the concerned switches and click Import into fabric.

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.

## Inventory Management



Discover Existing Switches | PowerOn Auto Provisioning (POAP)

Discovery Information > Scan Details >

← Back Note: Preserve Config selection is 'yes'. Import into fabric

<input type="checkbox"/>	Name	IP Address	Model	Version	Status	Progress
<input type="checkbox"/>	80.80.80.1	80.80.80.1	cisco WS-...	12.2(55)SE5,	timeout	
<input checked="" type="checkbox"/>	n9k-12	80.80.80.62	N9K-C939...	7.0(3)I7(3)	manageable	
<input type="checkbox"/>	88.88.88.3	88.88.88.3	cisco WS-...	12.2(55)SE5,	not reachable	
<input checked="" type="checkbox"/>	n9k-7	80.80.80.57	N9K-C939...	7.0(3)I7(3)	manageable	
<input checked="" type="checkbox"/>	n9k-8-spine	80.80.80.58	N9K-C939...	7.0(3)I7(3)	manageable	
<input checked="" type="checkbox"/>	n9k-13	80.80.80.63	N9K-C939...	7.0(3)I7(3)	manageable	
<input type="checkbox"/>	host-26-vinci-mgmt...	0.0.0.0	VMware ESX	Releasebuild-799733	not reachable	
<input checked="" type="checkbox"/>	n9k-14-spine	80.80.80.64	N9K-C921...	7.0(3)I7(3)	manageable	
<input type="checkbox"/>	n9k-15-spine	80.80.80.65	N9K-C921...	7.0(3)I7(1)	manageable	

Close

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 import switches again) till 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 and initiate the import process again by clicking on Add Switches in the Actions panel.

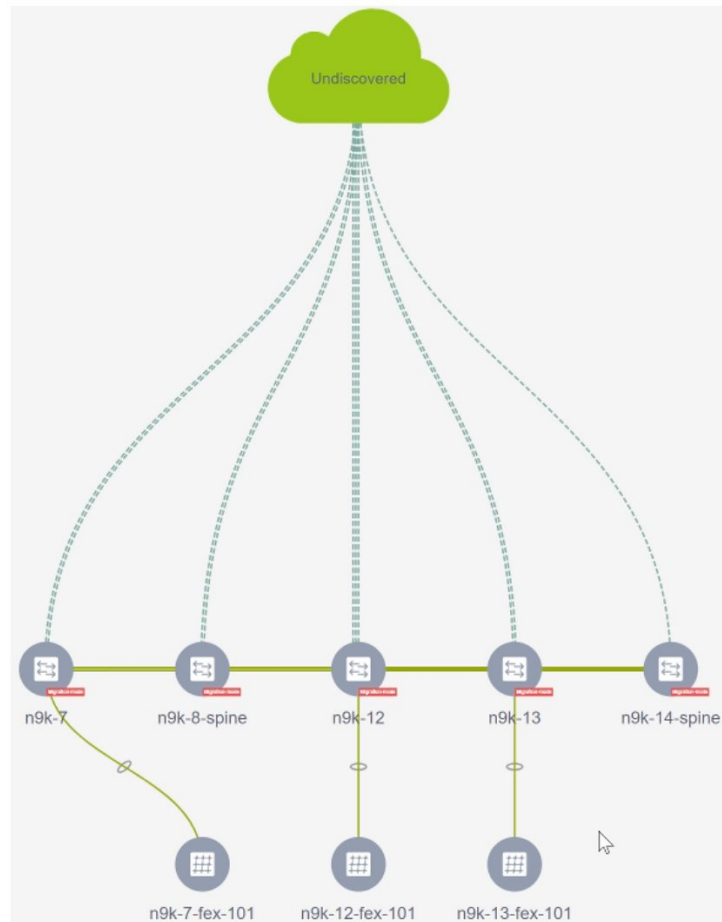
The screenshot shows the 'Inventory Management' window with the 'Discover Existing Switches' tab selected. The interface includes navigation links for 'Discovery Information' and 'Scan Details', a 'Back' button, a note 'Note: Preserve Config selection is 'yes'', and an 'Import into fabric' button. A table lists discovered switches with columns for Name, IP Address, Model, Version, Status, and Progress. The Progress column shows 'done' for several switches.

<input type="checkbox"/>	Name	IP Address	Model	Version	Status	Progress
<input type="checkbox"/>	80.80.80.1	80.80.80.1	cisco WS-...	12.2(55)SE5,	timeout	
<input checked="" type="checkbox"/>	n9k-12	80.80.80.62	N9K-C939...	7.0(3)I7(3)	manageable	done
<input type="checkbox"/>	88.88.88.3	88.88.88.3	cisco WS-...	12.2(55)SE5,	not reachable	
<input checked="" type="checkbox"/>	n9k-7	80.80.80.57	N9K-C939...	7.0(3)I7(3)	manageable	done
<input checked="" type="checkbox"/>	n9k-8-spine	80.80.80.58	N9K-C939...	7.0(3)I7(3)	manageable	done
<input checked="" type="checkbox"/>	n9k-13	80.80.80.63	N9K-C939...	7.0(3)I7(3)	manageable	done
<input type="checkbox"/>	host-26-vinci-mgmt...	0.0.0.0	VMware ESX	Releasebuild-799733	not reachable	
<input checked="" type="checkbox"/>	n9k-14-spine	80.80.80.64	N9K-C921...	7.0(3)I7(3)	manageable	done
<input type="checkbox"/>	n9k-15-spine	80.80.80.65	N9K-C921...	7.0(3)I7(1)	manageable	

Buttons: Close

After DCNM discovers all the switches, and the Progress column displays **done** for all switches, close the screen. The fabric topology screen comes up again. The switch is in Migration Mode now and the Migration mode label is displayed on the switch icons.

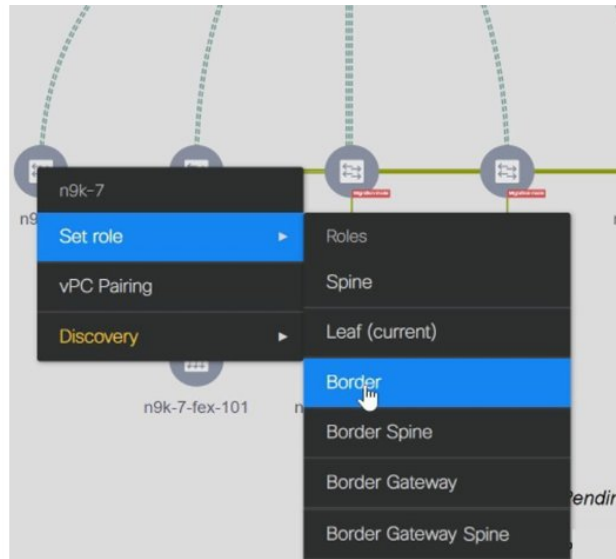
At this point, you must not try to add Greenfield or *new* switches. Support is not available for adding new switches during the migration process. It might lead to undesirable consequences for your network. However, you can add a new switch after the migration process is complete.



**Note** The switch discovery process might fail for a few switches, and the Discovery Error message displayed. However, such switches are still 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.

You must not proceed to the next step till all switches in the existing fabric are discovered in DCNM.

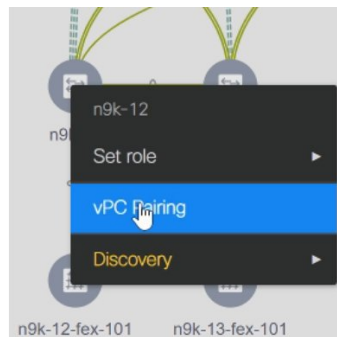
5. Each switch's role and vPC pairing must be set during the fabric migration process.  
Right-click the switch icon and use the Set role option (Leaf, Border, etc) to update switch role.



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.


**vPC Pairing** - The vPC pairing must be done for switches where the Layer 3 vPC peer-keep alive is used. The vPC configuration is automatically picked up from the switches when the vPC peer keep alive is established through the management option. This pairing reflects in the GUI only after the migration is complete.

- a. Right-click the switch icon and click vPC Pairing to set a vPC switch pair.



The Select vPC peer screen comes up. It lists potential vPC peer switches.

- b. Select the appropriate switch and click OK. The fabric topology comes up again. The vPC pair is formed now.







**Note** Check if you have added all switches from the current fabric. If you have missed adding switches, add them now. Once you are certain that you have imported all existing switches, move to the next step, the Save and Deploy option.

- Use the Save and Deploy option (at the top right part of the screen) to sync configurations between the switch and DCNM.

The Saving Fabric Configuration message comes up immediately. This indicates that overlay and underlay network migration, and switch and port channel settings migration to DCNM is initiated.

If there are configuration mismatches, error messages are displayed. Update changes in the fabric settings or the switch configuration as needed, and click Save and Deploy again.

After the migration of underlay and overlay networks, the Configuration Deployment screen comes up.

### Config Deployment ✕

Switch Name	IP Address	Switch Serial	Preview Config	Status	Re-sync	Progress
n9k-13	80.80.80.63	SAL18422FXE	Computing...	Fetching switch con...		30%
n9k-12	80.80.80.62	SAL18422FX8	Computing...	Fetching switch con...		30%
n9k-7	80.80.80.57	SAL1833YM64	Computing...	Fetching switch con...		30%
n9k-14-spine	80.80.80.64	SAL2016NXXB	Computing...	Fetching switch con...		30%
n9k-8-spine	80.80.80.58	SAL1833YM0V	Computing...	STARTED		0%

Deploy Config

The Preview Config column is updated with entries denoting a specific number of lines.

## Config Deployment



Step 1. Configuration Preview &gt; Step 2. Configuration Deployment Status &gt;

Switch Name	IP Address	Switch Serial	Preview Config	Status	Re-sync	Progress
n9k-13	80.80.80.63	SAL18422FXE	2106 lines	Out-of-sync		100%
n9k-12	80.80.80.62	SAL18422FX8	2106 lines	Out-of-sync		100%
n9k-7	80.80.80.57	SAL1833YM64	1939 lines	Out-of-sync		100%
n9k-14-spine	80.80.80.64	SAL2016NXXB	1 lines	Out-of-sync		100%
n9k-8-spine	80.80.80.58	SAL1833YM0V	11 lines	Out-of-sync		100%

We strongly recommend that you preview the configuration before proceeding to deploy it on the switches. Click the Preview Config column entry. The Config Preview screen 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.

## Config Preview - Switch 80.80.80.63

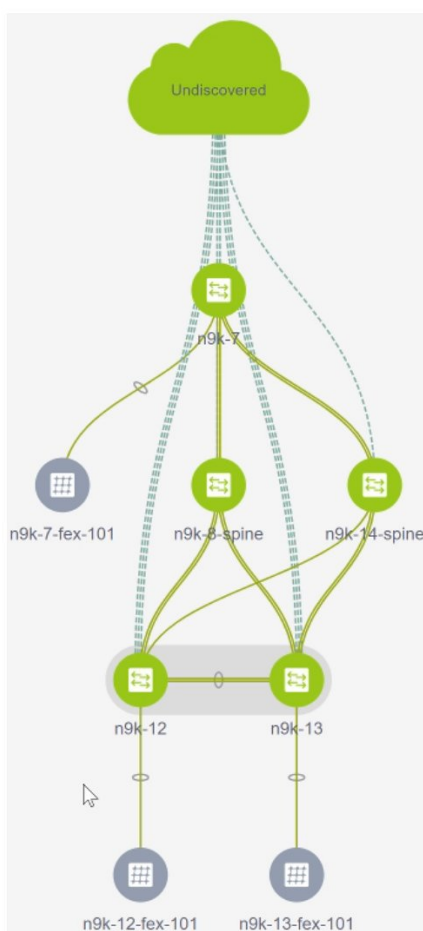


Pending Config	Expected Config	Current Config	Side-by-side Comparison
4 vdc n9k-13 id 1			
5 allow feature-set fex			
6 limit-resource vlan minimum 16 maximum 4094			
7 limit-resource vrf minimum 2 maximum 4096			
8 limit-resource port-channel minimum 0 maximum 256			
9 limit-resource u4route-mem minimum 248 maximum 248			
10 limit-resource u6route-mem minimum 96 maximum 96			
11 limit-resource m4route-mem minimum 58 maximum 58			
12 limit-resource m6route-mem minimum 8 maximum 8			
13 feature-set fex	feature-set fex		
14 feature nxapi	feature nxapi		
15 cfs eth distribute	cfs eth distribute		
16 nv overlay evpn	nv overlay evpn		
17 feature ospf	feature ospf		
18 feature bgp	feature bgp		
19 feature interface-vlan	feature interface-vlan		
20 feature vn-segment-vlan-based	feature vn-segment-vlan-based		
21	feature dhcp		
22 feature lacp	feature lacp		

Close the preview 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.

The progress bar shows 100% for each switch. After correct provisioning and successful configuration compliance, close the screen. In the fabric topology screen that comes up, all imported switch instances are displayed in green color, indicating successful configuration. Also, the Migration Mode label is not displayed on any switch icon.



**Post-transitioning of VXLAN fabric management to DCNM** - This completes the transitioning process of VXLAN fabric management to DCNM. Now, you can add new switches and provision overlay networks for your fabric. For details, refer the respective section in the Fabrics topic in the configuration guide.

#### Fabric Options

- **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 screen. DCNM tries to make a guess and intelligently place nodes that belong together in close proximity.
  - **Custom saved layout** - You can drag nodes around to your liking. Once you have the positions as how you like, you can click Save Layout to remember the positions. Next time you come to the topology, DCNM will draw the nodes based on your last saved layout positions.
- **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 [Restore Fabric, on page 40](#) 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 screen. 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.

## Migrating an MSD Fabric with Border Gateway Switches

When you migrate an existing MSD fabric with a border gateway switch into DCNM, make sure to note the following guidelines:

- **Underlay Multisite peering:** The eBGP peering and corresponding routed interfaces for underlay extensions between sites are captured in **switch\_freeform** and **routed\_interfaces**, and optionally in the **interface\_freeform** configs. This configuration includes all the global configs for multisite. Loopbacks for EVPN multisite are also captured via the appropriate interface templates.
- **Overlay Multisite peering:** The eBGP peering is captured as part of **switch\_freeform** as the only relevant config is under **router bgp**.
- **Overlays containing Networks or VRFs:** The corresponding intent is captured with the profiles on the Border Gateways with **extension\_type = MULTISITE**.

This ensures that the brownfield migration will be complete with no CC diff, and there will be no traffic disruption.

Perform the following steps after you migrate the member fabrics into DCNM:

Before you begin, ensure member fabrics have the correct **Site ID** in the fabric settings.

1. Create an MSD. For more information, see [Creating an MSD Fabric, on page 91](#).

2. Ensure that the fabric settings for MSD are correct including settings such as profile selection, the multisite loopback ID, and anycast GW MAC.
3. Move the member fabrics into the MSD. For more information, see [Moving the Member1 Fabric Under MSD-Parent-Fabric, on page 96](#).



**Note** The networks or VRFs definitions should be symmetric. Otherwise, you will not be able to deploy Multi-Site. If there are any errors based on conflicting definitions for VRFs or networks, you need to resolve before deployment.

4. Create multisite overlay IFC. For more information, see *Configuring Multi-Site Overlay IFCs*.

Multisite overlay IFCs need to be created if **Multi-Site Overlay IFC Deployment Method** is set to **Manual** under the **DCI** tab for the MSD fabric settings.

If **Multi-Site Overlay IFC Deployment Method** is set to **Direct\_To\_BGWS**, then overlay IFCs are created after brownfield migration, and associated with appropriate **MULTISITE\_OVERLAY** policy.

The intent generated by this IFC should match what was captured in the freeform for the **MULTISITE\_IFC** for BGP peering.

Repeat the above step for each BGW **MULTISITE\_OVERLAY** IFC and for each member fabric. After the Multi-Site overlay IFCs are successfully created, the intent for the eBGP multisite overlay peering captured in the freeform policy templates for the BGWs can be removed. Otherwise, the intent for the eBGP multisite overlay peering is captured twice.

Note that there is no need to create **MULTISITE\_UNDERLAY** IFCs as they have already been captured in the intent.

5. To verify, you can select networks or VRFs and corresponding BGWs, and see the expected configurations. You can now manage all the networks or VRFs for BGWs by using the regular top-down workflow.

## Post DCNM 10.4(2) or 11.0(1) to DCNM 11.1(1) Upgrade for VXLAN BGP EVPN and MSD Fabrics

Note the following guidelines after you upgrade DCNM Release 10.4(2) or 11.0(1) to DCNM 11.1(1):

- 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.
- After you upgrade DCNM Release 10.4(2) or 11.0(1) to Release 11.1(1), perform the following steps to use the LAN fabric features of DCNM 11.1(1):
  - Update or save all the Easy Fabrics with the new Easy Fabric Template, that is, **Easy\_Fabric\_11\_1**. Then click **Save & Deploy** to deploy each updated Easy fabric.
  - Update or save all the MSD Fabrics with new MSD Template, that is, **MSD\_Fabric\_11\_1**. Then click **Save & Deploy** to deploy each updated MSD fabric.




---

**Note** Under the **Resources** tab for each Easy Fabric, the Loopback IP Ranges should not be a duplicate of any other Easy Fabric Loopback IP Ranges.

---

After you upgrade DCNM Release 10.4(2) to Release 11.1(1) with custom VRF templates, do the following steps to use MSD feature:

1. For BGP ASN and multicast Group variables, edit the template. Refer [Modifying a Template, on page 165](#).
2. Add an attribute **isFabricInstance=true** in the custom VRF and network templates.

Otherwise while deploying, a network/VRFs created for a member fabric will have bgp ASN and router bgp values to null.

## 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 137](#).

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\_config** 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\_config** 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\_config**.
6. Add or update the CLIs in the **Freeform Config CLI** box.

Copy-paste the intended configuration with correct indentation, as seen in the running configuration on the Nexus switches. For more information, see [Resolving Freeform Config Errors in Switches, on page 137](#).

**7. Click **Save**.**

After the policy is saved, it gets added to the intended configurations for that switch.

**8. Close the policy screens. The Fabric Topology screen comes up again.**

**9. Right click the switch and click **Deploy Config**.**

The **Save & Deploy** option can also be used for deployment. However, the **Save & Deploy** option will identify mismatch between the intended and running configuration *across all* fabric switches.

**Pointers for switch\_freeform\_config Policy Configuration:**

- You can create multiple instances of the policy.
- For a vPC switch pair, create consistent **switch\_freeform\_config** policies on both the vPC switches.
- When you edit a **switch\_freeform\_config** 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\_config** 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 <b>Fabric</b> , <b>Device</b> , <b>DeviceInterface</b> , <b>DevicePair</b> , <b>Fabric</b> , and <b>Link</b> .
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 <b>TOP_DOWN_VRF_LAN</b> , <b>TOP_DOWN_NETWORK_VLAN</b> , <b>LOOPBACK_ID</b> , <b>VPC_ID</b> , and so on.
Is Allocated?	Specifies if the resource is allocated or not. The value is set to <b>True</b> if the resource is permanently allocated to the given entity. The value is set to <b>False</b> 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.
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.  <b>Note</b> 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.	<ul style="list-style-type: none"> <li>• CLI</li> <li>• POAP</li> </ul> <p><b>Note</b> POAP option is not applicable for Cisco DCNM LAN Fabric deployment</p> <ul style="list-style-type: none"> <li>• POLICY</li> <li>• SHOW</li> <li>• PROFILE</li> <li>• FABRIC</li> <li>• ABSTRACT</li> </ul>	Yes

Property Name	Description	Valid Values	Optional?
templateSubType	Specifies the sub type associated with the template.		

Property Name	Description	Valid Values	Optional?
		<ul style="list-style-type: none"> <li>• CLI                             <ul style="list-style-type: none"> <li>• N/A</li> </ul> </li>   <li>• POAP                             <ul style="list-style-type: none"> <li>• N/A</li> <li>• VXLAN</li> <li>• FABRICPATH</li> <li>• VLAN</li> <li>• PMN</li> </ul> </li>   <li><b>Note</b> POAP option is not applicable for Cisco DCNM LAN Fabric deployment.</li>   <li>• POLICY                             <ul style="list-style-type: none"> <li>• VLAN</li> <li>• NIERFACE_VLAN</li> <li>• INTERFACE_VPC</li> <li>• NIERFACE_HRNET</li> <li>• INTERFACE_BD</li> <li>• NIERFACE&gt;NNL</li> <li>• INTERFACE_FC</li> <li>• NIERFACE_MGMT</li> <li>• NIERFACE_OOBAC</li> <li>• INTERFACE_NVE</li> <li>• INTERFACE_VFC</li> <li>• NIERFACE&gt;NNL</li> <li>• DEVICE</li> <li>• FEX</li> <li>• NIRA_FABRIC_LINK</li> <li>• NIER_FABRIC_LINK</li> </ul> </li> </ul>	



Property Name	Description	Valid Values	Optional?
		<ul style="list-style-type: none"> <li>• INTERFACE</li> <li>• SHOW               <ul style="list-style-type: none"> <li>• VLAN</li> <li>• INTERFACE_VLAN</li> <li>• INTERFACE_VPC</li> <li>• INTERFACE_ETH</li> <li>• INTERFACE_BD</li> <li>• <del>INTERFACE_CHANNEL</del></li> <li>• INTERFACE_FC</li> <li>• INTERFACE_MGMT</li> <li>• <del>INTERFACE_COBACK</del></li> <li>• INTERFACE_NVE</li> <li>• INTERFACE_VFC</li> <li>• <del>INTERFACE_CHANNEL</del></li> </ul> </li> <li>• DEVICE               <ul style="list-style-type: none"> <li>• FEX</li> <li>• NIRA_FABRIC_LINK</li> <li>• NIER_FABRIC_LINK</li> </ul> </li> <li>• INTERFACE</li> <li>• PROFILE               <ul style="list-style-type: none"> <li>• VXLAN</li> </ul> </li> <li>• FABRIC               <ul style="list-style-type: none"> <li>• NA</li> </ul> </li> </ul>	

Property Name	Description	Valid Values	Optional?
		<ul style="list-style-type: none"> <li>• ABSTRACT</li> <li>• VLAN</li> <li>• INTERFACE_VLAN</li> <li>• INTERFACE_VPC</li> <li>• INTERFACE_EHRNET</li> <li>• INTERFACE_BD</li> <li>• <del>INTERFACE_CHANNEL</del></li> <li>• INTERFACE_FC</li> <li>• INTERFACE_MGMT</li> <li>• <del>INTERFACE_COBACK</del></li> <li>• INTERFACE_NVE</li> <li>• INTERFACE_VFC</li> <li>• <del>INTERFACE_CHANNEL</del></li> <li>• DEVICE</li> <li>• FEX</li> <li>• <del>NIRA_FABRIC_LINK</del></li> <li>• <del>NIER_FABRIC_LINK</del></li> <li>• INTERFACE</li> </ul>	

Property Name	Description	Valid Values	Optional?
contentType		<ul style="list-style-type: none"> <li>• CLI <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> </ul> </li> <li>• POAP <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> </ul> </li> </ul> <p><b>Note</b> POAP option is not applicable for Cisco DCNM LAN Fabric deployment.</p> <ul style="list-style-type: none"> <li>• POLICY <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> <li>• PYTHON</li> </ul> </li> <li>• SHOW <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> </ul> </li> <li>• PROFILE <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> <li>• PYTHON</li> </ul> </li> <li>• FABRIC <ul style="list-style-type: none"> <li>• PYTHON</li> </ul> </li> <li>• ABSTRACT <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> <li>• PYTHON</li> </ul> </li> </ul>	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	<p>You can have a list of IPv4, IPv6, or a combination of both types of addresses.</p> <p>Example 1: 172.22.31.97, 172.22.31.99, 172.22.31.105, 172.22.31.109</p> <p>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</p> <p>Example 3: 172.22.31.97, 172.22.31.99,  2001:0db8:85a3:0000:0000:8a2e:0370:7334,  172.22.31.254</p>	Yes
ipAddressWithoutPrefix	<p>Example: 192.168.1.1</p> <p>or</p> <p>Example: 1:2:3:4:5:6:7:8</p>	No
ipV4Address	IPv4 address	No
ipV4AddressWithSubnet	Example: 192.168.1.1/24	No
ipV6Address	IPv6 address	No
ipV6AddressWithPrefix	<p>Example: 1:2:3:4:5:6:7:8</p> <p>22</p>	No
ipV6AddressWithSubnet	IPv6 Address with Subnet	No
ISISNetAddress	<p>Example: 49.0001.00a0.c96b.c490.00</p>	No
long	Example: 100	No
macAddress	14 or 17 character length MAC address format	No
string	<p>Free text, for example, used for the description of a variable</p> <p>Example: string scheduledTime {  regularExpr="^([01]\d 2[0-3]):([0-5]\d)\$"; }</p>	No

Variable Type	Valid Value	Iterative?
string[]	Example: {a,b,c,str1,str2}	Yes
struct	<p>Set of parameters that are bundled under a single variable.</p> <pre> struct &lt;structure name declaration &gt; { &lt;parameter type&gt; &lt;parameter 1&gt;; &lt;parameter type&gt; &lt;parameter 2&gt;; ... } [&lt;structure_inst1&gt;] [, &lt;structure_inst2&gt;] [, &lt;structure_array_inst3 []&gt;;  struct interface_detail { string inf_name; string inf_description; ipAddress inf_host; enum duplex { validValues = auto, full, half; }; }myInterface, myInterfaceArray[]; </pre>	<p>No</p> <p><b>Note</b> If the struct variable is declared as an array, the variable is iterative.</p>
wwn (Available only in Cisco DCNM Web Client)	<p>Example:</p> <p>20:01:00:08:02:11:05:03</p>	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							
float Range	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							
integer Range	Range of signed numbers Example: 50-65	Yes	Yes		Yes	Yes							
interface	specifies interface Example: Ethernet 5/10	Yes	Yes				Yes	Yes	Yes	Yes			
interface Range		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
ipAddressList	<p>You can have a list of IPv4, IPv6, or a combination of both types of addresses.</p> <p>Example 1:                      122.3.9,                      122.3.9,                      1223.15,                      1223.10</p> <p>Example 2:                      10.1.1.1,                      10.1.1.2,                      2008:1:1:1</p> <p>Example 3:                      122.3.9,                      122.3.9,                      10.1.1.1,                      1223.15</p> <p><b>Note</b> Separate the addresses in the list using commas and not hyphens.</p>	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
<del>ipAddr</del>	IPv4 or IPv6 Address (does not require prefix)												
<del>ip4Addr</del>	IPv4 address	Yes											
<del>ip4Subnet</del>	IPv4 Address with Subnet	Yes											
<del>ip6Addr</del>	IPv6 address	Yes											
<del>ip6Prefix</del>	IPv6 Address with prefix	Yes											
<del>ip6Subnet</del>	IPv6 Address with Subnet	Yes											
<del>ip4Subnet</del>	Example: 192.168.1.0/24												
long	Example: 100	Yes			Yes	Yes							
<del>macAddr</del>	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  statement { ... }	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	<p>Set of <del>params</del> that are bundled under a single variable.</p> <pre> struct &lt;structure name declaration&gt; {   &lt;parameter type&gt;   &lt;parameter 1&gt;;   &lt;parameter type&gt;   &lt;parameter 2&gt;;   ... } [&lt;struct1&gt; [, &lt;struct2&gt; [, &lt;struct3&gt; [ ]&gt;]; </pre>												
wwn	WWN address												

### Example: Meta Property Usage

```

##template variables

integer VLAN_ID {
min = 100;
max= 200;
};

string USER_NAME {
defaultValue = admin123;
minLength = 5;
};

struct interface_a{

```

```

string inf_name;
string inf_description;
ipAddress inf_host;
enum duplex {
    validValues = auto, full, half;
};
}myInterface;

##

```

## Variable Annotation

You can configure the variable properties marking the variables using annotations.



**Note** Variable Annotations are available for POAP only. However, the annotations do not impact on the template type 'CLI'.

The following annotations can be used in the template variable section.

Annotation Key	Valid Values	Description
DataDepend	Text	
Description	Text	Description of the field appearing in the window
DisplayName	Text <b>Note</b> Enclose the text with quotes, if there is space.	Display name of the field appearing in the window
Enum	Text1, Text2, Text3, and so on	Lists the text or numeric values to select from
IsAlphaNumeric	"true" or "false"	Validates if the string is alphanumeric
IsAsn	"true" or "false"	
IsDestinationDevice	"true" or "false"	
IsDestinationFabric	"true" or "false"	
IsDestinationInterface	"true" or "false"	
IsDestinationSwitchName	"true" or "false"	
IsDeviceID	"true" or "false"	
IsDot1qId	"true" or "false"	
IsFEXID	"true" or "false"	

Annotation Key	Valid Values	Description
IsGateway	“true” or “false”	Validates if the IP address is a gateway
IsInternal	“true” or “false”	Makes the fields internal and does not display them on the window  <b>Note</b> Use this annotation only for the ipAddress variable.
IsManagementIP	“true” or “false”  <b>Note</b> 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”	
IsSourceSwitchName	“true” or “false”	

Annotation Key	Valid Values	Description
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

### Example: DisplayName Annotation

```
##template variables
@(DisplayName="Host Name", Description = "Description of the host")
String hostname;
@(DisplayName="Host Address", Description = " test description" IsManagementIP=true)
IPAddress hostAddress;
##
```

**Example: IsMandatory Annotation**

```
##template variables
@(IsMandatory="ipv6!=null")
ipV4Address ipv4;
@(IsMandatory="ipv4!=null")
ipV6Address ipv6;
##
```

**Example: IsMultiLineString Annotation**

```
##template variables
@(IsMultiLineString=true)
string EXTRA_CONF_SPINE;
##
```

**IsShow Annotation**

```
##template variables
boolean isVlan;
@(IsShow="isVlan==true")
integer vlanNo;
##
```

```
##template variables
boolean enableScheduledBackup;
@(IsShow="enableScheduledBackup==true",Description="Server time")
string scheduledTime;
##
The condition "enableScheduledBackup==true" evaluates to true/false
```

```
##template variables
@(Enum="Manual,Back2BackOnly,ToExternalOnly,Both")
string VRF_LITE_AUTOCONFIG;
@(IsShow="VRF_LITE_AUTOCONFIG!=Manual", Description="Target Mask")
integer DCI_SUBNET_TARGET_MASK
##
The condition "VRF_LITE_AUTOCONFIG!=Manual" matches string comparison to evaluate to true
or false
```

## Templates Content

This section includes the configuration commands and any parameters that you want to include in the template. These commands can include the variables declared in the template variables section. During the command generation process the variable values are substituted appropriately in the template content.



**Note** You must specify the commands that you include as if you were entering them in the global configuration command mode on any device. You must consider the command mode when you include commands.

Template content is governed by the usage of variables.

- **Scalar variables:** does not take a range or array of values which cannot be used for iteration (In the variable types table those marked iterate-able as 'No'). Scalar variables must be defined inside the template content.

```
Syntax: $$<variable name>$$
Example: $$USER_NAME$$
```

- **Iterative variables:** used for block iteration. These loop variable must be accessed as shown below inside the iteration block.

```
Syntax: @<loop variable>
Example:
foreach val in $$INTEGER_RANGE_VALUE$$ {
@val
}
```

- **Scalar Structure Variable:** Structure member variables can be accessed inside the template content.

```
Syntax: $$<structure instance name>.<member variable name>$$
Example: $$myInterface.inf_name$$
```

- **Array Structure Variable:** Structure member variables can be accessed inside the template content.

```
Syntax: $$<structure instance name>.<member variable name>$$
Example: $$myInterface.inf_name$$
```

In addition to the template variables, you can use the conditional and iterative command generation using the following statements:

- **if-else if-else Statement:** makes a logical decision in inclusion/exclusion of set of configuration command based on the value assigned for the variable in it.

```
Syntax: if(<operand 1> <logical operator> <operand 2>){
command1 ..
command2..
..
}
else if (<operand 3> <logical operator> <operand 4> )
{
Command3 ..
Command4..
..
}
else
{
Command5 ..
Command6..
..
}
Example: if-else if-else statement
if($$USER_NAME$$ == 'admin'){
Interface2/10
no shut
}
else {
Interface2/10
shut
}
```

- **foreach Statement:** used for iterating a block of commands. The iteration is performed based on the assigned loop variable value.

```
Syntax:
foreach <loop index variable> in $$<loop variable>$$ {
@<loop index variable> ..
}
```



```
Example: foreach Statement
foreach ports in $$MY_INF_RANGE$$ {
interface @ports
no shut
}
```

- **Optional parameters:** By default all parameters are mandatory. To make a parameter optional, you must annotate the parameter.

In the variable section, you can include the following command:

- **@(IsMandatory=false)**
- **Integer frequency;**

In the template content section, a command can be excluded or included without using “if” condition check, by assigning a value to the parameter. The optional command can be framed as below:

- **probe icmp [frequency frequency-value] [timeout seconds] [retry-count retry-count-value]**

## 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 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 165](#).
- 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.
-



## 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.  <b>Note</b> If Failover is triggered in Native HA, the Task Id sequence number is incremented by 32.
Task Type	Specifies the type of task. <ul style="list-style-type: none"> <li>• Compatibility</li> <li>• Upgrade</li> </ul>
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. <ul style="list-style-type: none"> <li>• Planned</li> <li>• In Progress</li> <li>• Completed</li> <li>• Completed with Exceptions</li> </ul>
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:

## 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.

- 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. DCCM discovers polling cycles in order to display the new version of the switch on the Cisco DCCM 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.
-

## 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.

---

## 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. <ul style="list-style-type: none"> <li>• Planned</li> <li>• In Progress</li> <li>• Completed</li> </ul>
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 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
- Support real-time and historical search for all endpoints across various search filters such as VRF, Network, Layer-2 VNI, Layer-3 VNI, Switch, IP, MAC, port, VLAN, and so on.
- Support for real-time and historical dashboards for insights such as endpoint lifetime, network, endpoint, VRF daily views, and operational heat map.
- Support for high availability
- Support for endpoint data that is stored for up to 180 days, amounting to a maximum of 5 G storage space.
- Support for optional flush of the endpoint data in order to start afresh.
- Supported scale: 10K endpoints

For more information about EPL, refer to the following sections:



## Configuring Endpoint Locator

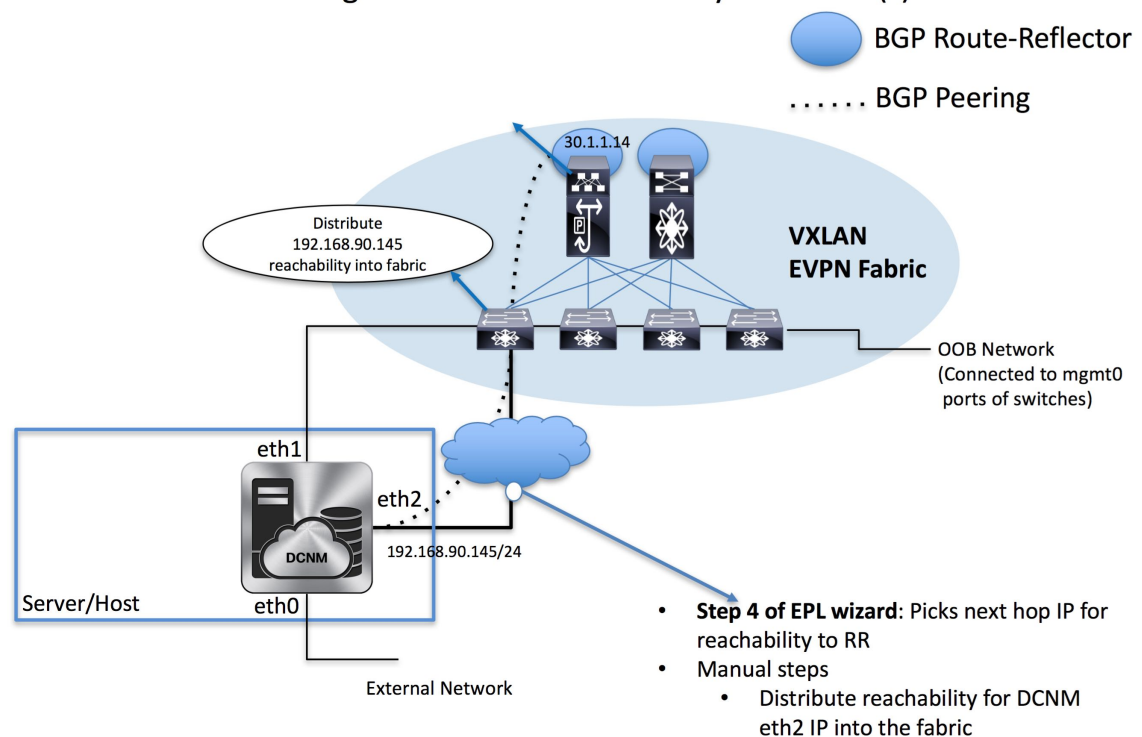
The DCNM OVA or the ISO installation comes with 3 interfaces—eth0 interface for external access to the DCNM, eth1 interface that is used primarily for fabric management, and eth2 interface for in-band network connectivity to Cisco DCNM. In most deployments the eth1 interface is part of the same network on which the mgmt0 interfaces of the Cisco Nexus switches reside (Out-of-band or OOB network). This allows DCNM to perform out-of-band management of these devices including POAP.

BGP peering between the Cisco DCNM and the Route-Reflector is required for EPL. Since the BGP process on Nexus devices typically runs on the non-management VRF, specifically default VRF, it requires an in-band IP connectivity from the Cisco DCNM to the fabric. For this purpose, the eth2 interface can be configured using the **appmgr setup inband** command. The user will be prompted to specify an IP address, netmask and gateway IP. On the fabric side if the DCNM eth2 port is directly connected to one of the front-end interfaces on a switch then the front-end interface can be configured using the *epl\_routed\_intf* template.

After the in-band connectivity is established between the physical or virtual DCNM and the fabric, BGP peering can be established. There is a simple wizard for enabling Endpoint Locator.

## Configuration

The Server Hosting DCNM has IP connectivity to BGP RR(s)



During the EPL configuration using the wizard, the route reflectors (RRs) are configured to accept DCNM as a BGP peer. During the same configuration, the DCNM is also configured by adding routes to the BGP loopback IP via the eth2 gateway. The DCNM can be directly attached to a ToR, or leaf, that in turn provides reachability to the RR. Also, DCNM can have simple IP connectivity via a gateway to the fabric in any case the gateway of eth2 should be appropriately configured when setting up the eth2 port on DCNM.



**Note** Cisco DCNM queries the BGP RR to glean information for establishment of the peering, like ASN, RR, IP, and so on.

It should be noted that it is very important to configure eth2 interface properly, if it is a native HA setup then eth2 on active and standby Cisco DCNMs must be in the same subnet, which means they should have the same gateway addresses.

The screenshot shows the Cisco DCNM interface configuration for leaf1 Ethernet1/24. The 'Edit Configuration' dialog is open, showing the following settings:

- Name: leaf1 Ethernet1/24
- Policy: trunk\_host
- int\_subif:  Enable spanning-tree fpp/guard
- nfm\_monitored:  Enable spanning-tree edge port behavior
- nfm\_trunk\_host: jumbo
- MTU for the interface: none
- Trunk Allowed Vians:  Admin state of the interface

The background table shows the following interface configurations:

Device Name	Name	Admin	Oper	Reason	Policy	Overlay Network	Status	PC	vPC ID	Speed	MTU	Mode	VL
leaf1	Ethernet1/29	↑	↓	XCVR not inserted	trunk_host	NA	✓			10Gb		trunk	
leaf1	Ethernet1/30	↑	↓	XCVR not inserted	trunk_host	NA	✓			10Gb		trunk	
leaf1	Ethernet1/31	↑	↓	XCVR not inserted	trunk_host	NA	✓			10Gb		trunk	
leaf1	Ethernet1/32	↑	↓	XCVR not inserted	trunk_host	NA	✓			10Gb		trunk	
leaf1	Ethernet1/33	↑	↑	ok	int_fabric_p2p	NA	✓			10Gb		routed	
leaf1	Ethernet1/34	↑	↓	XCVR not inserted	trunk_host	NA	✓			10Gb		trunk	

The screenshot shows the Cisco DCNM interface configuration for leaf1 Ethernet1/24. The 'Edit Configuration' dialog is open, showing the following settings:

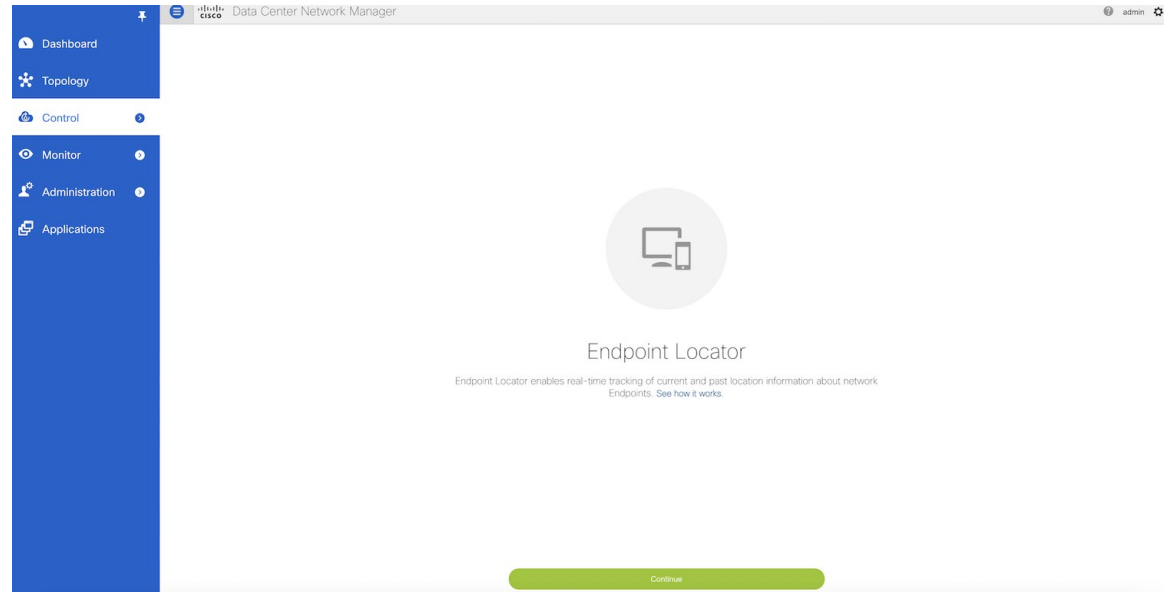
- Name: leaf1 Ethernet1/24
- Policy: epl\_routed\_intf
- General:
  - Interface IP: 192.168.94.1 (IP address of the interface)
  - IP Netmask Length: 24 (IP netmask length used with the IP address)
  - LS Routing Protocol: ospf (Select link-state routing protocol)
  - Link-State Routing Tag: UNDERLAY (Link-state routing protocol tag)
  - Interface Admin State:  Admin state of the interface

The background table shows the following interface configurations:

Device Name	Name	Admin	Oper	Reason	Policy	Overlay Network	Status	PC	vPC ID	Speed	MTU	Mode	VL
leaf1	Ethernet1/29	↑	↓	XCVR not inserted	trunk_host	NA	✓			10Gb		trunk	
leaf1	Ethernet1/30	↑	↓	XCVR not inserted	trunk_host	NA	✓			10Gb		trunk	
leaf1	Ethernet1/31	↑	↓	XCVR not inserted	trunk_host	NA	✓			10Gb		trunk	
leaf1	Ethernet1/32	↑	↓	XCVR not inserted	trunk_host	NA	✓			10Gb		trunk	
leaf1	Ethernet1/33	↑	↑	ok	int_fabric_p2p	NA	✓			10Gb		routed	
leaf1	Ethernet1/34	↑	↓	XCVR not inserted	trunk_host	NA	✓			10Gb		trunk	

## Procedure

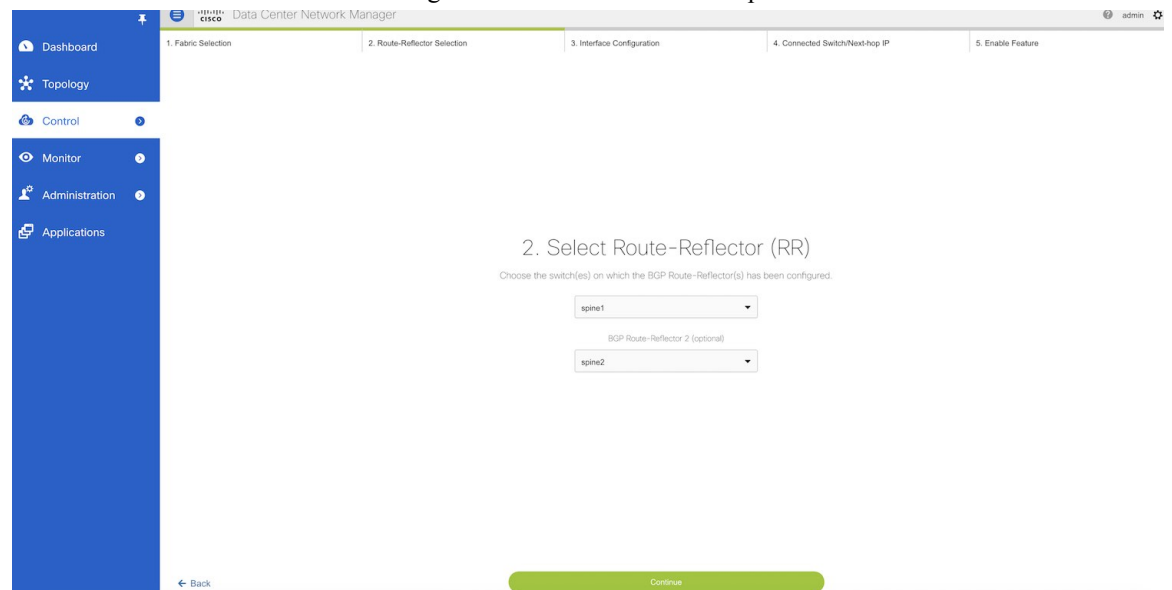
**Step 1** From the Cisco DCNM home page, choose **Control > Endpoint Locator > Configure**. The **Endpoint Locator** page appears with a **See how it works** help link.



**Step 2** Click **Continue**.

**Step 3** Select the appropriate fabric on which the endpoint locator feature should be enabled to track endpoint activity. EPL can only be enabled for one fabric. It can be DFA or EVPN.

**Step 4** Select the switches on the fabric hosting the RRs. Cisco DCNM will peer with the RRs.



**Step 5** Check DCNM eth2 configuration for IP reachability to the RR.

3. Verify DCNM In-band Interface

Choose the Ethernet interface on the DCNM that will provide reachability to the BGP Route-Reflector(s) within the fabric.

eth2

Interface IP

192.168.94.124 / 24

← Back Continue

**Step 6** Check Next-hop IP, and ensure the gateway IP is correct. If there is an error go to command line and reconfigure the eth2 port using the **appmgr setup inband** command.

4. Next-hop IP

Provide the next-hop IP that provides reachability to the BGP Route-Reflector (RR)

Configure my fabric

Next-hop IP

192.168.94.1

← Back Continue

**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.

The screenshot shows the Cisco Data Center Network Manager interface. The left sidebar contains navigation options: Dashboard, Topology, Control (selected), Monitor, Administration, and Applications. The main content area displays a progress bar with five steps: 1. Fabric Selection, 2. Route-Reflector Selection, 3. Interface Configuration, 4. Connected Switch/Next-hop IP, and 5. Enable Feature. The current step is '5. Review and Enable Endpoint Locator'. The configuration form includes the following fields:

Field	Value
Fabric	epi-test
DCNM Interface	eth2 (192.168.94.124/24)
Fabric configuration	Configure my fabric
Route-Reflector 1	spine1 (24.0.80.204)
Next-hop IP	192.168.94.1
* Collect additional information (Port, VLAN, etc.)	
Route-Reflector 2	spine2 (24.0.80.201)

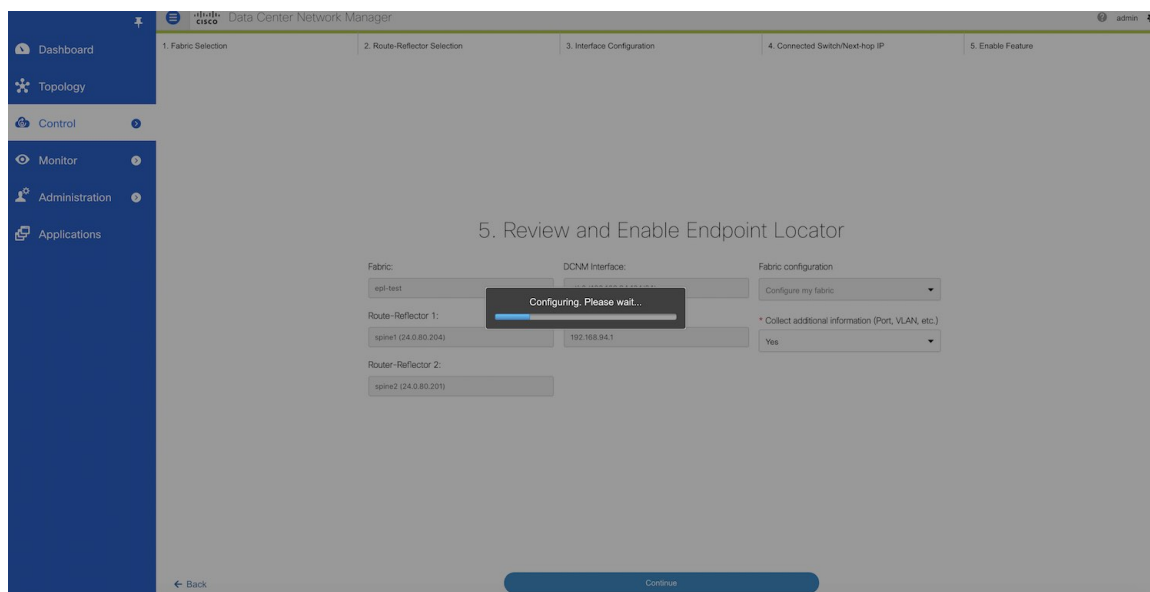
At the bottom of the form, there are 'Back' and 'Continue' buttons.

However, if the **Yes** option is selected in the drop down, a warning pop-up appears that feature NX-API must be supported and enabled on the switches/ToRs/leafs to gather this information. Otherwise this additional information cannot be fetched or reported.


This screenshot shows the same configuration step as the previous image, but with a warning pop-up displayed. The pop-up message reads: "This option requires NX-API feature to be enabled on the switches. Please ensure this step is done for the Endpoint Locator feature to fetch additional information. Are you sure you want to continue?" with 'Yes' and 'No' buttons. In the background, the 'Fabric configuration' dropdown menu is now set to 'Yes'.

### 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 during the enablement, the enable process will abort and the appropriate error message will be displayed. Otherwise, EPL will be successfully enabled and on clicking **OK**, the screen will be automatically redirected to the EPL dashboard.



When the Endpoint Locator feature is enabled, there are a number of steps that occur in the background. DCNM will contact the selected RRs and determine the ASN, determine whether the fabric is L3VPN or EVPN enabled, and also determine the interface IP that is bound to the BGP process. Also, appropriate BGP neighbor statements are added on the RR(s), to get it ready to accept the BGP connection that will be initiated from the DCNM. The neighbor address will be the same as that of the eth2 interface shown in step 2. In order to provide reachability to the RR, a static route will be added to DCNM. This ensures that DCNM has connectivity to the RR. Once EPL is successfully enabled, the user is automatically redirected to the EPL

dashboard that depicts operational and exploratory insights into the endpoints that are present in the fabric. For more information, refer to *Section Exploring Endpoint Locator Details*.

## Flushing the Endpoint Database


To flush the all the Endpoint information, perform the following steps:

### Procedure

**Step 1** From Cisco DCNM home page, choose **Control > Endpoint Locator > Configure**, and then click the **clean up** link.

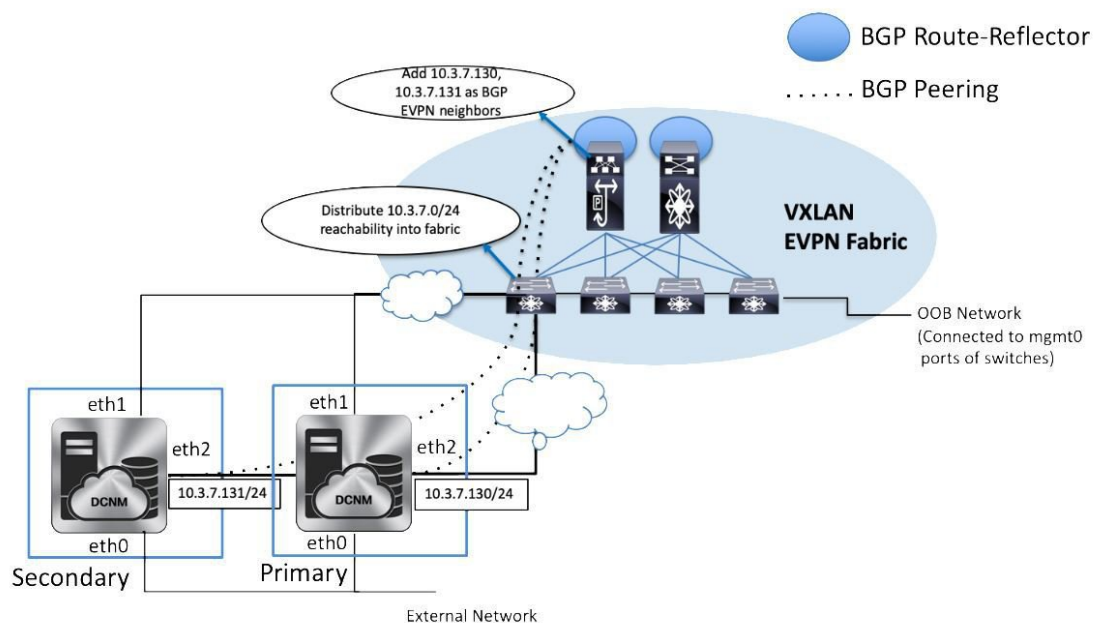
The screenshot shows the Cisco Data Center Network Manager (DCNM) interface. The left sidebar contains navigation options: Dashboard, Topology, Control, Monitor, Administration, and Applications. The main content area is titled 'Endpoint Locator' and contains configuration fields for Fabric (epi-test), DCNM Interface (eth2 (192.168.94.124/24)), Fabric configuration (Configure my fabric), Route-Reflector 1 (spine1 (24.0.80.204)), Next-hop IP (192.168.94.1), and Route-Reflector 2 (spine2 (24.0.80.201)). A 'clean up' link is visible in the bottom right corner of the configuration area.

This shows a warning message indicating that all the endpoint information from the database will be flushed.



**Step 2** Click **Delete** to continue or **Cancel** in case the user wants to abort.

## Configuring Endpoint Locator in DCNM High Availability Mode



The following example shows a sample output for the `apmgrp setup inband` command for a Cisco DCNM Native HA Appliance.

On Cisco DCNM Primary appliance:



```
[root@dcnm-primary]# appmgr setup inband
Configuring Interface for InBand Connectivity...
Please enter the information as prompted:
InBand Physical IP [e.g. 2.2.2.69]: 10.3.7.130
InBand Network Mask [e.g. 255.255.255.0]: 255.255.255.0
InBand Gateway [e.g. 2.2.2.1]: 10.3.7.1
InBand Virtual IP for HA setup [e.g. 2.2.2.60]: 10.3.7.254
InBand Virtual Network Mask [mandatory for HA setup] [e.g. 255.255.255.0]: 255.255.255.0
Peer Inband IP [mandatory for HA setup] [e.g. 2.2.2.59]: 10.3.7.131
Validating Inputs ...

You have entered these values..
PIP=10.3.7.130
NETMASK=255.255.255.0
GATEWAY=10.3.7.1
VIP=10.3.7.254
VIP_NETMASK=255.255.255.0
PEER_ETH2=10.3.7.131

Press 'y' to continue configuration, 'n' to discontinue [y] y

Done.
[root@dcnm-primary]#
```

On Cisco DCNM Secondary appliance:

```
[root@dcnm-secondary]# appmgr setup inband
Configuring Interface for InBand Connectivity...
Please enter the information as prompted:
InBand Physical IP [e.g. 2.2.2.69]: 10.3.7.131
InBand Network Mask [e.g. 255.255.255.0]: 255.255.255.0
InBand Gateway [e.g. 2.2.2.1]: 10.3.7.1
InBand Virtual IP for HA setup [e.g. 2.2.2.60]: 10.3.7.254
InBand Virtual Network Mask [mandatory for HA setup] [e.g. 255.255.255.0]: 255.255.255.0
Peer Inband IP [mandatory for HA setup] [e.g. 2.2.2.59]: 10.3.7.130
Validating Inputs ...

You have entered these values..
PIP=10.3.7.131
NETMASK=255.255.255.0
GATEWAY=10.3.7.1
VIP=10.3.7.254
VIP_NETMASK=255.255.255.0
PEER_ETH2=10.3.7.130

Press 'y' to continue configuration, 'n' to discontinue [y] y
HA Role is Active {"ResponseType":0,"Response":"Refreshed"}
Done.

[root@dcnm-secondary]#
```

After the in-band connectivity is established from both the Primary and Secondary nodes to the Fabric, to configure endpoint locator in DCNM HA mode from the Cisco DCNM Web UI, perform the following steps:

## Procedure

**Step 1** Choose **Control > Endpoint Locator > Configure**.

The **Endpoint Locator** window appears and the fabric configuration details are displayed.

- Step 2** In the Select a fabric to configure endpoint locator in DCNM HA mode.
- Step 3** Click **Continue**.
- Step 4** Select one or two Route-Reflectors (RRs).
- Step 5** Click **Continue**.
- Step 6** Verify the Ethernet interfaces on both primary and standby DCNM nodes.
- Step 7** Click **Continue**.
- Step 8** Verify the next-hop IP address on the primary and standby DCNM.  
Note that the next-hop IP corresponds to the eth2 gateway which should be the same on both the DCNMs.
- Step 9** Click **Continue**.
- Step 10** After selecting the NX-API enable or disable option and verifying the other information provided in the prior steps, click **Continue**.

---

### What to do next

After you configure the Endpoint Locator in HA mode, you can view details such as Endpoint Activity and Endpoint History in the Endpoint dashboard. To view these details, navigate to **Monitor > Endpoint Locator > Explore**.

## Adding High Availability Node to Endpoint Locator Configuration

A standalone DCNM setup can be converted into a native HA deployment at a subsequent time. If EPL is enabled on the standalone DCNM, you can enable EPL for Cisco DCNM Native HA deployment. To add a HA node to Endpoint Locator from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Endpoint Locator > Configure**.  
The **Endpoint Locator** page appears and the fabric configuration details are displayed.
- Step 2** Click the **Add HA node** link.
- Step 3** In the **Configure Standby DCNM Interface** page, choose the Ethernet interface on DCNM that provides reachability to the BGP Route-Reflectors within the fabric.
- Step 4** Click **Continue**.
- Step 5** In the Next-Hop page check the value of the next-hop IP.
- Step 6** Click **Configure HA Node**.  
The configuration details are displayed on the Endpoint Locator page.
- 

## Configuring Endpoint Locator for External Fabrics

In addition to Easy fabrics, DCNM Release 11.2(1) allows you to enable EPL for VXLAN EVPN fabrics comprising of switches that are imported into the external fabric. The external fabric can be in managed mode or monitored mode, based on the selection of **Fabric 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 neighborhood is added to the spines or RRs via some other means.

## Disabling Endpoint Locator

To disable endpoint locator from the Cisco DCNM Web UI, perform the following steps:

### Procedure

- 
- Step 1** Choose **Control > Endpoint Locator > Configure**.
- The **Endpoint Locator** window appears and the fabric configuration details are displayed.
- Step 2** Click **Disable Feature**.
- 

## Troubleshooting Endpoint Locator

There may be multiple reasons why enabling the Endpoint Locator feature may fail. Typically, if the appropriate devices are selected and the IP addresses to be used are correctly specified, the connectivity of the DCNM to the BGP RR may not be present due to which the feature cannot be enabled. This is a sanity check that is present to ensure that basic IP connectivity is available. The following image shows an example error scenario that was encountered during an attempt to enable the EPL feature.

The logs for EPL are located at the following location: `/usr/local/cisco/dcm/fm/logs`. The log that provides further details on what all occurred when the EPL feature is enabled or disabled, are present in the file `epl.log`. The following example provides a snapshot of the log that provides the user further information on when EPL enablement failed.

The following example helps you to understand why Endpoint Locator is unable to Connect to a switch.

```
#tail -f epl.log
2017.04.08 07:47:05 INFO [epl] Running script: [/sbin/appmgr, status, epls]
2017.04.08 07:47:05 INFO [epl] Received response:
2017.04.08 07:47:05 INFO [epl]
2017.04.08 07:47:05 INFO [epl] >>> Sat Apr 8 07:47:05 PDT 2017
2017.04.08 07:47:05 INFO [epl] appmgr status epls
2017.04.08 07:47:05 INFO [epl]
2017.04.08 07:47:05 INFO [epl] EPLS is stopped...
2017.04.08 07:47:08 INFO [epl] Running command: ifdown eth2
2017.04.08 07:47:08 INFO [epl] Received response:
2017.04.08 07:47:08 INFO [epl] EPL disabled successfully
2017.04.08 08:00:06 INFO [epl] Enable End Point Locator
2017.04.08 08:00:13 ERROR [epl] Failed to connect to switch 192.169.6.2:java.lang.Exception:
  Authentication failed : Ssh/Telnet failed to connect with the switch
2017.04.08 08:00:13 INFO [epl] Failed to Enabled End Point Locator. Trying to removing
configuration
2017.04.08 08:00:13 INFO [epl] Disable EPL
2017.04.08 08:00:19 ERROR [epl] Failed to connect to switch 192.169.6.2:java.lang.Exception:
  Authentication failed : Ssh/Telnet failed to connect with the switch
2017.04.08 08:00:19 ERROR [epl] Failed to connect to switch: 192.169.6.2
2017.04.08 08:00:21 ERROR [epl] Failed to unconfigure BGP neighbor or failed to connect to
switch or fabric information not provided.
2017.04.08 08:00:21 ERROR [epl] Failed to unconfigure BGP neighbor or failed to connect to
switch or fabric information not provided.
2017.04.08 08:00:21 INFO [epl] Received response: configure terminal
```

```

Interface Ethernet1/1
no ip address
switchport
end
Enter configuration commands, one per line. End with CNTL/Z.
(config)# Interface Ethernet1/1
(config-if)# no ip address
(config-if)# switchport
(config-if)# end
# from 192.169.6.45
2017.04.08 08:00:21 ERROR [epl] Failed to disable EndPoint locator:
java.lang.NullPointerException
2017.04.08 08:00:21 INFO [epl] EPL disabled succesfully

```

In this example, the LAN credentials set in DCNM for accessing the switch are incorrect. There may be other reasons for which enablement of the EPL feature may fail. In all scenarios, an appropriate error message is displayed. You can fetch additional context information from `epl.log`.

After the EPL is enabled successfully, all the debug, error, and info logs associated with endpoint information are stored in `bgp.log`. Depending on the scale of the network and the number of endpoint events, the file size will increase. Therefore, there is a restriction on the maximum number and size of `bgp.log`. Up to 10 such files will be stored with each file size of maximum of 10MB.

The EPL relies on BGP updates to get endpoint information. In order for this to work, the switch loopback or VTEP interface IP addresses must be discovered on the DCNM for all switches that have endpoints. To validate, navigate to the Cisco DCNM **Web UI > Dashboard > Switch > Interfaces** tab, and verify if the IP address and the prefix associated with the corresponding Layer-3 interfaces (typically loopbacks) are displayed correctly.

In a Cisco DCNM Cluster deployment, if EPL cannot establish BGP peering and the active DCNM is able to ping the loopback IP address of the spine, while the EPL container cannot, it implies that the eth2 port group for Cisco DCNM and its computes does not have Promiscuous mode set to **Accept**. After changing this setting, the container can ping the spine and EPL will establish BGP.

In a large-scale setup, it may take more than 30 seconds (default timer set in Cisco DCNM) to get this information from the switch. If this occurs, the `ssh.read-wait-timeout` property (in the **Administration > DCNM Server > Server Properties**) must be changed from 30000 (default) to 60000 or a higher value.

## Streaming Telemetry for LAN Deployments

In today's data center environments, granular visibility and tracking of network events has become critical. The traditional polling-based methods that pull the network state in predefined intervals need a fork-lift upgrade. More advanced streaming approaches are required that provide network event visibility in closer to real time through a push method. Streaming telemetry not only allows data to be pushed out at a much finer granularity with a lower cadence (shorter interval) but it also enables event-based notifications. While getting relevant data in a timely fashion is highly desirable, the data needs to be analyzed and converted into actionable insights.

As a first step toward LAN analytics, DCNM 11.0(1) enables subscriptions for environmental metrics through streaming telemetry for consumption and analysis. The environmental metrics that are streamed include CPU, Memory, Power, Temperature, and Fan Speed; all these are enabled with a single click. DCNM allows you to configure the streaming interval for these metrics. The default streaming interval for CPU, Memory is set to 30 seconds, and those for Power, Temperature, and Fan Speed is set to 300 seconds (5 minutes).

Starting from DCNM 11.1(1), subscriptions are enabled for Interface, Transceivers, Control Plane, and Resource summary metrics through streaming telemetry for consumption and analysis. DCNM allows you to

configure the streaming interval for these metrics. The default streaming interval for Interface, Transceivers, Control Plane, and Resource summary metrics is set to 30 seconds.

The per-metric real-time streaming dashboards allow filtering on a per fabric and per switch level including a per-switch drill-down where applicable. Streaming telemetry is currently supported on the Nexus 9000 platforms.

## Guidelines and Recommendations

- **Cisco DCNM LAN Telemetry is a preview-only feature. Do not enable this feature in the production environments.**
- In a cluster mode, a minimum of three compute nodes have to be up for LAN Telemetry to start properly. However, LAN Telemetry functions properly if any one of the three compute nodes is intermittently down.
- If two compute nodes go down, both nodes have to be restored for Zoo Keeper and Kafka Connect to bootstrap correctly and resume data transmission.
- We recommend using the LAN Telemetry feature for up to 30 switches.
- The LAN Telemetry feature is not supported on switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images.

## Pre-Requisites for Enabling the LAN Telemetry Feature

- The Cisco Nexus 9000 switches and Cisco DCNM need to be time synchronized (NTP is recommended).
- Minimum software version on the Nexus 9000 switches must be 7.0(3)I6(1) or higher.
- In the LAN Classic mode, you need to manually enable the following configurations on all the switches before enabling telemetry:
  - **feature nxapi**
  - **nxapi http port 80**



---

**Note** If the preceding configurations are unavailable on the switches, the telemetry health on Cisco DCNM does not show the configurations and the connection status for the telemetry-enabled switches. The preceding commands can be manually defined in a new template, and then pushed to all the switches in the fabric from Cisco DCNM. Use an unused port (for example, port 80) configure nxapi.


---

# Enabling the Streaming Telemetry Feature

## Procedure

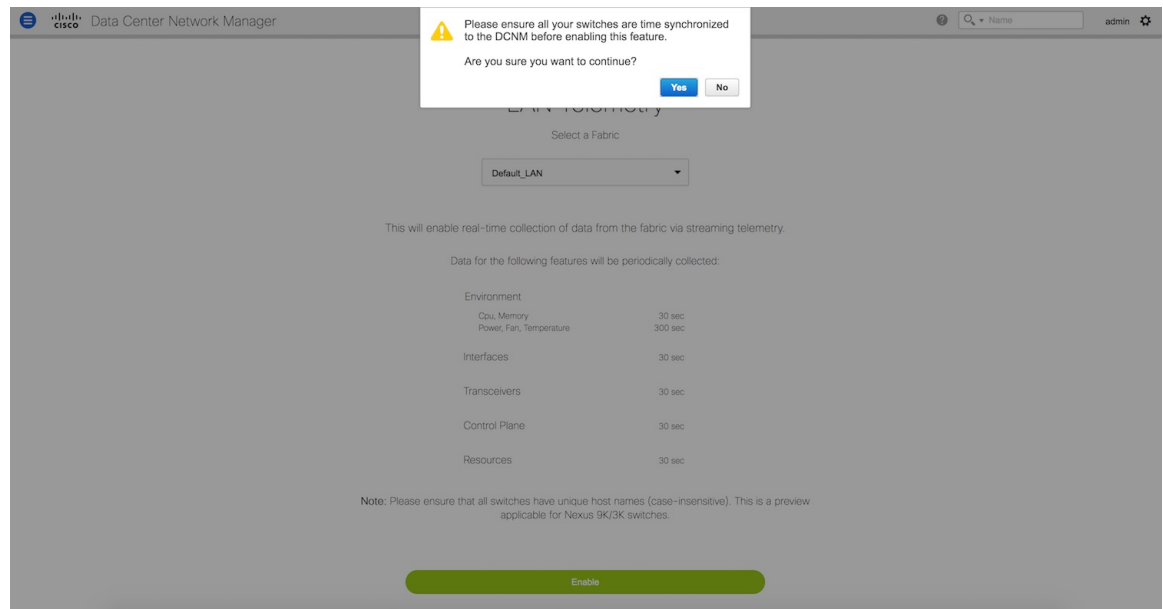
---

- Step 1** Choose **Control > LAN Telemetry > Configure**. Select the fabric for which LAN Telemetry has to be enabled. Then press the **Enable** button.



A warning message appears to indicate that the Cisco DCNM and switches need to be time-synchronized before this feature is enabled. Recall, that this is a prerequisite for this feature. If the prerequisite is met, acknowledge by clicking **Yes**.

**Note** When Telemetry is enabled, the NTP configuration is done on the switches for LAN Classic deployment, wherein the NTP server address is set to DCNM's out-of-band interface's IPv4 address. In case of HA setups, the NTP server address is set to the VIP address of the out-of-band interface. Ensure that the NTP configurations are not removed/modified from the switches.

**Step 2**

Once this feature is enabled, a message appears indicating the initialization process has begun, which takes a couple of minutes. This time is needed for the streaming configuration to be pushed to the switches. The initial data to be streamed out from the switches, which are consumed by DCNM, and depicted on the LAN telemetry dashboard.

Once the LAN telemetry preview feature is enabled, DCNM updates the switch telemetry configuration for the environmental metrics. Every switch that does not conform to the telemetry requirements (must be Cisco Nexus 9000) is excluded from the configuration update. The status of the switch configuration can be monitored by choosing **Control > LAN Telemetry > Health**.

Once the jobs are successfully executed, the required telemetry configuration has been applied to the switches and the streaming data appears once received and processed.

## LAN Telemetry Health

The LAN Telemetry Health window provides a detailed break-down of how much data is streamed out by each switch per feature for the last 24 hours. This window shows the status of the configuration for every switch, apart from showing the statistics of the received data for every metric from every switch. The Connection Status indicates the status of the connection used to transport telemetry data between the switch and DCNM.

When the connection status is shown as **Disconnected** the port configuration may not be accepted by the switch correctly. On the switch image 7.0(3)I7(6), if a switch already had **nxapi** configuration, and later it was managed by DCNM and telemetry was enabled on that fabric, DCNM pushes **http port 80** configuration so that it could query some NXAPI commands such as **show telemetry transport** and **show telemetry data collector details**, to monitor telemetry connection statistics. In this case, the switch does not update **http port 80** in its configuration even though the command was executed correctly. In such a scenario, issue the following commands on the switch:

```
switch# configure
switch(config)# no feature nxapi
switch(config)# feature nxapi
switch(config)# http port 80
```



**Note** You cannot configure ICAM on the Cisco Nexus 9000 Series Switches Release 7.0(3)I7(6), and therefore, the telemetry will fail until the switch issue is resolved.

To view the LAN Telemetry Health, perform the following steps:


## Procedure

### Step 1 Choose **Control > LAN Telemetry > Health**.

Name	Description	Additional Information	Update Period (seconds)	Packets Sent	Configuration Status	Connection Status
ste-n9k-10	N9K-C9396PX NXOS 9.2(1)	SAL18422FXL Default_LAN...		121744	✓ SUCCESS	✓ Connected
ste-n9k-11	N9K-C9396PX NXOS 9.2(1)	SAL18432P11 Default_LAN...		121375	✓ SUCCESS	✓ Connected
ste-n9k-18-deep	N9K-C9396PX NXOS 9.2(2)	SAL18432P61 Default_LAN...		100718	✓ SUCCESS	✓ Connected
ste-n9k-9	N9K-C9396PX NXOS 7.0(3)...	SAL1833YM60 Default_LA...		92293	✓ SUCCESS	✓ Connected
ste-n9k-bg1	N9K-C93180YC-EX NXOS ...	FDO21061Q4W Default_LA...		0	✓ SUCCESS	✓ Connected

**Step 2** Click the **Top Streamers** tab to view the graphs that depicts the top five streaming switches and has a drill-down capability for a feature-wise break-down.



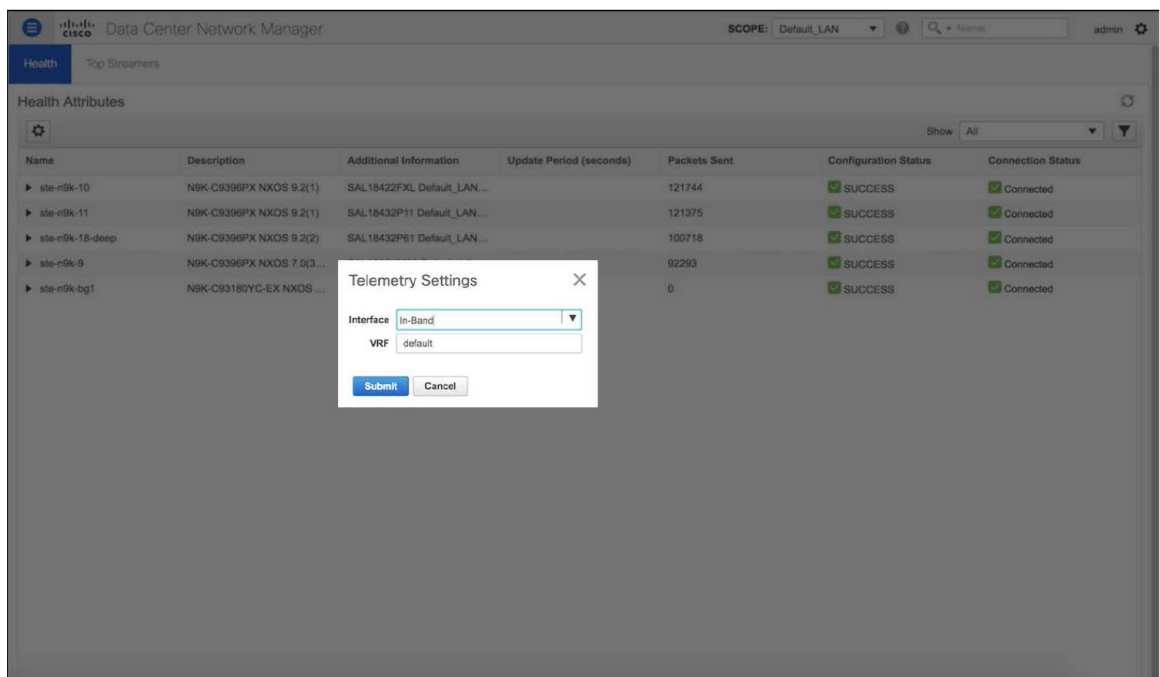
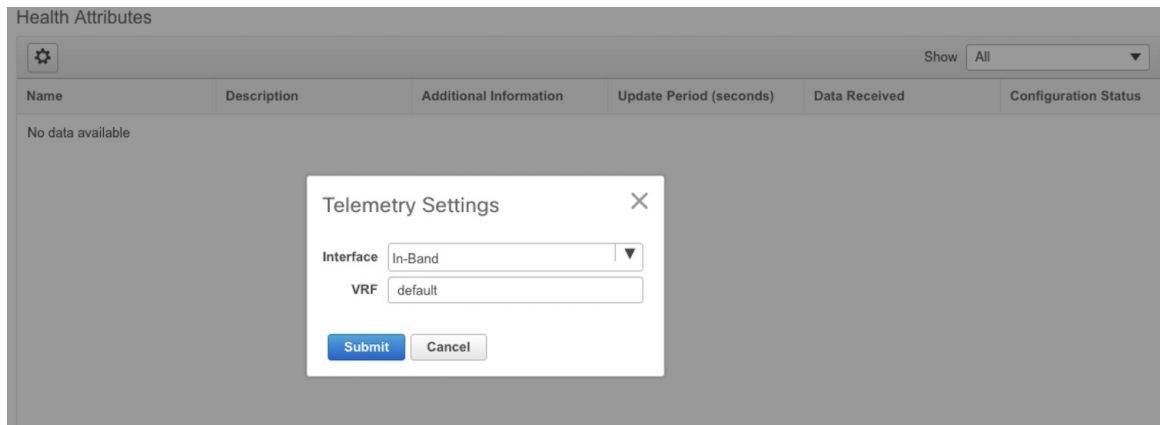


## Telemetry Streaming Interface

Telemetry data, by default is streamed through the management interface of the switches to the Cisco DCNM. This is the Out-of-Band network. This is a global configuration for all fabrics or switch-groups in DCNM. The switches can also stream the Telemetry data through their front panel ports to DCNM assuming there's connectivity from the switches to the DCNM. This is the In-band network. To use the in-band network, do the following:

### Procedure

- Step 1** Disable Telemetry on all the Enabled fabrics.
- Step 2** Go to the Health window and change the settings by clicking on the gear icon on the Health window. In the Telemetry Settings window that comes up, select **In-Band** from the Interface drop-down list. The VRF option is set to default. Click Submit.



The VRF option is set to default. You can also change this to any other VRF assuming that the switches are configured with the VRF and there is availability to the in-band interface of the DCNM through that VRF.

**Note** If Telemetry is already enabled for some fabrics, you should first disable Telemetry on all the enabled fabrics and only then modify the Telemetry network setting. After modifying the Telemetry network settings, you can enable Telemetry on the fabrics. Now, Telemetry data start coming through the in-band interface.