



## **Cisco DCNM LAN Fabric Configuration Guide, Release 11.0(1)**

### **Americas Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

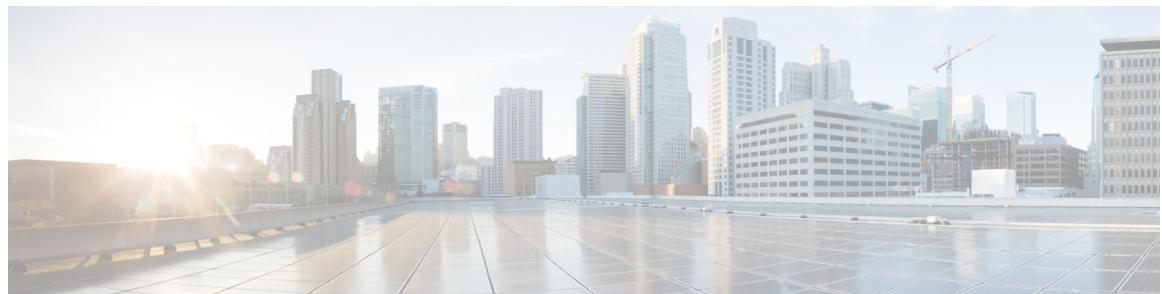
NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2019 Cisco Systems, Inc. All rights reserved.



## CONTENTS

---

### CHAPTER 1

#### Overview 1

Cisco Data Center Network Manager 1

---

### CHAPTER 2

#### Dashboard 3

Dashboard 3

Dashlets 4

---

### CHAPTER 3

#### Topology 7

Topology 7

Status 7

Scope 8

Searching 8

    Quick Search 8

    Host name (vCenter) 8

    Host IP 9

    Host MAC 9

    Multicast Group 9

    VXLAN ID (VNI) 9

    VLAN 9

    VXLAN OAM 9

    Show Panel 10

    Layouts 11

    Zooming, Panning, and Dragging 12

    Switch Slide-Out Panel 12

        Beacon 12

        Tagging 12

More Details	12
Link Slide-Out Panel	12
24-Hour Traffic	13
vCenter Compute Visualization	13
Enabling vCenter Compute Visualization	14
Using vCenter Compute Visualization	15
Troubleshooting vCenter Compute Visualization	19

---

**CHAPTER 4****Control** **21**

Fabrics	21
VXLAN BGP EVPN Fabrics Provisioning	21
Create a New VXLAN BGP EVPN Fabric	22
Add Switch Instances to the Fabric	30
Return Material Authorization (RMA)	40
Interfaces	45
Adding Interfaces	47
Editing Interfaces	48
Deleting Interfaces	49
Shutting Down and Bringing Up Interfaces	49
Viewing Interface Configuration	49
Rediscovering Interfaces	50
Viewing Interface History	50
Deploying Interface Configurations	50
Networks and VRFs Creation and Deployment in a Standalone Fabric	51
Post DCNM 10.4(2) to DCNM 11.0(1) Upgrade Procedure for VXLAN BGP EVPN Fabrics	78
Multi-Site Domain for VXLAN BGP EVPN Fabrics	85
Movement of a Standalone Fabric (With Existing Networks and VRFs) to an MSD Fabric	107
NFM Fabric Migration to a DCNM Fabric	107
Migrate an NFM-Managed VXLAN BGP EVPN Fabric to DCNM 11	107
Upgrade from DCNM 10.4(2) with NFM Overlay Migrations to DCNM 11	122
Post Migration Operations	122
Freeform Configurations on Fabric Switches	126
Management	134
Resources	134

Adding, Editing, Re-Discovering and Removing VMware Servers	134
Adding a Virtual Center Server	134
Deleting a VMware Server	135
Editing a VMware Server	135
Rediscovering a VMware Server	135
Template Library	136
Template Structure	137
Template Format	137
Template Variables	144
Variable Meta Property	145
Variable Annotation	150
Templates Content	153
Advanced Features	154
Adding a Template	157
Modifying a Template	157
Copying a Template	158
Deleting a Template	159
Importing a Template	159
Exporting a Template	159
Image Management	160
Image Upload	160
Deleting an Image	160
Image Upload	160
Install & Upgrade	161
Upgrade History	161
Switch Level History	166
Endpoint Locator	167
Endpoint Locator	167
Configuring Endpoint Locator	168
Configuring the In-band Port	176
Configuring Endpoint Locator in DCNM High Availability Mode	177
Adding High Availability Node to Endpoint Locator Configuration	179
Configuring Endpoint Locator for External Fabrics	179
Disabling Endpoint Locator	179

Troubleshooting Endpoint Locator	179
Streaming Telemetry for LAN Deployments	181
Guidelines and Recommendations	181
Pre-Requisites for Enabling the LAN Telemetry Feature	181
Enabling the Streaming Telemetry Feature	182
LAN Telemetry Health	183
Telemetry Streaming Interface	186

---

**CHAPTER 5****Monitor** 189

Inventory	189
Viewing Inventory Information for Switches	189
Viewing System Information	191
VXLAN	191
FEX	192
VDCs	195
Switch On-Board Analytics	202
Viewing Inventory Information for Modules	207
Viewing Inventory Information for Licenses	207
Monitoring Switch	208
Viewing Switch CPU Information	208
Viewing Switch Memory Information	208
Viewing Switch Traffic and Errors Information	209
Viewing Switch Temperature	209
Enabling Temperature Monitoring	210
Viewing Accounting Information	210
Viewing Events Information	210
Monitoring LAN	211
Monitoring Performance Information for Ethernet	211
Monitoring ISL Traffic and Errors	212
Monitoring a vPC	213
Monitoring vPC Performance	214
Monitoring Endpoint Locator	215
Exploring Endpoint Locator Details	215
LAN Telemetry	222

Monitoring LAN Telemetry	223
Alarms	230
Viewing Alarms and Events	230
Monitoring and Adding Alarm Policies	231
Activating Policies	233
Deactivating Policies	234
Importing Policies	234
Exporting Policies	234
Editing Policies	234
Deleting Policies	235

---

**CHAPTER 6****Administration** 237

DCNM Server	237
Starting, Restarting, and Stopping Services	237
Viewing Log Information	238
Server Properties	238
Modular Device Support	239
Managing Licenses	239
License Assignments	240
Server License Files	241
Native HA	242
Multi Site Manager	243
Management Users	244
Remote AAA	244
Local	244
Radius	245
TACACS+	245
Switch	245
LDAP	245
Managing Local Users	246
Adding Local Users	246
Deleting Local Users	246
Editing a User	247
User Access	247

Managing Clients	247
Performance Setup	248
Performance Setup LAN Collections	248
Event Setup	248
Viewing Events Registration	248
Notification Forwarding	249
Adding Notification Forwarding	249
Removing Notification Forwarding	251
Event Suppression	251
Add Event Suppression Rules	251
Delete Event Suppression Rule	252
Modify Event Suppression Rule	252
Credentials Management	253
LAN Credentials	253

---

**CHAPTER 7****Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - Multi-Site** **257**

Prerequisites	257
Limitations	258
Sample Scenario	259
EVPN Multi-Site Configuration	261
EVPN Multi-Site Extensions from BGW_3 to RS_1	261
Underlay Extension from BGW_3 to RS_1	261
Overlay Extension from BGW_3 to RS_1	270
Other EVPN Multi-Site Configurations	275
Deploying Networks and VRF Instances	277
Deploying Networks on the BGWs	277
Configurations in site1	281
Additional References	282
Appendix	282
Route Server Configurations	282

---

**CHAPTER 8****Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - VRF Lite** **285**

Prerequisites	285
Sample Scenario	286

VRF Lite Configuration	<b>288</b>
VRF Lite Configuration (on BL-1 towards ER-1 in 9K-FABRIC)	<b>288</b>
Extension from BL-1 to ER-1	<b>288</b>
VRF Lite Configuration (on BL-2 towards ER-1 in 9K-FABRIC)	<b>293</b>
Edge Router Configurations	<b>295</b>
Deploying VRF Instances on Border Leafs	<b>295</b>
Resources	<b>300</b>
Undeploying VRF Instances on the Border Leafs	<b>300</b>
Resources Update	<b>303</b>
Remove VRF Lite Inter-fabric configuration on vPC border leafs	<b>303</b>
Additional References	<b>305</b>
Appendix	<b>305</b>
Edge Router Configurations	<b>305</b>





# CHAPTER 1

## Overview

---

Cisco Data Center Network Manager (Cisco DCNM) automates the infrastructure of Cisco Nexus 5000, 6000, 7000, and 9000 Series Switches and Cisco MDS 9000 Series switches. Cisco DCNM enables you to manage multiple devices, while providing ready-to-use control, automation, monitoring, visualization, and troubleshooting capabilities.

The Cisco DCNM home page contains a navigation pane to the left, and shortcuts to a few Cisco DCNM features in the middle pane.

This guide provides comprehensive information about the UI functionalities for the LAN Fabric deployment functionalities.

The top pane displays the following UI elements:

- **Help (?)**: Launches the context-sensitive online help.
- *User Role*: Displays the role of the user who is currently logged in, for example, admin.
- **Gear** icon: Displays information about Cisco DCNM, enables you to change the Cisco DCNM UI password, and allows you to log out from Cisco DCNM UI.

For more information about Cisco DCNM, see:

<https://www.cisco.com/c/en/us/support/cloud-systems-management/data-center-network-manager-11-model.html>.

- [Cisco Data Center Network Manager, on page 1](#)

## Cisco Data Center Network Manager

Cisco Data Center Network Manager automates the infrastructure of Cisco Nexus 5000, 6000, 7000, and 9000 Series Switches and Cisco MDS 9000 Series switches. Cisco DCNM enables you to manage multiple devices, while providing ready-to-use control, automation, monitoring, visualization, and troubleshooting capabilities.

The Cisco DCNM home page contains a navigation pane to the left, and shortcuts to a few Cisco DCNM features in the middle pane.

This guide provides comprehensive information about the LAN Fabric deployment functionalities.

The top pane displays the following UI elements:

- **Help (?)**: Launches the context-sensitive online help.

- **User Role:** Displays the role of the user who is currently logged in, for example, admin.
- **Gear icon:** Displays information about DCNM, enables you to change Cisco DCNM UI password, and allows you to log out from the Cisco DCNM UI.

For more information about Cisco DCNM, see:

<https://www.cisco.com/c/en/us/support/cloud-systems-management/data-center-network-manager-11/model.html>.



## CHAPTER 2

# Dashboard

---

This chapter contains the following topics:

- [Dashboard, on page 3](#)

## Dashboard

The intent of **Dashboard** is to enable network and storage administrators to focus on particular areas of concern around the health and performance of data center switching. This information is provided as 24-hour snapshots. The functional view of LAN switching consists of six dynamic dashlets that display information in the context of the selected scope by default. The scope can be adjusted in the upper right corner of the window to display focused information that is particular to the managed domain. It offers details of a specific topology or set of topologies that is a part of the data center scope.

The various scopes that are available on the Cisco Data Center Network Manager (DCNM) web interface are:

- **Data Center**
- **Default\_SAN**
- **Default\_LAN**
- Each SAN Fabric
- Custom scopes that you create

From the left menu bar, choose **Dashboard**. The **Dashboard** window displays the default dashlets.

The following are the default dashlets that appear in the **Dashboard** window:

- Events
- Data Center
- Network Map
- Link Traffic
- Audit Log
- Server Status

From the **Dashlets** drop-down list, you can choose more dashlets so that they are added to the dashboard.

The panels can be added, removed, and dragged around to reorder.

## Dashlets

By default, a subset of the available dashlets is automatically displayed in the dashboard. To add a dashlet that is not automatically displayed in a dashboard, from the Cisco DCNM Web UI, perform the following steps:

### Procedure

**Step 1** Choose **Dashboard**.

**Step 2** From the **Dashlets** drop-down list, choose the dashlet that you want to add in the dashboard.

In the **Dashlets** drop-down list, an icon appears before the selected dashlet.

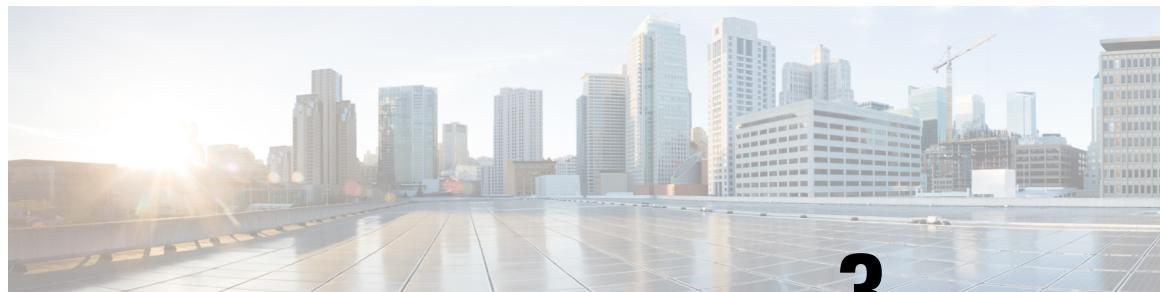
The following table lists the dashlets that you can add on the **Dashboard** window.

Dashlet	Description
Events	Displays events with <b>Critical</b> , <b>Error</b> , and <b>Warning</b> severity. In this dashlet, click the <b>Show Acknowledged Events</b> link to go to the <b>Monitor &gt; Switch &gt; Events</b> .
Link Traffic	Displays a diagram of Inter-Switch Link (ISL) and saturation link for transmitting and receiving in the data center.
Data Center	Displays the number of access, spine and leaf devices, and a generic health score for each switch group in the current scope. Devices are aggregated by type within a switch group.
Audit Log	Displays the accounting log table of Cisco DCNM.
Network Map	<p>Displays the populated switch groups that are visible in your Role Based Access Control (RBAC) scope on a world map. If you use the scope selector, it limits the set of switch groups displayed. If you use the pop-up option, the map opens in a new tab and can be configured.</p> <ul style="list-style-type: none"> <li>• The network map dialog box has properties that are different from the Summary dashboard view:</li> <li>• You can click and drag nodes to move them around the map. The map saves their new positions.</li> <li>• You can double click a node to trigger a slider that contains the summary inventory information pertaining to a specific switch group.</li> </ul>

Dashlet	Description
	<ul style="list-style-type: none"> <li>You can upload an image of your choice as the background to the network map.</li> </ul> <p><b>Note</b> You will be prompted to upload an image file with recommended dimension, which is the current window size. Reset returns the network map to its default state, resetting the position of the nodes and clearing the custom image.</p>
Server Status	Displays the status of DCNM and federation servers, and the health check status for the components.
Top ISLs/Trunks	Displays the performance data for the top ten performing ISLs, trunk ports or both. Each entry shows the current average receive and transmit percentage, with a graph depicting the percentage of time each trunk spent exceeding the currently configured thresholds.
Top SAN End Ports (SAN only)	<p>Displays the performance data for the top ten performing SAN host and storage ports. Each entry shows the current receive and transmit percentage, with a graph depicting the percentage of time each trunk spent exceeding the currently configured thresholds.</p> <p><b>Note</b> This dashlet is only for SAN.</p>
Top CPU	Displays CPU utilization for the discovered switches over the last 24 hours, with a red bar displaying the high watermark for that 24-hour period.
Top Temperature	<p>Displays the module temperature sensor details of switches.</p> <p><b>Note</b> This dashlet is only for LAN.</p>
Health	<p>Displays the health summary that contains two columns displaying the summary of problems and summary of events for the past 24 hours.</p> <p>Click the count adjacent to the warnings pertaining to switches, ISLs, hosts, or storage (other than 0) to view the corresponding inventory for that fabric.</p> <p>Click the count adjacent to the event severity levels (Emergency, Alert, Critical, Error, Warning, Notice, Info, or Debug) to view a summary of the corresponding events and descriptions.</p>

Dashlet	Description
Errors	Displays the error packets for the selected interface. This information is retrieved from the Errors > In-Peak and Errors > Out-Peak columns of the Monitor > LAN / Ethernet page.
Discards	Displays the error packets that are discarded for the selected interface.  <b>Note</b> The Discards dashlet is only for LAN.
Inventory (Ports)	Displays the ports inventory summary information.
Inventory (Modules)	Displays the switches on which the modules are discovered, the models name and the count.
Inventory (ISLs)	Displays the ISLs inventory summary information, such as the category and count of ISLs.
Inventory (Logical)	Displays the logical inventory summary information, such as the category and count of logical links.
Inventory (Switches)	Displays the switches inventory summary information such as the switch models and the corresponding count.
Inventory (Port Capacity)	Displays the port capacity inventory summary information such as the tiers, the number and percentage of the available ports, and the remaining days.

**Note** To restore the default dashlets in the dashboard page, click the **Default Set** link in the **Dashlet** drop-down list.



## CHAPTER 3

# Topology

- [Topology, on page 7](#)

## Topology

The Topology window displays color-encoded nodes and links that correspond to various network elements, including switches, links, fabric extenders, port-channel configurations, virtual port-channels, and more. For information about each of these elements, hover your cursor over the corresponding element. Also, click a node or the line for a link. A slide-in pane appears from the right side of the window. This pane displays detailed information about either the switch or the link.



### Note

You can open multiple tabs simultaneously and can function side by side to facilitate comparison and troubleshooting.

## Status

The color coding of each node and link corresponds to its state. The colors and what they indicate are described in the following list:

- Green: Indicates that the element is in good health and functioning as intended.
- Yellow: Indicates that the element is in warning state and requires attention to prevent any further problems.
- Red: Indicates that the element is in critical state and requires immediate attention.
- Gray: Indicates lack of information to identify the element or the element has been discovered.



### Note

- In the **Topology** window, FEX appears in gray (**Unknown** or **n/a**) because health is not calculated for FEX.

Similarly, in the **Fabric Builder** topology window there is no configuration sync status for the FEX and it appears as **n/a**.)

- Black: Indicates that the element is down.

## Scope

You can search the topology based on the scope. The default scopes available from the **SCOPE** drop-down list is: **DEFAULT\_LAN**

The following search options are available for **DEFAULT\_LAN**:

- Quick Search
- Host name (vCenter)
- Host IP
- Host MAC
- Multicast Group
- VXLAN ID (VNI)
- VLAN
- FabricPath
- VXLAN OAM

## Searching

When the number of nodes is large, it quickly becomes difficult to locate the intended switches and links. You can quickly find switches and links by performing a search. You are also able to search for VM tracker and generic setups. Searching feature enables you to see which leaf the host is connected to.

The following searches are available:



**Note**

By default, Quick Search is selected.

## Quick Search

**Quick Search** enables you to search for devices by name, IP address, model, serial number, and switch role. As you enter a search parameter in the **Search** field, the corresponding switches are highlighted in the topology. To perform a search for multiple nodes and links, separate multiple keywords using a comma, for example, ABCD12345, N7K, sw-dc4-12345, core, 172.23.45.67. Cisco DCNM supports wildcard searches too. If you know a serial number or switch name partially, you can build a search based on these partial terms that are preceded by an asterisk, for example, ABCD\*, sw\*12345, core, and so on.

To limit the scope of your search to a parameter, enter the parameter name followed by a space and the parameter in the Search field, for example, name=sw\*12345, serialNumber=ABCD12345, and so on.

## Host name (vCenter)

The host name search enables you to search for hosts by using vCenter.

## Host IP

You can search the topology using host IP addresses. The **Host IP** searches the switches in the scope to locate the hosts that match the IP address that you enter in the **Search** field. The **Host IP** search supports IPv4 and IPv6 addresses. From the Search drop-down list, choose **Host IP** to search the topology using the IP Address of the host device. Enter a host IP address in the **Search** field and press **Enter**. Click **Details** to view the corresponding host details.

## Host MAC

You can search a topology using host MAC addresses. The **Host MAC** searches the switches in the scope to locate the hosts that match the MAC address that you enter in the **Search** field. From the Search drop-down list, choose **Host MAC** to search the topology using a host MAC address. Enter a host MAC address in the **Search** field and press **Enter**. Click **Details** to view the corresponding host details.

## Multicast Group

The **Multicast Group** search is limited to the VXLAN context, VXLAN tunnel endpoint or VTEP switches, to get VXLAN IDs (VNIs) associated with this multicast address.

Select the **Multicast Group** search from the drop-down list, enter a multicast address in the search field, and press **Enter**. Click the **Details** link next to the search field to get the detailed multicast address table. The table displays switches, which have the searched multicast address configured on them, along with associated VNI, VNI status, and mapped VLAN.

You can also hover over switches that are highlighted to view details about the search you have performed.

## VXLAN ID (VNI)

The VXLAN ID or the VNI search lets you search the topology by VNI. Select the **VXLAN ID (VNI)** search from the drop-down list. Enter a VNI in the search field and press **Enter**. Click the **Details** link next to the search field to view the detailed VNI table. The table displays the switches that have VNI configured on them along with associated multicast address, VNI status, and mapped VLAN.

## VLAN

Search by a given VLAN ID. VLAN search provides the search for the VLAN configured on the switch or the links. If STP is enabled, then it provides information that is related to the STP protocol and the STP information for links.

## VXLAN OAM

You can track details such as reachability and actual path of the flows in a VXLAN EVPN based-fabric topology by choosing the **VXLAN OAM** option from the **Search** drop-down list or by entering **VXLAN OAM** in the **Search** field. This displays the **Switch to switch** and **Host to host** tabs. DCNM highlights the route on the topology between the source and destination switch for these two options.

The **Switch to switch** option provides the VXLAN OAM ping and traceroute test results for the VTEP-to-VTEP use-case. Provide the following values to enable search by using the **Switch to switch** option:

- From the **Source Switch** drop-down list, choose the source switch.
- From the **Destination Switch** drop-down list, choose the destination switch.
- From the **VRF** drop-down list, choose or enter the VRF details.

- Check the **All Path Included** check box to include all the paths in the search results.

The **Host to host** option provides the VXLAN OAM pathtrace results for the exact path that is taken by a given flow from the VTEP or switch that is connected to the source host to VTEP or switch that is connected to the destination host. For the **Host to host** use-case, there are two suboptions:

- VRF or SVI for a network is instantiated on the switches in the VXLAN EVPN fabric. In such a scenario, the IP address information of the end hosts is required.
- Layer 2 configuration for a given network is instantiated on the switches in the VXLAN EVPN fabric. In such a scenario, both the MAC and IP address information of the end hosts are required.

Provide the following values to enable search using the **Host to host** option:

- In the **Source IP** field, enter the IP address of the source host.
- In the **Destination IP** field, enter the IP address of the destination host.
- In the **VRF** field, choose VRF from the drop-down list or enter the VRF name that is associated with the hosts.
- (Optional) In the **Source Port** field, choose Layer 4 source port number from the drop-down list or enter its value.
- (Optional) In the **Destination Port** field, choose destination port number or enter its value.
- (Optional) In the **Protocol** field, choose the protocol value from the drop-down list or enter its value. This is the Layer 4 protocol, usually TCP or UDP.
- Click the **Interchange/Swap Source and Destination IPs (and MACs if applicable)** icon to interchange the source and destination IP addresses. This interchange allows a quick trace of the reverse path without reentering the host IP addresses or MAC addresses.
- Check the **Layer-2 only** check box to search the VXLAN-EVPN fabric that is deployed in Layer 2 only mode for some networks, that is, Layer 2 VNIs. Note that no SVIs or VRFs should be instantiated in the fabric for these networks when you use this search option.

Enter values for the following additional fields:

## Show Panel

You can choose to view your topology based on the following options:

- Auto Refresh:** Check this check box to automatically refresh the topology.
- Switch Health:** Check this check box to view the switch's health status.
- FEX:** Check this check box to view the Fabric Extender.



**Note**

The FEX feature is available only on LAN devices. Therefore, checking this check box displays only the Cisco Nexus switches that support FEX.

**Note**

FEX is also not supported on Cisco Nexus 1000V devices. Therefore, such devices will not be displayed in the topology when you check the **FEX** check box.

- **Links:** Check this check box to view links in the topology. The following options are available:
    - **Errors Only:** Click this radio button to view only links with errors.
    - **All:** Click this radio button to view all the links in the topology.
    - **VPC Only:** Check this check box to view only vPC peer-links and vPCs.
    - **Bandwidth:** Check this check box to view the color coding based on the bandwidth that is consumed by the links.
  - **OTV:** Check this check box to show the Overlay Transport Virtualization (OTV) topology with the cloud icon and the dotted links from the OTV edge devices. Hovering the cursor over the cloud and the links shows the relevant information for OTV topology, such as control group, extended VLANs, and so on. The OTV search field appears below the filter field. Use the OTV search field to search the shown OTV topology that is based on **Overlay ID** and **Extended VLAN ID**. The searched virtual links based on the **Overlay ID** and **Extended VLAN ID** are marked green.
- A **Details** link appears after you check the **OTV** check box. Clicking the link shows the OTV topology data. The **Overlay Network** column shows whether the particular topology is multicast based or unicast based. The **Edge Device** column displays the edge switches in the particular OTV topology. The other columns display the corresponding overlay interface, extended VLANs, join interface, and data group information.
- **UI controls:** Check the check box to show or hide the various controls on the **Topology** window.
  - **Refresh:** You can also perform a topology refresh by clicking the **Refresh** icon in the upper-right corner of this panel.

## Layouts

The topology supports different layouts along with a **Save Layout** option that remembers how you positioned your topology.

- **Hierarchical and Hierarchical Left-Right:** Provide an architectural view of your topology. Various switch roles can be defined that will draw the nodes on how you configure your CLOS topology.

**Note**

When running a large-scale setup, being able to easily view all your switches on a leaf-tier can become difficult. To mitigate this, DCNM splits your leaf-tier every 16 switches.

- **Random:** Nodes are placed randomly on the window. DCNM tries to make a guess and intelligently place nodes that belong together in close proximity.
- **Circular and Tiered-Circular:** Draw nodes in a circular or concentric circular pattern.

## Zooming, Panning, and Dragging

- **Custom saved layout:** Nodes can be dragged around according to your preference. After you position as required, click **Save** to retain the positions. The next time you come to the topology, DCNM will draw the nodes based on your last saved layout positions.

Before a layout is chosen, DCNM checks if a custom layout is applied. If a custom layout is applied, DCNM uses it. If a custom layout is not applied, DCNM checks if switches exist at different tiers, and chooses the Hierarchical layout or the Hierarchical Left-Right layout. Force-directed layout is chosen if all the other layouts fail.

## Zooming, Panning, and Dragging

You can zoom in and zoom out using the controls that are provided at the bottom left of the windows or by using your mouse's wheel.

To pan, click and hold anywhere in the whitespace and drag the cursor up, down, left, or right.

To drag switches, click, hold, and move the cursor around the whitespace region of the topology.

## Switch Slide-Out Panel

You can click on the switch to display the configured switch name, IP address, switch model, and other summary information such as status, serial number, health, last-polled CPU utilization, and last-polled memory utilization.

### Beacon

This button will be shown for switches that support the **beacon** command. After beaconing starts, the button will show a countdown. By default, the beaconing will stop after 60 seconds, but you can stop it immediately by clicking **Stop Beacon**.




---

**Note** The default time can be configured in `server.properties` file. Search for **beacon.turnOff.time**. The time value is in milliseconds. Note that this requires a server restart to take effect.

---

### Tagging

Tagging is a powerful yet easy way to organize your switches. Tags can be virtually any string, for example, *building 6, floor 2, rack 7, problem switch, and Justin debugging*.

Use the search functionality to perform searches based on tags.

### More Details

Click **Show more details**; detailed information appears in the switch's dashboard.

## Link Slide-Out Panel

You can click a link to view the status and the port or switches that describe the link.

## 24-Hour Traffic

This feature requires **Performance Monitoring** to be turned **ON**. When **Performance Monitoring** is **ON**, traffic information is collected and the aggregate information is displayed along with a graph showing traffic utilization.

## vCenter Compute Visualization

In virtualized environments, any kind of troubleshooting starts with identifying the network attachment point for the virtual machines. This means that a quick determination of the server, virtual switch, port group, VLAN, associated network switch, and physical port is critical. This requires multiple touch points and interactions between the server and the network administrator as well as reference to multiple tools (compute orchestrator, compute manager, network manager, network controller, and so on).

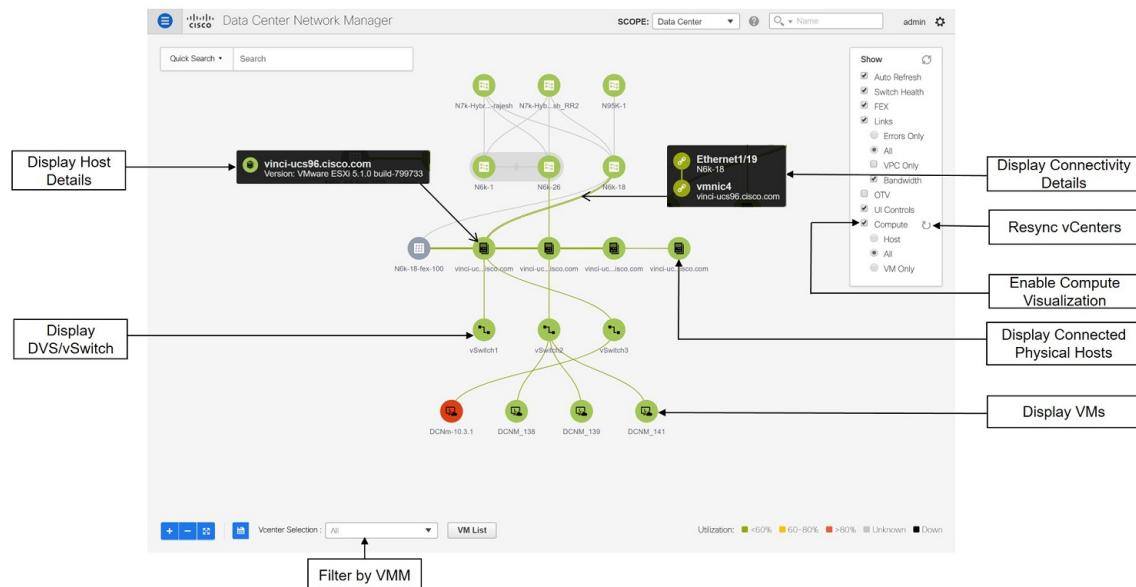
This allows you to visualize the vCenter-managed hosts and their leaf switch connections on the **Topology** window. The visualization options include viewing only the attached physical hosts, only the VMs, or both. When you select both, the topology all the way from the leaf switches to the VMs, including the virtual switches are displayed. The VM Search option highlights the path of the VM. Hover the cursor over a host or a connected uplink to view key information relevant to that entity. Up to four vCenters are supported.



### Note

- The vCenter Compute Visualization feature is supported on both the LAN Classic and Easy Fabrics installations for the vCenter-managed computes.
- It is not recommended to use special characters in a VM name as vCenter does not escape special characters used in display names. For more information, see <https://vss-wiki.eis.utoronto.ca/display/VSSPublic/Virtual+Machine+Naming>.

**Figure 1: vCenter Compute Visualization**



## Enabling vCenter Compute Visualization

To enable the vCenter Compute Visualization feature from the Cisco DCNM Web UI, perform the following steps.

### Procedure

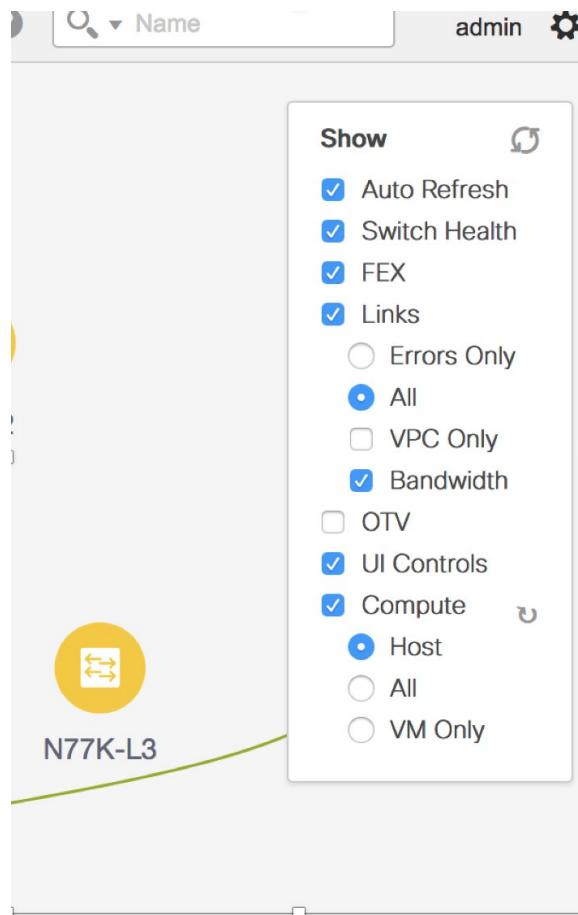
**Step 1** Choose Control > Management > Virtual Machine Manager.

The **Control > Management > Virtual Machine Manager** window appears.

**Step 2** Click the + icon to add a new VMware vSphere vCenter.

**Step 3** Enter the server IP address, username, and password to the vCenter. vCenter version 5.5 or later is required.

After the initial discovery, the information that is received from the vCenter is appropriately organized and displayed on the main **Topology** window. An extra menu item labeled **Compute** appears on the **Show** pane.



## Using vCenter Compute Visualization

To use the vCenter Compute Visualization feature from the Cisco DCNM Web UI, perform the following steps.

### Procedure

**Step 1** Choose **Topology**.

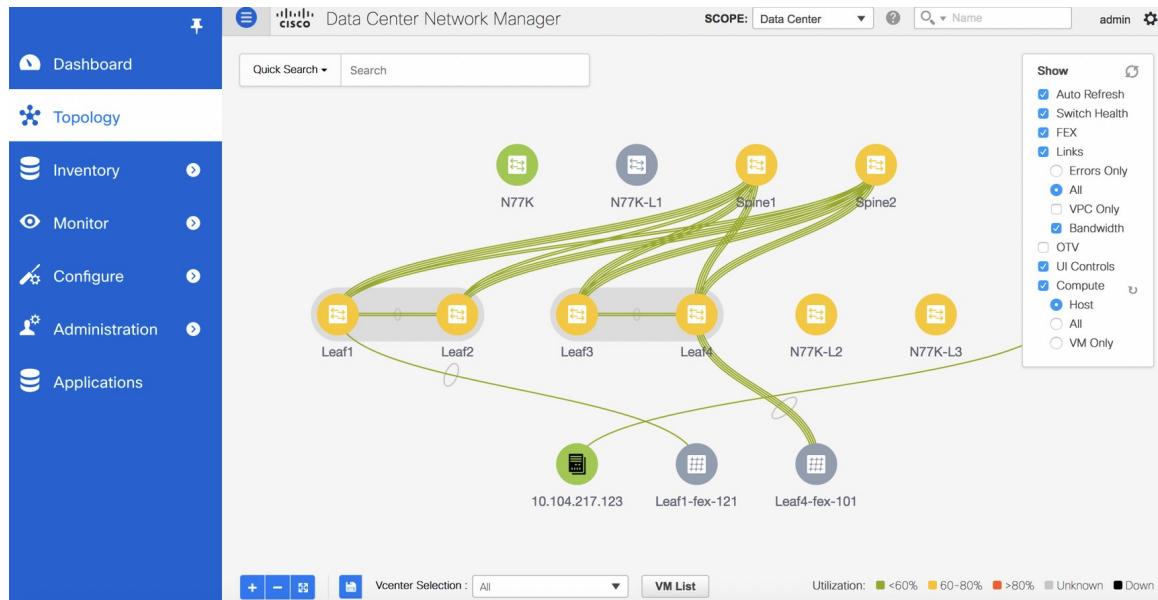
**Step 2** In the **Show** list, select **Compute** to enable the compute visibility.

By default, the **Host** check box is selected. This implies that the topology shows the VMWare vSphere ESXi hosts (servers), that are attached to the network switches.

The following options are available in the Compute Visualization feature.

- **Host**
- **All**
- **VM Only**

## Using vCenter Compute Visualization



In the **All** mode, you can see double-arrows that help you to extend a node. If you double-click this node, you can see all the hidden child nodes.

### Step 3 Click a specific ESXi host to view additional information.

The expanded topology displayed in the following figure, shows the virtual switches (both vSwitch and Distributed Virtual Switch) that are configured on the specific ESXi host.



### Step 4 When changing from the **Host** suboption to the **All** suboption, all the compute resources are expanded.

When **All** is selected, an expanded view of all the hosts, virtual switches, and virtual machines that are part of the topology are displayed. If a VM is powered off, it is shown in red color; otherwise, it is shown in green color.

**Note** The vCenter search is unavailable when compute visualization is not enabled. Also, this search is available only when you select the **All** option.

### Step 5 Instead of browsing through the large set of available information, to focus on a specific VM.

Enter a host name (vCenter) in the **Search** field at the top-left. When you start entering the characters, the topology is instantaneously updated with matching objects.

## Using the Virtual Machine List

The **Virtual Machine List** allows you to view the complete list of virtual machines.

### Procedure

#### Step 1 Choose Topology.

#### Step 2 Click VM List.

VM Name	VLAN	Virtual Sw...	Physical ...	Host	Switch	Switch Int...	Connection
1 DCNM_138	0	vSwitch2	vmnic3	vinci-ucs14...	N6k-26	Ethernet1/35	connected
2 DCNM_139	0	vSwitch2	vmnic3	vinci-ucs14...	N6k-26	Ethernet1/35	connected
3 DCNM_141	0	vSwitch2	vmnic3	vinci-ucs14...	N6k-26	Ethernet1/35	connected
4 DCNm-10.3.1	0	vSwitch3	vmnic3	vinci-ucs96...	N6k-18	Ethernet1/18	connected

Click **Export** to export the list of virtual machines into a .csv file.

Click on the name of a VM to view additional information about that virtual machine.

## Resynchronizing Virtual Machines

**VM List**

VM Name	VLAN	Virtual Sw...	Physical ...	Host	Switch	Switch Int...	Connection
1 DCNM_138	0	vSwitch2	vmnic3	vinci-ucs14...	N6k-26	Ethernet1/35	connected
	0	vSwitch2	vmnic3	vinci-ucs14...	N6k-26	Ethernet1/35	connected
	0	vSwitch2	vmnic3	vinci-ucs14...	N6k-26	Ethernet1/35	connected
	0	vSwitch3	vmnic3	vinci-ucs96...	N6k-18	Ethernet1/18	connected

**DCNM\_138**

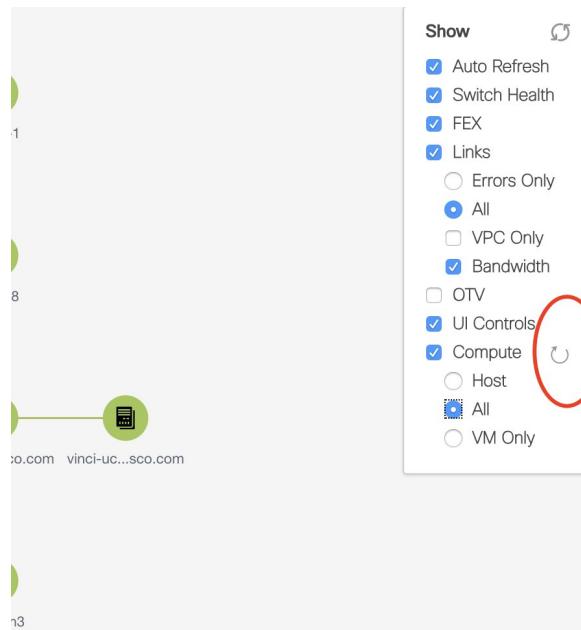
Network Adapter 1 – 00:59:56:93:8f:ds – Disconnected  
Network adapter 2 – 00:59:56:93:57:c2 – Disconnected  
Network adapter 3 – 00:59:56:93:32:96 – Disconnected  
No. of CPUs: 8  
Memory(MB): 24576  
Guest OS: Linux 3.10.0-693.21.1.el7.x86\_64 CentOS Linux release 7.4.1708 (Core)  
Guest IP: 172.28.10.138  
PRODUCT INFO:  
• Name: DCNM Virtual Appliance  
• Vendor: Cisco Systems, Inc.  
• Version: 11.0.0.700

**Note** When you export the VM List to a .CSV file, the .CSV file may appear correct. However, when the .CSV file is imported into Microsoft Excel, it might get reformatted, for example, the VLAN column 1-1024 could be reformatted to a date 1/1/2019. Therefore ensure that columns are formatted correctly in Microsoft Excel while importing the .CSV file.

## Resynchronizing Virtual Machines

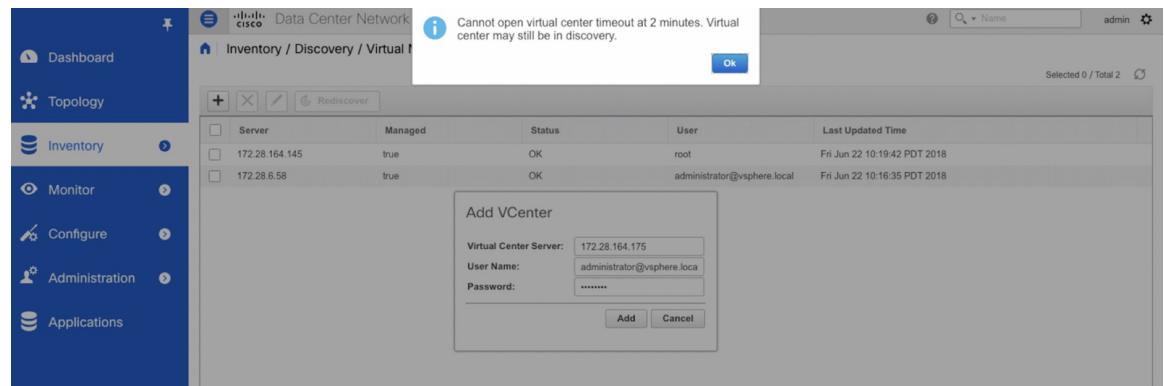
### Procedure

- Step 1** Choose **Topology**.
- Step 2** Click **Resync vCenters** icon next to **Compute**.



## Troubleshooting vCenter Compute Visualization

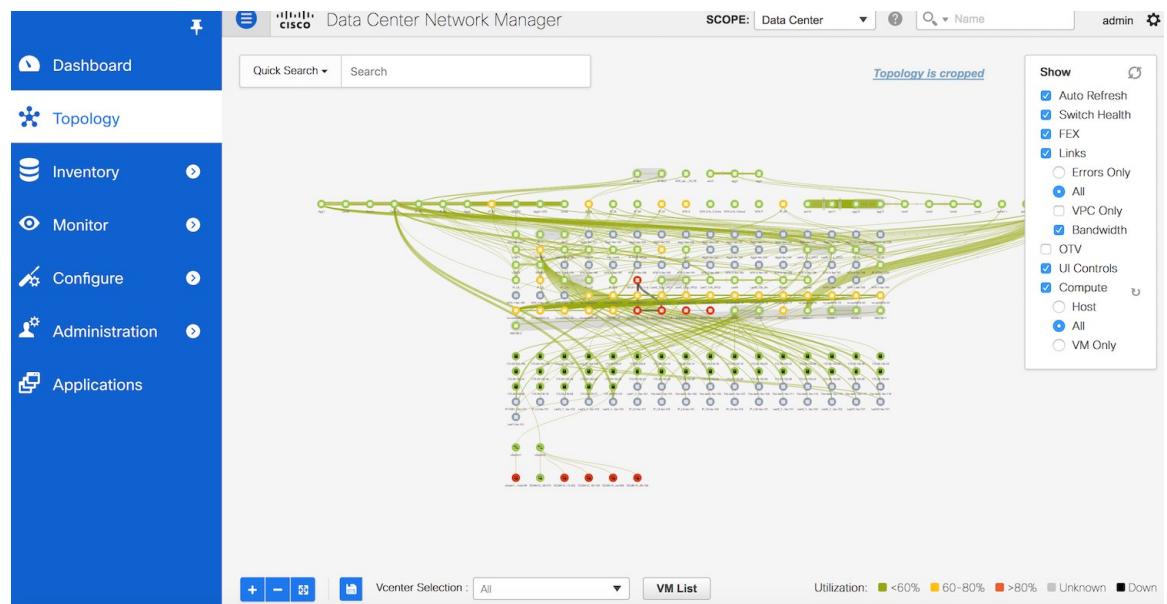
The following error window appears when the vCenter times out. This error might occur when the discovery of the vCenter is in progress.



## Viewing Topology in Scale Mode

The following window shows how the **Topology** window appears after about 200 devices are available in the topology. Note that the topology graph is trimmed down at scale.

## Viewing Topology in Scale Mode





## CHAPTER 4

# Control

---

- [Fabrics, on page 21](#)
- [Management, on page 134](#)
- [Template Library, on page 136](#)
- [Image Management, on page 160](#)
- [Endpoint Locator, on page 167](#)
- [Streaming Telemetry for LAN Deployments, on page 181](#)

## Fabrics

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

### VXLAN BGP EVPN Fabrics Provisioning

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

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

Fabric creation and updation:

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

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

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

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

## Create a New VXLAN BGP EVPN Fabric

Underlay provisioning:

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

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

Overlay network provisioning.

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

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

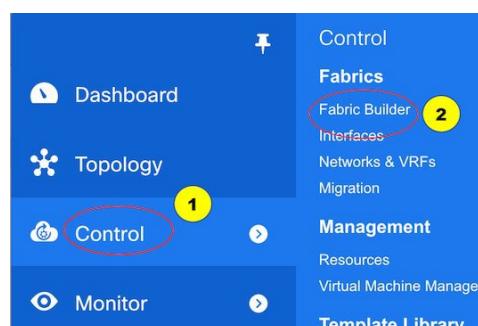
NFM fabric migration to the VXLAN BGP EVPN fabric.

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

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

## Create a New VXLAN BGP EVPN Fabric

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



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

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

**Create Fabric**

**Fabrics (4)**

- Easy60000** (Switch\_Fabric, ASN: 60000, Replication Mode: Multicast, Technology: VXLANFabric)
- Easy7200** (Switch\_Fabric, ASN: 7200, Replication Mode: Multicast, Technology: VXLANFabric)
- External65000** (External, ASN: 65000)
- MSD** (MSD, Member Fabrics: Easy60000)

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

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

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

Choose **Easy\_Fabric**. The fabric creation window for creating a standalone fabric comes up.

**Add Fabric**

\* Fabric Name :

\* Fabric Template :

**General** **Advanced** **Resources** **Manageability** **Bootstrap**

* BGP ASN	<input type="text"/> <small>1-4294967295   1-65535[.0-65535]</small>
* Fabric Interface Numbering	<input type="text" value="p2p"/> <small>Unnumbered or Numbered (Point-To-Point)</small>
* Link-State Routing Protocol	<input type="text" value="ospf"/> <small>Supported routing protocols (OSPF/IS-IS)</small>
* Replication Mode	<input type="text" value="Multicast"/> <small>Replication Mode for BUM Traffic</small>
* Multicast Group Subnet	<input type="text" value="239.1.0/25"/> <small>Multicast address with prefix 25 to 30</small>
* Anycast Gateway MAC	<input type="text" value="2020.0000.00aa"/> <small>Shared MAC address for all leafs (xxxx.xxxx.xxx)</small>
NX-OS Software Image Version	<input type="text"/> <small>If Set, Image Version Check Enforced On All Sw</small>

**Save** **Cancel**

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



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

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

## Create a New VXLAN BGP EVPN Fabric

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

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

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

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

**Multicast Group Subnet :** Multicast group address of the network.

**Anycast Gateway MAC :** Anycast gateway MAC address.

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

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

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

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

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

## Add Fabric

\* Fabric Name :

\* Fabric Template : Easy\_Fabric

General Advanced Resources Manageability Bootstrap

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

\* Power Supply Mode ps-redundant Default Power Supply Mode For The Fabric

\* CoPP Profile strict Fabric Wide CoPP Policy

Enable VXLAN OAM  For Operations And Management Of VXLAN Fabrics

Enable Tenant Routed Multicast  For Overlay Multicast Support In VXLAN Fabrics

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

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

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

**Save** **Cancel**

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

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

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

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

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

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

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

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

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

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

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

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



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

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

**Enable vPC Advertise PIP** - Enables the Advertise PIP feature.

*Freeform CLIs* - Fabric level freeform CLIs can be added while creating or editing a fabric. They are applicable to switches across the fabric. You must add the configurations as displayed in the running configuration, without indentation. Switch level freeform configurations such as VLAN, SVI, and interface

## Create a New VXLAN BGP EVPN Fabric

configurations should only be added on the switch. Refer the *Freeform Configurations on Fabric Switches* topic for a detailed explanation and examples.

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

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

- Click the **Resources** tab.

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

The fields in this tab are:

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

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

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

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

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

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

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



### Note

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

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

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

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

- Click the **Manageability** tab.

Add Fabric

The screenshot shows the 'Add Fabric' dialog box with the 'Manageability' tab selected. It includes fields for DNS and NTP server details, and a note about AAA server configuration.

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

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

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.

- Click the **Bootstrap** tab.

## Create a New VXLAN BGP EVPN Fabric

Add Fabric

\* Fabric Name :

\* Fabric Template

General Advanced Resources Manageability **Bootstrap**

Enable DHCP  Automatic IP Assignment For POAP  
 Start Address For Switch Out-of-Band POAP  
 End Address For Switch Out-of-Band POAP  
 Default Gateway For Mgmt VRF On The Switch  
 Prefix For Mgmt0 Interface On The Switch (Min:8 Max:30)

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

**Save** **Cancel**

The fields on this tab are:

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

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

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

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

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

Add Fabric

\* Fabric Name :

\* Fabric Template

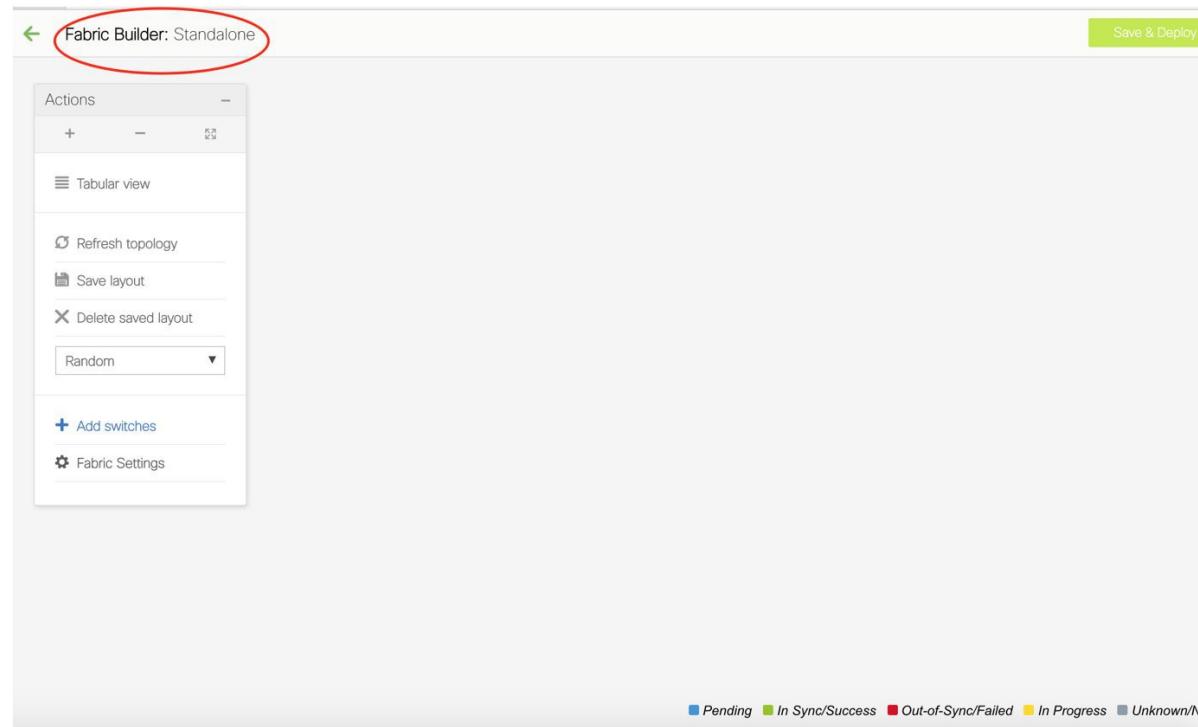
General Advanced Resources Manageability **Bootstrap**

Enable DHCP  Automatic IP Assignment For POAP  
 Start Address For Switch Out-of-Band POAP  
 End Address For Switch Out-of-Band POAP  
 Default Gateway For Mgmt VRF On The Switch  
 Prefix For Mgmt0 Interface On The Switch (Min:8 Max:30)

\* DHCP Scope Start Address   
\* DHCP Scope End Address   
\* Switch Management Default Gateway   
\* Switch Management Subnet Prefix

**Save** **Cancel**

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



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

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

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

### Delete a Fabric

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

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

## Add Switch Instances to the Fabric

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

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

### Discovering Existing Switches

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

## Inventory Management

[Discover Existing Switches](#) [PowerOn Auto Provisioning \(POAP\)](#)

Discovery Information > Scan Details >

Seed IP

172.23.244.91

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

Authentication Protocol

MD5

Username

admin

Password

.....

Max Hops

2

hop(s)

Preserve Config

no  yes

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

[Start discovery](#)

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

## Inventory Management

[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...		

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

## Add Switch Instances to the Fabric

Inventory Management

Discover Existing Switches | PowerOn Auto Provisioning (POAP)

Discovery Information > Scan Details >

**1** Import into fabric

	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...		

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

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

Standalone

Save & Deploy

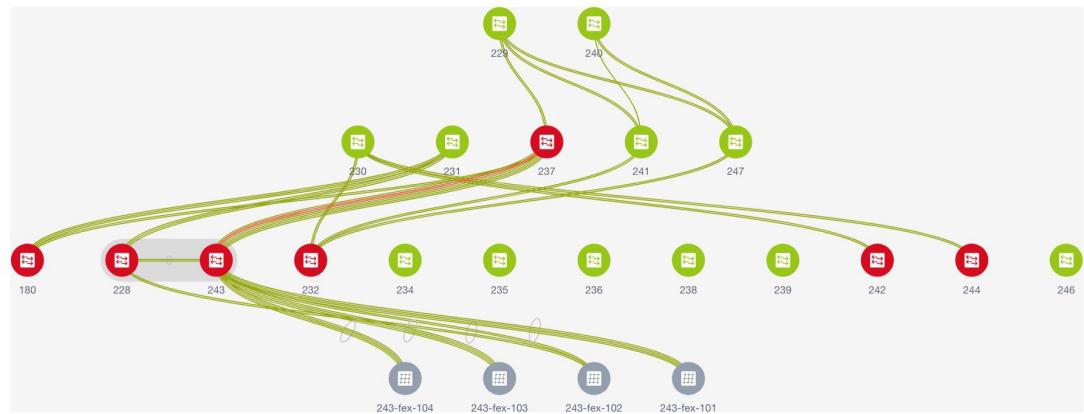
Actions

- + -
- Tabular view
- Refresh topology
- Save layout
- Delete saved layout
- Random
- Add switches
- Fabric Settings

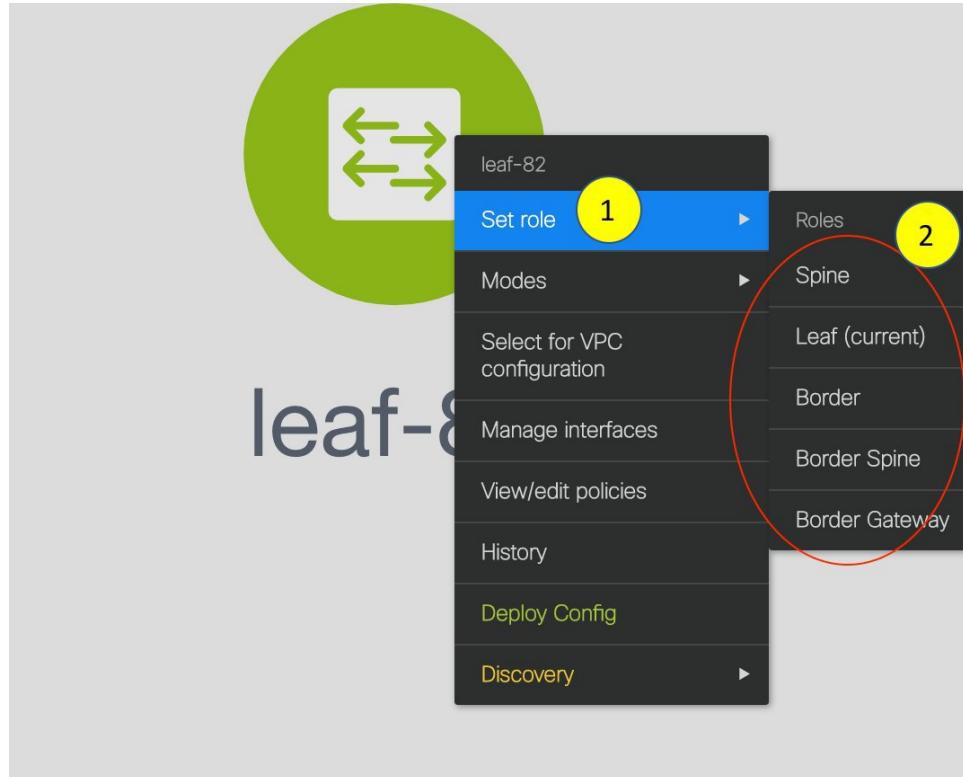
**leaf-91**

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

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



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



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

## Add Switch Instances to the Fabric


**Note**

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

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



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

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

Also, freeform CLIs that were entered earlier are deployed.

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

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

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

## Config Deployment

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

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

Deploy Config

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

## Config Preview - Switch 172.23.244.91

Pending Config	Expected Config	Current Config
<pre> router ospf UNDERLAY   no router-id 10.0.0.3 router-id 10.0.0.2 router bgp 65002 router-id 10.0.0.2 no apply profile MyVRF_50010 no apply profile MyNetwork_30010 no configure profile MyVRF_50000 configure terminal no configure profile MyVRF_50010 configure terminal no configure profile MyNetwork_30010 configure terminal interface loopback0   no ip address 10.0.0.3/32 ip address 10.0.0.2/32   ip router ospf UNDERLAY area 0.0.0.0   ip pim sparse-mode   no shutdown interface loopback1   no ip address 10.0.0.5/32 ip address 10.0.0.3/32   ip router ospf UNDERLAY area 0.0.0.0 </pre>		

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

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

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

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

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

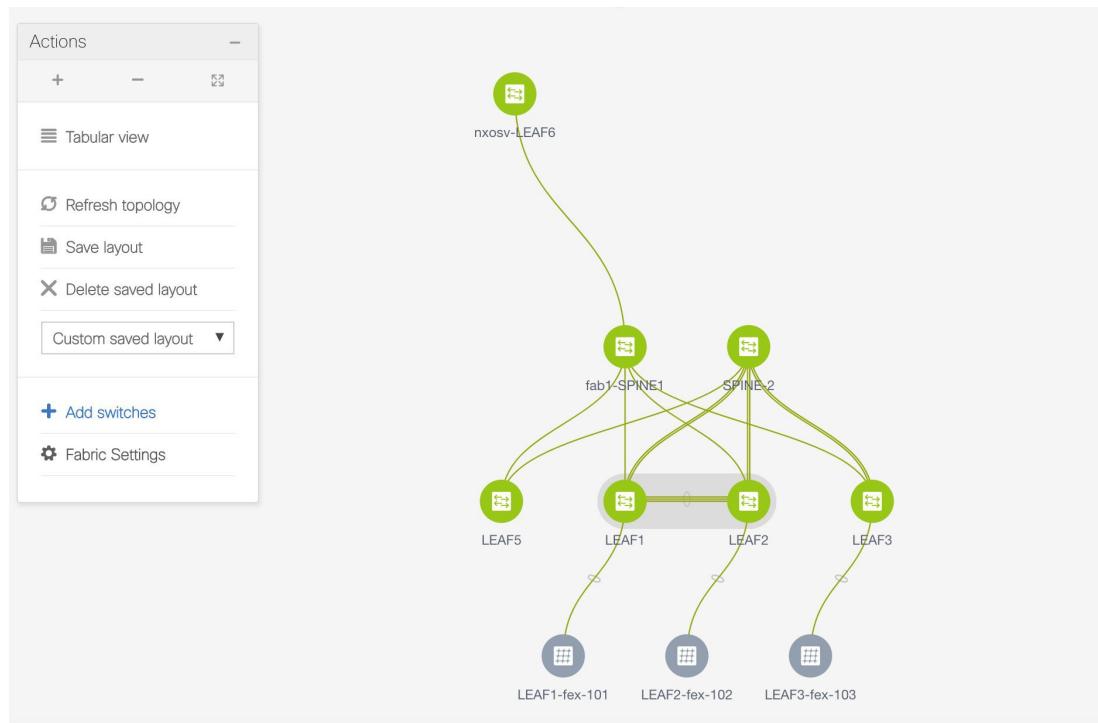
When a leaf switch boots up after a switch reload or RMA operation, DCNM provisions configurations for the switch and FEX devices connected to it. Occasionally, FEX connectivity comes up after DCNM provisions FEX (host interface) configurations, resulting in a configuration mismatch. To resolve the mismatch,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

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

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

## Discovering New Switches

1. Power on the new switch after ensuring that it is cabled to the DCNM server. Boot the Cisco NX-OS and setup switch credentials.
2. Execute the **write erase** and **reload** commands on the switch.  
Click **Yes** to both the CLI commands that prompt you to choose **Yes** or **No**.
3. Set the boot variable to the image that you want to POAP. DCNM uses this image to POAP. Also, DCNM injects an information script into the switch to collect the device onboarding information.
4. In the DCNM GUI, go to the *Standalone* fabric (Click **Control > Fabric Builder** and click the fabric *Standalone*). The fabric topology is displayed.



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

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

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

## Add Switch Instances to the Fabric



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

### Inventory Management

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

Close

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

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

7. Click **Refresh Topology** to get updated information. The added switch goes through the POAP cycle. Monitor and check the switch for POAP completion.
8. After the added switch completes POAP, the fabric builder topology page is refreshed with the added switch with some physical connections. However, the switch icon is in red color indicating that the fabric is *Out-Of-Sync* and you must click **Save & Deploy** on the fabric builder topology to deploy pending configurations (such as template and interface configurations) onto the switches.



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

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

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

10. Click **Close** to return to the fabric builder topology.
11. Click **Refresh Topology** to view the update. All switches must be in green color indicating that they are functional.
12. The switch and the link are discovered in DCNM. Configurations are built based on various policies (such as fabric, topology, and switch generated policies). The switch image (and other required) configurations are enabled on the switch.
13. In the DCNM GUI, the discovered switches can be seen in the *Standalone* fabric topology. Up to this step, the POAP is completed with basic settings. All the interfaces are set to trunk ports. You must setup interfaces through the **Control > Interfaces** option for any additional configurations, but not limited to the following:
  - vPC pairing.
  - Breakout interfaces.
  - Port channels, and adding members to ports.

**Note**

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

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

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

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

**Note**

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

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

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

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

## Return Material Authorization (RMA)

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

### Prerequisites

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

### Guidelines and Limitations

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

### POAP RMA Flow

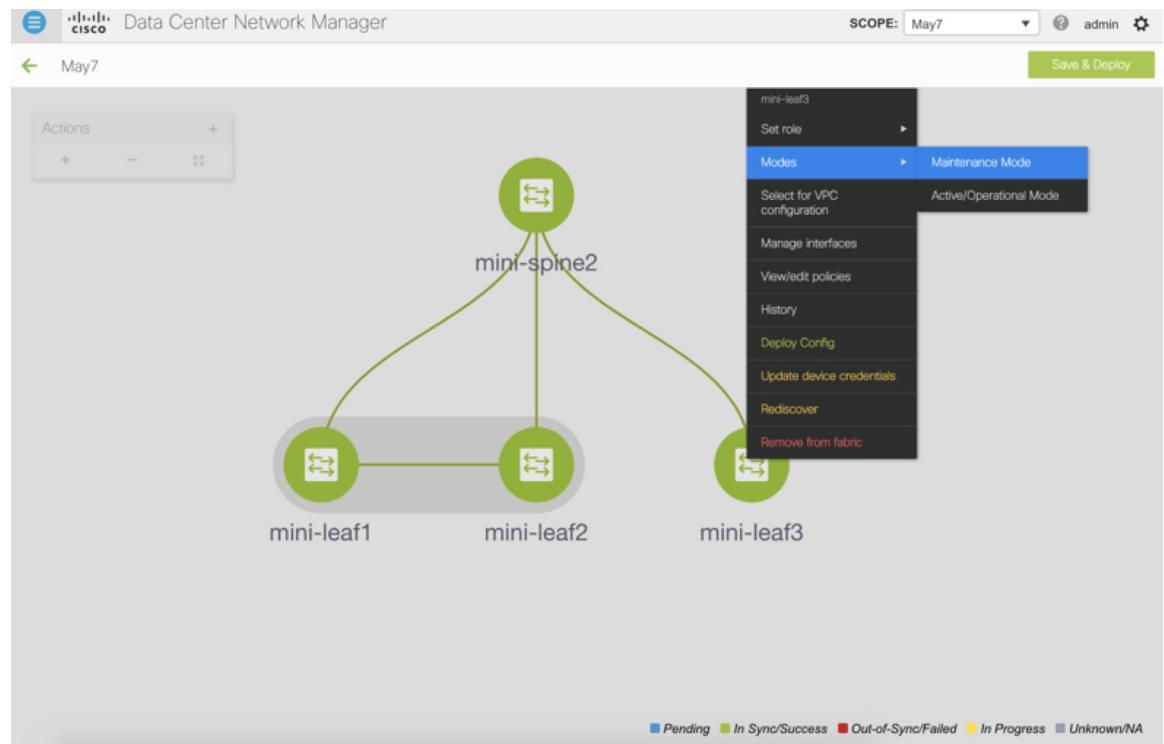
#### Procedure

---

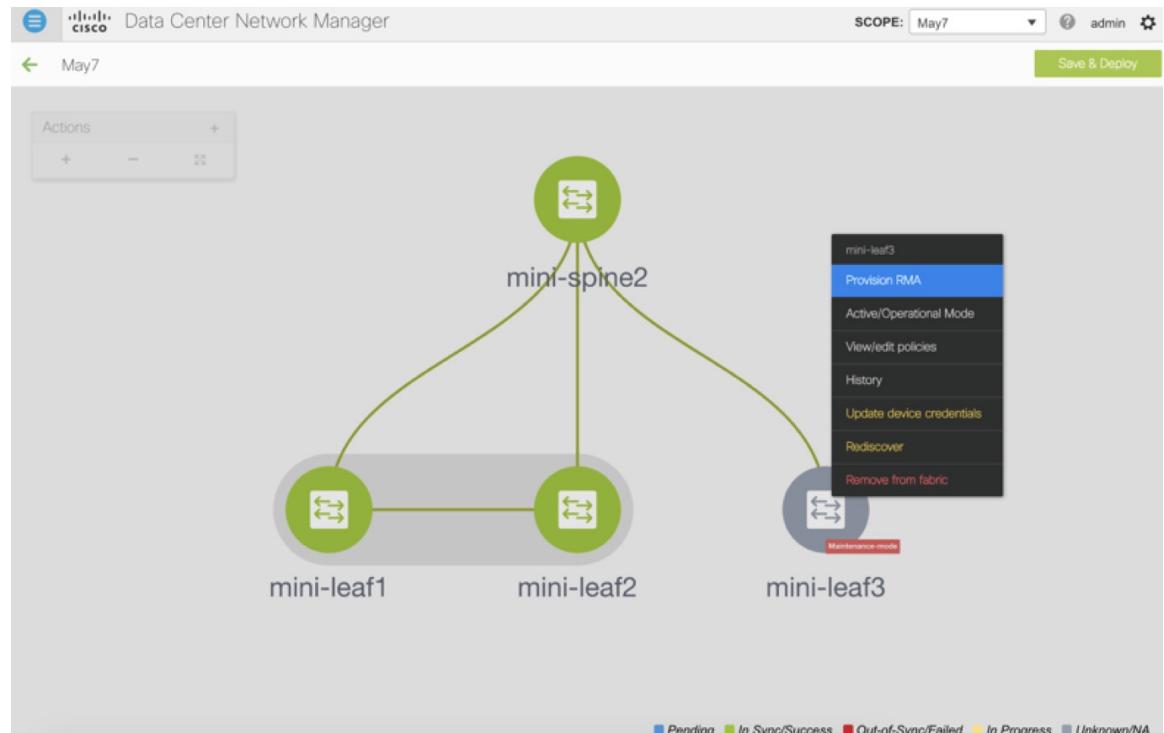
**Step 1** Choose **Control > Fabric Builder**.

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

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

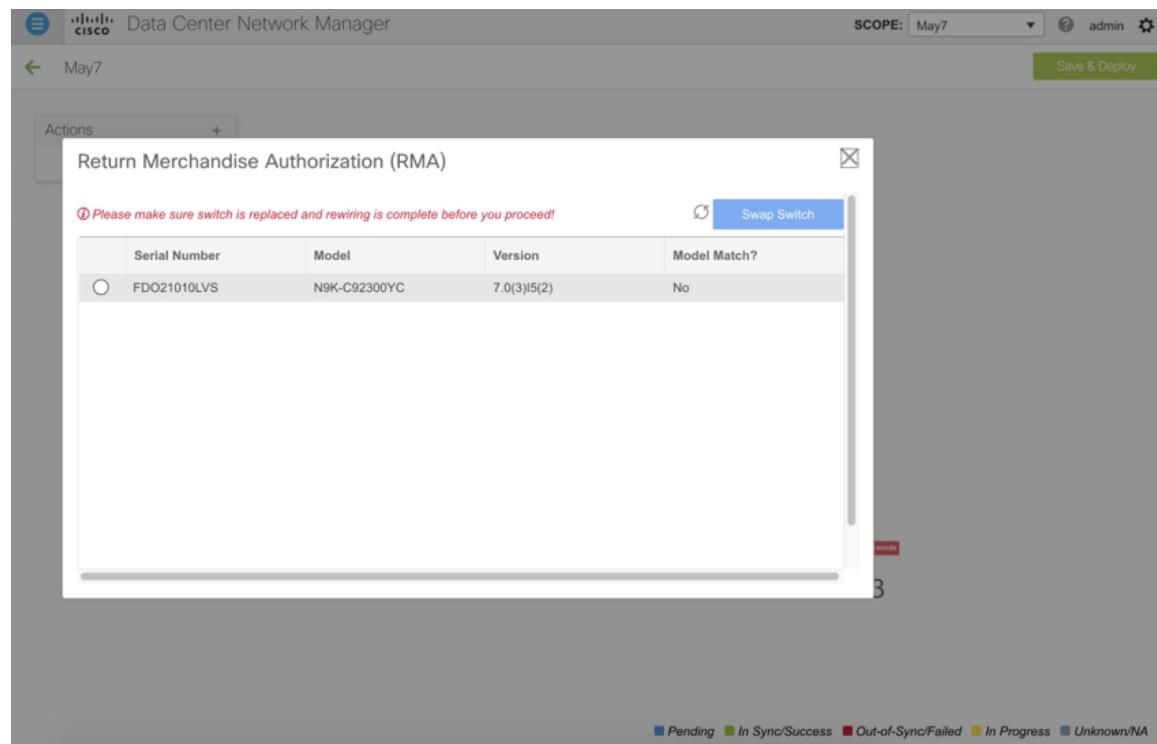


- Step 4** Physically replace the device in the network. Physical connections should be made in the same place on the replacement switch as they existed on the original switch.
- Step 5** Provision RMA flow and select the replacement device.

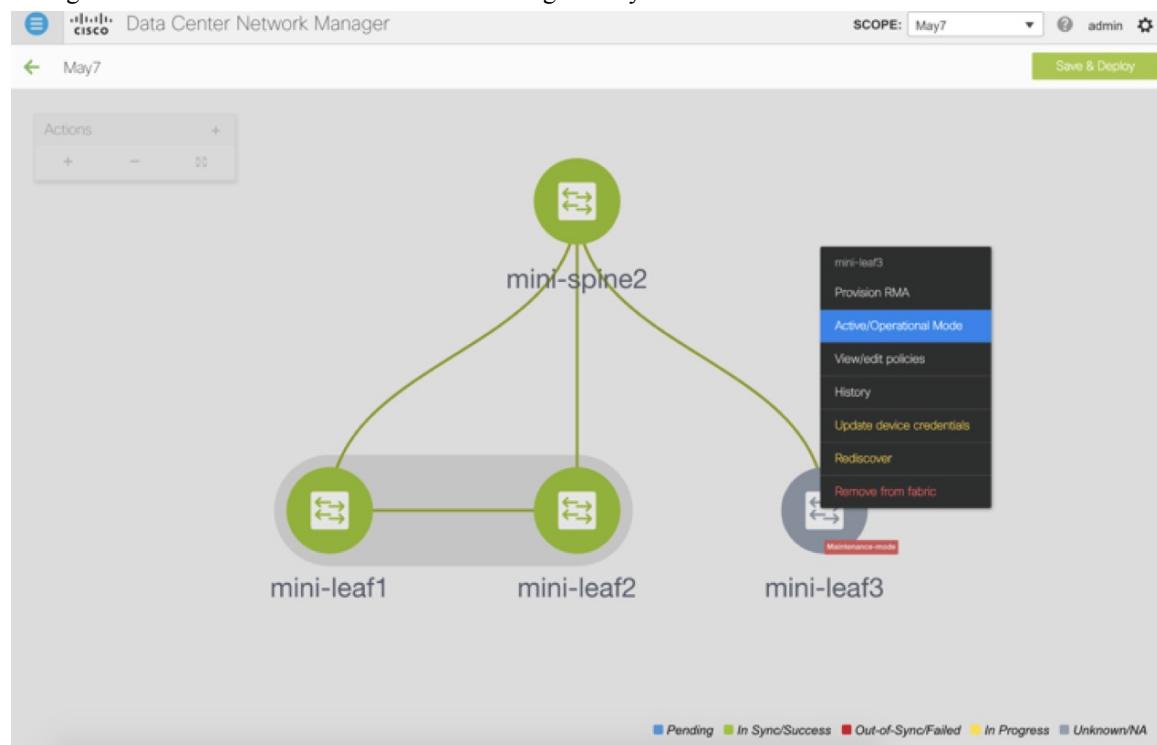


- Step 6** The Provision RMA UI will show the replacement device 5-10 minutes after it is powered on.

## POAP RMA Flow



- 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.

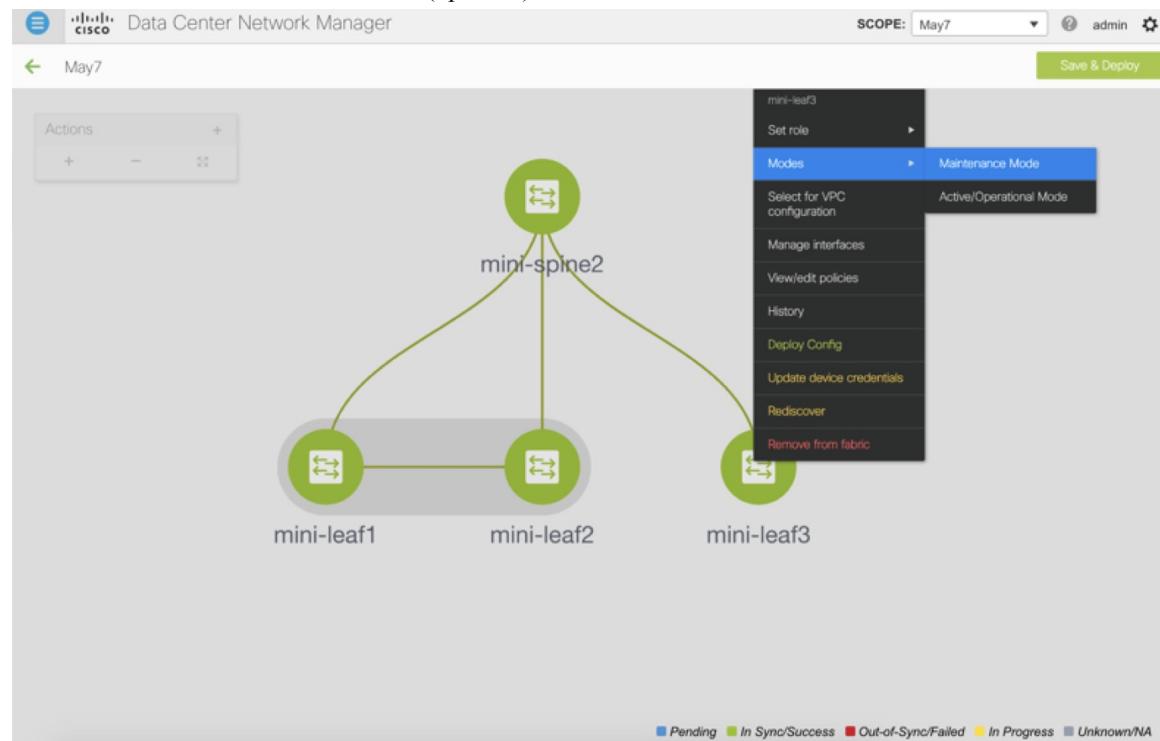


## Manual RMA Flow

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**.

## Manual RMA Flow

The screenshot shows the 'Config Deployment' window in Cisco DCNM. The table lists four switches:

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

At the bottom is a 'Deploy Config' button.

**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.

The screenshot shows the fabric topology in Cisco DCNM. The switches are represented as nodes:

- mini-spine2 (top, green)
- mini-leaf1 (bottom-left, grey)
- mini-leaf2 (bottom-right, grey)
- mini-leaf3 (right, blue)

A context menu is open over mini-leaf3, listing the following options:

- Provision RMA
- Active/Operational Mode (highlighted in blue)
- View/edit policies
- History
- Update device credentials
- Rediscover
- Remove from fabric

At the bottom is a legend for status colors: Pending (blue), In Sync/Success (green), Out-of-Sync/Failed (red), In Progress (yellow), Unknown/NA (grey).

## RMA for User with Local Authentication



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

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

### Procedure

- 
- Step 1** After the new switch comes online, SSH into the switch and reset the local user passwords with the cleartext password using the “username” command. Reset the local user passwords to resync the SNMP password. The password is stored in the configuration file in a nontransferable form.
- Step 2** Wait for the RMA to complete.
- Step 3** Update Cisco DCNM switch\_snmp\_user policy for the switch with the new SNMP MD5 key from the switch.
- 

## Interfaces

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

You can use the following functions:

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



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

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

- Blue: Pending
- Green: In Sync/Success

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

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



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

Field	Description
Add	Allows you to add a logical interface such as a port channel, vPC, Straight-through FEX, Active-Active FEX, loopback and subinterface.
Breakout, Unbreakout	Allows you to <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.
Rediscover	Allows you to rediscover or recalculate the compliance status on the selected interfaces.
Interface History	Allows you to display the interface deployment history details.

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

This section contains the following:

## Adding Interfaces

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

### Procedure

---

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

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

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

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

The **Add Interface** window appears.

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

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

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

You can create a vPC using the **vpc\_trunk\_host** policy. However, you cannot associate a VLAN to the vPC (as an allowed VLAN) from the **Interfaces** option. When you deploy an overlay network on the fabric switches, you should associate the corresponding VLAN to port channels (of the vPC domain), as applicable. Refer the *Networks Deployment in the Standalone Fabric* topic for network deployment details.

- When adding a subinterface, you must select a routed interface from the interface table before clicking the Add button.

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

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

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

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

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

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

You must not create a *\_upg* interface policy. For example, you should not create a policy using the **vpc\_trunk\_host\_upg**, **port\_channel\_aa\_fex\_upg**, **port\_channel\_trunk\_host\_upg**, and **trunk\_host\_upg** options.

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

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

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

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

The newly added interface appears in the screen.

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

## Editing Interfaces

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



**Note**

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

### Procedure

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

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

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

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

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

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

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

## Deleting Interfaces

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



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

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

### Procedure

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

**Step 2** Select the interfaces.

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

## Shutting Down and Bringing Up Interfaces

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

### Procedure

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

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

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

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

## Viewing Interface Configuration

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

### Procedure

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

Select the interface whose configurations you want to view.

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

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

---

## Rediscovering Interfaces

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

### Procedure

---

- Step 1** Choose **Control > Interfaces**.
- Step 2** Select the interfaces that you want to rediscover.
- Step 3** Click **Rediscover** to rediscover the selected interfaces. For example, after you edit or enable an interface, you can rediscover the interface.
- 

## Viewing Interface History

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

### Procedure

---

- Step 1** Choose **Control > Interfaces**.
- Step 2** Select the interface.
- Step 3** Click **Interface History** to view the configuration history on the interface.
- Step 4** Click **Status** to view each command that is configured for that configuration instance.
- 

## Deploying Interface Configurations

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

### Procedure

---

Choose **Deploy** to deploy and redeploy configurations that are saved for an interface.  
You can select multiple interfaces and deploy pending configurations.

---

# Networks and VRFs Creation and Deployment in a Standalone Fabric

The steps for overlay networks and VRFs provisioning are:

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



**Note**

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

The two steps are explained:

## Create Networks for the Fabric

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

## Select a Fabric

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

OR

Fabric Extension Setup

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

## Networks and VRFs Creation and Deployment in a Standalone Fabric

Fabric Selected: Standalone

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

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

### Create Network

**Network Information**

* Network ID	30000
* Network Name	MyNetwork_30000
* VRF Name	MyVRF_50000
* Layer 2 Only	<input type="checkbox"/>
* Network Template	Default_Network
* Network Extension Template	Default_Network_Extension
VLAN ID	

**Network Profile**

General	IPv4 Gateway/NetMask	example 192.0.2.1/24
Advanced	IPv6 Gateway/Prefix	example 2001:db8::1/64
	Interface Description	(?)

**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 (-). The corresponding Layer 3 VNI (or VRF VNI) is generated along with VRF creation.

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

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

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

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

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

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

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

#### General tab

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

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

Specify the anycast gateway IP address for transporting the L3 traffic from a server belonging to MyNetwork\_30000 and a server from another virtual network. By default the anycast gateway IP address is the same for MyNetwork\_30000 on all switches of the fabric that have the presence of the network.

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

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

- ARP Suppression
- Ingress Replication



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

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

A sample of the Create Network page:

## Networks and VRFs Creation and Deployment in a Standalone Fabric

Create Network

**Network Information**

* Network ID	30000
* Network Name	MyNetwork_30000
* VRF Name	MyVRF_50000
* Layer 2 Only	<input type="checkbox"/>
* Network Template	Default_Network
* Network Extension Template	Default_Network_Extension
VLAN ID	2400

**Network Profile**

General

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

Advanced

**Advanced Settings**

ARP Suppression	<input type="checkbox"/>	?
Ingress Replication	<input type="checkbox"/>	?
Multicast Group Address	239.1.1.8	?
* DHCPv4 Server	20.20.20.10	?
* DHCPv4 Server VRF	VRF_DHCP	?
MTU for L3 interface	[68-9216]	?

**Create Network**

- 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 Selection > Network Deployment >

Fabric Selected: Standalone

Selected 1 / Total 1

Networks

<input type="checkbox"/>	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
<input checked="" type="checkbox"/>	MyNetwork_30000	30000	MyVRF_50000	12.12.12.10/24		NA	2400

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

## Create VRFs for a Standalone Fabric

- 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.

The screenshot shows the VRFs page with the following details:

- Header: Fabric Selection > Network Selection > Network Deployment > Network View | Continue
- Title: Fabric Selected: Standalone
- Section: VRFs
- Buttons: +, Edit, Delete
- Table Headers: VRF Name, VRF ID, Status
- Table Data: MyVRF\_50000 (selected), VRF ID: 50000, Status: NA
- Filter: Show All
- Actions: Selected 1 / Total 1, Refresh, Settings

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

The screenshot shows the Create VRF dialog box with the following sections:

- VRF Information** (Collapsed):
 

* VRF ID	50001
* VRF Name	MyVRF_50001
* VRF Template	Default_VRF
* VRF Extension Template	Default_VRF_Extension
- VRF Profile** (Collapsed): A large empty area for profile configuration.
- Create VRF** (Blue button at the bottom right)

The fields in this screen are:

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



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

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

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

3. Click **Create VRF**.

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

Fabric Selected: Standalone

Selected 1 / Total 2 | Continue

VRFs

VRF Name	VRF ID	Status
MyVRF_50000	50000	NA
MyVRF_50001	50001	NA

### Networks Deployment in the Standalone Fabric

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

1. Go to the Select a Fabric page.

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

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

OR

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

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

Fabric Selected: Standalone

Selected 1 / Total 2 | Continue

Networks

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

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

**Note**

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

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

The screenshot shows the 'Networks' table with two rows selected. Column 1 is labeled 'Network Name' and contains 'MyNetwork\_30000' and 'MyNetwork\_30001'. Column 2 is 'Network ID' with values 30000 and 30001. Column 3 is 'VRF Name' with values 'MyVRF\_50000' and 'MyVRF\_50000'. Column 4 is 'IPv4 Gateway/Subnet' with values '12.12.12.10/24' and '12.12.10.20/24'. Column 5 is 'IPv6 Gateway/Prefix' and column 6 is 'Status', both of which are 'NA'. Column 7 is 'VLAN ID' with values 2400 and 2401. The 'Selected 2 / Total 2' status is shown at the top right. A yellow circle with the number '1' is on the first row, and another yellow circle with the number '2' is on the 'Continue' button.

Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
<input type="checkbox"/> MyNetwork_30000	30000	MyVRF_50000	12.12.12.10/24		NA	2400
<input checked="" type="checkbox"/> MyNetwork_30001	30001	MyVRF_50000	12.12.10.20/24		NA	2401

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

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

The screenshot shows the 'Network Deployment' page. At the top, it displays 'Fabric Name: Standalone' and 'Network(s) Selected'. Below this, there are two circular icons representing devices: 'leaf-91' and 'leaf-84', each with a double-headed arrow symbol inside. At the bottom left is a link to 'Device Selection Options', and at the bottom right are three colored status indicators: blue for 'Pending', green for 'In Sync/Success', and red for 'Out-of-Sync'.

- Fabric Name: Standalone
- Network(s) Selected

(i) [Device Selection Options](#)

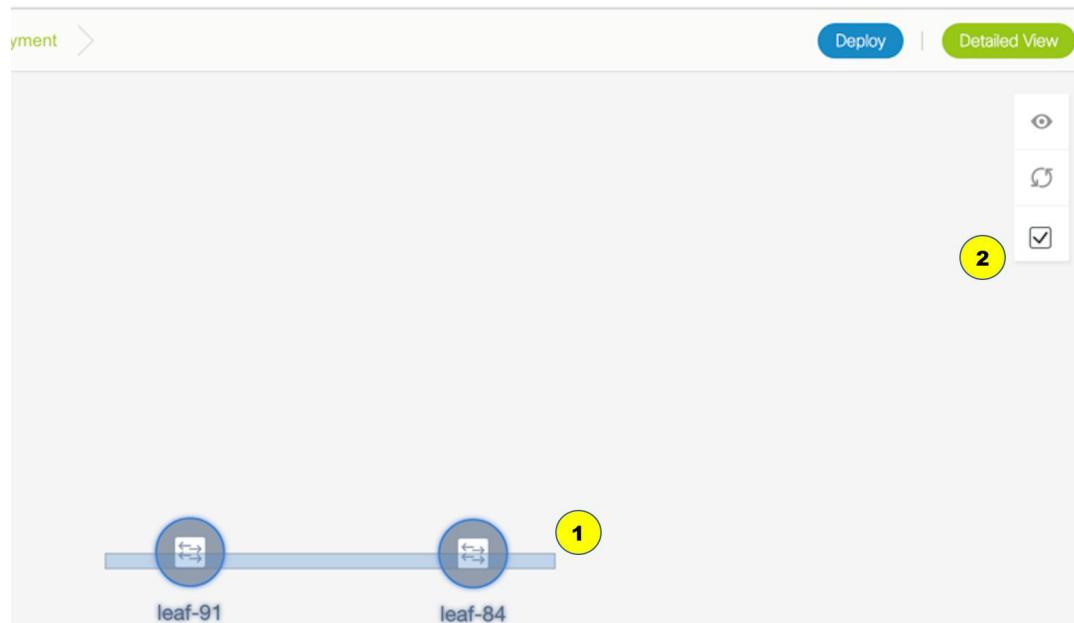
Pending In Sync/Success Out-of-Sync

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

The overlay networks (/VRFs) provisioning status is context-specific. It is a combination of networks that you chose for provisioning and the relevant switches in the topology. In this example, it means that the networks *MyNetwork\_30000* and *MyNetwork\_30001* are yet to be deployed on any switch in this fabric.

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

4. Double-click a switch (or use the *Multi-Select* option) to deploy the networks on it. For deployment of networks on multiple switches (like in this case, deploying *MyNetwork\_30000* and *MyNetwork\_30001* on leaf switches leaf-84 and leaf-91), do the following:
  1. Click *Multi-Select* from the panel at the top right part of the screen. The topology freezes to a static state.
  2. Drag the cursor across the switches.



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

## Switches Deploy

**Fabric Name:** *Standalone*

	Switch	VLAN	Interfaces	Status
<input type="checkbox"/>	leaf-84	2400	...	NA
<input type="checkbox"/>	leaf-91	2400	...	NA

A tab represents each network (the first network, *MyNetwork\_30000*, is displayed by default) that is being deployed. In each network tab, the switches are displayed. Each row represents a switch.

Click the checkbox next to the **Switch** column to select the switches. Both the switch check boxes are selected automatically. The network *MyNetwork\_30000* is ready to be provisioned on the switches leaf-84 and leaf-91.

Select the other network tab and make the same selections.

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



**Note**

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

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

Preview the configurations by clicking *Preview* (the eye icon above the Multi-Select option). Since *MyNetwork\_30000* and *MyNetwork\_30001* are networks of VRF 50000, the configurations contain VRF configurations followed by the network configurations.

## Preview Configuration



Select a Switch:

leaf-84

Select a Network

MyNetwork\_30000

## Generated Configuration:

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

VRF  
configurations

## Preview Configuration

Select a Switch:

leaf-84

Select a Network

MyNetwork\_30000

### Generated Configuration:

```
router bgp 65002
vrf myvrf_50000
address-family ipv4 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redist-subnet
maximum-paths ibgp 2
address-family ipv6 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redist-subnet
maximum-paths ibgp 2
interface nve1
member vni 50000 associate-vrf
configure terminal
apply profile MyVRF_50000
```

```
configure profile MyNetwork_30000
vlan 2400
vn-segment 30000
interface vlan2400
description Ethernet 1/1
vrf member myvrf_50000
ip address 12.12.12.10/24 tag 12345
ip dhcp relay address 20.20.20.10 use-vrf vrf_dhcp
fabric forwarding mode anycast-gateway
no shutdown
interface nve1
member vni 30000
mcast-group 239.1.1.0
evpn
vni 30000 I2
rd auto
route-target import auto
route-target export auto
configure terminal
apply profile MyNetwork_30000
```

**MyNetwork\_30000 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 View* appears.

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



**Note**

When you select multiple networks on the *Topology View* screen and proceed to the deployment screen, the switch color reflects the status of the first network in the selected list of networks. In this example, the switch color turns green when *MyNetwork\_30000* is provisioned on the switch. Go to the Networks page to view the individual status for all networks.

## Networks and VRFs Creation and Deployment in a Standalone Fabric

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

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

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

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

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



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

### VRFs Deployment in the Standalone Fabric

- 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).

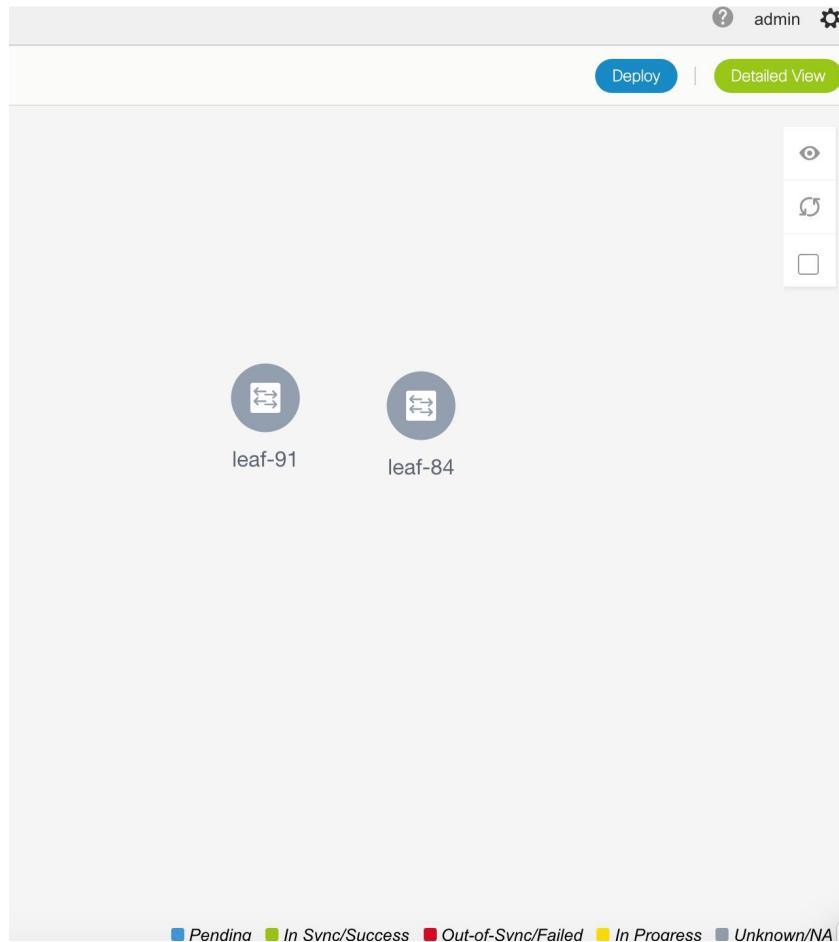
The VRFs page comes up. The list of VRFs created for the *Standalone* fabric are displayed in this screen.

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

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

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

The following example shows you how to deploy the *MyVRF\_50001* the VRF on the leaf switches leaf-84 and leaf-91. You can deploy VRFs simultaneously on multiple switches but of the same role (*Leaf*, *Border Gateway*, and so on).



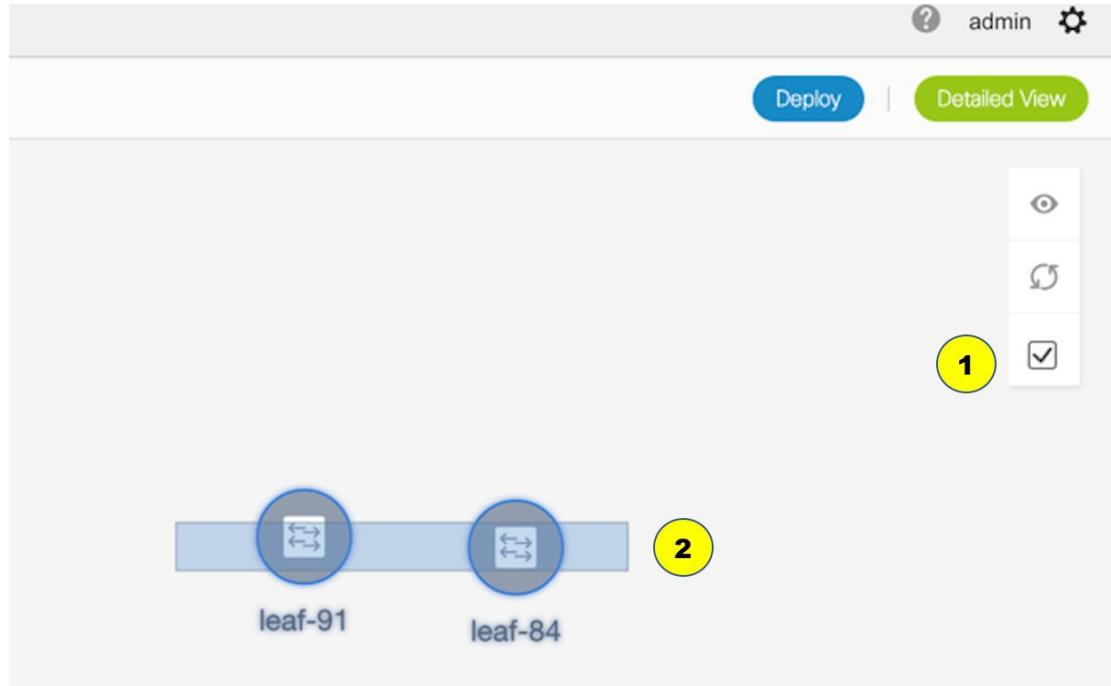
At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly (Blue for *Pending* state, yellow for *In Progress* state when the provisioning is in progress, red for failure state, green when successfully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

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

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

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

2. Drag the cursor across the switches.



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

## Switches Deploy

*Fabric Name:* **Standalone**

MyVRF\_50001

*Deploy Options:*

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

<input type="checkbox"/>	Switch	VLAN	Status	
<input type="checkbox"/>	leaf-84	2001	NA	
<input type="checkbox"/>	leaf-91	2001	NA	

Save

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

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

Select the other VRF tab and make the same selections.

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

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

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

## Preview Configuration

### Select a Switch:

leaf-84

### Select a VRF

MyVRF\_50001

### Generated Configuration:

```

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

```

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

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

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

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

The screenshot shows the 'Fabric Selection' > 'Network Selection' > 'Network Deployment' path. The main area is titled 'Fabric Name: Standalone VRF(s) Selected'. It features a table with columns: Name, Switch, Ports, and Status. There are three rows: 'MyVRF\_50001' (leaf-84, OUT-OF-SYNC), 'MyVRF\_50001' (leaf-91, OUT-OF-SYNC), and another row for 'MyVRF\_50001' which is currently selected. Buttons for 'Deploy', 'Preview', and 'History' are at the top of the table. A 'Topology View' button is in the top right corner. A 'Selected 0 / Total 2' indicator is also present.

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

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



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

### Undeploying Networks

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

1. Choose **Control > Networks and VRFs**.
2. In the **Select a Fabric** page, click **Continue** (at the top right part of the screen). The Networks page comes up.
3. Select the networks that you want to undeploy and click **Continue**. The *Topology View* page comes up.
4. On the **Topology View** page, select the **Multi-Select** button if you are undeploying the networks from multiple switches. The *Switches Deploy* screen comes up.  
(For a single switch, double-click the switch and the *Switches Deploy* screen comes up).
5. In the *Switches Deploy* screen, the **Status** column for the deployed networks is displayed as DEPLOYED. Unselect the check boxes next to the switches, as needed. Ensure that you repeat this on all tabs since each tab represents a network.
6. Click **Save** (at the bottom right part of the screen) to initiate the undeployment of the networks. The *Topology View* comes up again.



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

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

## Undeploying VRFs

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

1. Choose **Control > Networks and VRFs**.
2. In the Select a Fabric page, click **Continue** (at the top right part of the screen). The Networks page comes up.
3. Click the **VRF View** button (at the top right part of the screen) to go to the VRFs screen.
4. Select the VRFs that you want to undeploy and click **Continue**. The *Topology View* page comes up.
5. On the *Topology View* page, select the *Multi-Select* option if you are undeploying the VRFs from multiple switches. The Switches Deploy screen comes up.  
(For a single switch, double-click the switch and the Switches Deploy screen comes up).
6. In the Switches Deploy screen, the **Status** column for the deployed VRFs is displayed as DEPLOYED. Unselect the check boxes next to the switches, as needed. Ensure that you repeat this on all tabs since each tab represents a VRF.
7. Click **Save** (at the bottom right part of the screen) to initiate the undeployment of the VRFs. The *Topology View* comes up again.

**Note**

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

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

## Deleting Networks and VRFs in the MSD Fabric

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

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

## Creating an External Fabric

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

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

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

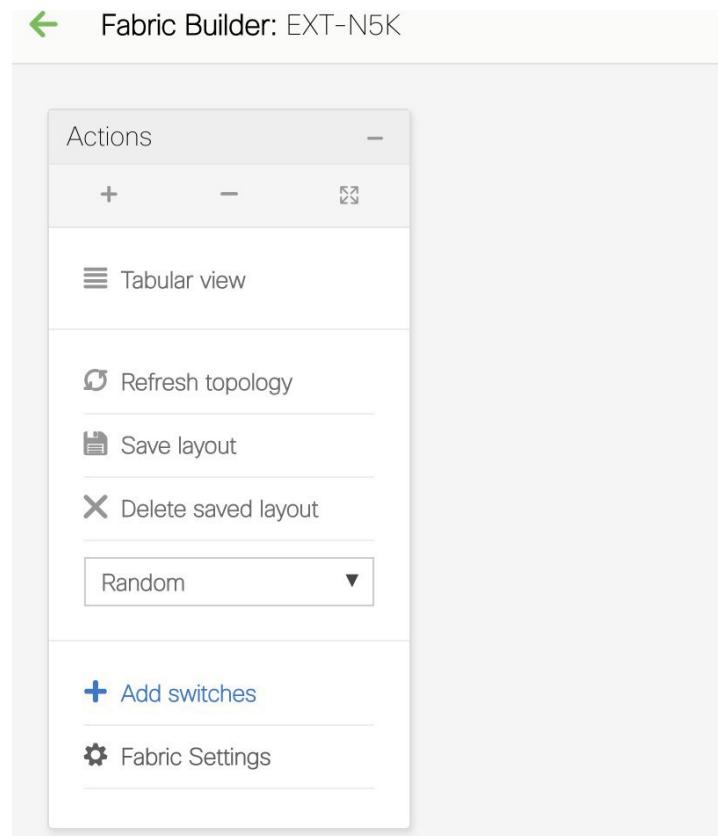
Follow these steps to create an external fabric from Fabric Builder.

1. Click **Control > Fabric Builder**. The Fabric Builder page comes up.
  2. Click the **Create Fabric** button. The Add Fabric screen comes up. The fields in this screen are:
    - Fabric Name** - Enter the name of the external fabric.
    - Fabric Template** - Choose *External\_Fabric*.
- When you choose the fabric template, the fabric creation screen for creating an external fabric comes up.
3. Enter the BGP AS number and click **Save**.

The screenshot shows the 'Add Fabric' configuration dialog box. At the top, it says 'Add Fabric' and has a close button 'X'. Below that, there are two input fields: one for 'Fabric Name' containing 'EXT-N5K' and another for 'Fabric Template' containing 'External\_Fabric'. A dropdown arrow is next to the template field. Below these, there are two tabs: 'General' (which is selected) and another tab that is partially visible. At the bottom of the dialog, there is a field for 'BGP AS #' with the value '555'. To the right of this field is a help icon and the text '1-4294967295 | 1-65535[0-65535]'. At the very bottom right are two buttons: 'Save' (in blue) and 'Cancel'.

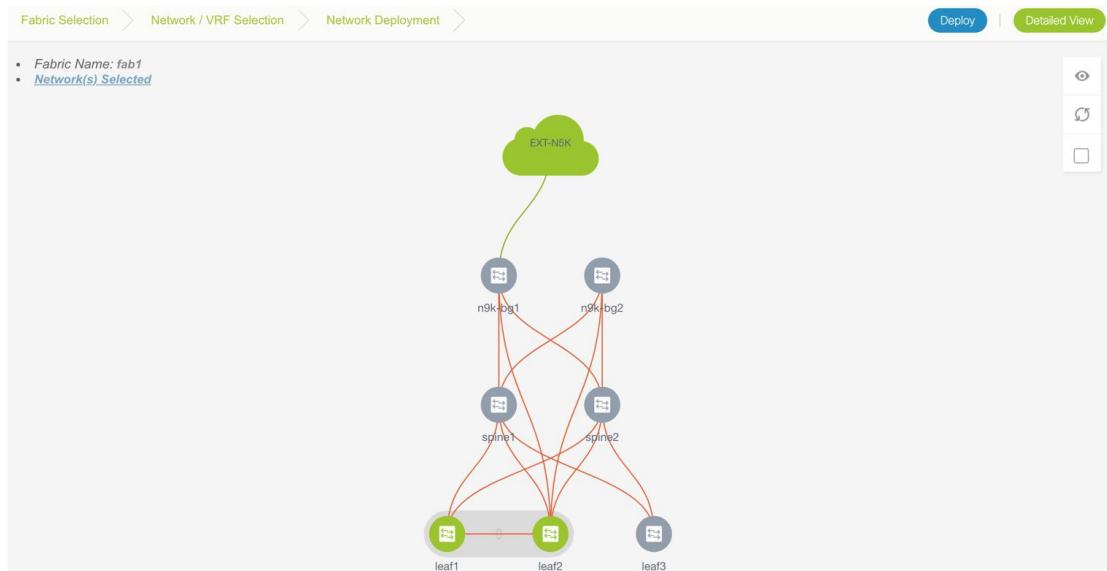
When you create an Inter-Fabric Connection from a VXLAN fabric to this external fabric, the BGP AS number is referenced as the external or neighbor fabric AS Number.

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



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

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



### Adding Fabric Extensions

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

1. Click **Control > Networks & VRFs** (under Fabrics submenu). The LAN Fabric Provisioning page comes up.
2. Click **Continue**. The Select a Fabric page is displayed. From the **Select a Fabric** drop down box, select the source fabric from which you want to connect to the other fabric.
3. Click **Fabric Extension Setup**.

## Select a Fabric

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

Easy7200
1
▼

OR

2
  
Fabric Extension Setup

The Fabric Extension screen comes up.

Type	Source Fabric	Source Device	Source Interface	Destination Fa...	Destination De...	Destination Int...	Configuration	Status
<input type="radio"/> MULTISITE_OVERLAY	Easy7200	N9K-3	Loopback100	Easy6000	N9K-15	Ethernet1/1	<a href="#">View Config</a>	DEPLOYED
<input type="radio"/> MULTISITE_UNDERLAY	Easy7200	N9K-3	Ethernet1/48	External	n7k1-BorderLeaf1	Ethernet7/1/4	<a href="#">View Config</a>	DEPLOYED
<input type="radio"/> VRF_LITE	Easy7200	N9K-4	Ethernet1/47	External	n7k1-BorderLeaf1	Ethernet7/4/1	<a href="#">View Config</a>	DEPLOYED

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

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

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



#### Note

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

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

## Add Inter-Fabric Connections

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

• • •

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

(i) VRF\_LITE: Set switch role - Border; MULTISITE: Set switch role - External

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

## Add Inter-Fabric Connections

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

• • •

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

(i) VRF\_LITE: Set switch role - Border; MULTISITE: Set switch role - Multisite

Previous    Next    Save & Deploy    Cancel

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

## Add Inter-Fabric Connections

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

• • •

---

▼ Network Profile

General

* IF_NAME	Ethernet1/2	?
* IP_MASK		?
* NEIGHBOR_IP		?
* NEIGHBOR ASN	65000	?
* Extension Type	VRF_LITE	?

Previous      Next      Save & Deploy      Cancel

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

## Add Inter-Fabric Connections

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

• • •

---

▼ Network Profile

General	<p>* IF_NAME <input type="text" value="Ethernet1/2"/> ?</p> <p>* IP_MASK <input type="text" value="10.2.3.4/24"/> ?</p> <p>* NEIGHBOR_IP <input type="text" value="10.2.3.10"/> ?</p> <p>* NEIGHBOR ASN <input type="text" value="65000"/> ?</p> <p>* Extension Type <input type="text" value="VRF_LITE"/> ?</p>
---------	--

**Previous** **Next** **Save & Deploy** **Cancel**

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

## Add Inter-Fabric Connections

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

Switch: N9K-3

Generated Configuration:

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

Previous Next Save & Deploy Cancel

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

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

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

Fabric Extension

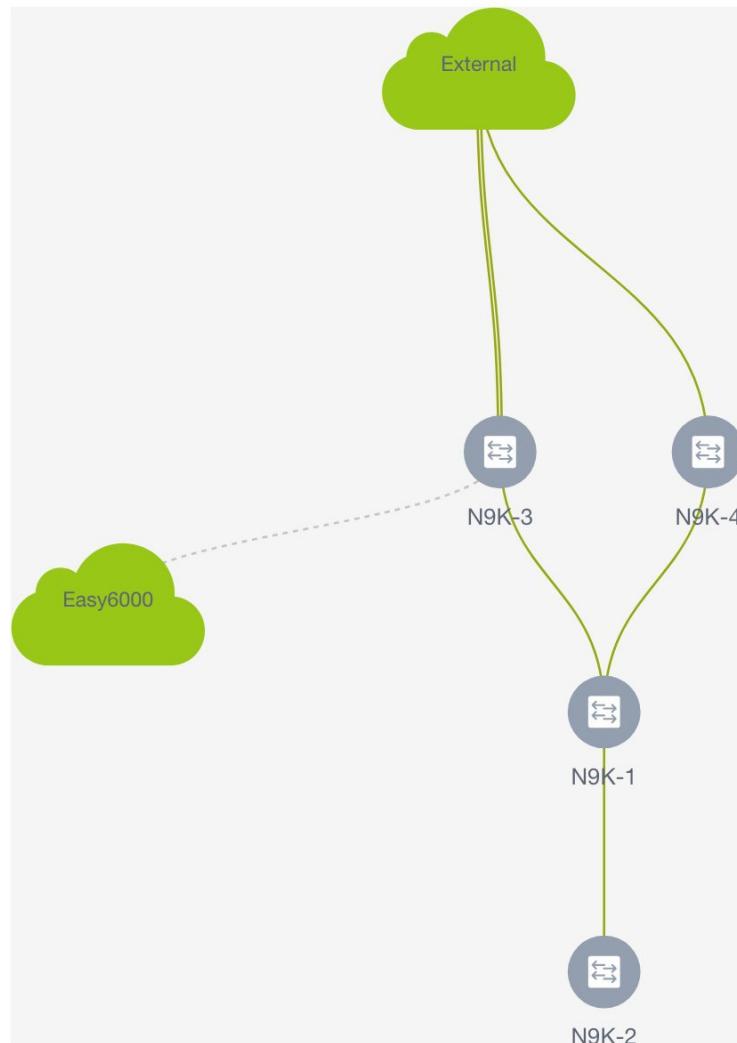
Inter-Fabric Connections

Selected 0 / Total 4

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

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

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



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

**Note**

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

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

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

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



### Note

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

### Guidelines and Limitations

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

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

- vPC switches – Ensure that the following configurations are present on vPC switches as is expected for a typical functioning vPC pair in a VXLAN BGP EVPN fabric.

- Secondary IP address on loopback1 (the loopback that is mapped to the NVE or VTEP interface).
- vPC peer link port channel and member interfaces.
- vPC peer link backup SVI and VLAN.

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

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

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

### Upgrade Procedure in the DCNM GUI

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

The screenshot shows the DCNM Fabric Builder interface. At the top, there is a blue circular icon with a gear and a star, followed by the text "Fabric Builder". Below this, a descriptive text reads: "Fabric Builder creates a managed and controlled SDN fabric. Select an existing fabric below or define a new VXLAN fabric, add switches using Power On Auto Provisioning (POAP), set the roles of the switches and deploy settings to devices." A blue button labeled "Create Fabric" is visible. Below this, a section titled "Fabrics (1)" displays a single fabric entry: "9K-FABRIC". The fabric details shown are: ASN: 2000, Replication Mode: Multicast, and Technology: VXLANFabric. There is also a small gear icon next to the fabric name.

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

Edit Fabric

\* Fabric Name : 9K-FABRIC

\* Fabric Template : Easy\_Fabric

General Advanced Resources Manageability Bootstrap

\* BGP ASN : 2000

\* Fabric Interface Numbering : p2p

\* Link-State Routing Protocol : ospf

\* Replication Mode : Multicast

\* Multicast Group Subnet : 239.1.1.0/25

\* Anycast Gateway MAC : 2020.0000.00aa

NX-OS Software Image Version

Save Cancel

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

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

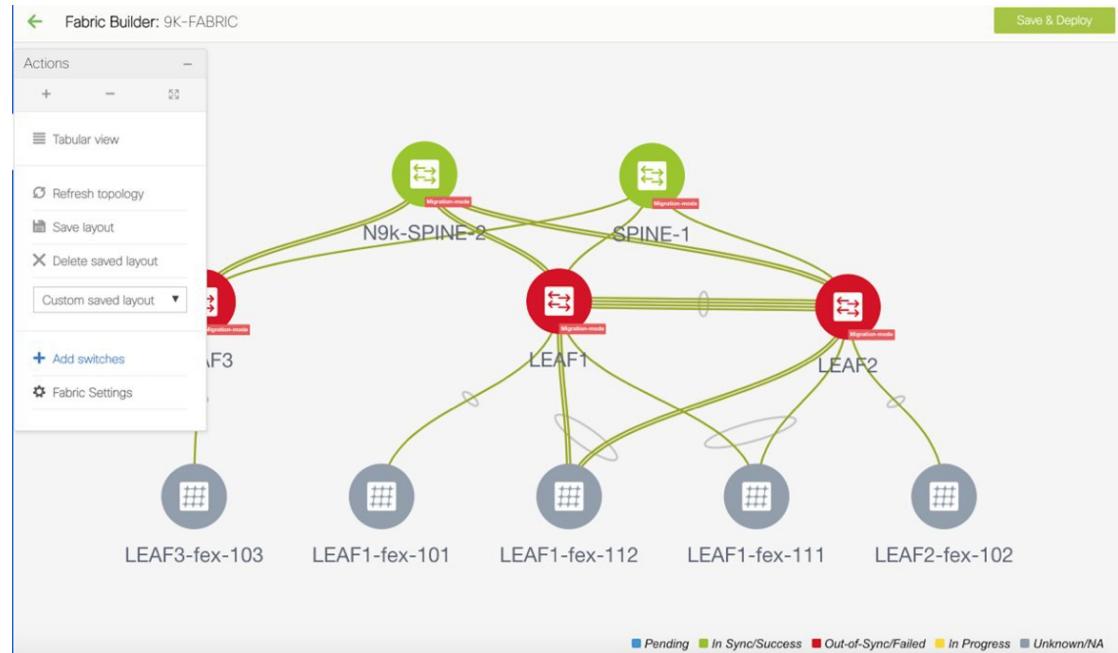
- BGP AS Number.
- Fabric underlay routing protocol (IS-IS or OSPF).
- Replication mode (Multicast or Ingress Replication).
- Fabric interface numbering (p2p or IP unnumbered).
- vPC peer link VLAN, if vPC is present.
- vPC delay restore time and other related parameters in the **Advanced** tab

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

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

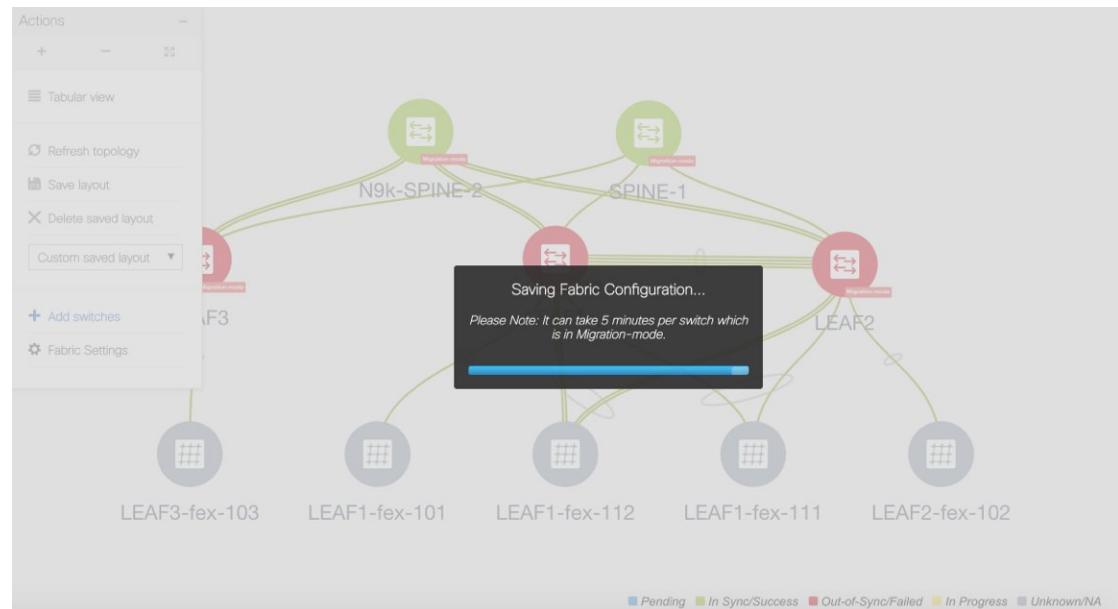
- Click **Save** to save the updated settings.

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



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

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



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

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

### Config Deployment

Switch Name	IP Address	Switch Serial	Preview Config	Status	Progress
N9k-SPINE-2	172.25.23.97	SAL2015NU0T	117 lines	Out-of-sync	<div style="width: 100%;">100%</div>
SPINE-1	172.25.23.81	FDO22062HL8	77 lines	Out-of-sync	<div style="width: 100%;">100%</div>
LEAF3	172.25.23.93	FDO20350MFK	50 lines	Out-of-sync	<div style="width: 100%;">100%</div>
LEAF2	172.25.23.92	FDO20350MHU	106 lines	Out-of-sync	<div style="width: 100%;">100%</div>
LEAF1	172.25.23.91	FDO20281K6K	106 lines	Out-of-sync	<div style="width: 100%;">100%</div>

[Deploy Config](#)

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

### Config Preview - Switch 172.25.23.97

Pending Config	Expected Config	Current Config
----------------	-----------------	----------------

```

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

```

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

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

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

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

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

### Config Deployment

The screenshot shows a modal dialog titled "Config Deployment". At the top, there are two tabs: "Step 1. Configuration Preview" and "Step 2. Configuration Deployment Status", with the latter being active. Below the tabs is a table with the following data:

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

At the bottom right of the dialog is a blue "Close" button.

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

### Config Deployment



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

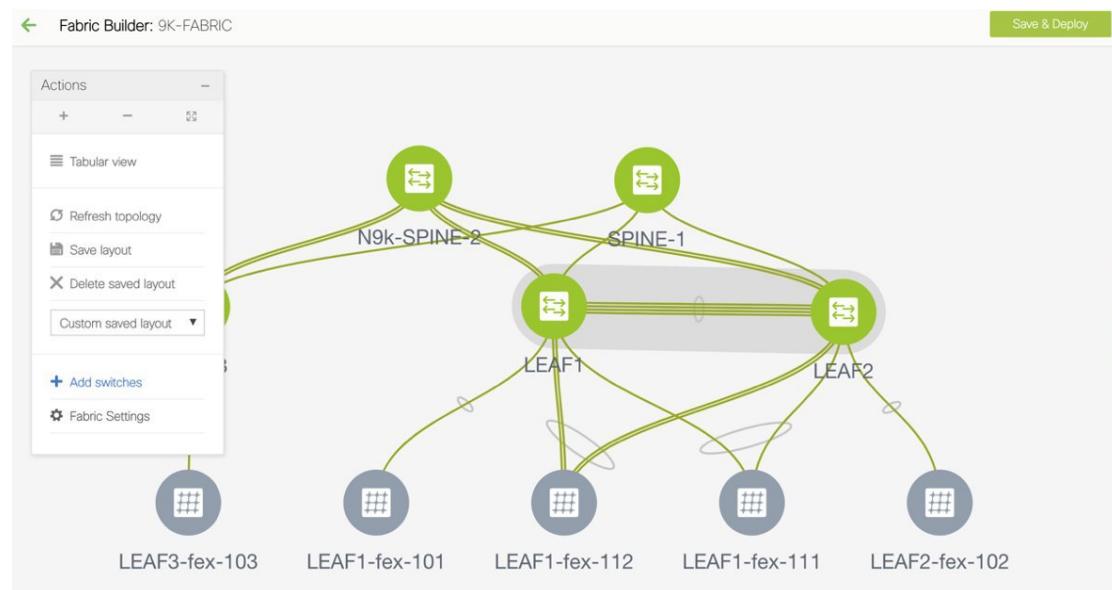
Switch Name	IP Address	Status	Status Description	Progress
LEAF3	172.25.23.93	COMPLETED	Deployed successfully	<div style="width: 100%;">100%</div>
SPINE-1	172.25.23.81	COMPLETED	Deployed successfully	<div style="width: 100%;">100%</div>
LEAF1	172.25.23.91	COMPLETED	Deployed successfully	<div style="width: 100%;">100%</div>
N9k-SPINE-2	172.25.23.97	COMPLETED	Deployed successfully	<div style="width: 100%;">100%</div>
LEAF2	172.25.23.92	COMPLETED	Deployed successfully	<div style="width: 100%;">100%</div>

Close

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

#### 9. Click Close.

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



## Multi-Site Domain for VXLAN BGP EVPN Fabrics

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

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



### Note

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

A few fabric-specific terms:

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

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

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

### Fabric and Switch Instance Variables

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

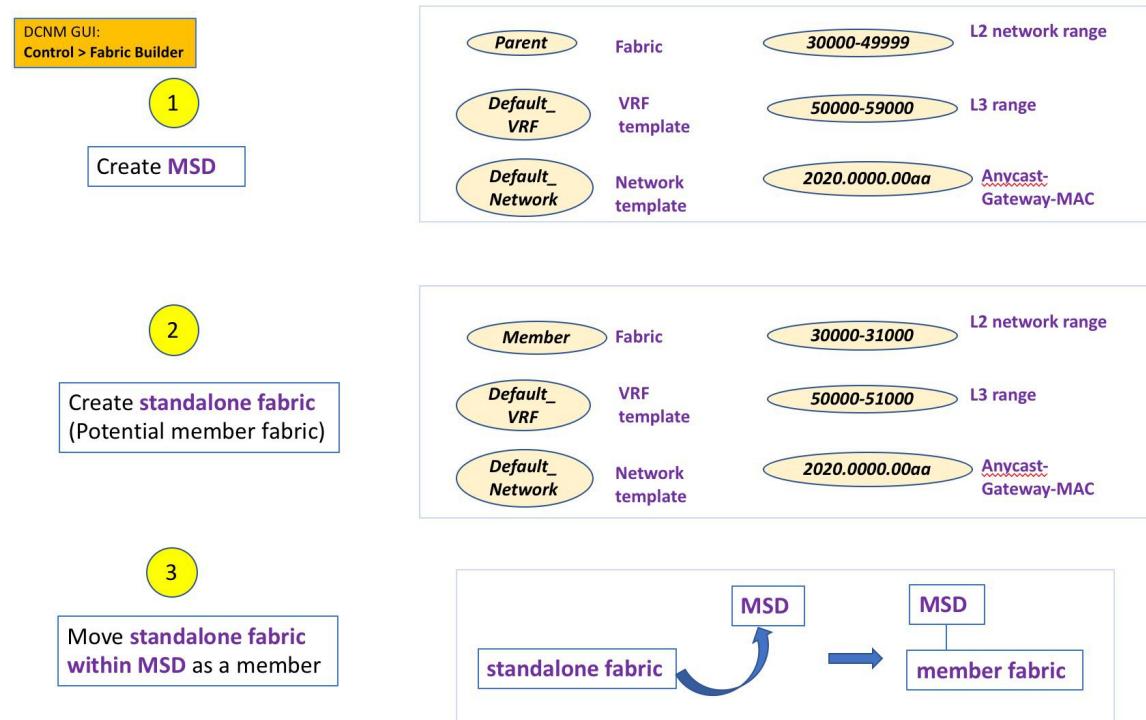
Fabric instance values can be edited in the fabric context. Specify fabric instance values for each fabric. For example, *multicast group subnet address*.

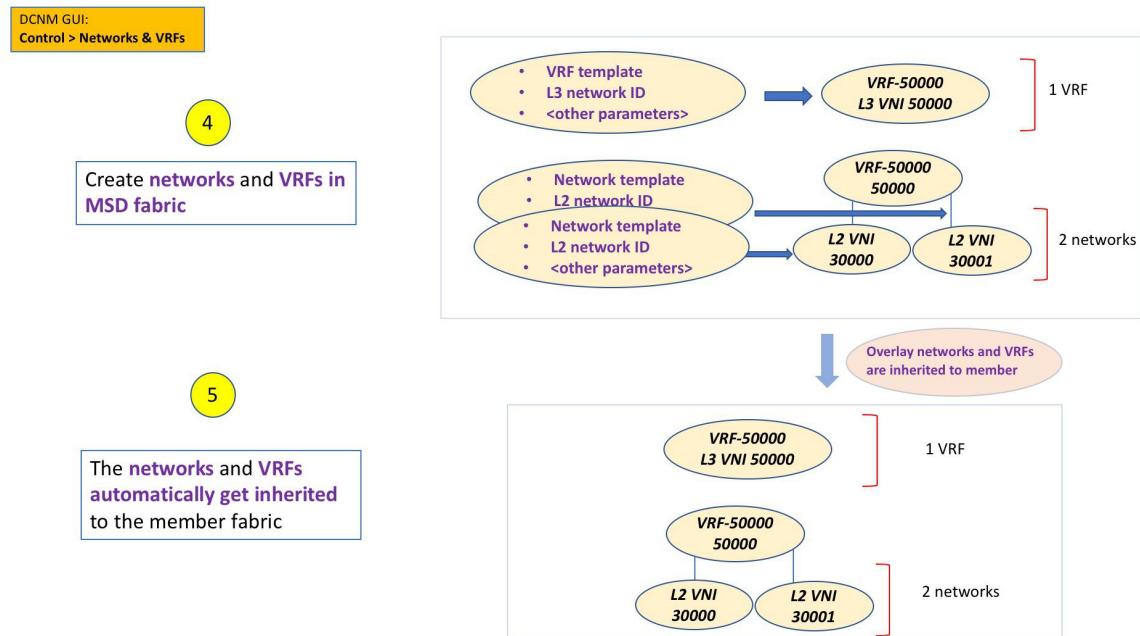
Switch instance values can be edited on deployment of the network on the switch. For example, *VLAN ID*.

### MSD and Member Fabric Process Flow

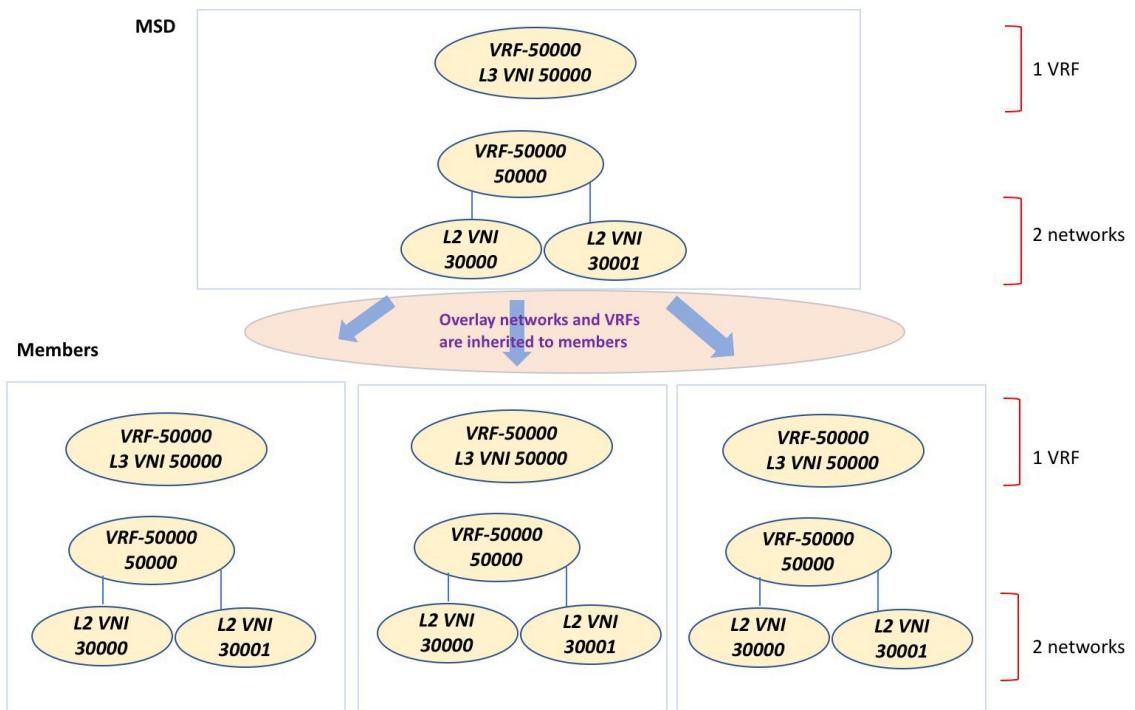
An MSD has multiple sites (and hence, multiple member fabrics under an MSD). VRFs and networks are created for the MSD and get inherited by the member fabrics. For example, VRF-50000 (and L3 network with ID 50000), and L2 networks with IDs 30000 and 30001 are created for the MSD, in one go.

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

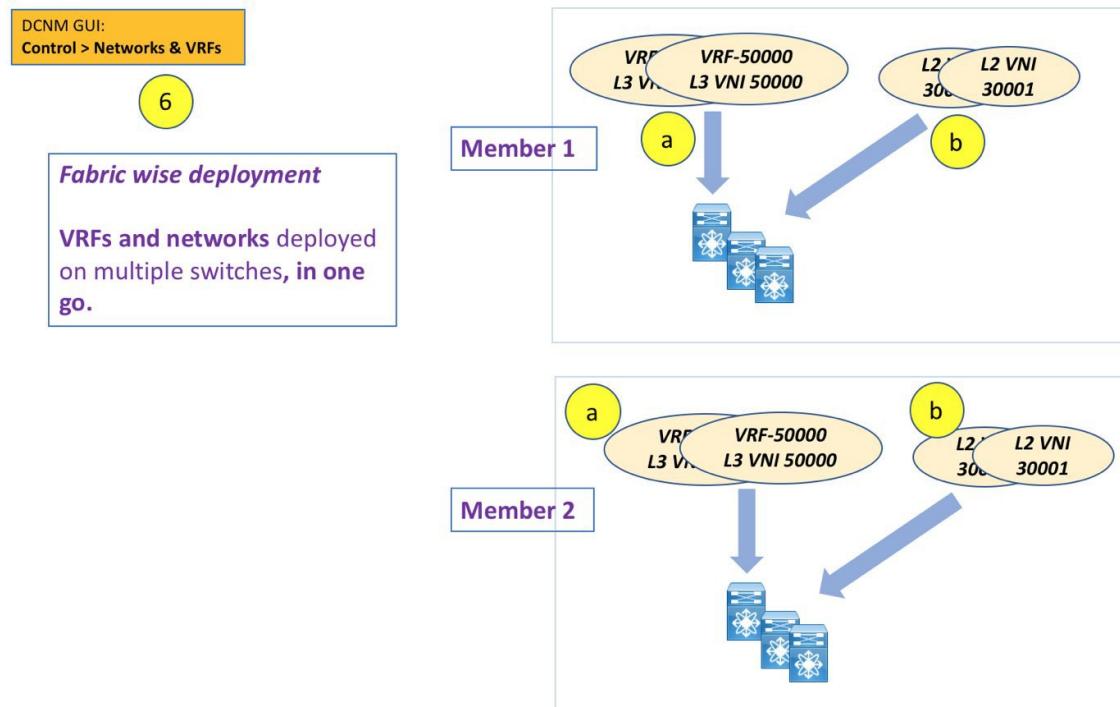




The sample flow explained the inheritance from the MSD to one member. An MSD has multiple sites (and hence, multiple member fabrics under an MSD). A sample flow from an MSD to multiple members:



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



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

Upcoming sections in the document explain the following:

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

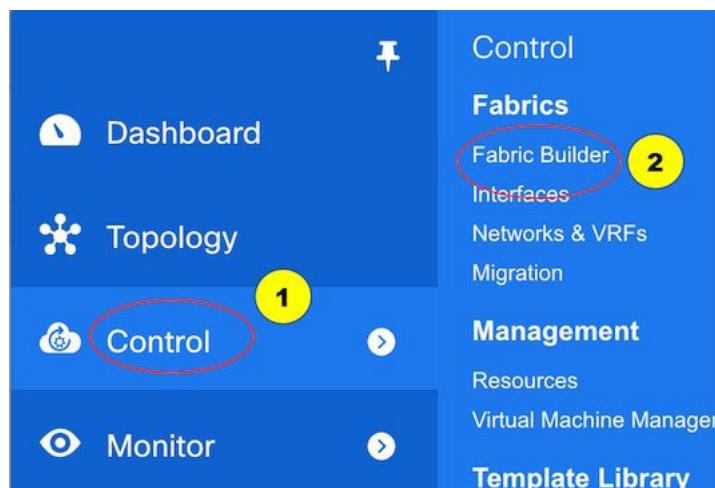
### Create an MSD Fabric and Associate Member Fabrics to It

The process is explained in two steps:

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

## Create an MSD Fabric

- Click Control > Fabric Builder.



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

Fabrics (4)			
Easy60000			Type: Switch_Fabric ASN: 60000 Replication Mode: Multicast Technology: VXLANFabric
Easy7200			Type: Switch_Fabric ASN: 7200 Replication Mode: Multicast Technology: VXLANFabric
External65000			Type: External ASN: 65000
MSD			Type: MSD Member Fabrics: Easy60000

A standalone or member fabric contains *Switch\_Fabric* in the **Type** field, its ASN 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.

- 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 X

\* Fabric Name :

\* Fabric Template :

**General**

L2 Segment ID Range	<input type="text" value="30000-49999"/>	<span>(?) L2 Segment ID Range</span>
L3 Partition ID Range	<input type="text" value="50000-59000"/>	<span>(?) L3 Partition ID Range</span>
* VRF Template	<input type="text" value="Default_VRF"/>	<span>(?) VRF Template</span>
* Default Network Template	<input type="text" value="Default_Network"/>	<span>(?) Network Template</span>
* VRF Extension Template	<input type="text" value="Default_VRF_Extension"/>	<span>(?) VRF Extension Template</span>
* Network Extension Template	<input type="text" value="Default_Network_Extension"/>	<span>(?) Network Extension Template</span>
Anycast-Gateway-MAC	<input type="text" value="2020.0000.00aa"/>	<span>(?) Shared MAC address for all leaves</span>

**Save** **Cancel**

The fields in the screen are explained:

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

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

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

**VRF Template** - Default VRF template.

**Default Network Template** - Default network template.

**VRF Extension Template** - Default VRF extension template.

**Network Extension Template** - Default network extension template.

**Anycast-Gateway-MAC** - Anycast gateway MAC address.

### 3. Click Save.

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

**Fabric Builder: MSD-Parent-Fabric**

**Actions**

- + -
- Refresh topology
- Save layout
- Delete saved layout
- Random
- Fabric Settings
- Move Fabrics

**Overlay Fabric Information**

- Total Networks: 0
- Total VRFs: 0

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

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



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

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

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

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

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

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

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

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

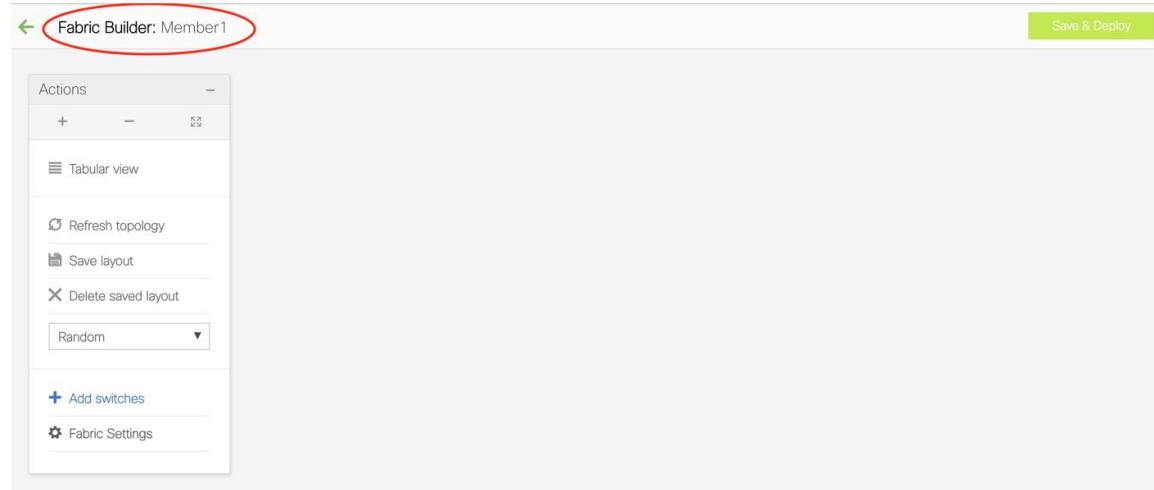
* Fabric Name :	<input type="text"/>
* Fabric Template	Easy_Fabric
<input type="button" value="General"/> <input type="button" value="Advanced"/> <input type="button" value="Resources"/> <input type="button" value="Manageability"/> <input type="button" value="Bootstrap"/>	
* Underlay Multicast Loopback IP Range <input type="text" value="10.254.254.0/24"/> <span>(?) Anycast Or Phantom RP IP Address Range</span>	
* Underlay Subnet IP Range <input type="text" value="10.3.0.0/16"/> <span>(?) Address range to assign P2P and Peer Link SVI</span>	
* Layer 2 VXLAN VNI Range <input type="text" value="30000-49000"/> <span>(?) Overlay Network Identifier Range (Min:1, Max:16777)</span>	
* Layer 3 VXLAN VNI Range <input type="text" value="50000-59000"/> <span>(?) Overlay VRF Identifier Range (Min:1, Max:16777)</span>	
* Network VLAN Range <input type="text" value="2300-2999"/> <span>(?) Per Switch Overlay Network VLAN Range (Min:2, Max:4095)</span>	
* VRF VLAN Range <input type="text" value="2000-2299"/> <span>(?) Per Switch Overlay VRF VLAN Range (Min:2, Max:4095)</span>	
* Subinterface Dot1q Range <input type="text" value="2-511"/> <span>(?) Per Border Dot1q Range For VRF Lite Connectivity</span>	
<input type="button" value="Save"/> <input type="button" value="Cancel"/>	

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

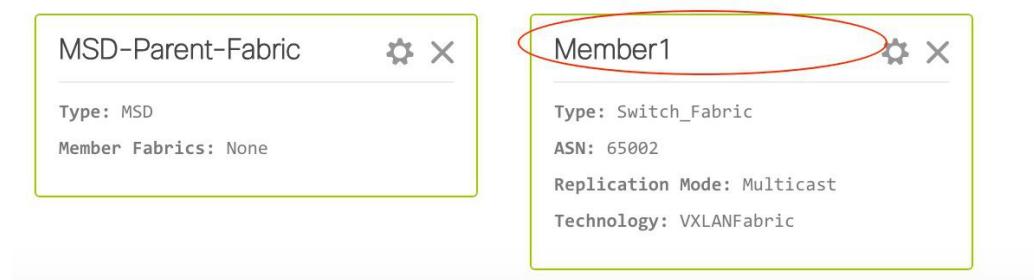
Other pointers:

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

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



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



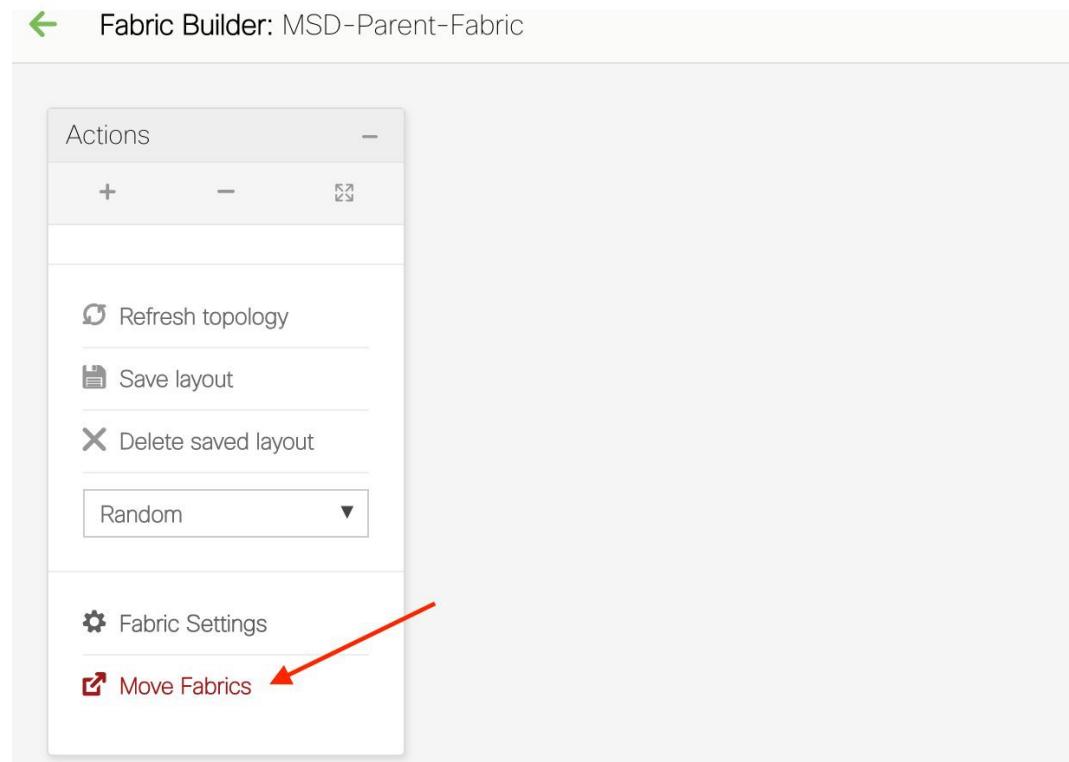
### Move the Member1 Fabric Under MSD-Parent-Fabrics

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

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

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

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



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

### 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 will not be displayed here.

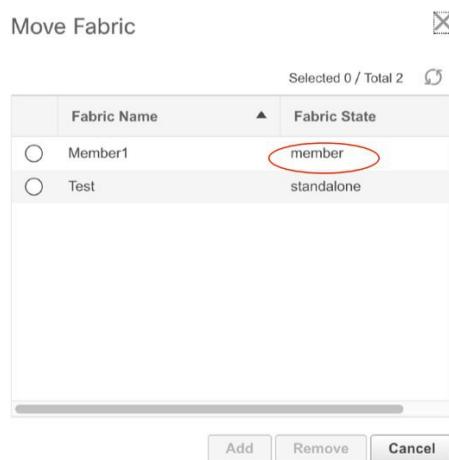
The *Member1* fabric is still a standalone fabric as seen in the image. A fabric is considered a member fabric of an MSD fabric only when you associate it with the MSD fabric. Also, each standalone fabric is a candidate for being an MSD fabric member, until you associate it to one of the MSD fabrics.

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

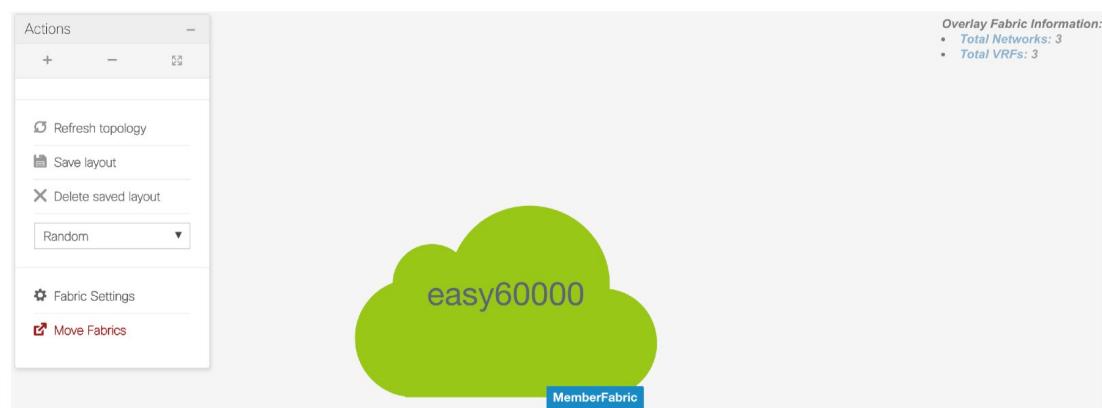
3. Click **Add**.

Immediately, a message appears at the top of the screen indicating that the *Member1* fabric is now associated with the MSD fabric *MSD-Parent-Fabric*. Now, the MSD-Parent-Fabric fabric page appears again.

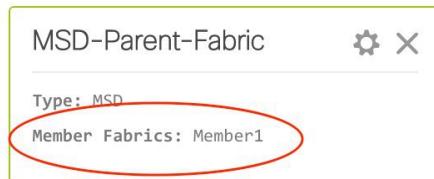
4. Click the **Move Fabrics** option to check the fabric status. You can see that the fabric status has changed from standalone to member.



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



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.

### Networks and VRFs Creation and Deployment in an MSD Fabric

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

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

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



#### Note

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

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

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

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

#### Create Networks in the MSD Fabric

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



## Select a Fabric

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

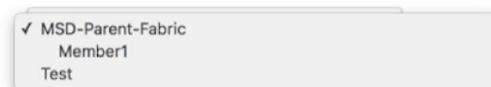
MSD-Parent-Fabric ▾

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



## Select a Fabric

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



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



## Select a Fabric

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

MSD-Parent-Fabric

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

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								

- 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	<input type="text" value="30000"/>
* Network Name	<input type="text" value="MyNetwork_30000"/>
* VRF Name	<input type="text" value="MyVRF_50000"/> <input type="button" value="+"/>
* Layer 2 Only	<input type="checkbox"/>
* Network Template	<input type="text" value="Default_Network"/>
Network Extension Template	<input type="text"/>
VLAN ID	<input type="text"/>

## ▼ Network Profile

<input checked="" type="radio"/> General
<input type="radio"/> Advanced

IPv4 Gateway/NetMask	<input type="text"/>	example 192.0.2.1/24
IPv6 Gateway/Prefix	<input type="text"/>	example 2001:db8::1/64
Interface Description	<input type="text"/>	<input type="button" value="?"/>

**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 will be blank. If you want to create a new VRF, click the + button. The VRF name should not contain any white spaces or special characters except underscore (\_), hyphen (-), and colon (:).

**Note**

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

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

**Network Template** - Allows you to select a network template.

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

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

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

**General tab**

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

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

**Interface Description** - Specifies the description for the interface.

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

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

A sample of the Create Network screen is given below.

Create Network

Network Information

- \* Network ID: 30000
- \* Network Name: MyNetwork\_30000
- \* VRF Name: MyVRF\_50000
- \* Layer 2 Only:
- \* Network Template: Default\_Network
- \* Network Extension Template: Default\_Network\_Extension
- VLAN ID: 2400

Network Profile

General	IPv4 Gateway/NetMask: 12.12.12.10/24	(example 192.0.2.1/24)
Advanced	IPv6 Gateway/Prefix:	(example 2001:db8::1/64)
	Interface Description: Interface vlan 2400	(?)

Create Network

**Advanced** tab:

▼ Network Profile

<input type="radio"/> General <input checked="" type="radio"/> Advanced	<b>ARP Suppression</b> <input type="checkbox"/> <b>* DHCPv4 Server</b> 20.20.20.10 <span style="float: right;">? <b>DHCP Relay IP</b></span> <b>* DHCPv4 Server VRF</b> VRF_DHCP <span style="float: right;">? [68-9216]</span> <b>MTU for L3 interface</b> <span style="float: right;">? [68-9216]</span>
--	---

**Create Network**

- 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

Selected 1 / Total 1 VRF View | Continue

Networks

Show All

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

### Editing and Deleting Networks in the MSD Fabric

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

## Edit Network

### ▼ Network Information

\* **Network ID**

\* **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

**Interface Description**

Save

### ▼ Network Profile

**General**

**Advanced**

**ARP Suppression**  ?

**DHCPv4 Server**  ? *DHCP Relay IP*

**DHCPv4 Server VRF**  ?

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

Save

Cancel

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

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

## Network Inheritance from MSD-Parent-Fabric to Member1

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

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

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

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

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

## Editing Networks in the Member Fabric

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

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

Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix
MyNetwork_30000	30000	MyVRF_50000	1.1.1.1/24	
<input checked="" type="checkbox"/> Network_30001	30001	MyVRF_50002	2.2.2.2/24	
MyNetwork_30002	30002	MyVRF_50001	3.3.3.3/24	

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

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

### ▼ Network Profile

General	ARP Suppression <input type="checkbox"/>	<a href="#">?</a>
Advanced	Ingress Replication <input type="checkbox"/>	<a href="#">?</a> Read-only per network, Fabric-wide setting
	Multicast Group Address	239.1.1.8 <a href="#">?</a>
	DHCPv4 Server	<a href="#">?</a> DHCP Relay IP
	DHCPv4 Server VRF	<a href="#">?</a>
	MTU for L3 interface	[68-9216] <a href="#">?</a>

### Create VRFs in the MSD Fabric

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

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

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

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

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

Fabric Selected: MSD-Parent-Fabric		
Selected 0 / Total 1 <a href="#">?</a> <a href="#">?</a> <a href="#">?</a>		
<a href="#">+</a>	<a href="#">Edit</a>	<a href="#">Delete</a>
VRF Name	VRF ID	Status
No data available		

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

## Create VRF



## ▼ VRF Information

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

## ▼ VRF Profile

**Create VRF**

The fields in this screen are:

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

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

**Note**

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

**VRF Template** - This is populated with the *Default\_VRF* template.

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

**3. Click **Create VRF**.**

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

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

Fabric Selected: MSD-Parent-Fabric

Selected 1 / Total 2

Show All

VRFs

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

## Editing and Deleting VRFs in the MSD Fabric

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

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

### VRF Inheritance from MSD-Parent-Fabric to Member1

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

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

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

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

VRF Name	VRF ID	Status
MyVRF_50000	50000	NA
MyVRF_50001	50001	NA

## Editing and Deleting VRFs in the Member Fabric

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

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

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

## Deployment and Undeployment of Networks and VRFs in Member Fabrics

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



### Note

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

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

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

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

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

## NFM Fabric Migration to a DCNM Fabric

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


**Note**

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

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

The two use cases involving NFM migration to DCNM are:

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


**Note**

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

Both use cases are explained in this document.

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

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

## Prerequisites for NFM Fabric Migration to DCNM

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

## Guidelines and Limitations

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

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

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

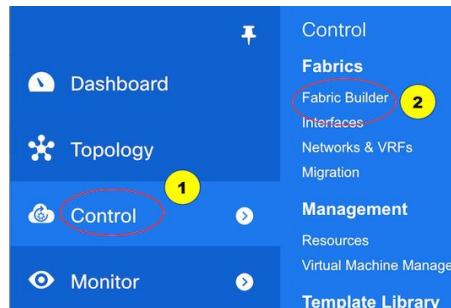
(SVI), VXLAN Network Identifiers (VNI), vPC port channels and so on, in the configurations applied to vPC pair devices.

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

### Create a VXLAN BGP EVPN Fabric Through DCNM

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

1. Click Control > Fabric Builder.



The Fabric Builder page comes up.

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



#### Note

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

The screenshot shows the 'Add Fabric' configuration dialog. At the top, there's a header 'Add Fabric' and a close button 'X'. Below the header, there are two required fields: 'Fabric Name' (set to 'NFM-Fabric') and 'Fabric Template' (set to 'NFM\_Fabric'). Under the 'General' tab, there are three more fields: 'BGP ASN' (set to '65535'), 'Anycast Gateway MAC' (set to 'aabb.bbbb.bbcc'), and 'NX-OS Software Image Version' (a dropdown menu). A note next to the MAC address field says 'Shared MAC address for all leafs (xxxx.xxxx.xxxx)'. Another note next to the NX-OS version field says 'If Set, Image Version Check Enforced On All Switches'.

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

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

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



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

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

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

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

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

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

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

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



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

## Add Fabric

\* Fabric Name : NFM-Fabric  
 \* Fabric Template NFM\_Fabric

General Bootstrap Resources Advanced

Enable DHCP  ⓘ Automatic IP Assignment For POAP  
 \* DHCP Scope Start Address 11.0.1.1 ⓘ Start Address For Switch Out-of-Band POAP  
 \* DHCP Scope End Address 11.0.1.254 ⓘ End Address For Switch Out-of-Band POAP  
 \* Switch Management Default Gateway 11.0.1.0 ⓘ Default Gateway For Mgmt VRF On The Switch  
 \* Switch Management Subnet Prefix 24 ⓘ Prefix For Mgmt0 Interface On The Switch (Min:8 Max:32)

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

## Add Fabric

\* Fabric Name : NFM-Fabric  
 \* Fabric Template NFM\_Fabric

General Bootstrap Resources Advanced

\* Underlay Routing Loopback IP Range 10.1.0.0/22 ⓘ Typically Loopback501 IP Address Range  
 \* Underlay VTEP Loopback IP Range 10.2.0.0/22 ⓘ Typically Loopback500 IP Address Range  
 \* Underlay Subnet IP Range 10.3.0.0/16 ⓘ Address range to assign P2P and Peer Link SVI  
 \* Layer 2 VXLAN VNI Range 30000-49000 ⓘ Overlay Network Identifier Range (Min:1, Max:16777)  
 \* Layer 3 VXLAN VNI Range 50000-59000 ⓘ Overlay VRF Identifier Range (Min:1, Max:16777)  
 \* Network VLAN Range 2300-2999 ⓘ Per Switch Overlay Network VLAN Range (Min:2, Max:4094)  
 \* VRF VLAN Range 2000-2299 ⓘ Per Switch Overlay VRF VLAN Range (Min:2, Max:4094)  
 \* Subinterface Dot1q Range 2-511 ⓘ Per Border Dot1q Range For VRF Lite Connectivity

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

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

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

**Note**

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

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

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

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



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

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

### Advanced tab.

Add Fabric

* Fabric Name :	NFM-Fabric	
* Fabric Template :	NFM_Fabric	
<input type="button" value="General"/> <input type="button" value="Bootstrap"/> <input type="button" value="Resources"/> <input type="button" value="Advanced"/>		
* VRF Template :	Default_VRF	Default Overlay VRF Template For Leafs
* Network Template :	Default_Network	Default Overlay Network Template For Leafs
* VRF Extension Template :	Default_VRF_Extension	Default Overlay VRF Template For Borders
* Network Extension Template :	Default_Network_Extension	Default Overlay Network Template For Borders
Site ID :	16777	For EVPN Multi-Site Support (Min:1, Max:16777)
Fabric MTU :	9216	MTU for fabric interfaces <576-9216>
* OSPF Routing Tag :	100	OSPF Routing Tag
<input type="checkbox"/> Enable OSPF Authentication		

The fields in this tab are:

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

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

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

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



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

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

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

**Fabric MTU** - Specifies the MTU for the fabric interfaces.

**OSPF Routing Tag** - Specifies the OSPF routing tag.

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

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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



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

**Enable vPC Advertise PIP** - Enables the Advertise PIP feature.

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

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

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

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

### Fabric Underlay Migration

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

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

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

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

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

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



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

## Inventory Management

The screenshot shows the 'Discovery Information' tab selected in the navigation bar. The configuration fields include:

- Seed IP:** 172.23.244.91 (with placeholder text: Ex: \*2.2.2.20\*, \*10.10.10.40-60\*, \*2.2.2.20, 2.2.2.21\*)
- Authentication Protocol:** MD5
- Username:** admin
- Password:** (redacted)
- Max Hops:** 2 hop(s)
- Preserve Config:** A toggle switch set to 'yes' (highlighted with a red oval). Below it is the note: "Selecting 'no' will clean up the configuration on switch(es)"

A blue 'Start discovery' button is at the bottom.

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



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

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



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

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

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

## Fabric Overlay Migration

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

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


**Note**

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

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

---

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


**Note**

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

Each overlay network migration consists of the following steps:

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

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

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



## Select a Fabric

Choose the fabric where migration needs to be performed.

NFM-Fabric-1 1 2

The NFM fabric migration page comes up.

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

Discovered VNI	NetworkName	VLAN	Discovered VRF	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status
20003	Auto_Net_VLAN3_VNI20003	3	NA			DISC...
20006	Auto_Net_VLAN6_VNI20006	6	NA			DISC...
20131	Auto_Net_VLAN131_VNI20...	131	VRF4	131.1.1.1/24	2131:1::1/64	DISC...
20132	Auto_Net_VLAN132_VNI20...	132	VRF4	132.1.1.1/24	2132:1::1/64	DISC...
20133	Auto_Net_VLAN133_VNI20...	133	VRF4	133.1.1.1/24		DISC...
20134	Auto_Net_VLAN134_VNI20...	134	VRF4	134.1.1.1/24		DISC...
20140	Auto_Net_VLAN140_VNI20...	140	VRF5	140.1.1.1/24		DISC...

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

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

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

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

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

The screenshot shows the 'Migration' step of the migration process. At the top, it displays 'Fabric Selected: fb1\_nfm' and 'Status: DISCOVERY COMPLETED'. Below this is a table titled 'Rediscover' containing 104 discovered networks. The columns include 'Discovered VNI', 'NetworkName', 'VLAN', 'Discovered VRF', 'IPv4 Gateway/Subnet', 'IPv6 Gateway/Prefix', and 'Status'. All entries show a 'Status' of 'DISC...'. A note at the bottom right indicates 'Selected 0 / Total 104'.

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



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

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

3. After the discovery process is completed, select the networks that you want to migrate to the DCNM fabric.
4. Click **Continue** at the top right part of the screen. The page that appears next has some additional options that allow you to preview existing configurations and the configurations that are going to be deployed on the switches.

The screenshot shows the 'Migration' step of the 'Fabric Selection > Network Selection > Migration' process. The fabric selected is 'fb1' with status 'MIGRATION PENDING'. A table lists four switches: LEAF1, LEAF2, LEAF3, and LEAF5. A modal window titled 'Preview Configuration' displays the configurations on the selected switches and the configurations to be deployed. The configurations include network definitions, VLAN settings, and EVPN routes.

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

**Preview Configuration**

Configurations on Switch:

```
#####
Network:Auto_Net_VLAN3001_VNI23001
#####
vlan 3001
  name 3001-L21
  vn-segment 23001
  interface nve1
    member vni 23001
    ingress-replication protocol bgp
  evpn
    vni 23001 12
    rd auto
    route-target import auto
    route-target export auto
```

Configurations to be Deployed:

```
#####
Network:Auto_Net_VLAN3001_VNI23001
#####
configure profile fb1-Auto_Net_VLAN3001_VNI
vlan 3001
  vn-segment 23001
  interface nve1
    member vni 23001
    ingress-replication protocol bgp
  evpn
```

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

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

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

### Viewing Overlay Migration Status

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

The screenshot shows the 'Migration Status' page with the fabric selected as 'Default\_LAN' and the status as 'MIGRATION PENDING'. The table lists network migrations with their status: Auto\_Net\_VLAN10\_VNI20010, Auto\_Net\_VLAN11\_VNI20011, Auto\_Net\_VLAN12\_VNI20012, Auto\_Net\_VLAN13\_VNI20013, Auto\_Net\_VLAN14\_VNI20014, Auto\_Net\_VLAN15\_VNI20015, and Auto\_Net\_VLAN16\_VNI20016. The first four are marked as 'COMPLETED', while the last three are in 'NETWORK ATTACH IN PROGRESS'.

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

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

## Migration History for Network 'Auto\_Net\_VLAN13\_VNI20013'

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

## Troubleshooting Cisco NFM to Cisco DCNM Migration

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

### Fabric Underlay Troubleshooting

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

The screenshot shows the Cisco DCNM Fabric Builder interface. On the left, there's a sidebar with actions like 'Tabular view', 'Refresh topology', 'Save layout', 'Delete saved layout', 'Random', 'Add switches', and 'Fabric Settings'. The main area is titled 'Fabric errors & warnings' and shows '18 Errors, 0 Warnings, 0 Info'. There are two error entries listed:

- Switch [FDO20270CPF] - OSPF Authentication Key ID [127] on interface [Ethernet1/49]. Fabric setting [126]**
  - Severity: Error
  - Category: Fabric
  - Entity type: Fabric\_Template
  - Entity name: preAdd:validatePreAddInfo:FDO20270CPF:OSPF\_AUTH
  - Reported: 2 minutes ago 2018-07-03 17:32:34
  - Details: [6]: [validatePreAddInfo:FDO20270CPF:OSPF\_AUTH], Line/Col:[0/0]. Msg = [Switch [FDO20270CPF] - OSPF Authentication Key ID [127] on interface [Ethernet1/49]. Fabric setting [126]]
- Switch [FDO20270CPF] - OSPF Authentication Key [6e41a8beb027c7362d365359b534a3eb] on interface [Ethernet1/49]. Fabric setting [5e41a8beb027c7362d365359b534a3eb]**
  - Severity: Error
  - Category: Fabric
  - Entity type: Fabric\_Template
  - Entity name: preAdd:validatePreAddInfo:FDO20270CPF:OSPF\_AUTH
  - Reported: 2 minutes ago 2018-07-03 17:32:34
  - Details: [5]: [validatePreAddInfo:FDO20270CPF:OSPF\_AUTH], Line/Col:[0/0]. Msg = [Switch [FDO20270CPF] - OSPF Authentication Key [6e41a8beb027c7362d365359b534a3eb] on interface [Ethernet1/49]. Fabric setting [5e41a8beb027c7362d365359b534a3eb]]

### Fabric Overlay Troubleshooting

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



## Network Migration Failures

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

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

### Migration Workflow Failures

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

### Migration Workflow Status Definitions

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

#### Discovery-related Status Definitions

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

DISCOVERY IN PROGRESS - The discovery is active.

DISCOVERY FAILED - The previous discovery failed.

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

DISCOVERY ABORTED - The previous discovery has been canceled.

DISCOVERY COMPLETED - The discovery has been completed successfully.

#### Migration-related Status Definitions

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

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

MIGRATION FAILED - The previous migration failed.

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

MIGRATION ABORTED - Migration has been canceled.

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

MIGRATION COMPLETED - All the networks have been migrated.

#### Network Migration Status Definitions

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

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

SWITCH MIGRATION PREPARATION FAILED - The switch preparation step failed.

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

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

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

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

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

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

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

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

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

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

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

**COMPLETED** - The network was migrated successfully.

#### **Network Migration History Definitions**

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

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

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

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

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


**Note**

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

1. Follow the recommended DCNM upgrade procedure and upgrade to DCNM 11.
2. After DCNM is reachable, click **Control > Fabric Builder**.  
The fabrics are listed in a distinct color.
3. Identify the NFM fabric and click **Edit**. Select the NFM fabric from the **Fabric Template** drop-down box. Many of the fabric settings have default values. Review all the settings to make sure that they match your fabric. Refer to the *Create a VXLAN BGP EVPN Fabric Through DCNM* section for information on the fabric settings.
4. Click **Save** at the bottom right part of the screen. All the switches are displayed with the **Migration Mode** tag.
5. Click **Save & Deploy** to complete the migration of the underlay networks .
6. The overlay networks do not need any additional migration action.

## **Post Migration Operations**

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

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

**Note**

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

Now the fabric is ready for use.

### Updating Switch Level Settings

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

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

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

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

## Edit Policy

Policy ID: POLICY-19130  
Entity Type: SWITCH  
\* Priority (1-1000):

Template Name: nfm\_switch\_settings  
Entity Name: SWITCH

General NTP Sylog CDP LLDP Advanced

\* Switch Name:  Host name of the switch

Variables:

**Update**

3. Make changes and click **Update** to update the settings.
4. A **Save & Deploy** pushes these configuration changes to the switch.

**Updating Fabric OSPF Authentication Parameters****Disabling OSPF Authentication**

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

**Note**

The task can cause traffic disruption.

**Enabling or Updating OSPF Authentication**

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

```
nfm-leaf(config)# interface loopback 999 [Pick a non-existent loopback id]
nfm-leaf(config-if)# ip ospf message-digest-key 127 md5 testPassword [Use the desired
```

```

key ID and password]

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

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

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

3. Click the Advanced tab and select the Enable OSPF Authentication check box if not already selected.

4. From the information that is collected earlier, enter the key ID into the OSPF Authentication Key ID
field and the 3DES key as-is into the OSPF Authentication Key field.

5. Click Save.

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

```



**Note** The task can cause traffic disruption.

### Updating Fabric BGP Authentication Parameters

#### Disabling BGP Authentication

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



**Note** The task can cause traffic disruption.

#### Enabling or Updating BGP Authentication

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

```

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

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

```

```
neighbor 1.1.1.1
password 3 f092f5f76d298504ca9b1ad0f1469ca8
```

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

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



**Note** The task can cause traffic disruption.

## Freeform Configurations on Fabric Switches

In DCNM, you can add custom configurations through freeform policies in the following ways:

1. Fabric-wide
  - On all leaf and border switches in the fabric, at once.
  - On all spine switches, at once.
2. On a specific switch.

Leaf switches are identified by the role *Leaf*, border switches by the role *Border* or *Border-Gateway* and spine switches by the role *Spine*.



**Note** You can deploy freeform CLIs when you create a fabric or when a fabric is already created. The following examples are for an existing fabric. However, you can use them as a reference for a new fabric.

### Deploy Fabric-Wide Freeform CLIs on Leaf and Spine Switches

1. Click **Control > Fabric Builder**. The Fabric Builder screen comes up. A rectangular box represents each fabric.
2. Click the **Settings** icon (located on the top right part of the rectangular box) for adding custom configurations to an existing fabric. The **Edit Fabric** screen comes up.  
(If you are creating a fabric for the first time, click **Create Fabric**).
3. Click the **Advanced** tab and update the following fields:

**Leaf Freeform Config** – In this field, add configurations for all leaf and border switches in the fabric. For example, you can add NTP, TACAS, and AAA configurations in this field.

Don't add VLAN, SVI, and interface-specific configurations.

**Spine Freeform Config** - In this field, add configurations for all spine switches in the fabric.



**Note**

Copy-paste the intended configuration with correct indentation, as seen in the running configuration on the Nexus switches. For more information, see [Resolving Freeform Config Errors in Switches, on page 133](#).

Edit Fabric

* Fabric Name :	green
* Fabric Template :	Easy_Fabric
<input type="button" value="General"/> <input type="button" value="Advanced"/> <input type="button" value="Resources"/> <input type="button" value="Manageability"/> <input type="button" value="Bootstrap"/>	
<b>CoPP Profile :</b> strict <span style="float: right;">Fabric Wide CoPP Policy</span>	
<input checked="" type="checkbox"/> <span style="float: right;">For Operations And Management Of VXLAN Fabrics</span>	
<input type="checkbox"/> <span style="float: right;">For Overlay Multicast Support In VXLAN Fabrics</span>	
<input type="checkbox"/> <span style="float: right;">For Primary VTEP IP Advertisement As Next-Hop Of Prefix Routes</span>	
<b>Leaf Freeform Config</b> <span style="float: right;">Additional CLIs For All Leafs As Captured From Sh</span>	
feature bash-shell feature scp-server	
<b>Spine Freeform Config</b> <span style="float: right;">Additional CLIs For All Spines As Captured From S</span>	
<input type="button" value="Save"/> <input type="button" value="Cancel"/>	

4. Click **Save**. The Fabric Builder screen comes up again.
5. Click within the box that represents the fabric. The Fabric Topology screen comes up.
6. Click **Save & Deploy** at the top right part of the screen to save and deploy configurations.

Configuration Compliance functionality will ensure that that intended configuration as expressed by those CLIs are present on the switches and if they are removed or there is a mismatch, then it will flag it as a mismatch and indicate that the device is OUT-OF-SYNC.

*Incomplete Configuration Compliance* - On some Cisco Nexus 9000 Series switches, in spite of configuring pending switch configurations using the **Save & Deploy** option, configuration compliance is not successful. Add a **switch\_freeform\_config** policy to the affected switch (as explained in the *Deploy Freeform CLIs on a Specific Switch* section) to resolve the issue. For example, consider the following persistent pending configurations:

```
line vty
logout-warning 0
```

After adding the above configurations in a **switch\_freeform\_config** policy and saving the updates, click **Save and Deploy** in the topology screen to complete the deployment process.

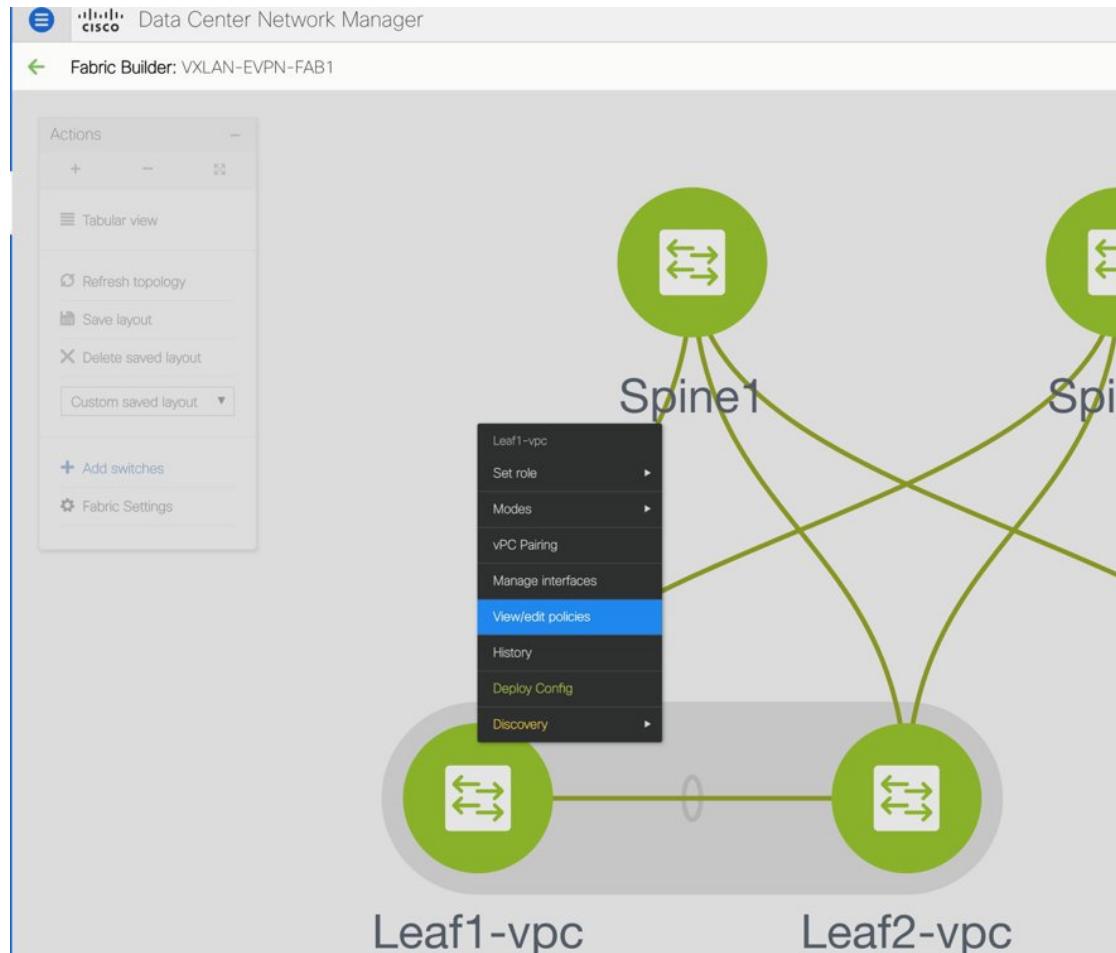
### Deploy Freeform CLIs on a Specific Switch

1. Click **Control > Fabric Builder**. The Fabric Builder screen comes up.
2. Click on the rectangular box that represents the fabric. The Fabric Topology screen comes up.



**Note** To provision freeform CLIs on a new fabric, you have to create a fabric, import switches into it, and then deploy freeform CLIs.

3. Right-click the switch icon and select the **View/edit policies** option.



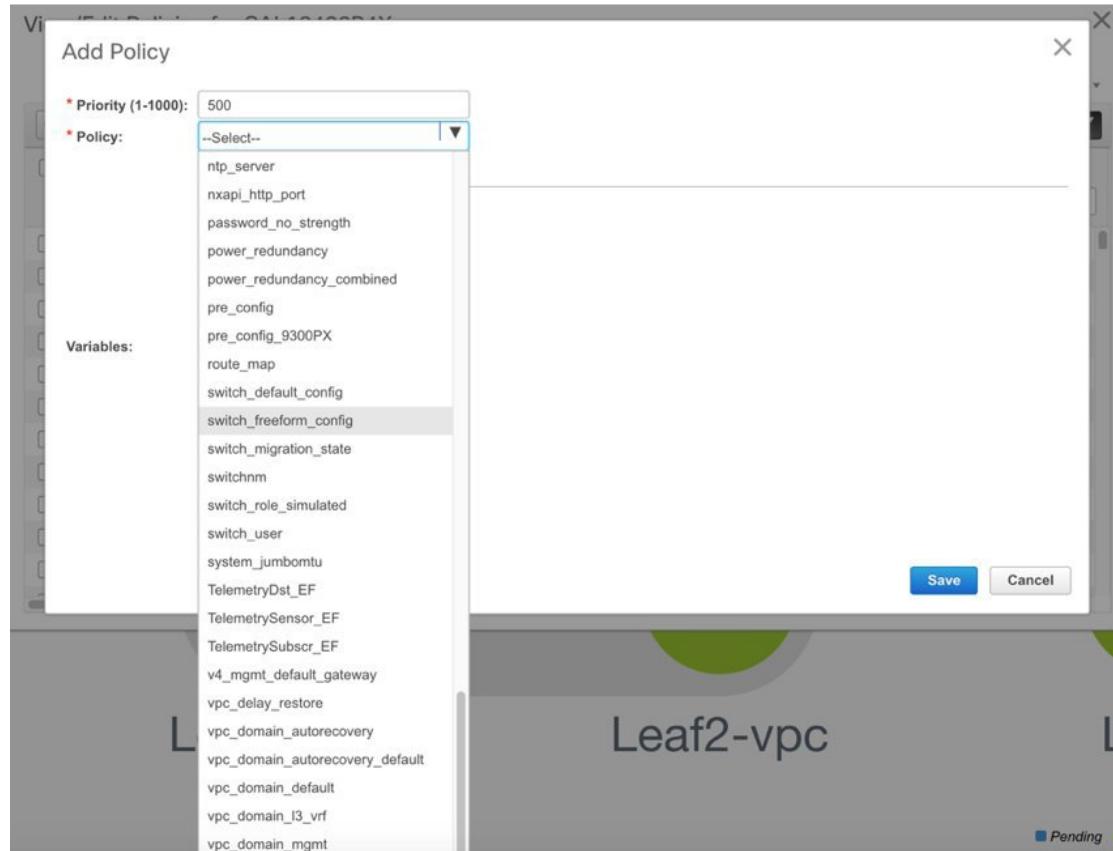
The **View/Edit Policies** screen comes up.

View/Edit Policies for SAL18432P4X							
		Priority	Editable	Entity Type	Entity Name	Source	Policy ID
<input type="checkbox"/>	Template						
<input type="checkbox"/>	switch_role_simulated	10	true	SWITCH	SWITCH		POLICY-290290
<input type="checkbox"/>	host	50	true	SWITCH	SWITCH		POLICY-277130
<input type="checkbox"/>	nfm_switch_user	100	true	SWITCH	SWITCH		POLICY-277110
<input type="checkbox"/>	ntp_server	100	true	SWITCH	SWITCH		POLICY-277200
<input type="checkbox"/>	power_redundancy	100	true	SWITCH	SWITCH		POLICY-277220
<input type="checkbox"/>	aaa_radius_use_vrf	151	true	SWITCH	SWITCH		POLICY-277210
<input type="checkbox"/>	feature_tacacs	50	false	SWITCH	SWITCH	UNDERLAY	POLICY-277260
<input type="checkbox"/>	feature_pim	50	false	SWITCH	SWITCH	UNDERLAY	POLICY-277270
<input type="checkbox"/>	feature_ngoam	50	false	SWITCH	SWITCH	UNDERLAY	POLICY-277280
<input type="checkbox"/>	copp_policy	50	false	SWITCH	SWITCH	UNDERLAY	POLICY-289830
<input type="checkbox"/>	base_feature_vpc	50	false	SWITCH	SWITCH	UNDERLAY	POLICY-290410

4. Click +. The **Add Policy** screen comes up.

In the **Priority** field, the priority is set to 500 by default. You can choose a higher priority (by specifying a lower number) for CLIs that need to appear higher up during deployment. For example, a command to enable a feature should appear earlier in the list of commands.

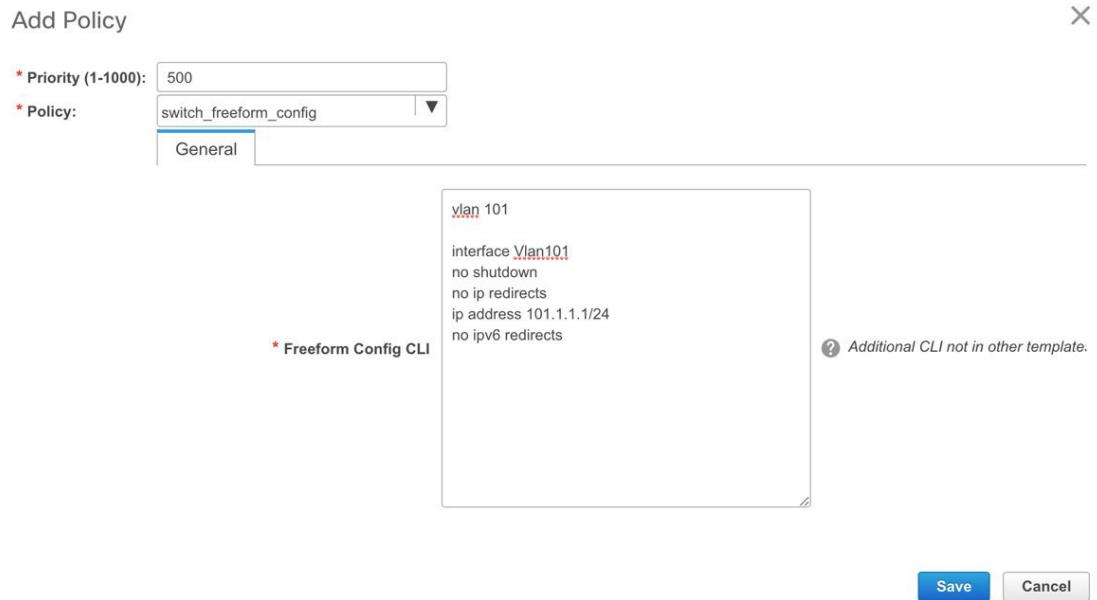
5. From the **Policy** field, select **switch\_freeform\_config**.



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

Copy-paste the intended configuration with correct indentation, as seen in the running configuration on the Nexus switches. For more information, see [Resolving Freeform Config Errors in Switches, on page 133](#).

A **switch\_freeform\_config** policy example for VLAN and corresponding SVI instantiation is given below.



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

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

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

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

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

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

- You can create multiple instances of the policy.
- You can add VLAN, SVI and other features. A specific VLAN and corresponding SVI instantiation should be configured through an individual **switch\_freeform\_config** policy.
- For a vPC switch pair, create consistent **switch\_freeform\_config** policies on both the vPC switches.
- Depending on the Cisco Nexus 9000 Series platform type (required for EX, FX, and FX2 platform types), you should include the **system nve infra-vlans 101** command in the policy.

### Freeform CLI Configuration Examples



**Note**

Refer the *Deploy Fabric-Wide Freeform CLIs on Leaf and Spine Switches* section and *Deploy Freeform CLIs on a Specific Switch* section for complete steps.

#### Console line configuration

This example involves deploying some fabric-wide freeform configurations (for all leaf, and spine switches), and individual switch configurations.

Fabric-wide session timeout configuration:

```
line console
  exec-timeout 1
```

Console speed configuration on a specific switch:

```
line console
  speed 115200
```

On the switch where the console speed was updated, both types of configurations are displayed:

```
N9k-switch # show run | b console

line console
  exec-timeout 0
  speed 115200
```

### ACL configuration

ACL configurations are typically configured on specific switches and not fabric-wide (leaf/spine switches). When you configure ACLs as freeform CLIs on a switch, you should include sequence numbers. Else, there will be a mismatch between the switch and DCNM. A configuration sample with sequence numbers:

```
ip access-list ACL_VTY
  10 deny tcp 172.29.171.67/32 172.29.171.36/32
  20 permit ip any any
ip access-list vlan65-acl
  10 permit ip 69.1.1.201/32 65.1.1.11/32
  20 deny ip any any

interface Vlan65
  ip access-group vlan65-acl in
line vty
  access-class ACL_VTY in
```

If you have configured ACLs without sequence numbers in a **switch\_freeform\_config** policy, update the policy with sequence numbers *as displayed in the switch*. After updating, use the per switch **Deploy Config** option by right clicking the device. Alternatively, use the **Save and Deploy** option in the topology screen so that configuration compliance is triggered again and inconsistencies resolved.

### Negotiation, speed and duplex port configuration

Consider the following commands configured on a leaf switch whose Ethernet1/10 interface is connected to a spine switch. The ethernet port speed, duplex mode and disabling of automatic negotiation of speed and duplex abilities over the link are configured for the interface.

```
interface Ethernet1/10
  speed 100000
  duplex full
  no negotiate auto
```

This can be configured as a **switch\_freeform\_config** policy on a switch.

If the above parameters are the same for all leaf switches (interface 1/10 on each leaf switch has the same settings and connected to a switch), then you can update fabric-wide CLIs for all leaf switches.

In the same way, you can configure all spine switches with the same port name and speed, duplex mode and negotiation settings.



- 
- Note** If you are enabling freeform configurations on all leaf or spine switches, as a best practice, ensure that all switches are connected through the same type of cable. For example, Active Optical Cables or Direct Attach Copper cables.
- 

### Resolving Freeform Config Errors in Switches

Copy-paste the running-config to the freeform config with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. Otherwise, configuration compliance in DCNM marks switches as out-of-sync.

Let us see an example of the freeform config of a switch.

```
feature bash-shell
feature telemetry

clock timezone CET 1 0
# Daylight saving time is observed in Metropolitan France from the last Sunday in March
(02:00 CET) to the last Sunday in October (03:00 CEST)
clock summer-time CEST 5 Sunday March 02:00 5 Sunday October 03:00 60
clock protocol ntp

telemetry
  destination-profile
    use-vrf management
```

The highlighted line about the daylight saving time is a comment that is not displayed in the **show running config** command output. Therefore, configuration compliance marks the switch as out-of-sync because the intent does not match the running configuration.

Let us check the running config in the switch for the clock protocol.

```
spine1# show run all | grep "clock protocol"
clock protocol ntp vdc 1
```

You can see that **vdc 1** is missing from the freeform config.

In this example, let us copy-paste the running config to the freeform config.

Here is the updated freeform config:

```
feature bash-shell
feature telemetry

clock timezone CET 1 0
clock summer-time CEST 5 Sunday March 02:00 5 Sunday October 03:00 60
clock protocol ntp vdc 1

telemetry
```

```
destination-profile
use-vrf management
```

After you copy-paste the running config and deploy, the switch will be in-sync. When you click **Save & Deploy**, the **Side-by-side Comparison** tab in the **Config Preview** window provides you information about the difference between the defined intent and the running config.

# Management

The Management menu includes the following submenus:

## Resources

Cisco DCNM allows you to manage the resources. The following table describes the fields that appear on this page.

Field	Description
Scope Type	Specifies the scope level at which the resources are managed. The scope types can be <b>Fabric</b> , <b>Device</b> , <b>DeviceInterface</b> , or <b>DevicePair</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 that are managed with device, device interface, or fabric. Valid values are ID type, subnet, or IP addresses.
Allocated To	Specifies the purpose of resource allocation.
Resource Type	Specifies the resource type. The valid values are <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 .

You see the list of VMware servers (if any) that are managed by Cisco DCNM-LAN in the table.

**Step 2** Click **Add**.

You see the **Add VCenter** window.

**Step 3** Enter the **Virtual Center Server** IP address for this VMware server.

**Step 4** Enter the **User Name** and **Password** for this VMware server.

**Step 5** Click **Add** to begin managing this VMware server.

---

## Deleting a VMware Server

You can remove a VMware server from the Cisco DCNM.

---

**Procedure**

---

**Step 1** Choose .

**Step 2** Select the check box next to the VMware server that you want to remove and click **Delete** to discontinue data collection for that VMware server.

---

## Editing a VMware Server

You can edit a VMware server from Cisco DCNM Web Client.

---

**Procedure**

---

**Step 1** Choose .

**Step 2** Check the check box next to the VMware server that you want to edit and click **Edit virtual center icon**.

You see the **Edit VCenter** dialog box.

**Step 3** Enter a the **User Name** and **Password**.

**Step 4** Select managed or unmanaged status.

**Step 5** Click **Apply** to save the changes.

---

## Rediscovering a VMware Server

You can rediscover a VMware server from Cisco DCNM.

### Procedure

---

- Step 1** Choose .
- Step 2** Select the check box next to the VMware that you want to rediscover.
- Step 3** Click **Rediscover**.  
A dialog box with warning "Please wait for rediscovery operation to complete." appears.
- Step 4** Click **OK** in the dialog box.
- 

## Template Library

You can add, edit, or delete templates that are configured across different Cisco Nexus and Cisco MDS platforms using Cisco DCNM Web client. From Cisco DCNM Web client home page, choose **Control > Template Library > Templates**. The following parameters are displayed for each template that is configured on Cisco DCNM Web client. Templates support JavaScript. You can use the JavaScript function in a template to perform arithmetic operations and string manipulations in the template syntax.

The following table describes the fields that appear on this page.

**Table 1: Templates Operations**

Field	Description
Add Template	Allows you to add a new template.
Modify/View Template	Allows you to view the template definition and modify as required.
Save Template As	Allows you to save the selected template in a different name. You can edit the template as required.
Delete Template	Allows you to delete a template
Import Template	Allows you to import a template from your local directory, one at a time.
Export template	Allows you to export the template configuration to a local directory location.
Import Template Zip File	Allows you to import .zip file, that contains more than one template that is bundled in a .zip format  All the templates in the ZIP file are extracted and listed in the table as individual templates.

**Table 2: Template Properties**

Field	Description
Template Name	Displays the name of the configured template.

Field	Description
Template Description	Displays the description that is provided while configuring templates.
Tags	Displays the tag that is assigned for the template and aids to filter templates based on the tags.
Implements	Displays the abstract template to be implemented.
Dependencies	Specifies the specific feature of a switch.
Supported Platforms	Displays the supported Cisco Nexus platforms compatible with the template. Check the check box of platforms that are supported with the template. <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.
Published	Specifies if the template is published or not.
Imports	Specifies the base template for importing.

In addition, from the menu bar, choose **Control > Template Library > Templates** and you can also:

- Click **Show Filter** to filter the templates that is based on the headers.
- Click **Print** to print the list of templates.
- Click **Export to Excel** to export the list of template to a Microsoft Excel spreadsheet.

This section contains the following:

## Template Structure

The configuration template content mainly consists of four parts. Click the **Help** icon next to the **Template Content** for information about editing the content of the template.

This section contains the following:

## Template Format

This section describes the basic information of the template. The possible fields are as detailed in the table below.

Property Name	Description	Valid Values	Optional?
name	The name of the template	Text	No

Property Name	Description	Valid Values	Optional?
description	Brief description about the template	Text	Yes
userDefined	Indicates whether the user created the template. Value is ‘true’ if user created.	“true” or “false”	Yes
supportedPlatforms	List of device platforms supports this configuration template. Specify ‘All’ to support all platforms.	N1K, N3K, N3500, N4K, N5K, N5500, N5600, N6K, N7K, N9K, MDS, VDC, N9K-9000v, All list separated by comma.	No
templateType	Specifies the type of Template used.	<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</li> <li>• N/A</li>   <li>• POAP</li> <li>• N/A</li> <li>• VXLAN</li> <li>• FABRICPATH</li> <li>• VLAN</li> <li>• PMN</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>• VLAN</li> <li>• INTERFACE_VLAN</li> <li>• INTERFACE_ETHERNET</li> <li>• INTERFACE_BD</li> <li>• INTERFACE_PORT_CHANNEL</li> <li>• INTERFACE_FC</li> <li>• INTERFACE_MGMT</li> <li>• INTERFACE_LOOPBACK</li> <li>• INTERFACE_NVE</li> <li>• INTERFACE_VFC</li> <li>• INTERFACE_SAN_PORT_CHANNEL</li> <li>• DEVICE</li> <li>• FEX</li> <li>• INTERFACE</li>   <li>• SHOW</li> <li>• VLAN</li> <li>• INTERFACE_VLAN</li> </ul>	

Property Name	Description	Valid Values	Optional?
		<ul style="list-style-type: none"><li>• INTERFACE_VPC</li><li>• INTERFACE_ETHERNET</li><li>• INTERFACE_BD</li><li>• INTERFACE_PORT_CHANNEL</li><li>• INTERFACE_FC</li><li>• INTERFACE_MGMT</li><li>• INTERFACE_LOOPBACK</li><li>• INTERFACE_NVE</li><li>• INTERFACE_VFC</li><li>• INTERFACE_SAN_PORT_CHANNEL</li><li>• DEVICE</li><li>• FEX</li><li>• INTERFACE</li><li>• PROFILE</li><li>• VXLAN</li><li>• FABRIC</li><li>• NA</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_ETHERNET</li> <li>• INTERFACE_BD</li> <li>• INTERFACE_PORT_CHANNEL</li> <li>• INTERFACE_FC</li> <li>• INTERFACE_MGMT</li> <li>• INTERFACE_LOOPBACK</li> <li>• INTERFACE_NVE</li> <li>• INTERFACE_VFC</li> <li>• INTERFACE_SAN_PORT_CHANNEL</li> <li>• DEVICE</li> <li>• FEX</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
timestamp	Shows the template modified time	Modified date and time in the format YYYY-MM-DD HH:MM:SS	Yes

## Template Variables

This section contains declared variables, the data type, default values, and valid values conditions for the parameters that are used in the template. These declared variables are used for value substitution in the template content section during the dynamic command generation process. Also these variables are used in decision making and in iteration blocks in the template content section. Variables have predefined data types. You can also add a description about the variable. The following table describes the syntax and usage for the available datatypes.

Variable Type	Valid Value	Iterative?
boolean	true false	No
enum	Example: running-config, startup-config	No
float	Floating number format	No
floatRange	Example: 10.1,50.01	Yes
Integer	Any number	No
integerRange	Contiguous numbers separated by “-“ Discrete numbers separated by “;”  Example: 1-10,15,18,20	Yes
interface	Format: <if type><slot>[<sub slot>]</port>  Example: eth1/1, fa10/1/2 etc.	No
interfaceRange	Example: eth10/1/20-25, eth11/1-5	Yes
ipAddress	IPv4 OR IPv6 address	No
ipAddressList	Example: 172.22.31.97, 172.22.31.99, 172.22.31.105, 172.22.31.109	Yes
ipAddressWithoutPrefix	Example: 192.168.1.1  or  Example: 1:2:3:4:5:6:7:8	No
ipV4Address	IPv4 address	No
ipV4AddressWithSubnet	Example: 192.168.1.1/24	No
ipV6Address	IPv6 address	No
ipV6AddressWithPrefix	Example: 1:2:3:4:5:6:7:8  22	No
ipV6AddressWithSubnet	IPv6 Address with Subnet	No

Variable Type	Valid Value	Iterative?
ISISNetAddress	Example: 49.0001.00a0.c96b.c490.00	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
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;];</pre> <p>struct interface_detail {              string inf_name;              string inf_description;              ipAddress inf_host;              enum duplex {                  validValues = auto, full, half;              };              myInterface, myInterfaceArray[];</p>	<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)	Example: 20:01:00:08:02:11:05:03	No

## Variable Meta Property

Each variable that is defined in the template variable section has a set of meta properties. The meta properties are mainly the validation rules that are defined for the variable.

The following table describes the various meta properties applicable for the available variable types.

## Variable Meta Property

Variable Type	Definition	Variable Meta Property											
		default Value	valid Values	decimal Length	min	max	min Slot	max Slot	min Port	max Port	min Length	max Length	regular Expr
boolean	A boolean value. Example: true	Yes											
enum			Yes										
float	signed real number Example: 75.56, -8.5	Yes	Yes	Yes	Yes	Yes							
faRange	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							
intRange	Range of signed numbers Example: 50-65	Yes	Yes		Yes	Yes							
interface	specifies interface Example: Ethernet 5/10	Yes	Yes				Yes	Yes	Yes	Yes			
intfRange		Yes	Yes				Yes	Yes	Yes	Yes			

Variable Type	Description	Variable Meta Property											
		default Value	valid Values	decimal Length	min	max	min Slot	max Slot	min Port	max Port	min Length	max Length	regular Expr
pAddr	IP address in IPv4 or IPv6 format	Yes											
pAddrList	Example: <u>192.168.1.1, 192.168.1.2</u>  Note	Yes											
pAdyRef	IPv4 or IPv6 Address (does not require pAddr)												
pV4Adrs	IPv4 address	Yes											
pV4Sub	IPv4 Address with Subnet	Yes											
pV6Adrs	IPv6 address	Yes											
pV6Sub	IPv6 Address with prefix	Yes											

## Variable Meta Property

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
IPv4Subnet	IPv6 Address with Subnet	Yes											
IPv4Subnet	Example: 4006:5:20												
long	Example: 100	Yes			Yes	Yes							
macAddress	MAC address												
string	literal string  Example for string  Regular expression string  stringLiteral {  stringLiteral }	Yes									Yes	Yes	Yes
string[]	string literals that are separated by a comma (,)  Example: {string1, string2}	Yes											

Variable Type	Description	Variable Meta Property											
		default Value	valid Values	decimal Length	min	max	min Slot	max Slot	min Port	max Port	min Length	max Length	regular Expr
struct	Set of <b>parts</b> that are bundled under a single variable.  struct <structure name declaration> { <parameter type>  <parameter 1>; <parameter type>  <parameter 2>; .... > [ <parameter> [ <parameter> []>];												
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”	
IsSiteId	“true” or “false”	
IsSourceDevice	“true” or “false”	
IsSourceFabric	“true” or “false”	
IsSourceInterface	“true” or “false”	
IsSourceSwitchName	“true” or “false”	
IsSwitchName	“true” or “false”	
IsRMID	“true” or “false”	

**Variable Annotation**

<b>Annotation Key</b>	<b>Valid Values</b>	<b>Description</b>
IsVPCDomainID	“true” or “false”	
IsVPCID	“true” or “false”	
IsVPCPeerLinkPort	“true” or “false”	
IsVPCPeerLinkPortChannel	“true” or “false”	
IsVPCPortChannel	“true” or “false”	
Password	Text	Validates the password field
UsePool	“true” or “false”	
UseDNSReverseLookup		
Username	Text	Displays the username field on the window

**Example: DisplayName Annotation**

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

**Example: IsMandatory Annotation**

```
##template variables
@{IsMandatory="ipv6!=null"}
ipV4Address ipv4;
@{IsMandatory="ipv4!=null"}
ipV6Address ipv6;
##
```

**Example: IsMultiLineString Annotation**

```
##template variables
@{IsMultiLineString=true}
string EXTRA_CONF_SPINE;
##
```

**IsShow Annotation**

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

##template variables
boolean enableScheduledBackup;
@(IsShow="enableScheduledBackup==true",Description="Server time")
```

```

string scheduledTime;
##
The condition "enableScheduledBackup==true" evaluates to true/false

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

```

## Templates Content

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



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

Template content is governed by the usage of variables.

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

Syntax: \$\$<variable name>\$\$  
Example: \$\$USER\_NAME\$\$

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

Syntax:@<loop variable>  
Example:  
foreach val in \$\$INTEGER\_RANGE\_VALUES\$\$ {  
@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.

## Advanced Features

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

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

```
Syntax:
foreach <loop index variable> in $$<loop variable>$$ {
    @<loop index variable> ..
}
Example: foreach Statement
foreach ports in $$MY_INF_RANGE$${
    interface @ports
    no shut
}
```

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

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

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

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

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

## Advanced Features

The following are the advanced features available to configure templates.

- Assignment Operation

Config template supports assignment of variable values inside the template content section. The values are validated for the declared data type of the variable. If there is a mismatch, the value is not assigned.

Assignment operation can be used under the following guidelines:

- The operator on the left must be any of the template parameters or a for loop parameter.
- The operator on the right values can be any of the values from template parameters, for loop parameters, literal string values surrounded by quotes or simple string values.

If a statement does not follow these guidelines, or if it does not suit this format, it will not be considered as assignment operation. It is substituted during command generation like other normal lines.

```
Example: Template with assignment operation
##template properties
name =vlan creation;
userDefined= true;
supportedPlatforms = All;
templateType = CLI;
published = false;
##
##template variables
integerRange vlan_range;
@(internal=true)
integer vlanName;
##
##template content
foreach vlanID in $$vlan_range$${
vlan @vlanID
$$vlanName$$=@vlanID
name myvlan$$vlanName$$
}
##
```

- Evaluate methods

Config template uses the Java runtime provided Java script environment to perform arithmetic operations (such as ADD, SUBTRACT, and so on), string manipulations, and so on.

Locate the JavaScript file in the template repository path. This file contains primary set of arithmetic, string functions. You can also add custom JavaScript methods.

These methods can be called from config template content section in below format:

```
Example1:
$$somevar$$ = evalscript(add, "100", $$anothervar$$)
```

Also the *evalscript* can be called inside if conditions as below:

```
if($$range$$ > evalscript(sum, $$vlan_id$$, -10)) {
do something...
}
```

You can call a method that is located at the backend of the Java script file.

- Dynamic decision

Config template provides a special internal variable “LAST\_CMD\_RESPONSE”. This variable stores the last command response from the device during the execution of the command. This can be used in the config template content to make dynamic decisions to deliver the commands that are based on the device condition.



**Note** The if block must be followed by an else block in a new line, which can be empty.

An example use case to create a VLAN, if it does not exist on the device.

```
Example: Create VLAN
##template content
show vlan id $$vlan_id$$
if($$LAST_CMD_RESPONSE$$ contains "not found"){
vlan $$vlan_id$$
}
else{
}
##
```

This special implicit variable can be used only in the “IF” blocks.

- Template referencing

You can have a base template with all the variables defined. This base template can be imported to multiple templates. The base template content is substituted in the appropriate place of the extending template. The imported template parameters and the contents can be accessed inside the extending template.

```
Example: Template Referencing
Base template:
##template properties
name =a vlan base;
userDefined= true;
supportedPlatforms = All;
templateType = CLI;
published = false;
timestamp = 2015-07-14 16:07:52;
imports = ;
##
##template variables
integer vlan_id;
##
##template content
vlan $$vlan_id$$
##

Derived Template:
##template properties
name =a vlan extended;
userDefined= true;
supportedPlatforms = All;
templateType = CLI;
published = false;
timestamp = 2015-07-14 16:07:52;
imports = a vlan base,template2;
##
##template variables
interface vlanInterface;
##
##template content
<substitute a vlan base>
interface $$vlanInterface$$
<substitute a vlan base>
##
```

When you launch the extended template, the parameter inputs for the base template are also obtained. In addition, the substituted content is used for complete CLI command generation.

## Adding a Template

To add user-defined templates and schedule jobs from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

**Step 1** Choose **Control > Template Library**.

The **Templates** window is displayed with the name of the template along with its description, supported platforms, and tags.

**Step 2** Click **Add** to add a new template.

**Step 3** Specify a **Template Name**, **Template Description**, **Tags**, **Implements**, and **Dependencies** for the new template. Specify a template name, description, tags, and supported platforms for the new template.

**Step 4** Select the supported platforms that the template must support.

**Step 5** Specify a **Template Type** for the template.

**Step 6** Select a **Template Sub Type** and **Template Content Type** for the template. Select **Published** to make the template read-only. You cannot edit a published template.

**Step 7** Click **Template Content** to edit the template syntax. For information about the structure of the Configuration Template, see the *Template Structure* section.

**Step 8** From the **Imports > Template Name** list, check the template check box.

The base template content is displayed in the **Template Content** window. The base template displays the template properties, template variables, and template content. This template can be imported in to another template and the base template content is substituted in the appropriate place of the extending template. When you launch the extended template, the parameter inputs for the base template are also obtained. Also, the substituted content is used for complete CLI command generation.

**Note** The base templates are CLI templates.

**Step 9** Click **Validate Template Syntax** to validate the template values.

If an error or a warning message appears, you can check the validation details in **Validation Table**.

**Note** You can continue to save the template if there are warnings only. However, if there is an error, you must edit the templates to fix the errors before you proceed.

**Step 10** Click **Save** to save the template.

**Step 11** Click **Save and Exit** to save the configuration and go back to the configuring templates screen.

---

## Modifying a Template

You can edit the user-defined templates. However, the predefined templates and templates that are already published cannot be edited.

**Procedure**

- 
- Step 1** From **Control > Template Library**, select a template.
- Step 2** Click **Modify/View template**.
- Step 3** Edit the template description and tags.  
The edited template content is displayed in a pane on the right.
- Step 4** From the **Imports > Template Name** list, check the template check box.  
The base template content is displayed in the **Template Content** window. You can edit the template content based on your requirement in the **Template Content** window. Click the help icon next to the **Template Content** window for information about editing the content of the template.
- Step 5** Edit the supported platforms for the template.
- Step 6** Click **Validate Template Syntax** to validate the template values.
- Step 7** Click **Save** to save the template.
- Step 8** Click **Save and Exit** to save the configuration and go back to the configuring templates screen.
- 

## Copying a Template

To copy a template from the Cisco DCNM Web UI, perform the following steps:

**Procedure**

- 
- Step 1** Choose **Control > Template Library**, and select a template.
- Step 2** Click **Save Template As**.
- Step 3** Edit the template name, description, tags, and other parameters.  
The edited template content is displayed in the right-hand pane.
- Step 4** From the **Imports > Template Name** list, check the template check box.  
The base template content is displayed in the **Template Content** window. You can edit the template content that is based on your requirement in the **Template Content** window. Click the help icon next to the **Template Content** window for information about editing the content of the template.
- Step 5** Edit the supported platforms for the template.
- Step 6** Click **Validate Template Syntax** to validate the template values.
- Step 7** Click **Save** to save the template.
- Step 8** Click **Save and Exit** to save the configuration and go back to the configuring templates screen.
-

## Deleting a Template

You can delete the user-defined templates. However, you cannot delete the predefined templates. From Cisco DCNM Release 11.0(1), you can delete multiple templates at once.

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

### Procedure

---

**Step 1** Choose **Control > Template Library**.

**Step 2** Use the check box to select a template and click **Remove template** icon.

The template is deleted without any warning message.

---

### What to do next

The template is deleted from the list of templates on the DCNM Web UI. When you restart the DCNM services, the deleted templates are displayed on the **Control > Template Library** page.

To delete the template permanently, delete the template that is located in your local directory: `Cisco Systems\dcm\dcnm\data\templates\`.

## Importing a Template

To import a template from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

**Step 1** Choose **Control > Template Library** and click **Import Template**.

**Step 2** Browse and select the template that is saved on your computer.

You can edit the template parameters, if necessary. For information, see [Modifying a Template, on page 157](#).

**Note** The “`n`” in the template is considered as a new line character when imported and edited, but it works fine when imported as a ZIP file.

**Step 3** Click **Validate Template Syntax** to validate the template.

**Step 4** Click **Save** to save the template or **Save and Exit** to save the template and exit.

---

## Exporting a Template

To export a template from the Cisco DCNM Web UI, perform the following steps:

**Procedure**

- 
- Step 1** Choose **Control > Template Library**.
- Step 2** Use the check box to select a template and click **Export Template**.  
The browser requests you to open or save the template to your directory.
- 

# Image Management

The **Image Management** menu includes the following options:

## Image Upload

This feature allows you to upload or delete images that are used during POAP and switch upgrade.

### Deleting an Image

To delete an image from the repository from the Cisco DCNM Web UI, perform the following steps:

**Procedure**

- 
- Step 1** On the **Smart Image Management** window, select an existing image from the list, and click **Delete**.
- Step 2** In the delete notification, click **Yes** to delete the image.
- Note** The default SCP Repository cannot be deleted.
- 

## Image Upload

To upload different types of images to the server from the Cisco DCNM Web UI, perform the following steps:



**Note** Devices use these images during POAP.

**Procedure**

- 
- Step 1** On the **Smart Image Management** window, check the server name check box to select the server for uploading images.  
The **Select Image File** window appears.
- Step 2** Click **Browse** to select the image file from the directory.

- Step 3** From the **Platform** drop-down list, select the device to which you must upload this image.
- Step 4** From the **Type** drop-down list, select the type of the image you are uploading to the device.
- Step 5** Click **OK**.
- The image is uploaded to the repository.
- 

## Install & Upgrade

The **Install & Upgrade** menu includes the following submenus:

### Upgrade History

This feature enables you to upgrade the Cisco Nexus Platform Switches using In-Service Software Upgrade (ISSU). This upgrade procedure may be disruptive or non-disruptive based on the device configuration. You can select the Kickstart, System, or NX-OS images from image repository or the file system on the device. To select the images from the repository, the same needs to be uploaded from **Control > Image Management > Image upload** tab.

The following table describes the fields that appear on **Control > Image Management > Upgrade History**.

Field	Description
Task Id	Specifies the serial number of the task. The latest task will be listed in the top. <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.

Field	Description
Scheduled At	Specifies the time when the task is specified to be executed. You can also choose to schedule a task to be executed at a later time.
Comment	Shows any comments that the Owner has added while performing the task.



**Note** After a fresh Cisco DCNM installation, this page will have no entries.

You can perform the following:

## New Installation

To upgrade the devices that are discovered from the Cisco DCNM, perform the following steps:

### Procedure

- Step 1** Choose **Control > Image Management > Install & Upgrade > Upgrade History**, click **New Installation** to install, or upgrade the kickstart and the system images on the devices.  
The devices with default VDCs are displayed in the **Select Switches** window.
- Step 2** Select the check box to the left of the switch name.  
You can select more than one device and move the devices to the right column.
- Step 3** Click **Add** or **Remove** icons to include the appropriate switches for upgrade.  
The selected switches appear in a column on the right.
- Step 4** Click **Next** to navigate to the **Specify Software Images** window. This tab displays the switches that you selected in the previous screen and allows you to choose the images for upgrade.
- The **Auto File Selection** check box enables you to specify an image version, and a path where you can apply the upgraded image to the selected devices.
  - **Select File Server** is disabled, and the default server is used.
  - In the **Image Version** field, specify the image version as displayed in **Image Upload** screen.
  - The **Path** field is disabled, and the default image path is used.

- Step 5** Click **Select Image** in the **Kickstart image** column.

The **Software Image Browser** dialog box appears.

- Note**
- Cisco Nexus 9000 Series Switches require only the system image to load the Cisco NX-OS operating system. Therefore, the option to select kickstart images for these devices is disabled.
  - If there is an issue in viewing the **Software Image Browser** dialog box, reduce the font size of your browser and retry.

- Step 6** Click **Select Image** in the **System Image** column.
- The **Software Image Browser** dialog box appears.
- Step 7** On the **Software Image Browser** dialog box, you can choose the image from **File Server** or **Switch File System**.
- If you choose **File Server**:
- From the **Select the File server** list, choose the appropriate file server on which the image is stored.
  - 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.
- Click **OK** to choose the kickstart image or **Cancel** to revert to the **Specify Software Images** window.
- If you choose **Switch File System**:
- From the **Select Image** list, choose the appropriate image that is located on the flash memory of the device.
  - Click **OK** to choose the kickstart image or **Cancel** to revert to the **Specify Software Images** dialog box.
- Step 8** The **Vrf** column indicates the name of the virtual routing and forwarding (VRF).
- Step 9** In the **Available Space** column, specify the available space for the **Primary Supervisor** and **Secondary Supervisor** modules of the switch.
- Available Space** column shows the available memory in MB on the switch (for less than 1 MB, it is shown and marked as KB).
- Bootflash browser shows the filename, size, and last modified date for all the files and directories on the switch bootflash. You can delete files by selecting them and clicking **Delete** to increase the available space on the switch.
- Step 10** **Selected Files Size** column shows the size of images that are selected from the server.
- If the total size of selected images is greater than available space on a switch, the file size is marked in red. We recommend that you create more space on the switch to copy images to it and install.
- Step 11** Drag and drop the switches to reorder the upgrade task sequence.
- Step 12** Select **Skip Version Compatibility** if you are sure that the version of the Cisco NX-OS software on your device is compatible with the upgraded images that you have selected.
- Step 13** Select **Select Parallel Line Card upgrade** to upgrade all the line cards at the same time.
- Upgrading a parallel line card is not applicable for Cisco MDS devices.
- Step 14** Select **Options** under the **Upgrade Options** column to choose the type of upgrade.
- Upgrade Options** window appears with two upgrade options. The drop-down list for **Upgrade Option 1** has the following options:
- **NA**
  - **bios-force**
  - **non-disruptive**

**NA** is the default value.

The drop-down list for **Upgrade Option 2** has the following options:

- NA
- bios-force

When **NA** is selected under **Upgrade Option 1**, **Upgrade Option 2** is disabled.

When **bios-force** is selected under **Upgrade Option 1**, **NA** is the only option under **Upgrade Option 2**.

When **non-disruptive** is selected under **Upgrade Option 1**, you can choose **NA** or **bios-force** under **Upgrade Option 2**.

Check the **Use this Option for all other selected devices** check box to use the selected option for all the selected devices and click **OK**.

**Note**

- The upgrade options are applicable only for Cisco Nexus 3000 Series and 9000 Series switches.
- Selecting the non-disruptive option for upgrading does not ensure a non-disruptive upgrade. Perform a compatibility check to ensure that the device supports non-disruptive upgrade.

**Step 15** Click **Next**.

If you did not select **Skip Version Compatibility**, the Cisco DCNM performs a compatibility check.

You can choose to wait until the check is complete or click **Finish Installation Later**.

The installation wizard is closed and a compatibility task is created in **Control > Image Management > Install & Upgrade > Upgrade History** tasks.

The time that is taken to check the image compatibility depends on the configuration and the load on the device.

The **Version Compatibility Verification** status column displays the status of verification.

If you skip the version compatibility check by choosing **Skip Version Compatibility**, Cisco DCNM displays only the name of the device, the **Current Action** column displays **Completed**, and the **Version Compatibility Verification** column displays **Skipped**.

**Step 16** Click **Finish Installation Later** to perform the upgrade later.

**Step 17** Click **Next**.

**Step 18** Check the check box to save the running configuration to the startup configuration before upgrading the device.

**Step 19** You can schedule the upgrade process to occur immediately or later.

1. Select **Deploy Now** to upgrade the device immediately.
2. Select **Choose time to Deploy** and specify the time in MMM/DD/YYYY HH:MM:SS format to perform the upgrade later.

This value is relative to the server time. If the selected time to deploy is in the past, the job is executed immediately.

**Step 20** You can choose the execution mode based on the devices and the line cards you have chosen to upgrade.

1. Select **Sequential** to upgrade the devices in the order in which they were chosen.
2. Select **Concurrent** to upgrade all the devices at the same time.

**Step 21** Click **Finish** to begin the upgrade process.

The Installation wizard closes and a task to Upgrade is created on the **Control > Image Management > Install & Upgrade > Upgrade History** page.

---

### What to do next

After you complete the ISSU on the switch, ensure that you wait for 20 minutes to allow the switch to reboot, and stabilize the SNMP agent. Cisco DCNM will discovery polling cycles in order to display the new version of the switch on the Cisco DCNM Web UI.

## Finish Installation

You can choose to complete the installation for tasks which was completed on the **Compatibility Check** page. Perform the following task to complete the upgrade process on the devices.

### Procedure

---

**Step 1** Choose **Control > Image Management > Install & Upgrade > Upgrade History**, select a task for which the compatibility check is complete.

Select only one task at a time.

**Step 2** Click **Finish Installation**.

**Software Installation Wizard** appears.

**Step 3** Check the check box to save the running configuration to the startup configuration before upgrading the device.

**Step 4** Check the check box to put a device in maintenance mode before upgrade. This option is valid only for the devices that support maintenance mode.

**Step 5** You can schedule the upgrade process to occur immediately or later.

1. Select **Deploy Now** to upgrade the device immediately.

2. Select **Choose time to Deploy** and specify the time in DD/MM/YYYY HH:MM:SS format to perform the upgrade later.

**Step 6** You can choose the execution mode that is based on the devices and the line cards that you have chosen to upgrade.

1. Select **Sequential** to upgrade the devices in the order in which they were chosen.

2. Select **Concurrent** to upgrade the devices at the same time.

**Step 7** Click **Finish** to complete the upgrade process.

---

## View

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

**Delete****Procedure**

- Step 1** Choose **Control > Image Management > Install & Upgrade > Upgrade History**, check the task ID check box.  
Select only one task at a time.
- Step 2** Click **View**.  
The **Installation Task Details** window is displayed.
- Step 3** Click **Settings**. Select **Columns** and choose the column details options.  
This window displays the location of the kickstart and system images, compatibility check status, installation status, descriptions, and logs.
- Step 4** Select the device.  
The detailed status of the task is displayed. For the completed tasks, the response from the device is displayed.  
If the upgrade task is in progress, a live log of the installation process appears.
- Note** This table is refreshed every 30 secs for jobs in progress, when you are on this window.

**Delete**

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

**Procedure**

- Step 1** Choose **Control > Image Management > Install & Upgrade > Upgrade History**, and check the **Task ID** check box.
- Step 2** Click **Delete**.
- Step 3** Click **OK** to confirm deletion of the job.

**Switch Level History**

You can view the history of the upgrade process at a switch level. You can view the current version of the switch and other details.

The following table describes the fields that appear on **Control > Image Management > Install & Upgrade > Switch Level History**.

Field	Description
Switch Name	Specifies the name of the switch
IP Address	Specifies the IP Address of the switch
Platform	Specifies the Cisco Nexus switch platform

Field	Description
Current Version	Specifies the current version on the switch software

Click the radio button next to a switch name to select the switch and view its upgrade history. Click **View** to view the upgrade task history for the selected switch.

The following table describes the fields that appear on **Control > Image Management > Install & Upgrade > Switch Level History > View Device Upgrade Tasks**:

Field	Description
Owner	Specifies the owner who initiated the upgrade.
Job Status	Specifies the status of the job. <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 menu includes the following submenus:

### Endpoint Locator

The Endpoint Locator (EPL) feature allows real-time tracking of endpoints within a data center. The tracking includes tracing the network life history of an endpoint and getting insights into the trends that are associated with endpoint additions, removals, moves, and so on. An endpoint is anything with at least one IP address and MAC address. In that sense, an endpoint can be a virtual machine (VM), container, bare-metal server, service appliance and so on.

**Important**

- EPL is supported for VXLAN BGP EVPN fabric deployments only in the DCNM LAN fabric installation mode. The VXLAN BGP EVPN fabric can be deployed as Easy fabric, Easy eBGP fabric, or an External fabric (managed or monitored mode). EPL is not supported for 3-tier access-aggregation-core based network deployments.
- EPL displays endpoints that have at least one IP address (IPv4 and/or IPv6). Also, these endpoints must be residing in networks where the gateway or SVI is configured on the network switches within the VXLAN EVPN fabric. In other words, EPL cannot determine the identity (IPv4/IPv6 address) of the endpoints for networks that are deployed as Layer-2 Only within the fabric.

EPL relies on BGP updates to track endpoint information. Hence, typically the DCNM needs to peer with the BGP Route-Reflector (RR) to get these updates. For this purpose, IP reachability from the DCNM to the RR is required. This can be achieved over in-band network connection to the DCNM eth2 interface.

Some key highlights of the Endpoint Locator are:

- Support for dual-homed and dual-stacked (IPv4 + IPv6) endpoints
- Support for up to two BGP Route Reflectors
- Support real-time and historical search for all endpoints across various search filters such as VRF, Network, Layer-2 VNI, Layer-3 VNI, Switch, IP, MAC, port, VLAN, and so on.
- Support for real-time and historical dashboards for insights such as endpoint lifetime, network, endpoint, VRF daily views, and operational heat map.
- Support for high availability
- Support for endpoint data that is stored for up to 180 days, amounting to a maximum of 5 G storage space.
- Support for optional flush of the endpoint data in order to start afresh.
- Supported scale: 10K endpoints

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

## Configuring Endpoint Locator

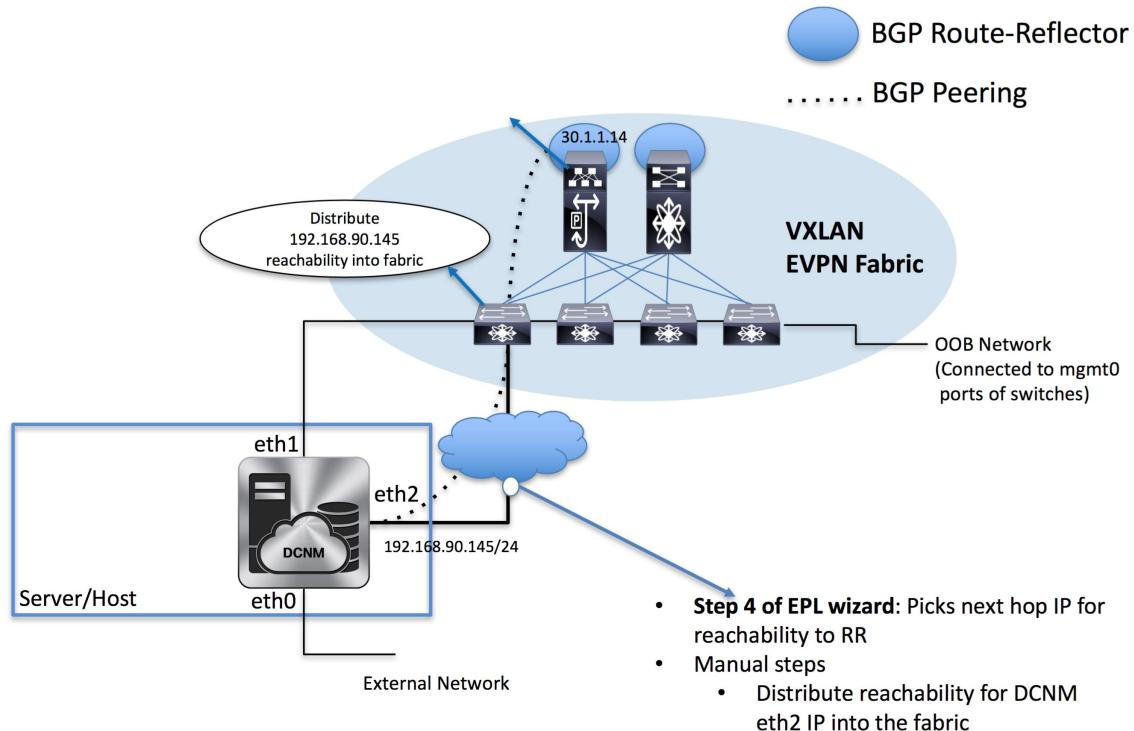
The DCNM OVA or the ISO installation comes with 3 interfaces—eth0 interface for external access to the DCNM, eth1 interface that is used primarily for fabric management, and eth2 interface for in-band network connectivity to Cisco DCNM. In most deployments the eth1 interface is part of the same network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

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

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

# Configuration

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



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.

## Configuring Endpoint Locator

The screenshots illustrate the configuration of an endpoint locator (EP) on a Cisco Data Center Network Manager (DCNM) interface. Both screenshots show the 'Edit Configuration' dialog for a selected interface.

**Top Screenshot (Trunk Host Policy):**

- Name:** leaf1-Ethernet1/24
- Policy:** trunk\_host (selected)
- Access Host Options:**
  - Enable spanning-tree bpdu guard
  - Enable spanning-tree edge port behavior
  - MTU for the interface (jumbo selected)
  - Trunk Allowed Vlans (none selected)
- Interface Admin State:**

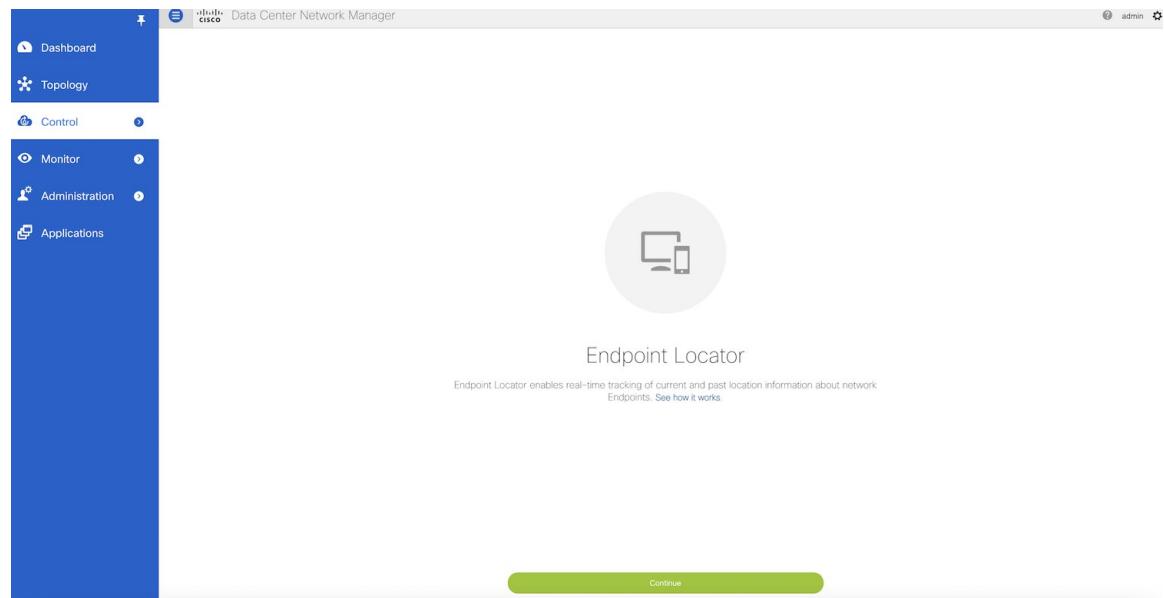
**Bottom Screenshot (ep1\_routed\_intf Policy):**

- Name:** leaf1-Ethernet1/24
- Policy:** ep1\_routed\_intf (selected)
- General Settings:**
  - Interface IP:** 192.168.94.1
  - IP Netmask Length:** 24
  - LS Routing Protocol:** ospf (selected)
  - Link-State Routing Tag:** UNDERLAY
  - Interface Admin State:**

## Procedure

### Step 1

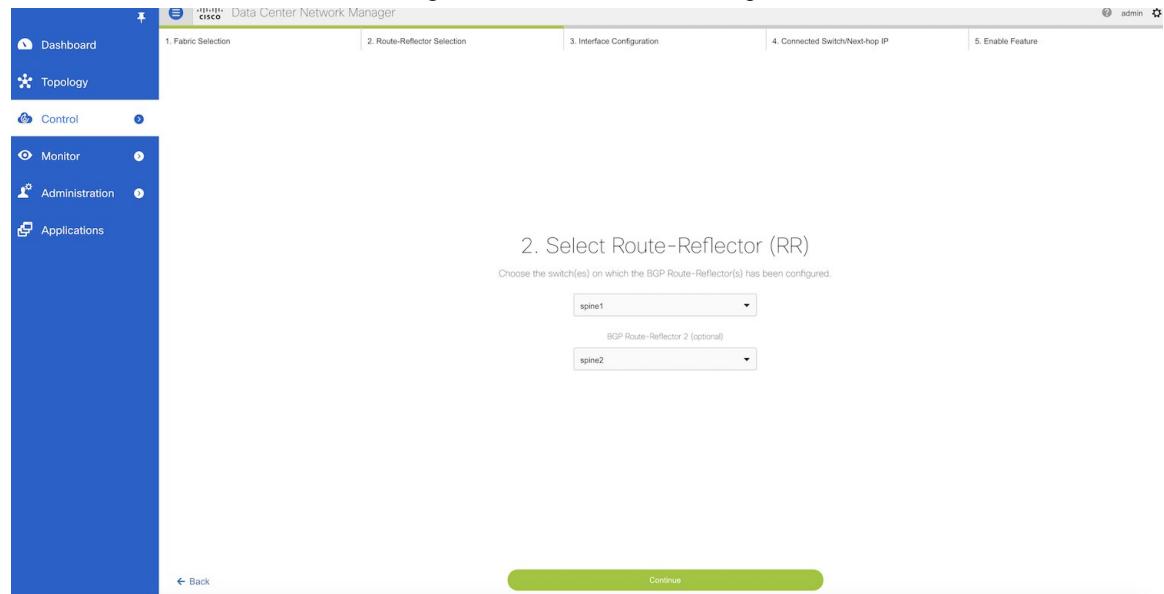
From the Cisco DCNM home page, choose **Control > Endpoint Locator > Configure**. The **Endpoint Locator** page appears with a [See how it works](#) help link.



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

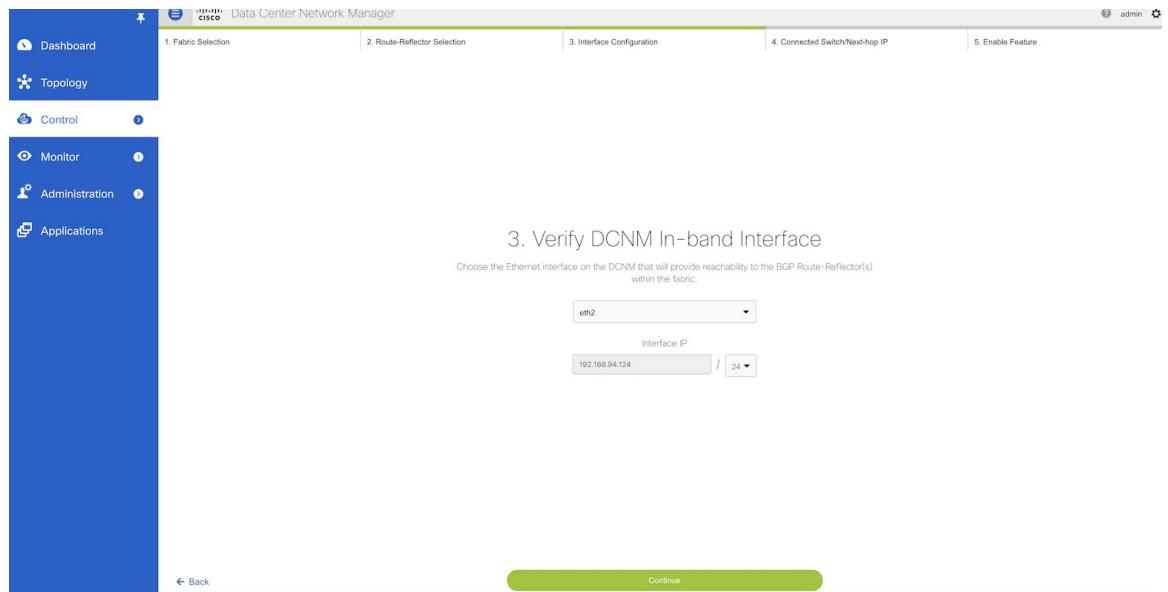
**Step 3** Select the appropriate fabric on which the endpoint locator feature should be enabled to track endpoint activity. EPL can only be enabled for one fabric. It can be DFA or EVPN.

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



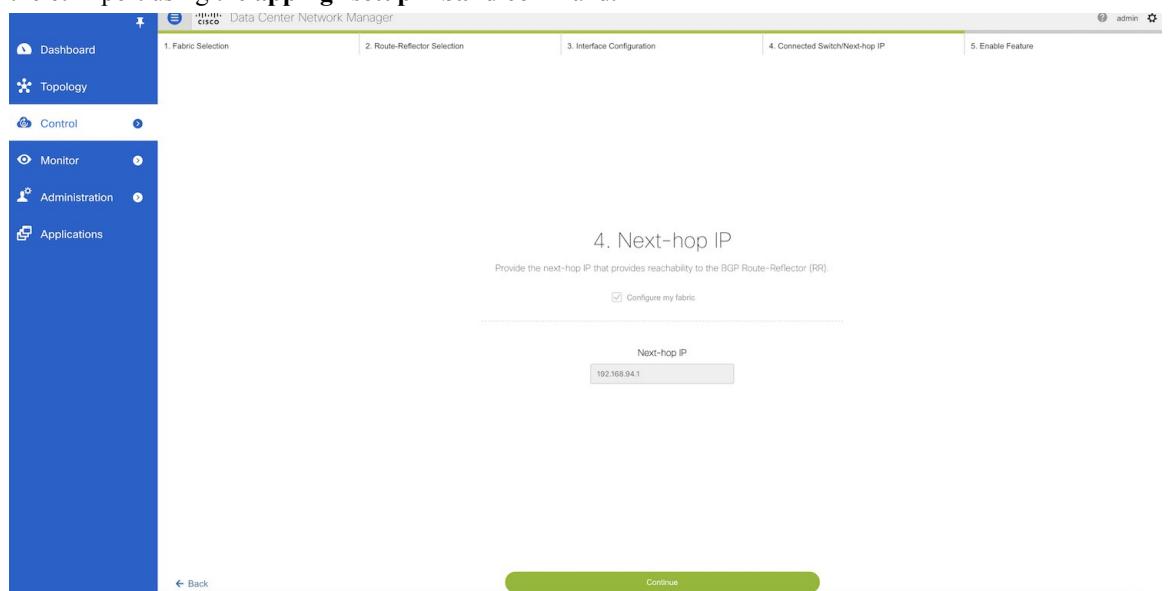
**Step 5** Check DCNM eth2 configuration for IP reachability to the RR.

## Configuring Endpoint Locator



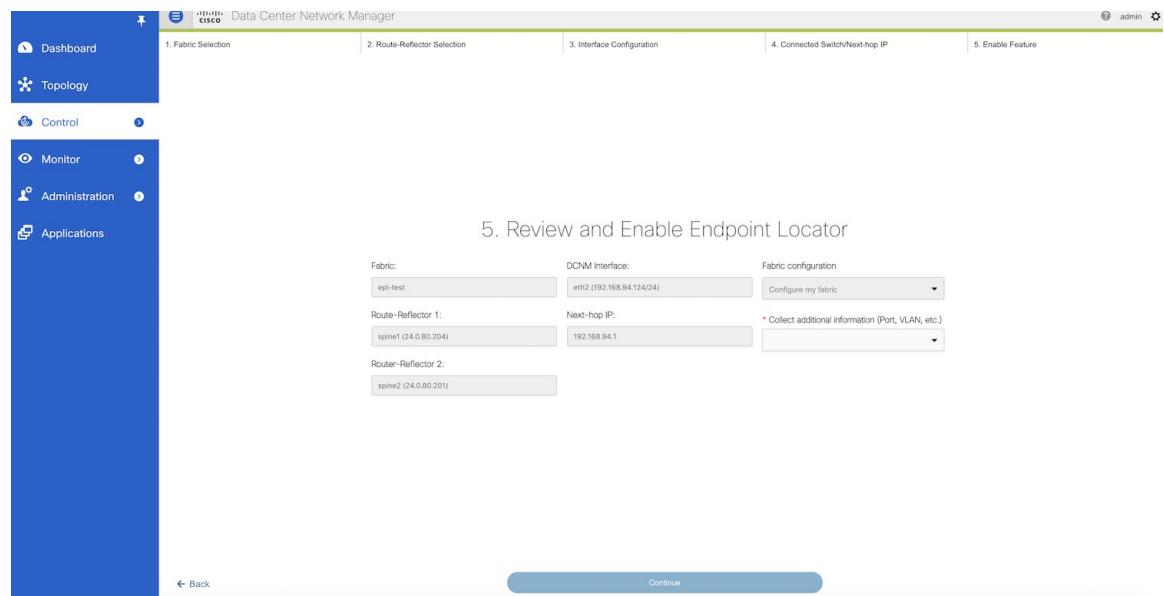
### Step 6

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.

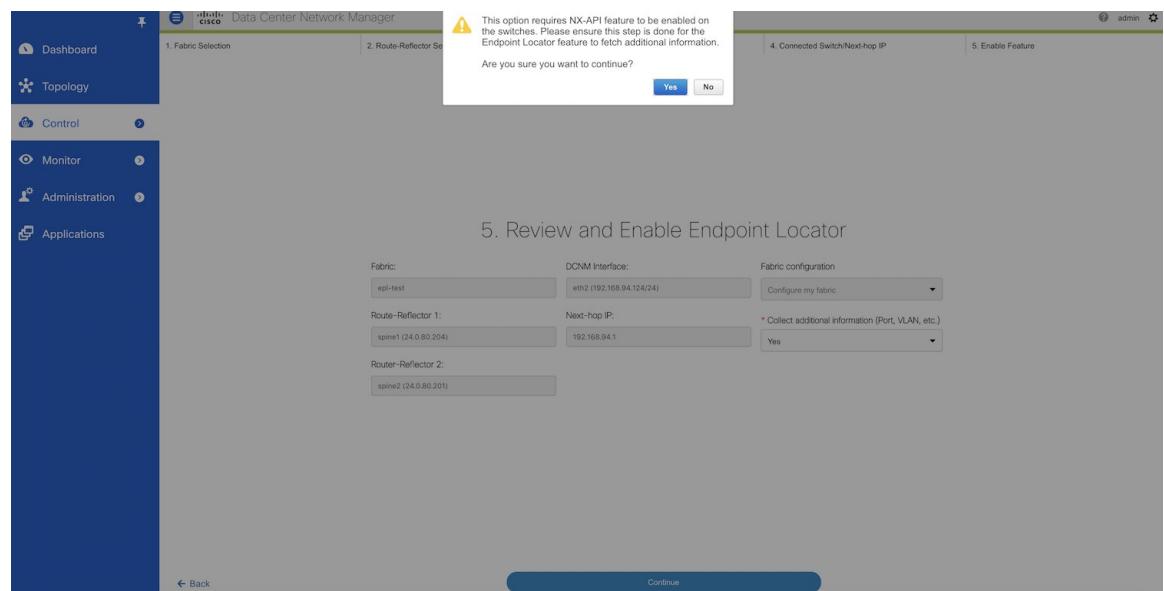


### Step 7

The last step provides a summary of the information entered in the previous steps. The wizard view allows navigation to any particular step whereby one can make necessary changes/edits. In this step, one must specify whether additional information such as PORT, VLAN, VRF etc. is required when enabling the EPL feature. If the **No** option is selected, then this information will not be collected and reported by EPL.

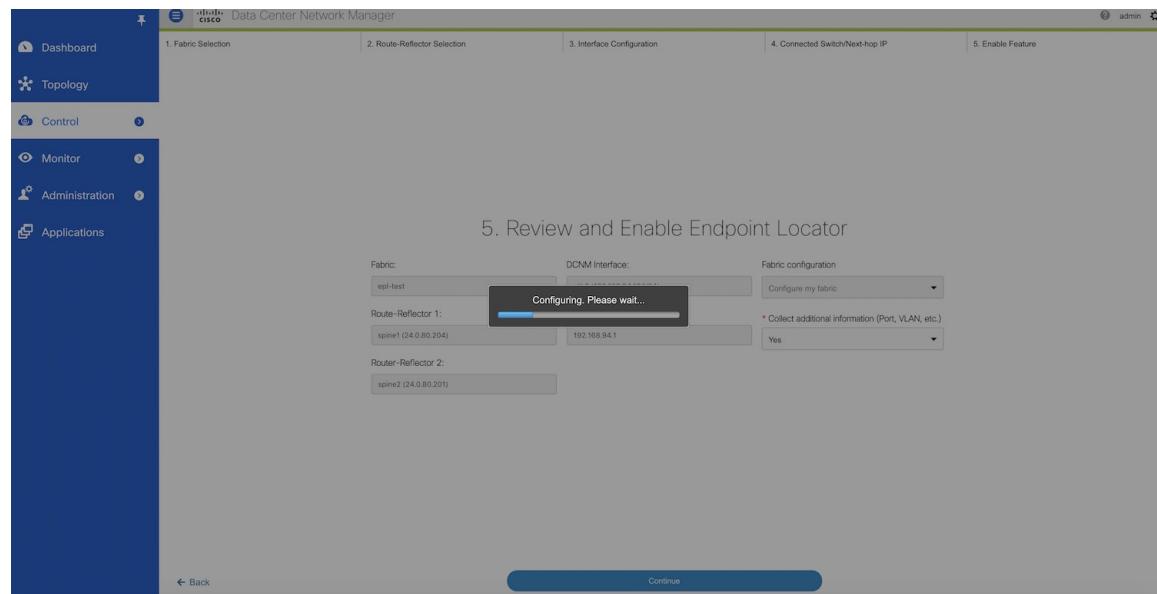


However, if the **Yes** option is selected in the drop down, a warning pop-up appears that feature NX-API must be supported and enabled on the switches/ToRs/leafs to gather this information. Otherwise this additional information cannot be fetched or reported.

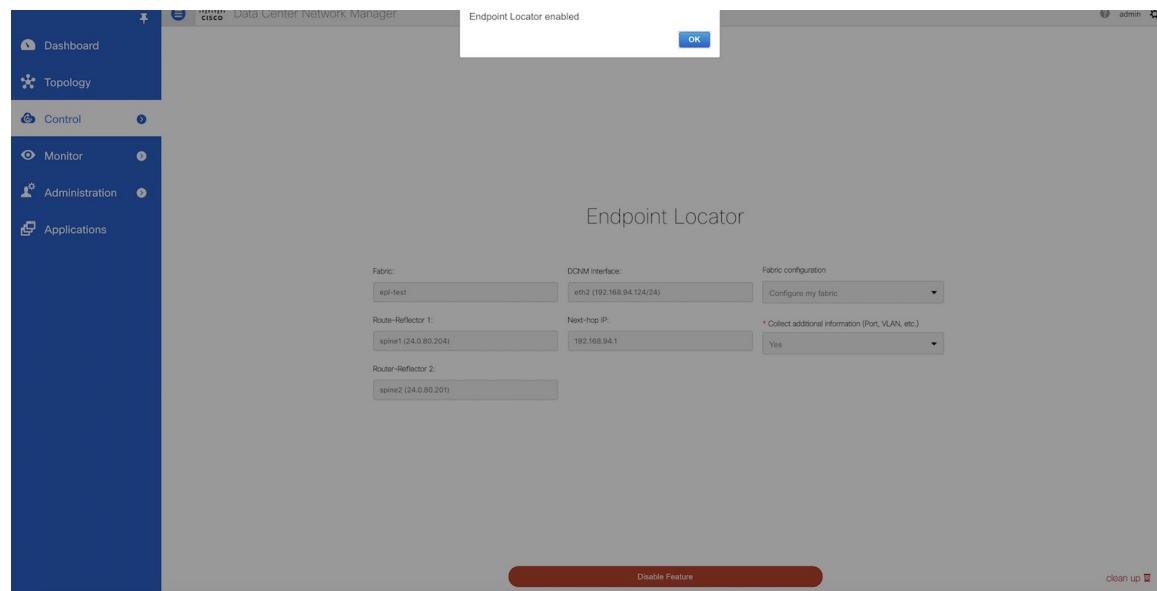


- Step 8** Once the appropriate selections are made and various inputs have been reviewed, click **Continue** to enable EPL. A progress bar will appear indicating the status of the EPL feature enable process.

## Configuring Endpoint Locator



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

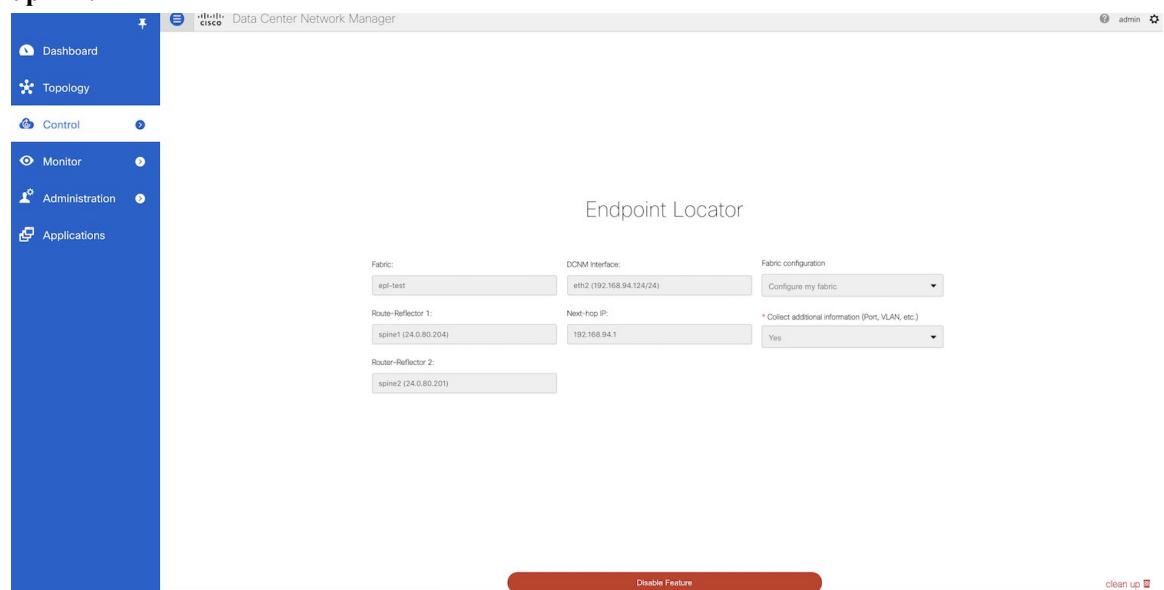
dashboard that depicts operational and exploratory insights into the endpoints that are present in the fabric. For more information, refer to *Section Exploring Endpoint Locator Details*.

## Flushing the Endpoint Database

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

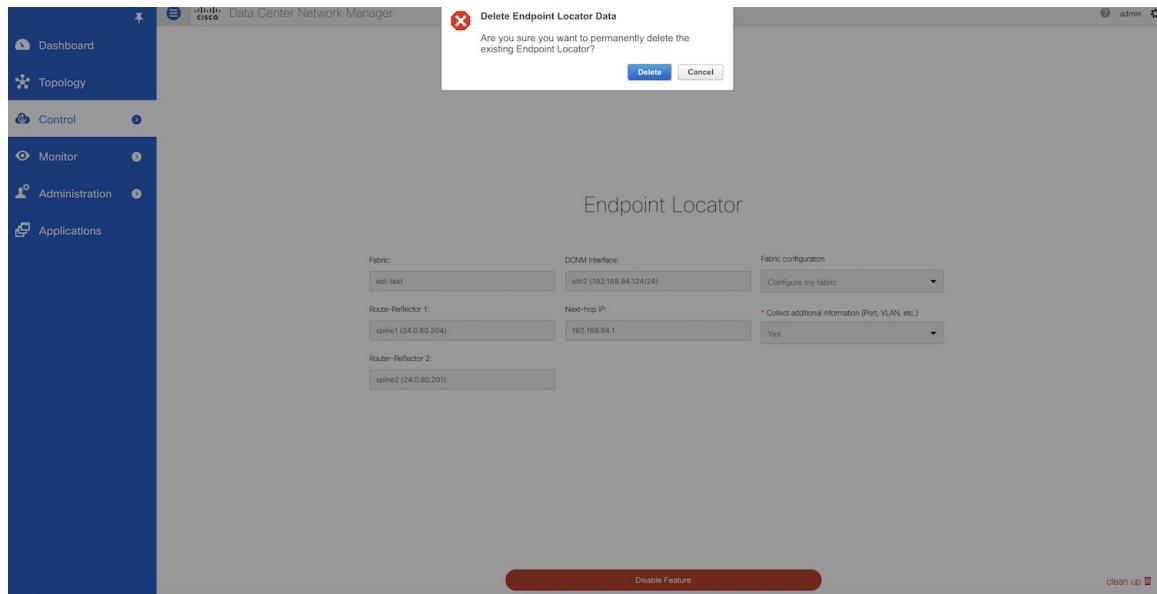
### Procedure

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



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

## **Configuring the In-band Port**



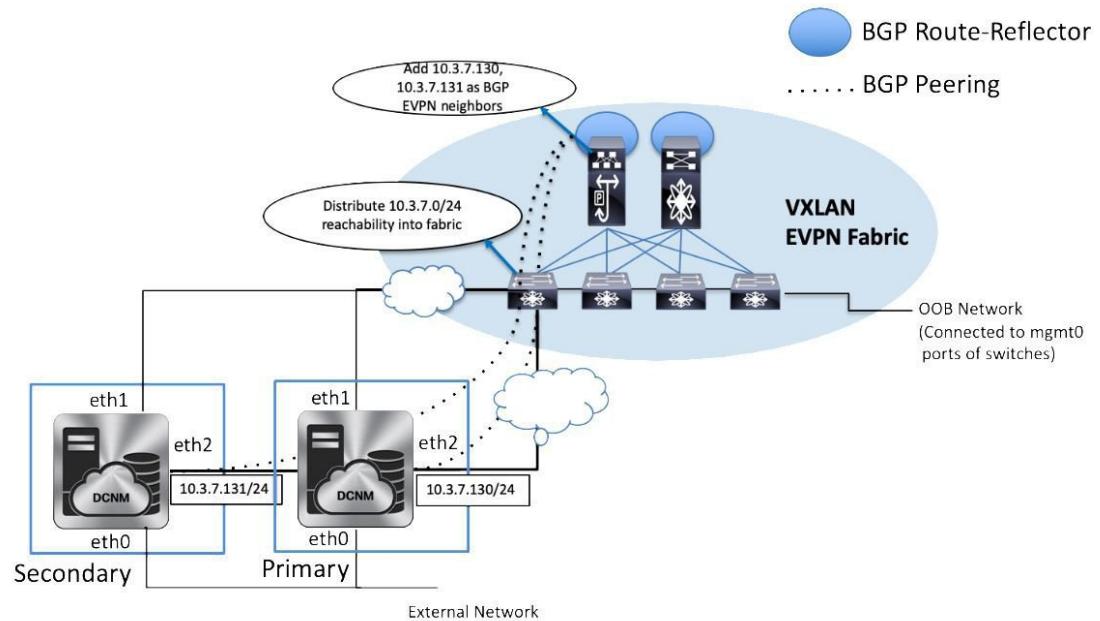
- Step 2** Click **Delete** to continue or **Cancel** in case the user wants to abort.

## Configuring the In-band Port

## **Procedure**

To configure the in-band port in Cisco DCNM, enter the **appmgr setup inband** command. See the example below.

## Configuring Endpoint Locator in DCNM High Availability Mode



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

On Cisco DCNM Primary appliance:

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

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

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

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

On Cisco DCNM Secondary appliance:

```
[root@dcnm-secondary]# appmgr setup inband
```

## Configuring Endpoint Locator in DCNM High Availability Mode

```

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

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

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

[root@dcnm-secondary]#

```

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

### Procedure

---

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

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

**Step 2** In the Select a fabric to configure endpoint locator in DCNM HA mode.

**Step 3** Click **Continue**.

**Step 4** Select one or two Route-Reflectors (RRs).

**Step 5** Click **Continue**.

**Step 6** Verify the Ethernet interfaces on both primary and standby DCNM nodes.

**Step 7** Click **Continue**.

**Step 8** Verify the next-hop IP address on the primary and standby DCNM.

Note that the next-hop IP corresponds to the eth2 gateway which should be the same on both the DCNMs.

**Step 9** Click **Continue**.

**Step 10** After selecting the NX-API enable or disable option and verifying the other information provided in the prior steps, click **Continue**.

---

### What to do next

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

## Adding High Availability Node to Endpoint Locator Configuration

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

### Procedure

---

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

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

- Step 2** Click the **Add HA node** link.

- Step 3** In the **Configure Standby DCNM Interface** page, choose the Ethernet interface on DCNM that provides reachability to the BGP Route-Reflectors within the fabric.

- Step 4** Click **Continue**.

- Step 5** In the Next-Hop page check the value of the next-hop IP.

- Step 6** Click **Configure HA Node**.

The configuration details are displayed on the Endpoint Locator page.

---

## Configuring Endpoint Locator for External Fabrics

In addition to Easy fabrics, DCNM Release 11.2(1) allows you to enable EPL for VXLAN EVPN fabrics comprising of switches that are imported into the external fabric. The external fabric can be in managed mode or monitored mode, based on the selection of **Fabric Mode Monitor Flag** in the **External Fabric Settings**. In case the monitor or read-only fabric option is selected for the fabric, while enabling EPL, the **Configure my fabric** option must be unchecked; because, the EPL neighborship is added to the spines or RRs via some other means.

## Disabling Endpoint Locator

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

### Procedure

---

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

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

- Step 2** Click **Disable Feature**.
- 

## Troubleshooting Endpoint Locator

There may be multiple reasons why enabling the Endpoint Locator feature may fail. Typically, if the appropriate devices are selected and the IP addresses to be used are correctly specified, the connectivity of the DCNM to the BGP RR may not be present due to which the feature cannot be enabled. This is a sanity check that is

present to ensure that basic IP connectivity is available. The following image shows an example error scenario that was encountered during an attempt to enable the EPL feature.

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

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

```
#tail -f epl.log
2017.04.08 07:47:05 INFO [epl] Running script: [/sbin/appmgr, status, epls]
2017.04.08 07:47:05 INFO [epl] Received response:
2017.04.08 07:47:05 INFO [epl]
2017.04.08 07:47:05 INFO [epl] >>> Sat Apr 8 07:47:05 PDT 2017
2017.04.08 07:47:05 INFO [epl] appmgr status epls
2017.04.08 07:47:05 INFO [epl]
2017.04.08 07:47:05 INFO [epl] EPLS is stopped...
2017.04.08 07:47:08 INFO [epl] Running command: ifdown eth2
2017.04.08 07:47:08 INFO [epl] Received response:
2017.04.08 07:47:08 INFO [epl] EPL disabled successfully
2017.04.08 08:00:06 INFO [epl] Enabled 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 successfully
```

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

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

The EPL relies on BGP updates to get endpoint information. In order for this to work, the switch loopback or VTEP interface IP addresses must be discovered on the DCNM for all switches that have endpoints. To validate, navigate to the Cisco DCNM Web UI > Dashboard > Switch > Interfaces tab, and verify if the IP

address and the prefix associated with the corresponding Layer-3 interfaces (typically loopbacks) are displayed correctly.

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

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

## Streaming Telemetry for LAN Deployments

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

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

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

## Guidelines and Recommendations

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

## Pre-Requisites for Enabling the LAN Telemetry Feature

- The Cisco Nexus 9000 switches and Cisco DCNM need to be time synchronized (NTP is recommended).
- Minimum software version on the Nexus 9000 switches must be 7.0(3)I6(1) or higher.

## Enabling the Streaming Telemetry Feature

- 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**

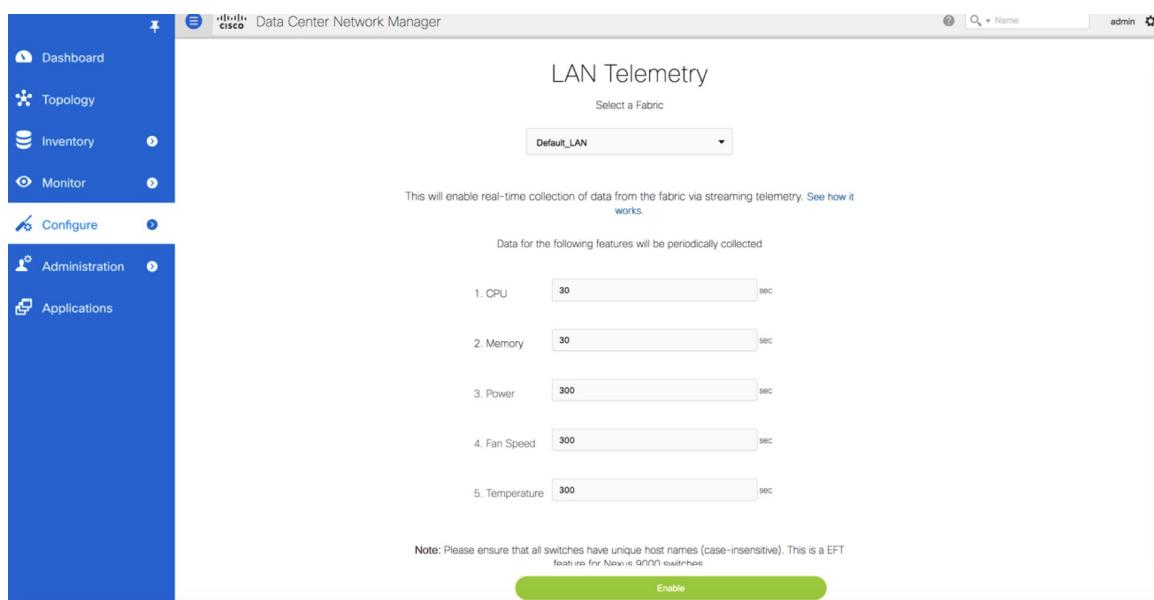
If the preceding configurations are unavailable on the switches, the telemetry health on Cisco DCNM does not show the configurations and the connection status for the telemetry-enabled switches. The preceding commands can be manually defined in a new template, and then pushed to all the switches in the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

## Enabling the Streaming Telemetry Feature

### Procedure

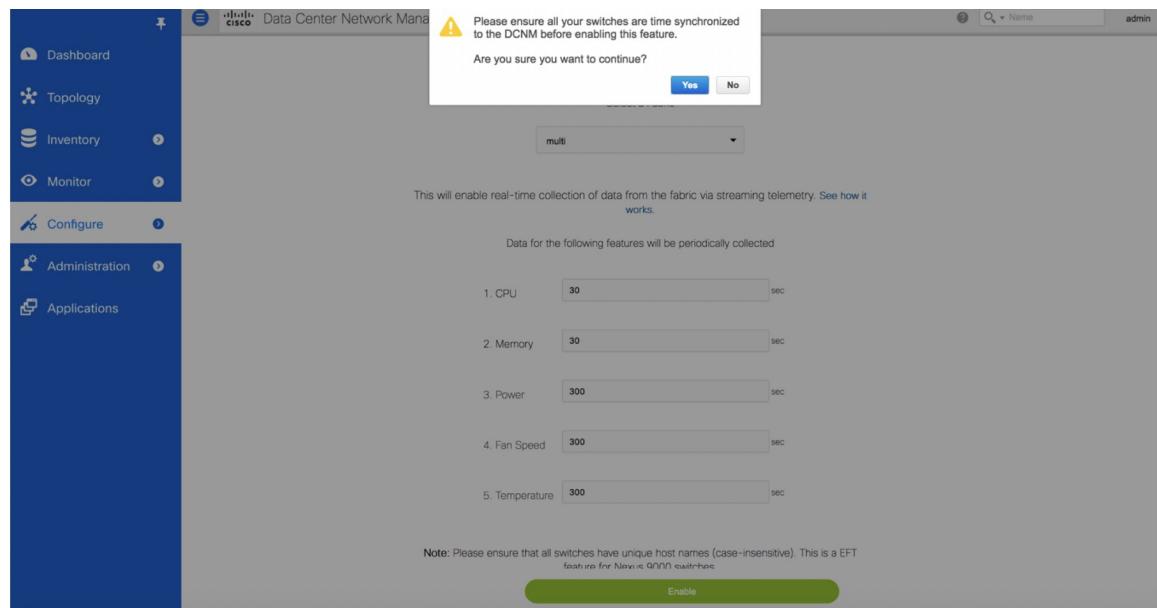
**Step 1**

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.

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:

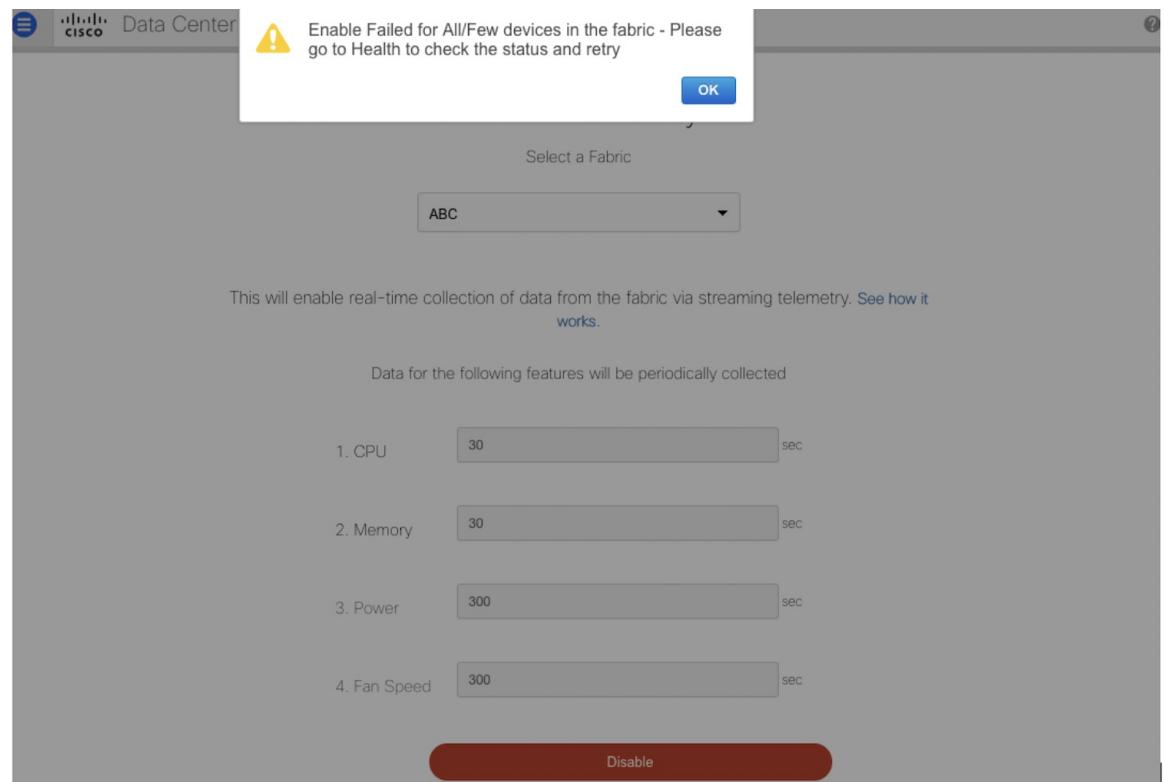
### Procedure

**Step 1**

Choose Control > LAN Telemetry > Health.

Name	Description	Additional Information	Update Period (seconds)	Data Received	Configuration Status
leaf3	N9K-C9396PX NXOS 7.0...	SAL18432P4S Default_L...		94.6 MB	<span style="color: green;">✓ SUCCESS</span>
fan	Fan Speed		300	515.7 KB	<span style="color: green;">✓ SUCCESS</span>
cpu	Per Process CPU Utilization		30	50.8 MB	<span style="color: green;">✓ SUCCESS</span>
resources	Overall System Resource...		30	1.5 MB	<span style="color: green;">✓ SUCCESS</span>
mem	Memory Utilization		30	41.3 MB	<span style="color: green;">✓ SUCCESS</span>
temp	Switch Temperature Data		300	319.8 KB	<span style="color: green;">✓ SUCCESS</span>
power	Power Consumption		300	255.8 KB	<span style="color: green;">✓ SUCCESS</span>
n9k-bg1	N9K-C93180YC-EX NXO...	FDO210721L3 Default_L...		167.0 MB	<span style="color: green;">✓ SUCCESS</span>
n9k-bg2	N9K-C93180YC-EX NXO...	FDO210705NY Default_L...		164.1 MB	<span style="color: green;">✓ SUCCESS</span>
spine1	N9K-C9396PX NXOS 7.0...	SAL1833YM11 Default_L...		148.4 MB	<span style="color: green;">✓ SUCCESS</span>
spine2	N9K-C9396PX NXOS 7.0...	SAL18422FUR Default_L...		144.1 MB	<span style="color: green;">✓ SUCCESS</span>
leaf1	N9K-C9396PX NXOS 7.0...	SAL18432P4X Default_L...		130.8 MB	<span style="color: green;">✓ SUCCESS</span>
leaf2	N9K-C9396PX NXOS 7.0...	SAL18432P5Q Default_L...		127.9 MB	<span style="color: green;">✓ SUCCESS</span>

When Telemetry is enabled or disabled, there is a chance that enabling or disabling can fail in some or all the switches. When that happens, a pop-up similar to the following screen appears.



There are two possible options:

1. You can go to the Health page, and retry the configuration for those switches that failed. When a configuration cannot be applied or removed on any switch, **Configuration Status** in the health page, appears as FAILED. Upon clicking the ‘FAILED’ link, a pop-up would show the reason for the failure. After you correct the failure, the configuration can be retried by clicking on the retry button appearing next to the Configuration Status for every switch. The screen-shot for that is also shown below.

Name	Description	Additional Information	Update Period (seconds)	Data Received	Configuration Status
n9k-101	N9K-C9396PX NXOS 7.0(...	SAL1833YM8Q ABC 10.1...		301.5 KB	<span style="color: green;">✓ SUCCESS</span>
n9k-102	N9K-C9396PX NXOS 7.0(...	SAL18422FU6 ABC 10.10....		111.4 KB	<span style="color: red;">✗ FAILED</span>

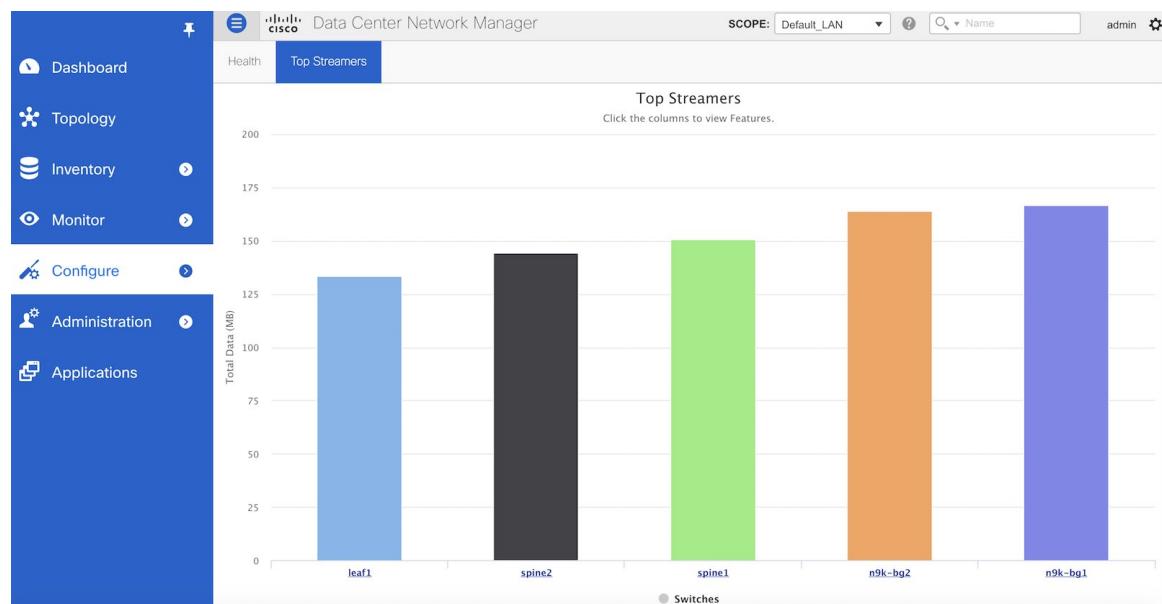
2. You can stay in the main **Telemetry > Configure** page. It would display a dialogue box with the failed message. Then you can reverse the configuration for the successfully configured switches. In other words:
  - When “Enable” fails for some or all switches, the screen has a Red button with “disable” option. This means, for those switches, wherein enabling Telemetry was successful, you can disable Telemetry

## Telemetry Streaming Interface

on those switches. If “Enable” failed on all switches, you will still see the Red button with “disable” option. Clicking on “disable” is a no-op. In both the cases, you will see the green button with the “enable” option in a few seconds after disabling is completed. This removes the “retry” option from the health page since you want to “disable” telemetry and there is nothing to retry.

- Similarly, when “Disable” fails for some or all switches, the screen has a Green button with “Enable” option. This means, for those switches, wherein disabling Telemetry was successful, you can Enable Telemetry on those switches. If “Disable” failed on all switches, you will still see the Green button with “Enable” option. Clicking on “Enable” is a no-op. In both the cases, you will see the Red button with the “Disable” option in a few seconds after Enabling is completed. Doing this, removes the “retry” option from the health page since you want to “enable” telemetry and there’s nothing to retry.

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



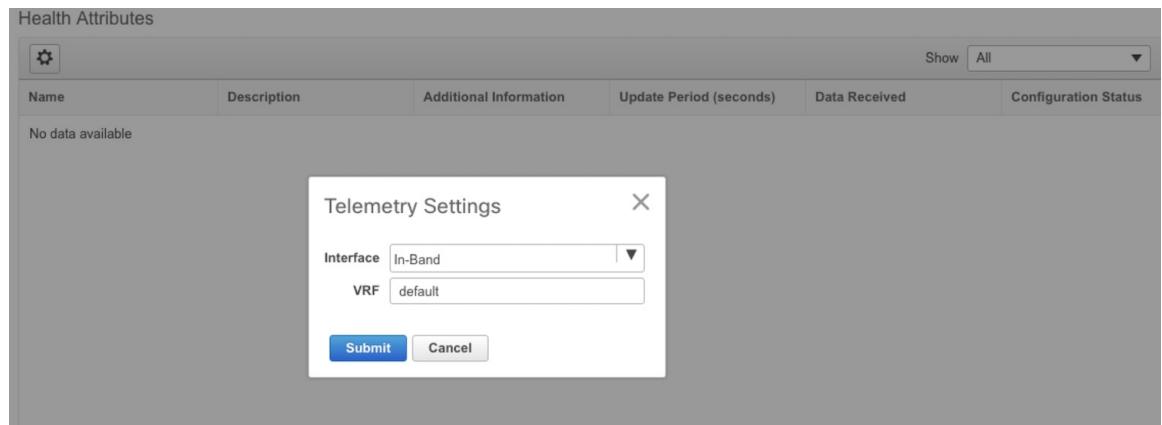
## Telemetry Streaming Interface

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

### Procedure

- Step 1** Disable Telemetry on all the Enabled fabrics.

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



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.

---





## CHAPTER 5

# Monitor

---

This chapter contains the following topics:

- [Inventory, on page 189](#)
- [Monitoring Switch, on page 208](#)
- [Monitoring LAN, on page 211](#)
- [Monitoring Endpoint Locator, on page 215](#)
- [LAN Telemetry, on page 222](#)
- [Alarms, on page 230](#)

# Inventory

This chapter contains the following topics:

## Viewing Inventory Information for Switches

To view the inventory information for switches from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

**Step 1** Choose **Monitor > Inventory > Switches**.

The **Switches** window with a list of all the switches for a selected Scope is displayed.

**Step 2** You can also view the following information.

- **Group** column displays the switch group to which the switch belongs.
- In the **Device Name** column, select a switch to display the Switch Dashboard.
- **IP Address** column displays the IP address of the switch.
- **WWN/Chassis ID** displays the Worldwide Name (WWN) if available or chassis ID.
- **Health** displays the health situation of the switch.

**Note** To refresh and recalculate the latest health data for all the switches on Cisco DCNM, click the **Recalculate Health** button above the switches table.

- **Status** column displays the status of the switch.
- **# Ports** column displays the number of ports.
- **Model** column displays the model name of the switch.
- **Serial No.** column displays the serial number of the switch.
- **Release** column displays the switch version.
- **License** column displays the DCNM license that is installed on the switch.
- **Up Time** column displays the time period for which the switch is active.

**Step 3** In the **Health** column, the switch health is calculated by the capacity manager based on the following formula in the server.properties file.

The function to implement is:

```
# calculate(x, x1, y, y1, z).
# @param x: Total number of modules.
# @param x1: Total number of modules in warning.
# @param y: Total number of switch ports.
# @param y1: Total number of switch ports in warning.
# @param z: Total number of events with severity of warning or above.
```

**Step 4** The value in the **Health** column is calculated based on the following default equation.

$$((x-x1)*1.0/x) *0.4 + ((y-y1)*1.0/y)*0.3 + ((z*1.0/1000>=1) ? 0: ((1000-z)*1.0/1000)*0.3).$$

In the above formula, the switch health value is calculated based on the following:

- Percentage of Warning Modules (Contributes 40% of the total health).
- Percentage of Warning Ports (Contributes 30% of the total health).
- Percentage of events with severity of Warning or above (Contributes 30% of the total health. If there are more than 1000 warning events, the event health value is 0).

You may also have your own health calculation formula by implementing the common interface class: com.cisco.dcbu.sm.common.rif.HealthCalculatorRif. Add the .jar file to the DCNM server and modify the health.calculator property to point to the class name you have created.

The default Java class is defined as: health.calculator=com.cisco.dcbu.sm.common.util.HealthCalculator.

- Capacity Manager calculates health only for the license switches. If the health column does not display a value, the switch either does not have a license or it has missed the capacity manager daily cycle.
- If the switch is unlicensed, click **Unlicensed** in the DCNM License column. The **Administration > License** window appears which allows you to assign a license to the user.
- The capacity manager runs two hours after the DCNM server starts. So, if you discover a device after two hours of the DCNM start time, the health will be calculated 24 hours after this DCNM start time

## Viewing System Information

The switch dashboard displays the details of the selected switch.

### Procedure

---

- Step 1** From the Cisco DCNM home page, choose **Monitor > Inventory > Switches**.

An inventory of all the switches that are discovered by Cisco Prime DCNM Web Client is displayed.

- Step 2** Click a switch in the **Device Name** column.

The **Switch** dashboard that corresponds to that switch is displayed along with the following information:

- Step 3** Click the **System Information** tab. This tab displays detailed system information such as group name, health, module, time when system is up, serial number, the version number, contact, location, DCNM license, status, system log sending status, CPU and memory utilization, and VTEP IP address are displayed. Click **Health** to access the Health score screen, which includes health score calculation and health trend. The popup contains Overview, Modules, Switch Ports, and Events tabs.

- (Optional) Click **SSH** to access the switch through Secure Shell (SSH).
  - (Optional) Click **Device Manager** to view a graphical representation of a Cisco MDS 9000 Family switch chassis, a Cisco Nexus 5000 Series switch chassis, a Cisco Nexus 7000 Series switch chassis, or a Cisco Nexus 9000 Series switch chassis including the installed switching modules, the supervisor modules, the status of each port within each module, the power supplies, and the fan assemblies.
  - (Optional) Click **HTTP** to access the switch through Hypertext Transfer Protocol (HTTP) for that switch.
  - (Optional) Click **Accounting** to go to the [Viewing Accounting Information, on page 210](#) window pertaining to this switch.
  - (Optional) Click **Backup** to go to the [Viewing a Configuration window](#).
  - (Optional) Click **Events** to go to the [Viewing Events Registration, on page 248](#) window.
  - (Optional) Click **Show Commands** to display the device show commands. The Device Show Commands page helps you to view commands and execute them.
  - (Optional) Click **Copy Running Config to Startup Config** to copy the running configuration to the startup configuration.
- 

## VXLAN

You can view VXLANS and their details under the **VXLAN** tab.

To view VXLANS, choose **Inventory > View > Switches**, and then click a switch name in the **Device Name** column.

The following table describes the fields that are displayed:

**Table 3: The VXLAN Tab**

Field	Description
VNI	Displays the Layer 2 (network) or Layer 3 (VRF) VXLAN VNI that is configured on a switch.
Multicast address	Displays the multicast address that is associated with the Layer 2 VNI, if applicable.
VNI Status	Displays the status of the VNI.
Mode	Displays the VNI modes: Control Plane or Data Plane.
Type	Displays whether the VXLAN VNI is associated with a network (Layer 2) or a VRF (Layer 3).
VRF	Displays the VRF name that is associated with the VXLAN VNI if it is a Layer 3 VNI.
Mapped VLAN	Displays the VLAN or Bridge domain that is mapped to VNI.

**FEX**

The Fabric Extender feature allows you to manage a Cisco Nexus 2000 Series Fabric Extender and its association with the Cisco NX-OS switch that it is attached to. A Fabric Extender is connected to the switch through physical Ethernet interfaces or a Port Channel. By default, the switch does not allow the attached Fabric Extender to connect until it has been assigned a chassis ID and is associated with the connected interface. You can configure a Fabric Extender host interface port as a routed or Layer 3 port. However, no routing protocols can be tied to this routed interface.



**Note** FEX feature is available on LAN devices only. Therefore, you will see FEX on Cisco DCNM **Inventory Switches**. FEX is also not supported on Cisco Nexus 1000V devices.



**Note** 4x10G breakout for FEX connectivity is not supported on Cisco Nexus 9500 Switches.



**Note** The Fabric Extender may connect to the switch through several separate physical Ethernet interfaces or at most one port channel interface.

This section describes how to manage Fabric Extender (FEX) on Cisco Nexus Switches through Cisco DCNM. You can create and manage FEX from Cisco DCNM **Inventory > Switches**.

**Note**

FEX tab is visible only if you choose a LAN device.

The following table describes the fields that appear on this page.

**Table 4: FEX Operations**

Field	Description
Add	Click to add a new FEX to a Cisco Nexus Switch.
Edit	Select any active FEX radio button and click Edit to edit the FEX configuration.  You can create an edit template and use it for editing FEX. Select template type as POLICY and sub type as FEX.
Delete	Select the FEX radio button, and click Delete icon to delete the FEX associated with the switch.
Show	Allows you to view various configuration details for the selected FEX ID. You can select the following from the drop-down list. <ul style="list-style-type: none"> <li>• show_diagnostic</li> <li>• show_fex</li> <li>• show_fex_detail</li> <li>• show_fex_fabric</li> <li>• show_fex_inventory</li> <li>• show_fex_module</li> </ul> The variables for respective show commands are displayed in the Variables area. Review the Variables and click <b>Execute</b> . The output appears in the <b>Output</b> area.  You can create a show template for FEX. Select template type as SHOW and sub type as FEX.
FEX History	Allows you to view the history of the FEX configuration tasks for a particular FEX. You can review the Event Type, Policy Name, Status, Time of Execution, User Name for the selected FEX.

**Table 5: FEX Field and Description**

Field	Description
Fex Id	Uniquely identifies a Fabric Extender that is connected to a Cisco NX-OS device.
Fex Description	Description that is configured for the Fabric Extender.

Field	Description
Fex Version	Specifies the version of the FEX that is associated with the switch.
Pinning	An integer value that denotes the maximum pinning uplinks of the Fabric Extender that is active at a time.
State	Specifies the status of the FEX as associated with the Cisco Nexus Switch.
Model	Specifies the model of the FEX.
Serial No.	Specifies the configured serial number. <b>Note</b> If this configured serial number and the serial number of the Fabric Extender are not the same, the Fabric Extender will not be active.
Port Channel	Specifies the port channel number to which the FEX is physically connected to the Switch.
Ethernet	Refers to the physical interfaces to which the FEX is connected.
vPC ID	Specifies the vPC ID configured for FEX.

This chapter includes the following sections:

## Add FEX

To add single-home FEX from the Cisco DCNM Web UI, perform the following steps:

### Before you begin

You can add a Fabric Extender (FEX) to the Cisco Nexus Switches through the Cisco DCNM Web Client. If the FEX is physically connected to the switch, FEX will become online after it is added. If the FEX is not physically connected to the switch, the configuration is deployed to the switch, which in turn enables FEX when connected.



#### Note

You can create only single homed FEX through **Inventory > Switches > FEX > Add FEX**. To create a dual-homed FEX, use the vPC wizard through **Configure > Deploy > vPC**.

Ensure that you have successfully discovered LAN devices and configured LAN credentials before you configure FEX.

### Procedure

---

#### Step 1

Choose **Inventory > Switches > FEX**.

The **FEX** window is displayed.

#### Step 2

Click the **Add FEX** icon.

- Step 3** In the General tab, in the **PORTCHANNEL** field, enter the interface port channel number which is connected to the FEX.
- Step 4** In the **INT\_RANGE** field, enter the interface range within which the FEX is connected to the switch.
- Note** Do not enter the interface range, if the interfaces are already a part of port channel.
- Step 5** In the **FEX\_ID** field, enter the ID for FEX that is connected to a Cisco NX-OS device.  
The identifier must be an integer value between 100 to 199.
- Step 6** Click **Add**.  
The configured Single-home FEX appears in the list of FEXs associated to the device.

## Edit FEX

To edit and deploy FEX from the Cisco DCNM Web UI, perform the following steps:

### Procedure

- Step 1** Choose **Inventory > Switches > FEX**.  
The **FEX** window is displayed.
- Step 2** Select the FEX radio button that you must edit. Click **Edit FEX** icon.
- Step 3** In the Edit Configuration window, from the Policy drop-down list, select **Edit\_FEX** to edit the FEX configuration.
- Step 4** Edit the **pinning** and **FEX\_DESC** fields, as required.
- Note** If you initially configured port 33 on the parent switch as your only fabric interface, all 48 host interfaces are pinned to this port. If you provision another port, for example 35, then you must perform this procedure to redistribute the host interfaces. All host interfaces are brought down and host interfaces 1 to 24 are pinned to fabric interface 33 and host interfaces 25 to 48 are pinned to fabric interface 35.
- Step 5** Click **Preview**.  
You can view the generated configuration for the selected FEX ID. The following is a configuration example for FEX ID 101.
- ```
fex 101
pinning max-links 1
description test
```
- Step 6** After you review the configuration summary on the Preview window, on the Edit Configuration screen, click **Deploy** to deploy the FEX for the switch.

## VDCs

This section describes how to manage Virtual Device Contexts (VDCs) on Cisco Nexus 7000 Switches through Cisco DCNM.

Users with the network administrator (network-admin) role can create Virtual Device Contexts (VDCs). VDC resource templates limit the amount of physical device resources available to the VDC. The Cisco NX-OS software provides a default resource template, or you can create resource templates.

You can create and manage VDCs from Cisco DCNM **Inventory > Switches > VDCs**. As Cisco DCNM supports DCNM on Cisco Nexus 7000 Series only, click an active Cisco Nexus 7000 Switch. After you create a VDC, you can change the interface allocation, VDC resource limits, and the high availability (HA) policies.

The following table describes the fields that appear on this page.

**Table 6: Vdc Operations**

| Field      | Description                                                                                                                                                                                                                                                                                                                                   |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Add        | Click to add a new VDC.                                                                                                                                                                                                                                                                                                                       |
| Edit       | Select any active VDC radio button and click Edit to edit the VDC configuration.                                                                                                                                                                                                                                                              |
| Delete     | Allows you to edit the VDC configuration. Select any active VDC radio button and click Edit to edit the VDC configuration.                                                                                                                                                                                                                    |
| Resume     | Allows you to delete the VDC. Select any active VDC radio button and click Delete to remove the VDC associated with the device.                                                                                                                                                                                                               |
| Suspend    | Allows you to suspend an active non-default VDC. Save the VDC running configuration to the startup configuration before suspending the VDC. Otherwise, you will lose the changes to the running configuration.<br><br><b>Note</b> You cannot suspend the default VDC.<br><br><b>Caution</b> Suspending a VDC disrupts all traffic on the VDC. |
| Rediscover | Allows you to resume a non-default VDC from the suspended state. The VDC resumes with the configuration that is saved in the startup configuration.                                                                                                                                                                                           |
| Show       | Allows you to view the Interfaces and Resources that are allocated to the selected VDC.<br><br>In the Interface tab, you can view the mode, admin-status, and operational status for each interface associated with the VDC.<br><br>In the Resource tab, you can view the allocation of resources and current usage of these resources.       |

**Table 7: Vdc Table Field and Description**

| <b>Field</b>                                                                                               | <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name                                                                                                       | Displays the unique name for the VDC                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Type                                                                                                       | Specifies the type of VDC. The two types of VDCs are: <ul style="list-style-type: none"> <li>• Ethernet</li> <li>• Storage</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Status                                                                                                     | Specifies the status of the VDC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Resource Limit-Module Type                                                                                 | Displays the allocated resource limit and module type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| HA-Policy <ul style="list-style-type: none"> <li>• Single Supervisor</li> <li>• Dual Supervisor</li> </ul> | <p>Specifies the action that the Cisco NX-OS software takes when an unrecoverable VDC fault occurs.</p> <p>You can specify the HA policies for single supervisor module and dual supervisor module configurations when you create the VDC. The HA policy options are as follows:</p> <p><b>Single supervisor module configuration:</b></p> <ul style="list-style-type: none"> <li>• Bringdown—Puts the VDC in the failed state. To recover from the failed state, you must reload the physical device.</li> <li>• Reload—Reloads the supervisor module.</li> <li>• Restart—Takes down the VDC processes and interfaces and restarts it using the startup configuration.</li> </ul> <p><b>Dual supervisor module configuration:</b></p> <ul style="list-style-type: none"> <li>• Bringdown—Puts the VDC in the failed state. To recover from the failed state, you must reload the physical device.</li> <li>• Restart—Takes down the VDC processes and interfaces and restarts it using the startup configuration.</li> <li>• Switchover—Initiates a supervisor module switchover.</li> </ul> <p>The default HA policies for a non-default VDC that you create is restart for a single supervisor module configuration and switchover for a dual supervisor module configuration. The default HA policy for the default VDC is reload for a single supervisor module configuration and switchover for a dual supervisor module configuration.</p> |

| Field                                                                                                        | Description                                                                                                |
|--------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|
| Mac Address                                                                                                  | Specifies the default VDC management MAC address.                                                          |
| Management Interface <ul style="list-style-type: none"> <li>• IP Address Prefix</li> <li>• Status</li> </ul> | Specifies the IP Address of the VDC Management interface. The status shows if the interface is up or down. |
| SSH                                                                                                          | Specifies the SSH status                                                                                   |

This chapter includes the following sections:

## Add VDCs

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

### Before you begin

Ensure that you have discovered the physical device using a username that has the network-admin role.

Obtain an IPv4 or IPv6 address for the management interface (mgmt 0) if you want to use out-of-band management for the VDC.

Create a storage VDC to run FCoE. The storage VDC cannot be the default VDC and you can have one storage VDC on the device.

### Procedure

---

**Step 1** Choose **Inventory > Switches > VDC**.

The **VDC** window is displayed.

**Step 2** Click the **Add VDC** icon.

**Step 3** From the drop-down list, select the VDC type.

You can configure the VDC in two modes.

- Ethernet VDC
- Storage VDC

The default VDC type is Ethernet.

**Step 4** Click **OK**.

---

### Configuring Ethernet VDCs

To configure VDC in Ethernet mode from the Cisco DCNM Web UI, perform the following steps:

## Procedure

- Step 1** In the General Parameter tab, specify the **VDC Name**, **Single supervisor HA-policy**, **Dual supervisor HA-policy**, and **Resource Limit - Module Type**.
- Step 2** In the Allocate Interface tab, select the network interfaces (dedicated interfaces membership) to be allocated to the VDC.
- Click **Next**.
- Step 3** In the Allocate Resource tab, specify the resource limits for the VDC.
- Select the radio button and choose **Select a Template from existing Templates** or **Create a New Resource Template**. VDC resource templates describe the minimum and maximum resources that the VDC can use. If you do not specify a VDC resource template when you create a VDC, the Cisco NX-OS software uses the default template, vdc-default.
- If you choose Select a Template from existing Templates, from the **Template Name** drop-down list, you can select **None**, **global-default**, or **vdc-default**.
- The template resource limits are detailed in the following below:

**Table 8: Template Resource Limits**

| Resource                                    | Minimum | Maximum                            |
|---------------------------------------------|---------|------------------------------------|
| Global Default VDC Template Resource Limits |         |                                    |
| Anycast Bundled                             |         |                                    |
| IPv6 multicast route memory                 | 8       | 8<br>Route memory is in megabytes. |
| IPv4 multicast route memory                 | 48      | 48                                 |
| IPv6 unicast route memory                   | 32      | 32                                 |
| IPv4 unicast route memory                   |         |                                    |
| VDC Default Template Resource Limits        |         |                                    |
| Monitor session extended                    |         |                                    |
| Monitor session mx exception                |         |                                    |
| Monitor SRC INBAND                          |         |                                    |
| Port Channels                               |         |                                    |
| Monitor DST ERSPAN                          |         |                                    |
| SPAN Sessions                               |         |                                    |
| VLAN                                        |         |                                    |
| Anycast Bundled                             |         |                                    |

| Resource                    | Minimum | Maximum |
|-----------------------------|---------|---------|
| IPv6 multicast route memory |         |         |
| IPv4 multicast route memory |         |         |
| IPv6 unicast route memory   |         |         |
| IPv4 unicast route memory   |         |         |
| VRF                         |         |         |

- If you choose Create New Resource Template, enter a unique **Template Name**. In the Resource Limits area, enter the minimum and maximum limits, as required for the resources.

You can edit individual resource limits for a single VDC through the Cisco DCNM Web Client > Inventory > Switches > VDC.

Click **Next**.

- Step 4** In the Authenticate tab, you can allow the Admin to configure the password and also authenticate users using AAA Server Groups.

In the Admin User Area:

- Check the **Enable Password Strength Check** checkbox, if necessary.
- In the **Password** field, enter the admin user password.
- In the **Confirm Password** field, reenter the admin user password.
- In the **Expiry Date** field, click the down arrow and choose an expiry date for the admin user from the Expiry Date dialog box. You can also select **Never** radio button not to expire the password.

In the AAA Server Groups area:

- In the **Group Name** field, enter an AAA server group name.
- In the **Servers** field, enter one or more host server IPv4 or IPv6 addresses or names, which are separated by commas.
- In the **Type** field, choose the type of server group from the drop-down list.

Click **Next**.

- Step 5** In the Management Ip tab, enter IPv4 or IPv6 Address information.

Click **Next**.

- Step 6** In the Summary tab, review the VDC configuration.

Click **Previous** to edit any parameters.

Click **Deploy** to configure VDC on the device.

- Step 7** In the Deploy tab, the status of the VDC deployment is displayed.

A confirmation message appears. Click **Know More** to view the commands that are executed to deploy the VDC.

Click **Finish** to close the VDC configuration wizard and revert to view the list of VDCs configured on the device.

## Configuring Storage VDCs

To configure VDCs in storage mode from the Cisco DCNM Web UI, perform the following steps:

### Before you begin

Create a separate storage VDC when you run FCoE on the device. Only one of the VDCs can be a storage VDC, and the default VDC cannot be configured as a storage VDC.

You can configure shared interfaces that carry both Ethernet and Fibre Channel traffic. In this specific case, the same interface belongs to more than one VDC. The shared interface is allocated to both an Ethernet and a storage VDC.

### Procedure

**Step 1** In the General Parameter tab, specify the **VDC Name**, **Single supervisor HA-policy**, **Dual supervisor HA-policy**, and **Resource Limit - Module Type**.

**Step 2** In the Allocate FCoE Vlan tab, select the available **Ethernet Vdc** from the drop-down list.

The existing Ethernet VLANs range is displayed. Select **None** not to choose any available Ethernet VDCs.

You can allocate specified FCoE VLANs to the storage VDC and specified interfaces.

Click **Next**.

**Step 3** In the Allocate Interface tab, add the dedicated and shared interfaces to the FCoE VDC.

**Note** The dedicated interface carries only FCoE traffic and the shared interface carries both the Ethernet and the FCoE traffic.

You can configure shared interfaces that carry both Ethernet and Fibre Channel traffic. In this specific case, the same interface belongs to more than one VDC. FCoE VLAN and shared interface can be allocated from same Ethernet VDC.

Click **Next**.

**Step 4** In the Authenticate tab, you can allow the Admin to configure the password and also authenticate users using AAA Server Groups.

In the Admin User Area:

- Check the **Enable Password Strength Check** checkbox, if necessary.
- In the **Password** field, enter the admin user password.
- In the **Confirm Password** field, reenter the admin user password.
- In the **Expiry Date** field, click the down arrow and choose an expiry date for the admin user from the Expiry Date dialog box. You can also select **Never** radio button not to expire the password.

In the AAA Server Groups area:

- In the **Group Name** field, enter an AAA server group name.
- In the **Servers** field, enter one or more host server IPv4 or IPv6 addresses or names, which are separated by commas.
- In the **Type** field, choose the type of server group from the drop-down list.

Click **Next**.

**Step 5** In the Management Ip tab, enter IPv4 or IPv6 Address information.

Click **Next**.

**Step 6** In the Summary tab, review the VDC configuration.

Click **Previous** to edit any parameters.

Click **Deploy** to configure VDC on the device.

**Step 7** In the Deploy tab, the status of the VDC deployment is displayed.

A confirmation message appears. Click **Know More** to view the commands that are executed to deploy the VDC.

Click **Finish** to close the VDC configuration wizard and revert to view the list of VDCs configured on the device.

---

## Edit VDC

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

### Procedure

---

**Step 1** Choose **Inventory > Switches > VDC**.

The **VDC** window is displayed.

**Step 2** Select the VDC radio button that you must edit. Click the **Edit VDC** icon.

**Step 3** Modify the parameters as required.

**Step 4** After you review the configuration summary on the Summary tab, click **Deploy** the VDC with the new configuration.

---

## Switch On-Board Analytics

For the selected switch, the **Switch On-Board Analytics** dashboard displays the following charts:



### Note

The graph data cannot be retrieved if correct certificates are not added to the Switch. Ensure that the certificates are valid for nxapi feature and SAN analytics to function properly.

- Top 10 Slowest Ports

- Top 10 Slowest Target Ports
- Top 10 Slowest Flows
- Top 10 Slowest ITLs
- Top 10 Port Traffic
- Top 10 Target Ports Traffic
- Top 10 Flow Traffic
- Top 10 ITL Traffic

The following metrics are supported by the Switch On-Board Analytics charts:

- Read and Write Completion Time—Time that is taken for an IO to complete successfully, that is, the time gap between IO status from a Target and IO command from an Initiator. The following metrics are supported:

- Read Completion Time Min
- Read Completion Time Max
- Write Completion Time Min
- Write Completion Time Max

The IO engine tracks the maximum and minimum IO completion time for read and write commands in the context of a switch's port, target port, flows, initiators, and LUNs.

- Read and Write Initiation Time—Time that is taken for an IO to initiate, that is, the time gap between the first response packet from a Target and IO Command from Initiator. The following metrics are supported:

- Read Initiation Time Min
- Read Initiation Time Max
- Write Initiation Time Min
- Write Initiation Time Max

The IO engine tracks the maximum and minimum IO initiation time for read and write commands in the context of a switch's port, target port, flows, initiators, and LUNs.

- Read and Write IO Bandwidth—Read and write command bandwidth observed in the context of a switch's port traffic, target port traffic, flow traffic, initiators, and LUNs. The IO bandwidth is computed at every four second time interval based on the number of bytes read or written.

- Read and Write IO Rate—Read and write command IO rate observed in the context of a switch's port traffic, target port traffic, flow traffic, initiators, and LUNs. The IO rate is computed at every four second time interval that is based on the number of IO performed.

- Read and Write IO Size—Read and write command IO size observed in the context of a switch's port traffic, target port traffic, flow traffic, initiators, and LUNs. The following metrics are supported:

- Read IO Size Min
- Read IO Size Max

## Viewing Switch On-Board Analytics

- Write IO Size Min
- Write IO Size Max

The IO engine tracks the maximum and minimum IO size for read and write commands.

## Viewing Switch On-Board Analytics

You can view the switch on-board analytics information from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

**Step 1** Choose **Inventory > View > Switches**.

The discovered switches are displayed.

**Step 2** Click a switch name in the **Device Name** column.

The **Switch** dashboard that corresponds to that switch is displayed.

**Step 3** Click the **Switch On-Board Analytics** tab.

This tab displays the Switch On-Board Analytics charts.

---

## Configuring Settings for the Switch On-Board Analytics Charts

Perform the following actions to configure the settings for the switch on-board analytics charts:

- From the **Show Time as** drop-down list, choose time to be shown in the charts. You can choose one of the following options:
  - **Microseconds**
  - **Milliseconds**
  - **Seconds**

By default, **Microseconds** is chosen.



**Note**

The **Show Time** drop-down list is applicable only for the top ten slowest ports, target ports, flows, and ITLs.

- From the **Show Flow From** drop-down list, choose whether to show flows from a **Target** or from an **Initiator**. By default, flows from a **Target** are chosen.



**Note**

The **Show Flow From** drop-down list is applicable only for the charts displaying flows and ITLs.

- From the **Show bandwidth and Size** as drop-down list, choose the traffic information to be shown in the charts. You can choose one of the following options:

- Bytes
- KB
- MB

By default, **Bytes** is chosen.

- Check the **Filter results** check box, and click either the **by fc port** or **by VSAN** radio button and specify the appropriate values to filter the chart results. The FC port value must be in the **fc slot/port** format and the VSAN value must be a digit within the allowed VSAN range.

Click the Filter icon next to the **by fc port** to apply changes.



**Note** Filtering results by VSAN is not applicable for the **Top 10 Slowest Ports** or **Top 10 Port Traffic** charts.

- Check the **Single Column** check box to display the charts in a single column instead of double columns.
- Click the **Refresh** icon in the upper-right corner to refresh the charts.

## Viewing Switch On-Board Analytics Charts

Perform the following actions to view the charts under the **Switch On-Board Analytics** tab:

- View the charts for the top ten slowest ports, target ports, flows, and ITLs by choosing one of the following variables from the drop-down list:
  - **Read Completion Time**—The read command completion time observed in the context of a switch's port.
  - **Write Completion Time**—The write command completion time observed in the context of a switch's port.
  - **Read Initiation Time**—The read command initiation time observed in the context of a switch's port.
  - **Write Initiation Time**—The write command initiation time observed in the context of a switch's port.



**Note**

- By default, **Read Completion Time** is selected and all the units for time are in **Microseconds**.
- Each chart contains a legend that provides information about the variable displayed. Each variable has a check box. Unselecting the check box removes the variable data from the chart or table.

## Viewing Switch On-Board Analytics Charts

- View the charts for the top ten port traffic, target port traffic, flow traffic, and ITL traffic by choosing one of the following variables from the drop-down list:
  - **Read IO Rate**—The read command data observed in the context of a switch's port.
  - **Write IO Rate**—The write command observed in the context of a switch's port.
  - **Read IO Size**—The read command size observed in the context of a switch's port.
  - **Write IO Size**—The write command size observed in the context of a switch's port.
  - **Read IO Bandwidth**—The read command bandwidth observed in the context of a switch's port.
  - **Write IO Bandwidth**—The write command bandwidth observed in the context of a switch's port.



### Note

- By default, **Read IO Rate** is selected. The **Read IO Rate** is IO per second. Both **Rate** and **Bandwidth** units are per second over an 8-second range. The **Size** value is for the life of the switch or since the last clear command was run from the CLI.
- The **Read IO Size** and **Read IO Bandwidth** units are in bytes per second. You can change this unit by using the **Show Bandwidth and Size** drop-down list. You can choose from the three options: **Bytes**, **KB**, and **MB**.
- Each chart contains a legend that provides information about the variable displayed. Each variable has a check box. Unselecting the check box removes the variable data from the chart or table.

- 
- Choose the format to display information from the **Show** drop-down list. You can choose one of the following formats:

- **Chart**

- **Table**

- **Chart and Table**



### Note

- To display information in the **Chart and Table** format, enlarge your browser window or check the **Single Column** check box on the upper right corner.
- The default for Top ten Slowest Ports and Top 10 Port Traffic is **Chart and Table**.

- 
- Use the **Chart Type** drop-down list to display information in the **Bar Chart** or **Stacked Bar Chart**.
  - Use the **Actions** drop-down list to export information in a CSV or PDF, or print the required information.
  - To view a chart or a table in a new window, click the **Detach** icon on the upper-right corner of a chart or a table. After detaching a chart or table, you can view the top 25 slowest ports, target ports, flows, ITLs, or their traffic.

## Viewing Inventory Information for Modules

To view the inventory information for modules from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

**Step 1** Choose **Inventory > View > Modules**.

The **Modules** window is displayed with a list of all the switches and its details for a selected Scope.

**Step 2** You can view the following information.

- **Group** column displays the group name of the module.
  - **Switch** column displays the switch name on which the module is discovered.
  - **Name** displays the module name.
  - **ModelName** displays the model name.
  - **SerialNum** column displays the serial number.
  - **2nd SerialNum** column displays the second serial number.
  - **Type** column displays the type of the module.
  - **Slot** column displays the slot number.
  - **Hardware Revision** column displays the hardware version of the module.
  - **Software Revision** column displays the software version of the module.
  - **Asset ID** column displays the asset id of the module.
  - **OperStatus** column displays the operation status of the module.
- 

## Viewing Inventory Information for Licenses

To view the inventory information for licenses from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

**Step 1** Choose **Inventory > View > Licenses**.

The **Licenses** window is displayed based on the selected Scope.

**Step 2** You can view the following information.

- **Group** column displays the group name of switches.
- **Switch** column displays the switch name on which the feature is enabled.
- **Feature** displays the installed feature.

- **Status** displays the usage status of the license.
  - **Type** column displays the type of the license.
  - **Warnings** column displays the warning message.
- 

## Monitoring Switch

The Switch menu includes the following submenus:

### Viewing Switch CPU Information

To view the switch CPU information from the Cisco DCNM Web UI, perform the following steps:

#### Procedure

---

**Step 1** Choose **Monitor > Switch > CPU**.

The **CPU** window is displayed. This window displays the CPU information for the switches in that scope.

**Step 2** You can use the drop-down to filter the view by Last 10 Minutes, Last Hour, Last Day, Last Week, Last Month, and Last Year.

**Step 3** In the **Switch** column, click the switch name to view the Switch Dashboard.

**Step 4** Click the chart icon in the **Switch** column to view the CPU utilization.

You can also change the chart timeline to Last 10 Minutes, Last Hour, Last Day, Last Week, Last Month, and Last Year. You can choose the chart type and chart options to show as well.

---

### Viewing Switch Memory Information

To view the switch memory information from the Cisco DCNM Web UI, perform the following steps:

#### Procedure

---

**Step 1** Choose **Monitor > Switch > Memory**.

The memory panel is displayed. This panel displays the memory information for the switches in that scope.

**Step 2** Use the drop-down to filter the view by Last 10 Minutes, Last Hour, Last Day, Last Week, Last Month, and Last Year.

**Step 3** Click the chart icon in the **Switch** column to see a graph of the memory usage of the switch.

**Step 4** In the **Switch** column, click the switch name to view the Switch Dashboard.

- Step 5** You can use the drop-down to view the chart in different time lines. Use the chart icons to view the memory utilization chart in varied views.

## Viewing Switch Traffic and Errors Information

To view the switch traffic and errors information from the Cisco DCNM Web UI, perform the following steps:

### Procedure

- Step 1** Choose **Monitor > Switch > Traffic**.

The **Switch Traffic** panel is displayed. This panel displays the traffic on that device for the past 24 hours.

- Step 2** Use the drop-down to filter the view by 24 hours, Week, Month, and Year.

- Step 3** Click the **Export** icon in the upper-right corner to export the data into a spreadsheet.

- Step 4** Click **Save**.

- Step 5** Click the switch name to view the **Switch Dashboard** section.

## Viewing Switch Temperature

Cisco DCNM includes the module temperature sensor monitoring feature, using which you can view the sensor temperature of a switch. You can choose an interval by which to filter the sensor list. The default interval is **Last Day**. Only sensors that have historical temperature data is shown in the list. You can choose between Last ten Minutes, Last Hour, Last Day, Last Week, and Last Month.



- Note** It is not necessary to configure the LAN credentials under the **Configure > Credentials Management > LAN Credentials** screen to fetch the temperature monitoring data from the switches.

To view the switch temperature information from the Cisco DCNM Web UI, perform the following steps:

### Procedure

- Step 1** Choose **Monitor > Switch > Temperature**.

The **Switch Temperature** window is displayed with the following columns.

- **Scope**: The sensor belongs to a switch, which is part of a fabric. The fabric that it belongs to is shown as its scope. When the scope selector at the top of Cisco DCNM is used, the sensor list is filtered by that scope.
- **Switch**: Name of the switch the sensor belongs to.
- **IP Address**: IP Address of the switch.
- **Temperature Module**: The name of the sensor module.

## Enabling Temperature Monitoring

- **Avg/Range:** The first number is the average temperature over the interval that is specified at the top of the table. The second set of numbers is the range of the temperature over that interval.
- **Peak:** The maximum temperature over the interval

- Step 2** From this list, each row has a chart icon, which you can click. A chart is displayed, which shows historical data for the sensor. The interval for this chart can be changed as well, between 24 hours, 1 week, and 1 month.
- 

## Enabling Temperature Monitoring

You can enable the temperature monitoring feature for LAN switches from the LAN Collections screen, and for the SAN switches by setting a few properties under Administration > DCNM Server > Server Properties screens.

### Enabling Temperature Monitoring for LAN Switches

1. From the menu bar, choose **Administration > Performance Setup > LAN Collections**.
2. Select the **Temperature Sensor** check box.
3. Select the type of LAN switches for which you want to collect performance data.
4. Click **Apply** to save the configuration.

## Viewing Accounting Information

To view the accounting information from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Monitor > Switch > Accounting**.

The fabric name or the group name along with the accounting information is displayed.

- Step 2** Select **Advanced Filter** beside the filter icon to search the accounting information by **Source**, **Username**, **Time**, and **Description**. Or select **Quick Filter** to search under each column.

- Step 3** You can also select a row and click the **Delete** icon to delete accounting information from the list.

- Step 4** You can use the **Print** icon to print the accounting details and use the **Export** icon to export the data to a Microsoft Excel spreadsheet.
- 

## Viewing Events Information

To view the events and syslog from the Cisco DCNM Web UI, perform the following steps:

---

**Procedure****Step 1** Choose **Monitor > Switch > Events**.

The fabrics along with the switch name and the events details are displayed.

The **Count** column displays the number of times the same event has occurred during the time period as shown in the **Last Seen** and **First Seen** columns.

Click a switch name in the **Switch** column to view the switch dashboard.

**Step 2** Select an event in the table and click the **Add Suppressor** icon to open the shortcut of adding an event suppressor rule.**Step 3** Select one or more events from the table and click the **Acknowledge** icon to acknowledge the event information for the fabric.

- After you acknowledge the event for a fabric, the acknowledge icon is displayed in the **Ack** column next to the fabric.

**Step 4** Select the fabric and click the **Unacknowledge** icon to cancel an acknowledgment for a fabric.**Step 5** Select **Advanced Filter** beside the filter icon to search the accounting information by **Source**, **Username**, **Time**, and **Description**. Or select **Quick Filter** to search under each column.**Step 6** Select a fabric and use the **Delete** icon to delete the fabric and event information from the list.**Step 7** Click the **Print** icon to print the event details.**Step 8** Click the **Export to Excel** icon to export the data.

---

## Monitoring LAN

The LAN menu includes the following submenus:

### Monitoring Performance Information for Ethernet

To monitor the performance information for ethernet from the Cisco DCNM Web UI, perform the following steps:

---

**Procedure****Step 1** Choose **Monitor > LAN > Ethernet**.

The **Ethernet** window is displayed.

**Step 2** You can use the drop-down to filter the view by Last 10 Minutes, Last Hour, Last Day, Last Week, Last Month, and Last Year.

There are variations to this procedure. In addition to these basic steps, you can also perform the following steps:

- Select the name of an Ethernet port from the **Name** column to see a graph of the traffic across that Ethernet port for the past 24 hours. You can change the time range for this graph by selecting it from the drop-down list in the upper-right corner.
- To export the data into a spreadsheet, click the **Export** icon in the upper-right corner and click **Save**.
- Use the chart icons to view the traffic chart in varied views. You can also use the icons to **Append**, **Predict**, and **Interpolate Data**.
- For the Rx/Tx calculation, see the following Rx/Tx calculation.

**Note** The conversion for Fabrics is 10 bit = 1 byte and for LAN traffic, the conversion is 8 bit = 1 byte.

- Average Rx/Tx % = Average Rx/Tx divided by Speed \* 100
- Peak Rx/Tx % = Peak Rx/Tx divided by Speed \* 100

**Note** If the performance tables do not contain any data, see the Thresholds section to turn on performance data collection.

---

## Monitoring ISL Traffic and Errors

To monitor the ISL traffic and errors from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

#### Step 1 Choose **Monitor > LAN > Link**.

The **ISL Traffic and Errors** window is displayed. This panel displays the ISL information for the end devices in that scope. You can reduce or expand the scope of what is displayed by using the scope menu.

#### Step 2 You can use the drop-down to filter the view by Last 10 Minutes, Last Hour, Last Day, Last Week, Last Month, and Last Year.

**Note** **NaN** (Not a Number) in the data grid means that the data is not available.

There are variations to this procedure. In addition to these basic steps, you can perform the following steps to view detailed information for ISLs:

- To change the time range for this graph, select it from the drop-down list in the upper-right corner.
- To view the detailed information for a specific period, drag the slider control to choose the time interval for which you need the information.
- Use the chart icons to view the traffic chart in varied views. You can also use the icons to **Append**, **Predict**, and **Interpolate Data**. To view real-time information, choose **Real Time** from the drop-down list in the **Chart** menu.
- To export the data into a spreadsheet, choose **Export** from the drop-down list in the **Chart** menu and then click **Save**.
- For the Rx/Tx calculation, see the following Rx/Tx calculation.

**Note** The conversion for Fabrics is 10 bit = 1 byte and for LAN traffic, the conversion is 8 bit = 1 byte.

- Average Rx/Tx % = Average Rx/Tx divided by Speed \* 100
- Peak Rx/Tx % = Peak Rx/Tx divided by Speed \* 100

**Note** If the performance tables do not contain any data, see the Performance Setup Thresholds section to turn on performance.

---

## Monitoring a vPC

The virtual port channel (vPC) feature enables you to view the links that are physically connected to different devices as a single port channel. A vPC is an extended form of a port channel which allows you to create redundancy and increase bisectional bandwidth by enabling multiple parallel paths between nodes and allowing load balancing traffic. Traffic is distributed among two single device vPC endpoints. If there is an inconsistency in the vPC configurations, the vPC does not function correctly.



**Note** To view the vPC in **vPC Performance**, both primary and secondary device should be designated to the user. If either one kind of switch is not designated, vPC information is displayed.

Cisco DCNM **Web Client > Monitor > vPC** displays only consistent vPCs displays both the consistent and inconsistent vPCs.

You can identify the inconsistent vPCs and resolve the inconsistencies in each vPC by using the Cisco DCNM **Web UI > Configure > Deploy > vPC Peer** and **Web Client > Configure > Deploy > vPC**.

[Table 9: vPC Performance, on page 213](#) displays the following vPC configuration details in the data grid view.

**Table 9: vPC Performance**

| Column                                          | Description                                                                  |
|-------------------------------------------------|------------------------------------------------------------------------------|
| Search box                                      | Enter any string to filter the entries in their respective column.           |
| <b>vPC ID</b>                                   | Displays vPC ID's configured device.                                         |
| <b>Domain ID</b>                                | Displays the domain ID of the vPC peer switches.                             |
| <b>Multi Chassis vPC EndPoints</b>              | Displays the multi-chassis vPC endpoints for each vPC ID under a vPC domain. |
| <b>Primary vPC Peer - Device Name</b>           | Displays the vPC Primary device name.                                        |
| <b>Primary vPC Peer - Primary vPC Interface</b> | Displays the primary vPC interface.                                          |
| <b>Primary vPC Peer - Capacity</b>              | Displays the capacity for the primary vPC peer.                              |
| <b>Primary vPC Peer - Avg. Rx/sec</b>           | Displays the average receiving speed of primary vPC peer.                    |
| <b>Primary vPC Peer - Avg. Tx/sec</b>           | Displays the average sending speed of primary vPC peer.                      |

| Column                                  | Description                                                     |
|-----------------------------------------|-----------------------------------------------------------------|
| <b>Primary vPC Peer - Peak Util%</b>    | Displays the peak utilization percentage of primary vPC peer.   |
| <b>Secondary vPC Peer - Device Name</b> | Displays the vPC secondary device name.                         |
| <b>Secondary vPC Interface</b>          | Displays the secondary vPC interface.                           |
| <b>Secondary vPC Peer - Capacity</b>    | Displays the capacity for the secondary vPC peer.               |
| <b>Secondary vPC Peer - Avg. Rx/sec</b> | Displays the average receiving speed of secondary vPC peer.     |
| <b>Secondary vPC Peer - Avg. Tx/sec</b> | Displays the average sending speed of secondary vPC peer.       |
| <b>Secondary vPC Peer - Peak Util%</b>  | Displays the peak utilization percentage of secondary vPC peer. |

You can use this feature as following:

## Monitoring vPC Performance

You can view the relationship among consistent virtual port channels (vPCs). You can view the statistics of all member interfaces and the aggregate of the statistics at the port channel level.



**Note** This tab only displays consistent vPCs.

To view the VPC performance information from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

**Step 1** Choose **Monitor > LAN > vPC**.

The **vPC Performance** statistics is displayed. The aggregated statistics of all vPCs are displayed in a tabular manner.

**Step 2** Click the **vPC ID**.

The vPC topology, **vPC Details**, **Peer-link Details**, and **Peer-link Status** is displayed.

The **vPC Consistency**, **Peer-link Consistency**, and **vPC Type2 Consistency** for the vPC is displayed.

- Click the **vPC Details** tab, you can view the parameter details of vPC **Basic Setting** and **Layer 2 Settings** for both Primary and Secondary vPC devices.
- Click the **Peer-link Details** tab, to view the parameter details of peer-link **vPC Global Setting** and **STP Global Settings** for both Primary and Secondary vPC devices.
- Click the **Peer-link Status** tab, the **vPC Consistency**, and **Peer-Link Consistency** status is displayed. The parameter details of **Role Status** and **vPC Peer keep-alive Status** for both Primary and Secondary vPC devices is also displayed.

**Step 3** Click the peer-link icon in front of the **Device Name** in the **Primary vPC peer** or **Secondary vPC peer** column to view its member interface.

**Step 4** Click the **Show Chart** icon of the corresponding interface to view its historical statistics.

The traffic distribution statistics appear at the bottom of the vPC window. By default, the Cisco DCNM Web Client displays the historical statistics for 24 hours.

There are variations to this procedure. In addition to these basic steps, you can also perform the following steps to view detailed information for flows:

- To change the time range for this graph, select it from the drop-down list in the upper right corner.
- To view the detailed information for a specific period, drag the slider control to choose the time interval for which you need the information.
- Use the chart icons to view the traffic chart in varied views.
- You can also use the icons to **Append**, **Predict**, and **Interpolate Data**.
- To print the vPC Utilization data, click the **Print** icon in the upper-right corner. The vPC Utilization page appears.
- To export the data into a spreadsheet, click the **Export** icon in the upper-right corner and click **Save File**.

**Note** If the performance tables do not contain any data, see the Thresholds section to turn on performance data collection.

---

## Monitoring Endpoint Locator

The Endpoint Locator menu includes the following submenus:

### Exploring Endpoint Locator Details

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

#### Procedure

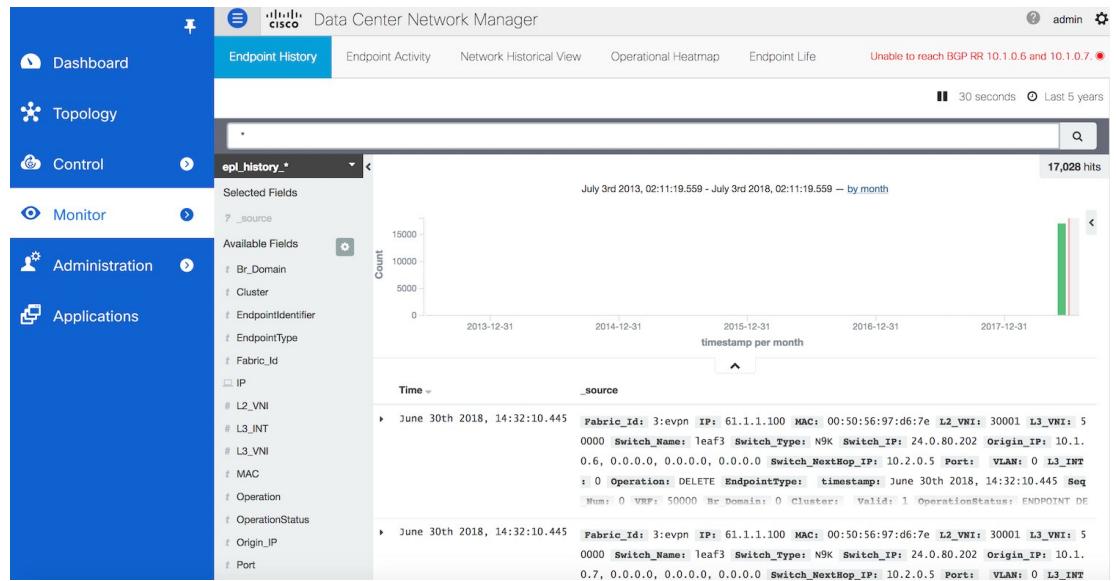
---

Choose **Monitor > Endpoint Locator > Explore**. The Endpoint Locator dashboard appears.

The Endpoint Locator Dashboard displays the following information:

- **Endpoint History**—Real time plot displaying Endpoint events for the period specified in the relative or absolute date range. A user can search for a specific metric value in the search bar. Search is supported on any of the fields as specified under the “Available Fields” column on the menu on the left. A sample screenshot of the endpoint history based on an IP address specified in the search field is depicted below.

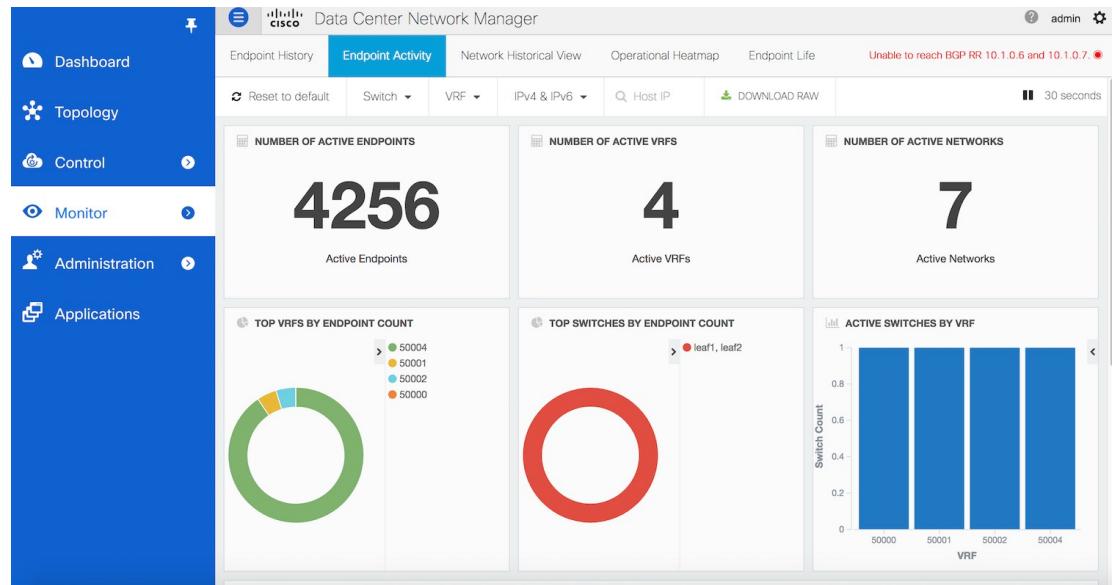
## Exploring Endpoint Locator Details



- **Endpoint Activity**—This view displays the current state of the active endpoints in the fabric.

**Filters** - You can filter and view results for a switch, VRF, IPv4 and IPv6 type of address and IP address of an end point. The entire dashboard view across all tiles and the data table, are updated as soon as the search filters are applied.

**Tiles** - The number of active endpoints including the number of active VRFs and active networks are listed in the top 3 tiles, just below the filters. The break-up of active endpoints is also available on a per VRF as well as a per switch basis. If there is at least one active endpoint in a given VRF behind a switch, then that VRF is considered as active on that switch. Note that the VRF may be configured on a number of switches but it is only considered active and justifies burning resources on the switch, if there is at least one active endpoint in that VRF behind that switch. In that sense, the “ACTIVE SWITCHES BY VRF” tile can provide a good insight for the network administrator into removing extraneous VRF configurations from switches where it may not be needed. At the bottom of the dashboard, there is a data table named LIST OF ACTIVE ENDPOINTS which provides a list of endpoints with context information such as the VRF, IP, MAC, Switch, VLAN, Port etc. By default, the endpoint information is refreshed every 30 seconds. However, the refresh interval may be changed as desired.



Search results can be downloaded in csv format by clicking on the “DOWNLOAD RAW” icon at the top left part of the screen. A sample snippet of the downloaded csv file from a search result is shown below:

| 1  | Fabric_Id | IP         | MAC           | L2_VNI | L3_VNI | Switch_Nam | Switch_Type | Switch_IP | Origin_IP   | Switch_Port          | VLAN | L3_INT | Operation | EndpointType | timestamp        | Seq_Num | VRF   | B_r_Domain | Cluster    | Valid | Op |
|----|-----------|------------|---------------|--------|--------|------------|-------------|-----------|-------------|----------------------|------|--------|-----------|--------------|------------------|---------|-------|------------|------------|-------|----|
| 2  | 3:evpn    | 1.0.14.114 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.6 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 3  | 3:evpn    | 1.0.14.113 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.6 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 4  | 3:evpn    | 1.0.14.114 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.6 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 5  | 3:evpn    | 1.0.14.113 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.6 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 6  | 3:evpn    | 1.0.14.113 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.7 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 7  | 3:evpn    | 1.0.14.112 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.7 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 8  | 3:evpn    | 1.0.14.111 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.7 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 9  | 3:evpn    | 1.0.14.111 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.7 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 10 | 3:evpn    | 1.0.14.109 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.7 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 11 | 3:evpn    | 1.0.14.110 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.7 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 12 | 3:evpn    | 1.0.14.110 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.7 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 13 | 3:evpn    | 1.0.14.109 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.7 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 14 | 3:evpn    | 1.0.14.108 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.7 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 15 | 3:evpn    | 1.0.14.108 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.7 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 16 | 3:evpn    | 1.0.14.107 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.7 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |
| 17 | 3:evpn    | 1.0.14.107 | 00:00:00:2f:c | 30003  | 50004  | leaf1      | leaf2       | N9K       | 24.0.80.203 | 10.1.0.6 10.10.2.0.1 | 0    | 0      | 0 ACTIVE  |              | 2018-06-30 12:31 | 0       | 50004 | 0          | 10.2.0.1:0 | 1     |    |

It is possible to search based on any of the fields describing the information of each endpoint. For example, if the user wants to know the list of endpoints in a given network, that can be achieved as follows. Recall that each network is represented by a unique 24-bit identifier. This parameter is represented by the field L2\_VNI. Here are the steps:

1. Go to the LIST OF ACTIVE ENDPOINTS data table and click on any row. This will expand the row as shown below:

| LIST OF ACTIVE ENDPOINTS       |       |            |                   |              |      |      |
|--------------------------------|-------|------------|-------------------|--------------|------|------|
| Time                           | VRF   | IP         | MAC               | Switch_Name  | Port | VLAN |
| ▶ June 30th 2018, 12:31:11.675 | 50004 | 1.0.14.114 | 00:00:00:2f:09:a1 | leaf1, leaf2 |      | 0    |
| ▶ June 30th 2018, 12:31:11.675 | 50004 | 1.0.14.113 | 00:00:00:2f:09:ff | leaf1, leaf2 |      | 0    |
| ▶ June 30th 2018, 12:31:11.624 | 50004 | 1.0.14.114 | 00:00:00:2f:09:a1 | leaf1, leaf2 |      | 0    |
| ▶ June 30th 2018, 12:31:11.624 | 50004 | 1.0.14.113 | 00:00:00:2f:09:ff | leaf1, leaf2 |      | 0    |
| ▶ June 30th 2018, 12:31:11.429 | 50004 | 1.0.14.112 | 00:00:00:2f:09:fd | leaf1, leaf2 |      | 0    |
| ▶ June 30th 2018, 12:31:11.409 | 50004 | 1.0.14.112 | 00:00:00:2f:09:fd | leaf1, leaf2 |      | 0    |

## Exploring Endpoint Locator Details

| LIST OF ACTIVE ENDPOINTS         |             |            |                   |             |                |      |  | 1-6 of 6 < > |
|----------------------------------|-------------|------------|-------------------|-------------|----------------|------|--|--------------|
| Time                             | VRF         | IP         | MAC               | Switch_Name | Port           | VLAN |  |              |
| November 17th 2018, 01:54:00.901 | myvrf_50000 | 60.1.1.134 | 00:50:56:97:d3:30 | leaf3       | Ethernet1/48   | 600  |  |              |
| November 17th 2018, 00:28:38.867 | myvrf_50000 | 60.1.1.135 | 00:50:56:97:3f:5b | leaf1       | port-channel48 | 600  |  |              |
| November 17th 2018, 00:28:38.845 | myvrf_50000 | 60.1.1.135 | 00:50:56:97:3f:5b | leaf2       | port-channel48 | 600  |  |              |

View surrounding documents | View single document

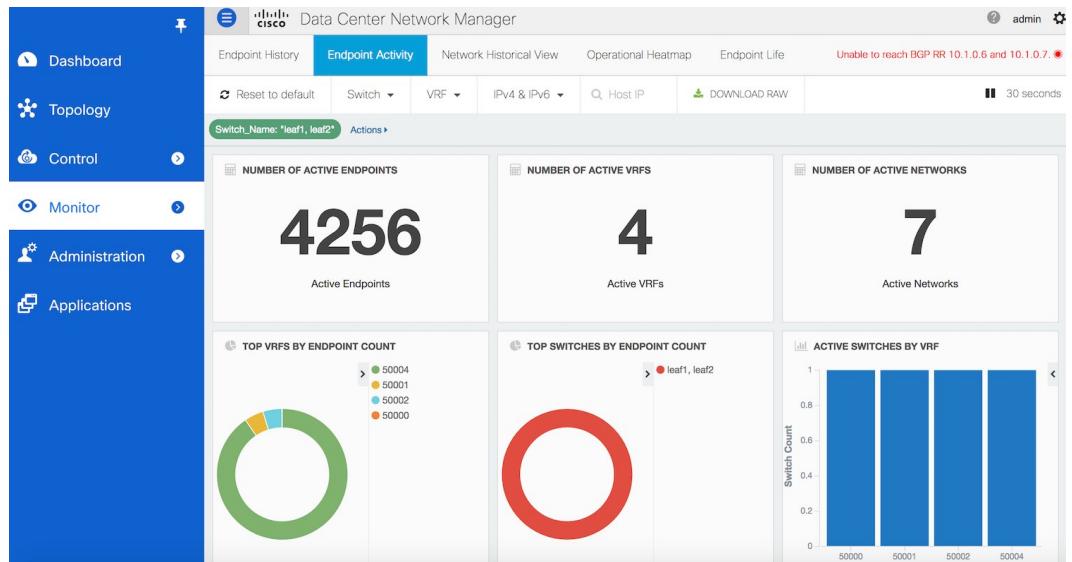
[Table](#) [JSON](#)

```

t Br_Domain    Q Q * 600
t Cluster      Q Q * 11.3.0.1:0
t EndpointIdentifier Q Q * IPv4:60.1.1.135:30000
t EndpointType   Q Q *
t Fabric_Id     Q Q * 4:evpn
t IP           Q Q * 60.1.1.135
t IPVersion     Q Q * IPv4
# L2_VNI       Q Q * 30,000
# L3_INT       Q Q * 600
# L3_VNI       Q Q * 50,000
t MAC          Q Q * 00:50:56:97:3f:5b
t Operation     Q Q * ACTIVE
t OperationStatus Q Q *

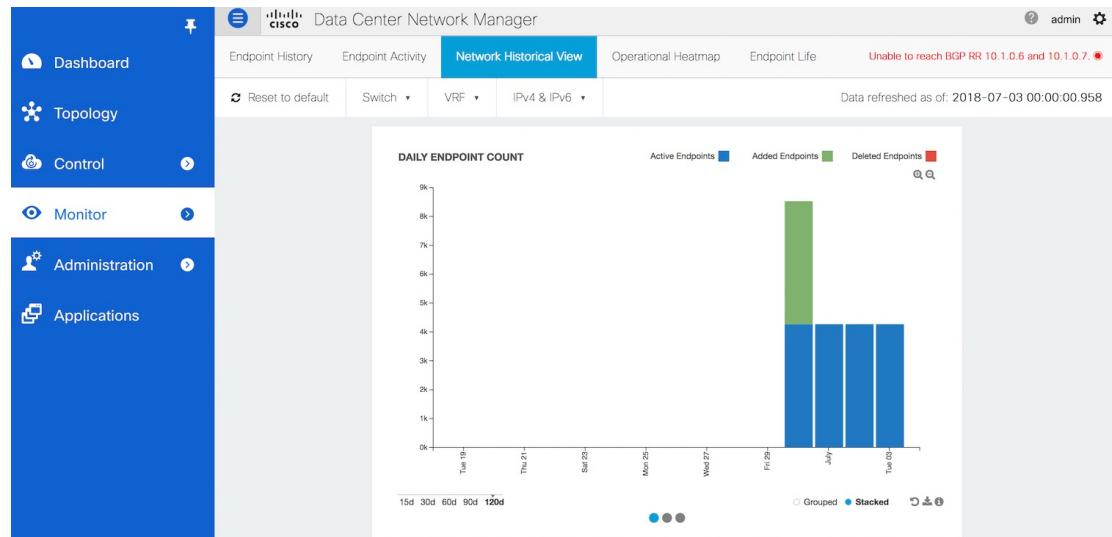
```

- Click the **Filter for value +** icon next to the L2\_VNI field. This selects the highlighted value (30000 in this example) and filters the search results based on that. In other words, the information of all active endpoints in the network associated with L2\_VNI 30000 is displayed on the dashboard. If instead, all endpoints that are not in the network L2\_VNI are required, click the – icon next to the L2\_VNI value of 30000. In the same manner, one can choose any combination of fields to get the set of endpoints matching the corresponding selected filter criteria.

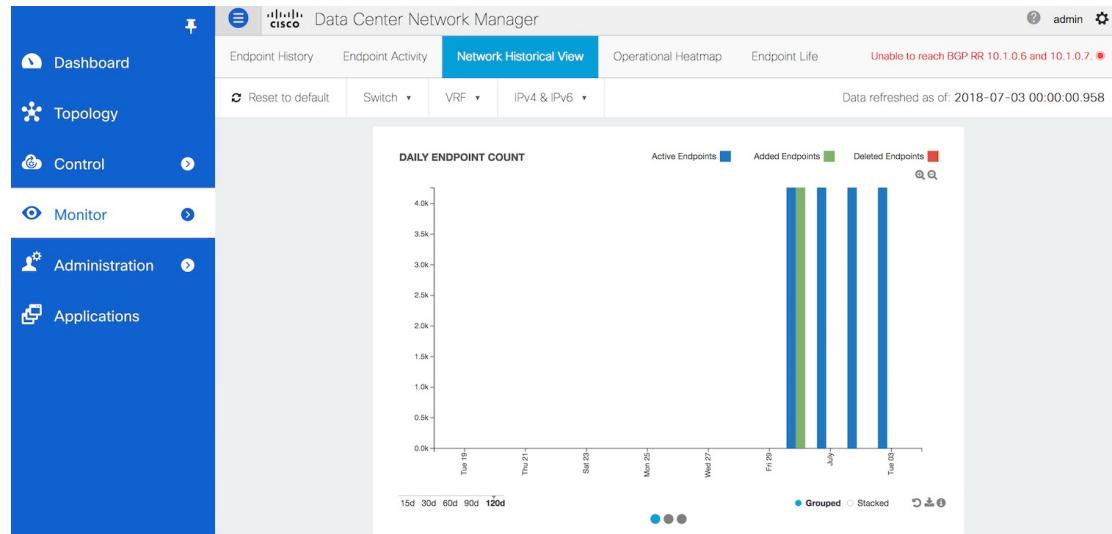


- Network Historical View**— The NHV view displays historical information of endpoints, networks, and VRFs (tenants) captured on a daily basis. These graphs are updated once a day at mid-night based on the DCNM server time. The time at which the data is refreshed/updated is listed at the top right. The idea is to provide a daily report of the Active, Added (New) and Deleted endpoints, networks, and VRFs respectively. If the same endpoint is added and removed on a day, then that contributes to an add count of 1 and a delete count of 1. Users can select one of the 3 dots at the bottom to toggle between the endpoints, networks, & VRF views. There are options to zoom in/out using zoom icons on top right. The users can also select the type of visualization with the choices being – Grouped or Stacked (shown below). Daily reports up to 180 days in the past can be displayed. Active endpoints/networks/VRFs are shown in blue color, deleted ones are shown in red color while the added ones are shown in green color. Every block in all screens is ‘clickable’ and the complete dataset associated with the selection, can be downloaded in csv format.

The historic endpoint count in ‘Stacked’ format is shown below:

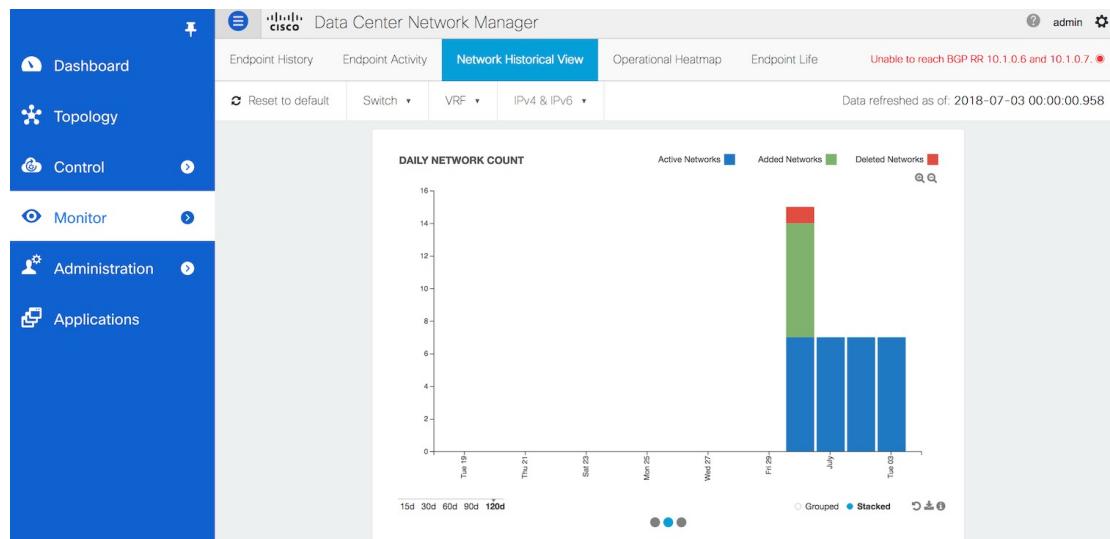


The same representation with the Grouped visualization selection is shown below:

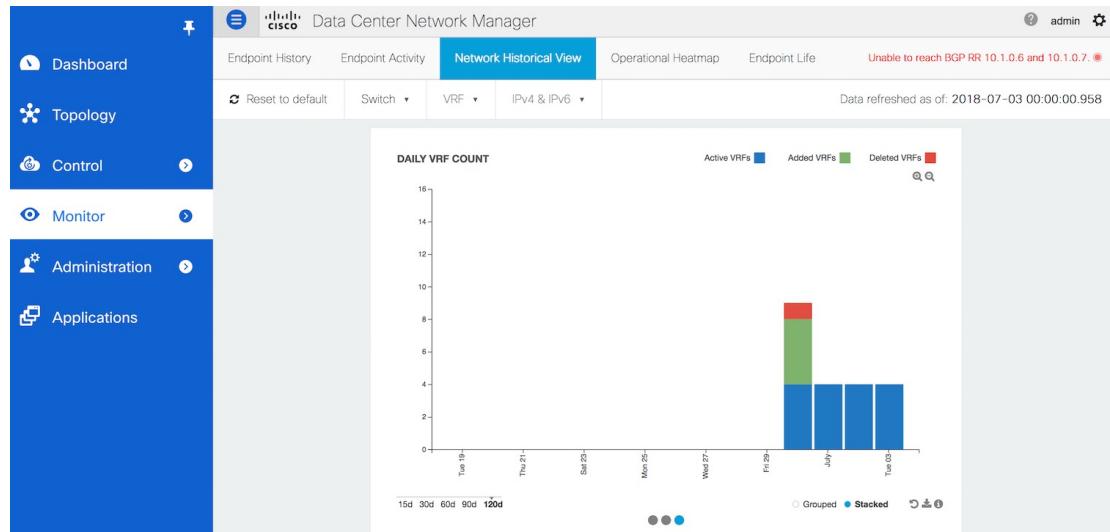


Similarly, the figure below depicts the historic network count in stacked format:

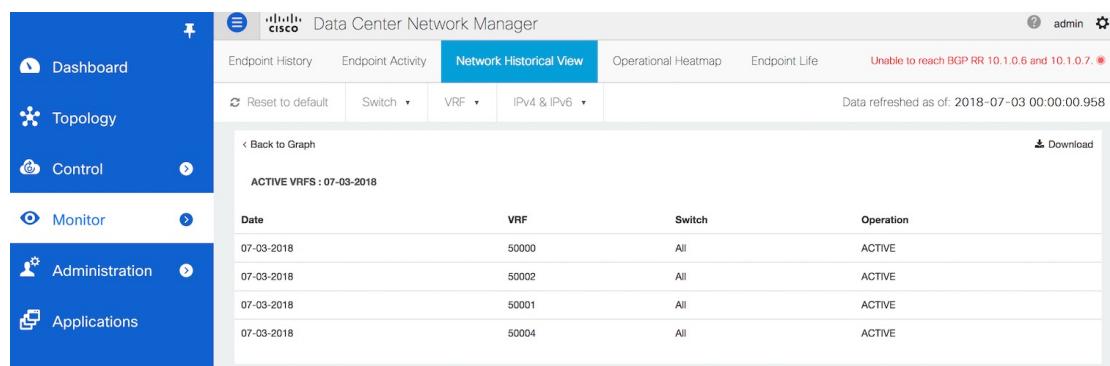
## Exploring Endpoint Locator Details



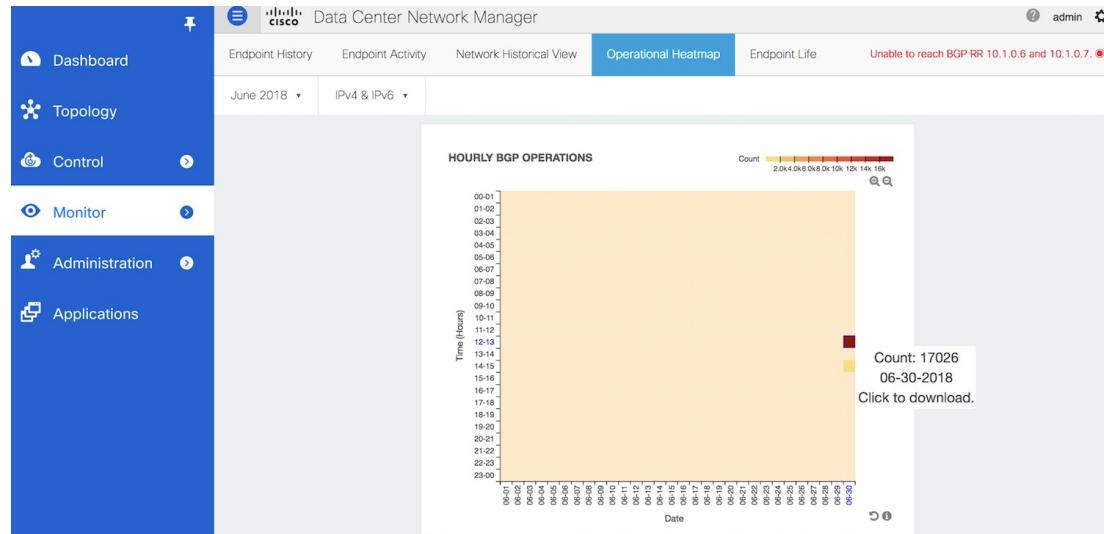
Along the same lines, the figure below depicts the historic vrf count:



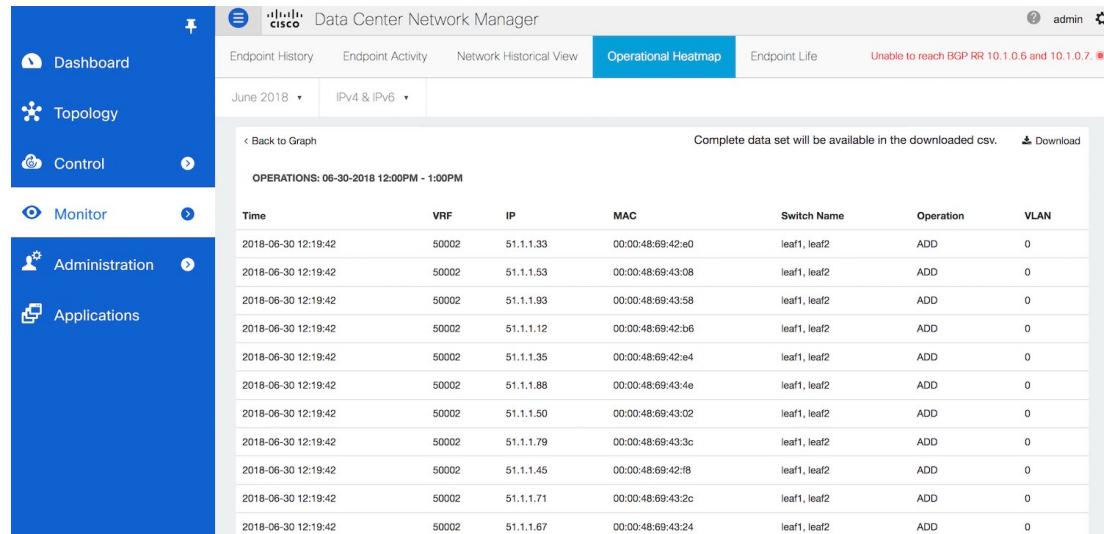
The figure below provides a sample screenshot of the endpoints added on 07-25-2019 obtained by clicking on the blue bar for that day.



- **Operational Heatmap**—This view displays a heat-map of all endpoint operations occurring in the fabric.



The heat-map is color coded and the intensity of the color varies based on the number of endpoint operations captured on an hourly basis. The break down is available per hour across dates, and user can see the details of operations that occurred during a particular hour on a particular day by clicking on the appropriate square. The figure below depicts the endpoint operations reported by BGP on 01-02-2018 between 12 and 1pm.

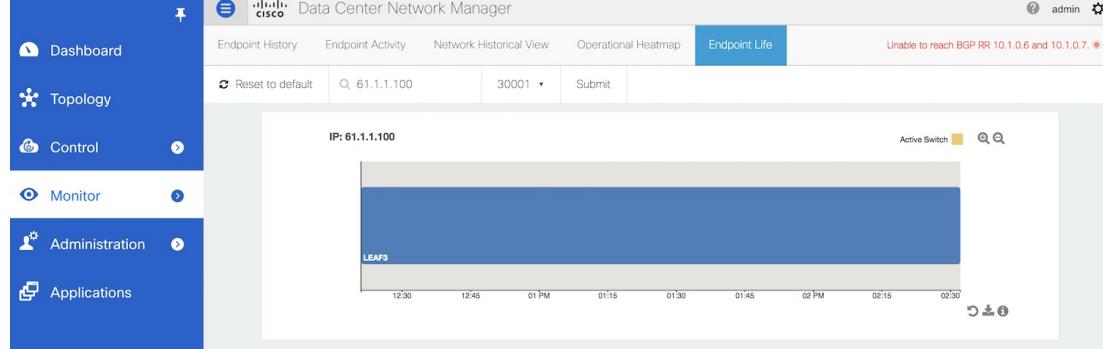


Again, as with the other views, the complete data set can be downloaded in csv format using the Download option. A sample screenshot of a downloaded csv file is shown below:

## LAN Telemetry

| 1  | Fabric_Id | IP        | MAC          | L2_VNI | L3_VNI | Switch_Nam  | Switch_Type | Switch_IP   | Origin_IP_0 | Origin_IP_1 | Origin_IP_2 | Origin_IP_3 | Switch_Next_Port | VLAN | L3_INT | Operation | EndpointType  | Timestamp | Seq_Num | VRF | Br_Domain | Clust |
|----|-----------|-----------|--------------|--------|--------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|------------------|------|--------|-----------|---------------|-----------|---------|-----|-----------|-------|
| 2  | 3:evpn    | 51.1.1.33 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 3  | 3:evpn    | 51.1.1.53 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 4  | 3:evpn    | 51.1.1.93 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 5  | 3:evpn    | 51.1.1.60 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 6  | 3:evpn    | 51.1.1.39 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 7  | 3:evpn    | 51.1.1.88 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 8  | 3:evpn    | 51.1.1.50 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 9  | 3:evpn    | 51.1.1.79 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 10 | 3:evpn    | 51.1.1.45 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 11 | 3:evpn    | 51.1.1.71 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 12 | 3:evpn    | 51.1.1.67 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 13 | 3:evpn    | 51.1.1.38 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 14 | 3:evpn    | 51.1.1.27 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 15 | 3:evpn    | 51.1.1.94 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 16 | 3:evpn    | 51.1.1.90 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 17 | 3:evpn    | 51.1.1.49 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 18 | 3:evpn    | 51.1.1.56 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 19 | 3:evpn    | 51.1.1.60 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 20 | 3:evpn    | 51.1.1.83 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 21 | 3:evpn    | 51.1.1.18 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 22 | 3:evpn    | 51.1.1.57 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 23 | 3:evpn    | 51.1.1.61 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 24 | 3:evpn    | 51.1.1.12 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 25 | 3:evpn    | 51.1.1.19 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 26 | 3:evpn    | 51.1.1.65 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |
| 27 | 3:evpn    | 51.1.1.75 | 00:00:48:69x | 30009  | 50002  | leaf1_leaf2 | NRK         | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 0.0.0.0     | 10.2.0.1         | 0    | 0 ADD  |           | Sun Jul 01 2C | 0         | 50002   | 0   | 10.2.1    |       |

- Endpoint Life**—This view displays a time line of a particular endpoint in its entire existence within the fabric. Specifically, given an identity of an endpoint in terms of its IP address and VRF/Network-identifier, the output displays the list of switches that an endpoint was present under including the associated start and end dates. This view is essentially the network life view of an endpoint. If the endpoint is viewed as active by the network, it will have a band here. If an endpoint is dual-homed, then there will be 2 horizontal bands reporting the endpoint existence, one band for each switch (typically the vPC pair of switches). As endpoints move within the network, for example with VM move, this view provides a succinct and intuitive pictorial view of this activity.



The underlying data that drives this view can also be downloaded in csv format (shown below) by clicking on download icon on right bottom corner.

| A  | B           | C           | D                     | E                                                       | F                                                       |
|----|-------------|-------------|-----------------------|---------------------------------------------------------|---------------------------------------------------------|
| 1  | Switch Name | VRF         | EndPointIdentifier    | Start Timestamp                                         |                                                         |
| 2  | n9k-12-vpc  | Beer:Corona | IPv4:60.1.1.134.30007 | Wed Dec 27 2017 21:41:33 GMT+0530 (India Standard Time) | Tue Jan 02 2018 18:56:32 GMT+0530 (India Standard Time) |
| 3  | n9k-13-vpc  | Beer:Corona | IPv4:60.1.1.134.30007 | Wed Dec 27 2017 21:41:49 GMT+0530 (India Standard Time) | Tue Jan 02 2018 18:56:33 GMT+0530 (India Standard Time) |
| 4  | n9k-12-vpc  | Beer:Corona | IPv4:60.1.1.134.30007 | Tue Jan 02 2018 20:54:21 GMT+0530 (India Standard Time) | Wed Jan 03 2018 14:25:02 GMT+0530 (India Standard Time) |
| 5  | n9k-13-vpc  | Beer:Corona | IPv4:60.1.1.134.30007 | Tue Jan 02 2018 20:54:21 GMT+0530 (India Standard Time) | Wed Jan 03 2018 14:24:45 GMT+0530 (India Standard Time) |
| 6  | n9k-12-vpc  | Beer:Corona | IPv4:60.1.1.134.30007 | Wed Jan 03 2018 14:35:40 GMT+0530 (India Standard Time) | Wed Jan 03 2018 16:09:09 GMT+0530 (India Standard Time) |
| 7  | n9k-13-vpc  | Beer:Corona | IPv4:60.1.1.134.30007 | Wed Jan 03 2018 14:35:44 GMT+0530 (India Standard Time) | Wed Jan 03 2018 16:09:10 GMT+0530 (India Standard Time) |
| 8  | n9k-12-vpc  | Beer:Corona | IPv4:60.1.1.134.30007 | Wed Jan 03 2018 16:15:18 GMT+0530 (India Standard Time) | Wed Jan 03 2018 18:02:49 GMT+0530 (India Standard Time) |
| 9  | n9k-13-vpc  | Beer:Corona | IPv4:60.1.1.134.30007 | Wed Jan 03 2018 16:15:18 GMT+0530 (India Standard Time) | Wed Jan 03 2018 18:02:48 GMT+0530 (India Standard Time) |
| 10 | n9k-12-vpc  | Beer:Corona | IPv4:60.1.1.134.30007 | Wed Jan 03 2018 18:35:09 GMT+0530 (India Standard Time) |                                                         |
| 11 | n9k-13-vpc  | Beer:Corona | IPv4:60.1.1.134.30007 | Wed Jan 03 2018 18:35:12 GMT+0530 (India Standard Time) | TRUE<br>TRUE                                            |

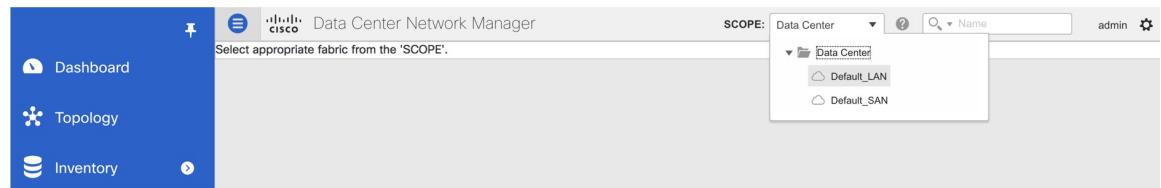
## LAN Telemetry

The LAN Telemetry menu includes the following submenus:

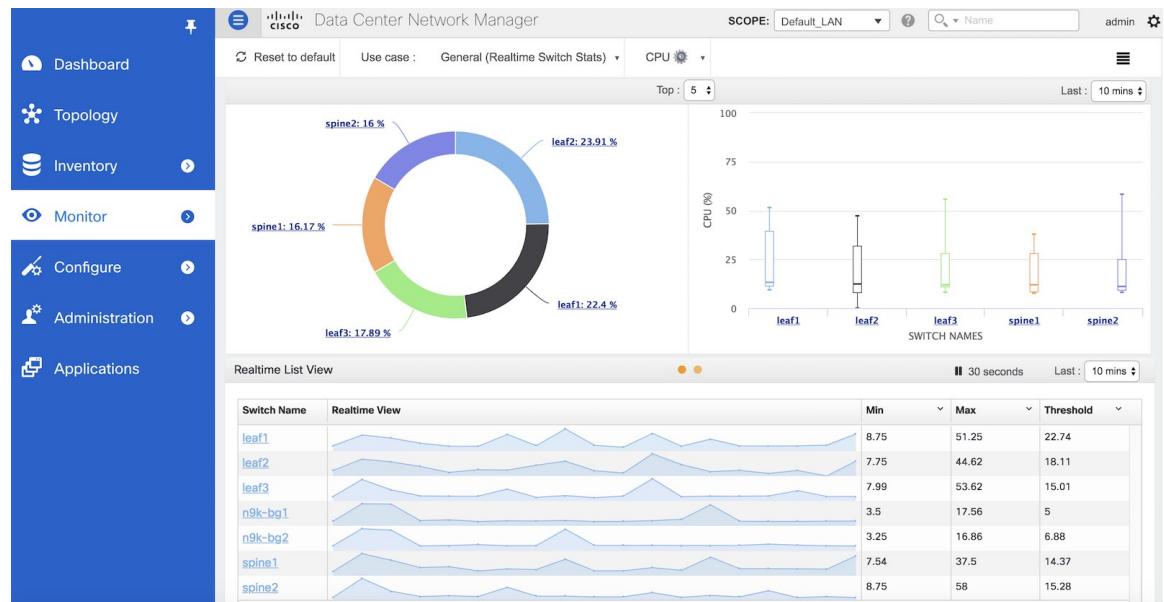
# Monitoring LAN Telemetry

## Procedure

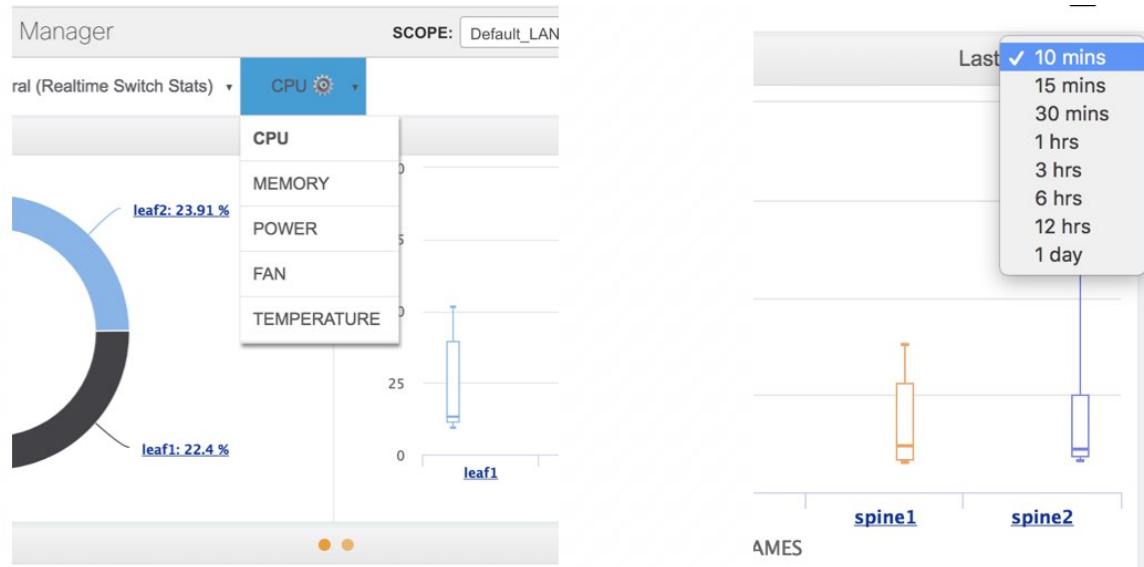
- Step 1** Once LAN telemetry has been successfully enabled, a new Telemetry Explore screen is available. You can navigate to the Telemetry Explore screen by choosing **Monitor > LAN Telemetry > Explore**. Select the fabric for which LAN telemetry has been enabled through the SCOPE at the top. After selecting one of the fabrics (Default\_LAN), LAN insights will appear.



- Step 2** There are three insights shown below through interactive visualizations depicting different aspects of Switch metric data based on the metric selected.



- Step 3** You can select metric from the metric dropdown as shown below(left):



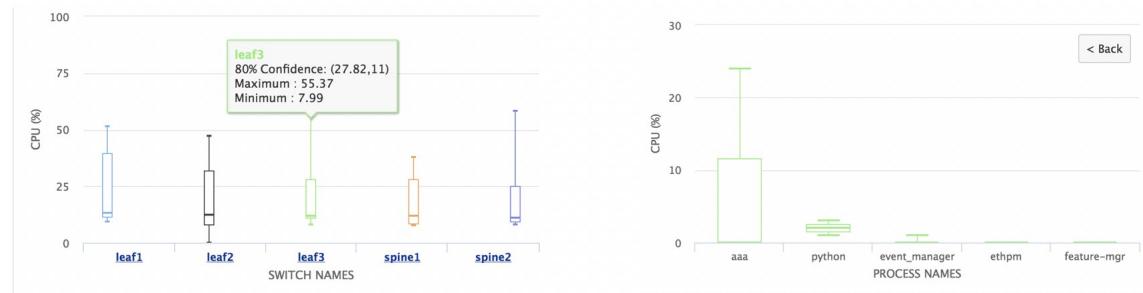
- Step 4** You can select the time-interval from Start With Last dropdown (shown above). The first two insights (Donut chart & Box plot) show the data for the selected time-interval.
- Step 5** You can select Top N switches from Top dropdown which is applicable for the first row insights (Donut chart & Box plot). For example if 5 is selected, Donut chart & Box plot is plotted for top five switches (based on selected metric value).



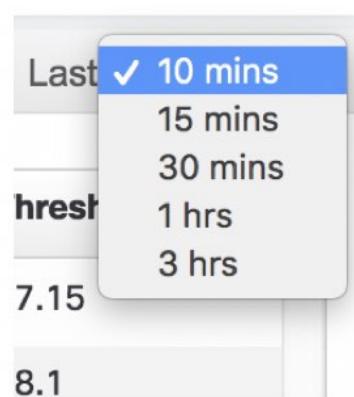
First insight is a Donut Chart showing numerical proportion of Top N switches based on selected metric. For example, donut chart below (left) shows the proportion of top five switches based on CPU usage values. When hovered, it shows the switch name and the corresponding metric value. Drilldown functionality is implemented in this insight too. Clicking on any of the slices goes to second level as shown below (right), which is a process-level donut chart of the selected switch showing numerical proportion of Top N processes of the selected switch based on selected metric. You can go back to switch-level donut chart using the back button that is provided in the top right corner.



Second insight is a Box Plot showing the variation of the selected metric for Top N switches in the selected time interval. Each Box in the visualization corresponds to one of the top N switches. It shows MIN, MAX metric value along with 80% Confidence interval of the selected metric. For example, Box plot below shows the top five switches and the variation of the CPU usage for each switch. When hovered, it shows Minimum and Maximum CPU Usage value with 80% Confidence interval. Drilldown functionality is implemented in this insight too. Clicking any of the boxes goes to second level as shown below (right), which is a process-level box plot of the selected switch showing MIN, MAX metric value, and 80% Confidence interval of Top N processes of the selected switch based on selected metric. You can go back to switch-level box plot using the Back button that is provided in the top right corner.



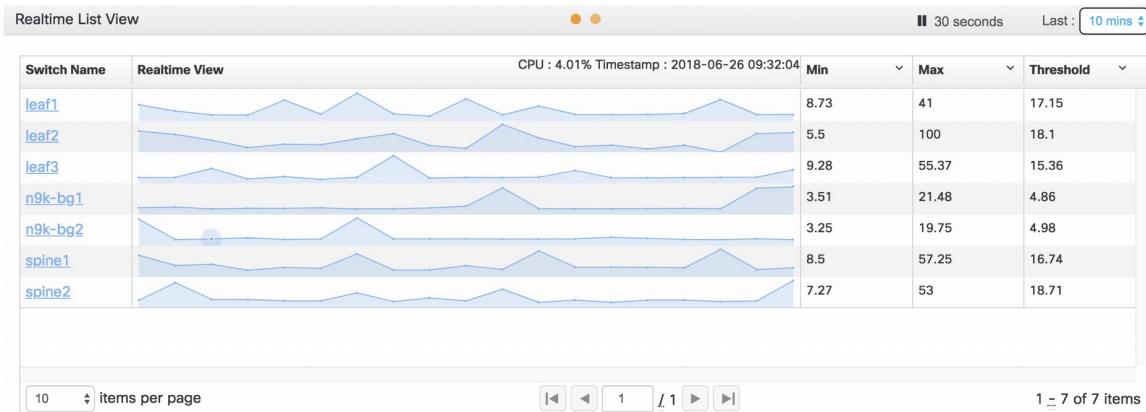
Third component is a Real time list view of all the switches, which were active for the selected metric for last 10 minutes (by default). You can change the time interval from the Last dropdown which will show the list of switches available for the selected time interval.



The data-table is updated every 30 seconds, which you can stop by enabling the pause button.

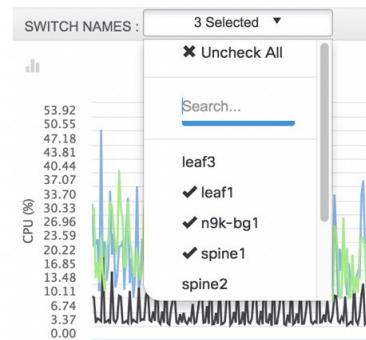
Real time List View is a tabular representation of switches for the selected metrics which shows the Minimum, Maximum and Threshold values of respective switches along with their real time view which is updated on every 30seconds. On hovering over the real-time view, you can view the exact usage per timestamp for the particular switch.

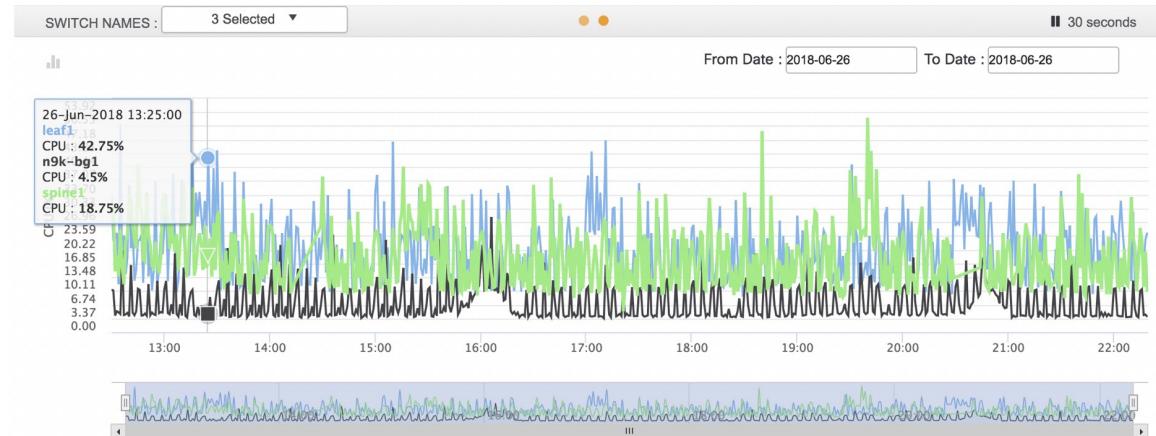
You can also sort the data-table in ascending/descending order on Min, Max & Threshold columns. The list of switches, which you can view per page, can also be customized. By default, only 10 switches per page are shown but you can change the “Items per page” that he wants to view by selecting the desired values from the dropdown.



Fourth component is a multiline graph, which is activated when you click the second orange dot present next to the “Realtime List View” heading.

This multiline graph shows real time selected metric data (shown below) for the selected switches where each line corresponds to a switch. Switches can be selected from the Switch Name dropdown (shown in right) which shows all switches in the fabric. You can use the time slider below to select any time range or choose from predefined Zoom options (1M, 5M, 10M and so on). You can also pause and resume the live data refreshing every 30 seconds using the pause button provide in the top right corner. When hovered, it shows the switch name, metric value and the corresponding timestamp.

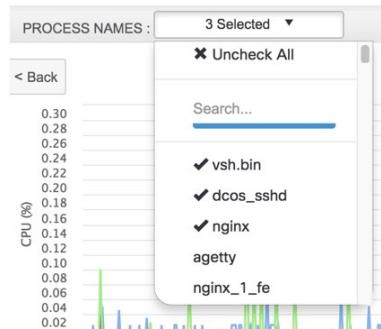




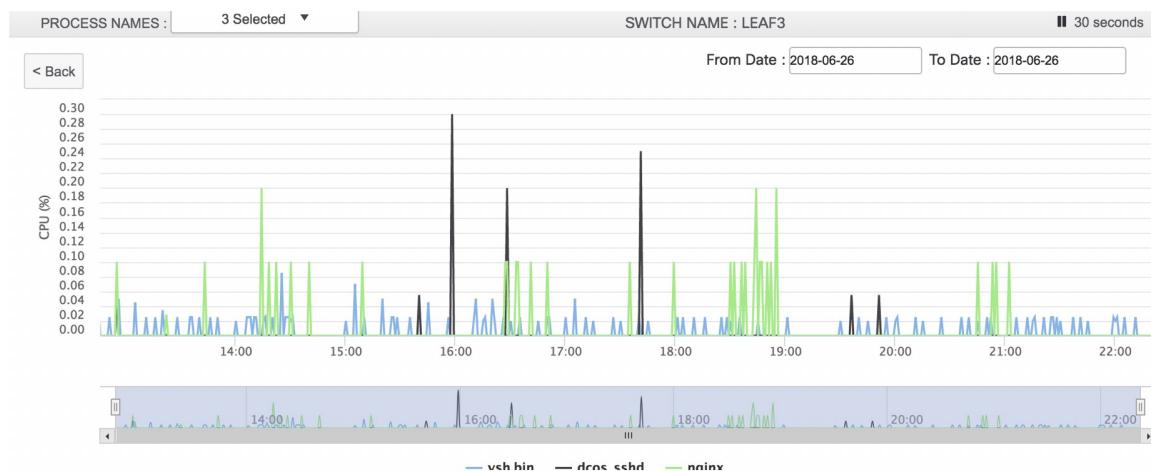
Advanced setting icon placed on the top left corner of the graph gives you an option to view events and analytics pertaining to the list of selected switches. This button is disabled by default and you have to click it to activate the button. Advanced setting icon is currently only available for CPU and MEMORY metrics. You can select up to 2 items from the advanced setting list.

You can also change the dates from the "From Date" and "To Date" date picker and select the respective dates for which he wants to view the real time data. By default, data for current date is displayed.

Drilldown functionality is implemented in this insight also and clicking on any of the lines goes to 2nd level as shown below, which is process level multiline graph of the selected switch showing real time selected metric data for the selected processes of the clicked switch. Processes can be selected from the Process Name dropdown (shown in right) which is populated by all process names of the clicked switch. You can go back to switch level multiline graph using the Back button that is provided in the top right corner.



## Monitoring LAN Telemetry

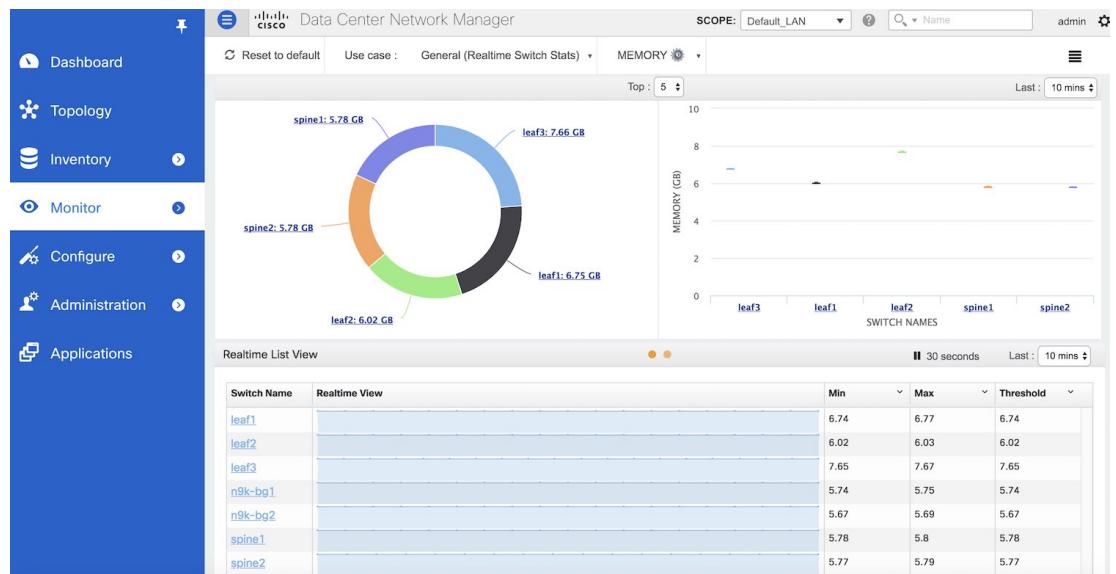


Drilldown functionality is available for all metrics and the second level metric value for each switch is as follows:

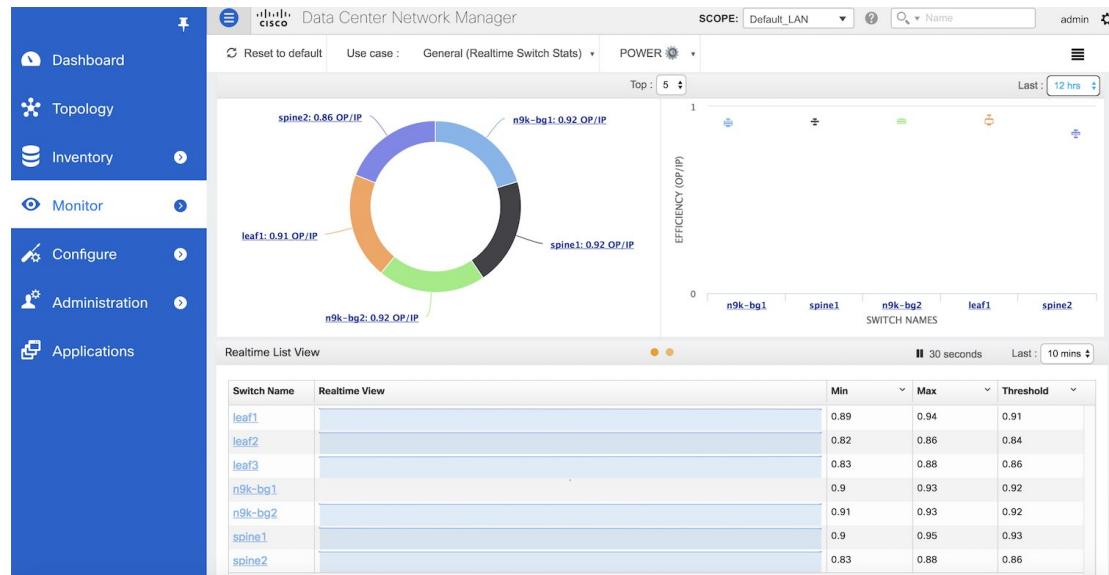
- CPU: Switch > Process level CPU Usage
- Memory: Switch > Process level Memory Usage
- Fan: Switch > Fan level Speed
- Power: Switch > Model level Efficiency
- Temperature: Switch > Temperature level Sensor

**Step 6** For the remaining four features, the functionality remains same as CPU, but the data that are shown on the page vary depending on the feature that you choose in the feature that you have selected.

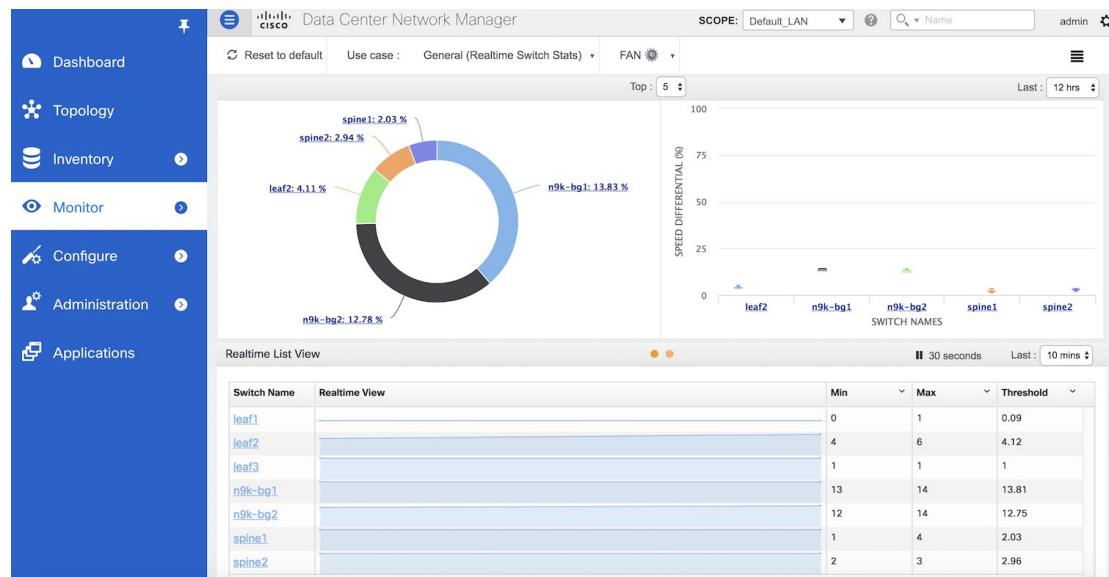
- Memory Data View: The memory dashboard depicts the actual memory consumption (RAM) on every switch in Gigabytes (GB). The per-process memory consumption is available at a later stage.



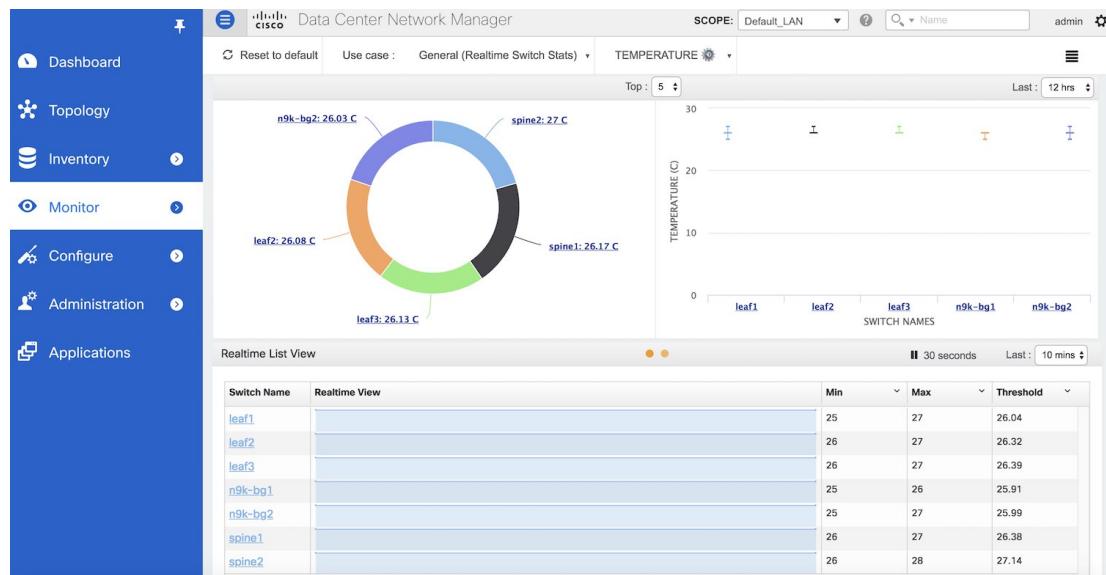
- Power Data View: The top-level view for the power dashboard depicts the efficiency of the power supplies. By definition, efficiency is Output-Power/Input-Power, which therefore results in a maximum efficiency of 1.0.



- Fan Data View—The top-level view for the fan dashboard depicts the speed difference between the various fans in the system. In the typical case, the expectation is that all the fans in the same tray, operate at more or less the same speed.



- Temperature Data View—The top-level view for the temperature dashboard displays information about the temperature level sensor.



## Alarms

The Alarms menu includes the following submenus:

### Viewing Alarms and Events

You can view the alarms, cleared alarms, and events.

#### Procedure

**Step 1** Choose .

**Step 2** Choose any of the following tabs.

- **Alarms:** This tab displays the alarms that are generated for various categories. This tab displays information such as ID (optional), Severity, Failure Source, Name, Category, Acknowledged, Creation Time, Last Updated (optional), Policy, and Message. You can specify the **Refresh Interval** in this tab. You can select one or more alarms and then acknowledge or unacknowledge their status using the **Change Status** drop-down list. In addition, you can select one or more alarms and then click the **Delete** button to delete them.
- **Cleared Alarms:** This tab displays the cleared alarms. This tab displays information such as ID (optional), Severity, Failure Source, Name, Category, Acknowledged, Creation Time, Cleared At (optional), Cleared By, Policy, and Message. You can select one or more alarms and then click the **Delete** button to delete them.
- **Events:** This tab displays the events that are generated for the switches. This tab displays information such as **Ack**, **Acknowledged user**, **Group**, **Switch**, **Severity**, **Facility**, **Type**, **Count**, **Last Seen**, and **Description**. You can select one or more events and then acknowledge or unacknowledge their status.

using the **Change Status** drop-down list. In addition, you can select one or more alarms and then click the **Delete** button to delete them. If you want to delete all events, click the **Delete All** button.

## Monitoring and Adding Alarm Policies

You can add alarm policies for the following:

- **Device Health:** Device health policies enable you to create alarms when Device ICMP Unreachable, Device SNMP Unreachable, or Device SSH Unreachable. Also, these policies enable you to monitor chassis temperature, CPU, and memory usage.
- **Interface Health:** Interface health policies enable you to monitor Up or Down, Packet Discard, Error, Bandwidth details of the interfaces. By default all interfaces are selected for monitoring.
- **Syslog Alarm:** Syslog Alarm Policy defines a pair of Syslog messages formats; one which raises the alarm, and one which clears the alarm.

### Before you begin

If you have created a self-signed certificate or imported an SSL certificate to the keystore, you must copy the new fmserver.jks located at /usr/local/cisco/dcm/wildfly-10.1.0.Final/standalone/configuration to /etc/elasticsearch. If you do not copy the fmserver.jks file to the elasticsearch directory, you will not be able to get the Alarms and Policies. As the elasticsearch database will be stabilizing, you cannot configure any Alarm Policy on the Cisco DCNM Web UI **Monitor > Alarms > Alarm Policies**.

### Procedure

**Step 1** Choose **Monitor > Alarms > Alarm Policies**.

**Step 2** Select the **Enable Alarms** check box to enable alarm policies.

**Step 3** From the **Add** drop-down list, choose any of the following:

- Device Health Policy: Select the devices for which you want to create policies. Specify the policy name, description, CPU Utilization parameters, Memory Utilization parameters, Environment Temperature parameters, device availability, and device features.
- Interface Health Policy: Select the devices for which you want to create policies. Specify the policy name, description, link-state, Bandwidth (In/Out), Inbound errors, Outbound errors, Inbound Discards, and Outbound Discards.
- Syslog Alarm Policy: Select the devices for which you want to create policies and then specify the following parameters.
  - Devices: Define the scope of this policy. Select individual devices or all devices to apply this policy.
  - Policy Name: Specify the name for this policy. It must be unique.
  - Description: Specify a brief description for this policy.
  - Severity: Define the severity level for this syslog alarm policy. Choices are: Critical, Major, Minor, and Warning.

- Identifier: Specify the identifier portions of the raise & clear messages.
- Raise Regex: Define the format of a syslog raise message. The syntax is as follows:  
**Facility-Severity-Type: Message**
- Clear Regex: Define the format of a syslog clear message. The syntax is as follows:  
**Facility-Severity-Type: Message**

**Table 10: Example 1**

| <b>Identifier</b> | <b>ID1-ID2</b>                                                              |
|-------------------|-----------------------------------------------------------------------------|
| Raise Regex       | ETHPORT-5-IF_ADMIN_UP: Interface Ethernet15/1 is admin up .                 |
| Clear Regex       | ETHPORT-5-IF_DOWN_NONE: Interface Ethernet15/1 is down (Transceiver Absent) |

In the above example, the regex expressions are part of the syslog messages that appear in the terminal monitor.

**Table 11: Example 2**

| <b>Identifier</b> | <b>ID1-ID2</b>                                         |
|-------------------|--------------------------------------------------------|
| Raise Regex       | ETH_PORT_CHANNEL-5-PORT_DOWN: \$(ID1): \$(ID2) is down |
| Clear Regex       | ETH_PORT_CHANNEL-5-PORT_UP: \$(ID1): \$(ID2) is up     |

**Table 12: Example 3**

| <b>Identifier</b> | <b>ID1-ID2</b>                                                             |
|-------------------|----------------------------------------------------------------------------|
| Raise Regex       | ETHPORT-5-IF_SFP_WARNING: Interface \$(ID1), High Rx Power Warning         |
| Clear Regex       | ETHPORT-5-IF_SFP_WARNING: Interface \$(ID1), High Rx Power Warning cleared |

**Step 4** Click **OK** to add the policy.

### Syslog Messages in Terminal Monitor and Console

The following examples show how the syslog messages appear in the terminal monitor and the console. The regex expression is matched with the part of the syslog messages after the % sign.

```
leaf-9516# terminal monitor
leaf-9516# conf t
leaf-9516(config)# int e15/1-32
leaf-9516(config-if-range)# no shut
```

```

2019 Aug 2 04:41:27 leaf-9516 %ETHPORT-5-IF_ADMIN_UP: Interface
Ethernet15/1 is admin up .
2019 Aug 2 04:41:27 leaf-9516 %ETHPORT-5-IF_DOWN_NONE: Interface
Ethernet15/1 is down (Transceiver Absent)
2019 Aug 2 04:41:27 leaf-9516 %ETHPORT-5-IF_ADMIN_UP: Interface
Ethernet15/2 is admin up .
2019 Aug 2 04:41:27 leaf-9516 %ETHPORT-5-IF_DOWN_NONE: Interface
Ethernet15/2 is down (Transceiver Absent)
2019 Aug 2 04:41:28 leaf-9516 %ETHPORT-5-IF_ADMIN_UP: Interface
Ethernet15/3 is admin up .

```

The syslog messages in the console have a similar format as they would appear in the terminal monitor, except for the additional port information enclosed in the %% signs. However, the regex expression is matched with the part of the syslog messages after the last % sign.

```

SR-leaf1# 2019 Aug 26 23:55:45 SR-leaf1 %% VDC-1 %% %PLATFORM-1-
PFM_ALERT: FAN_BAD: fan6
2019 Aug 26 23:56:15 SR-leaf1 %% VDC-1 %% %PLATFORM-1-PFM_ALERT:
FAN_BAD: fan6
2019 Aug 26 23:56:18 SR-leaf1 %% VDC-1 %% %ASCII-CFG-2-CONF_CONTROL:
System ready
2019 Aug 26 23:56:25 SR-leaf1 %% VDC-1 %% %PLATFORM-1-PFM_ALERT:
FAN_BAD: fan6
2019 Aug 26 23:56:35 SR-leaf1 %% VDC-1 %% %PLATFORM-1-PFM_ALERT:
FAN_BAD: fan6
2019 Aug 26 23:56:39 SR-leaf1 %% VDC-1 %% %VMAN-2-ACTIVATION_STATE:
Successfully activated virtual service 'guestshell'
2019 Aug 26 23:56:39 SR-leaf1 %% VDC-1 %% %VMAN-2-GUESTSHELL_ENABLED:
The guest shell has been enabled. The command 'guestshell' may be used
to access it, 'guestshell destroy' to remove it.
2019 Aug 26 23:56:45 SR-leaf1 %% VDC-1 %% %PLATFORM-2-FAN_REMOVED: Fan
module 5 (Serial number ) Fan5(sys_fan5) removed
2019 Aug 26 23:56:45 SR-leaf1 %% VDC-1 %% %PLATFORM-1-PFM_ALERT:
System will shutdown in 2 minutes 0 seconds due to fan policy
_pfmd_fanabsent_any_singlefan.
2019 Aug 26 23:56:45 SR-leaf1 %% VDC-1 %% %PLATFORM-1-PFM_ALERT:
FAN_BAD: fan6
2019 Aug 26 23:56:54 SR-leaf1 %% VDC-1 %% %PLATFORM-1-PFM_ALERT:
System will shutdown in 1 minutes 40 seconds due to fan policy
_pfmd_fanabsent_any_singlefan.
2019 Aug 26 23:56:54 SR-leaf1 %% VDC-1 %% %PLATFORM-1-PFM_ALERT:
FAN_BAD: fan6
2019 Aug 26 23:57:03 SR-leaf1 %% VDC-1 %% %PLATFORM-2-FANMOD_FAN_OK:
Fan module 5 (Fan5(sys_fan5) fan) ok
2019 Aug 26 23:57:03 SR-leaf1 %% VDC-1 %% %PLATFORM-1-PFM_ALERT:
FAN_BAD: fan6

```

## Activating Policies

After you create new alarm policies, activate them.

### Procedure

- 
- Step 1** Choose **Monitor > Alarms > Policies**.
  - Step 2** Select the policies that you want to activate and then click the **Activate** button.
-

## Deactivating Policies

You can deactivate the active alarm policies.

### Procedure

- 
- Step 1** Choose **Monitor > Alarms > Policies**.
  - Step 2** Select the policies that you want to deactivate and then click the **Deactivate** button.
- 

## Importing Policies

You can create alarm policies using the import functionality.

### Procedure

- 
- Step 1** Choose **Monitor > Alarms > Policies** and then click the **Import** button.
  - Step 2** Browse and select the policy file saved on your computer.
- You can only import policies in text format.
- 

## Exporting Policies

You can export the alarm policies into a text file.

### Procedure

- 
- Step 1** From the menu bar, choose **Monitor > Alarms > Policies**.
  - Step 2** Click the **Export** button and then select a location on your computer to store the exported file.
- 

## Editing Policies

### Procedure

- 
- Step 1** From the menu bar, choose **Monitor > Alarms > Policies**.
  - Step 2** Select the policy that you want to edit.
  - Step 3** Click the **Edit** button and then make necessary changes.
  - Step 4** Click the **OK** button.
-

## Deleting Policies

### Procedure

---

- Step 1** From the menu bar, choose **Monitor > Alarms > Policies**.
  - Step 2** Select the policy that you want to delete.
  - Step 3** Click the **Delete** button. The policy is deleted.
-





# CHAPTER 6

## Administration

---

This chapter contains the following topics:

- [DCNM Server, on page 237](#)
- [Management Users, on page 244](#)
- [Performance Setup, on page 248](#)
- [Event Setup, on page 248](#)
- [Credentials Management, on page 253](#)

## DCNM Server

The DCNM Server menu includes the following submenus:

### Starting, Restarting, and Stopping Services

To clean up the performance manager database (PM DB) stale entries, start, restart, or stop a service, from the Cisco DCNM Web UI, perform the following steps:

#### Procedure

---

**Step 1** Choose **Administration > DCNM Server > Server Status**.

The **Status** window appears that displays the server details.

**Step 2** In the **Actions** column, click the **Re(start)** icon to start or restart services, and click the **Stop** icon to stop services.

**Step 3** In the **Actions** column, click the **Delete** icon to clean up PM DB stale entries.

**Step 4** You can see the latest status in the **Status** column.

---

#### What to do next

See the latest status in the **Status** column.

#### Using the Commands Table

## Viewing Log Information

The commands table contains links to commands that launch new dialog boxes to provide information about the server status and server administrative utility scripts. These commands can be directly executed on the server CLI as well.

- **ifconfig:** click this link to view information about interface parameters, IP address, and netmask used on the Cisco DCNM server.
- **appmgr status all:** click this link to view the DCNM server administrative utility script that checks the status of different services currently running.
- **clock:** click this link to view information about the server clock details such as time, zone information.



### Note

The commands section is applicable only for the OVA or ISO installations.

## Viewing Log Information

You can view the logs for performance manager, SAN management server, SME server, web reports, web server, and web services. These processes have no corresponding GUI that allows you to view information about these log files. If you see errors, preserve these files for viewing.



### Note

Logs cannot be viewed from a remote server in a federation.

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

### Procedure

#### Step 1 Choose Administration > DCNM Server > Logs.

You see a tree-based list of logs in the left column. Under the tree, there is a node for every server in the federation. The log files are under the corresponding server node.

#### Step 2 Click a log file under each node of the tree to view it on the right.

#### Step 3 Double-click the tree node for each server to download a ZIP file containing log files from that server.

#### Step 4 Click the **Print** icon on the upper right corner to print the logs.

## Server Properties

You can set the parameters that are populated as default values in the DCNM server.

The backup configuration files are stored in the following path:  
`/usr/local/cisco/dcm/dcsm/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 Cisco DCNM LAN Fabric installation, the backup is taken per fabric and not per device. If the number of backup files exceeds the value entered in the field, the first version of the backup is deleted

to accommodate the latest version. For example, if the value entered in the field is **50** and when the 51<sup>st</sup> version of the fabric is backed up, the first backup file is deleted.

To set the parameters of the DCNM server from the Cisco DCNM Web UI, perform the following steps:

#### Procedure

---

- Step 1** Choose **Administration > DCNM Server > Server Properties**.
- Step 2** Click **Apply Changes** to save the server settings.
- 

## Modular Device Support

To support any new hardware that does not require many major changes, a patch can be delivered instead of waiting for the next DCNM release. **Modular Device Support** helps to deliver and apply the DCNM patch releases. An authorized DCNM administrator can apply the patch to the production setup. Patch releases are applicable for the following scenarios:

- Support any new hardware, like chassis or line cards
- Support latest NX-OS versions
- Support critical fixes as patches

To view the patch details from Cisco DCNM Web UI, perform the following steps:

#### Procedure

---

- Step 1** Choose **Administration > DCNM Server > Modular Device Support**.

You see the **DCNM Servers** column on the left in the window and **Modular Device support information** window on the right.

- Step 2** Expand **DCNM Servers** to view all the DCNM servers.

It includes the list of patches installed along with the version number, corresponding platforms supported, chassis supported, NX-OS version supported, PID supported, backup directory and the last patch deployment time in the **Modular Device support information** table.

---

#### What to do next

For more details about how to apply and rollback a patch, go to <http://www.cisco.com/go/dcnm> for more information.

## Managing Licenses

You can view the existing Cisco DCNM licenses by choosing **Administration > DCNM Server > License**. You can view and assign licenses in the following tabs:

**License Assignments**

- License Assignments
- Server License Files



**Note** By default, the **License Assignments** tab appears.

The following table displays the SAN and LAN license information.

| Field                            | Description                                                                                      |
|----------------------------------|--------------------------------------------------------------------------------------------------|
| License                          | Specifies SAN or LAN.                                                                            |
| Free/Total Server-based Licenses | Specifies the number of free licenses that are purchased out of the total number of licenses.    |
| Unlicensed/Total (Switches/VDCs) | Specifies the number of unlicensed switches or VDCs out of the total number of switches or VDCs. |
| Need to Purchase                 | Specifies the number of licenses to be purchased.                                                |

This section includes the following topics:

## License Assignments

The following table displays the license assignment details for every switch or VDC.

| Field          | Description                                                                                                                                                                                                                                    |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Group          | Displays if the group is fabric or LAN.                                                                                                                                                                                                        |
| Switch Name    | Displays the name of the switch.                                                                                                                                                                                                               |
| WWN/Chassis ID | Displays the world wide name or Chassis ID.                                                                                                                                                                                                    |
| Model          | Displays the model of the device. For example, DS-C9124 or N5K-C5020P-BF.                                                                                                                                                                      |
| License State  | Displays the license state of the switch that can be one of the following: <ul style="list-style-type: none"> <li>• Permanent</li> <li>• Eval</li> <li>• Unlicensed</li> <li>• Not Applicable</li> <li>• Expired</li> <li>• Invalid</li> </ul> |
| License Type   | Displays if the license is a switch-based embedded license or a server-based license.                                                                                                                                                          |

| Field            | Description                                                                                                                                              |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Eval Expiration  | Displays the expiry date of the license.<br><b>Note</b> Text under the <b>Eval Expiration</b> column is in red for licenses, which expire in seven days. |
| Assign License   | Select a row and click this option on the toolbar to assign the license.                                                                                 |
| Unassign License | Select a row and click this option on the toolbar to unassign the license.                                                                               |
| Assign All       | Click this option on the toolbar to refresh the table and assign the licenses for all the items in the table.                                            |
| Unassign All     | Click this option on the toolbar to refresh the table and unassign all the licenses.                                                                     |



**Note** You must have network administrator privileges to assign or unassign licenses.

When the fabric is first discovered and if the switch does not have a valid switch-based license, a license is automatically assigned to the fabric from the file license pool until no more licenses are left in the pool. If you have an existing fabric and a new switch is added to the fabric, the new switch is assigned a license if one is available in the file license pool and if it does not already have a switch-based license.

## Server License Files

### Server License Files

The following table displays the Cisco DCNM server license fields.

| Field            | Description                                                                                                                                         |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Filename         | Specifies the license file name.                                                                                                                    |
| Feature          | Specifies the licensed feature.                                                                                                                     |
| PID              | Specifies the product ID.                                                                                                                           |
| LAN (Free/Total) | Displays the number of free versus total licenses for LAN.                                                                                          |
| Eval Expiration  | Displays the expiry date of the license.<br><b>Note</b> Text in the <b>Eval Expiration</b> field is in Red for licenses that expires in seven days. |

### Adding Cisco DCNM Licenses

To add Cisco DCNM licenses from Cisco DCNM, perform the following steps:

#### Before you begin

You must have network administrator privileges to complete the following procedure.

### Procedure

---

**Step 1**

Choose **Administration > DCNM Server > License** to start the license wizard.

**Step 2**

Choose the **Server License Files** tab.

The valid Cisco DCNM-LAN license files are displayed.

Ensure that the security agent is disabled when you load licenses.

**Step 3**

Download the license pack file that you received from Cisco into a directory on the local system.

**Step 4**

Click **Add License File** and select the license pack file that you saved on the local machine.

The file is uploaded to the server machine, which is saved into the server license directory, and then loaded on to the server.

**Note**

Ensure that you do not edit the contents of the .lic file or the Cisco DCNM software ignores any features that are associated with that license file. The contents of the file are signed and must remain intact. When you accidentally copy, rename, or insert the license file multiple times, the duplicate files are ignored, but the original is counted.

## Native HA

### Procedure

---

**Step 1**

By default, DCNM is bundled with an embedded database engine PostgreSQL. The native DCNM HA is achieved by two DCNMs running as **Active / Warm Standby**, with their embedded databases synchronized in real time. So once the active DCNM is down, the standby takes over with the same database data and resume the operation. The *standby host database down* scenario is documented after this procedure.

**Step 2**

From the menu bar, choose **Administration > DCNM Server > Native HA**.

You see the **Native HA** window.

**Step 3**

You can allow manual failover of DCNM to the standby host by clicking the **Failover** button, and then click **OK**.

- Alternatively, you can initiate this action from the Linux console.

1. SSH into the DCNM active host.

2. Enter " "/usr/share/heartbeat/hb\_standby"

**Step 4**

You can allow manual syncing database and disk files to standby host by clicking **Force Sync**, and then click **OK**.

**Step 5**

You can test or validate the HA setup by clicking **Test** and then click **OK**.

## What to do next

Some HA troubleshooting scenarios are noted in this sub section.

**The standby host database is down:** Typically, the DCNM database (PostgreSQL) is up on the active and standby hosts. In DCNM 10.1 and earlier versions, the standby database can be down due to a database synchronization failure.

- Enter “ps -ef | grep post”. You should see multiple postgres processes running. If not, it indicates that the database is down.
- Restore database data from a backup file that is created at the beginning of database synchronization. Change directory to “/usr/local/cisco/dcm/db”
- Check existence of file replication/ pgsql-standby-backup.tgz. If the file exists, restore database data files:

```
rm -rf      data/*
tar -zxf      replication/ pgsql-standby-backup.tgz      data
/etc/init.d/postgresql-9.4 start
ps -ef | grep post
```

The active DCNM host will synchronize the two databases.

**The TFTP server is not bound to the eth1 VIP address on the active host:** The TFTP server should run on the active host (not on the standby host), and it should be bound to the eth1 VIP address. In some setups, the bind address is not the VIP address, as per the TFTP configuration file, and this could cause issues when switches try to use TFTP.

- Enter “grep bind /etc/xinetd.d/tftp” to check if the TFTP configuration file has the right bind address. If the displayed IP address is not the eth1 VIP address, then change the bind address to the VIP address. Repeat the procedure for the standby host. Update the bind address to the VIP address.
- Enter “/etc/init.d/xinetd restart” on the active host to restart TFTP.



**Note** The TFTP server can be started or stopped with the “appmgr start/stop ha-apps” command.

# Multi Site Manager

## Procedure

**Step 1** Multi-Site-Manager (MsM) provides a single pane for users to search for switches that are managed by DCNM globally. MsM can do realtime search to find out which switch globally handles the traffic for a given virtual machine based on IP address, name or mac address, and supporting VXLAN basing on segment ID as well. It provides hyperlink to launch the switch only. This window also plays the role of remote site registration. The registration only allows the current DCNM server to access the remote DCNM server or site. For the remote site to access the current DCNM server, registration is required on the remote site as well.

**Step 2** Choose **Administration > DCNM Server > Multi Site Manager**.

The MsM window displays the overall health or status of the remote site and the application health.

- Step 3** You can search by **Switch, VM IP, VM Name, MAC, and Segment ID**.
- Step 4** You can add a new DCNM server by clicking **+Add DCNM Server**. The **Enter Remote DCNM Server Information** window opens. Fill in the information that is required and click **OK** to save.
- Step 5** Click **Refresh All Sites** to display the updated information.

---

## Management Users

The Management Users menu includes the following submenus:

### Remote AAA

To configure remote AAA from the Cisco DCNM Web UI, perform the following steps:

#### Procedure

---

- Step 1** Choose **Administration > Management Users > Remote AAA Properties**.  
The AAA properties configuration window appears.
- Step 2** Use the radio button to select one of the following authentication modes:
- **Local**: In this mode the authentication authenticates with the local server.
  - **Radius**: In this mode the authentication authenticates against the RADIUS servers specified.
  - **TACACS+**: In this mode the authentication authenticates against the TACAS servers specified.
  - **Switch**: In this mode the authentication authenticates against the switches specified.
  - **LDAP**: In this mode the authentication authenticates against the LDAP server specified.
- Step 3** Click **Apply**.
- Note**      Restart the Cisco DCNM LAN services if you update the Remote AAA properties.

---

### Local

#### Procedure

---

- Step 1** Use the radio button and select **Local** as the authentication mode.  
**Step 2** Click **Apply** to confirm the authentication mode.

## Radius

### Procedure

- 
- Step 1** Use the radio button and select **Radius** as the authentication mode.
- Step 2** Specify the Primary server details and click **Test** to test the server.
- Step 3** (Optional) Specify the Secondary and Tertiary server details and click **Test** to test the server.
- Step 4** Click **Apply** to confirm the authentication mode.
- 

## TACACS+

### Procedure

- 
- Step 1** Use the radio button and select **TACACS+** as the authentication mode.
- Step 2** Specify the Primary server details and click **Test** to test the server.
- Step 3** (Optional) Specify the Secondary and Tertiary server details and click **Test** to test the server.
- Step 4** Click **Apply** to confirm the authentication mode.
- 

## Switch

### Procedure

- 
- Step 1** Use the radio button to select **Switch** as the authentication mode.  
DCNM also supports LAN switches with the IPv6 management interface.
- Step 2** Specify the Primary Switch name and click **Apply** to confirm the authentication mode.
- Step 3** (Optional) Specify the names for Secondary and Tertiary Switches.
- Step 4** Click **Apply** to confirm the authentication mode.
- 

## LDAP

### Procedure

- 
- Step 1** Use the radio button and select **LDAP** as the authentication mode.
- Step 2** In the **Host** field, enter DNS address of the host.
- Step 3** Click **Test** to test the AAA server. The **Test AAA Server** window pops out.
- Step 4** Enter a valid **Username** and **Password** in the **Test AAA Server** window.

A dialog box appears confirming the status of the AAA server test. If the test has failed, the **LDAP Authentication Failed** dialog box appears.

- Step 5** In the **Port** field, enter a port number.
  - Step 6** (Optional) Select the **SSL Enabled** check box, if SSL is enabled on the AAA server.
  - Step 7** In the **Base DN** field, enter the base domain name.
  - Step 8** In the **Filter** field, specify the filter parameters.
  - Step 9** Choose an option to determine a role by either **Attribute** or **Admin Group Map**.
  - Step 10** In the **Role Admin Group** field, enter the name of the role.
  - Step 11** In the **Map to DCNM Role** field, enter the name of the role to be mapped.
  - Step 12** In the **Access Map** field, enter the Role Based Access Control (RBAC) group to be mapped.
  - Step 13** Click **Apply Changes** icon on the upper right corner to apply the LDAP configuration.
- 

## Managing Local Users

As an admin user, you can use Cisco DCNM Web UI to create a new user, assign the role and associate one or more groups or scope for the user.

This section contains the following:

### Adding Local Users

#### Procedure

- Step 1** From the menu bar, choose **Administration > Management Users > Local**. You see the **Local Users** page.
  - Step 2** Click **Add User**.  
You see the **Add User** dialog box.
  - Step 3** Enter the username in the **User name** field.  
**Note** The username is case sensitive, but the username guest is a reserved name, which is not case sensitive. The guest user can only view reports. The guest user cannot change the guest password, or access the Admin options in DCNM Web Client.
  - Step 4** From the **Role** drop-down list, select a role for the user.
  - Step 5** In the **Password** field, enter the password.
  - Step 6** In the **Confirm Password** field, enter the password again.
  - Step 7** Click **Add** to add the user to the database.
  - Step 8** Repeat Steps 2 through 7 to continue adding users.
- 

### Deleting Local Users

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

**Procedure**

- 
- Step 1** Choose **Administration > Management Users > Local**.  
The **Local Users** page is displayed.
- Step 2** Select one or more users from the **Local Users** table and click the **Delete User** button.  
**Step 3** Click **Yes** on the warning window to delete the local user. Click **No** to cancel deletion.
- 

## Editing a User

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

**Procedure**

- 
- Step 1** Choose **Administration > Management Users > Local**.
- Step 2** Use the checkbox to select a user and click the **Edit User** icon.
- Step 3** In the **Edit User** window, the **Username** and **Role** are mentioned by default. Specify the **Password** and **Confirm Password**.
- Step 4** Click **Apply** to save the changes.
- 

## User Access

To control the local users to access the specific groups from the Cisco DCNM Web UI, perform the following steps:

**Procedure**

- 
- Step 1** Choose **Administration > Management Users > Local**.  
The **Local Users** window is displayed.
- Step 2** Select one user from the **Local Users** table. Click **User Access**.  
The **User Access** selection window is displayed.
- Step 3** Select the groups allowed to access for the user and click **Apply**.
- 

## Managing Clients

You can use Cisco DCNM to disconnect DCNM Client Servers.

**Procedure**

- 
- Step 1** Choose **Administration > Management Users > Clients**.

A list of DCNM Servers are displayed.

- Step 2** Use the check box to select a DCNM server and click **Disconnect Client** to disconnect the DCNM server.

**Note** You cannot disconnect a current client session.

---

## Performance Setup

The Performance Setup menu includes the following submenus:

### Performance Setup LAN Collections

If you are managing your switches with the Performance Manager, you must set up an initial set of flows and collections on the switch. You can use Cisco DCNM to add and remove performance collections. License the switch and kept it in the **Managed Continuously** state before creating a collection for the switch.

To add a collection, follow these steps:

**Procedure**

- 
- Step 1** Choose **Administration > Performance Setup > LAN Collections**.

- Step 2** For all the licensed LAN switches, use the check boxes to enable performance data collection for **Trunks, Access, Errors & Discards, and Temperature Sensor**.

- Step 3** Use the check boxes to select the types of LAN switches for which you want to collect performance data.

- Step 4** Click **Apply** to save the configuration.

- Step 5** In the confirmation dialog box, click **Yes** to restart the performance collector.
- 

## Event Setup

The Event Setup menu includes the following submenus:

### Viewing Events Registration

To enable **Send Syslog**, **Send Traps** and **Delayed Traps** you must configure the following in the DCNM SAN client:

- Enabling **Send Syslog**: Choose **Physical Attributes > Events > Syslog > Servers**. Click **Create Row**, provide the required details, and click **Create**.

- Enabling **Send Traps**: Choose **Physical Attributes > Events > SNMP Traps > Destination**. Click **Create Row**, provide the required details, and click **Create**.
- Enabling **Delayed Traps**: Choose **Physical Attributes > Events > SNMP Traps > Delayed Traps**. In the **Feature Enable** column, use the check boxes to enable delayed traps for the switch and specify the delay in minutes.

### Procedure

---

- Step 1** Choose **Administration > Event Setup > Registration**.  
The SNMP and Syslog receivers along with the statistics information are displayed.
- Step 2** Check the **Enable Syslog Receiver** check box and click **Apply**, to enable the syslog receiver if it is disabled in the server property.  
To configure event registration or syslog properties, choose **Administration > DCNM Server > Server Properties** and follow the on-screen instructions.
- Step 3** Select **Copy Syslog Messages to DB** and click **Apply** to copy the syslog messages to the database.  
If this option is not selected, the events will not be displayed in the events page of the Web client.  
The columns in the second table display the following:
  - Switches sending traps
  - Switches sending syslog
  - Switches sending syslog accounting
  - Switches sending delayed traps

---

## Notification Forwarding

You can use Cisco DCNM Web UI to add and remove notification forwarding for system messages.

This section contains the following:

### Adding Notification Forwarding

Cisco DCNM Web UI forwards fabric events through email or SNMPv1 traps.

To add and remove notification forwarding for system messages from the Cisco DCNM Web UI, perform the following steps:



---

**Note**

Test forwarding works only for the licensed fabrics.

---

## Procedure

---

**Step 1** Choose **Administration > Event Setup > Forwarding**.

The events forwarding scope, the recipient email address, severity of the event and type of the event is displayed. The description Regex field is applicable only when the forwarding source is selected as Syslog while adding the events forwarder.

**Step 2** Check the **Enable** checkbox to enable events forwarding.

**Step 3** Specify the **SMTP Server** details and the **From** email address.

**Step 4** Click **Apply** to save the configuration, or in the **Apply and Test** icon, use the drop-down to select the fabric.

Click **Apply and Test** to save and test the configuration.

**Step 5** In the **Event Count Filter**, add a filter for the event count to the event forwarder.

The forwarding stops forwarding an event if the event count exceeds the limit as specified in the event count filter. In this field, you can specify a count limit. Before an event can be forwarded, the Cisco DCNM checks if its occurrence exceeds the count limit. If it does, the event will not be forwarded.

**Step 6** Select the **Snooze** checkbox and specify the **Start** date and time and the **End** date and time. Click **Apply** to save the configuration.

**Step 7** Under the **Event Forwarder Rules** table, click the + icon to add an event forwarder rule.

You see the **Add Event Forwarder Rule** dialog box.

**Step 8** In the **Forwarding Method**, choose either **E-mail** or **Trap**. If you choose **Trap**, a **Port** field is added to the dialog box.

**Step 9** If you choose the **E-mail** forwarding method, enter the IP address in the **Email Address** field. If you choose the **Trap** method, enter the trap receiver IP address in the **Address** field and specify the port number.

You can either enter an IPv4 or IPv6 addresses or DNS server name in the **Address** field.

**Step 10** For **Forwarding Scope**, choose the **Fabric/LAN** or **Port Groups** for notification.

**Step 11** In the **Source** field, select **DCNM** or **Syslog**.

If you select **DCNM**, then:

- a) From the **Type** drop-down list, choose an event type.
- b) Check the **Storage Ports Only** check box to select only the storage ports.
- c) From the **Minimum Severity** drop-down list, select the severity level of the messages to receive.
- d) Click **Add** to add the notification.

If you select **Syslog**, then:

- a) In the **Facility** list, select the syslog facility.
- b) Specify the syslog **Type**.
- c) In the **Description Regex** field, specify a description that matches with the event description.
- d) From the **Minimum Severity** drop-down list, select the severity level of the messages to receive.
- e) Click **Add** to add the notification.

**Note** The **Minimum Severity** option is available only if the **Event Type** is set to All.

The traps that are transmitted by Cisco DCNM correspond to the severity type. A text description is also provided with the severity type.

```
trap type(s) = 40990 (emergency)
40991 (alert)
40992 (critical)
40993 (error)
40994 (warning)
40995 (notice)
40996 (info)
40997 (debug)
textDescriptionOid = 1, 3, 6, 1, 4, 1, 9, 9, 40999, 1, 1, 3, 0
```

## Removing Notification Forwarding

You can remove notification forwarding.

### Procedure

- 
- Step 1** Choose **Administration > Event Setup > Forwarding**.
- Step 2** Select the check box in front of the notification that you want to remove and click **Delete**.
- 

## Event Suppression

Cisco DCNM allows you to suppress the specified events that are based on the user-specified suppressor rules. Such events will not be displayed on the Cisco DCNM Web UI. The events will neither be persisted to DCNM database, nor forwarded via email or SNMP trap.

You can view, add, modify, and delete suppressor rules from the table. You can create a suppressor rule from the existing event table. Select a given event as the template, and invoke the rule dialog window. Event details are automatically ported from the selected event in the event table to the input fields of the rule creation dialog window.

This section includes the following:

### Add Event Suppression Rules

To add rules to the Event Suppression from the Cisco DCNM Web UI, perform the following steps:

### Procedure

- 
- Step 1** Choose **Administration > Event Setup > Suppression**.  
The **Suppression** window is displayed.
- Step 2** Click the **Add** icon above the **Event Suppressors** table.  
The **Add Event Suppressor Rule** window is displayed.
- Step 3** In the **Add Event Suppressor Rule** window, specify the **Name** for the rule.
- Step 4** Select the required **Scope** for the rule that is based on the event source.

**Delete Event Suppression Rule**

In the Scope drop-down list, the LAN groups and the port groups are listed separately. You can choose **LAN**, **Port Groups** or **Any**. For **LAN**, select the scope of the event at the Fabric or Group or Switch level. You can only select groups for **Port Group** scope. If you select **Any** as the scope, the suppressor rule is applied globally.

**Step 5** Enter the **Facility** name or choose from the **LAN Switch Event Facility List**.

If you do not specify a facility, wildcard is applied.

**Step 6** From the drop-down list, select the **Event Type**.

If you do not specify the event type, wildcard is applied.

**Step 7** In the **Description Matching** field, specify a matching string or regular expression.

The rule matching engine uses regular expression that is supported by Java Pattern class to find a match against an event description text.

**Step 8** Check the **Active Between** box and select a valid time range during which the event is suppressed.

By default, the time range is not enabled, i.e., the rule is always active.

**Note** In general, you must not suppress accounting events. Suppressor rule for Accounting events can be created only for certain rare situations where Accounting events are generated by actions of DCNM or switch software. For example, lots of 'sync-snmp-password' AAA syslog events are automatically generated during the password synchronization between DCNM and managed switches. To suppress Accounting events, navigate to the **Suppressor table** and invoke the **Add Event Suppressor Rule** dialog window.

**Note** Choose **Monitor > Switch > Events** to create a suppressor rule for a known event. There is no such shortcut to create suppressor rules for Accounting events.

**Delete Event Suppression Rule**

To delete event suppressor rules from the Cisco DCNM Web UI, perform the following steps:

**Procedure**

**Step 1** Choose **Administration > Event Setup > Suppression**.

**Step 2** Select the rule from the list and click **Delete** icon.

**Step 3** Click **Yes** to confirm.

**Modify Event Suppression Rule**

To modify the event suppressor rules, do the following tasks:

**Procedure**

**Step 1** Choose **Administration > Event Setup > Suppression**.

**Step 2** Select the rule from the list and click **Edit**.

You can edit **Facility**, **Type**, **Description Matching** string, and **Valid time range**.

**Step 3** Click **Apply** to save the changes.

---

## Credentials Management

The Credential Management menu includes the following submenus:

### LAN Credentials

While changing the device configuration, Cisco DCNM uses the device credentials provided by you. However, if the LAN Switch credentials are not provided, Cisco DCNM prompts you to open the **Administration > Credentials Management > LAN Credentials** page to configure LAN credentials.

Cisco DCNM uses two sets of credentials to connect to the LAN devices:

- **Discovery Credentials**—Cisco DCNM uses these credentials during discovery and periodic polling of the devices.
- **Configuration Change Credentials**—Cisco DCNM uses these credentials when user tries to use the features that change the device configuration.

LAN Credentials Management allows you to specify configuration change credentials. Before changing any LAN switch configuration, you must furnish *Configuration Change* SSH credentials for the switch. If you do not provide the credentials, the configuration change action will be rejected.

These features get the device write credentials from LAN Credentials feature.

- Upgrade (ISSU)
- Maintenance Mode (GIR)
- Patch (SMU)
- Template Deployment
- POAP-Write erase reload, Rollback
- Interface Creation/Deletion/Configuration
- VLAN Creation/Deletion/Configuration
- VPC Wizard

You must specify the configuration change credentials irrespective of whether the devices were discovered initially or not. This is a one-time operation. Once the credentials are set, that will be used for any configuration change operation.

#### Default Credentials

Default credentials is used to connect all the devices that the user has access to. You can override the default credentials by specifying credentials for each of the devices in the Switch Table below.

Cisco DCNM tries to use individual switch credentials in the Switch Table, to begin with. If the credentials (username/password) columns are empty in the Switch Table, the default credentials will be used.

### Switch Table

Switch table lists all the LAN switches that user has access. You can specify the switch credentials individually, that will override the default credentials. In most cases, you need to provide only the default credentials.

You can perform the following operations on this screen.

- [Edit Credentials, on page 254](#)
- [Validate Credentials, on page 254](#)
- [Clear Switch Credentials, on page 254](#)

The LAN Credentials for the DCNM User table has the following fields.

| Field      | Description                                      |
|------------|--------------------------------------------------|
| Switch     | Displays the LAN switch name.                    |
| IP Address | Specifies the IP Address of the switch.          |
| User Name  | Specifies the username of the switch DCNM user.  |
| Password   | Displays the encrypted form of the SSH password. |
| Group      | Displays the group to which the switch belongs.  |

### Edit Credentials

Perform the following task to edit the credentials.

1. From the Cisco DCNM home page, choose **Administration > Credentials Management > LAN Credentials**, check the **Switch** check box for which you need to edit the credentials.
2. Click Edit icon.
3. Specify **User Name** and **Password** for the switch.

### Validate Credentials

Perform the following task to validate the credentials.

1. From the **Administration > Credentials Management > LAN Credentials**, check the **Switch** check box for which you need to validate the credentials.
2. Click **Validate**.

A confirmation message appears, stating if the operation was successful or a failure.

### Clear Switch Credentials

Perform the following task to clear the switch credentials.

1. From the **Administration > Credentials Management > LAN Credentials**, check the **Switch** check box for which you need to clear the credentials.

2. Click **Clear**.
3. Click **Yes** to clear the switch credentials from the DCNM server.





## CHAPTER 7

# Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - Multi-Site

This section explains how to connect two Virtual eXtensible Local Area Network (VXLAN) Border Gateway Protocol (BGP) Ethernet VPN (EVPN) fabrics through DCNM using the EVPN Multi-Site feature. The EVPN Multi-Site configurations are applied on the Border Gateways (BGWs) of the two fabrics.



**Note**

In Cisco® Data Center Network Manager (DCNM) 11.0(1), you can connect BGWs of two standalone fabrics or two member fabrics of an MSD.

Multi-Site Domain (MSD), introduced in DCNM 11.0(1) release, is a multifabric container that is created to manage multiple member fabrics. It is a single point of control for definition of overlay networks and VRFs that are shared across member fabrics. See *Multi-Site Domain for VXLAN BGP EVPN Fabrics* section in the *Control* chapter for more information on MSD.



**Note**

For a detailed explanation on the EVPN Multi-Site feature, see the [VXLAN BGP EVPN Multi-Site Design and Deployment](#) document.

- [Prerequisites , on page 257](#)
- [Limitations, on page 258](#)
- [Sample Scenario, on page 259](#)
- [EVPN Multi-Site Configuration , on page 261](#)
- [Deploying Networks and VRF Instances, on page 277](#)
- [Additional References, on page 282](#)
- [Appendix , on page 282](#)

## Prerequisites

- The EVPN Multi-Site feature requires Cisco Nexus 9000 Series NX-OS Release 7.0(3)I7(1) or later.
- Familiarity with VXLAN BGP EVPN data center fabric architecture and configuration through DCNM.
- Familiarity with MSD fabrics, if you are connecting member fabrics of an MSD.

**Limitations**

- Fully configured VXLAN BGP EVPN fabrics that are ready to be connected using the EVPN Multi-Site feature, external fabric(s) configuration through DCNM, and relevant external fabric devices' configuration (for example, route servers).
  - VXLAN BGP EVPN fabrics (and their interconnection) can be configured manually or using DCNM. This document explains the process to connect the fabrics through DCNM. So, you should know how to configure and deploy a VXLAN BGP EVPN fabric, and how to create an external fabric through DCNM. For more details, see the *VXLAN BGP EVPN Fabrics Provisioning* section in the *Control* chapter.
  - When you enable the EVPN Multi-Site feature on a BGW, ensure that there are no prior overlay deployments on it. Remove existing overlay profiles and then start provisioning Multi-Site extensions through DCNM.
  - Ensure that the role of the designated BGWs is *Border Gateway*. To verify, right-click the BGW and click **Set role**. You can see that *(current)* is added to the current role of the switch.
- If the current role is not *Border Gateway*, you should remove the device from the fabric and discover it again through DCNM using the POAP bootstrap option and re-provision the configurations for the device.
- To ensure consistency across fabrics, ensure the following:



**Note** These checks are done for member fabrics of an MSD when the fabrics are moved under the MSD fabric.

- The underlay IP addresses across the fabrics, the loopback 0 address and the loopback 1 address subnets should be unique.
- Each fabric should have a unique site ID and BGP AS number associated and configured.
- All fabrics should have the same Anycast Gateway MAC address.
- While the MSD provisions a global range of network and VRF values, some parameters are fabric-specific and some are switch-specific. You should specify fabric instance values for each fabric (for example, *multicast group subnet address*) and switch instance values for each switch (for example, *VLAN ID*).

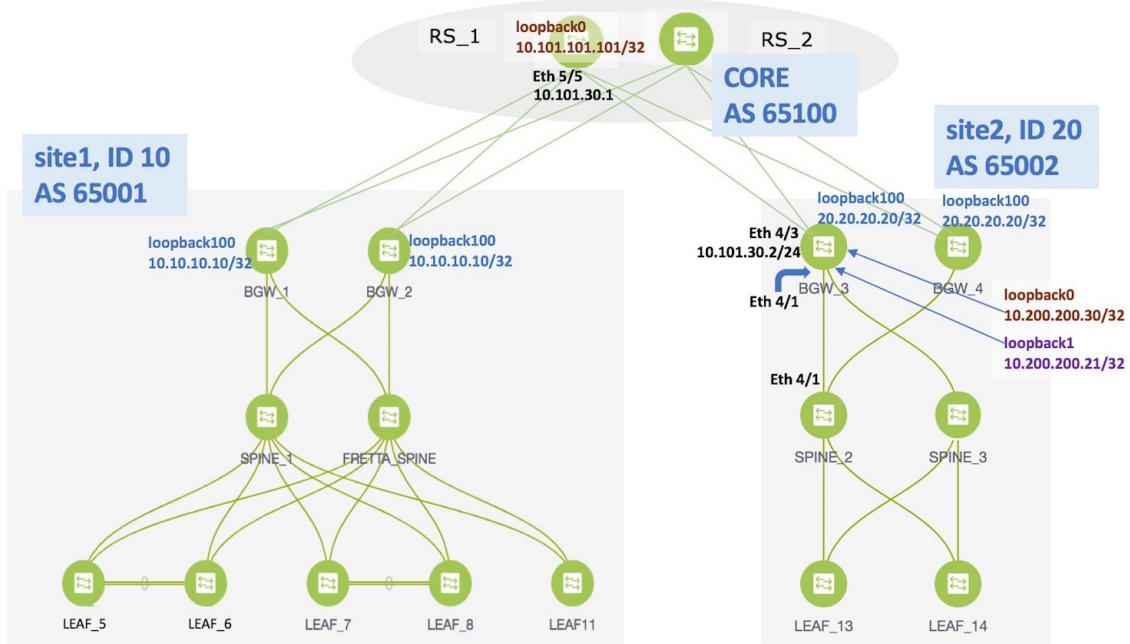
After completing the EVPN Multi-Site specific prerequisites, start EVPN Multi-Site configuration on *BGW\_3* with extensions to the route server *RS\_1*.

## Limitations

- BGWs cannot form a virtual port channel (vPC) switch pair.
- The VXLAN OAM feature in Cisco DCNM is only supported on a single fabric or site.

# Sample Scenario

The EVPN Multi-Site feature is explained through an example scenario. Consider two VXLAN BGP EVPN fabrics, *site1* and *site2* connected through devices in an external fabric, *CORE*. This document shows you how to enable end-to-end Layer 3 and Layer 2 traffic between hosts in *site1* and *site2*, through *CORE* the fabric.



Network configurations for the two VXLAN BGP EVPN fabrics are provisioned through DCNM software, 11.0(1) release. VXLAN BGP EVPN configurations are configured on the switches in the two fabrics. However, server traffic between the sites is only possible through a Data Center Interconnect (DCI) function. If a server in *site1* has to send traffic to a server in *site2* or vice versa, the DCI function (such as the Multi-Site feature, which is used for this example) should be configured on the BGWs of both the fabrics.

Route servers *RS\_1* and *RS\_2* are route servers that are directly attached to the two VXLAN BGP EVPN fabrics. From the VXLAN fabrics' point of view, the route servers belong to an external fabric, *CORE*, with a different AS number. For representation purposes, the *CORE* fabric is created as an external fabric through DCNM, and *RS\_1* and *RS\_2* are associated with it.



**Note** Though creating an external fabric is a prerequisite for this use case, steps are noted here for quick reference. To create an external fabric in DCNM, click **Control > Fabric Builder**. On the Fabric Builder page, click **Create Fabric**. On the Add Fabric page, enter the name of the fabric (*CORE*), select *External\_Fabric* as the fabric template, enter the AS Number and click **Save**.

The *CORE* fabric is created as an external fabric.

The steps that are involved to enable EVPN Multi-Site feature and traffic flow across the sites or fabrics are:

1. Top-Down deployment of the underlay for the IP core at the BGWs. This is a one-time configuration.

2. Top-Down deployment of the BGP overlay for the IP core. This is a one-time configuration for each BGW.
3. Deployment of networks and virtual routing and forwarding (VRF) instances on the leaf switches. This is a per network/VRF configuration.
4. Deployment of networks/VRFs at the BGWs. This is a per network/VRF configuration.

**EVPN Multi-Site Feature-**This requires setting up the BGW base configuration for enabling the EVPN Multi-Site feature on the BGWs and the underlay peering to the external devices. This is followed by establishing overlay peering from the BGW to appropriate external devices, either BGWs in other fabrics or route servers. Both the underlay and overlay peering are established over eBGP. While eBGP is mandatory for the overlay peering, you can use eBGP or an IGP for the underlay.



- Note** DCNM 11.0(1) Top-Down provisioning only supports eBGP underlay.

BGWs are special devices that allow clear control and data plane segregation from one site to another, allowing for policy enforcement points for any inter-fabric traffic. They allow the same data plane (VXLAN) and control plane (BGP EVPN) to be employed both for inter-fabric and intra-fabric traffic.

The end-to-end configurations can be split into these 2 steps:

1. **EVPN Multi-Site configurations on the BGWs (*BGW\_1*, *BGW\_2*, *BGW\_3* and *BGW\_4*).**
  1. EVPN Multi-Site feature on the BGWs on *site1*—Overlay and underlay connections between the BGWs *BGW\_1* and *BGW\_2*, and directly connected route servers *RS\_1* and *RS\_2* in the *CORE* fabric.
  2. EVPN Multi-Site feature on the BGWs on *site2*—This includes overlay and underlay connections between the BGWs *BGW\_3* and *BGW\_4*, and directly connected route servers *RS\_1* and *RS\_2* in the *CORE* fabric.
  3. Configurations on *RS\_1* and *RS\_2*—Configurations in the *CORE* fabric are not in the scope of DCNM provisioning and this document. For completeness, it is mentioned here, and sample configurations provided in the Appendix section.



- Note** The network interconnecting the BGWs can be more complex than just 2 switches. The proper configuration (routing protocol peering, MTU settings, etc) required in that network is a one-time initial infrastructure configuration that must be performed outside of DCNM.

For this example, *BGW\_3* EVPN Multi-Site configurations will be explained.

## 2. Deploying Networks and VRF Instances on the leaf switches and the BGWs

For this example, 2 networks are configured on the BGWs in *site2* (with the assumption that network deployment on leaf switches is already completed).

After successful deployment on both the sites, Layer 2 and Layer 3 traffic will flow between the two sites.

**Note**

In the DCNM GUI, the lines connecting devices that are managed by DCNM (for example, *LEAF\_5* to *SPINE\_1* and *SPINE\_1* to *BGW\_2*) symbolize a physical cable connection, and not that the connection is functional and network traffic flows between them.

To start with, let us consider EVPN Multi-Site provisioning on *BGW\_3* through DCNM Top-Down LAN Fabric Provisioning.

## EVPN Multi-Site Configuration

### EVPN Multi-Site Extensions from *BGW\_3* to *RS\_1*

1. Choose **Control > Networks & VRFs**. The LAN Fabric Provisioning page appears.
2. Click **Continue**. The Select a Fabric page appears.
3. Select **site2** from the drop-down box since you are configuring the BGW *BGW\_3* on *site2*.
4. Click **Fabric Extension Setup** since the purpose of this task is to allow *site2* to communicate to external fabrics through *RS\_1* and *RS\_2*. The Fabric Extension screen comes up.

The screenshot shows the 'Fabric Extension' screen with the 'Inter-Fabric Connections' section. The table has columns for Type, Source Fabric, Source Device, Source Interface, Destination Fa..., Destination De..., Destination Int..., Configuration, and Status. A note at the bottom says 'No data available'. The top right shows 'Selected 0 / Total 0' and filter options.

| Type              | Source Fabric | Source Device | Source Interface | Destination Fa... | Destination De... | Destination Int... | Configuration | Status |
|-------------------|---------------|---------------|------------------|-------------------|-------------------|--------------------|---------------|--------|
| No data available |               |               |                  |                   |                   |                    |               |        |

The Inter-Fabric Connections section lists previously created external connections from the BGWs on *site2*. Each line represents a physical or logical connection between a BGW in *site2* and an external device in another fabric. For each connection, the source fabric, source device, source interface, destination fabric, destination device, and destination interface are listed along with the type of external connectivity. This section is empty as this is the first time you are adding an external connection.

*To extend the fabric through EVPN Multi-Site, you should first create an underlay extension and then an overlay extension.*

### Underlay Extension from *BGW\_3* to *RS\_1*

1. Click the + icon to add a new external connection. The **Add Inter-Fabric Connection** screen appears.

**Underlay Extension from BGW\_3 to RS\_1**

## Add Inter-Fabric Connections

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

● ● ●

|                         |                  |
|-------------------------|------------------|
| * Extension Type        | VRF_LITE         |
| * Base Template         | ext_base_setup   |
| * Extension Template    | ext_fabric_setup |
| * Source Fabric         | SITE_2           |
| * Destination Fabric    |                  |
| * Source Device         |                  |
| * Source Interface      |                  |
| * Destination Device    |                  |
| * Destination Interface |                  |

(i) VRF\_LITE: Set switch role - Border; MULTISITE: Set switch role - "Border Gateway"

Previous    Next    Save & Deploy    Cancel

By default, **VRF\_LITE** is populated in the **Extension Type** field. Change the selection to **MULTISITE\_UNDERLAY**.

## Add Inter-Fabric Connections

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

• • •

|                         |                              |
|-------------------------|------------------------------|
| * Extension Type        | MULTISITE_UNDERLAY           |
| * Base Template         | ext_base_setup               |
| * Extension Template    | ext_multisite_underlay_setup |
| * Source Fabric         | SITE_2                       |
| * Destination Fabric    |                              |
| * Source Device         |                              |
| * Source Interface      |                              |
| * Destination Device    |                              |
| * Destination Interface |                              |

(i) VRF\_LITE: Set switch role - Border; MULTISITE: Set switch role - "Border Gateway"

Previous    Next    Save & Deploy    Cancel

**Base Template**-By default, the *ext\_base\_setup* base template is populated. This template is a one-time configuration that is pushed to the BGW.

**Extension Template**-*ext\_multisite\_underlay\_setup* is a setup template that contains the configuration that is generated and pushed to the BGW to set up the corresponding interfabric connection.

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

**Note**

You to 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 *site2* since the EVPN Multi-Site underlay connection is between *BGW\_3* in *site2* and *RS\_1* in the *CORE* fabric.

**Destination Fabric**-Choose *CORE*.

**Source Device** and **Source Interface**-Choose *BGW\_3* as the source device and an Ethernet interface that needs to be connected to *RS\_3*.

**Destination Device** and **Destination Interface**—Choose *RS\_1* as the destination device and the Ethernet interface that connects to the BGW *BGW\_3*.

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.

**Underlay Extension from BGW\_3 to RS\_1**

After filling up the Fabric Interconnect section, the screen looks like this:

### Add Inter-Fabric Connections

The screenshot shows the 'Add Inter-Fabric Connections' wizard in step 1: Fabric Interconnect. The configuration fields are as follows:

- \* Extension Type:** MULTISITE\_UNDERLAY
- \* Base Template:** ext\_base\_setup
- \* Extension Template:** ext\_multisite\_underlay\_setup
- \* Source Fabric:** SITE\_2
- \* Destination Fabric:** CORE
- \* Source Device:** BGW\_3
- \* Source Interface:** Ethernet4/3
- \* Destination Device:** RS\_1
- \* Destination Interface:** Ethernet5/5

At the bottom, there are buttons for **Previous**, **Next**, **Save & Deploy**, and **Cancel**.

- Click **Next** to go to the Define Variables section. The fields in this section are:

**IF\_NAME**—In this field, the interface name is autopopulated from the previous step.

**IP\_MASK**—Fill up this field with the IP address of the *BGW\_3* interface that connects to *RS\_1*.

**NEIGHBOR\_IP**—Fill up this field with the IP address of the *RS\_1* interface that connects to *BGW\_3*.

**NEIGHBOR ASN**—In this field, the AS number of *RS\_1* will be autopopulated.

A filled up screen is displayed:

## Add Inter-Fabric Connections

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

• • •

---

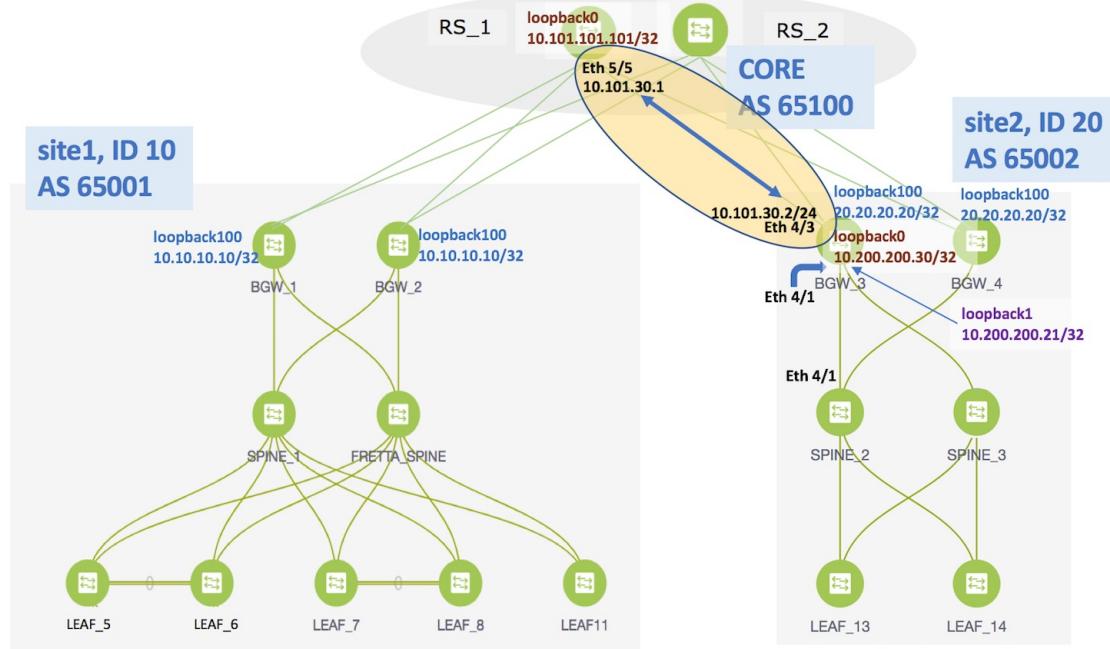
▼ Network Profile

|           |                                                                                                                                                                                                                                                       |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| General   |                                                                                                                                                                                                                                                       |
| MULTISITE | <p>* IF_NAME <input type="text" value="Ethernet4/3"/> ?</p> <p>* IP_MASK <input type="text" value="10.101.30.2/24"/> ?</p> <p>* NEIGHBOR_IP <input type="text" value="10.101.30.1"/> ?</p> <p>* NEIGHBOR ASN <input type="text" value="65100"/> ?</p> |

Previous Next Save & Deploy Cancel

The corresponding connection in the topology is displayed:

### Underlay Extension from BGW\_3 to RS\_1



### 3. Click the MULTISITE tab.

While the **General** tab contains external connection details, this tab contains intra-fabric information such as fabric IGP, fabric facing Ethernet interface, and so on.

The **MULTISITE** tab only appears the first time that you create an EVPN Multi-Site underlay on a device, since the details remain the same for subsequent connections. The next time you create an EVPN Multi-Site underlay connection *on the same device*, only the General tab will be available.

## Add Inter-Fabric Connections



1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

● ● ●

---

▼ Network Profile

| General                        | MULTISITE |
|--------------------------------|-----------|
| <b>*</b> Fabric Site ID        | 20        |
| <b>*</b> NVE Identifier        | 1         |
| <b>*</b> Fabric Interfaces     |           |
| <b>*</b> Multisite Loopback ID | 100       |
| <b>*</b> MultiSite Loopback IP |           |
| <b>*</b> Routing Protocol      | is-is     |
| <b>*</b> IS-IS/OSPF Router ID  | UNDERLAY  |
| <b>*</b> OSPF Area #           | 0         |

Previous Next Save & Deploy Cancel

**Fabric Site ID**—This is the identification for the VXLAN BGP EVPN fabric *site2* to which *BGW\_3* belongs. The site ID is auto populated from the fabric settings, but it is editable. It should be same on all BGWs in one fabric and distinct from all other fabrics.

**NVE Identifier**—This is the VXLAN overlay ID.

**Fabric Interfaces**—Fill up this field with the interfaces on *BGW\_3* that connects to other intra-fabric device ports. Since Ethernet 4/1 connects to *SPINE\_2* and Ethernet 4/2 connects to *SPINE\_3* in the topology, the interfaces should be entered over here.

**Multisite Loopback ID and Multisite Loopback IP**—These are the loopback ID and IP address of this EVPN Multi-Site instance. The loopback IP address should be same for all BGWs in one fabric and distinct from all other fabrics.

**Routing Protocol and Router ID**—This is the IGP and the IGP instance ID within the fabric. Note that, if the IGP used in your setup is OSPF, the field has to be updated to *OSPF*.

**OSPF AREA**—OSPF area ID within the fabric.

A fully filled screen looks like this.

**Underlay Extension from BGW\_3 to RS\_1****Add Inter-Fabric Connections**

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

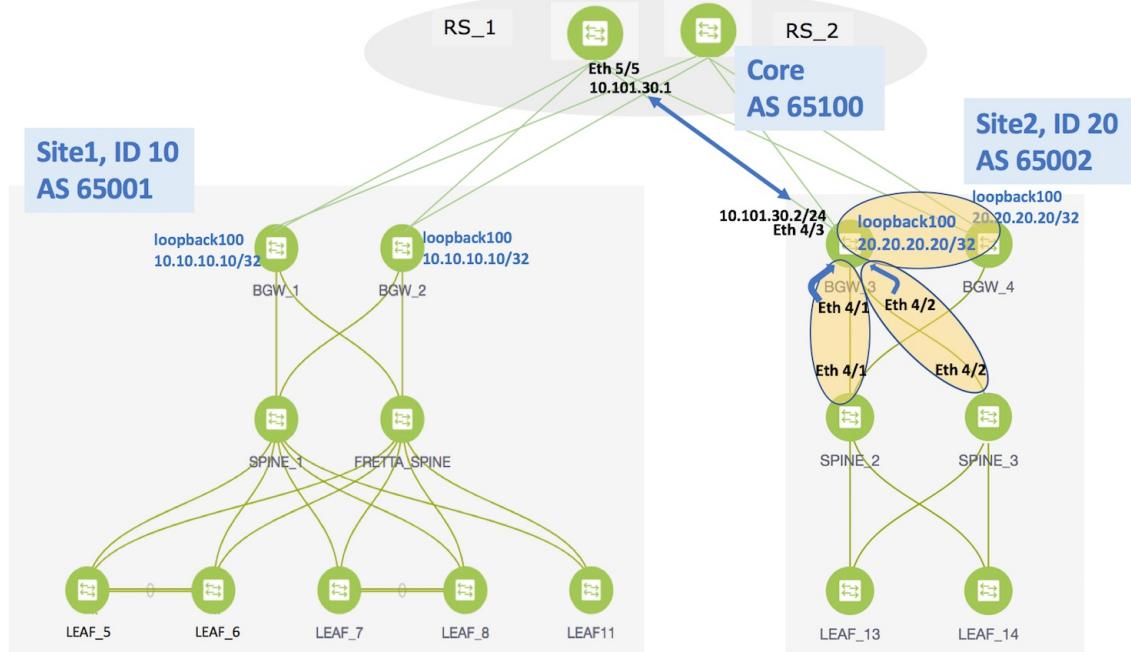
• • •

▼ Network Profile

| General | MULTISITE             |                |                              |
|---------|-----------------------|----------------|------------------------------|
| *       | Fabric Site ID        | 20             | ?                            |
| *       | NVE Identifier        | 1              | ?                            |
| *       | Fabric Interfaces     | Eth4/1, Eth4/2 | ? E.g. e1/1-4, e2/2          |
| *       | Multisite Loopback ID | 100            | ? [0-1023]                   |
| *       | MultiSite Loopback IP | 20.20.20.20    | ? IPv4 address               |
| *       | Routing Protocol      | ospf           | ? Select IGP (ospf or is-is) |
| *       | IS-IS/OSPF Router ID  | UNDERLAY       | ? String                     |
| *       | OSPF Area #           | 0              | ? String                     |

**Previous** **Next** **Save & Deploy** **Cancel**

The corresponding topology depiction is given below:



- Now that all the information is filled in, click **Next** to go to the **Preview and Deploy** section.

## Add Inter-Fabric Connections

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

Switch: BGW\_3

Generated Configuration:

```
route-map RMAP-REDIST-DIRECT permit 10
  match tag 54321

  evpn multisite border-gateway 20

  interface loopback100
    description Used for EVPN Multi-Site
    ip address 20.20.20.20/32 tag 54321

    ip router ospf UNDERLAY area 0
    no shutdown

  interface nve 1
    multisite border-gateway interface loopback100

  interface e4/1
    evpn multisite fabric-tracking
```

Previous Next Save & Deploy Cancel

Here, you can preview the configuration that will be deployed to *BGW\_3*. Note that no configuration will be pushed to the external device itself.

- Click **Save and Deploy** to complete the task. This results in the configuration getting pushed to *BGW\_3*. The external connection will appear in the Fabric Extension screen.

| Fabric Extension                         |        |               |               |                  |                 |                   |                             |               |        |
|------------------------------------------|--------|---------------|---------------|------------------|-----------------|-------------------|-----------------------------|---------------|--------|
| Inter-Fabric Connections                 |        |               |               |                  |                 |                   |                             |               |        |
|                                          | Type   | Source Fabric | Source Device | Source Interf... | Destination ... | Destination De... | Destination Interf...       | Configuration | Status |
| <input type="radio"/> MULTISITE_UNDERLAY | SITE_2 | BGW_3         | Ethernet4/3   | CORE             | RS_1            | Ethernet5/5       | <a href="#">View Config</a> | DEPLOYED      |        |

You can check the status of the deployment (Deployment Pending, Deployed, Failed) in the **Status** column.

In case of FAILED or UNDEPLOYMENT FAILED status, use the hyperlink in the **Status** column to check the error messages for failure.

To view the configurations, click on *View Config* in the **Configuration** field.

After the underlay configuration, you need to configure the overlay configuration from *BGW\_3* to *RS\_1* (the external device connected to *BGW\_3*), as shown in the next section.

## Overlay Extension from BGW\_3 to RS\_1



**Note** You can have multiple underlay connections to an external device but only one overlay connection from *BGW\_3* to each external device.

1. In the **Fabric Extension** page, click on the + icon to add an external overlay connection. The **Add Inter-Fabric Connections** screen appears.

By default, **VRF\_LITE** is populated in the **Extension Type** field. Change the selection to **MULTISITE\_OVERLAY**. The screen changes accordingly.

### Add Inter-Fabric Connections

The screenshot shows the 'Add Inter-Fabric Connections' wizard with three steps:

- Step 1: Fabric Interconnect** (highlighted with a blue border)
- Step 2: Define Variables**
- Step 3: Preview & Deploy**

Below the steps, there are several configuration fields:

|                         |                             |
|-------------------------|-----------------------------|
| * Extension Type        | MULTISITE_OVERLAY           |
| * Base Template         | ext_base_setup              |
| * Extension Template    | ext_multisite_overlay_setup |
| * Source Fabric         | SITE_2                      |
| * Destination Fabric    |                             |
| * Source Device         |                             |
| * Source Interface      |                             |
| * Destination Device    |                             |
| * Destination Interface |                             |

A small note at the bottom right of the form area states: (i) VRF\_LITE: Set switch role - Border; MULTISITE: Set switch role - "Border Gateway".

At the bottom of the form are buttons for **Previous**, **Next**, **Save & Deploy**, and **Cancel**.

**Base Template**—*ext\_base\_setup* is auto-populated in this field. The *ext\_base\_setup* base template is a one-time configuration pushed to the BGW.

**Extension Template**—*ext\_multisite\_overlay\_setup* is a setup template that contains the configuration that will be generated and pushed to the BGW to setup the corresponding inter-fabric connection. These templates are auto-populated with corresponding pre-packaged default templates based on your selection.

**Source Fabric**—This field is pre-populated with *site2* since you are deploying the configurations in *site2*.

**Destination Fabric**—For the destination fabric, select the fabric that contains *RS\_1*, *CORE*.

**Source Device**—Choose *BGW\_3* since the overlay connection is from *BGW\_3* to *RS\_1*.

**Source Interface**—Choose the source interface. Typically, a loopback interface is created for the overlay. The loopback IP address (*loopback0* in this example) is used for BGP peering with the destination interface.

**Destination Device**—Choose *RS\_1* since the overlay connection is from *BGW\_3* to *RS\_1*.

**Destination Interface**—Choose the destination interface. Choose the interface which is the BGP peer address. Note that the destination interface is not used in generating the configuration.

After filling up the Fabric Interconnect section, the screen looks like this.

### Add Inter-Fabric Connections

| Field                   | Value                       |
|-------------------------|-----------------------------|
| * Extension Type        | MULTISITE_OVERLAY           |
| * Base Template         |                             |
| * Extension Template    | ext_multisite_overlay_setup |
| * Source Fabric         | SITE_2                      |
| * Destination Fabric    | CORE                        |
| * Source Device         | BGW_3                       |
| * Source Interface      | Loopback0                   |
| * Destination Device    | RS_1                        |
| * Destination Interface | Loopback0                   |

Below the form:

- Previous
- Next
- Save & Deploy
- Cancel

- Click **Next** to go to the **Define Variables** section. The fields in this screen:

**NEIGHBOR ASN**—This field is populated with the *RS\_1*'s AS Number.

**Overlay Neighbor IP**—Enter the IP address on *RS\_1* that the overlay peers with. This is typically a loopback address.

**IF\_NAME**—In this field, the source interface is auto-populated from the previous step.

A fully filled screen looks like this:

**Overlay Extension from BGW\_3 to RS\_1**

## Add Inter-Fabric Connections

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

• • •

---

▼ Network Profile

|         |                                                                                                                       |
|---------|-----------------------------------------------------------------------------------------------------------------------|
| General | <p>* NEIGHBOR ASN 65100 ?</p> <p>* Overlay Neighbor IP 10.101.101.101 ? IPv4 address</p> <p>* IF_NAME Loopback0 ?</p> |
|---------|-----------------------------------------------------------------------------------------------------------------------|

Previous Next Save & Deploy Cancel

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

## Add Inter-Fabric Connections

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

Switch: BGW\_3

Generated Configuration:

```
router bgp 65002
  neighbor 10.101.101.101 remote-as 65100
    update-source Loopback0
    ebgp-multipath 5
    peer-type fabric-external
    address-family l2vpn evpn
      send-community
      send-community extended
      rewrite-evpn-rt-asn
```

**Previous** **Next** **Save & Deploy** **Cancel**

Here, you can preview the overlay configuration that will be deployed to *BGW\_3*. In this section, you can see that an overlay connection is being established from *Loopback0* on *BGW\_3* to the neighbor with AS Number 65100.

Note that no configuration will be pushed to the external device itself.

- Click **Save and Deploy** to complete the task. This results in the configuration getting pushed to *BGW\_3*. The external connection will appear in the Fabric Extension screen.

Fabric Extension

| Inter-Fabric Connections                 |              |               |                  |                   |                   |                    |                             |          |   |
|------------------------------------------|--------------|---------------|------------------|-------------------|-------------------|--------------------|-----------------------------|----------|---|
| Type                                     | Source Fa... | Source Device | Source Interface | Destination Fa... | Destination De... | Destination Int... | Configuration               | Status   |   |
| <input type="radio"/> MULTISITE_OVERLAY  | SITE_2       | BGW_3         | Loopback0        | CORE              | RS_1              | Loopback0          | <a href="#">View Config</a> | DEPLOYED | X |
| <input type="radio"/> MULTISITE_UNDERLAY | SITE_2       | BGW_3         | Ethernet4/3      | CORE              | RS_1              | Ethernet5/5        | <a href="#">View Config</a> | DEPLOYED |   |

## Overlay Extension from BGW\_3 to RS\_1

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

### IFC Pointers

- Extensions will need to be deleted and then reconfigured in case of deployment failures. Currently there is no option to edit or redeploy an overlay or underlay extension.
- To see the deployment history of a functioning IFC, click the **View Config** hyperlink in the **Configuration** column (step 1 in the image). The Inter-fabric Connections Deployment History page comes up. In this page, the **Source** column refers to the specific IFC number. Click the link in the **Status** column (step 2) to view commands executed for the IFC.

| Type               | Source Fabric | Source Device | Source Interface | Destination Fa... | Destination De... | Destination Int... | Configuration               | Status   |
|--------------------|---------------|---------------|------------------|-------------------|-------------------|--------------------|-----------------------------|----------|
| MULTISITE_OVERLAY  | Easy60000     | n9k-15-bgw    | Loopback0        | Easy7200          | n9k-3             | Loopback0          | <a href="#">View Config</a> | DEPLOYED |
| MULTISITE_UNDERLAY | Easy6...      |               |                  |                   |                   |                    |                             |          |

| Inter-fabric Connections Deployment History for n9k-15-bgw ( FDO20401LB4 ) |             |        |         |                       |                         |
|----------------------------------------------------------------------------|-------------|--------|---------|-----------------------|-------------------------|
| Entity Name                                                                | Entity Type | Source | Status  | Status Description    | Time of Completion      |
| FDO20401LB4                                                                | SWITCH      | IFC-4  | SUCCESS | Successfully deployed | 2018-07-17 08:51:35.892 |

| Command Execution Details for n9k-15-bgw ( FDO20401LB4 ) |         |              |
|----------------------------------------------------------|---------|--------------|
| Config                                                   | Status  | CLI Response |
| router bgp 60000                                         | SUCCESS |              |
| neighbor 10.1.0.4                                        | SUCCESS |              |
| remote-as 7200                                           | SUCCESS |              |
| update-source loopback0                                  | SUCCESS |              |
| ebgp-multihop 5                                          | SUCCESS |              |
| peer-type fabric-external                                | SUCCESS |              |
| address-family l2vpn evpn                                | SUCCESS |              |
| send-community                                           | SUCCESS |              |
| send-community extended                                  | SUCCESS |              |
| rewrite-evpn-if-asn                                      | SUCCESS |              |

You can only see functioning IFCs in this screen. To view functioning and deleted IFCs, you should right-click the switch and click **History** (steps 1 and 2 in the image below). In the Policy Deployment screen that comes up, filter the **Source** column for the IFC (step 3 - *IFC-8, IFC-4*, etc) and click the link in the **Status** column (step 4, below) for detailed information.

The screenshot shows the Cisco DCNM interface for managing EVPN configurations. At the top left, a context menu for the device 'n9k-15' is open, with item 'History' highlighted. A blue arrow points from this menu down to a table titled 'Policy Deployment History for n9k-15-bgw ( FDO20401LB4 )'. This table lists a single entry: 'FDO20401LB4' (Entity Name), 'SWITCH' (Entity Type), 'IFC-8' (Source), 'SUCCESS' (Status), and 'Successfully deployed' (Status Description). Another blue arrow points from this table down to another table titled 'Command Execution Details for n9k-15-bgw ( FDO20401LB4 )'. This table contains four rows of command logs:

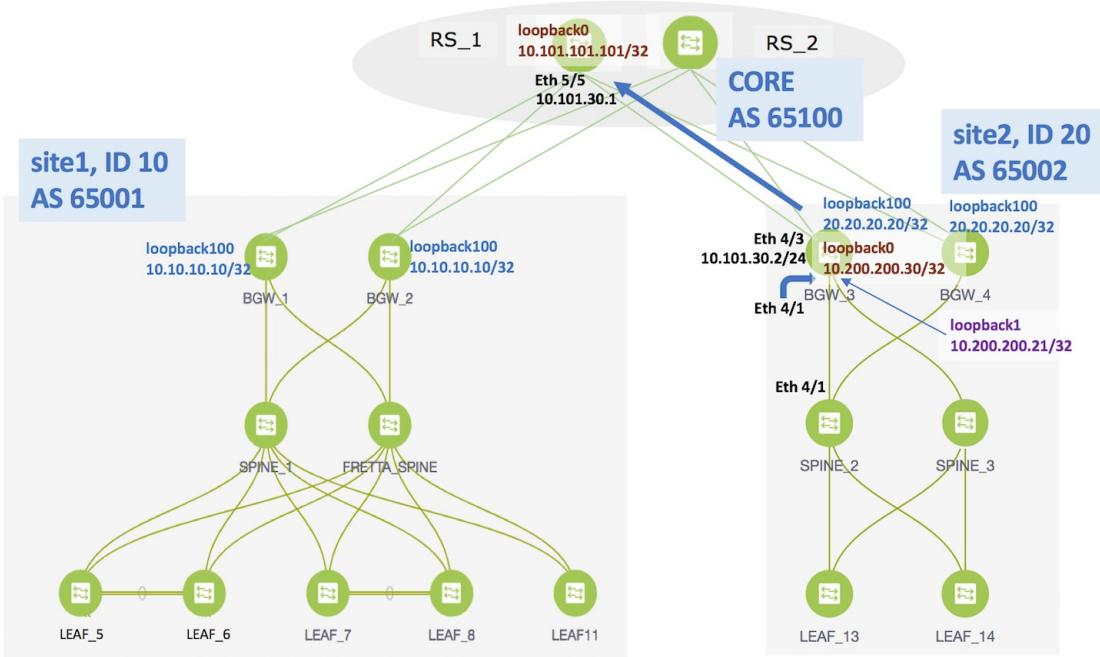
| Config                                   | Status  | CLI Response |
|------------------------------------------|---------|--------------|
| no route-map extcon-rmap-filter deny...  | SUCCESS |              |
| router bgp 60000                         | SUCCESS |              |
| address-family ipv4 unicast              | SUCCESS |              |
| redistribute direct route-map rmap-re... | SUCCESS |              |

- When a destination switch in an IFC is removed, and not available, in DCNM, you will still be able to delete a deployed IFC.

## Other EVPN Multi-Site Configurations

At this stage, overlay and underlay EVPN Multi-Site configurations are provisioned on *BGW\_3* toward *RS\_1* (as shown by the arrow in the figure).

## Other EVPN Multi-Site Configurations



To complete EVPN Multi-Site configurations between *site1* and *site2* using DCNM, you should also configure as follows:

- **On site2**

- EVPN Multi-Site configurations from *BGW\_3* to *RS\_2*.
- EVPN Multi-Site configurations from *BGW\_4* to *RS\_1* and *RS\_2*.

- **On site1**

- EVPN Multi-Site configurations from *BGW\_1* to *RS\_1* and *RS\_2*.
- EVPN Multi-Site configurations from *BGW\_2* to *RS\_1* and *RS\_2*.

- **On the route servers**

- Apart from the DCNM provisioning on the BGWs of *site1* and *site2*, you should enable appropriate configurations on *RS\_1* and *RS\_2* for connectivity between the route servers and the BGWs.

Sample *RS\_1* configurations are provided in the *Appendix* section for your reference.

As noted earlier, the end-to-end Multi-Site configurations through DCNM Top-Down provisioning include these two steps:

- (1) **Multi-Site configurations on the BGWs (*BGW\_1*, *BGW\_2*, *BGW\_3* and *BGW\_4*).**

- (2) **Deploying Networks and VRF Instances on the leaf switches and the BGWs.**

At this stage, the first step explanation is complete. In the next part of the document, the networks' configuration (second step), is explained. After appropriate network configurations on the leaf switches and BGWs, server traffic will flow across the two sites for the deployed and extended networks and VRFs.

# Deploying Networks and VRF Instances

Typically, you create a fabric in DCNM, then create and deploy networks and VRFs on devices within the fabric on leaf switches, and then configure the BGWs for external connectivity. Though the focus of the document is external connectivity with EVPN Multi-Site configurations on BGWs using DCNM, for completeness and right context, network deployment on the BGWs is explained in this section. When EVPN Multi-Site deployment is completed, server traffic from these networks and VRFs on *site2* will pass through a BGW (*BGW\_3* or *BGW\_4*) towards *site1*.


**Note**

For VRF deployment, refer the *Deploying VRF Instances on Border Leafs* section in the chapter *Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - VRF Lite*.

## Deploying Networks on the BGWs

*Before you begin* - In this scenario, we will deploy two networks in *site2*, *MyNetwork\_10000* and *MyNetwork\_10001*, on the BGWs *BGW\_3* and *BGW\_4*. You should ensure that you have already deployed the networks that you want to extend to *site1* on the leaf switches (*LEAF\_13* and *LEAF\_14* in this case).

After deploying the 2 networks on the leaf switches and the BGWs, the networks will be extended to *site1*. To know how to create a new fabric, network, and VRF, see the *Fabrics* section in the *Control* chapter in the *Cisco DCNM LAN Fabric User Guide, Release 11.0(1)*. The procedure:

1. In the Select a Fabric page, click the **Continue** button at the top right part of the screen. The **Networks** page comes up.

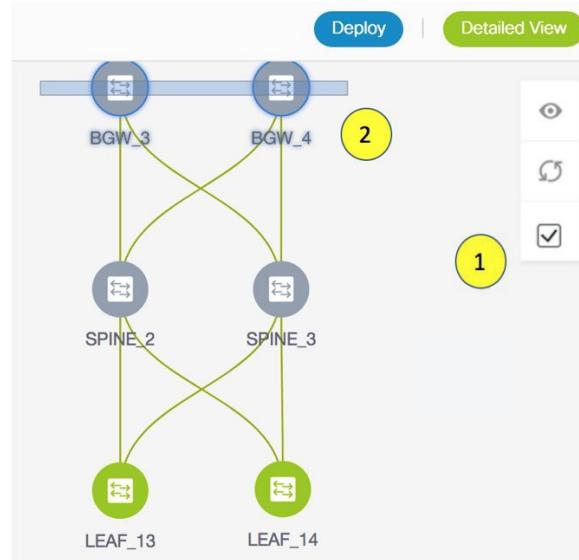
(To access the Select a Fabric page, click **Control > Networks & VRFs**. The LAN Fabric Provisioning page comes up. Click **Continue**. The Select a Fabric page comes up.)

2. We will deploy two new networks *MyNetwork\_10000* and *MyNetwork\_10001* on the BGWs. To do that, select the checkboxes (in the extreme left column).

|                                     | Network Name    | Network ID | VRF Name     | IPv4 Gateway/Subnet | IPv6 Gateway/Prefix | Status     | VLAN ID |
|-------------------------------------|-----------------|------------|--------------|---------------------|---------------------|------------|---------|
| <input checked="" type="checkbox"/> | MyNetwork_10000 | 10000      | MyVRF_200000 | 10.1.10.1/24        | 10:1:A::1/48        | UNDEPLOYED |         |
| <input checked="" type="checkbox"/> | MyNetwork_10001 | 10001      | MyVRF_200000 | 10.1.11.1/24        |                     | UNDEPLOYED | 11      |
| <input type="checkbox"/>            | MyNetwork_10002 | 10002      | MyVRF_200002 | 10.1.12.1/24        | 10:1:C::1/48        | UNDEPLOYED |         |
| <input type="checkbox"/>            | MyNetwork_10003 | 10003      | NA           | 10.1.13.1/24        |                     | UNDEPLOYED |         |

3. Click the **Continue** button at the top right part of the screen. The Network Deployment page (Topology View) comes up. You can deploy networks simultaneously on multiple switches. The selected devices should have the same role (*Leaf, Border Gateway, etc*). So, deploy the selected networks on the BGWs.
4. Select the multi-select check box available at the right part of the page. (displayed as step 1 in the image). Then, click your mouse (or track pad) and drag the cursor across *BGW\_3* and *BGW\_4*. (step 2).

## Deploying Networks on the BGWs



**Note** In the image, you can see that the networks are deployed on the leaf switches (green color indicates *deployed* status). Note that the color code (and hence the deployment state) on switches is contextual and specific to the selection. In this scenario, the deployed state only depicts that networks *MYNetwork10000* and *MYNetwork10001* are deployed on leaf switches *LEAF\_13* and *LEAF\_14*. It does not display information about other (networks and VRFs) deployment instances, if any.

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

## Switches Deploy



*Fabric Name: SITE\_2*

MyNetwork\_10000

MyNetwork\_10001

*Deploy Options:*

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

| <input type="checkbox"/> | Switch | ▲ | VLAN | Extend    | Status |
|--------------------------|--------|---|------|-----------|--------|
| <input type="checkbox"/> | BGW_3  |   | 10   | MULTISITE | NA     |
| <input type="checkbox"/> | BGW_4  |   | 10   | MULTISITE | NA     |

Save

A tab is displayed for each network. Click the checkbox next to the **Switch** column. Both the BGW check boxes will be selected automatically and the **Extension Details** section will appear at the bottom part of the screen.

In the **Extension Details** section, select the **Switch** checkbox (or ensure that you select the check box in each row) and click **Save** (bottom right part of your screen).

## Switches Deploy

Fabric Name: SITE\_2

MyNetwork\_10000    MyNetwork\_10001

Deploy Options:

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

| <input checked="" type="checkbox"/> | Switch | VLAN | Extend    | Status |
|-------------------------------------|--------|------|-----------|--------|
| <input checked="" type="checkbox"/> | BGW_3  | 10   | MULTISITE | NA     |
| <input checked="" type="checkbox"/> | BGW_4  | 10   | MULTISITE | NA     |

Extension Details

| <input checked="" type="checkbox"/> | Switch | Type      | IF_NAME   |
|-------------------------------------|--------|-----------|-----------|
| <input checked="" type="checkbox"/> | BGW_3  | MULTISITE | Loopback0 |
| <input checked="" type="checkbox"/> | BGW_3  | MULTISITE | Loopback0 |
| <input checked="" type="checkbox"/> | BGW_4  | MULTISITE | Loopback0 |
| <input checked="" type="checkbox"/> | DCNM_1 | MULTISITE | Loopback0 |

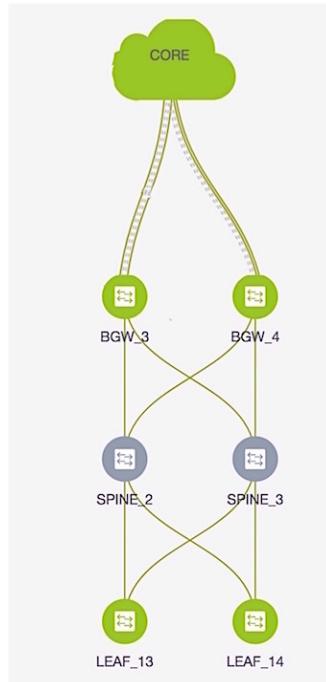
Save

After saving the details in this screen, the Network Deployment screen (Topology view) appears.

*BGW\_3* and *BGW\_4* will be displayed in blue color, indicating pending deployment. If you want to check your configurations again, click on the Preview (eye) icon.

5. After you verify that the configurations that are generated from the profiles are correct for the selected switches, click the **Deploy** button (on the top right part of the screen) to deploy the *MYNetwork10000* and *MYNetwork10001* network configurations on *BGW\_3* and *BGW\_4*.

DCNM shows the deployment status in the topology by highlighting the switch icons with different colors, yellow for *In Progress* and green for *Deployed*.



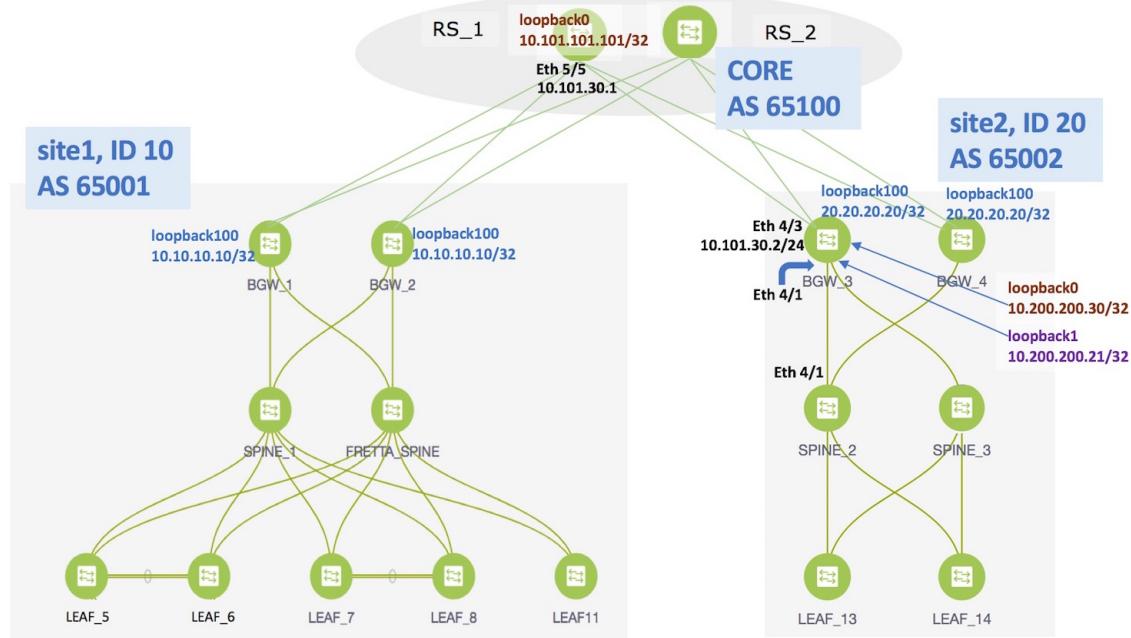
From the snapshot, you can see that the 2 networks *MYNetwork10000* and *MYNetwork10001* have been deployed on the leaf switches and BGWs.

6. After configurations in *site2* are complete, configure the following in *site1* too.

## Configurations in *site1*

Provision the networks *MYNetwork10000* and *MYNetwork10001* on the leaf switches (*LEAF\_5*, *LEAF\_6*, *LEAF\_7*, *LEAF\_8*, *LEAF\_11*) and the BGWs (*BGW\_1* and *BGW\_2*).

## Additional References



As noted in the *EVPN Multi-Site Configuration* section, enable the following for end-to-end configuration:

- Since DCNM does not provision configurations for *RS\_1* and *RS\_2* (devices directly connected to the BGWs), enable appropriate configurations on these devices.
- Configure the EVPN Multi-Site feature on the *site1* BGWs (as explained in this document) so that server traffic from the 2 networks can flow to *site2* and back.

## Additional References

| Document Title and Link                                                 | Document Description                                                      |
|-------------------------------------------------------------------------|---------------------------------------------------------------------------|
| <a href="#">VXLAN EVPN Multi-Site Design and Deployment White Paper</a> | This document explains Multi-Site design and deployment in detail.        |
| <a href="#">Configuring VXLAN EVPN Multi-Site</a>                       | This document explains manual configurations for the Multi-Site solution. |

## Appendix

### Route Server Configurations

***RS\_1* configuration example for the overlay**—The following configurations are enabled on *RS\_1*, and reproduced here for reference.

**Note**

*switch(config)#* refers to the global configuration mode. To access this mode, type the following on your switch: **switch# configure terminal**.

```

switch(config)#

route-map ALL-PATHS permit 100
  set path-selection all advertise
route-map RMAP-REDIST-DIRECT permit 10
  match tag 12345
route-map UNCHANGED permit 10
  set ip next-hop unchanged

switch(config)#

interface loopback0
  ip address 10.101.101.101/32 tag 12345
line vty
router bgp 65100
  router-id 10.101.101.101
  address-family ipv4 unicast
    redistribute direct route-map RMAP-REDIST-DIRECT
    maximum-paths 4
    additional-paths send
    additional-paths receive
    additional-paths selection route-map ALL-PATHS
  address-family l2vpn evpn
    retain route-target all
  template peer OVERLAY-PEERING
    update-source loopback0
    ebgp-multihop 5
    address-family l2vpn evpn
      send-community both
      route-map UNCHANGED out
  neighbor 10.100.100.10
    inherit peer OVERLAY-PEERING
    remote-as 65001
    address-family l2vpn evpn
      rewrite-evpn-rt-asn
      route-map UNCHANGED out
  neighbor 10.100.100.20
    inherit peer OVERLAY-PEERING
    remote-as 65001
    address-family l2vpn evpn
      rewrite-evpn-rt-asn
      route-map UNCHANGED out
  neighbor 10.101.11.2
    remote-as 65101
    update-source Ethernet5/1
    address-family ipv4 unicast
      next-hop-self
  neighbor 10.101.12.2
    remote-as 65101
    update-source Ethernet5/2
    address-family ipv4 unicast
      next-hop-self
  neighbor 10.101.13.2
    remote-as 65102
    update-source Ethernet5/3
    address-family ipv4 unicast
      next-hop-self

```

**Route Server Configurations**

```
neighbor 10.101.14.2
  remote-as 65102
  update-source Ethernet5/4
  address-family ipv4 unicast
    next-hop-self
neighbor 10.101.30.2
  remote-as 65002
  update-source Ethernet5/5
  address-family ipv4 unicast
    next-hop-self
neighbor 10.101.40.2
  remote-as 65002
  update-source Ethernet5/6
  address-family ipv4 unicast
    next-hop-self
neighbor 10.200.200.30
  remote-as 65002
  update-source loopback0
  ebgp-multipath 5
  address-family l2vpn evpn
    rewrite-evpn-rt-asn
    send-community both
    route-map UNCHANGED out
neighbor 10.200.200.40
  remote-as 65002
  update-source loopback0
  ebgp-multipath 5
  address-family l2vpn evpn
    rewrite-evpn-rt-asn
    send-community both
    route-map UNCHANGED out
```



## CHAPTER 8

# Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - VRF Lite

External connectivity from data centers is a prime requirement. Virtual eXtensible Local Area Network (VXLAN) Border Gateway Protocol (BGP) Ethernet VPN (EVPN) based data center fabrics provide east-west connectivity by distributing IP-MAC reachability information among various devices within the fabric. While the EVPN Multi-Site feature provides inter site connectivity, the VRF Lite feature is used for connecting the fabric to an external Layer 3 domain. Tenants, typically represented by virtual routing and forwarding instances (VRFs) can procure external connectivity via special nodes called borders. In this way, tenant workloads in one data center fabric can have Layer 3 connectivity to hosts within the same VRF in other fabrics. This chapter describes LAN Fabric provisioning of the Nexus 9000-based border devices through the Cisco® Data Center Network Manager (DCNM) for the VRF Lite use case. This use case covers VRF extension from border devices connected to edge routers that in turn provide connectivity to the external fabric.

- [Prerequisites , on page 285](#)
- [Sample Scenario, on page 286](#)
- [VRF Lite Configuration , on page 288](#)
- [Deploying VRF Instances on Border Leafs, on page 295](#)
- [Undeploying VRF Instances on the Border Leafs , on page 300](#)
- [Additional References, on page 305](#)
- [Appendix , on page 305](#)

## Prerequisites

- The VRF Lite feature requires Cisco Nexus 9000 Series NX-OS Release 7.0(3)I6(2) or later.
- Familiarity with VXLAN BGP EVPN data center fabric architecture and top-down based LAN fabric provisioning through the DCNM.
- Fully configured VXLAN BGP EVPN fabrics including underlay and overlay configurations on the various leaf and spine devices, external fabric configuration through DCNM, and relevant external fabric device configuration (edge routers, for example).
  - A VXLAN BGP EVPN fabric (and its connectivity to an external Layer 3 domain for north-south traffic flow) can be configured manually or using DCNM. This document explains the process to connect the fabric to an edge router (outside the fabric, towards the external fabric) through DCNM. So, you should know how to configure and deploy VXLAN BGP EVPN and external fabrics through

DCNM. For more details, see the **Control** chapter in the *Cisco DCNM LAN Fabric User Guide, Release 11.0(1)*.

- Ensure that the role of the designated border leaf switches is *Border*. To verify, right-click the switch and click **Set role**. You can see that *(current)* is added to the current role of the switch. If the current role is not *Border* or *Border Gateway*, you should remove the device from the fabric and discover it again through DCNM using the POAP bootstrap option and re-provision the configurations for the device.




---

**Note** For an explanation on the VRF Lite feature, see the [Cisco Programmable Fabric with VXLAN BGP EVPN Configuration Guide](#) document.

---

## Sample Scenario

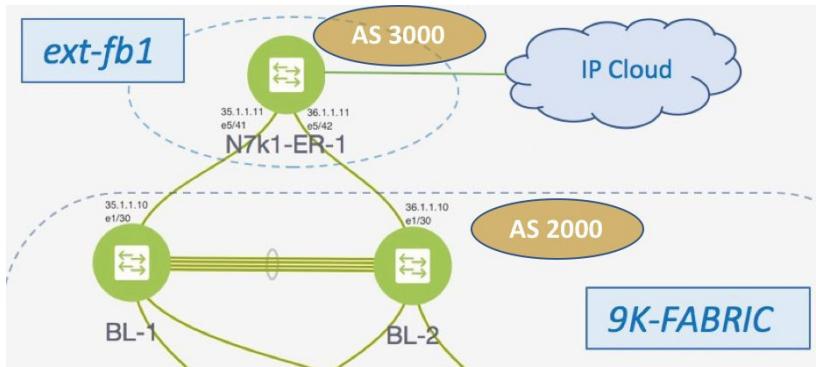
The VRF Lite feature is explained through an example scenario. Consider a VXLAN BGP EVPN fabric, *9K-FABRIC*, whose border devices *BL-1* and *BL-2* are connected through an edge router in the fabric *ext-fb1*, to a shared IP core. This document will show you how to enable Layer 3 north-south traffic between the VXLAN fabric border devices and the edge router.




---

**Note** In this scenario, DCNM allows provisioning for fabric switches and the border devices. The edge router connected to the border devices needs to be manually configured (the edge router *N7k1-ER1* [or *ER1*] in the *ext-fb1* fabric is connected to *BL1* and *BL2* in *9K-FABRIC*).

---



Network configurations for the fabric is provisioned through DCNM. For external Layer 3 reachability from hosts connected to leaf switches within the fabric, border devices need to be provisioned with the appropriate VRF configuration. Multiple border devices in the fabric ensure redundancy in the case of failures as well as effective load distribution.

*N7k1-ER1* (or *ER1*) is directly attached to the 2 border leafs. From the VXLAN fabric's point of view, the edge router belongs to an external fabric, *ext-fb1*, with a different AS number. For representation purposes, the *ext-fb1* fabric is created as an external fabric through DCNM, and *ER1* is associated with it in DCNM.

**Note**

External fabric creation is a prerequisite for this use case. To create an external fabric *ext-fb1* in DCNM, follow these steps.

- 1. Click Control > Fabric Builder.**

The Fabric Builder page comes up.

- 2. Click Create Fabric.**

The Add Fabric screen comes up.

- 3. Enter the fabric name (*ext-fb1*) and select *External\_Fabric* in the Fabric Template drop-down box.**

- 4. Enter the BGP AS number and click Save.**

- 5. The *ext-fb1* fabric is created as an external fabric.**

**VRF Lite**—This requires setting up the border leaf configuration for enabling the VRF Lite feature by establishing eBGP peering from the border leaf to appropriate external devices like the edge routers. In this context, border leafs are special devices that allow clear control and data plane segregation from the fabric domain to the external Layer 3 domain (while allowing for policy enforcement points for any inter-fabric traffic).

The steps involved to enable VRF Lite feature are:

- 1. Connecting the VXLAN BGP EVPN fabric with the edge router**—Top-Down deployment for the VRF Lite feature configures route maps and an eBGP session in the default VRF through an interface (parent interface) connected to the edge router. This is a one-time setup for each edge router connected to a border leaf.
- 2. VRF extensions**—For each VRF that is to be extended, a unique sub interface towards the edge router and an eBGP session through this sub interface is configured on the border leaf. This is a per-VRF configuration. The corresponding configurations have to be manually enabled on the edge router too.

The end-to-end configurations can be split into these 2 steps:

- 1. VRF Lite configurations on the border leafs (*BL-1, BL-2*)**

1. VRF Lite function on *BL-1* and *BL-2* in *9K-FABRIC* that are directly connected to *ER-1*.
2. Configurations on edge router *ER-1* - These configurations are not in the scope of DCNM provisioning and this document. It is mentioned here for completeness and sample configurations are provided in the *Appendix* section.

- 2. Deploying VRF instances on the border leafs (*BL-1, BL-2*)**

For this example, multiple VRFs will be configured on the border leafs in *9K-FABRIC*.

After successful VRF Lite deployment at the border leafs and on the edge router, traffic will flow between them.



**Note** In the DCNM topology view, the lines connecting devices managed by DCNM (for example, *BL-1* to *N7k1-SPINE-1*) symbolize a physical cable connection. They do not indicate that the connection is functional and traffic flows between them.

To start off with, let us consider VRF Lite provisioning on border leafs *BL-1* and *BL-2* through DCNM Top-Down LAN Fabric Provisioning.

## VRF Lite Configuration

### VRF Lite Configuration (on *BL-1* towards *ER-1* in *9K-FABRIC*)

1. Click **Control > Network & VRFs**. The LAN Fabric Provisioning page appears.
2. Click **Continue**. The Select a Fabric page comes up.
3. Select *9K-FABRIC* from the drop-down box since you are configuring border leaf *BL-1* in the fabric *9K-FABRIC*.

In the same page, click **Fabric Extension Setup** since the purpose of this task is to allow *9K-FABRIC* to communicate to the edge router in the external fabric.

The Fabric Extension screen comes up.

| Type              | Source Fabric | Source Device | Source Interface | Destination Fabric | Destination Device | Destination Interface | Configuration | Status |
|-------------------|---------------|---------------|------------------|--------------------|--------------------|-----------------------|---------------|--------|
| No data available |               |               |                  |                    |                    |                       |               |        |

The **Inter-Fabric Connections** section lists previously created external connections from the border leafs in *9K-FABRIC*. This section is empty as this is the first time you are adding an external connection. Each row represents a physical or logical connection between a border leaf in *9K-FABRIC* and the edge router in the *ext-fb1* fabric. For each connection, the source fabric, source device, source interface, destination fabric, destination device, and destination interface are listed along with the type of external connectivity.

*To extend the fabric through VRF-Lite, you should first create an extension.*

### Extension from *BL-1* to *ER-1*

1. Click on the + icon (at the top left part of the screen) to add a new external connection. The **Add Inter-Fabric Connection** screen appears.

## Add Inter-Fabric Connections

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

• • •

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

(i) VRF\_LITE: Set switch role - Border; MULTISITE: Set switch role - "Border Gateway"

By default, *VRF\_LITE* is populated in the **Extension Type** field.

**Base Template**—By default, the *ext\_base\_setup* base template is populated. This template represents a one-time configuration pushed to the border leaf *BL-1*.

**Extension Template**—*ext\_fabric\_setup*, as the name indicates, represents the template that outputs the configuration required to setup the connection between the border leaf and the edge router. As opposed to the configuration represented by the Base Template that is applied only once per border leaf, the Extension Template generated configuration is executed once for every connection between a border leaf and the edge router.



### Note

These templates are auto-populated with corresponding pre-packaged default templates based on your selection. You can add, edit or delete user-defined templates. For more details, see the *Template Library* section in the *Control* chapter .

**Source Fabric**—This field is pre-populated with *9K-FABRIC* since the VRF Lite connection is between *BL-1* in *9K-FABRIC* and *ER-1* in the *ext-fb1* fabric.

**Destination Fabric**—Choose *ext-fb1*.

**Source Device** and **Source Interface**—Choose *BL-1* as the source device and an Ethernet interface that needs to be connected to *ER-1*.

**Destination Device** and **Destination Interface**—Choose *ER-1* as the destination device and the Ethernet interface that connects to the border leaf *BL-1*.

## Extension from BL-1 to ER-1

Note that based on the selection of the source device and source interface, the destination information will be auto-populated based on CDP information if available. There is extra validation performed to ensure that the destination external device is indeed part of the destination fabric.



**Note** If the destination device is a non Cisco Nexus device (such as a Cisco ASR or Catalyst device), then you can manually type in the switch name and the interface name (for example, *Catalyst9400* and *Ethernet1/1*) in these fields.

After filling up the Fabric Interconnect section, the screen looks like this.

### Add Inter-Fabric Connections

|                                                                                                                                                                    |                  |                    |   |                    |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|--------------------|---|--------------------|
| 1 Fabric Interconnect                                                                                                                                              | →                | 2 Define Variables | → | 3 Preview & Deploy |
| • • •                                                                                                                                                              |                  |                    |   |                    |
| * Extension Type                                                                                                                                                   | VRF_LITE         |                    |   |                    |
| * Base Template                                                                                                                                                    | ext_base_setup   |                    |   |                    |
| * Extension Template                                                                                                                                               | ext_fabric_setup |                    |   |                    |
| * Source Fabric                                                                                                                                                    | 9K-FABRIC        |                    |   |                    |
| * Destination Fabric                                                                                                                                               | ext-fb1          |                    |   |                    |
| * Source Device                                                                                                                                                    | BL-1             |                    |   |                    |
| * Source Interface                                                                                                                                                 | Ethernet1/30     |                    |   |                    |
| (i) VRF_LITE: Set switch role - Border; MULTISITE: Set switch role - "Border"                                                                                      |                  |                    |   |                    |
| * Destination Device                                                                                                                                               | N7k1-ER-1        |                    |   |                    |
| * Destination Interface                                                                                                                                            | Ethernet5/41     |                    |   |                    |
| <input type="button" value="Previous"/> <input type="button" value="Next"/> <input type="button" value="Save &amp; Deploy"/> <input type="button" value="Cancel"/> |                  |                    |   |                    |

- Click **Next** to go to the **Define Variables** section. The fields are:

**IF\_NAME**—In this field, the interface name is auto-populated from the previous step.

**Interface IP\_MASK**—Fill up this field with the IP address and mask of the *BL-1* interface that connects to *ER-1*.

**NEIGHBOR\_IP**—Fill up this field with the IP address of the *ER-1* interface that connects to *BL-1*.

**NEIGHBOR ASN**—In this field, the AS number of *ER-1*'s fabric (*ext-fb1*) will be auto-populated.

**Extension Type**—In this field, **VRF\_Lite** will be auto-populated.

A sample screenshot of the fully filled up screen:

## Add Inter-Fabric Connections

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

Network Profile

General

- \* IF\_NAME: Ethernet1/30
- \* IP\_MASK: 35.1.1.10/24
- \* NEIGHBOR\_IP: 35.1.1.11
- \* NEIGHBOR ASN: 3000
- \* Extension Type: VRF\_LITE

Previous Next Save & Deploy Cancel

- Click **Next** to go to the **Preview and Deploy** section. The two sections of the screen are shown in the 2 images:

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

Switch: BL-1

Generated Configuration:

```

ip prefix-list default-route seq 5 permit 0.0.0.0/0 le 1
ip prefix-list host-route seq 5 permit 0.0.0.0/0 eq 32
route-map EXTCON-RMAP-FILTER deny 10
  match ip address prefix-list default-route
route-map EXTCON-RMAP-FILTER deny 20
  match ip address prefix-list host-route
route-map EXTCON-RMAP-FILTER permit 1000

ipv6 prefix-list default-route-v6 seq 5 permit 0::/0
ipv6 prefix-list host-route-v6 seq 5 permit 0::/0 eq 128
route-map EXTCON-RMAP-FILTER-V6 deny 10
  match ipv6 address prefix-list default-route-v6
route-map EXTCON-RMAP-FILTER-V6 deny 20
  match ipv6 address prefix-list host-route-v6
route-map EXTCON-RMAP-FILTER-V6 permit 1000

interface Ethernet1/30

```

Generated Config

Previous Next Save & Deploy Cancel

## Extension from BL-1 to ER-1

Add Inter-Fabric Connections

Switch: BL-1

Generated Configuration:

```

interface Ethernet1/30
  no switchport
  ip address 35.1.1.10/24
  no shutdown

router bgp 2000
  address-family ipv4 unicast
    redistribute direct route-map RMAP-REDIST-DIRECT
    neighbor 35.1.1.11 remote-as 3000
    update-source Ethernet1/30
    address-family ipv4 unicast
      next-hop-self

```

Previous Next Save & Deploy Cancel

**Note**

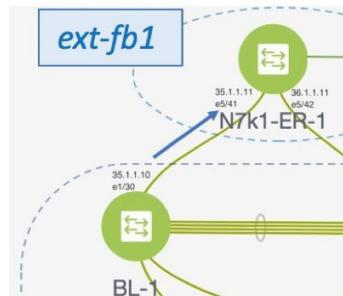
If a VXLAN BGP EVPN fabric border leaf is connected to more than one edge router, the prefix-list and route map configurations are pushed only for the first fabric extension instance. Similarly when deleting fabric extension instances on a border leaf, the global configurations (prefix-list and route-maps) are removed from the border leaf only after the last fabric extension instance is deleted.

In this screen, you can preview the configurations that will be deployed to *BL-1*. Note that no configuration will be pushed to the external device (edge router) itself.

A one-time configuration of route maps along with the parent interface connection is displayed. Also, you can see that BGP peering information in the default routing table is configured for *BL-1*. The corresponding BGP configurations should be manually enabled on *ER-1*.

- Click **Save and Deploy** to complete the task. This results in the configuration getting pushed to *BL-1*. The external connection will appear in the Fabric Extension screen, under **Inter-Fabric Connections**.

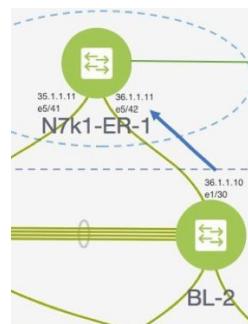
At this stage, an extension is enabled from *BL-1* to *ER-1*, as indicated by the arrow in the image.



Next, you need to enable an extension from *BL-2* to *ER-1* too.

### VRF Lite Configuration (on *BL-2* towards *ER-1* in 9K-FABRIC)

As described in the previous section, enable an extension from *BL-2* to *ER-1*. After configurations are pushed to *BL-2*, an extension will be enabled from *BL-2* to *ER-1*, as shown in the screen shot.



A preview of the configurations on *BL-2* is given in these 2 screen shots.

## VRF Lite Configuration (on BL-2 towards ER-1 in 9K-FABRIC)

## Add Inter-Fabric Connections



1 Fabric Interconnect

2 Define Variables

3 Preview &amp; Deploy



Switch: BL-2

## Generated Configuration:

```

ip prefix-list default-route seq 5 permit 0.0.0.0/0 le 1
ip prefix-list host-route seq 5 permit 0.0.0.0/0 eq 32
route-map EXTCON-RMAP-FILTER deny 10
  match ip address prefix-list default-route
route-map EXTCON-RMAP-FILTER deny 20
  match ip address prefix-list host-route
route-map EXTCON-RMAP-FILTER permit 1000

ipv6 prefix-list default-route-v6 seq 5 permit 0::/0
ipv6 prefix-list host-route-v6 seq 5 permit 0::/0 eq 128
route-map EXTCON-RMAP-FILTER-V6 deny 10
  match ipv6 address prefix-list default-route-v6
route-map EXTCON-RMAP-FILTER-V6 deny 20
  match ip address prefix-list host-route-v6
route-map EXTCON-RMAP-FILTER-V6 permit 1000

interface Ethernet1/30

```

Generated Config

[Previous](#)[Next](#)[Save & Deploy](#)[Cancel](#)

## Add Inter-Fabric Connections



1 Fabric Interconnect



2 Define Variables



3 Preview &amp; Deploy



Switch: BL-2

Generated Configuration:

```
interface Ethernet1/30
  no switchport
  ip address 36.1.1.10/24
  no shutdown

router bgp 2000
  address-family ipv4 unicast
    redistribute direct route-map RMAP-REDIST-DIRECT
    neighbor 36.1.1.11 remote-as 3000
    update-source Ethernet1/30
    address-family ipv4 unicast
    next-hop-self
```

[Previous](#)[Next](#)[Save & Deploy](#)[Cancel](#)

## Edge Router Configurations

Apart from the DCNM provisioning on the border leafs in the two fabrics, you should also enable appropriate configurations on *ER-1* for connectivity between the edge router and the border leafs. Sample *ER-1* configuration is provided in the *Appendix* section for your reference.

**What to do next**—As noted earlier, the end-to-end VRF-Lite configurations through DCNM Top-Down provisioning includes these 2 steps:

1. **VRF Lite configurations on the border leafs (*BL-1*, *BL-2*)**
2. **Deploying VRF Instances on the border leafs (*BL-1*, *BL-2*)**

At this stage, the first step explanation is complete. The next section explains how VRF extension configuration is pushed to the border leafs.

## Deploying VRF Instances on Border Leafs

*Before you begin*—In this scenario, we will deploy three VRF instances, *MyVRF-50016*, *MyVRF-50018*, and *MyVRF-50019* on the border leafs *BL-1* and *BL-2* in *9K-FABRIC*. You should ensure that you have already deployed the corresponding network(s) on the fabric's leaf switches.

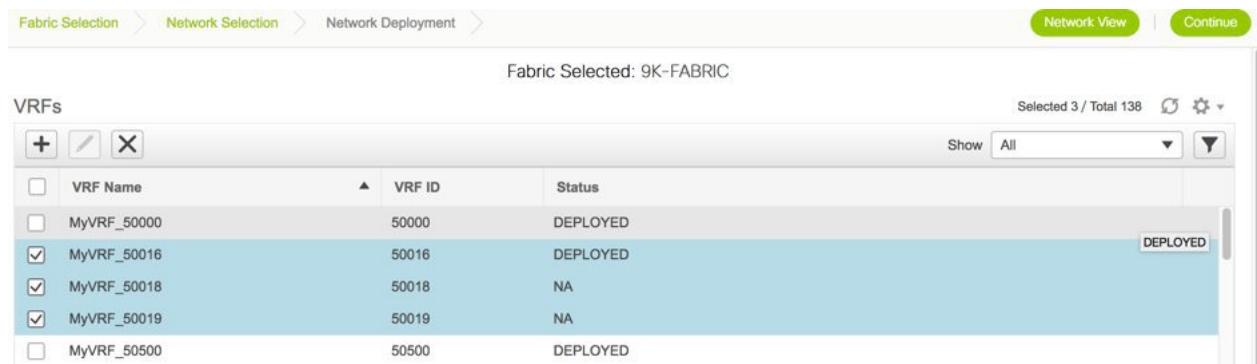
## Deploying VRF Instances on Border Leafs

After deploying one network on the leaf switches, you will have to deploy the associated VRF on the border leafs so that the network(s) can be extended from/to the *9K-FABRIC*. To know how to create a fabric, and networks and VRFs, see the *Control* chapter in the *Cisco DCNM LAN Fabric User Guide, Release 11.0(1)*.

In the Select a Fabric page, ensure that you select *9K-FABRIC* in the drop-down box and click **Continue** (at the top right part of the screen). After clicking **Continue**, the **Networks** page comes up.

Click on **VRF View**. The **VRFs** page comes up.

We will deploy 3 new VRF instances *MyVRF-50016*, *MyVRF-50018*, and *MyVRF-50019* on the border leafs. To do that, select the checkboxes (in the extreme left column).



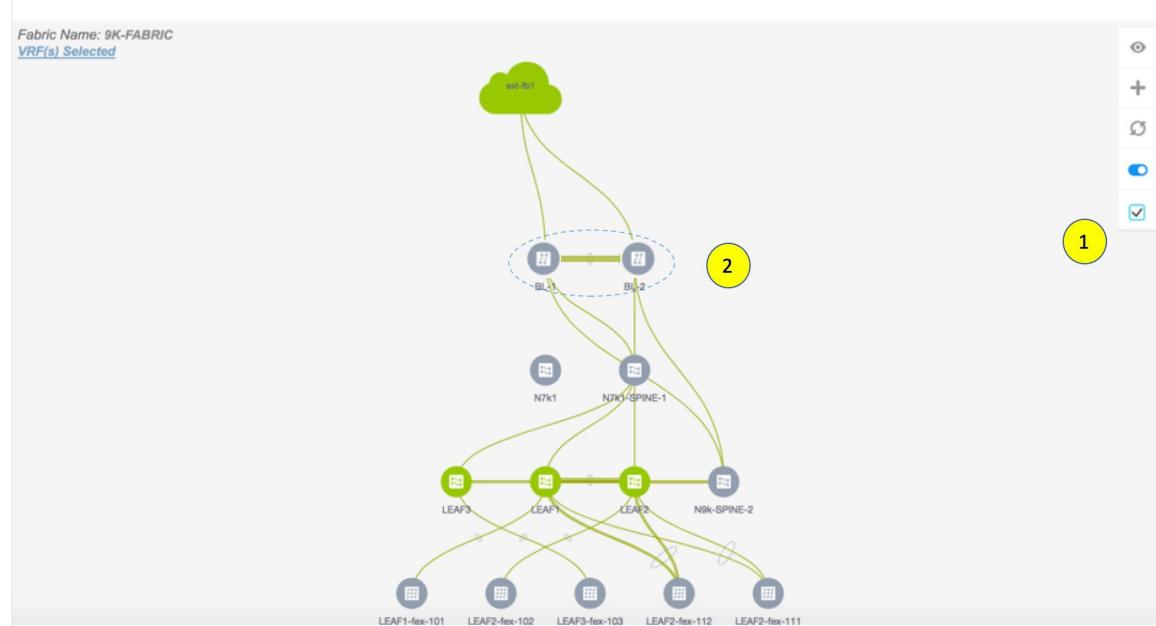
The screenshot shows the 'VRFs' section of the Cisco DCNM interface. At the top, it says 'Fabric Selected: 9K-FABRIC'. Below is a table with columns: 'VRF Name', 'VRF ID', and 'Status'. Three rows are selected, indicated by a blue background:

| VRF Name                                        | VRF ID | Status   |
|-------------------------------------------------|--------|----------|
| MyVRF_50000                                     | 50000  | DEPLOYED |
| <input checked="" type="checkbox"/> MyVRF_50016 | 50016  | DEPLOYED |
| <input checked="" type="checkbox"/> MyVRF_50018 | 50018  | NA       |
| <input checked="" type="checkbox"/> MyVRF_50019 | 50019  | NA       |
| MyVRF_50500                                     | 50500  | DEPLOYED |

Click the **Continue** button at the top right part of the screen. The VRF Deployment page (Topology View) comes up. You can deploy VRFs on multiple switches simultaneously, but with the same role. So, deploy the selected VRFs on the border leafs.



- Note** In the image, you can see that the VRF instances are deployed on the leaf switches (green color indicates deployed status). Note that the color code, and hence the deployment state on switches is contextual and specific to the selection. In this scenario, the deployed state only depicts that the 3 selected VRFs are deployed on leaf switches LEAF3, LEAF1 and LEAF2. It does not display information about other VRF deployment instances, if any.



Select the multi-select check box from the panel of options available (*Step 1* in the image).

Then, click your mouse (or track pad) and drag the cursor across *BL-1* and *BL-2* (*Step 2* in the image).

Immediately, the **Switches Deploy** screen (for VRFs) appears. A tab is displayed for each VRF.

Click the checkbox next to the **Switch** column. Both the border leaf check boxes will be selected automatically. Alternatively, you can select check boxes next to the switches.

#### Switches Deploy

*Fabric Name: 9K-FABRIC*

MyVRF\_50016   MyVRF\_50018   MyVRF\_50019

#### Deploy Options:

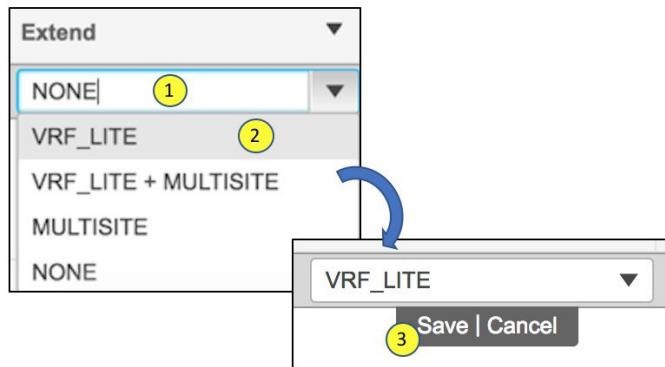
ⓘ Select the row and click on the cell to edit and save changes

| <input type="checkbox"/>            | Switch | VLAN | Extend | Status |
|-------------------------------------|--------|------|--------|--------|
| <input checked="" type="checkbox"/> | BL-1   | 2001 | NONE   | NA     |
| <input checked="" type="checkbox"/> | BL-2   | 2001 | NONE   | NA     |

Click on **NONE** in the **Extend** column, select **VRF\_LITE** and click on the **Save** button below it.

Repeat this action for the second row too. A sample screenshot:

## Deploying VRF Instances on Border Leafs



This creates a VRF Lite extension for this VRF, as seen in the **Extension Details** section that appears at the bottom part of the screen.

In the **Extension Details** section, select the **Source Switch** checkbox (or ensure that you select the check box in each row). This is how the screen looks when you select both the switches in the **Extension Details** section.

The corresponding dot1Q tag for the VRF is auto-populated in the **DOT1Q\_ID** field.

**Switches Deploy**

**Fabric Name:** 9K-FABRIC

MyVRF\_50016 MyVRF\_50018 MyVRF\_50019

**Deploy Options:**

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

| Switch                                   | VLAN | Extend   | Status  |
|------------------------------------------|------|----------|---------|
| <input checked="" type="checkbox"/> BL-1 | 2001 | VRF_LITE | PENDING |
| <input checked="" type="checkbox"/> BL-2 | 2001 | VRF_LITE | PENDING |

**Extension Details**

| Source Switch                            | Type     | IF_NAME      | DOT1Q_ID | IP_MASK      |
|------------------------------------------|----------|--------------|----------|--------------|
| <input checked="" type="checkbox"/> BL-1 | VRF_LITE | Ethernet1/30 | 3        | 35.1.1.10/24 |
| <input checked="" type="checkbox"/> BL-2 | VRF_LITE | Ethernet1/30 | 3        | 36.1.1.10/24 |

**Save**

Now, select the *MyVRF\_50018* and *MyVRF\_50019* and similarly update relevant parameters.

Click the **Save** button at the bottom right part of the Switches Deploy screen to save all VRFs' configurations on the selected switches. The VRF Deployment screen (Topology view) appears.

*BL-1* and *BL-2* icons will be displayed in blue color, indicating that a deployment is pending. If you want to check your configurations, click on the Preview (eye) icon.

**Preview Configuration**

Select a Switch: BL-1 Select a VRF: MyVRF\_50016

```
Generated Configuration:
maximum-paths ibgp 2

network 0::/0

neighbor 35.1.1.11 remote-as 3000
address-family ipv4 unicast
send-community both
route-map EXTCOM-RMAP-FILTER out

neighbor 35.1.1.12 remote-as 3000
address-family ipv6 unicast
send-community both
route-map EXTCOM-RMAP-FILTER-V6 out

interface Ethernet1/30.3
encapsulation dot1q 3
vrf member MyVRF_50016
ip address 35.1.1.10/24
ipv6 address 35.1.1.1:1::1/64
no shutdown

configure terminal
```

**Preview Configuration**

Select a Switch: BL-2 Select a VRF: MyVRF\_50019

```
Generated Configuration:
configure profile 9K-FABRIC-Default_VRF_Extension-50019
vlan 2003
vn-segment 50019
interface vlan 2003
vrf member MyVRF_50019
ip forward
ipv6 forward
no ip redirects
no ipv6 redirects
mtu 9216
no shut

interface nve 1
member vni 50019 associate-vrf

vrf context MyVRF_50019
vni 50019
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn

ip route 0/0 36.1.1.11
address-family ipv6 unicast
route-target both auto
route-target both auto evpn

ipv6 route 0.../0 36.1.1.1...0
```

You can select a switch and a VRF to view corresponding configurations. Configuration details of *MyVRF\_5016* that is pushed to *BL-1* are included in the Appendix section.

After you verify that the configurations that are generated from the profiles are correct for the selected switches, click the **Deploy** button (on the top right part of the Topology View screen) to deploy the *MyVRF\_50016*, *MyVRF\_50018*, and *MyVRF\_50019* VRF configurations on *BL-1* and *BL-2*.

DCNM shows the deployment status in the topology by highlighting the switch icons with different colors, yellow for In Progress, green for Deployed, and red for Out of sync status.

When the switch icons turn green, it indicates that the *MyVRF\_50016*, *MyVRF\_50018*, and *MyVRF\_50019* VRF configurations have been deployed on the border leafs of the *9K-FABRIC*. You can also click the **Detailed View** option to see the status.

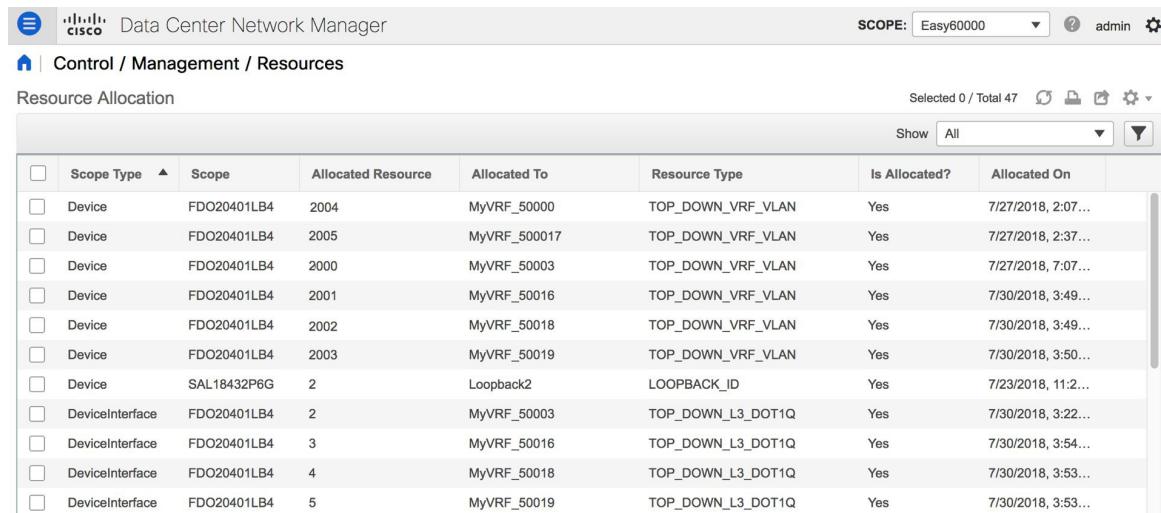
After configurations in *9K-FABRIC* are complete, you should enable configurations in *Fabric2* too.

## Resources

The Resources page (Resource Allocation section) gives information of all the resources allocated or deployed on each device per fabric. This includes the network VLANs, VRF VLANs, and the sub interface dot1q identifiers employed for the VRF Lite extension. Once a VRF is undeployed, the associated resources in the Resource Allocation section will be unallocated and updated immediately.

To access the Resource Allocation page, click **Control > Management > Resources**.

As we can see in the screenshot below, after deploying VRF instances *MyVRF\_50016*, *MyVRF\_50018* and *MyVRF\_50019* on the border leafs, the associated VLAN-VRF mapping is displayed in the Resource Allocation screen.



The screenshot shows the Cisco DCNM interface with the title "Data Center Network Manager". The top navigation bar includes "Control / Management / Resources", "SCOPE: Easy60000", and user "admin". The main content area is titled "Resource Allocation" with a sub-header "Selected 0 / Total 47". A table lists 14 rows of resource allocation details:

| <input type="checkbox"/> | Scope Type      | Scope       | Allocated Resource | Allocated To | Resource Type     | Is Allocated? | Allocated On       |
|--------------------------|-----------------|-------------|--------------------|--------------|-------------------|---------------|--------------------|
| <input type="checkbox"/> | Device          | FDO20401LB4 | 2004               | MyVRF_50000  | TOP_DOWN_VRF_VLAN | Yes           | 7/27/2018, 2:07... |
| <input type="checkbox"/> | Device          | FDO20401LB4 | 2005               | MyVRF_500017 | TOP_DOWN_VRF_VLAN | Yes           | 7/27/2018, 2:37... |
| <input type="checkbox"/> | Device          | FDO20401LB4 | 2000               | MyVRF_50003  | TOP_DOWN_VRF_VLAN | Yes           | 7/27/2018, 7:07... |
| <input type="checkbox"/> | Device          | FDO20401LB4 | 2001               | MyVRF_50016  | TOP_DOWN_VRF_VLAN | Yes           | 7/30/2018, 3:49... |
| <input type="checkbox"/> | Device          | FDO20401LB4 | 2002               | MyVRF_50018  | TOP_DOWN_VRF_VLAN | Yes           | 7/30/2018, 3:49... |
| <input type="checkbox"/> | Device          | FDO20401LB4 | 2003               | MyVRF_50019  | TOP_DOWN_VRF_VLAN | Yes           | 7/30/2018, 3:50... |
| <input type="checkbox"/> | Device          | SAL18432P6G | 2                  | Loopback2    | LOOPBACK_ID       | Yes           | 7/23/2018, 11:2... |
| <input type="checkbox"/> | DeviceInterface | FDO20401LB4 | 2                  | MyVRF_50003  | TOP_DOWN_L3_DOT1Q | Yes           | 7/30/2018, 3:22... |
| <input type="checkbox"/> | DeviceInterface | FDO20401LB4 | 3                  | MyVRF_50016  | TOP_DOWN_L3_DOT1Q | Yes           | 7/30/2018, 3:54... |
| <input type="checkbox"/> | DeviceInterface | FDO20401LB4 | 4                  | MyVRF_50018  | TOP_DOWN_L3_DOT1Q | Yes           | 7/30/2018, 3:53... |
| <input type="checkbox"/> | DeviceInterface | FDO20401LB4 | 5                  | MyVRF_50019  | TOP_DOWN_L3_DOT1Q | Yes           | 7/30/2018, 3:53... |

The VRF instances *MyVRF\_50016*, *MyVRF\_50018*, and *MyVRF\_50019* are deployed on *BL-I*, with their corresponding VLANs *2001*, *2002*, and *2003*.

Also, the corresponding dot1Q IDs *3*, *4*, and *5* are displayed

## Undeploying VRF Instances on the Border Leafs

VRFs can be deployed/undeployed on the border leafs. The following steps will demonstrate undeployment of VRFs on the border leafs.

For *9K-FABRIC*, navigate to the **Networks** page and click **VRF View**. The VRFs page will be displayed.

Select *MyVRF-50018* and *MyVRF-50019* and click **Continue**.

Fabric Selected: 9K-FABRIC

VRFs

| <input type="checkbox"/> VRF Name               | VRF ID | Status     |
|-------------------------------------------------|--------|------------|
| <input type="checkbox"/> MyVRF_50000            | 50000  | DEPLOYED   |
| <input type="checkbox"/> MyVRF_50016            | 50016  | DEPLOYED   |
| <input checked="" type="checkbox"/> MyVRF_50018 | 50018  | DEPLOYED   |
| <input checked="" type="checkbox"/> MyVRF_50019 | 50019  | DEPLOYED   |
| <input type="checkbox"/> MyVRF_50500            | 50500  | DEPLOYED   |
| <input type="checkbox"/> VRF 50011              | 50011  | UNDEPLOYED |

The Topology View page is displayed. Follow similar steps as described in the Deploying VRFs section on the border leafs.

Select *BL-1* and *BL-2* switches in the topology page. The **Switches Deploy** screen will be displayed.

A tab is displayed for each VRF. *MyVRF\_50018* is currently selected in the below screenshot.

Switches Deploy

Fabric Name: 9K-FABRIC

MyVRF\_50018    MyVRF\_50019

Deploy Options:

① Select the row and click on the cell to edit and save changes

| <input type="checkbox"/> Switch          | VLAN | Extend   | Status   |
|------------------------------------------|------|----------|----------|
| <input checked="" type="checkbox"/> BL-1 | 2002 | VRF_LITE | DEPLOYED |
| <input checked="" type="checkbox"/> BL-2 | 2002 | VRF_LITE | DEPLOYED |

Extension Details

| <input type="checkbox"/> Source Switch   | Type     | IF_NAME      | DOT1Q_ID | IP_MASK      |
|------------------------------------------|----------|--------------|----------|--------------|
| <input checked="" type="checkbox"/> BL-1 | VRF_LITE | Ethernet1/30 | 4        | 35.1.1.10/24 |
| <input checked="" type="checkbox"/> BL-2 | VRF_LITE | Ethernet1/30 | 4        | 36.1.1.10/24 |

Double click the checkbox next to the **Switch** column or uncheck the check box next to *BL-1* and *BL-2*. Both of the check boxes will be de-selected and the **Extension Details** section will disappear at the bottom part of the screen.

## Undeploying VRF Instances on the Border Leafs

Switches Deploy

**Fabric Name:** 9K-FABRIC

|             |             |
|-------------|-------------|
| MyVRF_50018 | MyVRF_50019 |
|-------------|-------------|

**Deploy Options:**

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

| <input type="checkbox"/> Switch | VLAN | Extend   | Status   |
|---------------------------------|------|----------|----------|
| <input type="checkbox"/> BL-1   | 2002 | VRF_LITE | DEPLOYED |
| <input type="checkbox"/> BL-2   | 2002 | VRF_LITE | DEPLOYED |

**Save**

Now, select *MyVRF\_50019* and update similarly.

Switches Deploy

**Fabric Name:** 9K-FABRIC

|             |             |
|-------------|-------------|
| MyVRF_50018 | MyVRF_50019 |
|-------------|-------------|

**Deploy Options:**

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

| <input type="checkbox"/> Switch | VLAN | Extend   | Status   |
|---------------------------------|------|----------|----------|
| <input type="checkbox"/> BL-1   | 2003 | VRF_LITE | DEPLOYED |
| <input type="checkbox"/> BL-2   | 2003 | VRF_LITE | DEPLOYED |

**Save**

Click on the **Save** button at the bottom right part of the Switches Deploy screen to undeploy all VRF configurations on the selected switches. The VRF Deployment screen (Topology view) appears.

Similar to the deployment process, the *BL-1* and *BL-2* switch icons will be displayed in blue color, indicating pending undeployment. You can preview the information by clicking the Preview (eye) icon.

The configurations for *MyVRF\_50018* on *BL-1* switch will be removed as displayed in the following screen. You can select a switch and VRF to view corresponding configurations.

## Preview Configuration



Select a Switch:

BL-1

Select a VRF

MyVRF\_50018

Generated Configuration:

```
configure terminal
no apply profile 9K-FABRIC-Default_VRF_Extension-50018
no configure profile 9K-FABRIC-Default_VRF_Extension-50018
```

After you verify that the configuration profiles that will be removed are correct for the selected switches, click the **Deploy** button (on the top right part of the screen) to undeploy the *MyVRF\_50018* and *MyVRF\_50019* configurations on *BL-1* and *BL-2*.

## Resources Update

To access the Resource Allocation page, click **Control > Management > Resources**. After undeploying the VRFs *MyVRF\_50018* and *MyVRF\_50019* on the border leafs, the Resource Allocation page has the associated VLAN-VRF mapping removed.

The screenshot shows the 'Resource Allocation' page in Cisco DCNM. The table lists various resources and their allocations. The columns are: Scope Type, Scope, Allocated Resource, Allocated To, Resource Type, Is Allocated?, and Allocated On. The data in the table is as follows:

| Scope Type      | Scope       | Allocated Resource | Allocated To | Resource Type     | Is Allocated? | Allocated On       |
|-----------------|-------------|--------------------|--------------|-------------------|---------------|--------------------|
| Device          | FDO20401LB4 | 2004               | MyVRF_50000  | TOP_DOWN_VRF_VLAN | Yes           | 7/27/2018, 2:07... |
| Device          | FDO20401LB4 | 2005               | MyVRF_500017 | TOP_DOWN_VRF_VLAN | Yes           | 7/27/2018, 2:37... |
| Device          | FDO20401LB4 | 2000               | MyVRF_50003  | TOP_DOWN_VRF_VLAN | Yes           | 7/27/2018, 7:07... |
| Device          | FDO20401LB4 | 2001               | MyVRF_50016  | TOP_DOWN_VRF_VLAN | Yes           | 7/30/2018, 3:49... |
| Device          | SAL18432P6G | 2                  | Loopback2    | LOOPBACK_ID       | Yes           | 7/23/2018, 11:2... |
| DeviceInterface | FDO20401LB4 | 2                  | MyVRF_50003  | TOP_DOWN_L3_DOT1Q | Yes           | 7/30/2018, 3:22... |
| DeviceInterface | FDO20401LB4 | 3                  | MyVRF_50016  | TOP_DOWN_L3_DOT1Q | Yes           | 7/30/2018, 3:54... |

In the screenshot, it shows that *MyVRF\_50018* and *MyVRF\_50019* that was deployed on *BL-1* with VLAN 2002 and 2003 are now removed/unallocated.

## Remove VRF Lite Inter-fabric configuration on vPC border leafs

VRF Lite configuration can also be removed in a similar manner as long as there are no VRF extensions enabled over that connection. The following steps will demonstrate removal of *BL-1* and *BL-2* VRF Lite connections.

Follow similar steps as described in the VRF Lite configuration for *BL-1* in *9K-FABRIC*.

1. Click **Control > Networks & VRFs**.
2. Select *9K-FABRIC* from the drop-down box and click **Fabric Extension Setup**. The **Fabric Extension** screen comes up

### Remove VRF Lite Inter-fabric configuration on vPC border leafs

| Fabric Extension               |               |               |                  |                   |                   |                    |                             |          |  |
|--------------------------------|---------------|---------------|------------------|-------------------|-------------------|--------------------|-----------------------------|----------|--|
| Inter-Fabric Connections       |               |               |                  |                   |                   |                    |                             |          |  |
| Type                           | Source Fabric | Source Device | Source Interface | Destination Fa... | Destination De... | Destination Int... | Configuration               | Status   |  |
| <input type="radio"/> VRF_LITE | 9K-FABRIC     | BL-1          | Ethernet1/30     | ext-fb1           | N7k1-ER-1         | Ethernet5/41       | <a href="#">View Config</a> | DEPLOYED |  |
| <input type="radio"/> VRF_LITE | 9K-FABRIC     | BL-2          | Ethernet1/30     | ext-fb1           | N7k1-ER-1         | Ethernet5/42       | <a href="#">View Config</a> | DEPLOYED |  |

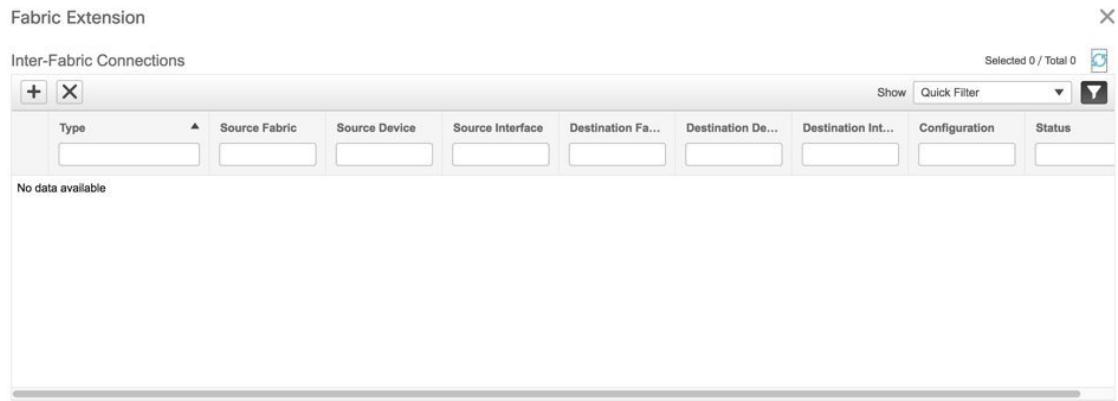
3. Click on the radio button next to **VRF\_LITE** in the first row with Source Device *BL-1*.
4. Click the **X** button to delete this entry.

| Fabric Extension                          |               |               |                  |                   |                   |                    |                             |            |  |
|-------------------------------------------|---------------|---------------|------------------|-------------------|-------------------|--------------------|-----------------------------|------------|--|
| Inter-Fabric Connections                  |               |               |                  |                   |                   |                    |                             |            |  |
| Type                                      | Source Fabric | Source Device | Source Interface | Destination Fa... | Destination De... | Destination Int... | Configuration               | Status     |  |
| <input checked="" type="radio"/> VRF_LITE | 9K-FABRIC     | BL-1          | Ethernet1/30     | ext-fb1           | N7k1-ER-1         | Ethernet5/41       | <a href="#">View Config</a> | DEPLOYMENT |  |
| <input type="radio"/> VRF_LITE            | 9K-FABRIC     | BL-2          | Ethernet1/30     | ext-fb1           | N7k1-ER-1         | Ethernet5/42       | <a href="#">View Config</a> | DEPLOYED   |  |

The next screen shows that the *BL-1* connection to *ER-1* is removed from the fabric extension list.

| Fabric Extension               |               |               |                  |                   |                   |                    |                             |          |  |
|--------------------------------|---------------|---------------|------------------|-------------------|-------------------|--------------------|-----------------------------|----------|--|
| Inter-Fabric Connections       |               |               |                  |                   |                   |                    |                             |          |  |
| Type                           | Source Fabric | Source Device | Source Interface | Destination Fa... | Destination De... | Destination Int... | Configuration               | Status   |  |
| <input type="radio"/> VRF_LITE | 9K-FABRIC     | BL-2          | Ethernet1/30     | ext-fb1           | N7k1-ER-1         | Ethernet5/42       | <a href="#">View Config</a> | DEPLOYED |  |

5. Similarly, select *BL-2* and click **X** to remove the *BL-2* connection to *ER-1*. After both *BL-1* and *BL-2* VRF Lite connections are removed, the Fabric Extension screen will have no entries.



## Additional References

| Document Title and Link                                                           | Document Description                                         |
|-----------------------------------------------------------------------------------|--------------------------------------------------------------|
| <a href="#">Cisco Programmable Fabric with VXLAN BGP EVPN Configuration Guide</a> | This document explains external connectivity using VRF Lite. |

## Appendix

### Edge Router Configurations

**ER-1 Configuration Example**—The following configurations are enabled on *ER-1* to connect to *BL-1* and *BL-2* (border leafs), and reproduced here for reference.


**Note**

*switch(config)#* refers to the global configuration mode. To access this mode, type the following on your switch: **switch# configure terminal**.

```

switch(config)#
interface Ethernet5/41    ## ER-1 interface to BL-1
  ip address 35.1.1.11/24
  no shutdown

interface Ethernet5/42    ## ER-1 interface to BL-2
  ip address 36.1.1.11/24
  no shutdown

router bgp 3000          ## eBGP sessions
  neighbor 35.1.1.10 remote-as 2000  ###Peering to BL-1 (eBGP)
    update-source Ethernet5/41
    address-family ipv4 unicast
      next-hop-self
  neighbor 36.1.1.10 remote-as 2000  ###Peering to BL-2 (eBGP)
    update-source Ethernet5/42

```

## Edge Router Configurations

```
address-family ipv4 unicast
  next-hop-self
```

The following configurations are manually enabled on *ER-1* for VRF extension to the border leafs:

```
configure profile 9K-FABRIC-Default_VRF_Extension-50016
  vrf context MyVRF_50016
    address-family ipv4 unicast
      route-target import 3000:3
      route-target export 3000:3
      rd 3000:3
    interface Ethernet5/41.3
      encapsulation dot1Q 3
      vrf member MyVRF_50016
      ip address 35.1.1.11/24
      ipv6 address 35:1:1:1::2/64
      no shutdown
    interface Ethernet5/42.3
      encapsulation dot1Q 3
      vrf member MyVRF_50016
      ip address 36.1.1.11/24
      ipv6 address 36:1:1:1::2/64
      no shutdown
    router bgp 3000
      vrf MyVRF_50016
        address-family ipv4 unicast
          maximum-paths ibgp 2
        neighbor 35.1.1.10 remote-as 2000
          address-family ipv4 unicast
            send-community both
        neighbor 36.1.1.10 remote-as 2000
          address-family ipv4 unicast
            send-community both
```

### Configurations Pushed to *BL-1* Through DCNM:

VRF extension pushed to *BL-1* through DCNM

```
### Route map
ip prefix-list default-route seq 5 permit 0.0.0.0/0 le 1
ip prefix-list host-route seq 5 permit 0.0.0.0/0 eq 32
route-map EXTCON-RMAP-FILTER deny 10
  match ip address prefix-list default-route
route-map EXTCON-RMAP-FILTER deny 20
  match ip address prefix-list host-route
route-map EXTCON-RMAP-FILTER permit 1000

ipv6 prefix-list default-route-v6 seq 5 permit 0::/0
ipv6 prefix-list host-route-v6 seq 5 permit 0::/0 eq 128
route-map EXTCON-RMAP-FILTER-V6 deny 10
  match ipv6 address prefix-list default-route-v6
route-map EXTCON-RMAP-FILTER-V6 deny 20
  match ip address prefix-list host-route-v6
route-map EXTCON-RMAP-FILTER-V6 permit 1000

### VRF-Lite interface of BL-1
interface Etherenet1/30
  no switchport
  ip address 35.1.1.10/24
  no shutdown
```

```
### External BGP (eBGP) session of BL-1
router bgp 2000
  address-family ipv4 unicast
    redistribute direct route-map RMAP-REDIST-DIRECT
    neighbor 35.1.1.11 remote-as 3000
    update-source Ethernet1/30
  address-family ipv4 unicast
    next-hop-self
```

The following configuration profile is pushed through DCNM when *MyVRF\_50016* is deployed on *BL-1*:

```
configure profile 9K-FABRIC-Default_VRF_Extension-50016
vlan 2001
  vn-segment 50016
  interface vlan 2001
    vrf member MyVRF_50016
    ip forward
    ipv6 forward
    no ip redirects
    no ipv6 redirects
    mtu 9216
    no shutdown

  interface nve 1
    member vni 50016 associate-vrf

  vrf context MyVRF_50016
    vni 50016
    rd auto
      address-family ipv4 unicast
        route-target both auto
        route-target both auto evpn
        ip route 0/0 35.1.1.11
      address-family ipv6 unicast
        route-target both auto
        route-target both auto evpn
        ipv6 route 0::/0 35.1.1.1.2

router bgp 2000
  vrf MyVRF_50016 ## bgp VRF configured
  address-family ipv4 unicast
    advertise l2vpn evpn
    redistribute direct route-map FABRIC-RMAP-REDIST-SUBNET
    maximum-paths ibgp 2
    network 0/0
  address-family ipv6 unicast
    advertise l2vpn evpn
    redistribute direct route-map FABRIC-RMAP-REDIST-SUBNET
    maximum-paths ibgp 2
    network 0::/0
  neighbor 35.1.1.11 remote-as 3000
    address-family ipv4 unicast
    send-community both
    route-map EXTCON-RMAP-FILTER out
  neighbor 35.1.1.1.2 remote-as 3000
    address-family ipv6 unicast
    send-community both
    route-map EXTCON-RMAP-FILTER-V6 out

  interface Ethernet1/30.3 #sub interface member of VRF deployed
  encapsulation dot1q 3
  vrf member MyVRF_50016
  ip address 35.1.1.10/24
```

**Edge Router Configurations**

```
ipv6 address 35:1:1:1::1/64
no shutdown

configure terminal
  apply profile 9K-FABRIC-Default_VRF_Extension-50016
```