



## **Cisco DCNM LAN Fabric Configuration Guide, Release 11.2(1)**

**First Published:** 2019-06-06

**Last Modified:** 2020-04-21

### **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 NON-INFRINGEMENT 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

---

|                  |                                   |          |
|------------------|-----------------------------------|----------|
| <b>CHAPTER 1</b> | <b>Overview</b>                   | <b>1</b> |
|                  | Cisco Data Center Network Manager | 1        |

---

|                  |                  |          |
|------------------|------------------|----------|
| <b>CHAPTER 2</b> | <b>Dashboard</b> | <b>3</b> |
|                  | Dashboard        | 3        |
|                  | Dashlets         | 4        |

---

|                  |                                |          |
|------------------|--------------------------------|----------|
| <b>CHAPTER 3</b> | <b>Topology</b>                | <b>9</b> |
|                  | Topology                       | 9        |
|                  | Status                         | 9        |
|                  | Scope                          | 10       |
|                  | Searching                      | 10       |
|                  | Quick Search                   | 11       |
|                  | Host name (vCenter)            | 11       |
|                  | Host IP                        | 11       |
|                  | Host MAC                       | 11       |
|                  | Multicast Group                | 11       |
|                  | VXLAN ID (VNI)                 | 11       |
|                  | VLAN                           | 12       |
|                  | VXLAN OAM                      | 12       |
|                  | Show Panel                     | 13       |
|                  | Layouts                        | 14       |
|                  | Zooming, Panning, and Dragging | 14       |
|                  | Switch Slide-Out Panel         | 16       |
|                  | Beacon                         | 16       |
|                  | Tagging                        | 16       |

- More Details 16
- Link Slide-Out Panel 17
  - 24-Hour Traffic 17
- vCenter Compute Visualization 17
  - Enabling vCenter Compute Visualization 18
  - Using vCenter Compute Visualization 20
  - Troubleshooting vCenter Compute Visualization 24

**CHAPTER 4**

**Control 27**

- Fabrics 27
  - VXLAN BGP EVPN Fabrics Provisioning 28
  - Creating a New VXLAN BGP EVPN Fabric 31
    - Adding Switches to a Fabric 45
    - Pre-provisioning a Device 59
    - Changing the TCAM Configuration on a Device 62
    - Adding a vPC L3 Peer Keep-Alive Link 63
  - Brownfield Deployment-Transitioning VXLAN Fabric Management to DCNM 66
  - Creating a New Fabric for EBGp-Based Underlay 67
    - Applying Policies On A Fabric With An eBGP Underlay 76
    - Deploying Fabric Underlay eBGP Policies 77
    - Deploying Fabric Overlay eBGP Policies 78
    - Deploying Spine Switch Overlay Policies 78
    - Deploying Leaf Switch Overlay Policies 79
    - Guidelines 80
  - Creating an External Fabric 80
    - Discovering New Switches 90
    - Pre-provisioning a Device 95
    - Creating a vPC Setup in the External Fabric 98
    - Undeploying a vPC Setup in the External Fabric 103
  - Multi-Site Domain for VXLAN BGP EVPN Fabrics 103
    - Removing a Fabric From an MSD 133
    - Moving a Standalone Fabric (With Existing Networks and VRFs) to an MSD Fabric 133
- SSH Key RSA Handling 134
- Switch Operations 135

|  |     |
|--|-----|
| Fabric Multi Switch Operations                         | 138 |
| Fabric Links   | 138 |
| Creating Intra-Fabric Links                            | 139 |
| Creating Inter-Fabric Links                            | 144 |
| Exporting Links  | 147 |
| Importing Links  | 148 |
| Viewing Details of Fabric Links                        | 148 |
| Viewing the Traffic Details of Fabric Links            | 149 |
| vPC Fabric Peering                                     | 150 |
| Creating a Virtual Peer Link                           | 151 |
| Converting a Physical Peer Link to a Virtual Peer Link | 154 |
| Converting a Virtual Peer Link to a Physical Peer Link | 155 |
| Viewing and Editing Policies                           | 156 |
| Viewing Policies                                       | 157 |
| Adding a Policy  | 157 |
| Deploying Policies                                     | 158 |
| Editing a Policy                                       | 158 |
| Current Switch Configuration                           | 159 |
| Retrieving the Authentication Key                      | 160 |
| Return Material Authorization (RMA)                    | 161 |
| Prerequisites  | 161 |
| Guidelines and Limitations                             | 162 |
| POAP RMA Flow  | 162 |
| Manual RMA Flow  | 164 |
| RMA for User with Local Authentication                 | 167 |
| Interfaces   | 167 |
| Adding Interfaces                                      | 170 |
| Editing Interfaces                                     | 171 |
| Deleting Interfaces                                    | 173 |
| Shutting Down and Bringing Up Interfaces               | 173 |
| Viewing Interface Configuration                        | 174 |
| Rediscovering Interfaces                               | 174 |
| Viewing Interface History                              | 174 |
| Deploying Interface Configurations                     | 175 |

|   |     |
|---|-----|
| Creating External Fabric Interfaces                                     | 175 |
| Creating and Deploying Networks and VRFs                                | 176 |
| Viewing Networks and VRFs for a Fabric                                  | 177 |
| Creating Networks for the Standalone Fabric                             | 177 |
| Editing Networks for the Standalone Fabric                              | 182 |
| Creating VRFs for the Standalone Fabric                                 | 182 |
| Editing VRFs for the Standalone Fabric                                  | 187 |
| Deploying Networks for the Standalone and MSD Fabrics                   | 188 |
| Deploying VRFs for the Standalone and MSD Fabrics                       | 197 |
| Undeploying Networks for the Standalone Fabric                          | 203 |
| Undeploying VRFs for the Standalone Fabric                              | 204 |
| Deleting Networks and VRFs  | 204 |
| Configuring Multiple VLAN IDs to a Single VNI                           | 204 |
| Fabric Backup and Restore   | 206 |
| Backing Up Fabrics  | 206 |
| Restoring Fabrics   | 208 |
| Deleting a VXLAN BGP EVPN Fabric  | 212 |
| Post DCNM 11.2(1) Upgrade for VXLAN BGP EVPN, External, and MSD Fabrics | 212 |
| Changing ISIS Configuration from Level 1 to Level 2                     | 213 |
| Configuration Compliance in DCNM  | 214 |
| Configuration Compliance in External Fabrics                            | 221 |
| Resolving Diffs for Case Insensitive Commands                           | 226 |
| Enabling Freeform Configurations on Fabric Switches                     | 227 |
| Management  | 231 |
| Resources   | 231 |
| Adding, Editing, Re-Discovering and Removing VMware Servers             | 231 |
| Adding a Virtual Center Server  | 231 |
| Deleting a VMware Server  | 232 |
| Editing a VMware Server   | 232 |
| Rediscovering a VMware Server   | 232 |
| Template Library  | 233 |
| Template Structure  | 234 |
| Template Format   | 234 |
| Template Variables  | 241 |

|  |     |
|--|-----|
| Variable Meta Property                             | 243 |
| Variable Annotation                                | 249 |
| Templates Content                                  | 253 |
| Advanced Features                                  | 255 |
| Adding a Template                                  | 257 |
| Modifying a Template                               | 258 |
| Copying a Template                                 | 259 |
| Deleting a Template                                | 260 |
| Importing a Template                               | 260 |
| Exporting a Template                               | 260 |
| Image Management                                   | 261 |
| 261  |     |
| Deleting an Image                                  | 261 |
| Image Upload                                       | 261 |
| Install & Upgrade                                  | 262 |
| Upgrade History                                    | 262 |
| Switch Level History                               | 268 |
| Endpoint Locator                                   | 269 |
| Endpoint Locator                                   | 269 |
| Configuring Endpoint Locator                       | 270 |
| Configuring Endpoint Locator in DCNM Cluster Mode  | 282 |
| Configuring Endpoint Locator for External Fabrics  | 284 |
| Configuring Endpoint Locator for eBGP EVPN Fabrics | 284 |
| EPL Connectivity Options                           | 287 |
| Disabling Endpoint Locator                         | 291 |
| Troubleshooting Endpoint Locator                   | 291 |
| Monitoring Endpoint Locator                        | 293 |
| Endpoint Locator Dashboard                         | 293 |
| LAN Telemetry Health                               | 297 |
| Health   | 297 |
| Top Streamers                                      | 304 |

---

**CHAPTER 5**
**Monitor 307**

## Inventory 307

- Viewing Inventory Information for Switches **307**
  - Viewing System Information **309**
    - Hosts **309**
    - Capacity **310**
    - Features **310**
    - VXLAN **310**
    - VLAN **311**
    - Switch Modules **312**
    - FEX **312**
    - VDCs **315**
    - Switch On-Board Analytics **322**
  - Viewing Inventory Information for Modules **326**
  - Viewing Inventory Information for Licenses **327**
- Monitoring Switch **328**
  - Viewing Switch CPU Information **328**
  - Viewing Switch Memory Information **328**
  - Viewing Switch Traffic and Errors Information **329**
  - Viewing Switch Temperature **329**
    - Enabling Temperature Monitoring **330**
  - Viewing Accounting Information **330**
  - Viewing Events Information **330**
- Monitoring LAN **331**
  - Monitoring Performance Information for Ethernet **331**
  - Monitoring ISL Traffic and Errors **332**
  - Monitoring a vPC **333**
    - Monitoring vPC Performance **334**
- Monitoring Endpoint Locator **335**
  - Exploring Endpoint Locator Details **335**
- Alarms **343**
  - Viewing Alarms and Events **343**
  - Monitoring and Adding Alarm Policies **344**
    - Activating Policies **347**
    - Deactivating Policies **347**
    - Importing Policies **347**

|                       |     |
|-----------------------|-----|
| Exporting Policies    | 347 |
| Editing Policies      | 348 |
| Deleting Policies     | 348 |
| Health Monitor Alarms | 348 |

---

**CHAPTER 6****Administration 351**

|   |     |
|---|-----|
| DCNM Server                                 | 351 |
| Starting, Restarting, and Stopping Services | 351 |
| Viewing Log Information                     | 352 |
| Server Properties                           | 353 |
| Modular Device Support                      | 353 |
| Managing Licenses                           | 354 |
| License Assignments                         | 355 |
| Smart License                               | 356 |
| Server License Files                        | 359 |
| Native HA                                   | 360 |
| Multi Site Manager                          | 361 |
| Device Connector                            | 364 |
| Management Users                            | 367 |
| Remote AAA                                  | 367 |
| Local                                       | 368 |
| Radius                                      | 368 |
| TACACS+                                     | 369 |
| Switch                                      | 369 |
| LDAP  | 369 |
| Managing Local Users                        | 372 |
| Adding Local Users                          | 372 |
| Deleting Local Users                        | 372 |
| Editing a User                              | 373 |
| User Access                                 | 373 |
| Managing Clients                            | 374 |
| Performance Setup                           | 374 |
| Performance Setup LAN Collections           | 375 |
| Event Setup                                 | 375 |

|   |                         |
|---|-------------------------|
| Viewing Events Registration                           | 375                     |
| Notification Forwarding                               | 376                     |
| Adding Notification Forwarding                        | 376                     |
| Removing Notification Forwarding                      | 378                     |
| Event Suppression                                     | 378                     |
| Add Event Suppression Rules                           | 378                     |
| Delete Event Suppression Rule                         | 379                     |
| Modify Event Suppression Rule                         | 379                     |
| Credentials Management                                | 379                     |
| LAN Credentials                                       | 380                     |
| <hr/>   |                         |
| <b>CHAPTER 7</b>                                      | <b>Applications 383</b> |
| Cisco DCNM in Unclustered Mode                        | 383                     |
| Cisco DCNM in Clustered Mode                          | 384                     |
| Requirements for Cisco DCNM Clustered Mode            | 384                     |
| Installing a Cisco DCNM Compute                       | 386                     |
| Networking Policies for OVA Installation              | 386                     |
| Enabling the Compute Cluster                          | 388                     |
| Managing Application Network Pools                    | 390                     |
| Adding Computes into the Cluster Mode                 | 390                     |
| Transitioning Compute Nodes                           | 392                     |
| Transitioning Compute nodes from VM to Service Engine | 392                     |
| Transitioning Compute nodes from Service Engine to VM | 393                     |
| Preferences   | 394                     |
| In-Band Telemetry Network and NTP Configuration       | 394                     |
| Telemetry Network and NTP Requirements                | 400                     |
| Installing and Deploying Applications                 | 401                     |
| Application Framework User Interface                  | 406                     |
| Catalog   | 407                     |
| Watch Tower   | 408                     |
| Alerts  | 409                     |
| Service Utilization                                   | 409                     |
| Compute Utilization                                   | 410                     |
| Compute   | 410                     |

|   |     |
|---|-----|
| Preferences                                     | 411 |
| In-Band Telemetry Network and NTP Configuration | 412 |
| Failure Scenario                                | 418 |
| Compute Node Disaster Recovery                  | 418 |

**CHAPTER 8**

|   |            |
|---|------------|
| <b>Connecting Cisco Data Center and a Public Cloud</b>                                | <b>419</b> |
| Connecting Cisco Data Center and a Public Cloud                                       | 419        |
| Topology Overview   | 420        |
| Guidelines and Limitations  | 421        |
| Prerequisites   | 421        |
| Task Summary  | 421        |
| Enabling the Preview Functionality  | 422        |
| Setting Up the On-premise External Fabric with CSR 1000v                              | 423        |
| Creating an External Fabric   | 423        |
| Discovering the On-Premises Core Router   | 424        |
| Setting Up the VXLAN EVPN Fabric  | 425        |
| Creating a VXLAN EVPN Fabric  | 425        |
| Assigning the BGW Role  | 425        |
| Setting Up the External Fabric with CSR in Azure                                      | 426        |
| Creating an External Fabric   | 426        |
| Discovering the Core Router   | 426        |
| Setting Up the MSD Fabric for Connectivity  | 427        |
| Creating an MSD Fabric  | 428        |
| Moving Other Fabrics into the MSD Fabric  | 428        |
| Setting Up Connections  | 429        |
| Connecting the On-Premises BGW and the On-Premises Core Router                        | 429        |
| Connecting the On-prem Core Router and the Public-cloud Core Router with IPsec Tunnel | 431        |
| Connecting the On-prem BGW and the Public-cloud Core Router using EVPN Peering        | 433        |
| Saving and Deploying Configurations   | 434        |
| Extending VRFs  | 435        |
| Deploying and Extending the VRF On-prem Core Router                                   | 436        |
| Creating and Deploying VRF on Public Cloud  | 437        |
| Configuring Default Gateway for the VM  | 438        |
| Verifying the Connectivity  | 439        |

Deploying Cisco CSR 1000v on Microsoft Azure 439

Viewing Links and Core Routers Details 443

Resetting Packet Counter Using API 443

---

**CHAPTER 9**

**Managing a Brownfield VXLAN BGP EVPN Fabric 445**

Overview 445

Prerequisites 446

Guidelines and Limitations 446

Fabric Topology Overview 448

DCNM Brownfield Deployment Tasks 449

Verifying the Existing VXLAN BGP EVPN Fabric 449

Creating a VXLAN BGP EVPN Fabric 452

Adding Switches and Transitioning VXLAN Fabric Management to DCNM 460

Verifying the Import of the VXLAN BGP EVPN Fabric 473

Verifying VXLANs and Commands on Switches 473

Verifying Resources 476

Verifying Networks 477

Migrating a Bottom-Up VXLAN Fabric to DCNM 480

Resolving Config Compliance Error on Switches with Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) Images 488

Modifying VLAN Names in a Switch with Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) Images 492

Changing a Brownfield Imported BIDIR Configuration 494

Manually Adding PIM-BIDIR Configuration for Leaf or Spine Post Brownfield Migration 495

Migrating an MSD Fabric with Border Gateway Switches 495

---

**CHAPTER 10**

**Template Usage in Cisco DCNM LAN Fabric Deployment 499**

Policy Template 499

Fabric Template 503

Profile Template 503

Viewing, Editing, and Adding Policies 504

Viewing Policies 505

Editing Policies 507

Adding Policies 508

|  |     |
|--|-----|
| Deploying New Configurations               | 508 |
| switch_freeform Template Usage             | 509 |
| Example: Create a switch_freeform policy   | 509 |
| Changing the Contents of a Template in Use | 512 |

**CHAPTER 11****Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - VRF Lite 515**

|  |     |
|--|-----|
| Prerequisites  | 515 |
| Sample Scenarios   | 518 |
| VRF Lite Through the DCNM GUI – From a BGW Device to a Nexus 7000 Series Edge Router | 519 |
| VRF Lite Through the DCNM GUI – From a BGW Device To a Non-Nexus Device              | 531 |
| Automatic VRF Lite (IFC) Configuration   | 538 |
| Deleting VRF Lite IFCs   | 541 |
| Additional References  | 543 |
| Appendix   | 543 |
| N9K-3-BGW Configurations   | 543 |

**CHAPTER 12****Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - Multi-Site 547**

|  |     |
|--|-----|
| Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - Multi-Site        | 547 |
| Prerequisites  | 548 |
| Limitations  | 549 |
| Save & Deploy Operation in the MSD Fabric                                  | 549 |
| EVPN Multi-Site Configuration  | 551 |
| Configuring Multi-Site Underlay IFCs - DCNM GUI                            | 552 |
| Configuring Multi-Site Underlay IFCs - Autoconfiguration                   | 553 |
| Configuring Multi-Site Underlay IFCs Towards a Non-Nexus Device - DCNM GUI | 554 |
| Configuring Multi-Site Overlay IFCs  | 556 |
| Configuring Multi-Site Overlay IFCs - Autoconfiguration                    | 558 |
| Configuring Multi-Site Overlay IFCs Towards a Non-Nexus Device - DCNM GUI  | 559 |
| Overlay and Underlay Peering Configurations on the Route Server N7k1-RS1   | 561 |
| Viewing, Editing and Deleting Multi-Site Overlays                          | 561 |
| Deleting Multi-Site IFCs   | 561 |
| Creating and Deploying Networks and VRFs in the MSD Fabric                 | 562 |
| Deploying a Legacy Site BGW (vPC-BGWs)                                     | 566 |
| Additional References  | 570 |

Appendix **570**

- Multi-Site Fabric Base Configurations – Box Topology **570**
- IBGP Configuration for the Box Topology in the Easy7200 Fabric **571**
- Route Server Configuration **572**
  - Multi-Site Overlay IFC Configuration **573**
  - Multi-Site Underlay IFC Configuration – Out-of-Box Profiles **573**



# CHAPTER 1

## Overview

---

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

## Cisco Data Center Network Manager

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 capabilities, such as, control, automation, monitoring, visualization, and troubleshooting.

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 functionality for Cisco DCNM LAN Fabric deployment.

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:** Click on the gear icon to see a drop-down list with the following options:
  - **Logged in as:** displays the user role of the current logged in user.
  - **Change Password:** Allows you to change the password for current logged in user.  
If you are a **network administrator** user, you can modify the passwords of the other users.
  - **About:** Displays the Version, Installation Type, and time since when the Web UI is operational.
  - **Logout:** Allows you to terminate the Web UI and returns to the login screen.

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:

- Data Center
- Inventory - Switches
- Inventory - Modules
- Top CPU
- Top ISLs/Trunks
- Link Traffic
- Events

- Server Status
- Audit Log

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  | 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 click detach option, the map opens in a new tab and can be configured. <ul style="list-style-type: none"> <li>• The network map dialog box has properties that are different from the Summary dashboard view:</li> </ul> |

| Dashlet         | Description  |
|-----------------|--|
|                 | <ul style="list-style-type: none"> <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> <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   | <p>Displays the status of DCNM and federation servers, and the health check status for the components.</p> <p>The following services, server, and status details are displayed under the <b>DCNM</b> tab.</p> <ul style="list-style-type: none"> <li>• Database Server</li> <li>• Search Indexer</li> <li>• Performance Collector</li> <li>• NTPD Server</li> <li>• DHCP Server</li> <li>• SNMP Traps</li> <li>• Syslog Server</li> </ul> <p>The following component status and details are displayed under the <b>Health Check</b> tab.</p> <ul style="list-style-type: none"> <li>• AMQP Server</li> <li>• DHCP Server</li> <li>• TFTP Server</li> <li>• EPLS</li> <li>• EPLC</li> </ul> |
| Top ISLs/Trunks | <p>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</p>  |

| Dashlet                      | Description  |
|------------------------------|--|
|                              | 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> |
| Errors                       | Displays the error packets for the selected interface. This information is retrieved from the <b>Errors &gt; In-Peak</b> and <b>Errors &gt; Out-Peak</b> columns of the <b>Monitor &gt; LAN / Ethernet</b> page.   |
| Discards                     | <p>Displays the error packets that are discarded for the selected interface.</p> <p><b>Note</b> The Discards dashlet is only for LAN.</p>  |
| 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.   |

| Dashlet                   | Description   |
|---------------------------|---|
| 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 9](#)

## 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**.)

- After moving a cable from one port to another port, the old fabric link is retained in the **Topology** window, and it is shown in the red color indicating that the link is down. The port movements are not updated in the **Topology** window. You need to rediscover the switch for the updated ports to be displayed in DCNM.

- 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.

## Pod Name (Container)

You can also click on the Pod List to view the information regarding all the pods running on the selected Cluster. If Cluster Selection is All, all the pods running on all the clusters in your topology is displayed. You can also export the Pod List data for further analysis.

## 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.
- **Compute:** Check the check box to enable the compute visibility 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.
- **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.

In VXLAN (standalone, MSD, and MSD member) fabrics and external fabrics, discovered links or connections (via CDP) to non-DCNM managed switches are represented by a cloud labelled **Undiscovered**.

### Undiscovered Cloud Display

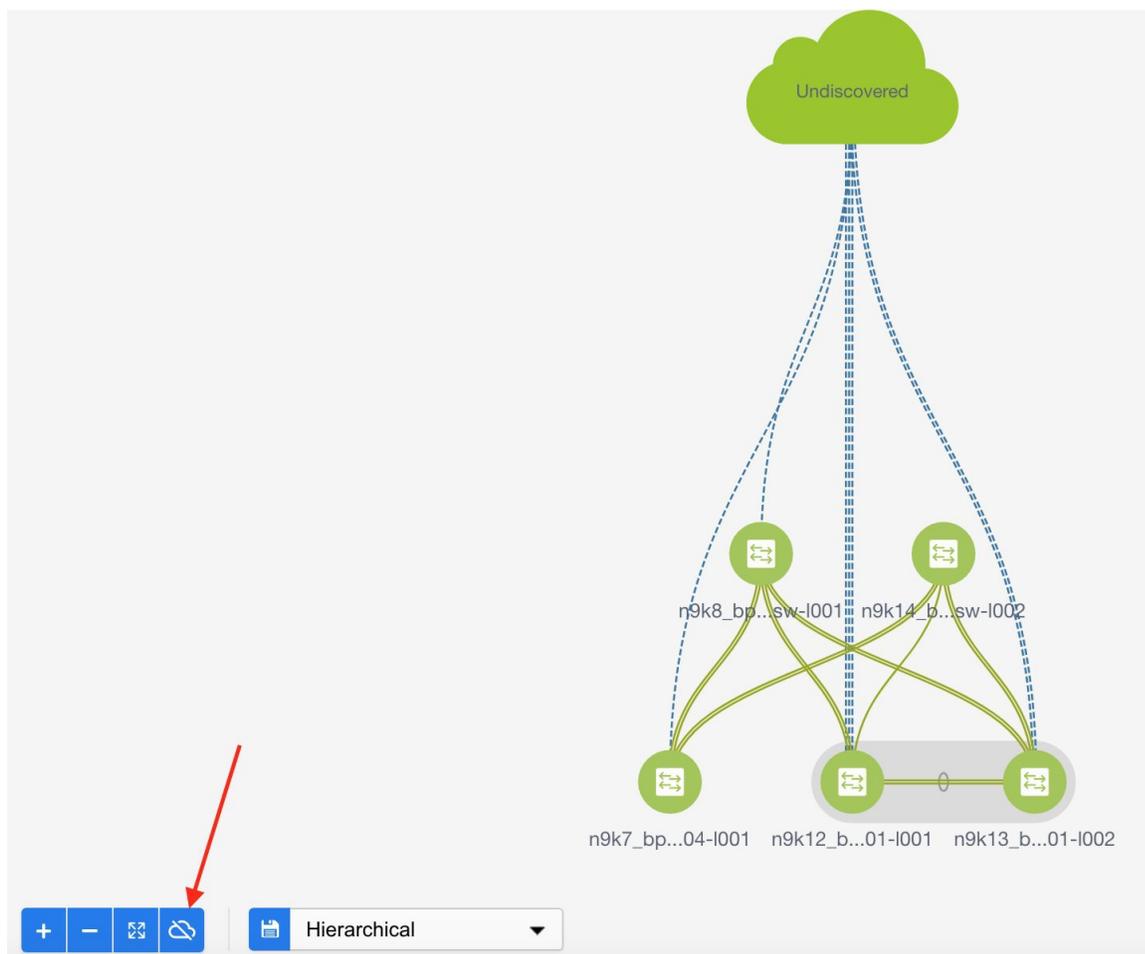
In the **Topology** screen, you can see an **Undiscovered** cloud at the top part of the image.



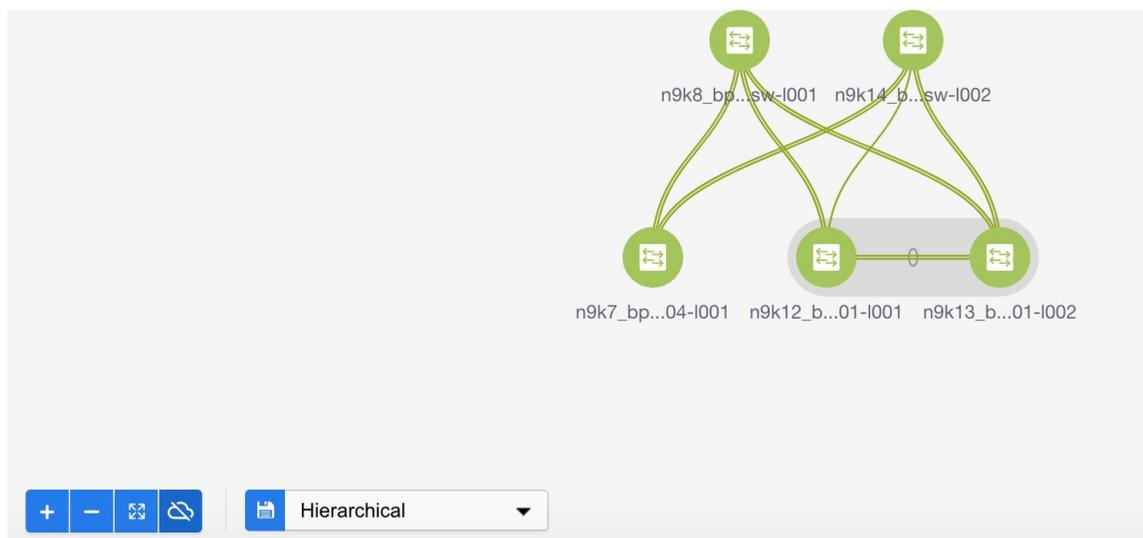

---

**Note** The Undiscovered cloud is hidden by default. You can display the Undiscovered cloud by clicking the Cloud icon (at the bottom left part of the screen).

---



Click again to stop the **Undiscovered** cloud from being displayed. You can see that the **Undiscovered** cloud and its links to the fabric devices are not displayed.



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

## 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-pollled CPU utilization, and last-pollled 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.

VMM supports computes connecting to a border spine. Border Spine is a new switch role managed by easy fabric in Cisco DCNM 11.1(1).

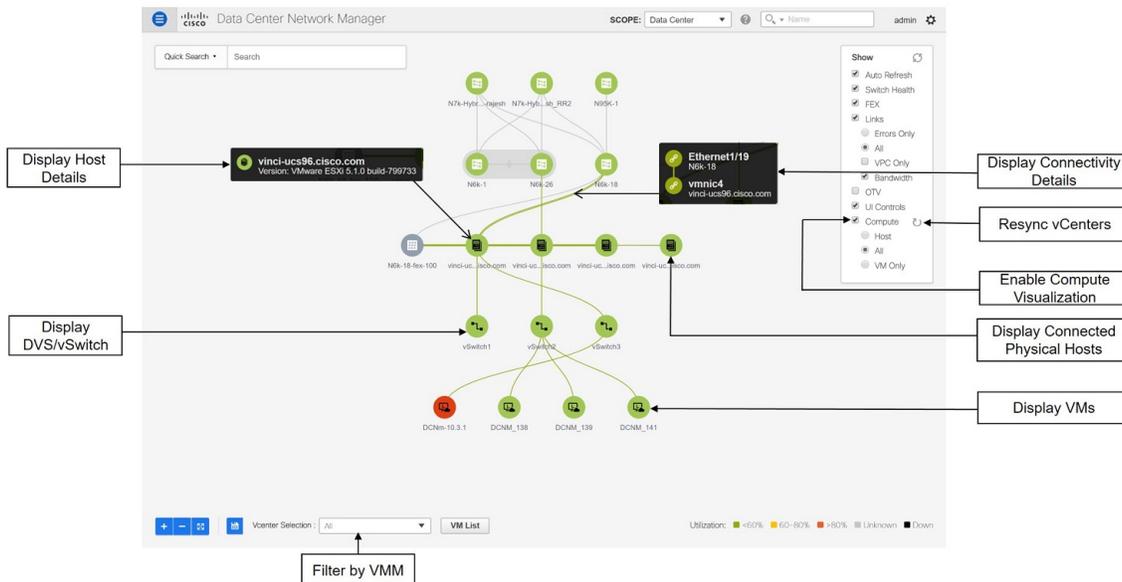


---

**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>.
  - Cisco DCNM doesnot support non-Cisco blade servers.
-

Figure 1: vCenter Compute Visualization



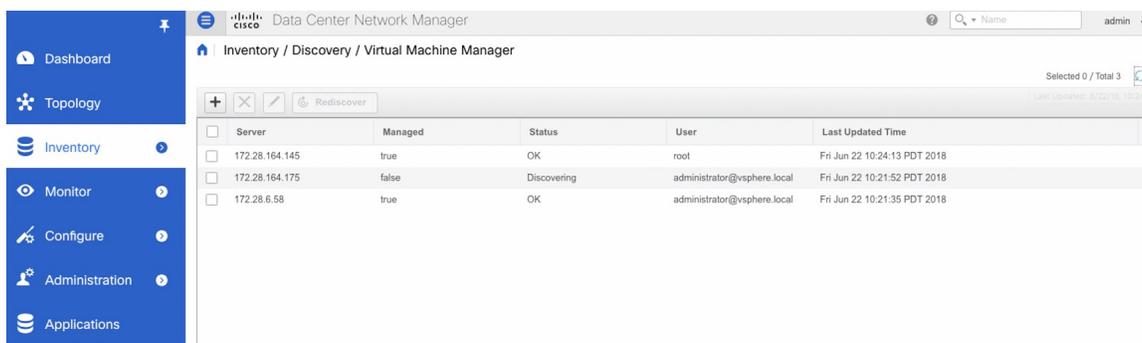
## Enabling vCenter Compute Visualization

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

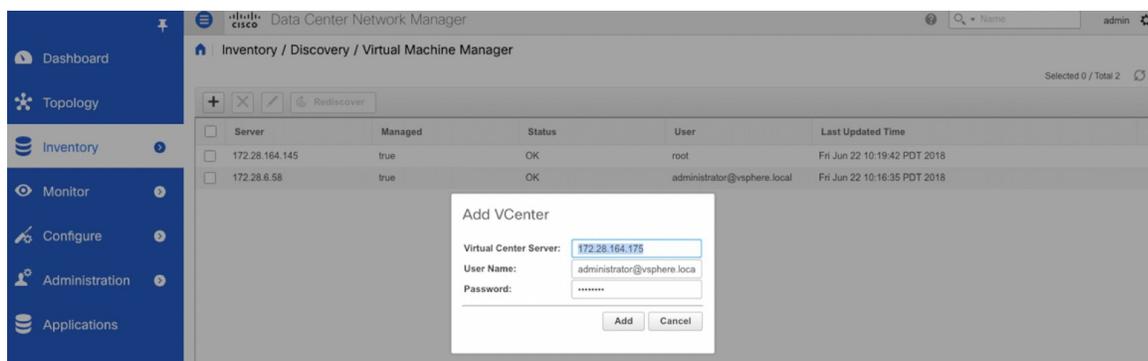
### Procedure

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

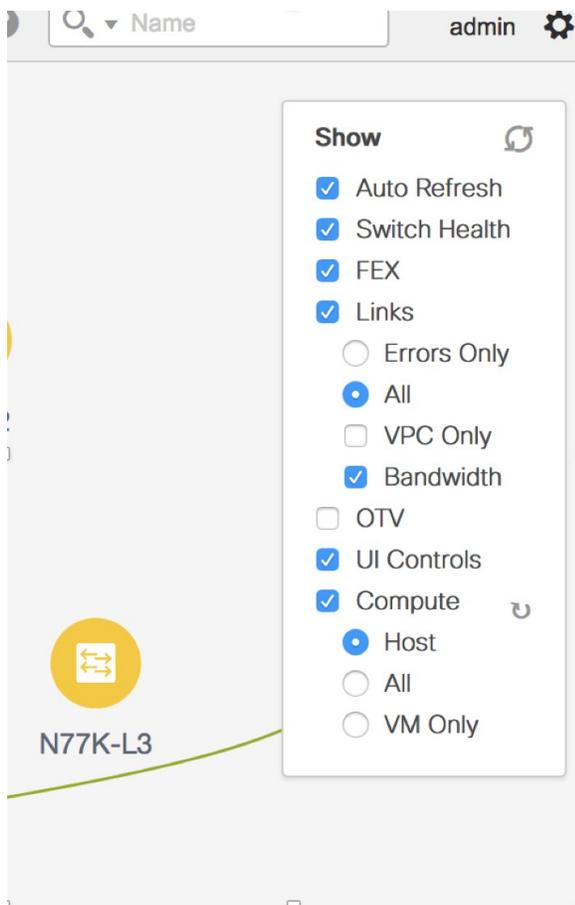
The **Control > Management > Virtual Machine Manager** window appears.



**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.



**Note** When you add a vCenter, you can run into a situation where an image upload is in progress, and thus the compute visualization is not complete. The **Topology** window displays the following message for more than 10 minutes:

"Compute visualization data fetch in progress - Please give some time."

Navigate to the **Applications** window, and confirm that the VMM application is not running. If it is running, indicated by the green or orange dot on the top left corner of the application icon, the problem is caused by a different scenario. Otherwise, delete the vCenter, wait for around 15 minutes, and re-add it. Verify the application status and continue with your DCNM tasks.

## Using vCenter Compute Visualization

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

### Procedure

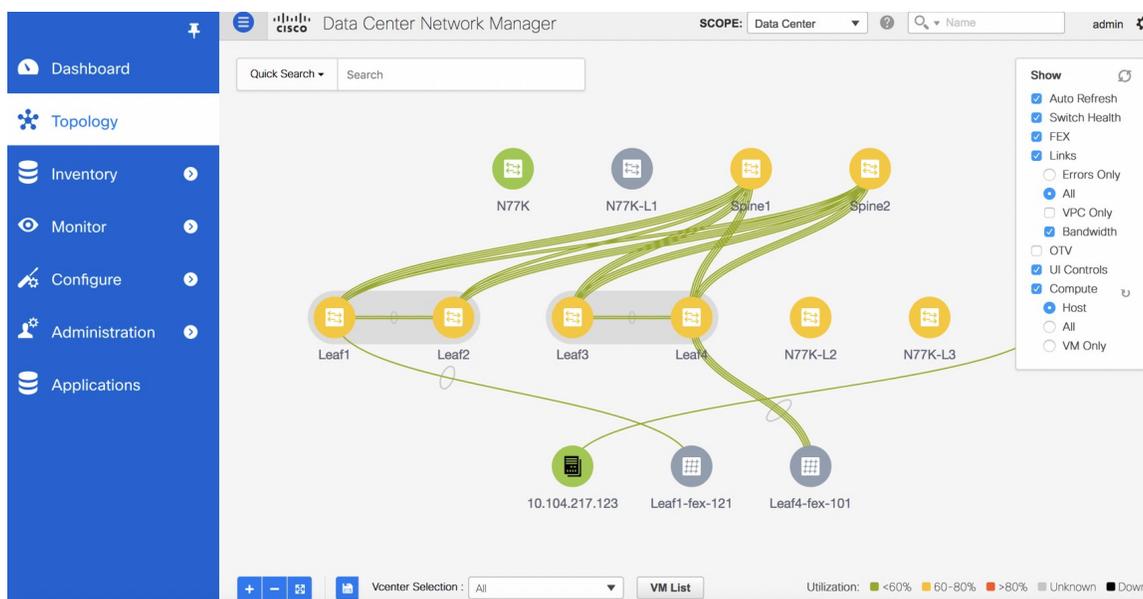
**Step 1** Choose **Topology**.

**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.

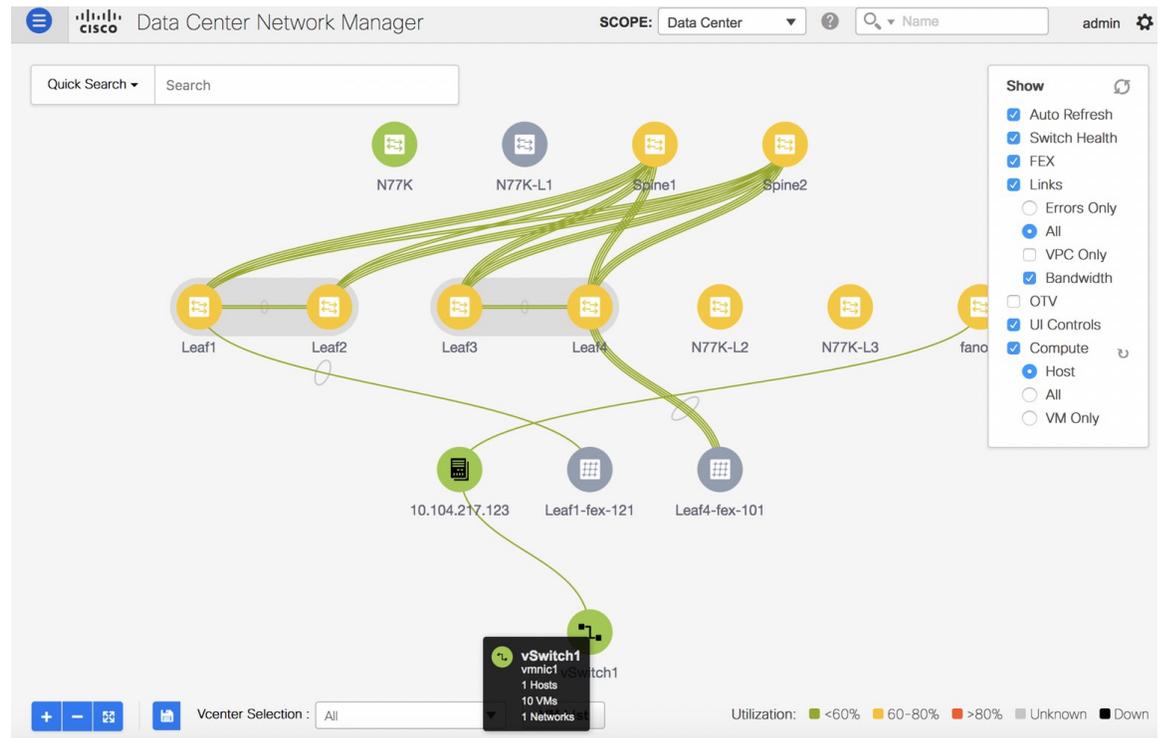
- **Host**
- **All**
- **VM Only**



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.

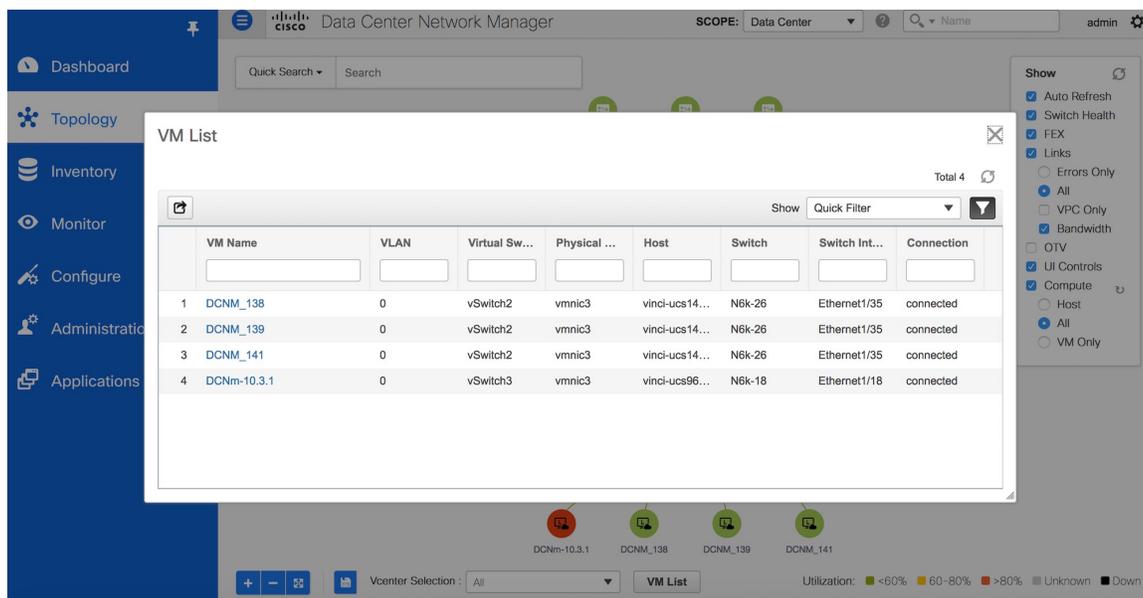
**Note** Ensure that you select the **FEX** checkbox when you are viewing Compute nodes. The Hosts or VMs behind the FEX will be dangling, otherwise.

## Using the Virtual Machine List

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

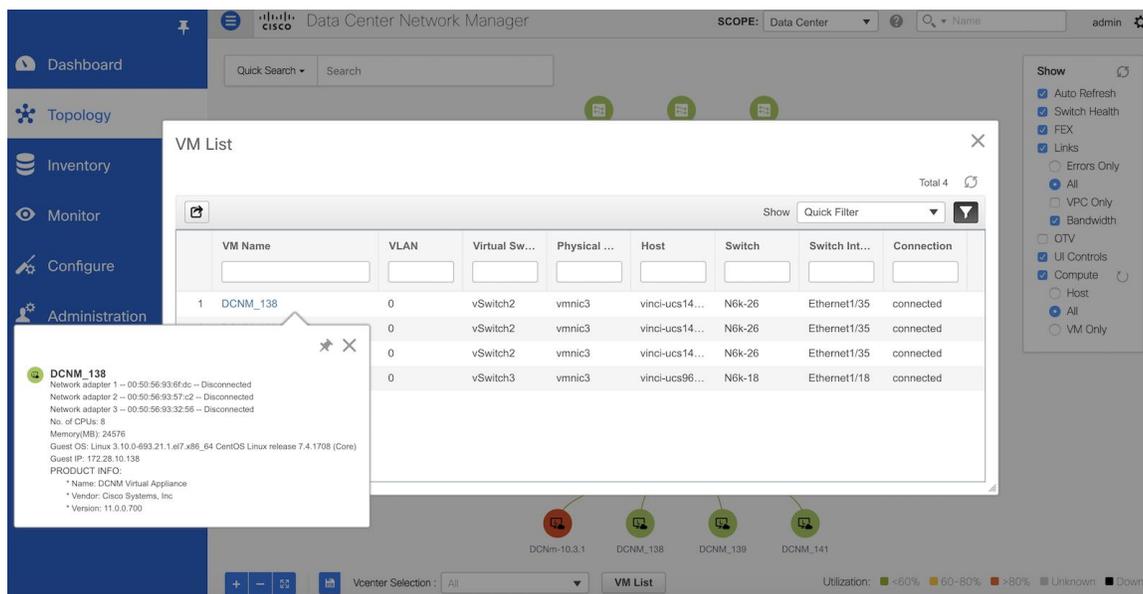
### Procedure

- Step 1** Choose **Topology**.
- Step 2** Click **VM List**.



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.

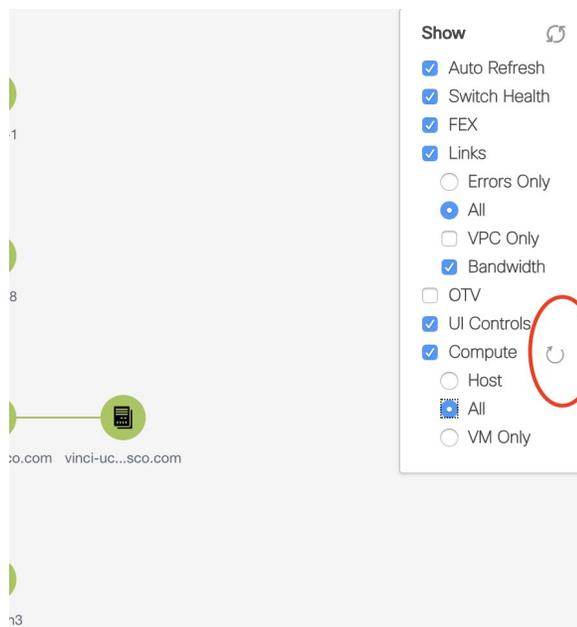


**Note** When you export the VM List to a .CSV file, the .CSV file may appear correct. However, when the .CSV file is imported into Microsoft Excel, it might get reformatted, for example, the VLAN column 1-1024 could be reformatted to a date 1/1/2019. Therefore ensure that columns are formatted correctly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

## Resynchronizing Virtual Machines

### Procedure

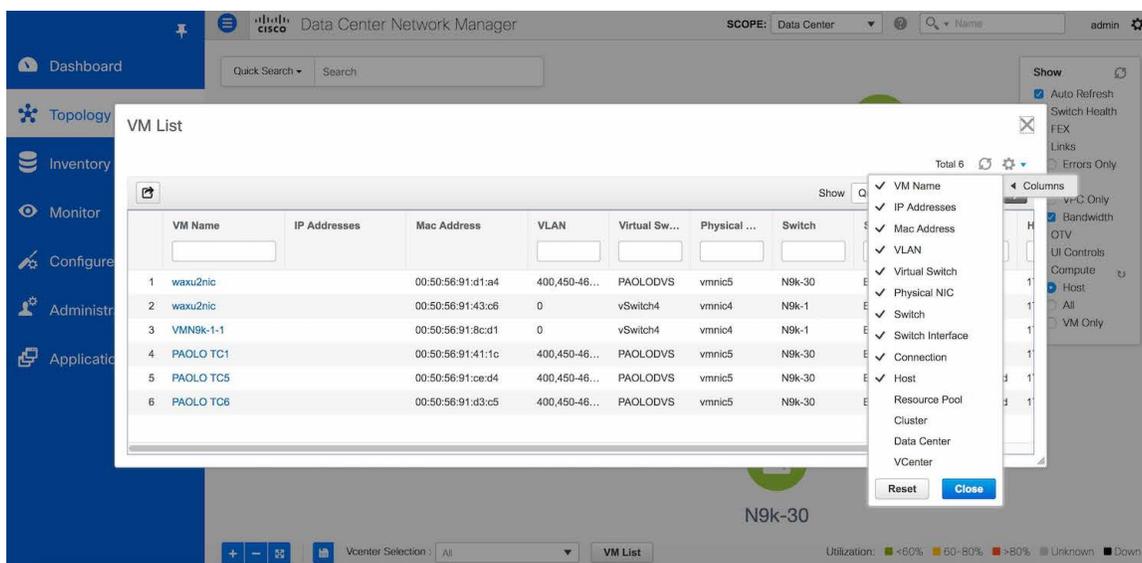
- Step 1** Choose **Topology**.
- Step 2** Click **Resync vCenters** icon next to **Compute**.



## Selecting a Column in the Virtual Machine List

### Procedure

- Step 1** In the **VM List** window, click the **Columns** under the gear icon drop-down list.

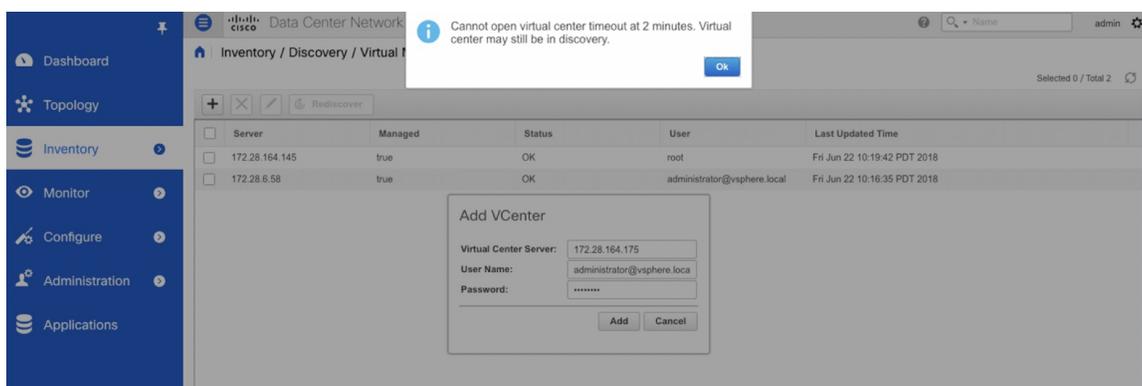


**Step 2** Select the columns that you want to display in the VM list table. If you select additional columns, click **Resync vCenters** icon to refresh and view the new columns.

Periodic resynchronization with the vCenter happens in the back-end. To configure the resync timer value, choose **Administration > DCNM Server > Server Properties**. In the **#GENERAL > DATA SOURCES VMWARE** section, specify the timer value in the **vmm.resync.timer** field. The default value is 60 (for 60minutes), and this value can be increased or decreased. If you enter a value that is less than 60 minutes, the feature is disabled.

## 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.

The screenshot displays the Cisco Data Center Network Manager interface in 'Scale Mode'. The top navigation bar includes the Cisco logo, the title 'Data Center Network Manager', the 'SCOPE: Data Center' dropdown, a search bar, and the user 'admin'. The left sidebar contains navigation links for Dashboard, Topology, Inventory, Monitor, Configure, Administration, and Applications. The main workspace shows a dense network topology with numerous nodes and interconnecting links. A 'Quick Search' bar is located at the top of the workspace. A 'Topology is cropped' warning is visible in the top right of the workspace. On the right side, a 'Show' panel provides various filtering options, including 'Auto Refresh', 'Switch Health', 'FEX', 'Links', 'Errors Only', 'VPC Only', 'Bandwidth', 'OTV', 'UI Controls', 'Compute', 'Host', and 'VM Only'. At the bottom of the workspace, there are zoom controls (+, -, fit), a 'Vcenter Selection' dropdown set to 'All', a 'VM List' button, and a legend for utilization levels: <60% (green), 60-80% (yellow), >80% (red), Unknown (grey), and Down (black).





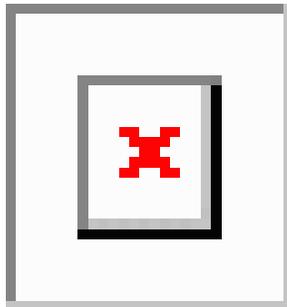
# CHAPTER 4

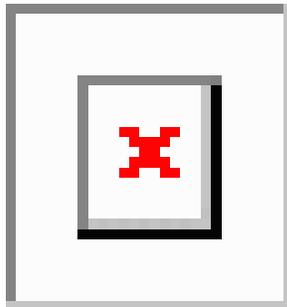
## Control

---



Note



Click  to view the contents of this chapter.

---

This chapter contains the following topics:

- [Fabrics, on page 27](#)
- [Management, on page 231](#)
- [Template Library, on page 233](#)
- [Image Management, on page 261](#)
- [Endpoint Locator, on page 269](#)
- [LAN Telemetry Health, on page 297](#)

## Fabrics

The following terms are referred to in the document:

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

This section contains the following topics:

## VXLAN BGP EVPN Fabrics Provisioning

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

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

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

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

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

Create, edit, and delete a fabric:

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

Fabric Membership changes

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

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

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

Transitioning VXLAN fabric management to DCNM

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

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

Underlay provisioning:

- Create, deploy, view, edit and delete a port-channel, vPC switch pair, straight through FEX, AA FEX, loopback, and subinterface.

- Create breakout and unbreakout ports.
- Shut down and bring up interfaces.
- Rediscover ports and view interface configuration history.
- Designate a switch interface as a routed port, trunk port, OSPF interface, and so on.




---

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

---

**Control** > **Networks** and **Control** > **VRFs** menu options (under the **Fabrics** sub menu).

Overlay network provisioning.

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

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

#### Guidelines for VXLAN BGP EVPN Fabrics Provisioning

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

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

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

The intent in DCNM contains the line:

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

The running config contains the line:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

## Creating a New VXLAN BGP EVPN Fabric

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

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

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

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

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

The fields are explained:

**Fabric Name** - Enter the name of the fabric.

**Fabric Template** - From the drop-down menu, choose the **Easy\_Fabric\_11\_1** fabric template. The fabric settings for creating a standalone fabric comes up.

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



### Note

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

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

Add Fabric ✕

\* Fabric Name :

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- a. Change the value in the drop-down box to 2.
- b. Identify the spine switches designated as route reflectors.

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

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

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

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

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

**Anycast Gateway MAC** : Specifies the anycast gateway MAC address.

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

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

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

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

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

The screenshot shows the 'Replication' configuration tab with the following fields and values:

- Replication Mode**: Multicast (dropdown menu)
- Multicast Group Subnet**: 239.1.1.0/25
- Enable Tenant Routed Multicast (TRM)**:  (checkbox)
- Default MDT Address for TRM VRFs**: (empty text field)
- Rendezvous-Points**: 2 (dropdown menu)
- RP Mode**: asm (dropdown menu)
- Underlay RP Loopback Id**: 254
- Underlay Primary RP Loopback Id**: (empty text field)
- Underlay Backup RP Loopback Id**: (empty text field)

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

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

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

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

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

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

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

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

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

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



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

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

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

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

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

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

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

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

| General | Replication | vPC                               | Advanced                                | Resources | Manageability   | Bootstrap | Configuration Backup |
|---------|-------------|-----------------------------------|---|-----------|---|-----------|----------------------|
|         |             | * vPC Peer Link VLAN              | <input type="text" value="3600"/>       |           | VLAN for vPC Peer Link SVI (Min:2, Max:3967)                      |           |                      |
|         |             | * vPC Peer Keep Alive option      | <input type="text" value="management"/> |           | Use vPC Peer Keep Alive with Loopback or Management               |           |                      |
|         |             | * vPC Auto Recovery Time          | <input type="text" value="360"/>        |           | Auto Recovery Time In Seconds (Min:240, Max:3600)                 |           |                      |
|         |             | * vPC Delay Restore Time          | <input type="text" value="150"/>        |           | vPC Delay Restore Time For vPC links in seconds (Min:1, Max:4096) |           |                      |
|         |             | vPC Peer Link Port Channel Number | <input type="text" value="500"/>        |           | Port Channel ID for vPC Peer Link (Min:1, Max:4096)               |           |                      |
|         |             | vPC IPv6 ND Synchronize           | <input checked="" type="checkbox"/>     |           | Enable IPv6 ND synchronization between vPC peers                  |           |                      |
|         |             | vPC advertise-pip                 | <input type="checkbox"/>                |           | For Primary VTEP IP Advertisement As Next-Hop Of Prefix Routes    |           |                      |

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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




---

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

---

**Enable OSPF Authentication** – Select the check box to enable OSPF authentication. Deselect the check box to disable it. If you enable this field, the OSPF Authentication Key ID and OSPF Authentication Key fields get enabled.

**OSPF Authentication Key ID** - The Key ID is populated.

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




---

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

---

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

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

**ISIS Authentication Key ID** - The Key ID is populated.

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




---

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

---

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

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

**Enable VXLAN OAM** - Enables the VXLAN OAM function for existing switches.

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

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




---

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

---

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




---

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

---

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




---

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

---

The BFD feature is disabled by default.

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

**Enable BGP Authentication** - Select the check box to enable BGP authentication. Deselect the check box to disable it. If you enable this field, the BGP Authentication Key Encryption Type and BGP Authentication Key fields are enabled.




---

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

---

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

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




---

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

---

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

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

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

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

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

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

**Spine Freeform Config** - Add CLIs that should be added to switches with a *Spine*, *Border Spine*, and *Border Gateway Spine* roles.

7. Click the **Resources** tab.

| General  | Replication | vPC             | Advanced | Resources  | Manageability | Bootstrap | Configuration Backup |
|--|-------------|-----------------|----------|--|---------------|-----------|----------------------|
| Manual Underlay IP Address Allocation <input type="checkbox"/> <small>Checking this will disable Dynamic Underlay IP Address Allocations</small> |             |                 |          |  |               |           |                      |
| * Underlay Routing Loopback IP Range   |             | 10.2.0.0/22     |          | <small>Typically Loopback0 IP Address Range</small>                              |               |           |                      |
| * Underlay VTEP Loopback IP Range  |             | 10.3.0.0/22     |          | <small>Typically Loopback1 IP Address Range</small>                              |               |           |                      |
| * Underlay RP Loopback IP Range  |             | 10.254.254.0/24 |          | <small>Anycast or Phantom RP IP Address Range</small>                            |               |           |                      |
| * Underlay Subnet IP Range   |             | 10.4.0.0/16     |          | <small>Address range to assign Numbered and Peer Link SVI IPs</small>            |               |           |                      |
| * Layer 2 VXLAN VNI Range  |             | 30000-49000     |          | <small>Overlay Network Identifier Range (Min:1, Max:16777214)</small>            |               |           |                      |
| * Layer 3 VXLAN VNI Range  |             | 50000-59000     |          | <small>Overlay VRF Identifier Range (Min:1, Max:16777214)</small>                |               |           |                      |
| * Network VLAN Range   |             | 2300-2999       |          | <small>Per Switch Overlay Network VLAN Range (Min:2, Max:3967)</small>           |               |           |                      |
| * VRF VLAN Range   |             | 2000-2299       |          | <small>Per Switch Overlay VRF VLAN Range (Min:2, Max:3967)</small>               |               |           |                      |
| * Subinterface Dot1q Range   |             | 2-511           |          | <small>Per Border Dot1q Range For VRF Lite Connectivity (Min:2, Max:511)</small> |               |           |                      |
| * VRF Lite Deployment  |             | Manual          |          | <small>VRF Lite Inter-Fabric Connection Deployment Options</small>               |               |           |                      |
| * VRF Lite Subnet IP Range   |             | 10.33.0.0/16    |          | <small>Address range to assign P2P DCI Links</small>                             |               |           |                      |
| * VRF Lite Subnet Mask   |             | 30              |          | <small>Mask for Subnet Range (Min:8, Max:31)</small>                             |               |           |                      |

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

- By default, DCNM allocates the underlay IP address resources (for loopbacks, fabric interfaces, etc) dynamically from the defined pools. If you select the check box, the allocation scheme switches to static, and some of the dynamic IP address range fields are disabled.
- For static allocation, the underlay IP address resources must be populated into the Resource Manager (RM) using REST APIs.  
Refer the Cisco DCNM REST API Reference Guide, Release 11.2(1) for more details. The REST APIs must be invoked after the switches are added to the fabric, and before you use the Save & Deploy option.
- The Underlay RP Loopback IP Range field stays enabled if BIDIR-PIM function is chosen for multicast replication.
- Changing from static to dynamic allocation keeps the current IP resource usage intact. Only future IP address allocation requests are taken from dynamic pools.

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

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

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

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

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

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

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

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

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

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

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

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



**Note** When you update a range of values, ensure that it does not overlap with other ranges. You should only update one range of values at a time. If you want to update more than one range of values, do it in separate instances. For example, if you want to update L2 and L3 ranges, you should do the following.

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

8. Click the **Manageability** tab.

| General | Replication | vPC | Advanced | Resources | Manageability                              | Bootstrap | Configuration Backup   |
|---------|-------------|-----|----------|-----------|--|-----------|--|
|         |             |     |          |           | DNS Server IP <input type="text"/>         |           | IP Address of DNS Server if used, server IP can be v4 or v6        |
|         |             |     |          |           | DNS Server VRF <input type="text"/>        |           | VRF to be used to contact DNS Server if used. VRF name can be defe |
|         |             |     |          |           | Second DNS Server IP <input type="text"/>  |           | IP Address of Second DNS Server if used, server IP can be v4 or v6 |
|         |             |     |          |           | Second DNS Server VRF <input type="text"/> |           | VRF to be used to contact Second DNS Server if used. VRF name car  |
|         |             |     |          |           | NTP Server IP <input type="text"/>         |           | IP Address of NTP Server if used, server IP can be v4 or v6        |
|         |             |     |          |           | NTP Server VRF <input type="text"/>        |           | VRF to be used to contact NTP Server if used. VRF name can be defa |
|         |             |     |          |           | Second NTP Server IP <input type="text"/>  |           | IP Address of Second NTP Server if used, server IP can be v4 or v6 |
|         |             |     |          |           | Second NTP Server VRF <input type="text"/> |           | VRF to be used to contact Second NTP Server if used. VRF name can  |

The fields in this tab are:

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

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

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

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

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

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

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

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

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

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

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



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

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

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

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

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

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

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

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

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

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

## 9. Click the **Bootstrap** tab.

| General  | Replication | vPC | Advanced | Resources | Manageability | Bootstrap | Configuration Backup |
|--|-------------|-----|----------|-----------|---------------|-----------|----------------------|
| <p><b>Enable Bootstrap</b> <input type="checkbox"/> ? <i>Automatic IP Assignment For POAP</i></p> <p>Enable Local DHCP Server <input type="checkbox"/> ? <i>Automatic IP Assignment For POAP From Local DHCP Server</i></p> <p>DHCP Scope Start Address <input type="text"/> ? <i>Start Address For Switch Out-of-Band POAP</i></p> <p>DHCP Scope End Address <input type="text"/> ? <i>End Address For Switch Out-of-Band POAP</i></p> <p>Switch Management Default Gateway <input type="text"/> ? <i>Default Gateway For Mgmt VRF On The Switch</i></p> <p>Switch Management Subnet Prefix <input type="text"/> ? <i>Prefix For Mgmt0 Interface On The Switch (Min:8, Max:30)</i></p> <p>Bootstrap Freeform Config <input type="text"/> ? <i>Note ! All configs should strictly match 'show run' output, with respect to case and newlines. Any mismatches will yield unexpected diffs during deploy.</i></p> <p>DHCP Multi Subnet Scope <input type="text"/> ? <i>Enter One Subnet Scope per line. Start_IP, End_IP, Gateway, Prefix e.g. 10.6.0.2, 10.6.0.9, 10.6.0.1, 24 10.7.0.2, 10.7.0.9, 10.7.0.1, 24</i></p> |             |     |          |           |               |           |                      |

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

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

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

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

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

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

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

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

*DHCP scope and management default gateway IP address specification* - If you specify the management default gateway IP address 10.0.1.1 and subnet mask 24, ensure that the DHCP scope is within the specified subnet, between 10.0.1.2 and 10.0.1.254.

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

Copy-paste the running-config to a **freeform config** field with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. For more information, see [Resolving Freeform Config Errors in Switches, on page 230](#).

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

The format of the scope should be defined as:

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

For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24

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

General | EVPN | vPC | Advanced | Manageability | Bootstrap | **Configuration Backup**

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

Scheduled Fabric Backup  ? Backup at Specified Scheduled Time

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

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

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

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

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

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

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

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

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



---

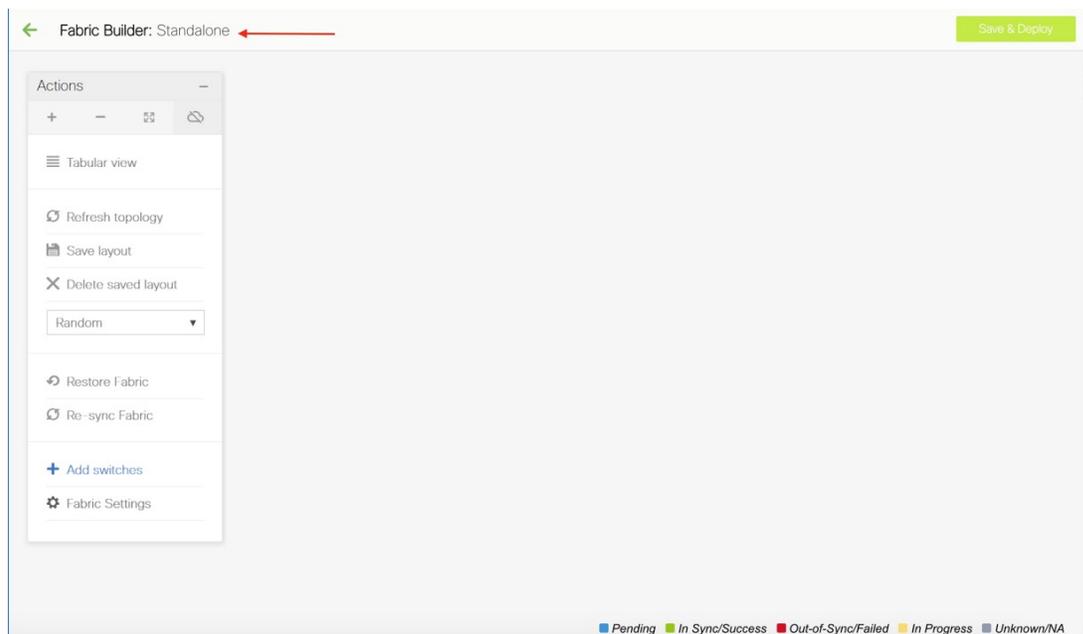
**Note** Hourly and scheduled backup processes happen only during the next periodic configuration compliance activity, and there can be a delay of up to an hour. To trigger an immediate backup, do the following:

- a. Choose **Control > Fabric Builder**. The Fabric Builder screen comes up.
- b. Click within the specific fabric box. The fabric topology screen comes up.
- c. From the **Actions** pane at the left part of the screen, click **Re-Sync Fabric**.

---

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

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

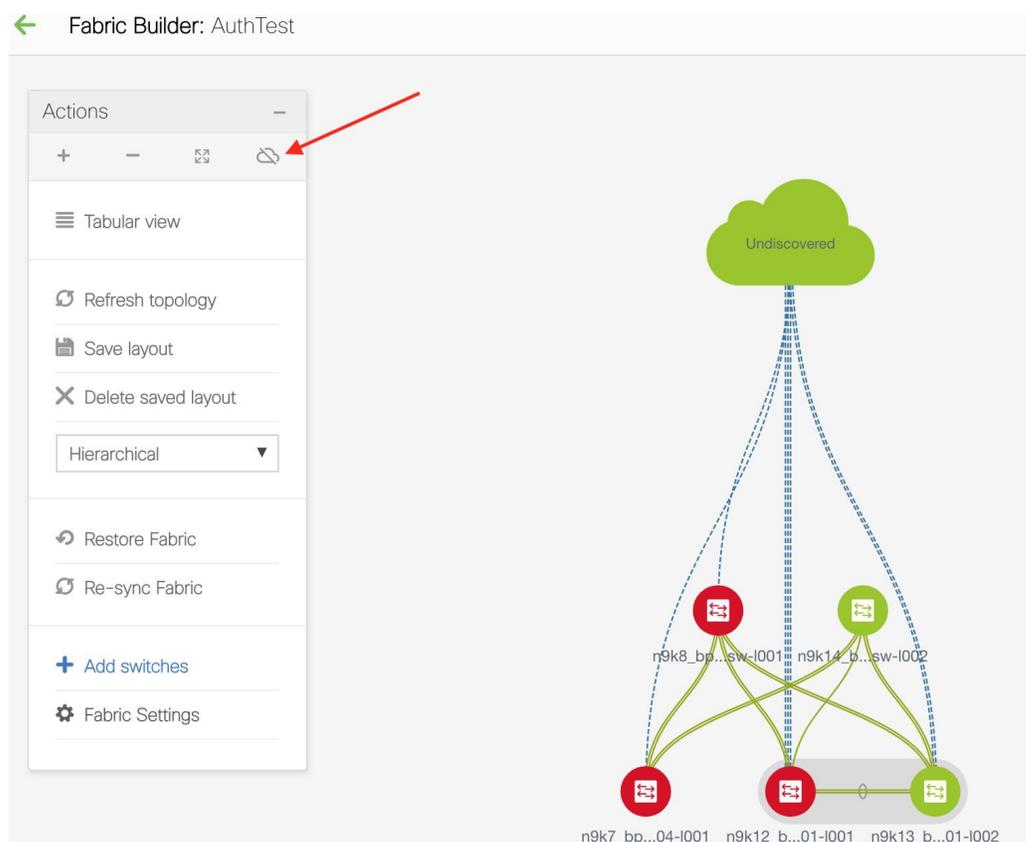


(At the same time, the newly created fabric instance appears on the **Fabric Builder** screen. To go to the **Fabric Builder** screen, click the left arrow (←) button above the **Actions** pane [to the left of the screen]).

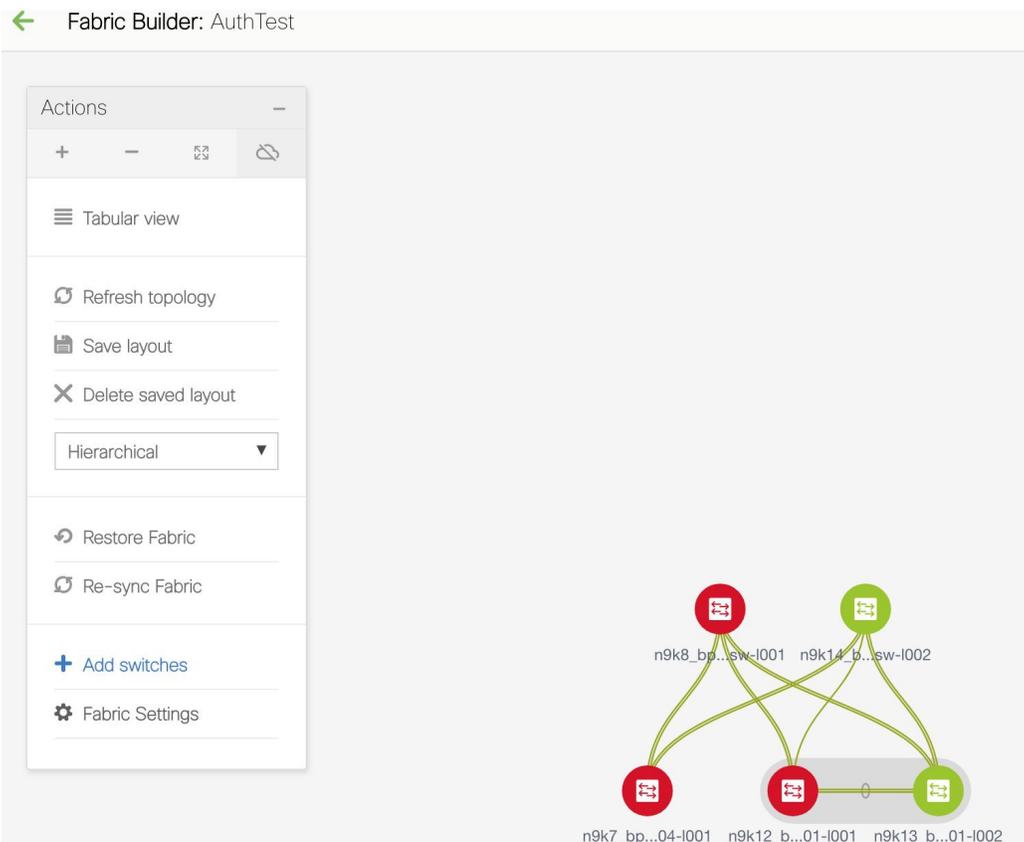
The **Actions** pane allows you to perform various functions. One of them is the **Add switches** option to add switches to the fabric. After you create a fabric, you should add fabric devices. The options are explained:

- **Tabular View** - By default, the switches are displayed in the topology view. Use this option to view switches in the tabular view.
- **Refresh topology** - Allows you to refresh the topology.
- **Save Layout** – Saves a custom view of the topology. You can create a specific view in the topology and save it for ease of use.
- **Delete saved layout** – Deletes the custom view of the topology
- **Topology views** - You can choose between Hierarchical, Random and Custom saved layout display options.
  - **Hierarchical** - Provides an architectural view of your topology. Various Switch Roles can be defined that draws the nodes on how you configure your CLOS topology.
  - **Random** - Nodes are placed randomly on the window. DCNM tries to make a guess and intelligently place nodes that belong together in close proximity.
  - **Custom saved layout** - You can drag nodes around to your liking. Once you have the positions as how you like, you can click Save Layout to remember the positions. Next time you come to the topology, DCNM will draw the nodes based on your last saved layout positions.
- **Restore Fabric** – Allows you to restore the fabric to a prior DCNM configuration state (one month back, two months back, and so on). For more information, see the *Restore Fabric* section.

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

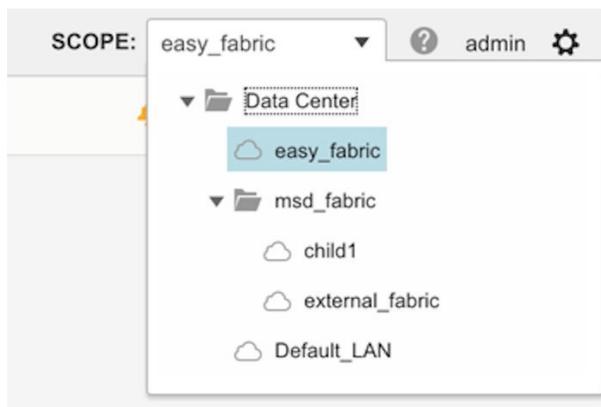


When you click the icon, the Undiscovered cloud and its links to the selected fabric topology are not displayed.



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

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



## Adding Switches to a Fabric

Switches in each fabric are unique, and hence, each switch can only be added to one fabric.

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

Additionally, you can pre-provision switches. For more information, see [Pre-provisioning a Device](#), on page 59.

## Discovering Existing Switches

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

### Inventory Management

Discover Existing Switches

PowerOn Auto Provisioning (POAP)

Discovery Information >

Scan Details >

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

Authentication Protocol  ▼

Username

Password

Max Hops  ▲  
▼ hop(s)

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

Start discovery

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

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... |          |

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

Inventory Management ✕

Discover Existing Switches | PowerOn Auto Provisioning (POAP)

Discovery Information > Scan Details >

← Back 2 Import into fabric

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

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

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



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

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

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



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

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

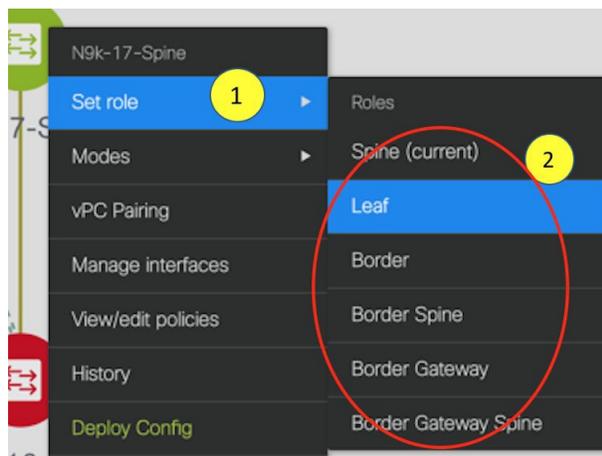
Device connectivity issue: Before proceeding further, wait for ten minutes for the switch-internal processes to complete. Else, you might encounter a device connectivity failure message at a later stage.

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

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



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



**Note**

- Starting from DCNM 11.1(1), switch roles can be changed if there are no overlays on the switches, but only as per the list of allowed switch role changes given at [Switch Operations, on page 135](#).
- After you upgrade to Cisco DCNM Release 11.1(1) with an existing fabric with the Easy\_Fabric template, you cannot set the Border Spine or Border Gateway Spine roles to switches, because these roles are not supported with the Easy\_Fabric template. You need to use the **Easy\_fabric\_11\_1** template to set these roles for switches in a fabric.

If you choose the Hierarchical layout for display (in the Actions panel), the topology automatically gets aligned as per role assignment, with the leaf switches at the bottom, the spine switches connected on top of them, and the border switches at the top.

**Note**

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

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

**Note**

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Config Deployment ✕

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

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

[Deploy Config](#)

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

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

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

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

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

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

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

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

#### 7. Close the screen.

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

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

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

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



---

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

---

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

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

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

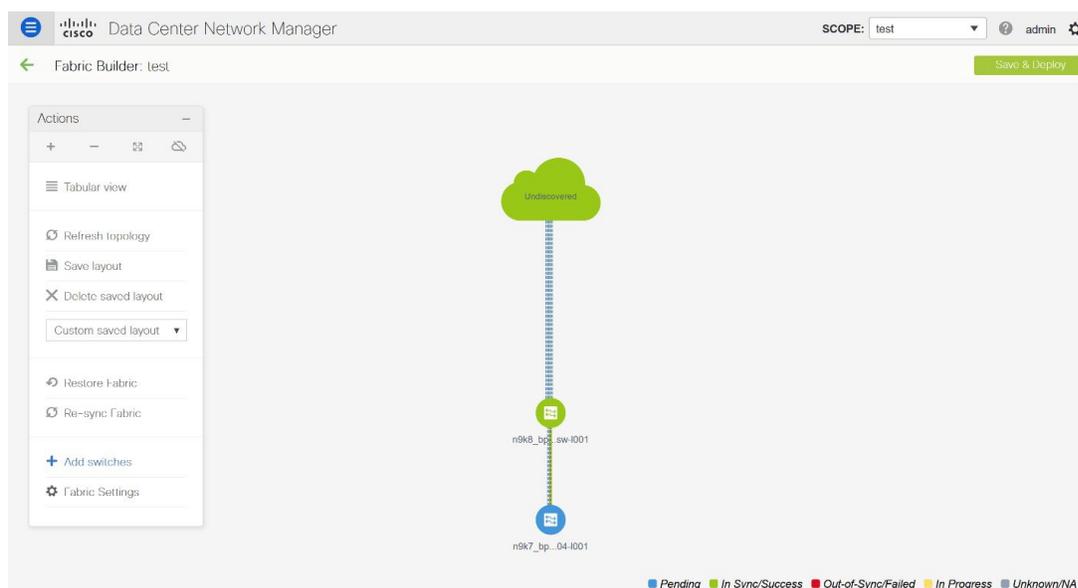
When a leaf switch boots up after a switch reload or RMA operation, DCNM provisions configurations for the switch and FEX devices connected to it. Occasionally, FEX connectivity comes up after DCNM provisions FEX (host interface) configurations, resulting in a configuration mismatch. To resolve the mismatch,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

An example of the **Deploy Config** option usage is for switch-level freeform configurations. Refer [Enabling Freeform Configurations on Fabric Switches](#) for details.

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

## Discovering New Switches

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



**Note** If you want to POAP with 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 window and click the **Add switches** option from the **Actions** panel. The Inventory Management window comes up.
6. Click the **POAP** tab.

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

**Note**

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

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

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

Inventory Management ✕

Discover Existing Switches PowerOn Auto Provisioning (POAP)

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

+ 📄 🔄 \* Admin Password  \* Confirm Admin Password  🔒

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

Close

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

Beginning with Release 11.2(1), you can provision devices in advance. To pre-provision devices, refer to [Pre-provisioning a Device](#), on page 59.

7. In the **Admin Password** and **Confirm Admin Password** fields, enter and confirm the admin password. This admin password is applicable for all the switches displayed in the POAP window.

**Note**

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

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

Inventory Management ✕

Discover Existing Switches | PowerOn Auto Provisioning (POAP)

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

+ 🔄 ↺ \* Admin Password  \* Confirm Admin Password  🔒

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

Close

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

Inventory Management ✕

Discover Existing Switches | PowerOn Auto Provisioning (POAP)

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

+ 🔄 ↺ \* Admin Password  \* Confirm Admin Password  🔒

Discovery Credentials ✕

\*Discovery Username:

\*Discovery Password:

\*Confirm Discovery Password:

OK Clear

Close

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

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

**Note**

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

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

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

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

**Note**

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

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

12. After the pending configurations are deployed, the **Progress** column displays 100% for all switches.
13. Click **Close** to return to the fabric builder topology.
14. Click **Refresh Topology** to view the update. All switches must be in green color indicating that they are functional.
15. The switch and the link are discovered in DCNM. Configurations are built based on various policies (such as fabric, topology, and switch generated policies). The switch image (and other required) configurations are enabled on the switch.
16. In the DCNM GUI, the discovered switches can be seen in the *Standalone* fabric topology. Up to this step, the POAP is completed with basic settings. All the interfaces are set to trunk ports. You must setup interfaces through the **Control > Interfaces** option for any additional configurations, but not limited to the following:
  - vPC pairing.
  - Breakout interfaces.
  - Port channels, and adding members to ports.

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

### Fabric errors & warnings



0 Errors, 2 Warnings, 0 Info

[X Delete all](#)

**Warning** The Secondary IP address of the NVE source interface has been modified for switch SN [FDO20260UEK] and peer SN [FDO20291AVQ] due to vpc feature configuration. Please make sure to shut/noshut the nve interfaces from DCNM Interface Manager Screen. X

|             |  |
|-------------|--|
| Severity    | warning  |
| Category    | Fabric   |
| Entity type | Fabric_Template  |
| Entity name | configSave:vpcPairing:FDO20260UEK:FDO20291AVQ  |
| Reported    | less than a minute ago 2019-03-17 09:30:00   |
| Details     | [2]: [vpcPairing:FDO20260UEK:FDO20291AVQ]. Line/Col:[0/0]. Msg = [The Secondary IP address of the NVE source interface has been modified for switch SN [FDO20260UEK] and peer SN [FDO20291AVQ] due to vpc feature configuration. Please make sure to shut/noshut the nve interfaces from DCNM Interface Manager Screen.] |

**Warning** The Secondary IP address of the NVE source interface has been modified for switch SN [FDO20291AVQ] and peer SN [FDO20260UEK] due to vpc feature configuration. Please make sure to shut/noshut the nve interfaces from DCNM Interface Manager Screen. X

|             |  |
|-------------|--|
| Severity    | warning  |
| Category    | Fabric   |
| Entity type | Fabric_Template  |
| Entity name | configSave:vpcPairing:FDO20291AVQ:FDO20260UEK  |
| Reported    | less than a minute ago 2019-03-17 09:30:00   |
| Details     | [1]: [vpcPairing:FDO20291AVQ:FDO20260UEK]. Line/Col:[0/0]. Msg = [The Secondary IP address of the NVE source interface has been modified for switch SN [FDO20291AVQ] and peer SN [FDO20260UEK] due to vpc feature configuration. Please make sure to shut/noshut the nve interfaces from DCNM Interface Manager Screen.] |

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

## Interfaces

|                                     | Device Name | Name        |
|-------------------------------------|-------------|-------------|
| <input type="checkbox"/>            | N9K-2-Leaf  | Ethernet2/0 |
| <input type="checkbox"/>            | N9K-2-Leaf  | Ethernet2/1 |
| <input type="checkbox"/>            | N9K-2-Leaf  | Ethernet2/2 |
| <input type="checkbox"/>            | N9K-2-Leaf  | Ethernet2/3 |
| <input type="checkbox"/>            | N9K-2-Leaf  | Ethernet2/4 |
| <input type="checkbox"/>            | N9K-2-Leaf  | Ethernet2/5 |
| <input type="checkbox"/>            | N9K-2-Leaf  | Ethernet2/6 |
| <input checked="" type="checkbox"/> | N9K-2-Leaf  | nve1        |

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

**Note**

- After discovering a switch (new or existing), at any point in time you can provision configurations on it again through the POAP process. The process removes existing configurations and provision new configurations. You can also deploy configurations incrementally without invoking POAP.
- You might encounter an issue with module discovery after bootstrap. In such cases, the discovery happens after a delay. If not, go through the discovery process again.

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

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

**Note**

- Changing of the switch role is allowed only before executing **Save & Deploy**.
- Starting from DCNM 11.1(1), switch roles can be changed if there are no overlays on the switches, but only as per the list of allowed switch role changes given at [Switch Operations, on page 135](#).
- After you upgrade to Cisco DCNM Release 11.1(1) with an existing fabric with the **Easy\_Fabric** template, you cannot set the Border Spine or Border Gateway Spine roles to switches, because these roles are not supported with the **Easy\_Fabric** template. You need to use the **Easy\_fabric\_11\_1** template to set these roles for switches in a fabric.

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

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

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

## Pre-provisioning a Device

In DCNM 11.2, you can provision devices in advance.




---

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

---

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

Note the following guidelines:

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

The examples of the values in the **Data** field are as follows:

- `{"modulesModel": ["N9K-C93180LC-EX"], "gateway": "10.1.1.1/24"}`
- `{"modulesModel": ["N9K-C93180LC-EX"], "breakout": "interface breakout module 1 port 1 map 10g-4x", "portMode": "hardware profile portmode 4x100G+28x40G", "gateway": "172.22.31.1/24"}`

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

## Procedure

---

- Step 1** 1. Click **Control > Fabric Builder**.
- The **Fabric Builder** screen is displayed.
- Step 2** Click within the fabric box.
- Step 3** From the Actions panel, click the **Add switches** option.
- The **Inventory Management** screen is displayed.
- Step 4** Click the **POAP** tab.
- Step 5** In the **POAP** tab, do the following:
- Click + from the top left part of the screen.  
The Add a new device screen comes up.
  - Fill up the device details as shown in the screenshot.
  - Click **Save**.

Inventory Management

Discover Existing Switches | PowerOn Auto Provisioning (POAP) | Move Neighbor Switches

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

\* Admin Password:  \* Confirm Admin Password:  Bootstrap

+ [Refresh] [Refresh] [Refresh]

| Serial Number | Model | Version | IP Address | Hostname |
|---------------|-------|---------|------------|----------|
|               |       |         |            |          |

Add a new device to pre-provisioning

\*Serial Number:

\*Model:

\*Version:

\*IP Address:

\*Hostname:

\*Data:   
JSON Object which contains model name of the Modules  
 Eg: {"modulesModel": "N9K-EX"}

Save Clear

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

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

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

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

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

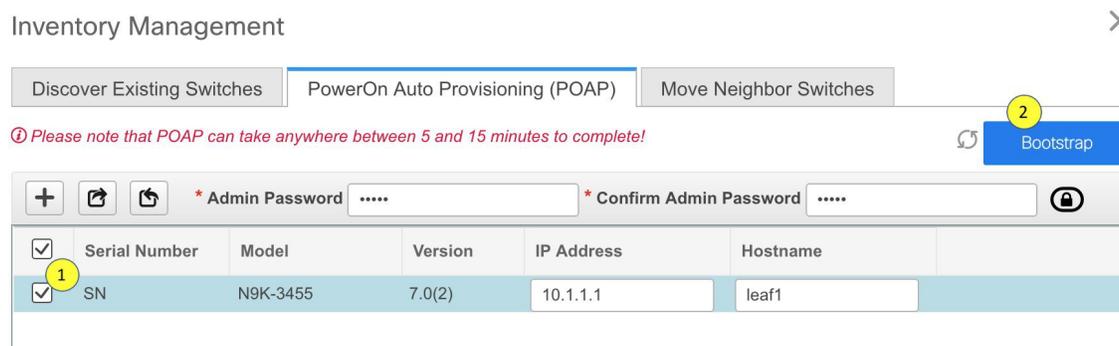
Add new devices' information in the .csv file with all the mandatory fields (SerialNumber, Model, version, IpAddress, Hostname and Data fields [JSON Object]).

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

|   | A                            | B                     | C                        | D                        | E         | F  | G |
|---|------------------------------|-----------------------|--------------------------|--------------------------|-----------|--|---|
| 1 | #SerialNumber(Eg:FD01344GH5) | #Model(Eg:N9k-C9236C) | #Version(Eg:7.0(3)12(3)) | #IPAddress of the device | #HostName | #Data(JSON Field contains model name of the modules) |   |
| 2 | Serial Number                | Model                 | Version                  | IP Address               | Hostname  | Data   |   |
| 3 | FDO21331SND                  | N9K-93180YC-EX        | 7.0(3)15(2)              | 1.1.1.1                  | leaf1     | {"modulesModel":["N9K-93180YC-EX"]}                  |   |
| 4 | FDO21351N3X                  | N9K-C9236C            | 7.0(3)14(1)              | 11.1.1.1                 | spine1    | {"modulesModel":["N9K-C9236C"]}                      |   |
| 5 | FDO21491A5K                  | N9K-C93240YC-FX2      | 7.0(3)17(3)              | 12.1.1.1                 | leaf2     | {"modulesModel":["N9K-C93240YC-FX2"]}                |   |
| 6 |                              |                       |                          |                          |           |  |   |

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

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



The leaf1 device appears in the fabric topology.

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

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

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

Create intra-link policy using one of the templates:

- **int\_pre\_provision\_intra\_fabric\_link** to automatically generate intra fabric interface configuration with DCNM allocated IP addresses
- **int\_intra\_fabric\_unnum\_link\_11\_1** if you are using unnumbered links
- **int\_intra\_fabric\_num\_link\_11\_1** if you want to manually assign IP addresses to intra-links

Click **Save & Deploy**.

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

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

For more information, see [Return Material Authorization \(RMA\)](#), on page 161.

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

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

## Changing the TCAM Configuration on a Device

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

- Cisco Nexus 9300 Series Switches: **tcam\_pre\_config\_9300** and **tcam\_pre\_config\_vxlan**
- Cisco Nexus 9500 Series Switches: **tcam\_pre\_config\_9500** and **tcam\_pre\_config\_vxlan**

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

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

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

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

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

Make sure that the priority of this **tcam\_config** policy is higher than the **tcam\_pre\_config\_vxlan** policy so that the config policy with **racl 1024** is configured before the **tcam\_pre\_config\_vxlan** policy.




---

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

---

## Adding a vPC L3 Peer Keep-Alive Link

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




---

**Note**

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

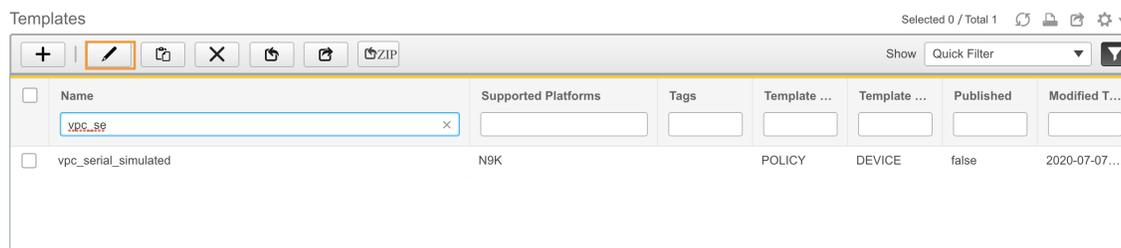
---

### Procedure

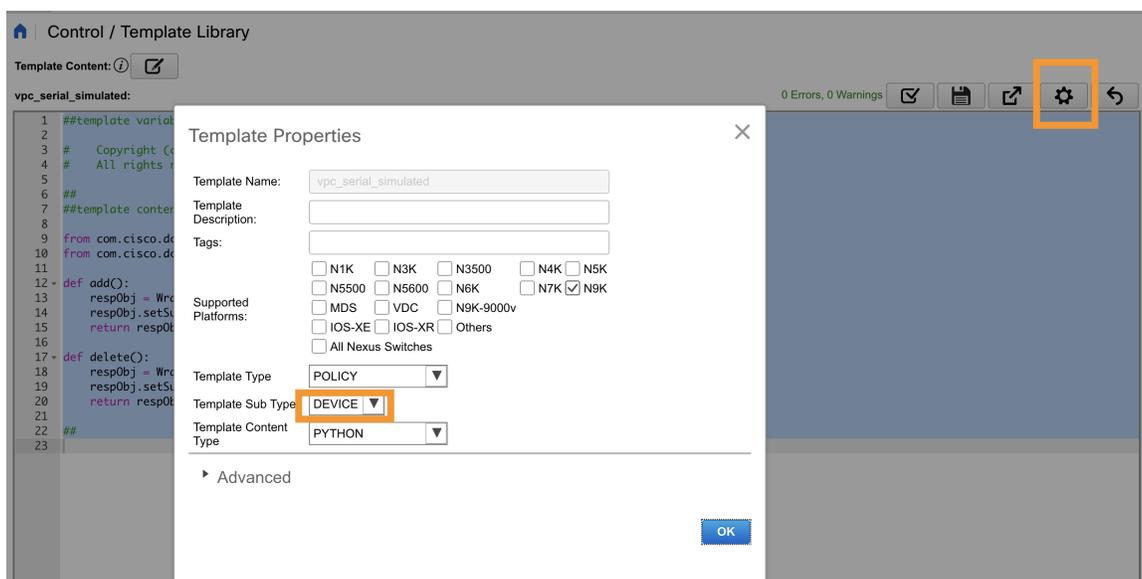
---

- Step 1** From DCNM, navigate to **Control > Template Library**.
- Step 2** Search for the **vpc\_serial\_simulated** policy, select it, and click the **Edit** icon.

## Adding a vPC L3 Peer Keep-Alive Link



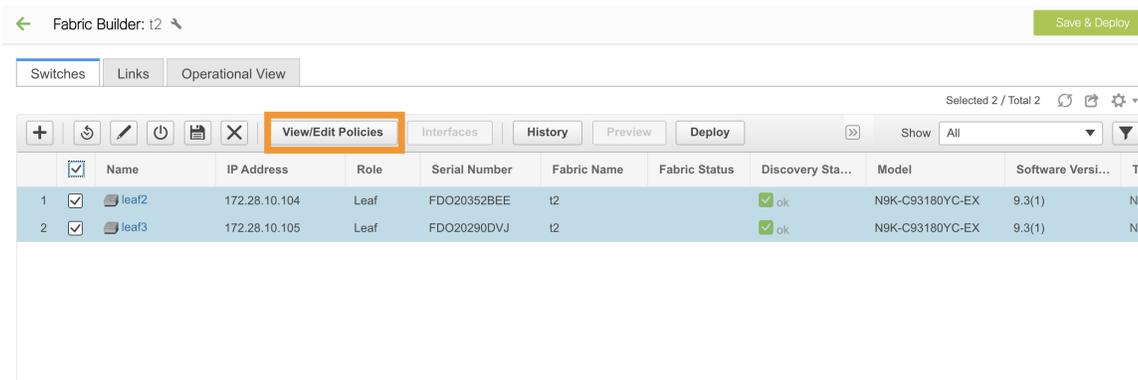
**Step 3** Edit the template properties and set the **Template Sub Type** to **Device** so that this policy appears in **View/Edit Policies**.



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

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

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



**Step 6** Click + to add policies.

**Step 7** From the **Policy** drop-down list, select **vpc\_serial\_simulated** policy and add priority. Click **Save**.

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

- Step 8** Navigate back to **Tabular View** and click the **Links** tab.
- Step 9** Select the link between vPC pair, which has to be a vPC peer keep alive and click **Edit**.
- Step 10** From the **Link Template** drop-down list, select **int\_intra\_vpc\_peer\_keep\_alive\_link\_11\_1**.  
Enter values for the remaining fields. Make sure to leave the field empty for the default VRF and click **Save**.

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

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

If VRF is non-default, use **switch\_freeform** to create the respective VRF.

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

The screenshot displays the Cisco Data Center Network Manager (DCNM) interface. The main area shows a topology view with two green circular nodes labeled 'leaf2' and 'leaf3' connected by a double line. A sidebar on the left lists actions such as 'Refresh topology', 'Save layout', and 'Add switches'. A detailed panel on the right shows the configuration for 'leaf2', including its IP address (172.28.10.104), serial number (FDO20352BEE), version (9.3(1)), and VPC Domain ID (1). The 'Peerlink State' is highlighted in orange, showing 'Peer is OK'. The 'Health' section shows a 95% overall status with metrics for Modules (91.67%), Switch ports (83.61%), and Alarms (100.00%).

## Brownfield Deployment-Transitioning VXLAN Fabric Management to DCNM

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

## Creating a New Fabric for eBGP-Based Underlay

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

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

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

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

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

The fields are explained:

**Fabric Name** - Enter the name of the fabric.

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

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

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

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

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

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

The fabric is identified by the spine switch AS number.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

4. Click **EVPN**. Most of the fields in this tab are auto-populated. The fields are:

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

You can convert a routed fabric to a VXLAN enabled fabric by selecting this option. When the fabric is VXLAN enabled, you can create and deploy overlay networks or VRFs. The procedure for creating and deploying networks or VRFs is the same as in Easy\_Fabric\_11\_1. For more information, see *Creating and Deploying Networks and VRFs* in the Control chapter in *Cisco DCNM LAN Fabric Configuration Guide*.

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

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

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




---

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

---

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

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

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

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




---

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

---

**Enable Tenant DHCP:** Enables tenant DHCP support.

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

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

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

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

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

The screenshot shows the configuration interface for the vPC section, specifically the Advanced tab. The tabs at the top are General, EVPN, vPC, Advanced, Manageability, Bootstrap, and Configuration Backup. The Advanced tab is active. The configuration fields are as follows:

- Power Supply Mode:** ps-redundant (dropdown menu)
- CoPP Profile:** strict (dropdown menu)
- Enable BFD:**  (checkbox)
- Greenfield Cleanup Option:** Disable (dropdown menu)
- Enable BGP Authentication:**  (checkbox)
- BGP Authentication Key Encryption ...:** (dropdown menu)
- BGP Authentication Key:** (text input field)
- Leaf Freeform Config:** (text area)
- Spine Freeform Config:** (text area)
- VRF Lite Subnet IP Range:** 10.33.0.0/16 (text input field)
- VRF Lite Subnet Mask:** 30 (text input field)

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

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

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



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

The BFD feature is disabled by default.

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

**Enable BGP Authentication:** Select the check box to enable BGP authentication. Deselect the check box to disable it. If you enable this field, the BGP Authentication Key Encryption Type and BGP Authentication Key fields are enabled.

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

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



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

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

**Spine Freeform Config** - Add CLIs that should be added to switches with a Spine, Border Spine, and Border Gateway Spine roles.

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

## 7. Click the **Manageability** tab.

| General | EVPN | vPC | Advanced | Manageability                 | Bootstrap | Configuration Backup  |
|---------|------|-----|----------|-------------------------------|-----------|---|
|         |      |     |          | DNS Server IP                 |           | ? IP Address of DNS Server if used, server IP can be v4 or v6                                       |
|         |      |     |          | DNS Server VRF                |           | ? VRF to be used to contact DNS Server if used. VRF name can be default, management, etc.           |
|         |      |     |          | Second DNS Server IP          |           | ? IP Address of Second DNS Server if used, server IP can be v4 or v6                                |
|         |      |     |          | Second DNS Server VRF         |           | ? VRF to be used to contact Second DNS Server if used. VRF name can be default, management, etc.    |
|         |      |     |          | NTP Server IP                 |           | ? IP Address of NTP Server if used, server IP can be v4 or v6                                       |
|         |      |     |          | NTP Server VRF                |           | ? VRF to be used to contact NTP Server if used. VRF name can be default, management, etc.           |
|         |      |     |          | Second NTP Server IP          |           | ? IP Address of Second NTP Server if used, server IP can be v4 or v6                                |
|         |      |     |          | Second NTP Server VRF         |           | ? VRF to be used to contact Second NTP Server if used. VRF name can be default, management, etc.    |
|         |      |     |          | AAA Server Type               | none      | ? radius, tacacs, or none if not using AAA  |
|         |      |     |          | AAA Server IP                 |           | ? IP Address of AAA Server if used, server IP can be v4 or v6                                       |
|         |      |     |          | AAA Shared Secret             |           | ? Shared secret (type-7 encrypted) if AAA Server is used (Max Size 63)                              |
|         |      |     |          | Second AAA Server IP          |           | ? IP Address of second AAA Server if used, server IP can be v4 or v6                                |
|         |      |     |          | Second AAA Shared Secret      |           | ? Shared secret (type-7 encrypted) if Second AAA Server is used (Max Size 63)                       |
|         |      |     |          | AAA Server VRF                |           | ? VRF to be used to contact AAA Server(s) if used. VRF name can be default, management, etc.        |
|         |      |     |          | Syslog Server IP              |           | ? IP Address of Syslog Server if used, server IP can be v4 or v6                                    |
|         |      |     |          | Syslog Server Severity        | 5         | ? Syslog severity   |
|         |      |     |          | Syslog Server VRF             |           | ? VRF to be used to contact Syslog Server if used. VRF name can be default, management, etc.        |
|         |      |     |          | Second Syslog Server IP       |           | ? IP Address of Second Syslog Server if used, server IP can be v4 or v6                             |
|         |      |     |          | Second Syslog Server Severity | 5         | ? Second Syslog Server severity   |
|         |      |     |          | Second Syslog Server VRF      |           | ? VRF to be used to contact Second Syslog Server if used. VRF name can be default, management, etc. |

The fields in this tab are:

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

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

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

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

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

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

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

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

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

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

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



---

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

---

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

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

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

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

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

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

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

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

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

8. Click the **Bootstrap** tab.

| General   | EVPN | vPC | Advanced | Manageability | Bootstrap | Configuration Ba                               |
|---|------|-----|----------|---------------|-----------|--|
| <b>Enable Bootstrap</b> <input type="checkbox"/> <span>?</span> <i>Automatic IP Assignment For POAP</i>                         |      |     |          |               |           |  |
| Enable Local DHCP Server <input type="checkbox"/> <span>?</span> <i>Automatic IP Assignment For POAP From Local DHCP Server</i> |      |     |          |               |           |  |
| DHCP Scope Start Address  |      |     |          |               |           | <span>?</span> <i>Start Address For Switc</i>  |
| DHCP Scope End Address  |      |     |          |               |           | <span>?</span> <i>End Address For Switc</i>    |
| Switch Management Default Gateway   |      |     |          |               |           | <span>?</span> <i>Default Gateway For Mg</i>   |
| Switch Management Subnet Prefix   |      |     |          |               |           | <span>?</span> <i>Prefix For Mgmt0 Interfa</i> |
| Bootstrap Freeform Config   |      |     |          |               |           |  |
| DHCP Multi Subnet Scope   |      |     |          |               |           |  |

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

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

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

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

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

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

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

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

*DHCP scope and management default gateway IP address specification* - If you specify the management default gateway IP address 10.0.1.1 and subnet mask 24, ensure that the DHCP scope is within the specified subnet, between 10.0.1.2 and 10.0.1.254..

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

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

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

The format of the scope should be defined as:

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

For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24

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

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

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

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

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

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

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



- Note** Hourly and scheduled backup processes happen only during the next periodic configuration compliance activity, and there can be a delay of up to an hour. To trigger an immediate backup, do the following:
- Choose **Control > Fabric Builder**. The Fabric Builder screen comes up.
  - Click within the specific fabric box. The fabric topology screen comes up.
  - From the **Actions** panel at the left part of the screen, click **Re-Sync Fabric**.

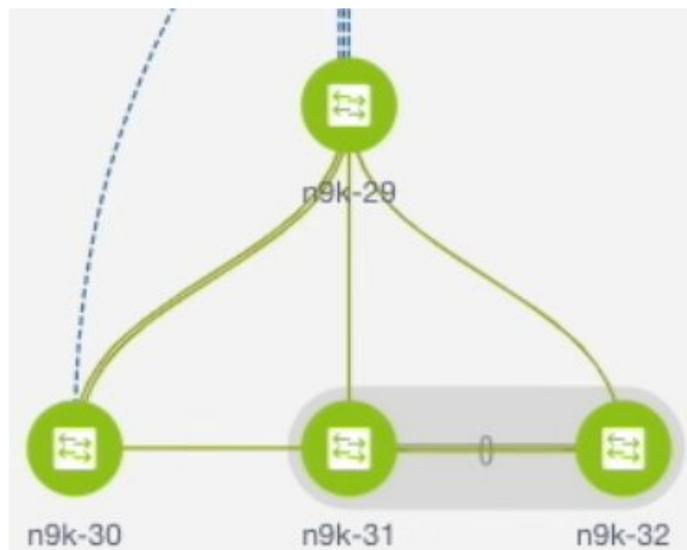
You can also initiate the fabric backup in the fabric topology window. Click **Backup Now** in the **Actions** pane.

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

### VXLAN Fabric With eBGP Underlay – Pointers

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

## Applying Policies On A Fabric With An eBGP Underlay



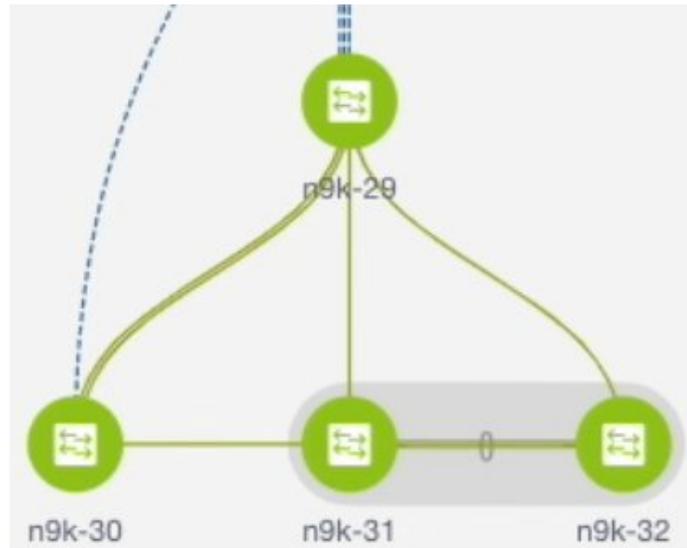
The topology shows a VXLAN fabric enabled with eBGP for the underlay. In DCNM, a fabric with the Easy\_Fabric\_eBGP template is created. One spine switch (n9k-29) and three leaf switches (n9k-30, and vPC switch pair n9k-31 and n9k-32) are imported to it.

This topic covers the following:

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

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

## Deploying Fabric Underlay eBGP Policies



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

The two different types of fabrics are:

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

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

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

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



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

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



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

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

## Deploying Fabric Overlay eBGP Policies

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

## Deploying Spine Switch Overlay Policies

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

Add Policy
✕

\* Priority (1-1000):

\* Policy:  ▼

General

\* Leaf IP List  ? list of leaf IP address for peering list e.g. 10.2.0.

\* Leaf BGP ASN  ? BGP ASN of each leaf, separated by ,

\* BGP Update-Source Interface  ? Source of BGP session and updates

Enable Tenant Routed Multicast  ? Tenant Routed Multicast setting needs to match the fabric setting

Variables:      Enable BGP Authentication  ? BGP Authentication needs to match the fabric setting

The fields on the screen are:

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

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

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



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

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

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

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

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

## Deploying Leaf Switch Overlay Policies

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

Add Policy
✕

\* Priority (1-1000):

\* Policy:  ▼

General

\* Spine IP List  ? list of spine IP address for peering list e.g. 10.2.

\* BGP Update-Source Interface  ? Source of BGP session and updates

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

Enable BGP Authentication  ? BGP Authentication needs to match the fabric setting

Variables:

The fields on the screen are:

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

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

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

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

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

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

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

## Guidelines

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

## Creating an External Fabric

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

- An external fabric is a monitor-only or managed mode fabric.
- You can import, remove, and delete switches for an external fabric.
- For Inter-Fabric Connection (IFC) cases, you can choose Cisco 9000, 7000 and 5600 Series switches as destination switches in the external fabric.
- You can use non-existing switches as destination switches.
- The template that supports an external fabric is External\_Fabric.
- If an external fabric is an MSD fabric member, then the MSD topology screen displays the external fabric with its devices, along with the member fabrics and their devices.

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

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

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

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

- External fabric in Monitored or Managed Mode

### Creating External Fabric from Fabric Builder

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

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

**Fabric Name** - Enter the name of the external fabric.

**Fabric Template** - Choose *External\_Fabric*.

When you choose the fabric template, the fabric creation screen for creating an external fabric comes up.

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

#### General tab

**BGP AS #** - Enter the BGP AS number.

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

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

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

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

#### Advanced tab

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

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

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

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

#### Resources tab

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

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

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

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

The fields on this tab are:

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

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

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

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

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

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

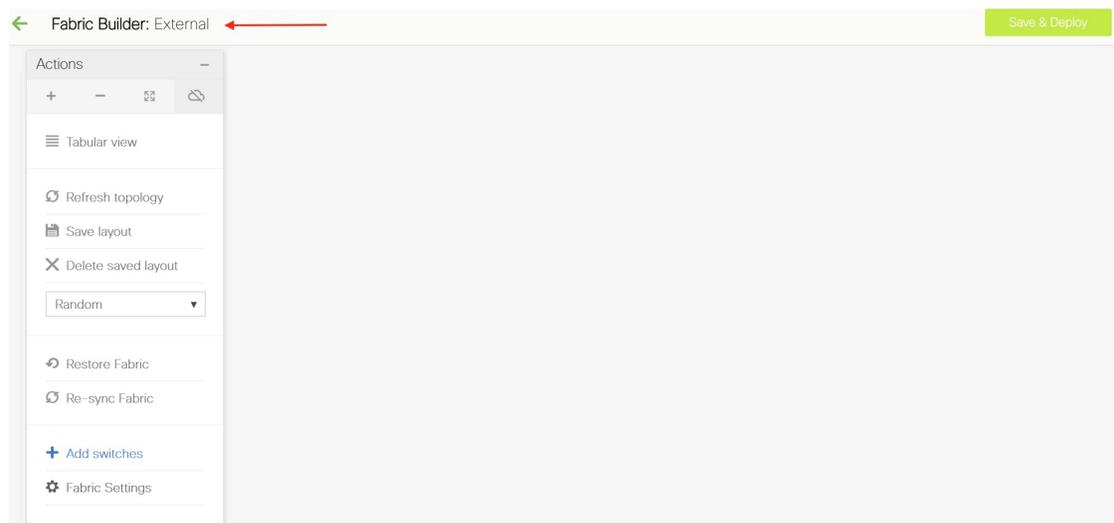
You can also initiate the fabric backup in the fabric topology window. Click **Backup Now** in the **Actions** pane.

Pointers for hourly and scheduled backup:

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

## 5. Click **Save**.

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



After creating the external fabric, add switches to it.

### Add Switches to the External Fabric

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

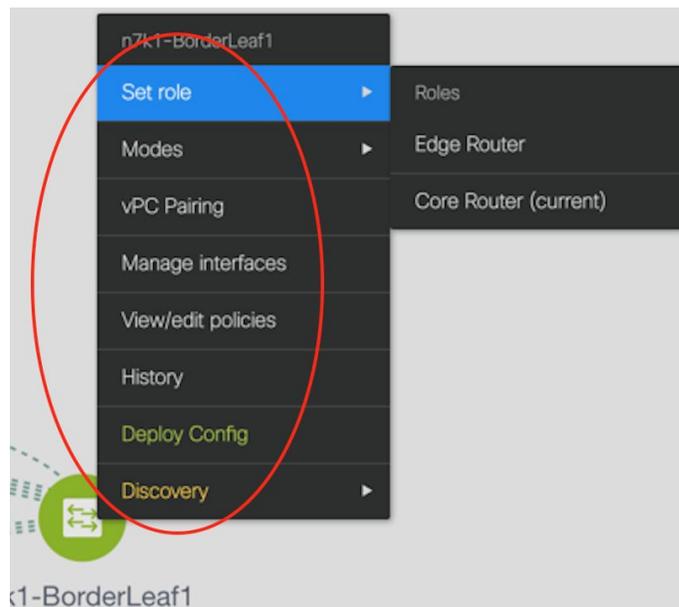
You can also add switches by clicking Tabular View > Switches > + .

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

You can discover multiple switches at the same time. The switches must be properly cabled and connected to the DCNM server and the switch status must be manageable.

The switch discovery process is initiated. The Progress column displays the progress. After DCNM discovers the switch, the screen closes and the fabric screen comes up again. The switch icons are seen at the centre of the fabric screen.

6. Click Refresh topology to view the latest topology view.
7. *External Fabric Switch Settings* - The settings for external fabric switches vary from the VXLAN fabric switch settings. Right-click on the switch icon and set or update switch options.



The options are:

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



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

**Modes** – Active/Operational mode.

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

**Manage Interfaces** – Deploy configurations on the switch interfaces.

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

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

**History** – View per switch deployment history.

**Deploy Config** – Deploy per switch configurations.

**Discovery** - You can use this option to update the credentials of the switch, reload the switch, rediscover the switch, and remove the switch from the fabric.

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

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

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



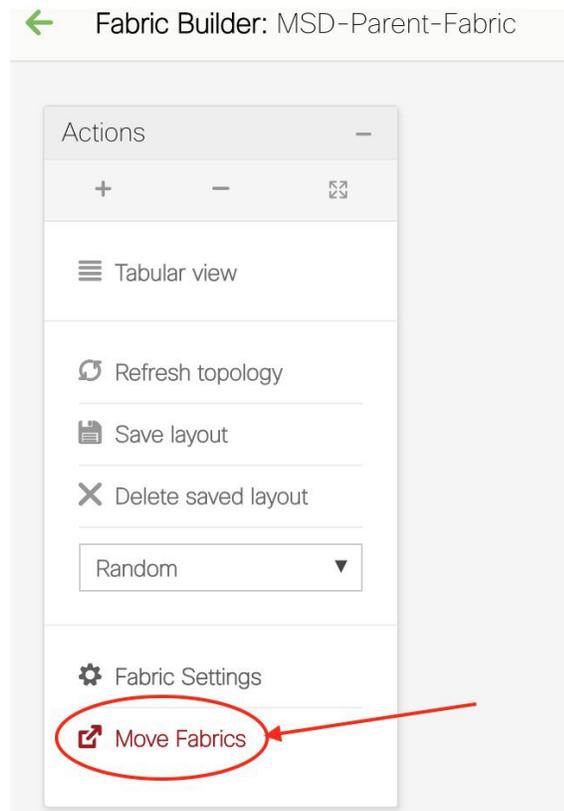
**Note** If a switch in an external fabric does not accept default credentials, you should perform one of the following actions:

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

### Move an External Fabric Under an MSD Fabric

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

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



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

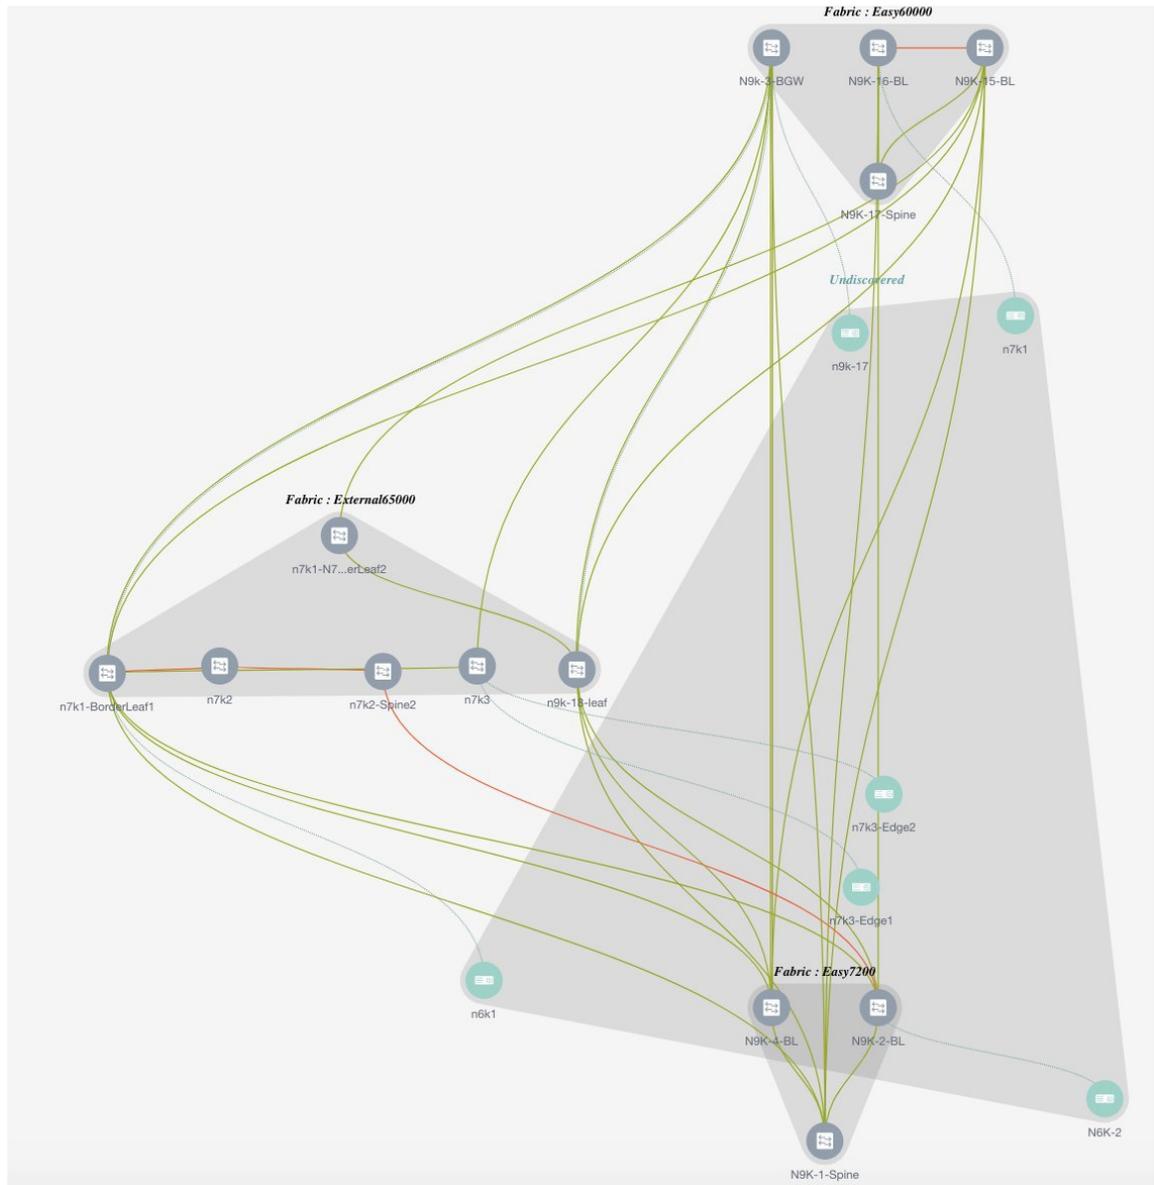
4. Select the radio button next to the external fabric and click Add.

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

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

#### **External Fabric Depiction in an MSD Fabric Topology**

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



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

### External Fabric Switch Operations

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

|   | <input type="checkbox"/>            | Name              | IP Address  | Role       | Serial Number              | Fabric Name   | Fabric Status | Discovery Status | Model     |
|---|-------------------------------------|-------------------|-------------|------------|----------------------------|---------------|---------------|------------------|-----------|
| 1 | <input checked="" type="checkbox"/> | n7k1-BorderLeaf1  | 111.0.0.78  | core ro... | TBM14299900:BorderLeaf1    | External65000 | In-Sync       | ✔ ok             | N7K-C7010 |
| 2 | <input type="checkbox"/>            | n7k1-N7K-1-Bor... | 111.0.0.150 | core ro... | TBM14299900:N7K-1-Borde... | External65000 | In-Sync       | ✔ ok             | N7K-C7010 |

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

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

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




---

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

---

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

### External Fabric Links

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

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

1. Go to the topology screen and click the Tabular view option in the Actions panel, at the left part of the screen.  
The Switches | Links screen comes up.
2. Choose one or more checkboxes and click the Delete icon at the top left.  
The links are deleted.

### Move Neighbor Switch to External Fabric

1. Click Add switches. The Inventory Management screen comes up.
2. Click Move Neighbor Switches tab.
3. Select the switch and click Move Neighbor at the top right part of the screen.  
To delete a neighbor, select a switch and click Delete Neighbor at the top right.

## Discovering New Switches

To discover new switches, perform the following steps:

### Procedure

- 
- Step 1** Power on the new switch in the external fabric after ensuring that it is cabled to the DCNM server.  
Boot the Cisco NX-OS and setup switch credentials.
- Step 2** Execute the **write**, **erase**, and **reload** commands on the switch.  
Choose **Yes** to both the CLI commands that prompt you to choose Yes or No.
- Step 3** On the DCNM UI, choose **Control > Fabric Builder**.  
The **Fabric Builder** screen is displayed. It contains a list of fabrics wherein a rectangular box represents each fabric.
- Step 4** Click **Edit Fabric** icon at the top right part of the fabric box.  
The **Edit Fabric** screen is displayed.
- Step 5** Click the **Bootstrap** tab and update the DHCP information.
- Step 6** Click **Save** at the bottom right part of the Edit Fabric screen to save the settings.
- Step 7** In the Fabric Builder screen, click within the fabric box.  
The fabric topology screen appears.
- Step 8** In the fabric topology screen, from the Actions panel at the left part of the screen, click **Add switches**.  
The Inventory Management screen comes up.
- Step 9** Click the **POAP** tab.  
In an earlier step, the reload command was executed on the switch. When the switch restarts to reboot, DCNM retrieves the serial number, model number, and version from the switch and displays them on the Inventory Management along screen. Also, an option to add the management IP address, hostname, and password are made available. If the switch information is not retrieved, refresh the screen using the Refresh icon at the top right part of the screen.
- Note** At the top left part of the screen, export and import options are provided to export and import the .csv file that contains the switch information. You can pre-provision a device using the import option too.

### Inventory Management ✕

Discover Existing Switches
PowerOn Auto Provisioning (POAP)
Move Neighbor Switches

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

↻ Bootstrap

+ ↻ ↺
\* Admin Password 
\* Confirm Admin Password 
🔒

| <input type="checkbox"/> | Serial Number | Model     | Version | IP Address                                | Hostname                                  |
|--------------------------|---------------|-----------|---------|---|---|
| <input type="checkbox"/> | TBM14299900   | N7K-C7010 | 8.0(1)  | <input style="width: 80px;" type="text"/> | <input style="width: 80px;" type="text"/> |

Close

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

Beginning with Release 11.2(1), you can provision devices in advance. To pre-provision devices, refer to [Pre-provisioning a Device](#), on page 59.

**Step 10** In the **Admin Password** and **Confirm Admin Password** fields, enter and confirm the admin password. This admin password is applicable for all the switches displayed in the POAP window.

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

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

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

Inventory Management ✕

Discover Existing Switches | PowerOn Auto Provisioning (POAP)

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

+ ↻ ↺ \* Admin Password  \* Confirm Admin Password  🔒

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

Close

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

Inventory Management ✕

Discover Existing Switches | PowerOn Auto Provisioning (POAP)

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

+ ↻ ↺ \* Admin Password  \* Confirm Admin Password  🔒

Serial Number Model

No Data available

Discovery Credentials ✕

\*Discovery Username:

\*Discovery Password:

\*Confirm Discovery Password:

OK Clear

Close

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

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

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

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

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

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

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

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

**Step 15** Monitor and check the switch for POAP completion.

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

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

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

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

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

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

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

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

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

## Policy Deployment History for N9k-16-leaf ( SAL18432P6G )

| Entity Name  | Entity Type | Source   | Status  | Status Description    | User  | Time of Completion      |
|--------------|-------------|----------|---------|-----------------------|-------|-------------------------|
| SAL18432P6G  | SWITCH      | DCNM     | SUCCESS | Successfully deployed | admin | 2019-03-29 07:55:25.521 |
| Ethernet1/1  | INTERFACE   | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-03-29 07:43:41.453 |
| Ethernet1/2  | INTERFACE   | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-03-29 07:43:39.642 |
| Ethernet1/3  | INTERFACE   | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-03-29 07:43:37.805 |
| Ethernet1/4  | INTERFACE   | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-03-29 07:43:35.993 |
| Ethernet1/11 | INTERFACE   | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-03-29 07:43:34.18  |
| Ethernet1/10 | INTERFACE   | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-03-29 07:43:32.562 |
| Ethernet1/13 | INTERFACE   | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-03-29 07:43:30.551 |

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

## Command Execution Details for N9k-16-leaf ( SAL18432P6G )

| Config                             | Status  | CLI Response   |
|------------------------------------|---------|--|
| interface ethernet1/2              | SUCCESS |  |
| shutdown                           | SUCCESS |  |
| switchport                         | SUCCESS |  |
| switchport mode trunk              | SUCCESS |  |
| switchport trunk allowed vlan none | SUCCESS |  |
| mtu 9216                           | SUCCESS |  |
| spanning-tree port type edge trunk | SUCCESS | Edge port type (portfast) should only be enabled on p... |
| shutdown                           | SUCCESS |  |

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

Up to this step, the POAP is completed with basic settings. All the interfaces are set to trunk ports. You must setup interfaces through the **Control > Interfaces** option for any additional configurations, but not limited to the following:

- vPC pairing.
- Breakout interfaces  
Support for breakout interfaces is available for 9000 Series switches.
- Port channels, and adding members to ports.

**Note** After discovering a switch (new or existing), at any point in time you can provision configurations on it again through the POAP process. The process removes existing configurations and provision new configurations. You can also deploy configurations incrementally without invoking POAP.

## Pre-provisioning a Device

In DCNM 11.2, you can provision devices in advance.




---

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

---

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

Note the following guidelines:

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

The examples of the values in the **Data** field are as follows:

- {"modulesModel": ["N9K-C93180LC-EX"], "gateway": "10.1.1.1/24"}
- {"modulesModel": ["N9K-C93180LC-EX"], "breakout": "interface breakout module 1 port 1 map 10g-4x", "portMode": "hardware profile portmode 4x100G+28x40G", "gateway": "172.22.31.1/24"}

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

## Procedure

---

- Step 1** 1. Click **Control > Fabric Builder**.
- The **Fabric Builder** screen is displayed.
- Step 2** Click within the fabric box.
- Step 3** From the Actions panel, click the **Add switches** option.
- The **Inventory Management** screen is displayed.
- Step 4** Click the **POAP** tab.
- Step 5** In the **POAP** tab, do the following:
- Click + from the top left part of the screen.  
The Add a new device screen comes up.
  - Fill up the device details as shown in the screenshot.
  - Click **Save**.

Inventory Management

Discover Existing Switches | PowerOn Auto Provisioning (POAP) | Move Neighbor Switches

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

+ | \* Admin Password | \* Confirm Admin Password | Bootstrap

Serial Number | Model | Version | IP Address | Hostname

Add a new device to pre-provisioning

\*Serial Number: SN

\*Model: N9K-3455

\*Version: 7.0(2)

\*IP Address: 10.1.1.1

\*Hostname: leaf1

\*Data: {"modulesModel": ["N9K-EX"]}   
*JSON Object which contains model name of the Modules*  
*Eg: {"modulesModel": ["N9K-EX"]}*

Save | Clear

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

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

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

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

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

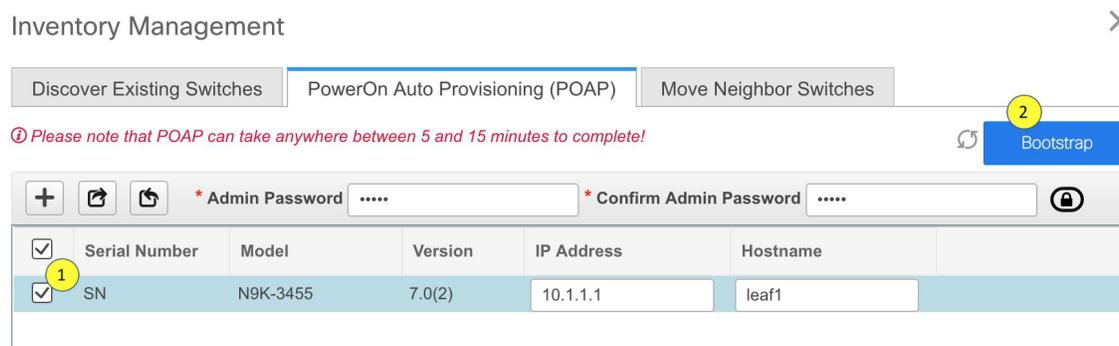
Add new devices' information in the .csv file with all the mandatory fields (SerialNumber, Model, version, IpAddress, Hostname and Data fields [JSON Object]).

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

|   | A                            | B                     | C                        | D                        | E         | F  | G |
|---|------------------------------|-----------------------|--------------------------|--------------------------|-----------|--|---|
| 1 | #SerialNumber(Eg:FD01344GH5) | #Model(Eg:N9k-C9236C) | #Version(Eg:7.0(3)12(3)) | #IPAddress of the device | #HostName | #Data(JSON Field contains model name of the modules) |   |
| 2 | Serial Number                | Model                 | Version                  | IP Address               | Hostname  | Data   |   |
| 3 | FDO21331SND                  | N9K-93180YC-EX        | 7.0(3)15(2)              | 1.1.1.1                  | leaf1     | {"modulesModel":["N9K-93180YC-EX"]}                  |   |
| 4 | FDO21351N3X                  | N9K-C9236C            | 7.0(3)14(1)              | 11.1.1.1                 | spine1    | {"modulesModel":["N9K-C9236C"]}                      |   |
| 5 | FDO21491A5K                  | N9K-C93240YC-FX2      | 7.0(3)17(3)              | 12.1.1.1                 | leaf2     | {"modulesModel":["N9K-C93240YC-FX2"]}                |   |
| 6 |                              |                       |                          |                          |           |  |   |

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

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



The leaf1 device appears in the fabric topology.

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

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

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

Create intra-link policy using one of the templates:

- **int\_pre\_provision\_intra\_fabric\_link** to automatically generate intra fabric interface configuration with DCNM allocated IP addresses
- **int\_intra\_fabric\_unnum\_link\_11\_1** if you are using unnumbered links
- **int\_intra\_fabric\_num\_link\_11\_1** if you want to manually assign IP addresses to intra-links

Click **Save & Deploy**.

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

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

For more information, see [Return Material Authorization \(RMA\)](#), on page 161.

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

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

## Creating a vPC Setup in the External Fabric

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

### Procedure

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

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

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

Select vPC peer for N5596-37 ✕

| 1                                | Switch name | Recommended | Reason                                    |
|----------------------------------|-------------|-------------|---|
| <input checked="" type="radio"/> | N5648-38    | true        | Switches are connected and have same role |

Note : Peer one = N5596-37,Peer two = N5648-38

vPC Pair Template

No Policy

vpc\_pair 2

No Policy

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

vPC Pair Template  ▼

vPC Domain | vPC Peerlink

\* vPC Domain ID  ? vPC

\* Peer-1 vPC Keep-alive Local IP Address  ? IP a

\* Peer-1 vPC Keep-alive Peer IP Address  ? IP a

\* Peer-2 vPC Keep-alive Local IP Address  ? IP a

\* Peer-2 vPC Keep-alive Peer IP Address  ? IP a

\* vPC Keep-alive VRF Name  ? Narr

vPC+  ? Check this if it's a vPC+ topology

\* Fabricpath switch id  ? Fabr

Configure VTEPS  ? Check this to configure NVE source loopbac

\* NVE interface  ? NVE

\* Peer 1 NVE source loopback interface  ? Peer

**vPC Domain tab:** Enter the vPC domain details.

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

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

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

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

| vPC Domain                               | vPC Peerlink  |
|--|---|
|  | * vPC Domain ID <input type="text" value="3"/> ? vPC                              |
| * Peer-1 vPC Keep-alive Local IP Address | <input type="text" value="10.10.10.2"/> ? IP ac                                   |
| * Peer-1 vPC Keep-alive Peer IP Address  | <input type="text" value="10.10.10.3"/> ? IP ac                                   |
| * Peer-2 vPC Keep-alive Local IP Address | <input type="text" value="10.10.10.4"/> ? IP ac                                   |
| * Peer-2 vPC Keep-alive Peer IP Address  | <input type="text" value="10.10.10.5"/> ? IP ac                                   |
| * vPC Keep-alive VRF Name                | <input type="text" value="vPC-VRF"/> ? Nam  |
| vPC+                                     | <input type="checkbox"/> ? Check this if it's a vPC+ topology                     |
| Fabricpath switch id                     | <input type="text"/> ? Fabr   |
| Configure VTEPS                          | <input checked="" type="checkbox"/> ? Check this to configure NVE source loopback |
| * NVE interface                          | <input type="text" value="nve1"/> ? NVE   |
| * Peer 1 NVE source loopback interface   | <input type="text" value="4"/> ? Peer   |
| * Peer 2 NVE source loopback interface   | <input type="text" value="4"/> ? Peer   |

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

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

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

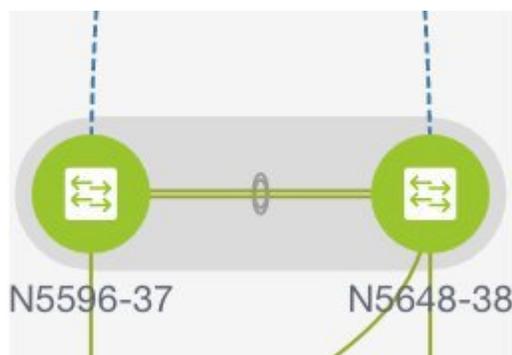
vPC Domain

vPC Peerlink

|  |  |   |
|--|--|---|
| Peer-1 Peerlink Port-Channel ID          | <input type="text" value="10"/>          | ? Peer-1  |
| Peer-2 Peerlink Port-Channel ID          | <input type="text" value="10"/>          | ? Peer-2  |
| Peer-1 Peerlink Member Interfaces        | <input type="text" value="e1/5,eth1/7"/> | ? A list of                                     |
| Peer-2 Peerlink Member Interfaces        | <input type="text" value="e1/5,eth1/7"/> | ? A list of                                     |
| Port Channel Mode                        | <input type="text" value="active"/>      | ? Channel                                       |
| Switch Port Mode                         | <input type="text" value="trunk"/>       | ? Switch  |
| Peer-1 Peerlink Port Channel Description | <input type="text"/>                     | ? Add de  |
| Peer-2 Peerlink Port Channel Description | <input type="text"/>                     | ? Add de  |
| Enable VPC Peerlink Port Channel         | <input checked="" type="checkbox"/>      | ? Uncheck to disable the vPC Peerlink port-chan |
| * Trunk Allowed Vlans                    | <input type="text" value="none"/>        | ? Trunk A                                       |
| Native Vlan                              | <input type="text" value="1"/>           | ? Native  |

**Step 3** Click **Save**.

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



To update vPC setup details, do the following:

- a. Right-click a vPC switch and choose vPC Pairing.  
The **vPC peer** dialog box comes up.
- b. Update the field(s) as needed.  
When you update a field, the **Unpair** icon changes to **Save**.
- c. Click **Save** to complete the update.

## Undeploying a vPC Setup in the External Fabric

### Procedure

---

- Step 1** Right-click a **vPC** switch and choose **vPC Pairing**.  
The vPC peer screen comes up.
- Step 2** Click **Unpair** at the bottom right part of the screen.  
The vPC pair is deleted and the fabric topology window appears.
- Step 3** Click **Save & Deploy**.  
The **Config Deployment** dialog box appears.
- Step 4** (Optional) Click the value under the **Preview Config** column.  
View the pending configuration in the **Config Preview** dialog box. The following configuration details are deleted on the switch when you unpair: vPC feature, vPC domain, vPC peerlink, vPC peerlink member ports, loopback secondary IPs, and host vPCs. However, the host vPCs and port channels are not removed. Delete these port channels from the **Interfaces** window if required.

**Note** Resync the fabric if it is out of sync.

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

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

---

## Multi-Site Domain for VXLAN BGP EVPN Fabrics

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

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

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

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

**Note**

- vPC support is added for BGWs in the DCNM 11.1(1) release.
- The MSD feature is unsupported on the switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images.
- The VXLAN OAM feature in Cisco DCNM is only supported on a single fabric or site.

A few fabric-specific terms:

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

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

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

### Fabric and Switch Instance Variables

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

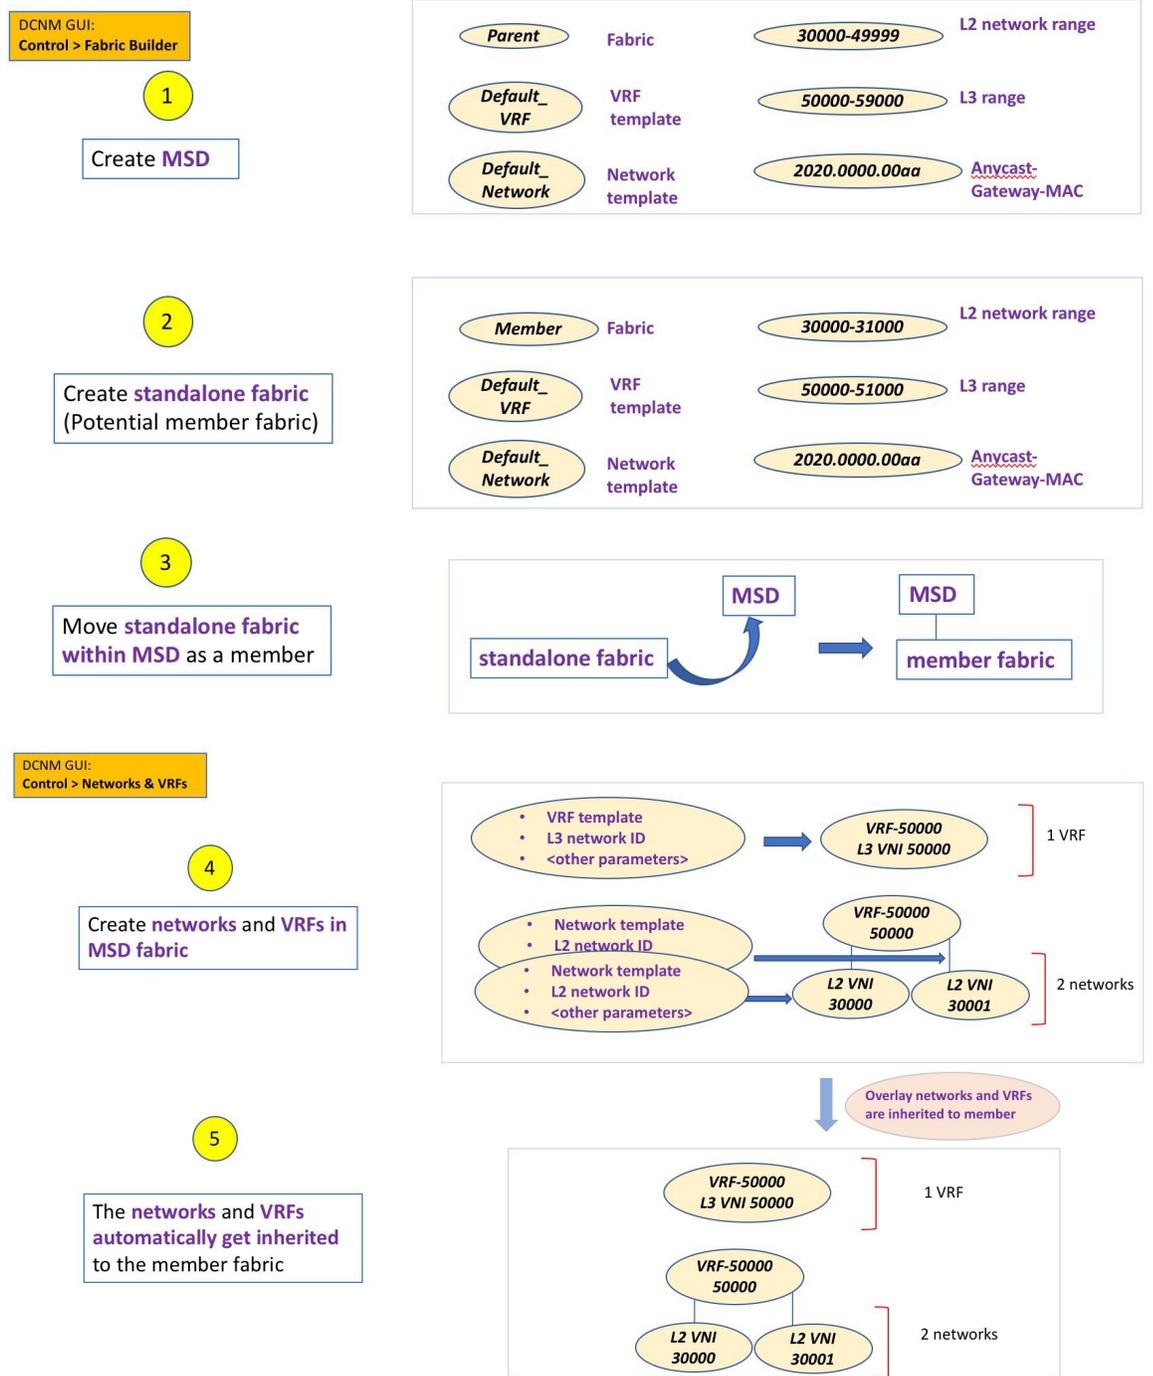
Fabric instance values can only be edited or updated in the fabric context from the VRFs and Networks window. The appropriate fabric should be selected in the **SCOPE** drop-down list to edit the fabric instance values. Some of the examples of fabric instance variables are BGP ASN, Multicast group per network or VRF, etc. For information about editing multicast group address, see [Editing Networks in the Member Fabric, on page 126](#).

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

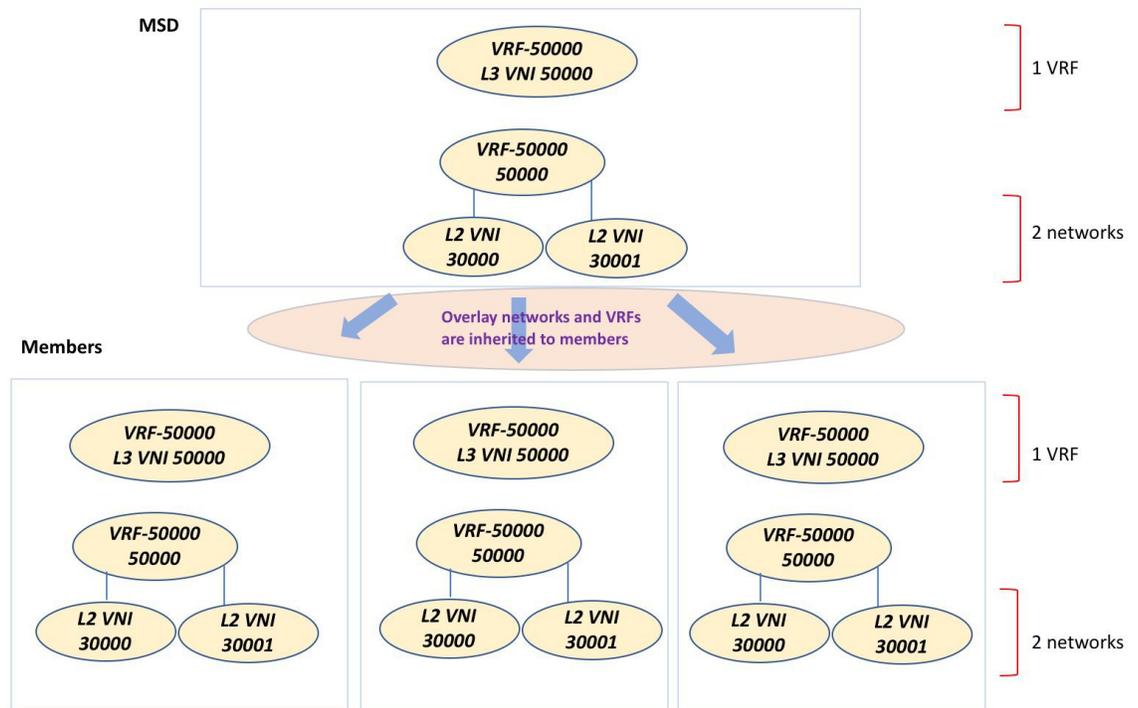
## MSD and Member Fabric Process Flow

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

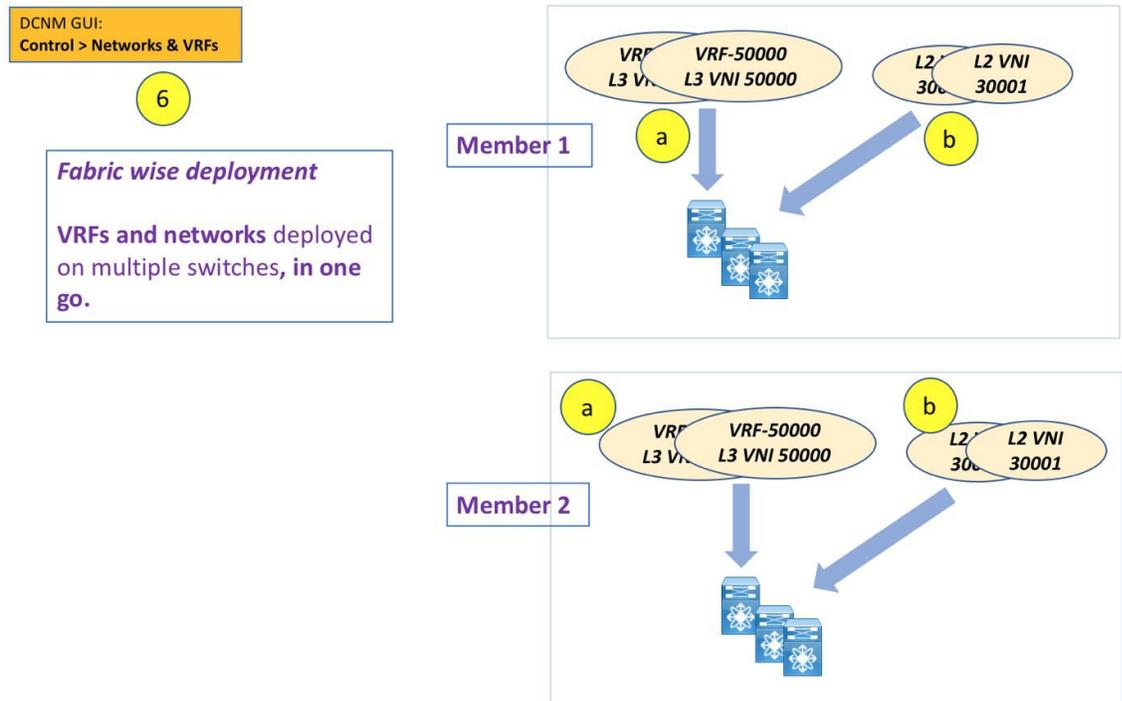
A high-level flow chart of the MSD and member fabric creation and MSD-to-member fabric inheritance process:



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



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



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



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

Upcoming sections in the document explain the following:

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

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

The process is explained in two steps:

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

## Creating an MSD Fabric

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

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



### Fabric Builder

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

Create Fabric

Fabrics (4)

|  |  |  |   |
|--|--|--|---|
| <p>External65000  </p> <p>Type: External<br/>ASN: 650000</p> | <p>Easy60000  </p> <p>Type: Switch_Fabric<br/>ASN: 60000<br/>Replication Mode: Multicast<br/>Technology: VXLANFabric</p> | <p>Easy7200  </p> <p>Type: Switch_Fabric<br/>ASN: 7200<br/>Replication Mode: Multicast<br/>Technology: VXLANFabric</p> | <p>MSD  </p> <p>Type: MSD<br/>Member Fabrics: External65000, Easy7200</p> |
|--|--|--|---|

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

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

**Fabric Name** - Enter the name of the fabric.

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

## Add Fabric



\* Fabric Name :

\* Fabric Template : MSD\_Fabric\_11\_1 ▼

---

General DCI Resources

\* Layer 2 VXLAN VNI Range  ? *Overlay Network Identifier Range (Min:1, Max:16777214)*

\* Layer 3 VXLAN VNI Range  ? *Overlay VRF Identifier Range (Min:1, Max:16777214)*

\* VRF Template  ▼ ? *Default Overlay VRF Template For Leafs*

\* Network Template  ▼ ? *Default Overlay Network Template For Leafs*

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

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

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

\* Multisite Routing Loopback Id  ? *0-512*

The fields in the screen are explained:

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

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

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

**VRF Template** - Default VRF template.

**Network Template** - Default network template.

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

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

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

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

3. Click the **DCI** tab.

## Add Fabric

\* Fabric Name :

\* Fabric Template :

General DCI Resources

DCI Subnet IP Range  ? Address range

Subnet Target Mask  ? Target Mask

\* Multi-Site Overlay IFC Deploy Met...  ? Manual/GU

MS Route Server List  ? Multi-Site R

BGP ASN of Route Server(s) one for...  ? 1-42949672

Multi-Site Underlay IFC Deploy Optio...  ? Clear for Manual, Check for Auto

The fields are:

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

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

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

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

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

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

4. Click the **Resources** tab.

General DCI Resources

\* MultiSite Routing Loopback IP Range 10.10.0.0/22 ? Typically Loopback100 IP Address Range

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

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

5. Click **Save**.

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

← Fabric Builder MSD-Parent-Fabric Save & Deploy

Actions

- + -
- Tabular view
- 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.

Fabrics (5)

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

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

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

### Creating and Moving a New Fabric Under the MSD Fabric as a Member

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

New fabric creation is explained in the Easy Fabric creation process. In the MSD document, fabric movement is covered. However, some pointers about a standalone (potential member) fabric:

| Field                                 | Value                    | Help Text  |
|---------------------------------------|--------------------------|--|
| Static Underlay IP Address Allocation | <input type="checkbox"/> | Checking this will disable Dynamic Underlay IP Address Allocations |
| * Underlay Routing Loopback IP Range  | 10.2.0.0/22              | Typically Loopback0 IP Address Range                               |
| * Underlay VTEP Loopback IP Range     | 10.3.0.0/22              | Typically Loopback1 IP Address Range                               |
| * Underlay RP Loopback IP Range       | 10.254.254.0/24          | Anycast or Phantom RP IP Address Range                             |
| * Underlay Subnet IP Range            | 10.4.0.0/16              | Address range to assign Numbered and Peer L...                     |
| * Layer 2 VXLAN VNI Range             | 30000-49000              | Overlay Network Identifier Range (Min:1, Max:16...                 |
| * Layer 3 VXLAN VNI Range             | 50000-59000              | Overlay VRF Identifier Range (Min:1, Max:1677...                   |
| * Network VLAN Range                  | 2300-2999                | Per Switch Overlay Network VLAN Range (Min:2...                    |

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

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

- You must update one range of values at a time. If you want to update more than one range of values, do it in separate instances. For example, if you want to update L2 and L3 ranges, you should do the following:
  1. Update the L2 range and click **Save**.
  2. Click the **Edit Fabric** option again, update the L3 range and click **Save**.

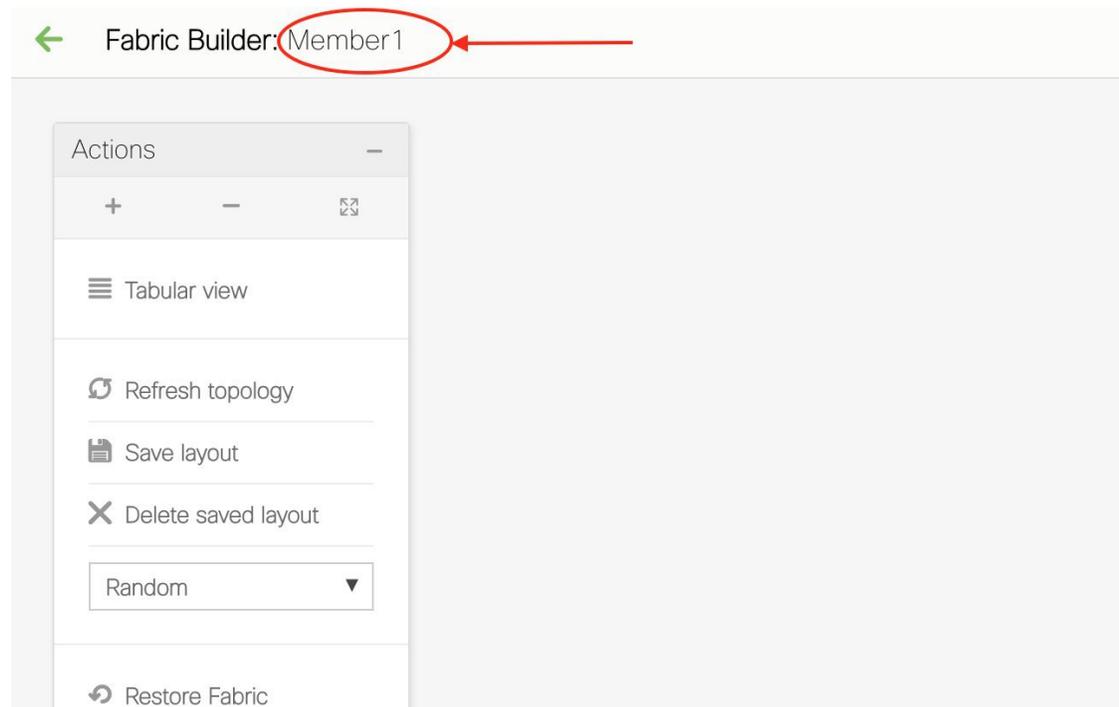
Ensure that the **Anycast Gateway MAC**, the **Network Template** and the **VRF Template** field values are the same as the MSD fabric. Else, member fabric movement to the MSD fail.

Other pointers:

- Ensure that the Anycast Gateway MAC, the Network Template and the VRF Template field values are the same as the MSD fabric. Else, member fabric movement to the MSD fail.
- The member fabric should have a Site ID configured and the Site ID must be unique among the members.
- The BGP AS number should be unique for a member fabric.
- The underlay subnet range for loopback0 should be unique.
- The underlay subnet range for loopback1 should be unique.

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

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



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



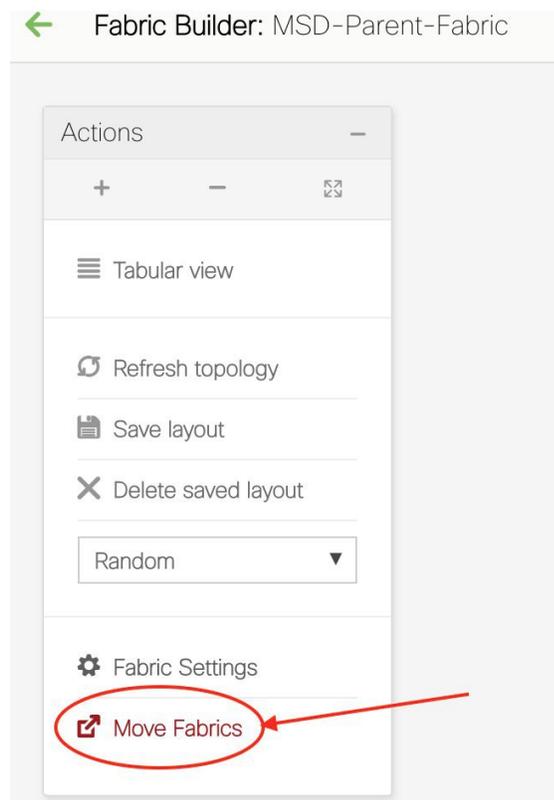
### Moving the Member1 Fabric Under MSD-Parent-Fabric

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

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

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

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



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

## Move Fabric



Selected 0 / Total 2

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

Add

Remove

Cancel

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

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

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

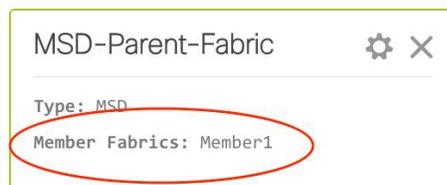
Immediately, a message appears at the top of the screen indicating that the *Member1* fabric is now associated with the MSD fabric *MSD-Parent-Fabric*. Now, the MSD-Parent-Fabric fabric page appears again.

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



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

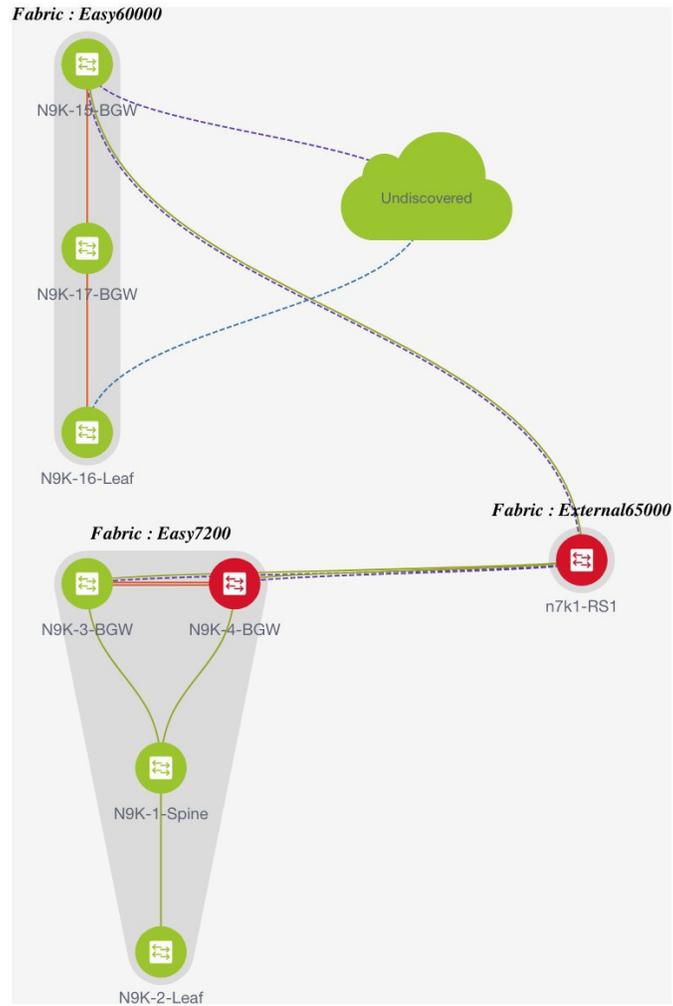
You can see that *Member1* is now added to MSD fabric and is displayed in the **Member Fabrics** field.



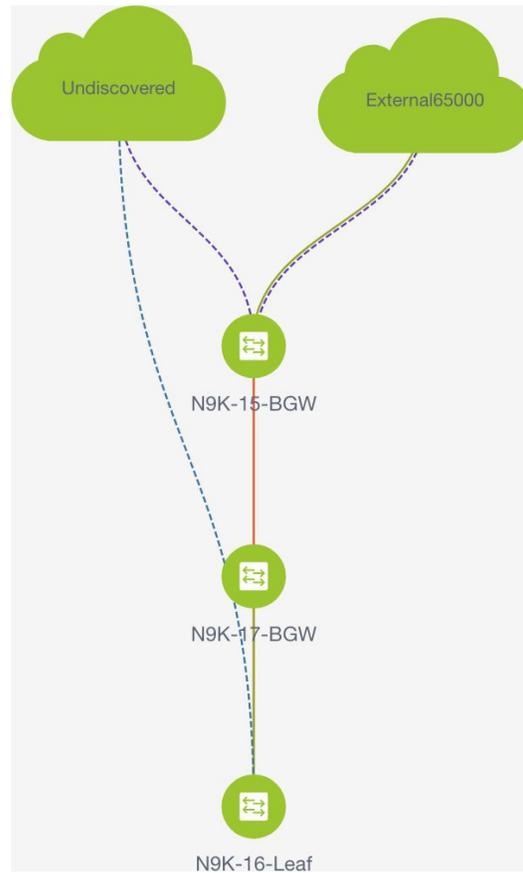
### MSD Fabric Topology View Pointers

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

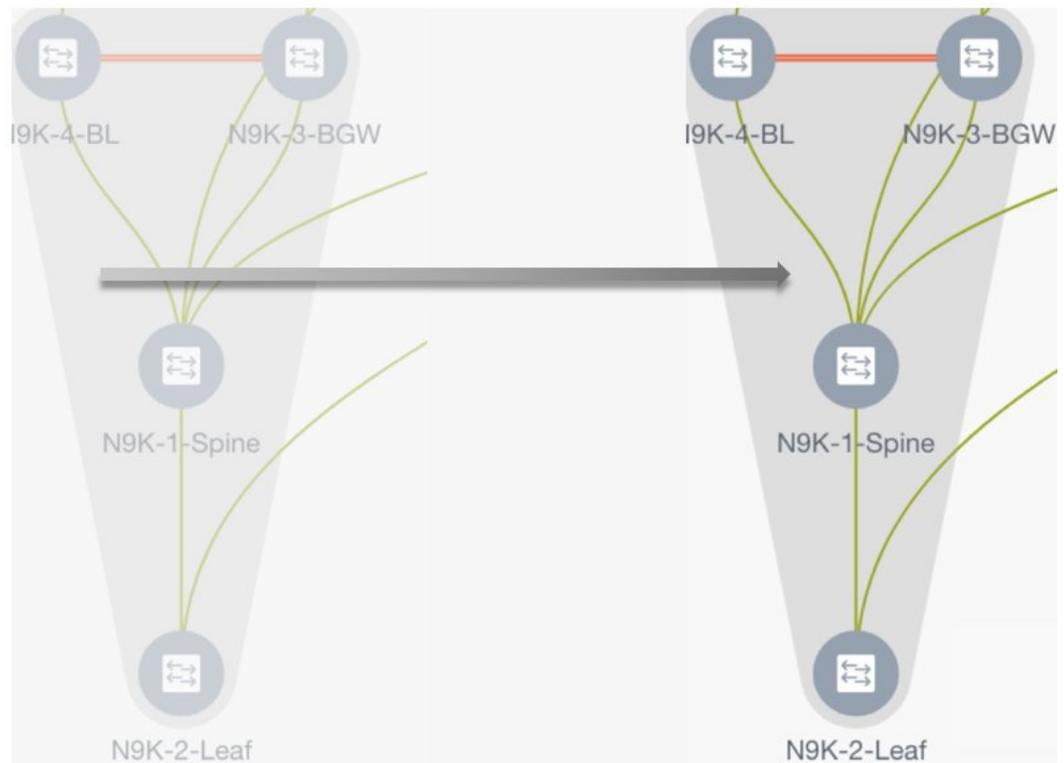
All links are displayed, including intra-fabric links and Multi-Site (underlay and overlay), and VRF Lite links to remote fabrics.



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



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



### Adding and Editing Links

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

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

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

### Creating and Deploying Networks and VRFs in an MSD Fabric

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

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

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

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



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

For example, if you create networks 30000 and 30001 for an MSD that contains two member fabrics, the networks are automatically created for the member fabrics and are available for deployment.

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

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

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

### Creating Networks in the MSD Fabric

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

The Networks screen comes up.

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

The screenshot shows the Cisco Data Center Network Manager interface. At the top, it says "Data Center Network Manager" and "SCOPE: bgp2". Below that, there are navigation tabs for "Network / VRF Selection" and "Network / VRF Deployment". A "Fabric Selected: bgp2" label is present. The main area is titled "Networks" and shows a table with the following data:

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

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

/ VRF Selection > Network / VRF Deployment > 2 Continue

## Select a Fabric

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

MSD-Parent-Fabric 1 ▼

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

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

Fabric Selected: MSD-Parent-Fabric

Networks Selected 0 / Total 0 ↻ ⚙

+
✎
✕
↻
↺

Show All ▼ ⌵

| <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

\* Network Name

\* VRF Name  +

Layer 2 Only

\* Network Template

\* Network Extension Template

VLAN ID

---

▼ Network Profile

General

Advanced

IPv4 Gateway/NetMask  ? example 192.0.2.1/24

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

Vlan Name  ?

Interface Description  ?

MTU for L3 interface  ? [68-9216]

The fields in this screen are:

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

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

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



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

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

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

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

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

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

**General** tab

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

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

**VLAN Name** - Enter the VLAN name.

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

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

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

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

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

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

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

A sample of the Create Network screen:

## Create Network

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

Network Profile
 

- General
  - IPv4 Gateway/NetMask: 20.10.1.1/24 ? example 192.0.2.1/24
  - IPv6 Gateway/Prefix:  ? example 2001:db8::1/64
  - Vlan Name: Drill ?
  - Interface Description:  ?
  - MTU for L3 interface:  ? [68-9216]
  - IPv4 Secondary GW1: 20.10.2.1/24 ? example 192.0.2.1/24
  - IPv4 Secondary GW2: 20.10.3.1/24 ? example 192.0.2.1/24

**Create Network**

## Advanced tab:

Network Profile
 

- Advanced
  - ARP Suppression:  ?
  - \* DHCPv4 Server 1: 20.20.20.1 ? DHCP Relay IP
  - DHCPv4 Server 2: 20.20.30.1 ? DHCP Relay IP
  - \* DHCPv4 Server VRF: Foo ?
  - Loopback ID for DHCP Relay interface: 4 ?
  - Routing Tag: 12345 ? [0-4294967295]
  - TRM Enable:  ? Enable Tenant Routed Multicast
  - L2 VNI Route-Target Both Enable:  ?

**Create Network**

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

Fabric Selected: MSD-Parent-Fabric

Networks Selected 1 / Total 1

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

## Editing Networks in the MSD Fabric

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

Fabric Selected: MSD-Parent-Fabric

Networks Selected 1 / Total 1

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

The Edit Network screen comes up.

### Edit Network

▼ Network Information

\* Network ID

\* Network Name

\* VRF Name

Layer 2 Only

\* Network Template

\* Network Extension Template

VLAN ID

---

▼ Network Profile

General

Advanced

IPv4 Gateway/NetMask  ? example 192.0.2.1/24

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

Vlan Name  ?

Interface Description  ?

MTU for L3 interface  ? [68-9216]

IPv4 Secondary GW1  ? example 192.0.2.1/24

IPv4 Secondary GW2  ? example 192.0.2.1/24

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

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

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

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

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

The Networks screen comes up.

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

The screenshot shows the Cisco Data Center Network Manager interface. At the top, the breadcrumb navigation is "Network / VRF Selection" > "Network / VRF Deployment". The "SCOPE" dropdown is set to "bgp2". Below the navigation, the text "Fabric Selected: bgp2" is displayed. The main area is titled "Networks" and shows a table with one network entry selected. The table has columns for Network Name, Network ID, VRF Name, IPv4 Gateway/Subnet, IPv6 Gateway/Prefix, Status, and VLAN ID.

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

### Editing Networks in the Member Fabric

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

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

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



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

### Deleting Networks in the MSD and Member Fabrics

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

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



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

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

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

### Creating VRFs in the MSD Fabric

1. From the MSD fabric's Networks page, click the **VRF View** button at the top right part of the screen to create VRFs.
  - a. Choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

SCOPE: bgp2

Network / VRF Selection > Network / VRF Deployment

Fabric Selected: bgp2

VRFs Selected 1 / Total 1

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

- b. Choose the MSD fabric (*MSD-Parent-Fabric*) from the drop-down box and click **Continue**. The Networks page comes up.
- c. Click **VRF View** at the top right part of the Networks page].

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

Fabric Selected: MSD-Parent-Fabric

VRFs Selected 0 / Total 0

| <input type="checkbox"/> | VRF Name | VRF ID | Status |
|--------------------------|----------|--------|--------|
| No data available        |          |        |        |

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

The fields in this screen are:

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

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



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

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

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

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

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

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

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

**TRM Enable** – Select the checkbox to enable TRM.

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

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

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

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

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

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

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

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

A sample screenshot:

Create VRF
✕

---

▼ VRF Information

\* VRF ID

\* VRF Name

\* VRF Template

\* VRF Extension Template

---

▼ VRF Profile

General

Advanced

VRF Vlan Name  ?

VRF Intf Description  ?

VRF Description  ?

**Advanced** tab:

▼ VRF Profile

General

Advanced

**Routing Tag**  ? [0-4294967295]

**Redistribute Direct Route Map**  ?

**Max BGP Paths**  ? [1-64]

**Max iBGP Paths**  ? [1-64]

**TRM Enable**  ? Enable Tenant Routed Multicast

**Is RP External**  ? Is RP external to the fabric?

**RP Address**  ? IPv4 Address

**RP Loopback ID**  ? 0-1023

**Underlay Mcast Add...**  ? IPv4 Multicast Address

**Overlay Mcast Groups**  ? 224.0.0.0/8 to 239.255.255.255/8

**Enable IPv6 link-loc...**  ? Enables IPv6 link-local Option under VRF SVI

**Advertise Host Routes**  ? Flag to Control Advertisement of /32 and /128 Routes to Edge Routers

**Advertise Default Route**  ? Flag to Control Advertisement of Default Route Internally

[Create VRF](#)

## 5. Click Create VRF.

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

Fabric Selected: MSD-Parent-Fabric

VRFs Selected 1 / Total 1

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

## Editing VRFs in the MSD Fabric

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

Fabric Selected: MSD-Parent-Fabric

VRFs Selected 1 / Total 1

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

The Edit VRF screen comes up.

Edit VRF
✕

---

▼ VRF Information

\* VRF ID

\* VRF Name

\* VRF Template

VRF Extension Template

---

▼ VRF Profile

General

Advanced

VRF Vlan Name  ?

VRF Intf Description  ?

VRF Description  ?

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

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

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

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

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

The screenshot shows the Cisco Data Center Network Manager interface. At the top, the breadcrumb is "Network / VRF Selection > Network / VRF Deployment". The "SCOPE" dropdown is set to "bgp2". Below this, the "Fabric Selected: bgp2" is displayed. The "VRFs" section shows a table with one VRF listed:

| VRF Name    | VRF ID | Status |
|-------------|--------|--------|
| MyVRF_50000 | 50000  | NA     |

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

Fabric Selected: Member1

VRFs Selected 0 / Total 1

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

### Deleting VRFs in the MSD and Member Fabrics

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

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



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

### Editing VRFs in the Member Fabric

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

### Deleting VRFs in the Member Fabric

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

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

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

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

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



**Note** The deployment (and undeployment) of networks and VRFs in member fabrics are the same as explained for standalone fabrics. Refer [Creating and Deploying Networks and VRFs](#) .

## Removing a Fabric From an MSD

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

### Before you begin

Make sure that there are no VRFs deployed on the border switches in the fabric that you want to remove. For more information, see [Deployment and Undeployment of Networks and VRFs in Member Fabrics](#), on page 132.



---

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

---

### Procedure

---

- Step 1** From the **Fabric Builder** window, click an MSD fabric.
- Step 2** Click **Move Fabric** in the **Actions** menu.
- Step 3** In the **Move Fabric** window, select the respective radio button of the fabric that you want to remove and click **Remove**.
- In the fabric removal notification window, click **Close**.
- Step 4** Click **Save & Deploy** for the MSD in the **Fabric Builder** window.
- Step 5** Click **Deploy Config** in the **Config Deployment** window.
- Click **Close**.
- Step 6** Navigate to the fabric that you removed from MSD and click **Save & Deploy**.
- Step 7** Click **Deploy Config** in the **Config Deployment** window.
- Click **Close**.
- 

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

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

- The MSD fabric inherits the networks and VRFs of the standalone fabric that do not exist in the MSD fabric. These networks and VRFs are in turn inherited by the member fabrics.
- The newly created member fabric inherits the networks and VRFs of the MSD fabric (that do not exist in the newly created member fabric).

- If there are conflicts between the standalone and MSD fabrics, validation ensures that an error message is displayed. After the updation, when you move the member fabric to the MSD fabric, the move will be successful. A message comes up at the top of the page indicating that the move is successful.

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

## SSH Key RSA Handling

### Bootstrap scenario

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

### Greenfield and Brownfield scenarios

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

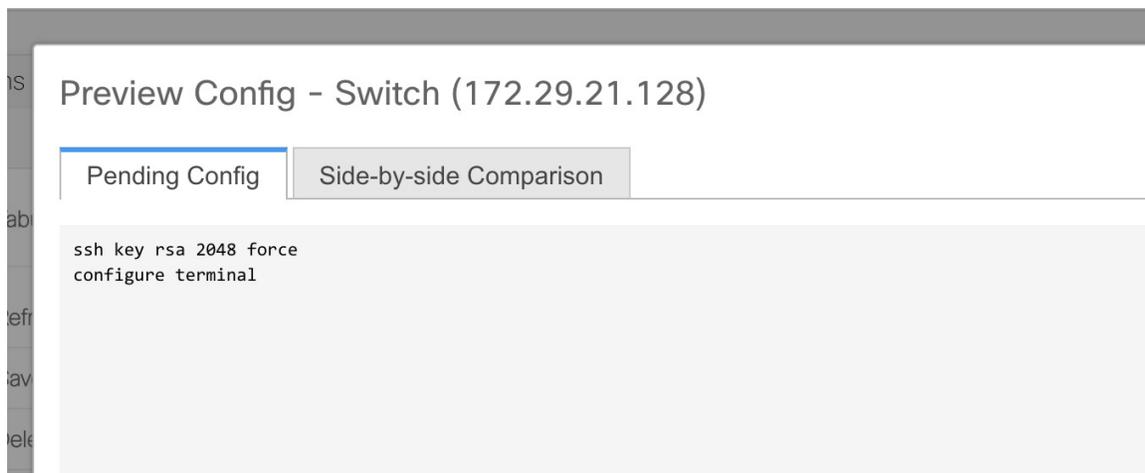
However, on Cisco Nexus 9000 Releases 9.3(1) and 9.3(2), the **ssh key rsa key-length force** command fails while the device is booting up during the ASCII replay process. For more information, refer [CSCvs40704](#).

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

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

**With Strict Config-Compliance mode:**

- Delete the Policy Template Instance (PTI) that has the **ssh key rsa 2048 force** command by clicking **View/Edit Policies** in the **Tabular View** of the **Fabric Builder** window.
- Create a new PTI with the **ssh key rsa 2040 force** command by clicking **View/Edit Policies**.

**Without Strict Config-Compliance mode:**

- Delete the PTI with the **ssh key rsa 2048 force** command in the intent by clicking **View/Edit Policies** in the **Tabular View** of the **Fabric Builder** window.
- Create a switch\_freeform PTI with the **ssh key rsa 2040 force** command in the intent to match the Out-Of-Band change from the device.

## Switch Operations

To view various options, right-click on switch:

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

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



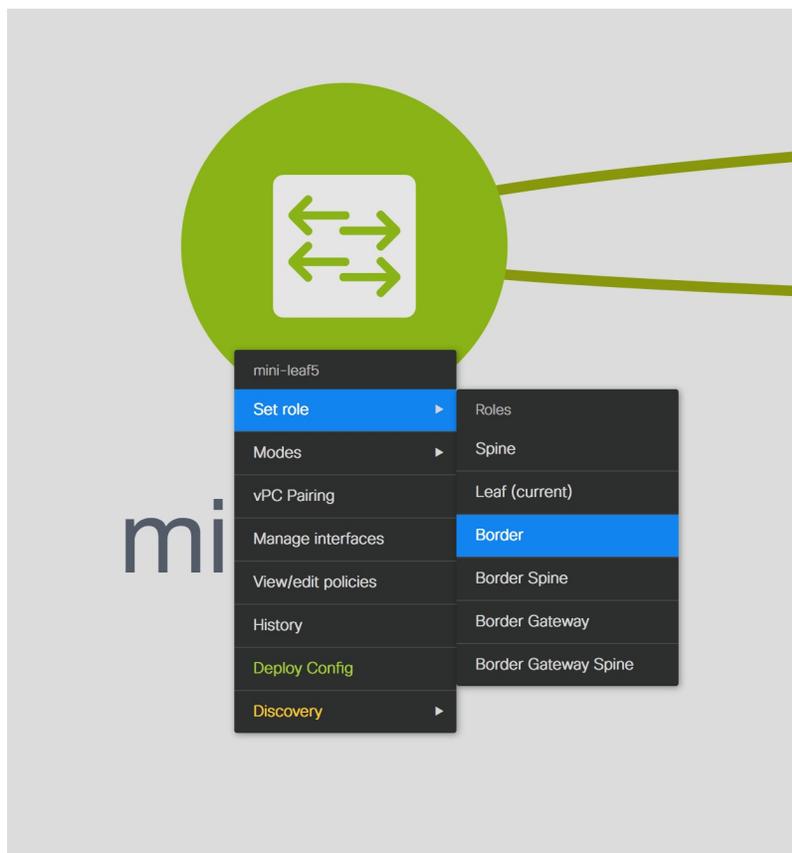
---

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

---

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

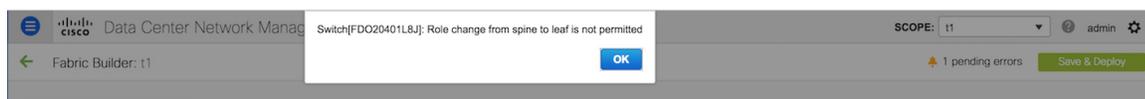
- Leaf to Border
- Border to Leaf
- Leaf to Border Gateway
- Border Gateway to Leaf
- Border to Border Gateway
- Border Gateway to Border
- Spine to Border Spine
- Border Spine to Spine
- Spine to Border Gateway Spine
- Border Gateway Spine to Spine
- Border Spine to Border Gateway Spine
- Border Gateway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



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

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

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



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

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

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

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



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

## Fabric Multi Switch Operations

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

|   | <input type="checkbox"/>            | Name             | IP Addr... | Role                 | Serial Number | Fabric N... | Fabric... ▲ | Di...                                  | Model           | Softwa...   | Last Updatec  |
|---|-------------------------------------|------------------|------------|----------------------|---------------|-------------|-------------|--|-----------------|-------------|---------------|
| 1 | <input checked="" type="checkbox"/> | N9K-16-Leaf      | 111.0.0.96 | leaf                 | SAL18432P6G   | Easy60000   | In-Sync     | <input checked="" type="checkbox"/> ok | N9K-C9396PX     | 7.0(3)17(4) | 6 minutes ago |
| 2 | <input type="checkbox"/>            | N9K-17-BGW-Spine | 111.0.0.97 | border gateway spine | FDO20401LEJ   | Easy60000   | In-Sync     | <input checked="" type="checkbox"/> ok | N9K-C93180YC-EX | 7.0(3)17(3) | 6 minutes ago |
| 3 | <input type="checkbox"/>            | N9K-15-BGW-Spine | 111.0.0.95 | border gateway spine | FDO20401LB4   | Easy60000   | Out-of-sync | <input checked="" type="checkbox"/> ok | N9K-C93180YC-EX | 7.0(3)17(4) | 6 minutes ago |

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

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

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



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

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

## Fabric Links

You can add links between border switches of different fabrics (inter-fabric links) or between switches in the same fabric (intra-fabric links). You can only create an inter-fabric connection (IFC) for a switch that is managed by DCNM.

There are scenarios where you might want to define links between switches before connecting them physically. The links could be inter-fabric or intra-fabric links. Doing so, you can express and represent your intent to

add links. The links with intent are displayed in a different colour till they are actually converted to functional links. Once you physically connect the links, they are displayed as connected.

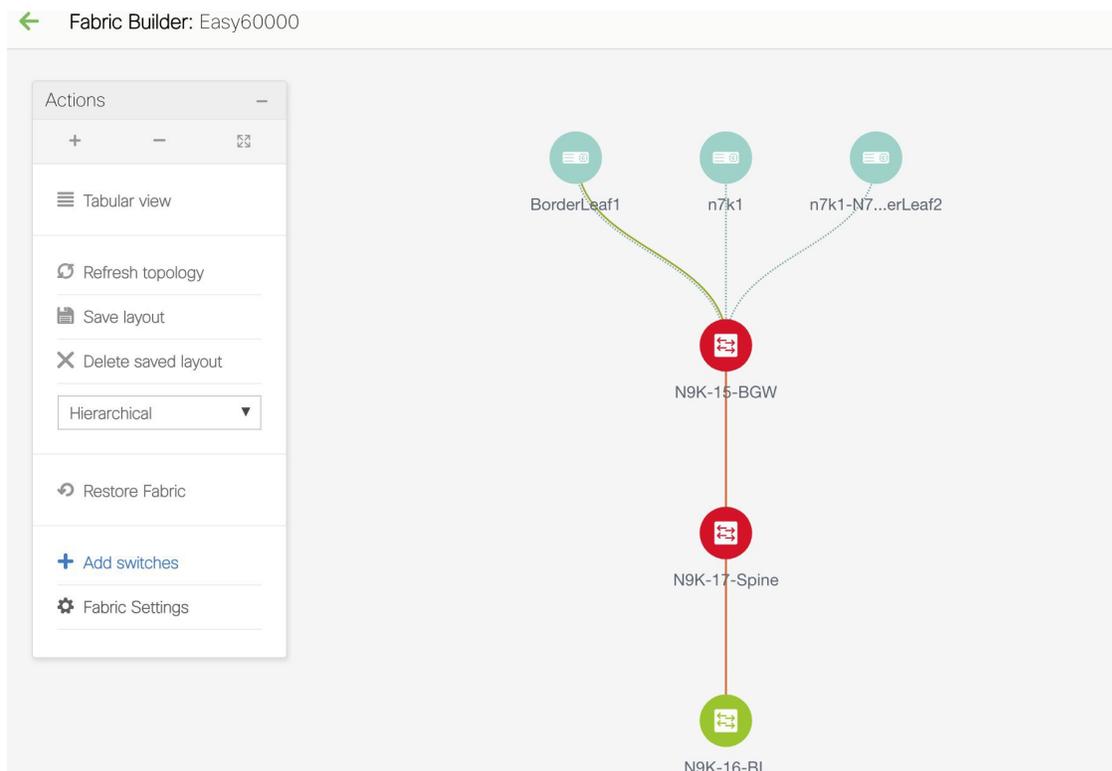
Management links might show up in the fabric topology as red colored links. To remove such links, right-click the link and click **Delete Link**.

From Cisco DCNM Release 11.1(1), the Border Spine and Border Gateway Spine roles are added to switch roles for border switches.

You can create links between existing and pre-provisioned devices as well by selecting the pre-provisioned device as the destination device.

## Creating Intra-Fabric Links

1. Click Control > Fabric Builder to go to the Fabric Builder screen.
2. Click within the rectangular box that represents the fabric. The fabric topology screen comes up.
3. Click Tabular view in the Actions panel that is displayed at the left part of the screen.



A screen with the tabs Switches and Links appears. They list the fabric switches and links in a table.

← Fabric Builder: Easy60000 Save & Deploy

Switches **Links**

View/Edit Policies Manage Interfaces History Deploy Show All

|   | <input type="checkbox"/> | Name         | IP Address | Role       | Serial Number | Fabric Name | Fabric Status | Discovery Status | Model           |
|---|--------------------------|--------------|------------|------------|---------------|-------------|---------------|------------------|-----------------|
| 1 | <input type="checkbox"/> | N9K-15-BGW   | 111.0.0.95 | border ... | FDO20401LB4   | Easy60000   | In-Sync       | ok               | N9K-C93180YC-EX |
| 2 | <input type="checkbox"/> | N9K-16-Leaf  | 111.0.0.96 | leaf       | SAL18432P6G   | Easy60000   | In-Sync       | ok               | N9K-C9396PX     |
| 3 | <input type="checkbox"/> | N9K-17-Spine | 111.0.0.97 | spine      | FDO20401LEJ   | Easy60000   | In-Sync       | ok               | N9K-C93180YC-EX |

4. Click the Links tab. You can see a list of links.

The list is empty when you are yet to create a link.

Switches **Links**

Show All

|    | <input type="checkbox"/> | Scope                     | Name  | Policy                            | Admin State | Oper State |
|----|--------------------------|---------------------------|---|-----------------------------------|-------------|------------|
| 1  | <input type="checkbox"/> | Easy60000                 | N9K-15-BGW-Ethernet1/3---n7k1-N7K-1-BorderLeaf2-Ethe... |                                   |             |            |
| 2  | <input type="checkbox"/> | Easy60000                 | N9K-16-Leaf-Ethernet2/1---n7k1-Ethernet7/8              |                                   |             |            |
| 3  | <input type="checkbox"/> | External65000<->Easy60000 | BorderLeaf1-Loopback0---N9K-15-BGW-loopback0            | multisite_overlay_setup_rs_test   |             |            |
| 4  | <input type="checkbox"/> | Easy7200<->Easy60000      | N9K-4-BGW-Ethernet1/2---N9K-15-BGW-Ethernet1/8          | ext_multisite_underlay_setup_test |             |            |
| 5  | <input type="checkbox"/> | Easy7200<->Easy60000      | N9K-3-BGW-Ethernet1/2---N9K-15-BGW-Ethernet1/7          | ext_multisite_underlay_setup_test |             |            |
| 6  | <input type="checkbox"/> | Easy60000                 | N9K-15-BGW-Ethernet1/5---N9K-17-Spine-Ethernet1/1       | int_intra_fabric_num_link_11_1    |             |            |
| 7  | <input type="checkbox"/> | Easy7200<->Easy60000      | N9K-1-Spine-Ethernet1/1---N9K-16-Leaf-Ethernet1/3       |                                   |             |            |
| 8  | <input type="checkbox"/> | Easy60000                 | N9K-17-Spine-Ethernet1/2---N9K-16-Leaf-Ethernet1/5      | int_intra_fabric_num_link_11_1    |             |            |
| 9  | <input type="checkbox"/> | Easy7200<->Easy60000      | N9K-2-Leaf-Ethernet1/2---N9K-16-Leaf-Ethernet1/4        |                                   |             |            |
| 10 | <input type="checkbox"/> | Easy60000                 | N9K-15-BGW-Ethernet1/2---N9K-16-Leaf-Ethernet1/2        |                                   |             |            |
| 11 | <input type="checkbox"/> | Easy60000<->Easy7200      | N9K-15-BGW-Ethernet1/4---N9K-1-Spine-Ethernet1/2        |                                   |             |            |
| 12 | <input type="checkbox"/> | Easy60000<->Easy7200      | N9K-15-BGW-Ethernet1/50---N9K-18-BGW-Ethernet1/7        |                                   |             |            |
| 13 | <input type="checkbox"/> | Easy60000<->External65000 | N9K-15-BGW-Ethernet1/49---n7k1-BorderLeaf1-Ethernet7/6  |                                   |             |            |

5. Click the Add (+) button at the top left part of the screen to add a link.

The Add Link screen comes up. By default, the Intra-Fabric option is chosen as the link type.

## Link Management - Add Link

The screenshot shows the 'Link Management - Add Link' configuration page. The 'Link Type' dropdown is highlighted with a red box and an arrow pointing to it. Below it are several other dropdown menus for Link Sub-Type, Link Template, Source Fabric, Destination Fabric, Source Device, Source Interface, Destination Device, and Destination Interface. A 'Link Profile' section is expanded, showing fields for FABRIC\_NAME, Source IP, Destination IP, Interface Admin State (checked), and MTU (9216). A 'Save' button is at the bottom right.

The fields are:

Link Type – Choose Intra-Fabric to create a link between two switches in a fabric.

Link Sub-Type – This field populates Fabric indicating that this is a link within the fabric.

Link Template: You can choose any of the following link templates.

- `int_intra_fabric_num_link_11_1`: If the link is between two ethernet interfaces assigned with IP addresses, choose `int_intra_fabric_num_link_11_1`.
- `int_intra_fabric_unnum_link_11_1`: If the link is between two IP unnumbered interfaces, choose `int_intra_fabric_unnum_link_11_1`.
- `int_intra_vpc_peer_keep_alive_link_11_1`: If the link is a vPC peer keep-alive link, choose `int_intra_vpc_peer_keep_alive_link_11_1`.
- `int_pre_provision_intra_fabric_link`: If the link is between two pre-provisioned devices, choose `int_pre_provision_intra_fabric_link`. After you click **Save & Deploy**, an IP address is picked from the underlay subnet IP pool.

Correspondingly, the Link Profile section fields is updated.

Source Fabric – The fabric name populates this field since the source fabric is known.

Destination Fabric – Choose the destination fabric. For an intra-fabric link, source and destination fabrics are the same.

Source Device and Source Interface – Choose the source device and interface.

Destination Device and Destination Interface – Choose the destination device and interface.



**Note** Select the pre-provisioned device as the destination device if you are creating a link between an existing device and a pre-provisioned device.

**General** tab in the Link Profile section

Interface VRF – Name of a non-default VRF for this interface.

Source IP and Destination IP – Specify the source and destination IP addresses of the source and destination interfaces, respectively.



**Note** The Source IP and Destination IP fields do not appear if you choose **int\_pre\_provision\_intra\_fabric\_link** template.

Interface Admin State – Check or uncheck the check box to enable or disable the admin state of the interface.

MTU – Specify the maximum transmission unit (MTU) through the two interfaces.

Link Management - Add Link
✕

\* Link Type

\* Link Sub-Type

\* Link Template

\* Source Fabric

\* Destination Fabric

\* Source Device

\* Source Interface

\* Destination Device

\* Destination Interface

▼ Link Profile

General

Advanced

\* FABRIC\_NAME  ? FABRIC NAME

\* Source IP  ? IP address of the source interface

\* Destination IP  ? IP address of the destination interface

Interface Admin State  ? Admin state of the interface

\* MTU  ? MTU for the interface

**Advanced** tab.

▼ Link Profile

General

Advanced

Source Interface Desc... Border Leaf to Route Reflector1 ? Add description to the source inte

Destination Interface ... Route Reflector1 to Border Leaf ? Add description to the destinior

Source Interface Free... ? Additional CLI for source Interfac

Destination Interface ... ? Additional CLI for destination Inte

Save

Source Interface Description and Destination Interface Description – Describe the links for later use. For example, if the link is between a leaf switch and a route reflector device, you can enter the information in these fields (Link from leaf switch to RR 1 and Link from RR 1 to leaf switch). This description will be converted into a config, but will not be pushed into the switch. After **Save & Deploy**, it will reflect in the running configuration.

Source Interface Freeform CLIs and Destination Interface Freeform CLIs: Enter the freeform configurations specific to the source and destination interfaces. You should add the configurations as displayed in the running configuration of the switch, without indentation. For more information, refer [Enabling Freeform Configurations on Fabric Switches](#).

- Click Save at the bottom right part of the screen.

The new link appears in the Links tab.

|   | <input type="checkbox"/> | Scope     | Name   | Policy                         | Admin State | Oper State |
|---|--------------------------|-----------|--|--------------------------------|-------------|------------|
| 1 | <input type="checkbox"/> | Easy60000 | N9K-16-BL-Ethernet1/40---N9K-17-Spine-Ethernet1/40 | int_intra_fabric_num_link_11_1 |             |            |
| 2 | <input type="checkbox"/> | Easy60000 | N9K-16-BL-Ethernet2/1---n7k1-Ethernet7/8           |                                |             |            |
| 3 | <input type="checkbox"/> | Easy60000 | N9K-15-BGW-Ethernet1/5---N9K-17-Spine-Ethernet1/1  | int_intra_fabric_num_link_11_1 |             |            |

- Click **Save & Deploy** to deploy the link configurations on the switches.

The Config Deployment screen comes up. It displays the configuration status on the switches. You can also view the pending configurations by clicking the respective link in the Preview Config column. When you click a link in the Preview Config column, the Config Preview window comes up. It lists the pending configurations on the switch. The Side-by-side Comparison tab displays the running configuration and expected configuration side-by-side.

- Close the preview screen and click Deploy Config. The pending configurations are deployed.
- After ensuring that the progress is 100% in all the rows, click Close at the bottom part of the screen. The Links screen comes up again.

Click <- at the top left part of the screen to go to the fabric topology. In the fabric topology, you can see that the link between the two devices is displayed.

## Creating Inter-Fabric Links

1. Click the Links tab in the Switches | Links page. The list of previously created links are displayed. The list contains intra-fabric links (between switches in a fabric), and inter-fabric links (between BGWs or border leaf/spine switches of different fabrics).

|    | <input type="checkbox"/> | Scope               | Name   | Policy                         | Admin State | Oper State |
|----|--------------------------|---------------------|--|--------------------------------|-------------|------------|
| 1  | <input type="checkbox"/> | Easy60000           | N9K-16-Leaf-Ethernet2/1---n7k1-Ethernet7/8               |                                |             |            |
| 2  | <input type="checkbox"/> | Easy60000           | N9K-15-bgw-Ethernet1/49---n7k1-BorderLeaf1-Ethernet7/6   |                                |             |            |
| 3  | <input type="checkbox"/> | Easy60000           | N9K-15-bgw-Ethernet1/3---n7k1-N7K-1-BorderLeaf2-Ether... |                                |             |            |
| 4  | <input type="checkbox"/> | Easy60000           | N9K-17-Spine-Ethernet1/2---N9K-16-Leaf-Ethernet1/5       | int_intra_fabric_num_link_11_1 |             |            |
| 5  | <input type="checkbox"/> | Easy60000           | N9K-15-bgw-Ethernet1/5---N9K-17-Spine-Ethernet1/1        | int_intra_fabric_num_link_11_1 |             |            |
| 6  | <input type="checkbox"/> | New7200<->Easy60000 | n9k-3-bgw-Ethernet1/2---N9K-15-bgw-Ethernet1/7           |                                |             |            |
| 7  | <input type="checkbox"/> | Easy60000<->New7200 | N9K-15-bgw-Ethernet1/50---n9k-18-bgw-Ethernet1/7         |                                |             |            |
| 8  | <input type="checkbox"/> | New7200<->Easy60000 | n9k-4-bgw-Ethernet1/2---N9K-15-bgw-Ethernet1/8           |                                |             |            |
| 9  | <input type="checkbox"/> | Easy60000           | N9K-15-bgw-Ethernet1/2---N9K-16-Leaf-Ethernet1/2         |                                |             |            |
| 10 | <input type="checkbox"/> | New7200<->Easy60000 | n9k-2-leaf-Ethernet1/2---N9K-16-Leaf-Ethernet1/4         |                                |             |            |
| 11 | <input type="checkbox"/> | New7200<->Easy60000 | n9k-1-spine-Ethernet1/1---N9K-16-Leaf-Ethernet1/3        |                                |             |            |
| 12 | <input type="checkbox"/> | Easy60000<->New7200 | N9K-15-bgw-Ethernet1/4---n9k-1-spine-Ethernet1/2         |                                |             |            |

2. Click the Add (+) button at the top left part of the screen to add a link. The Add Link screen comes up. By default, the Intra-Fabric option is chosen as the link type.

### Link Management - Add Link

Link Management - Add Link
✕

\* Link Type

\* Link Sub-Type

\* Link Template

\* Source Fabric

\* Destination Fabric

\* Source Device

\* Source Interface

\* Destination Device

\* Destination Interface

▼ Link Profile

General

Advanced

\* FABRIC\_NAME  ? FABRIC NAME

\* Source IP  ? IP address of the source interface

\* Destination IP  ? IP address of the destination interface

Interface Admin State  ? Admin state of the interface

\* MTU  ? MTU for the interface

- From the Link Type drop-down box, choose Inter-Fabric since you are creating an IFC. The screen changes correspondingly.

Link Management - Add Link ✕

\* Link Type: Inter-Fabric

\* Link Sub-Type: VRF\_LITE

\* Link Template: ext\_fabric\_setup\_test

\* Source Fabric: Easy60000

\* Destination Fabric:

\* Source Device:

\* Source Interface:

\* Destination Device:

\* Destination Interface:

---

▼ Link Profile

General

\* Local BGP AS #: 60000 ? Local BGP Autonomous System Nu

\* IP\_MASK: ?

\* NEIGHBOR\_IP: ?

\* NEIGHBOR\_ASN: ?

[Save](#)

The fields for inter-fabric link creation are explained:

Link Type – Choose Inter-Fabric to create an inter-fabric connection between two fabrics, via their border switches.

Link Sub-Type – This field populates the IFC type. Choose **VRF\_LITE**, **MULTISITE\_UNDERLAY**, or **MULTISITE\_OVERLAY** from the drop-down list.

The Multi-Site options are explained in the Multi-Site use case.

Link Template: The link template is populated.

The templates are autopopulated with corresponding pre-packaged default templates that are based on your selection.



**Note** You can add, edit, or delete user-defined templates. See *Template Library* section in the Control chapter for more details.

Source Fabric - This field is prepopulated with the source fabric name.

Destination Fabric - Choose the destination fabric from this drop-down box.

Source Device and Source Interface - Choose the source device and Ethernet interface that connects to the destination device.

Destination Device and Destination Interface—Choose the destination device and Ethernet interface that connects to the source device.

Based on the selection of the source device and source interface, the destination information is autopopulated based on Cisco Discovery Protocol information, if available. There is an extra validation performed to ensure that the destination external device is indeed part of the destination fabric.

**General** tab in the Link Profile section.

Local BGP AS# - In this field, the AS number of the source fabric is autopopulated.

IP\_MASK—Fill up this field with the IP address of the source interface that connects to the destination device.

NEIGHBOR\_IP—Fill up this field with the IP address of the destination interface.

NEIGHBOR\_ASN—In this field, the AS number of the destination device is autopopulated.

After filling up the Add Link screen, it looks like this:

Link Management - Add Link
✕

|                                |                       |   |
|--------------------------------|-----------------------|---|
| <b>* Link Type</b>             | Inter-Fabric          | ▼ |
| <b>* Link Sub-Type</b>         | VRF_LITE              | ▼ |
| <b>* Link Template</b>         | ext_fabric_setup_test | ▼ |
| <b>* Source Fabric</b>         | Easy60000             | ▼ |
| <b>* Destination Fabric</b>    | New7200               | ▼ |
| <b>* Source Device</b>         | N9K-15-bgw            | ▼ |
| <b>* Source Interface</b>      | Ethernet1/9           | ▼ |
| <b>* Destination Device</b>    | n9k-18-bgw            | ▼ |
| <b>* Destination Interface</b> | Ethernet1/9           | ▼ |

▼ Link Profile

General

|                         |             |                                  |
|-------------------------|-------------|----------------------------------|
| <b>* Local BGP AS #</b> | 60000       | ? Local BGP Autonomous System Nu |
| <b>* IP_MASK</b>        | 10.3.4.5/24 | ?                                |
| <b>* NEIGHBOR_IP</b>    | 10.3.4.7    | ?                                |
| <b>* NEIGHBOR_ASN</b>   | 7200        | ?                                |

Save

4. Click Save at the bottom right part of the screen.

The Switches|Links screen comes up again. You can see that the IFC is created and displayed in the list of links.

|   | <input type="checkbox"/> | Scope               | Name   | Policy                |
|---|--------------------------|---------------------|--|-----------------------|
| 1 | <input type="checkbox"/> | Easy60000           | N9K-16-Leaf~Ethernet2/1---n7k1~Ethernet7/8             |                       |
| 2 | <input type="checkbox"/> | Easy60000           | N9K-15-bgw~Ethernet1/49---n7k1~BorderLeaf1~Ethernet7/6 |                       |
| 3 | <input type="checkbox"/> | Easy60000<->New7200 | N9K-15-bgw~Ethernet1/9---n9k-18-bgw~Ethernet1/9        | ext_fabric_setup_test |

5. Click on Save & Deploy to deploy the link configurations on the switches.

The Config Deployment screen comes up. It displays the configuration status on the switches. You can also view the pending configurations by clicking the respective link in the Preview Config column. When you click a link in the Preview Config column, the Config Preview window comes up. It lists the pending configurations on the switch. The Side-by-side Comparison tab displays the running configuration and expected configuration side-by-side.

6. Close the preview screen and click Deploy Config. The pending configurations are deployed.
7. After ensuring that the progress is 100% in all the rows, click Close at the bottom part of the screen. The Links screen comes up again.
8. Click <- at the top left part of the screen to go to the fabric topology. In the fabric topology, you can see that the link between the two devices is displayed.

If the two fabrics are member fabric of an MSD, then you can see the link in the MSD topology too.

When you enable the VRF Lite function using the ToExternalOnly method or Multisite function via MSD fabric, IFCs are automatically created between the (VXLAN fabric) border/BGW device and connected (external fabric) edge router/core device. When you remove the ER/core/border/BGW device, the corresponding IFCs (link PTIs) to/from that switch are deleted on DCNM. Subsequently, DCNM removes the corresponding IFC configurations, if any, from the remaining devices on the next Save & Deploy operation. Also, if you want to remove a device that has an IFCs and overlay extensions over those IFCs, you should undeploy all overlay extensions corresponding to those IFCs for switch delete to be possible.

To undeploy VRF extensions, click Control > Networks & VRFs, select the VXLAN fabric and the extended VRFs, and undeploy the VRFs in the VRF deployment screen.

To delete the IFCs, click Control > Fabric Builder, go to the fabric topology screen, click Tabular view, and delete the IFCs from the Links tab.

Ensure that the fabric switch names are unique. If you deploy VRF extensions on switches with the same name, it leads to erroneous configuration.

The new fabric is created, the fabric switches are discovered in DCNM, the underlay networks provisioned on those switches, and the configurations between DCNM and the switches are synced. The remaining tasks are:

- Provision interface configurations such as vPCs, loopback interface, and subinterface configurations. Refer [Interfaces](#).
- Create overlay networks and VRFs and deploy them on the switches. Refer [Creating and Deploying Networks and VRFs](#).

## Exporting Links

1. Choose Control > Fabric Builder, and select a fabric.

The fabric topology window appears.

2. Click **Tabular view** in the **Actions** panel.

A window with the **Switches** and **Links** tabs appears.

3. Click the **Links** tab.

You can see a list of links. The list is empty when you are yet to create a link.

4. Click the **Export Links** icon to export the links in a CSV file.

The following details of links are exported: link template, source fabric, destination fabric, source device, destination device, source switch name, destination switch name, source interface, destination interface, and nvPairs. The nvPairs field consists of a JSON object.

## Importing Links

You can import a CSV file containing details of links to add new links to the fabric. The CSV file should have the following details of links: link template, source fabric, destination fabric, source device, destination device, source switch name, destination switch name, source interface, destination interface, and nvPairs.



### Note

- You cannot update existing links.
- The **Import Links** icon is disabled for external fabric.

1. Choose **Control > Fabric Builder**, and select a fabric.

The fabric topology window appears.

2. Click **Tabular view** in the **Actions** panel.

A window with the **Switches** and **Links** tabs appears.

3. Click the **Links** tab.

You can see a list of links. The list is empty when you are yet to create a link.

4. Click the **Import Links** icon.

The file server directory opens.

5. Browse the directory and select the CSV file that you want to import.

6. Click **Open**.

A confirmation screen appears.

7. Click **Yes** to import the selected file.

## Viewing Details of Fabric Links

You can view information about a fabric link, like IP subnet between links to deploy underlay, MTU, speed mismatch, and so on, in the topology view of a fabric builder. To view the details of a link from the Cisco DCNM Web client, perform the following steps:

## Procedure

---

**Step 1** Choose **Control > Fabrics > Fabric Builder** and select a fabric.

The topology view of the fabric appears.

**Step 2** Double-click any of the links.

The details window appears. You can view the devices that are connected using this link, summary, and the data traffic.

**Step 3** Click **Show more details**.

A comparison table of the two devices connected by the link appears. It includes the following parameters of the devices: device name, name, admin status, operation status, reason, policies, overlay network, status, PC, vPC ID, speed, MTU, mode, VLANs, IP or prefix, VRF, neighbor, and description.

- Note**
- You can view the traffic details of a fabric link by clicking the device name with hyperlink. Alternatively, you can view these traffic details in the details window. See *Viewing the Traffic Details of the Fabric Links* section for more information.
  - You can view the expected configuration of a fabric link by clicking the policy with the hyperlink.

**Step 4** Click the **Back** icon to go back to the details window.

**Note** You can click the **Close** icon to exit the details window.

---

## Viewing the Traffic Details of Fabric Links

In the details window of a fabric link, you can choose how you want to view the traffic details. You can view the traffic details based on the time duration, format, and export this information.

You can view the data traffic of a link for the following durations from the duration drop-down list:

- 24 Hours
- Week
- Month
- Year

**Show:** Click **Show**, and choose **Chart**, **Table**, or **Chart and Table** from the drop-down list to see how you want to view the traffic details. Enlarge your browser window to view the details in **Chart and Table** format.

If you choose **Chart**, hover over the traffic chart to view the Rx and Tx values, along the Y axis, for the corresponding time, along X axis. You can change the time duration values of the X axis by moving the sliders in the time range selector. You can choose the Y-axis values by checking or unchecking the Rx and Tx check boxes.



**Note** If you select **Week**, **Month**, or **Year** as the time duration, you can also view the Peak Rx and Peak Tx values along the Y axis.

Select **Table** to view the traffic information in tabular format.

**Chart Type and Chart Options:** Choose **Area Chart** or **Line Chart** from the **Chart Type** drop-down list.

You can choose the following chart options:

- **Show Fill Patterns**
- **Show Datamarkers**
- **Y Axis Log Scale**

**Actions:** Export or print the traffic information by choosing the appropriate options from the **Actions** drop-down list.

## vPC Fabric Peering

You can create a virtual peer link for two switches or change the existing physical peer link to a virtual peer link. Only greenfield deployments support vPC fabric peering in Cisco DCNM, Release 11.2(1). This feature is applicable for **Easy\_Fabric\_11\_1** and **Easy\_Fabric\_eBGP** fabric templates.

### Guidelines and Limitations

The following are the guidelines and limitations for vPC fabric pairing.

- vPC fabric peering is supported from Cisco DCNM Release 11.2(1) and Cisco NX-OS Release 9.2(3).
- Only Cisco Nexus N9K-C9332C Switch, Cisco Nexus N9K-C9364C Switch, Cisco Nexus N9K-C9348GC-FXP Switch as also the Cisco Nexus 9000 Series Switches that ends with FX, FX2, and FX2-Z support vPC fabric peering.
- If you use other Cisco Nexus 9000 Series Switches, a warning will appear during **Save & Deploy**. A warning appears in this case because these switches will be supported in future releases.
- If you try pairing switches that do not support vPC fabric peering, using the **Use Virtual Peerlink** option, a warning will appear when you deploy the fabric.
- You can convert a physical peer link to a virtual peer link and vice-versa with or without overlays.
- Switches with border gateway leaf roles do not support vPC fabric peering.
- vPC fabric peering is not supported for Cisco Nexus 9000 Series Modular Chassis and FEXs. An error appears during **Save & Deploy** if you try to pair any of these.
- Only greenfield deployments support vPC fabric peering.
- However, you can import switches that are connected using physical peer links and convert the physical peer links to virtual peer links after **Save & Deploy**. To update a TCAM region during the feature configuration, use the **hardware access-list tcam ingress-flow redirect 512** command in the configuration terminal.

### Fields and Description

To view the vPC pairing window of a switch, from the fabric topology window, right-click the switch and choose **vPC Pairing**. The vPC pairing window for a switch has the following fields:

| Field                | Description   |
|----------------------|---|
| Use Virtual Peerlink | Allows you to enable or disable the virtual peer linking between switches.  |
| Switch name          | Specifies all the peer switches in a fabric.<br><br><b>Note</b> When you have not paired any peer switches, you can see all the switches in a fabric. After you pair a peer switch, you can see only the peer switch in the vPC pairing window. |
| Recommended          | Specifies if the peer switch can be paired with the selected switch. Valid values are <b>true</b> and <b>false</b> . Recommended peer switches will be set to <b>true</b> .   |
| Reason               | Specifies why the vPC pairing between the selected switch and the peer switches is possible or not possible.  |
| Serial Number        | Specifies the serial number of the peer switches.   |

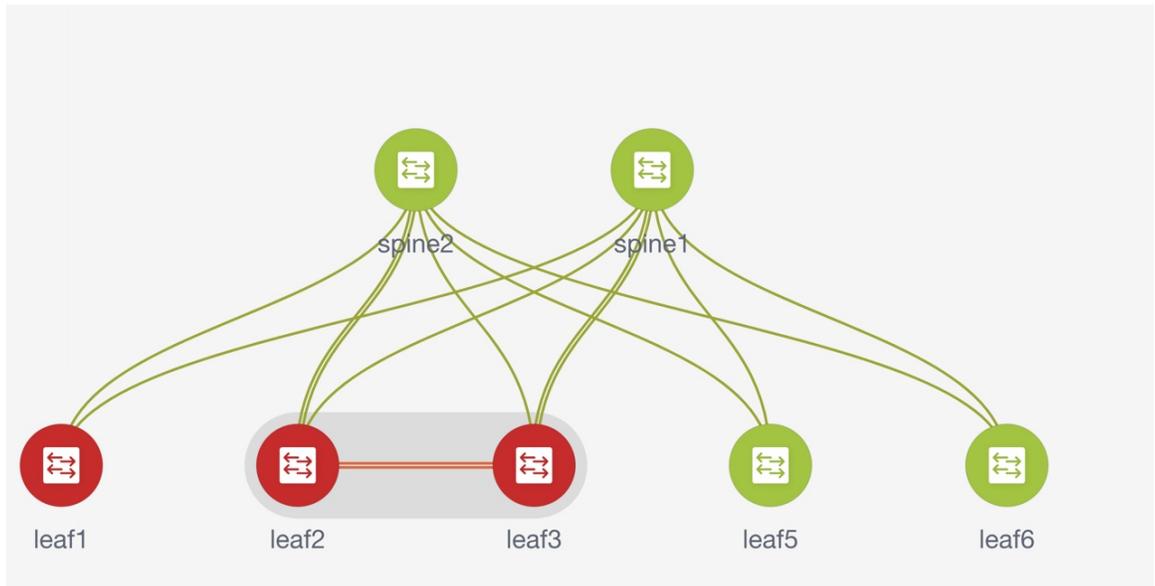
You can perform the following with the **vPC Pairing** option:

## Creating a Virtual Peer Link

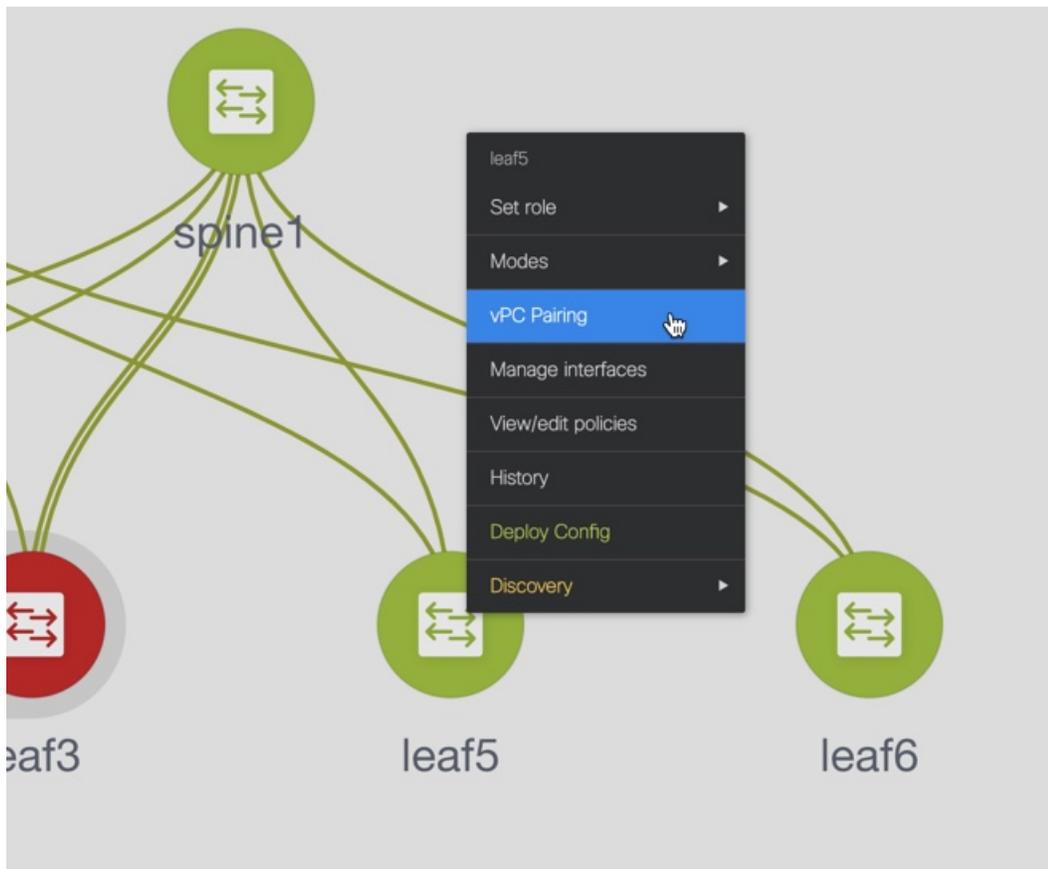
To create a virtual peer link from the Cisco DCNM Web UI, perform the following steps:

### Procedure

- 
- Step 1** Choose **Control > Fabrics**.  
The **Fabric Builder** window appears.
- Step 2** Choose a fabric with the **Easy\_Fabric\_11\_1** or **Easy\_Fabric\_eBGP** fabric templates.  
The fabric topology window appears.



**Step 3** Right-click a switch and choose **vPC Pairing** from the drop-down list.  
The window to choose the peer appears.



**Note** You will get the following error when you choose a switch with the border gateway leaf role.

<switch-name> has a Network/VRF attached. Please detach the Network/VRF before vPC Pairing/Unpairing

**Step 4** Check the **Use Virtual Peerlink** check box.

**Step 5** Choose a peer switch and check the **Recommended** column to see if pairing is possible.

If the value is **true**, pairing is possible. You can pair switches even if the recommendation is **false**. However, you will get a warning or error during **Save & Deploy**.

**Step 6** Click **Save**.

Select vPC peer for leaf5
✕

Use Virtual Peerlink

1

|   | Switch name                            | Recommended | Reason                                   | Serial Number |
|---|--|-------------|--|---------------|
| 2 | <input checked="" type="radio"/> leaf6 | true        | Switches have same role                  | FDO22360M0D   |
|   | <input type="radio"/> leaf3            | false       | Already paired with FDO20352BEE          | FDO20290DVJ   |
|   | <input type="radio"/> leaf1            | false       | N9K-C93180YC-EX doesn't support Virtu... | FDO2035283H   |
|   | <input type="radio"/> spine2           | false       | Switches have different roles            | FDO20352B6H   |
|   | <input type="radio"/> spine1           | false       | Switches have different roles            | FDO20401L8J   |
|   | <input type="radio"/> leaf2            | false       | Already paired with FDO20290DVJ          | FDO20352BEE   |

3

**Step 7** In the **Fabric Topology** window, click **Save & Deploy**.

The **Config Deployment** window appears.

**Step 8** Click the field against the switch in the **Preview Config** column.

The **Config Preview** window appears for the switch.

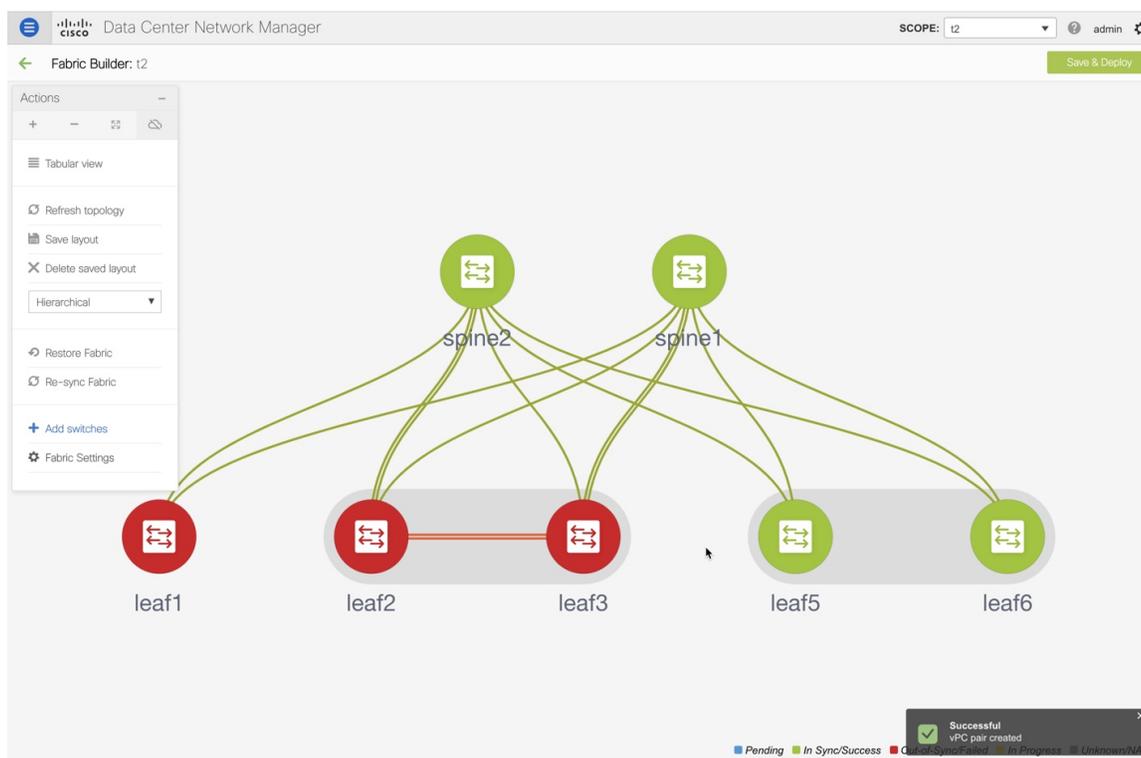
**Step 9** View the vPC link details in the pending configuration and the side-by-side configuration.

**Step 10** Close the window.

**Step 11** Click the pending errors icon next to the **Save & Deploy** icon to view errors and warnings, if any.

If you see any warnings that are related to TCAM, click the **Resolve** icon. A confirmation dialog box about reloading switches appears. Click **OK**. You can also reload the switches from **Tabular view** in the fabric topology window.

The switches that are connected through vPC fabric peering, are enclosed in a gray cloud.



## Converting a Physical Peer Link to a Virtual Peer Link

To convert a physical peer link to a virtual peer link from the Cisco DCNM Web UI, perform the following steps:

### Before you begin

- Plan the conversion from physical peer link to virtual peer link during the maintenance window of switches.
- Ensure the switches support vPC fabric peering. Only the following switches support vPC fabric peering:
  - Cisco Nexus N9K-C9332C Switch, Cisco Nexus N9K-C9364C Switch, and Cisco Nexus N9K-C9348GC-FXP Switch
  - Cisco Nexus 9000 Series Switches that ends with FX, FX2, and FX2-Z

### Procedure

- Step 1** Choose **Control > Fabrics**.  
The **Fabric Builder** window appears.
- Step 2** Choose a fabric with the **Easy\_Fabric\_11\_1** or **Easy\_Fabric\_eBGP** fabric templates.

- Step 3** Right-click the switch that is connected using the physical peer link and choose **vPC Pairing** from the drop-down list.
- The window to choose the peer appears.
- Note** You will get the following error when you choose a switch with the border gateway leaf role.
- ```
<switch-name> has a Network/VRF attached. Please detach the Network/VRF before vPC Pairing/Unpairing
```
- Step 4** Check the **Recommended** column to see if pairing is possible.
- If the value is **true**, pairing is possible. You can pair switches even if the recommendation is **false**. However, you will get a warning or error during **Save & Deploy**.
- Step 5** Check the **Use Virtual Peerlink** check box.
- The **Unpair** icon changes to **Save**.
- Step 6** Click **Save**.
- Note** After you click **Save**, the physical vPC peer link is automatically deleted between the switches even without deployment.
- Step 7** In the **Fabric Topology** window, click **Save & Deploy**.
- The **Config Deployment** window appears.
- Step 8** Click the field against the switch in the **Preview Config** column.
- The **Config Preview** window appears for the switch.
- Step 9** View the vPC link details in the pending configuration and the side-by-side configuration.
- Step 10** Close the window.
- Step 11** Click the pending errors icon next to the **Save & Deploy** icon to view errors and warnings, if any.
- If you see any warnings that are related to TCAM, click the **Resolve** icon. A confirmation dialog box about reloading switches appears. Click **OK**. You can also reload the switches from **Tabular view** in the fabric topology window.
- The physical peer link between the peer switches turns red. Delete this link. The switches are connected only through a virtual peer link and are enclosed in a gray cloud.

---

## Converting a Virtual Peer Link to a Physical Peer Link

To convert a virtual peer link to a physical peer link from the Cisco DCNM Web UI, perform the following steps:

### Before you begin

Connect the switches using a physical peer link before disabling the vPC fabric peering.

### Procedure

---

- Step 1** Choose **Control > Fabrics**.  
The **Fabric Builder** window appears.
- Step 2** Choose a fabric with the **Easy\_Fabric\_11\_1** or **Easy\_Fabric\_eBGP** fabric templates.
- Step 3** Right-click the switch that is connected through a virtual peer link and choose **vPC Pairing** from the drop-down list.  
The window to choose the peer appears.
- Step 4** Uncheck the **Use Virtual Peerlink** check box.  
The **Unpair** icon changes to **Save**.
- Step 5** Click **Save**.
- Step 6** In the **Fabric Topology** window, click **Save & Deploy**.  
The **Config Deployment** window appears.
- Step 7** Click the field against the switch in the **Preview Config** column.  
The **Config Preview** window appears for the switch.
- Step 8** View the vPC peer link details in the pending configuration and the side-by-side configuration.
- Step 9** Close the window.
- Step 10** Click the pending errors icon next to the **Save & Deploy** icon to view errors and warnings, if any.  
If you see any warnings that are related to TCAM, click the **Resolve** icon. The confirmation dialog box about reloading switches appears. Click **OK**. You can also reload the switches from **Tabular view** in the fabric topology window.  
The virtual peer link, represented by a gray cloud, disappears and the peer switches are connected through a physical peer link.
- 

## Viewing and Editing Policies

Cisco DCNM provides the ability to group a set of switches, and allows you to push a set of underlay configurations to the group. This release enables you to create a policy template, and apply it to multiple selected switches.

To view, add, deploy, or edit a policy, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Select any available fabric, and then click **Tabular view**.
- Step 3** Select multiple switches in switches tab, and click **View/Edit Policies**.

## Viewing Policies

### Procedure

- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Select any available fabric, and then click **Tabular view**.
- Step 3** Select multiple switches in the switches tab and click **View/Edit Policies**.

Policies are listed in view or edit policies table for multiple switches.

|   | <input type="checkbox"/>            | Name         | IP Address    | Role | Serial Number | Fabric Name | Fabric Status | Discovery Status                       | Model     |
|---|-------------------------------------|--------------|---------------|------|---------------|-------------|---------------|----------------------------------------|-----------|
| 1 | <input checked="" type="checkbox"/> | anm-host80   | 172.23.244.80 | leaf | SAL1925HA3U   | easy_fabric | In-Sync       | <input checked="" type="checkbox"/> ok | N9K-C9312 |
| 2 | <input checked="" type="checkbox"/> | EVPN-Spine81 | 172.23.244.81 | leaf | SAL1919ELJQ   | easy_fabric | Out-of-sync   | <input checked="" type="checkbox"/> ok | N9K-C9312 |

- Step 4** Select a policy and click the **View** button to view its configs.

## Adding a Policy

### Procedure

- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Select any available fabric, and then click Tabular view.
- Step 3** Select a single or multiple switches in the **Switches** tab, and click the **View/Edit Policies** button.
- Step 4** Click the **Add** icon.
- Step 5** Select a policy template and enter the mandatory parameters data and click **Save**. PTI is added per each device based on n-number of devices selection.

## Add Policy



\* Priority (1-1000):

\* Policy:  ▼

General

\* Banner  ? Banner

Variables:

Save

Cancel

**Policy:** Select a policy from this drop-down list.

**Priority:** Specify a priority for the policy. The applicable values are from 1 to 1000. The default value is 500. The lower number in the **Priority** field means that there is a higher priority for the generated configuration and POAP startup-configuration. For example, features are 50, route-maps are 100, and vpc-domain is 200.

## Deploying Policies

### Procedure

- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Select any available fabric, and then click Tabular view.
- Step 3** Select multiple switches in the switches tab, and click the **View/Edit Policies** button.
- Step 4** Select multiple polices, and then click **Push Config**. The selected PTI's configs are pushed to the group of switches.

## Editing a Policy



**Note** Multiple policy editing is not supported.

## Procedure

- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Select any available fabric, and then click **Tabular view**.
- Step 3** Select multiple switches in the switches tab, and click the **View/Edit Policies** button.
- Step 4** Select a PTI, click **Edit** to modify the required data, and then click **Save** to save the PTI.
- Step 5** Select a PTI, click **Edit** to modify the required data, and then click **Push Config** to push the policy config to the device.

- Note**
- A warning appears if you push config for a Python policy.
  - A warning appears if you edit, delete, or push config a mark-deleted policy. A mark-deleted policy is set to **true** under the **Mark Deleted** column. The switch freeform child policies of **Mark Deleted** policies appears in the **View/Edit Policies** dialog box. You can edit only **Python switch\_freeform** policies. You cannot edit **Template\_CLI switch\_freeform\_config** policies.

### Edit Policy

Policy ID: POLICY-5290  
Entity Type: SWITCH

Template Name: host\_11\_1  
Entity Name: SWITCH

\* Priority (1-1000):

General

\* Switch Name  ? Host name of the switch (Max Size 63)

Variables:

## Current Switch Configuration

### Procedure

- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Select any available fabric, and then click **Tabular view**.

**Step 3** Select multiple switches in the switches tab, and click **View/Edit Policies**.

**Step 4** Click **Current Switch Config**.

The current switch configuration appears in the **Running Config** dialog box.

**Note** The running configuration will not appear for the Cisco CSR 1000v when you click **Current Switch Config** if the user role cannot access the enable prompt by default.

## Retrieving the Authentication Key

### Retrieving the 3DES Encrypted OSPF Authentication Key

1. SSH into the switch.
2. On an unused switch interface, enable the following:

```
config terminal
  feature ospf
  interface Ethernet1/1
    no switchport
    ip ospf message-digest-key 127 md5 ospfAuth
```

In the example, **ospfAuth** is the unencrypted password.



**Note** This Step 2 is needed when you want to configure a new key.

3. Enter the **show run interface Ethernet1/1** command to retrieve the password.

```
Switch # show run interface Ethernet1/1
  interface Ethernet1/1
    no switchport
    ip ospf message-digest key 127 md5 3 sd8478f4fsw4f4w34sd8478fsdfw
    no shutdown
```

The sequence of characters after **md5 3** is the encrypted password.

4. Update the encrypted password into the **OSPF Authentication Key** field.

### Retrieving the Encrypted IS-IS Authentication Key

To get the key, you must have access to the switch.

1. SSH into the switch.
2. Create a temporary keychain.

```
config terminal
  key chain isis
  key 127
  key-string isisAuth
```

In the example, **isisAuth** is the plaintext password. This will get converted to a Cisco type 7 password after the CLI is accepted.

3. Enter the **show run | section “key chain”** command to retrieve the password.

```
key chain isis
  key 127
    key-string 7 071b245f5a
```

The sequence of characters after key-string 7 is the encrypted password. Save it.

4. Update the encrypted password into the ISIS Authentication Key field.
5. Remove any unwanted configuration made in Step 2.

### Retrieving the 3DES Encrypted BGP Authentication Key

1. SSH into the switch and enable BGP configuration for a non-existent neighbor.




---

**Note** Non-existent neighbor configuration is a temporary BGP neighbor configuration for retrieving the password.

---

```
router bgp
  neighbor 10.2.0.2 remote-as 65000
  password bgpAuth
```

In the example, **bgpAuth** is the unencrypted password.

2. Enter the show run **bgp** command to retrieve the password. A sample output:

```
neighbor 10.2.0.2
  remote-as 65000
  password 3 sd8478fswerdfw3434fsw4f4w34sdsd8478fswerdfw3434fsw4f4w3
```

The sequence of characters after password 3 is the encrypted password.

3. Update the encrypted password into the **BGP Authentication Key** field.
4. Remove the BGP neighbor configuration.

## Return Material Authorization (RMA)

This section describes how to replace a physical switch in a Fabric when using Cisco DCNM Easy Fabric mode.

### Prerequisites

- Fabric is assumed to be up and running, and minimal disruption is desired when replacing the switch. Also, the switch must be replaced with a switch of the same model (ASIC type) and physical port configuration.
- To use the POAP RMA flow, you must configure the fabric for bootstrap (POAP).
- To copy the FEX configurations for the RMA of switches which have FEX deployed, you may need to perform the Save and Deploy operation one or two times.

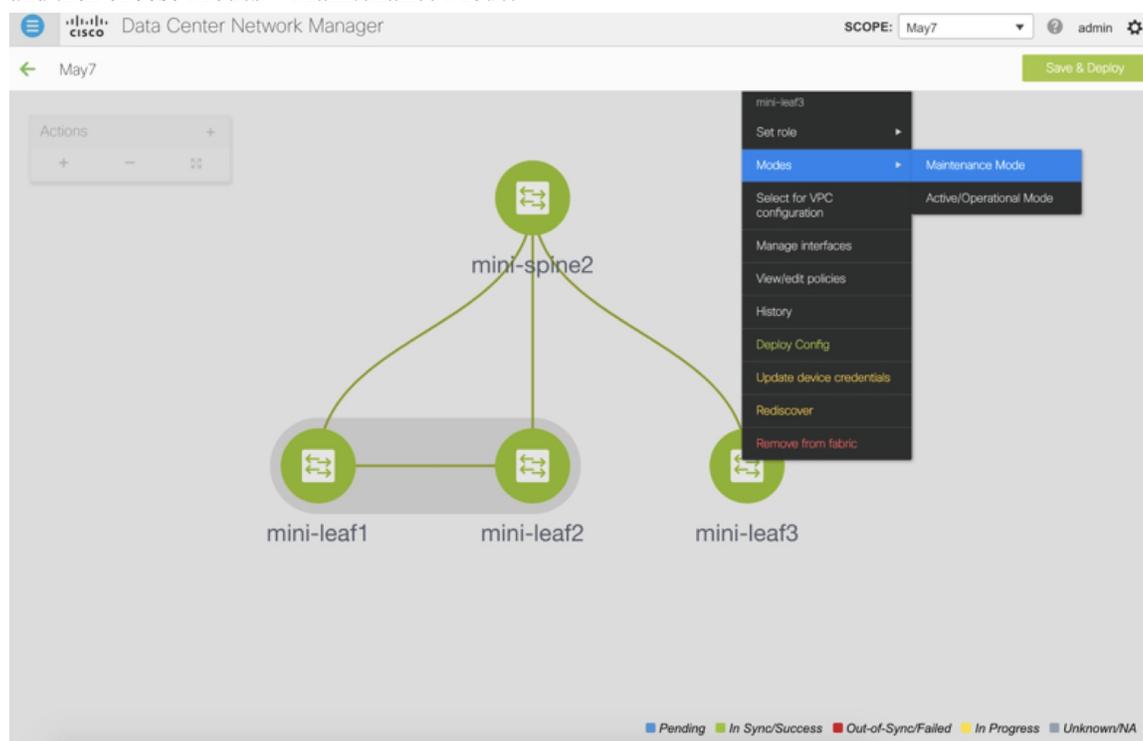
## Guidelines and Limitations

- The switch must be replaced with a switch of the same model (ASIC type) and physical port configuration. If not, the old switch must be removed and a new switch (replacement) added as a new switch into the fabric.

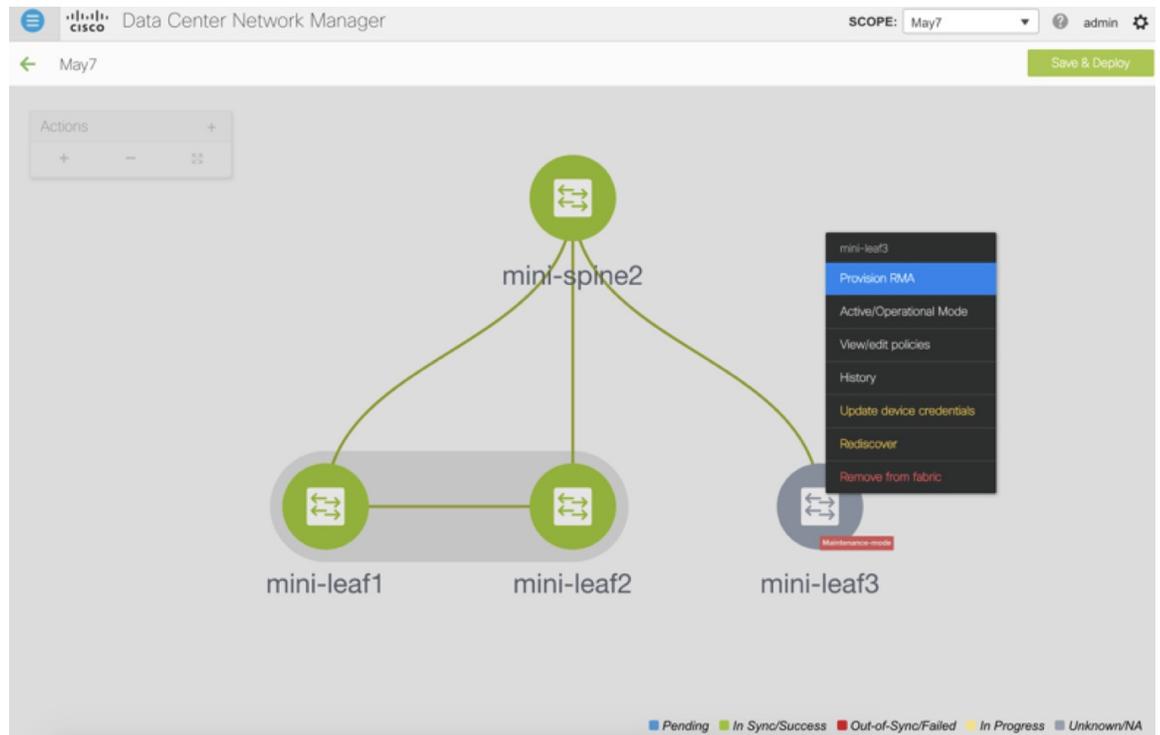
## POAP RMA Flow

### Procedure

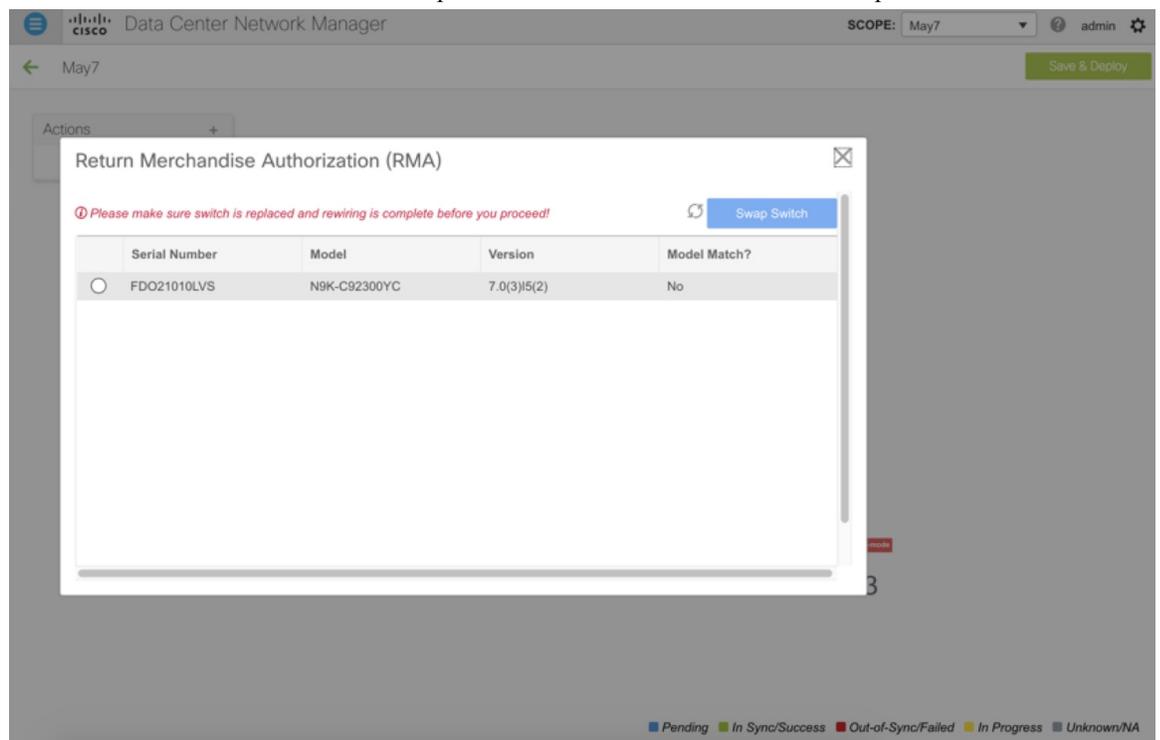
- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Click the Fabric where you want to perform RMA.
- Step 3** Move the device into maintenance mode. To move a device into maintenance mode, right-click on the device, and then choose **Modes > Maintenance Mode**.



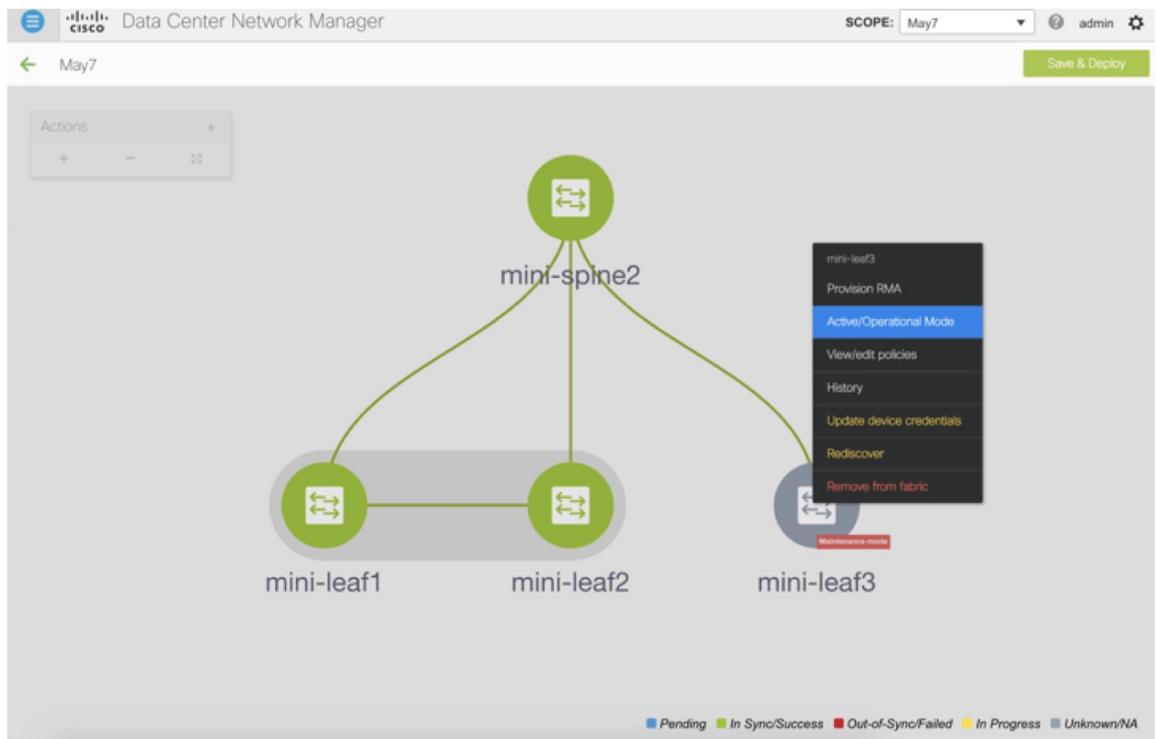
- Step 4** Physically replace the device in the network. Physical connections should be made in the same place on the replacement switch as they existed on the original switch.
- Step 5** Provision RMA flow and select the replacement device.



**Step 6** The Provision RMA UI will show the replacement device 5-10 minutes after it is powered on.



**Step 7** Select the correct replacement device and click **Swap Switch**. This begins POAP with the full “expected” configuration for that device. Total POAP time is generally around 10-15 minutes.

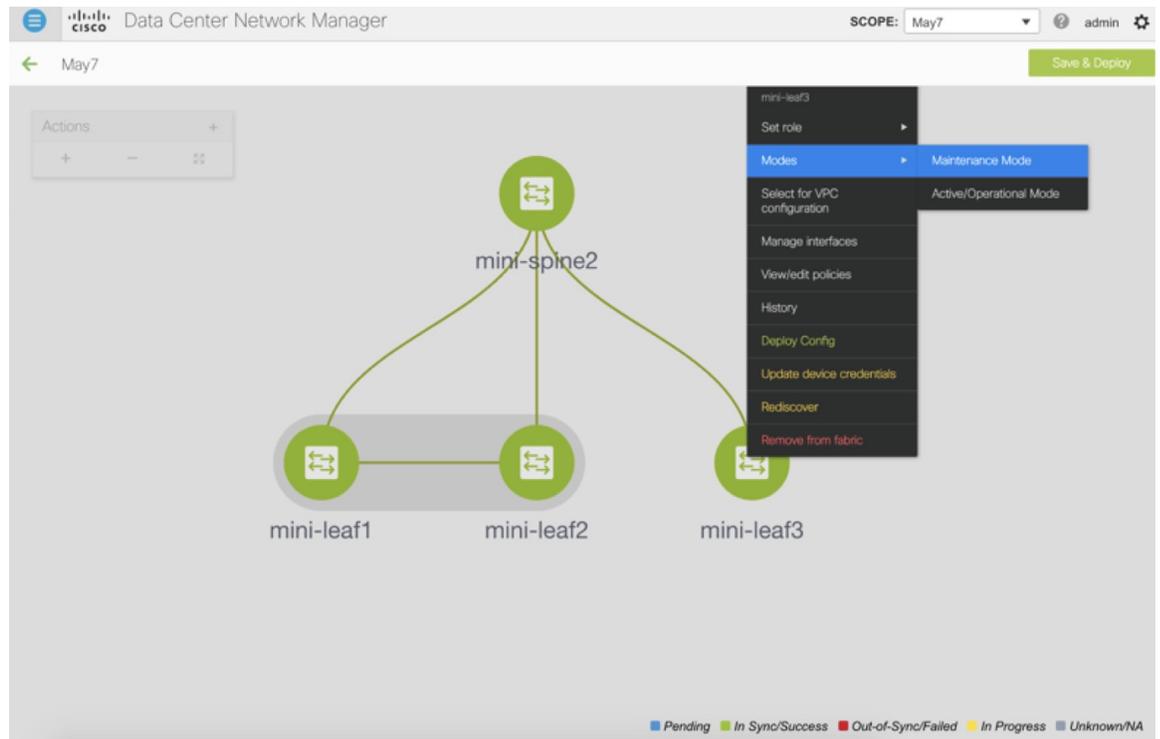


## Manual RMA Flow

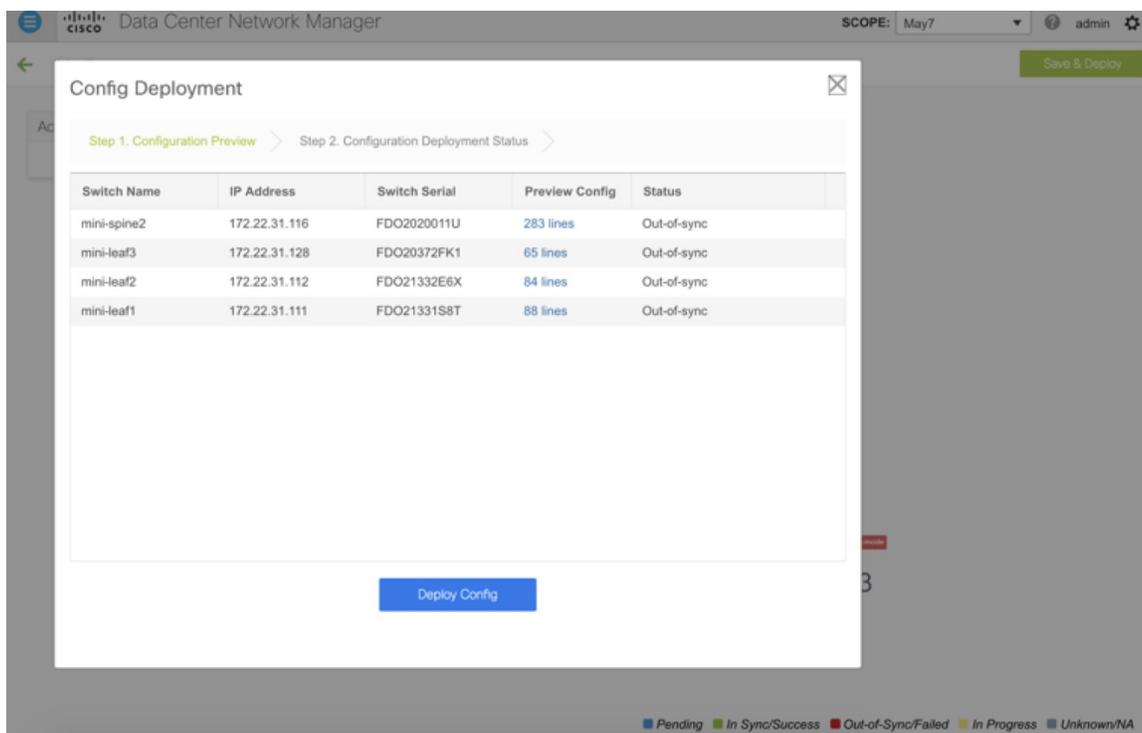
Use this flow when “Bootstrap” is not possible (or not desired), including cases that are *IPv6 only* for the initial Cisco DCNM 11.0(1) release.

### Procedure

- Step 1** Place the device in maintenance mode (optional).

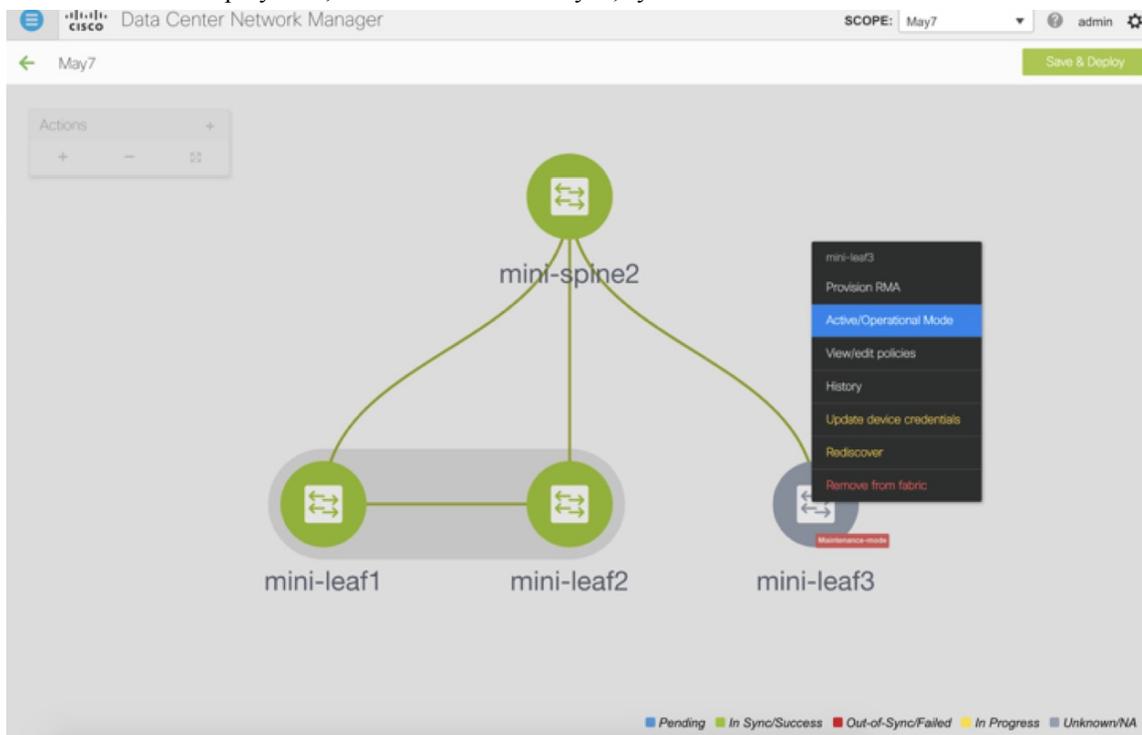


- Step 2** Physically replace the device in the network.
- Step 3** Log in through Console and set the Management IP and credentials.
- Step 4** The Cisco DCNM rediscovers the new device (or you can manually choose **Discovery > Rediscover**).
- Step 5** Deploy the expected configuration using **Deploy**.



**Step 6** Depending on the configuration, if breakout ports or FEX ports are in use, you have to deploy again to completely restore the configuration.

**Step 7** After a successful deployment, and the device is “In-Sync,” you must move the device back to Normal Mode.



## RMA for User with Local Authentication



**Note** This task is only applicable to non-POAP switches.

Use the following steps to perform RMA for a user with local authentication:

### Procedure

- 
- Step 1** After the new switch comes online, SSH into the switch and reset the local user passwords with the cleartext password using the “username” command. Reset the local user passwords to resync the SNMP password. The password is stored in the configuration file in a nontransferable form.
- Step 2** Wait for the RMA to complete.
- Step 3** Update Cisco DCNM switch\_snmp\_user policy for the switch with the new SNMP MD5 key from the switch.
- 

## Interfaces

The Interfaces option displays all the interfaces that are discovered for the switch, Virtual Port Channels (vPCs), and intended interfaces missing on the device.

You can use the following functions:

- Create, deploy, view, edit and delete a port channel, vPC, Straight-through FEX, Active-Active FEX, loopback, and subinterface.



### Note

- The following features are unsupported for the brownfield migration of switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images:

- FEX on switches other than Cisco Nexus 9300 Series switches and Cisco Nexus 9500 Series switches with X9500 line cards
- AA-FEX

For information about the platform support for FEX, refer to your platform and NX-OS documentation to check the feature compatibility.

- To edit interfaces associated with fabric links such as intra-fabric links and inter-fabric links, see [Editing Interfaces Associated with Links, on page 172](#).

- 
- Create tunnel interfaces for Cisco Cloud Services Router 1000v Series (Cisco CSR 1000v Series).
  - Create breakout and unbreakout ports.
  - Shut down and bring up interfaces.
  - Rediscover ports and view interface configuration history.

- Apply host policies on interfaces and vPCs. For example, int\_trunk\_host\_11\_1, int\_access\_host\_11\_1, and so on.
- View interface information such as its admin status, operation status, reason, policy, speed, MTU, mode, VLANs, IP/Prefix, VRF, port channel, and the neighbor of the interface.

**Note**

- The **Neighbor** column provides details of connected switches that are discovered, intent links, and Virtual Machine Manager (VMM) connectivity. You can navigate to the **Switch** dashboard of the corresponding switch by clicking it. However, intent links and VMM links are not hyperlinked and you cannot navigate to the corresponding **Switch** dashboard.
- Click the graph icon in the Name column to view the interface performance chart for the last 24 hours. However, note that performance data for VLAN interfaces that are associated with overlay networks is not displayed in this chart.

The **Status** column displays the following statuses of an interface:

- Blue: Pending
- Green: In Sync/Success
- Red: Out-of-Sync/Failed
- Yellow: In Progress
- Grey: Unknown/NA

You can filter and view information for any of the given fields (such as Device Name). The following table describes the buttons that appear on this page.

**Note**

- Ensure that appropriate configurations are deployed through the Fabric Builder option before deploying from the Interfaces option, including proper vPC pair configurations. If you add or edit an interface before fabric deployment, the configuration may fail on the device.
- You can also manage interfaces from the Fabric Builder topology screen. Right click the switch and on the Manage Interfaces option. You can manage the interfaces per switch. If the switch is part of a vPC Pair, then interfaces from both peers are displayed on the page.
- Deploy any underlays including vPC Pairing in the fabric before deploying any configurations from the interface manager.

| Field | Description                                                                                                                            |
|-------|----------------------------------------------------------------------------------------------------------------------------------------|
| Add   | Allows you to add a logical interface such as a port channel, vPC, Straight-through FEX, Active-Active FEX, loopback and subinterface. |

| Field                | Description                                                                                                                                                                        |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Breakout, Unbreakout | Allows you to <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.                                                                                                                    |
| Deploy               | Allows you to deploy or redeploy saved interface configurations.                                                                                                                   |

If you perform admin operations (shutdown/no shutdown) on SVI, which is part of a config profile, successive **Save & Deploy** operations generate **no interface vlan** command.

For SVI with no policy, on performing admin operation, that is, shutdown/no shutdown command pushed from **Interface Manager**, **int\_vlan\_admin\_state** policy is associated with the SVI.

For example, create and deploy the SVI from **switch\_freeform**.

```
interface vlan1234
  description test
  no shutdown
  no ip redirects
  no ipv6 redirects
```

If you shutdown the SVI from interface manager, the **int\_vlan\_admin\_state** policy is associated with the SVI.

Pending diff is shown as:

```
interface Vlan1234
  shutdown
  no ip redirects
  no ipv6 redirects
  description test
  no shutdown
```

Remove the **no shutdown** CLI from the free-form config.

If the user has performed admin operation on SVI, device will have interface in running config. Therefore, post network detach **interface vlan** will be still present and interface will be discovered. You need to manually delete the interface from **Interface Manager**.

This section contains the following:

## Adding Interfaces

To add the interfaces from the Cisco DCNM Web UI, perform the following steps:

### Procedure

**Step 1** Choose **Control > Interfaces**.

You see the **Scope** option at the top right. If you want to view interfaces for a specific fabric, select the fabric window from the list.

**Step 2** Click **Add** to add a logical interface.

The **Add Interface** window appears.

**Step 3** In the **Type** drop-down list, choose the type of the interface.

Valid values are Port Channel, virtual Port Channel (vPC), Straight-through (ST) FEX, Active-Active (AA) FEX, Loopback, Subinterface, Tunnel. The respective interface ID field (Port-channel ID, vPC ID, Loopback ID, Subinterface ID, or Tunnel ID) is displayed when you select an interface Type. For example, port channel, Straight-through FEX, Active-Active FEX, vPC, loopback, and subinterface.

- When you create a port channel through DCNM, add interfaces of the same speed. A port channel that is created from interfaces of varying speeds won't come up. For example, a port channel with two *10 Gigabit Ethernet* ports is valid. However, a port channel with a *10-Gigabit Ethernet + 25-Gigabit Ethernet* port combination isn't valid.
- To add vPC hosts, you must designate vPC switches in the fabric topology (through the Fabric Builder) and deploy vPC and peer-link configurations using the **Save and Deploy** option. Once the vPC pair configurations are deployed, it appears in the Select a vPC pair drop-down box.

You can create a vPC using the **int\_vpc\_trunk\_host\_11\_1** policy.

- When adding a subinterface, you must select a routed interface from the interface table before clicking the Add button.

**Step 4** In the **Select a Device** field, choose the device.

Devices are listed based on the fabric and interface type. External fabric devices aren't listed for ST FEX and AA FEX. In the case of vPC or Active to Active FEX, select the vPC switch pair.

**Step 5** Enter the ID value in the respective interface ID field (**Port-channel ID**, **vPC ID**, **Loopback ID** and **Subinterface ID**) that is displayed, based on the selected interface.

You can override this value. The new value is used only if it's available in the Resource Manager pool. Else, it results in an error.

**Step 6** In the **Policy** field, you can select the policy to be applied on an interface.

The field only lists the Interface Python Policy with tag *interface\_edit\_policy* and filtered based on the interface type.

You must not create a **\_upg** interface policy. For example, you shouldn't create a policy using the **vpc\_trunk\_host\_upg**, **port\_channel\_aa\_fex\_upg**, **port\_channel\_trunk\_host\_upg**, and **trunk\_host\_upg** options.

**Step 7** Click **Save** to save the configurations.

Only saved configurations are pushed to the device. While adding the interface, you can only modify the policy attribute after the first save. If you try to use an ID that is already used, you encounter the *Resource could not be allocated* error.

**Step 8** (Optional) Click the **Preview** option to preview the configurations to be deployed.

**Step 9** Click **Deploy** to deploy the specified logical interface.

The newly added interface appears in the screen.

**Breakout or Unbreakout:** You can break out and unbreakout an interface by using the **breakout** option at the top left.

## Editing Interfaces

To edit the interfaces from the Cisco DCNM Web UI, perform the following steps:



**Note** The **Edit Interface** allows you to change the policy and add or remove an interface from a port channel or vPC.

### Procedure

**Step 1** Choose **Control > Interfaces**.

You can break out and unbreak out an interface by using the breakout option at the top left part of the screen.

**Step 2** Select the interface check box to edit an interface or vPC.

Select corresponding check boxes for editing multiple interfaces. You cannot edit multiple port channels and vPC. You cannot edit interfaces of different types at the same time.

**Step 3** Click **Edit** to edit an interface.

The variables that are shown in the **Edit Configuration** window are based on the template and its policy. Select the appropriate policy. Preview the policy, save it and deploy the same. This window lists only Interface Python Policy with the tag *interface\_edit\_policy* and filtered based on the interface type.

In a vPC setup, the two switches are in the order the switch names are displayed in the edit window. For example, if Switch Name is displayed as *LEAF1:LEAF2*, then Leaf1 is peer switch one and Leaf2 is peer switch two.

During overlay network deployment on switches, the network can be associated with trunk interfaces. The trunk interface to network association is reflected in the **Interfaces** screen. You can update such interfaces.

For interface policies that are not created from the **Control > Interfaces** screen, you can edit some configurations but not change the policy itself. The policy and fields that cannot be edited are grayed out.

The following are some examples of policies that cannot be edited:

- Loopback interface policies - The `int_fabric_loopback_11_1` policy is used to create a loopback interface. You can edit the loopback IP address and description but not the `int_fabric_loopback_11_1` policy instance.
- Fabric underlay network interface policies (`int_fabric_num_11_1`, for example) and fabric overlay network interface (NVE) policies.
- Policies associated with port channels and member ports of port channels, including the port channels and member ports associated with a vPC.
- SVIs created during network and VRF creation. The associated VLANs appear in the interfaces list.

---

## Editing Interfaces Associated with Links

There are two types of links, namely intra-fabric links and inter-fabric links. As the name implies, intra-fabric links are set up between devices within the same Easy fabric and are typically used for spine-leaf connectivity. Inter-fabric links are set up between the Easy fabric, and typically other external or Easy fabrics. They are used for external WAN and/or DCI connectivity. A policy is associated with each link that effectively states the configuration that is applied to both ends of the link. In other words, the link policy becomes the parent of the individual child interface policies that are associated with the two interfaces that form the link. In this scenario, you must edit the link policy to edit the interface policy fields such as description, IP address, and any per interface freeform config. The following procedure shows how to edit the interfaces associated with links:

### Procedure

---

- Step 1** Choose **Control > Fabric Builder**, and select the fabric containing the link.
- Step 2** Click **Tabular view** in the **Actions** panel.  
A window with the **Switches** and **Links** tabs appears.
- Step 3** Click the **Links** tab.
- Step 4** Select the link that you want to edit and click the **Update Link** icon.

The screenshot shows the Cisco Data Center Network Manager (DCNM) interface. On the left, a table lists various links under the 'Links' tab. The 'Link Management - Edit Link' dialog is open, allowing configuration of a specific link. The dialog includes dropdown menus for Link Type (Intra-Fabric), Link Sub-Type (Fabric), Link Template (int\_intra\_fabric\_num\_link\_11\_1), Source Fabric (site2), Destination Fabric (site2), Source Device (ste-n9k-9), Source Interface (Ethernet1/39), Destination Device (ste-n9k-18-deep), and Destination Interface (Ethernet1/39). The Link Profile section is currently on the 'General' tab, showing fields for 'Source Interface Desc...', 'Destination Interface ...', and 'Source Interface Free...'. A note on the right states: 'Note 1 A, strictly match with respect I. Any mismatch: unexpected c'. A 'Save' button is at the bottom right of the dialog.

Update the link based on your requirements and click **Save**.

## Deleting Interfaces

To delete the interfaces from the Cisco DCNM Web UI, perform the following steps:



**Note** This option allows you to delete only logical ports, port channels, and vPCs. You can delete the interface if it does not have overlay or underlay policy attached.

When a port channel or vPC is removed, the corresponding member ports get the default policy associated. The Default Policy can be configured in `server.properties` file.

### Procedure

**Step 1** Choose **Control > Interfaces**.

**Step 2** Select the interfaces.

**Step 3** Click **Delete**.

You cannot delete logical interfaces created in the fabric underlay.

## Shutting Down and Bringing Up Interfaces

To shut down and bring up the interfaces from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Interfaces**.
  - Step 2** Select the interfaces that you want to shut down or bring up.
  - Step 3** Click **Shutdown** to disable the selected interfaces. For example, you may want to isolate a host from the network or a host that is not active in the network.
  - Step 4** Click **No Shutdown** to bring up the selected interfaces.
- 

## Viewing Interface Configuration

To view the interface configuration commands and execute them from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Interfaces**.  
Select the interface whose configurations you want to view.
  - Step 2** In the **Interface Show Commands** window, select the action from the **Show** drop-down box and click **Execute**. The interface configurations are displayed in the **Output** section, at the right of the screen.  
For Show commands, you must have corresponding **show** templates for interface or interface sub types like port channel or vPC, defined in the **Template Library**.
- 

## Rediscovering Interfaces

To rediscover the interfaces from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Interfaces**.
  - Step 2** Select the interfaces that you want to rediscover.
  - Step 3** Click **Rediscover** to rediscover the selected interfaces. For example, after you edit or enable an interface, you can rediscover the interface.
- 

## Viewing Interface History

To view the interface history from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Interfaces**.
  - Step 2** Select the interface.
  - Step 3** Click **Interface History** to view the configuration history on the interface.
  - Step 4** Click **Status** to view each command that is configured for that configuration instance.
- 

## Deploying Interface Configurations

To deploy the interface configuration from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Interfaces**.
  - Step 2** Choose an interface you want to deploy.
    - Note** You can select multiple interfaces and deploy pending configurations.
  - Step 3** Click **Deploy** to deploy or redeploy configurations that are saved for an interface.
- 

## Creating External Fabric Interfaces

You can add and edit port channel, vPC, subinterface, and loopback interfaces for external fabric devices. You cannot add Straight-through FEX and Active-Active FEX functions.

The Breakout port function is only supported for Cisco Nexus 9000 and 3000 series switches in the external fabric.

When you add an interface to an external fabric device, the Resource Manager is not in sync with the device. So, ensure that the value populated in the ID field (Port-channel ID, vPC ID, Loopback ID, etc) is not previously configured on the switch.

If you want to configure a portchannel in the external fabric, you should add and deploy the **feature\_lacp** policy on the switches where the portchannel will be configured.

### Add Policy ✕

\* Priority (1-1000):

\* Policy:  ▼

feature\_lACP

---

Variables:

When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on its switches. If you click **Save & Deploy** in the fabric topology screen, it displays an error message. However, the following settings (available when you right-click the switch icon) are allowed:

vPC pairing - You can designate a vPC switch pair, but it is only for reference.

View/edit policy - You can add a policy but you cannot deploy it on the switch.

Manage interfaces – You can only create intent for adding interfaces. If you try to deploy, edit, or delete interfaces, it results in an error message.

## Creating and Deploying Networks and VRFs

The steps for overlay networks and VRFs provisioning are:

1. Create networks and VRFs for the fabric.
2. Deploy the networks and VRFs on the fabric switches.



### Note

The undeployment and deletion of overlay networks and VRFs are explained after the explanation of deployment. Finally, creation of external fabrics and fabric extensions from VXLAN to external fabrics are documented.

You can navigate to the networks and VRFs window through any of the following options:

- From the home page: Click the **Networks & VRFs** button in the Cisco DCNM Web UI landing page.
- From the Control menu: From the home page of the Cisco DCNM Web UI, choose **Control > Fabrics > Networks** to navigate to the **Networks** window. Choose **Control > Fabrics > VRFs** to navigate to the **VRFs** window.

You can toggle between the network view and VRF view in both the windows by clicking the **VRF View** or **Network View** button. When you are in the networks or VRFs window, ensure you choose the appropriate fabric from the **Scope** drop-down list before you create any networks or VRFs.

## Viewing Networks and VRFs for a Fabric

- Click **Control** > **Networks** from the main menu.

The **Networks** screen comes up. The **SCOPE** drop down box (at the top right part of the screen) lists all fabrics managed by the DCNM instance, in alphabetical order. You can choose the correct fabric from **SCOPE**. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

Networks

Fabric Selected: bgp2

| Network Name                                        | Network ID | VRF Name | IPv4 Gateway/Subnet | IPv6 Gateway/Prefix | Status | VLAN ID |
|-----------------------------------------------------|------------|----------|---------------------|---------------------|--------|---------|
| <input checked="" type="checkbox"/> MyNetwork_30000 | 30000      | NA       |                     |                     | NA     |         |

- Click **Control** > **VRFs** from the main menu.

The **VRFs** screen comes up. The **SCOPE** drop down box (at the top right part of the screen) lists all fabrics managed by the DCNM instance, in alphabetical order. You can choose the correct fabric from **SCOPE**. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

VRFs

Fabric Selected: bgp2

| VRF Name                                        | VRF ID | Status |
|-------------------------------------------------|--------|--------|
| <input checked="" type="checkbox"/> MyVRF_50000 | 50000  | NA     |



**Note** The **Networks** or **VRFs** windows are applicable only for the Easy or MSD fabrics.

## Creating Networks for the Standalone Fabric

1. Click **Control** > **Networks** (under **Fabrics** submenu).

The **Networks** screen comes up.

2. Choose the correct fabric from **SCOPE**. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

SCOPE: bgp2 admin

Network / VRF Selection > Network / VRF Deployment > VRF View | Continue

Fabric Selected: bgp2

Networks Selected 1 / Total 1

|                                     | Network Name    | Network ID | VRF Name | IPv4 Gateway/Subnet | IPv6 Gateway/Prefix | Status | VLAN ID |
|-------------------------------------|-----------------|------------|----------|---------------------|---------------------|--------|---------|
| <input checked="" type="checkbox"/> | MyNetwork_30000 | 30000      | NA       |                     |                     | NA     |         |

- 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:  +
- Layer 2 Only:
- \* Network Template: Default\_Network\_Universal
- \* Network Extension Template: Default\_Network\_Extension\_Univer
- VLAN ID:

Network Profile

Generate Multicast IP *Please click only to generate a New Multicast Group Address and override the default value!*

General

- IPv4 Gateway/NetMask:  ? example 192.0.2.1/24
- IPv6 Gateway/Prefix:  ? example 2001:db8::1/64
- Vlan Name:  ?
- Interface Description:  ?
- MTU for L3 interface:  ? [68-9216]
- IPv4 Secondary GW:  ? example 192.0.2.1/24

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:** A universal template is autopopulated. This is only applicable for leaf switches.

**Network Extension Template:** A universal extension template is autopopulated. This allows you to extend this network to another fabric. The methods are VRF Lite, Multi Site, and so on. The template is applicable for border leaf switches and BGWs.

**VLAN ID:** Specifies the corresponding tenant VLAN ID for the network.

**Network Profile** section contains the *General* and *Advanced* tabs.

**General** tab

**IPv4 Gateway/NetMask:** Specifies the IPv4 address with subnet.



**Note** If the same IP address is configured in the IPv4 Gateway and IPv4 Secondary GW1 or GW2 fields of the network template, DCNM does not show an error, and you will be able to save this configuration. However, after the network configuration is pushed to the switch, it would result in a failure as the configuration is not allowed by the switch.

**IPv6 Gateway/Prefix:** Specifies the IPv6 address with subnet.

Specify the anycast gateway IP address for transporting the L3 traffic from a server belonging to MyNetwork\_30000 and a server from another virtual network. By default the anycast gateway IP address is the same for MyNetwork\_30000 on all switches of the fabric that have the presence of the network.

**VLAN Name** - Enter the VLAN name.

**Interface Description:** Specifies the description for the interface. This interface is a switch virtual interface (SVI).

**MTU for the L3 interface** - Enter the MTU for Layer 3 interfaces.

**IPv4 Secondary GW1** - Enter the gateway IP address for the additional subnet.

**IPv4 Secondary GW2** - Enter the gateway IP address for the additional subnet.

**Advanced** tab: Optionally, specify the advanced profile settings by clicking the **Advanced** tab:

**ARP Suppression** – Select the checkbox to enable the ARP Suppression function.

**Ingress Replication** - The checkbox is selected if the replication mode is Ingress replication.



**Note** Ingress Replication is a read-only option in the Advanced tab. Changing the fabric setting updates the field.

**Multicast Group Address-** The multicast IP address for the network is autopopulated.

Multicast group address is a per fabric instance variable. The number of underlay multicast groups supported is only 128. If all networks are deployed on all switches, you need not use a different multicast group per L2 VNI or a network. Therefore, multicast group for all networks in a fabric remains same. If a new multicast group address is required, you can generate it by clicking the **Generate Multicast IP** button.

**DHCPv4 Server 1** - Enter the DHCP relay IP address of the first DHCP server.

**DHCPv4 Server 2** - Enter the DHCP relay IP address of the next DHCP server.

**DHCPv4 Server VRF**- Enter the DHCP server VRF ID.

**Routing Tag** – The routing tag is autopopulated. This tag is associated with each gateway IP address prefix.

**TRM enable** – Select the checkbox to enable TRM.

**L2 VNI Route-Target Both Enable** - Select the check box to enable automatic importing and exporting of route targets for all L2 virtual networks.

**Enable L3 Gateway on Border** - Select the checkbox to enable a Layer 3 gateway on the border switches.

A sample of the Create Network screen is given below.

▼ Network Profile

Generate Multicast IP *ⓘ Please click only to generate a New Multicast Group Address and override the default value!*

|          |                       |              |                          |
|----------|-----------------------|--------------|--------------------------|
| General  | IPv4 Gateway/NetMask  | 20.10.1.1/24 | ? example 192.0.2.1/24   |
| Advanced | IPv6 Gateway/Prefix   |              | ? example 2001:db8::1/64 |
|          | Vlan Name             | Drill        | ?                        |
|          | Interface Description |              | ?                        |
|          | MTU for L3 interface  |              | ? [68-9216]              |
|          | IPv4 Secondary GW1    | 20.10.2.1/24 | ? example 192.0.2.1/24   |
|          | IPv4 Secondary GW2    | 20.10.3.1/24 | ? example 192.0.2.1/24   |

Create Network

▼ Network Profile

Generate Multicast IP *ⓘ Please click only to generate a New Multicast Group Address and override the default value!*

|          |                                      |                                     |                                              |
|----------|--------------------------------------|-------------------------------------|----------------------------------------------|
| General  | ARP Suppression                      | <input type="checkbox"/>            | ?                                            |
| Advanced | Ingress Replication                  | <input type="checkbox"/>            | ? Read-only per network, Fabric-wide setting |
|          | Multicast Group Address              | 239.1.1.0                           | ?                                            |
|          | DHCPv4 Server 1                      | 20.20.20.1                          | ? DHCP Relay IP                              |
|          | DHCPv4 Server 2                      | 20.20.30.1                          | ? DHCP Relay IP                              |
|          | DHCPv4 Server VRF                    | Foo                                 | ?                                            |
|          | Loopback ID for DHCP Relay interface | 4                                   | ?                                            |
|          | Routing Tag                          | 12345                               | ? [0-4294967295]                             |
|          | TRM Enable                           | <input type="checkbox"/>            | ? Enable Tenant Routed Multicast             |
|          | L2 VNI Route-Target Both Enable      | <input checked="" type="checkbox"/> | ?                                            |
|          | Enable L3 Gateway on Border          | <input checked="" type="checkbox"/> | ?                                            |

- Click **Create Network**. A message appears at the bottom right part of the screen indicating that the network is created.

The new network appears on the **Networks** page that comes up.

Fabric Selected: Standalone

Selected 1 / Total 1

| Network Name                                        | Network ID | VRF Name    | IPv4 Gateway/Subnet | IPv6 Gateway/Prefix | Status | VLAN ID |
|-----------------------------------------------------|------------|-------------|---------------------|---------------------|--------|---------|
| <input checked="" type="checkbox"/> MyNetwork_30000 | 30000      | MyVRF_50000 | 20.10.1.1/24        |                     | NA     |         |

The Status is *NA* since the network is created but not yet deployed on the switches. Now that the network is created, you can create more networks if needed and deploy the networks on the devices in the fabric.

## Export and Import Network Information

You can export network information to a .CSV file. The exported file contains information pertaining to each network, including the fabric it belongs to, the associated VRF, the network templates used to create the network, and all other configuration details that you saved during network creation.

In the Networks screen, click the Export icon to export network information as a .CSV file.

| A          | B           | C               | D         |
|------------|-------------|-----------------|-----------|
| fabric     | vrf         | networkName     | networkId |
| Standalone | MyVRF_50000 | MyNetwork_30000 | 30000     |
| Standalone | MyVRF_50000 | MyNetwork_30001 | 30001     |

You can use the exported .CSV file for reference or use it as a template for creating new networks. To import networks, do the following:

1. Update new records in the .CSV file. Ensure that the `networkTemplateConfig` field contains the JSON Object. A message at the bottom right part of the screen displays errors and success messages. This screenshot depicts two new networks being imported.

| A          | B           | C               | D         | E                         | F                                   | G                                                              | H | I | J | K | VLAN ID |
|------------|-------------|-----------------|-----------|---------------------------|-------------------------------------|----------------------------------------------------------------|---|---|---|---|---------|
| fabric     | vrf         | networkName     | networkId | networkTemplate           | networkExtensionTemplate            | networkTemplateConfig                                          |   |   |   |   |         |
| Standalone | MyVRF_50000 | MyNetwork_30002 | 30002     | Default_Network_Universal | Default_Network_Extension_Universal | ("suppressArp":"false","secondaryGW2":"","secondaryGW1":"","") |   |   |   |   |         |
| Standalone | MyVRF_50000 | MyNetwork_30003 | 30003     | Default_Network_Universal | Default_Network_Extension_Universal | ("suppressArp":"false","secondaryGW2":"","secondaryGW1":"","") |   |   |   |   |         |

2. In the Networks screen, click the Import icon and import the .CSV file into DCNM.

You can see that the imported networks are displayed in the Networks screen.

Networks Selected 0 / Total 4

| <input type="checkbox"/> | Network Name    | Network ID | VRF Name    | IPv4 Gateway/Subnet | IPv6 Gateway/Prefix | Status | VLAN ID |
|--------------------------|-----------------|------------|-------------|---------------------|---------------------|--------|---------|
| <input type="checkbox"/> | MyNetwork_30000 | 30000      | MyVRF_50000 | 20.10.1.1/24        |                     | NA     |         |
| <input type="checkbox"/> | MyNetwork_30001 | 30001      | MyVRF_50000 |                     |                     | NA     |         |
| <input type="checkbox"/> | MyNetwork_30002 | 30002      | MyVRF_50000 | 20.10.4.1/24        |                     | NA     |         |
| <input type="checkbox"/> | MyNetwork_30003 | 30003      | MyVRF_50000 |                     |                     | NA     |         |

## Editing Networks for the Standalone Fabric

To edit networks for standalone fabrics from Cisco DCNM Web UI, perform the following steps:

### Procedure

- Step 1** Click **Control > Networks**.  
The **Networks** window appears.
- Step 2** Choose a fabric from the **SCOPE** drop-down list.  
The **Networks** window refreshes and lists the networks in the fabric.
- Step 3** Choose a network.
- Step 4** Click the **Edit** icon.  
The **Edit Network** window appears.
- Step 5** Update the fields in the **General** and **Advanced** tabs of the **Network Profile** area as needed.
- Step 6** Click **Save** at the bottom right part of the window to save the updates.

## Creating VRFs for the Standalone Fabric

1. Click **Control > VRFs** (under **Fabrics** submenu).  
The VRFs screen comes up.
2. Choose the correct fabric from **SCOPE**. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

Data Center Network Manager SCOPE: bgp2 admin

Network / VRF Selection > Network / VRF Deployment > Network View | Continue

Fabric Selected: bgp2

VRFs Selected 1 / Total 1

| <input type="checkbox"/>            | VRF Name    | VRF ID | Status |
|-------------------------------------|-------------|--------|--------|
| <input checked="" type="checkbox"/> | MyVRF_50000 | 50000  | NA     |

3. Click the + button to add VRFs to the *Standalone* fabric. The Create VRF screen comes up. Most of the fields are autopopulated.

Create VRF
✕

---

▼ VRF Information

\* VRF ID

\* VRF Name

\* VRF Template  ▼

\* VRF Extension Template  ▼

---

▼ VRF Profile

General

Advanced

VRF Vlan Name  ?

VRF Intf Description  ?

VRF Description  ?

Create VRF

The fields in this screen are:

**VRF ID** and **VRF Name**: The ID and name of the VRF.



**Note** For ease of use, the VRF creation option is also available while you create a network.

**VRF Template**: This template is applicable for VRF creation, and only applicable for leaf switches.

**VRF Extension Template**: The template is applicable when you extend the VRF to other fabrics, and is applicable for border devices.

Fill the fields in the **VRF Profile** section.

**General** tab – Enter the VLAN ID of the VLAN associated with the VRF, the corresponding Layer 3 virtual interface, and the VRF ID.

**Advanced** tab – The fields in the tab are autopopulated.

**Routing Tag** – If a VLAN is associated with multiple subnets, then this tag is associated with the IP prefix of each subnet. Note that this routing tag is associated with overlay network creation too.

**Redistribute Direct Route Map** – Specifies the route map name for redistribution of routes in the VRF.

**Max BGP Paths** and **Max iBGP Paths** – Specifies the maximum BGP and iBGP paths.

**TRM Enable** – Select the checkbox to enable TRM.

If you enable TRM, then the RP address, the RP loopback ID and the underlay multicast address must be entered.

**Is RP External** – Enable this checkbox if the RP is external to the fabric.

**RP Address** and **RP Loopback ID** – Specifies the loopback ID and IP address of the RP.

**Underlay Multicast Address** – Specifies the multicast address associated with the VRF. The multicast address is used for transporting multicast traffic in the fabric underlay.



**Note** The multicast address in the **Multicast address for TRM** field in the fabric settings screen is populated in this field.

## Create VRF

### ▼ VRF Information

|                          |                                                              |
|--------------------------|--------------------------------------------------------------|
| * VRF ID                 | <input type="text" value="50000"/>                           |
| * VRF Name               | <input type="text" value="MyVRF_50000"/>                     |
| * VRF Template           | <input type="text" value="Default_VRF_Universal"/>           |
| * VRF Extension Template | <input type="text" value="Default_VRF_Extension_Universal"/> |

### ▼ VRF Profile

|          |                                                                                                                                               |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| General  | <input type="checkbox"/> <b>Is RP External</b> <small>? Is RP external to the fabric?</small>                                                 |
| Advanced | <input type="text" value=""/> <b>RP Address</b> <small>? IPv4 Address</small>                                                                 |
|          | <input type="text" value=""/> <b>RP Loopback ID</b> <small>? 0-1023</small>                                                                   |
|          | <input type="text" value="239.1.1.0"/> <b>Underlay Mcast Add...</b> <small>? IPv4 Multicast Address</small>                                   |
|          | <input type="text" value=""/> <b>Overlay Mcast Groups</b> <small>? 224.0.0.0/8 to 239.255.255.255/8</small>                                   |
|          | <input checked="" type="checkbox"/> <b>Enable IPv6 link-loc...</b> <small>? Enables IPv6 link-local Option under VRF SVI</small>              |
|          | <input type="checkbox"/> <b>Advertise Host Routes</b> <small>? Flag to Control Advertisement of /32 and /128 Routes to Edge Routers</small>   |
|          | <input checked="" type="checkbox"/> <b>Advertise Default Route</b> <small>? Flag to Control Advertisement of Default Route Internally</small> |

**Overlay Multicast Groups** – Specifies the multicast address for the VRF, used in the fabric overlay.

**Enable IPv6 link-local Option** - Select the check box to enable the IPv6 link-local option under the VRF SVI. If this check box is unchecked, IPv6 forward is enabled.

**Advertise Host Routes** – Enable the checkbox to control advertisement of /32 and /128 routes to Edge Routers.

**Advertise Default Route** – Enable the checkbox to control advertisement of default routes internally.

To allow inter-subnet communication between end hosts in different VXLAN fabrics, where the subnets are present in both fabrics, you must disable the **Advertise Default Route** feature (clear the **Advertise Default Route** checkbox) for the associated VRF. This will result in /32 routes for hosts being seen in both fabrics. For example, Host1 (VNI 30000, VRF 50001) in Fabric1 can send traffic to Host2 (VNI 30001, VRF 50001) in Fabric2 only if the host route is present in both fabrics. When a subnet is present in only one fabric then default route is sufficient for inter-subnet communication.

Sample screenshots of the Create VRF screen:

## Create VRF



## ▼ VRF Information

\* VRF ID

\* VRF Name

\* VRF Template

\* VRF Extension Template

## ▼ VRF Profile

| General | Advanced                                                                                                                                                                                                     |
|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|         | <p>VRF Vlan Name <input type="text" value="vlan 2500"/> ?</p> <p>VRF Intf Description <input type="text" value="interface vlan 2500"/> ?</p> <p>VRF Description <input type="text" value="coke:vrf1"/> ?</p> |

Create VRF

Advanced tab:

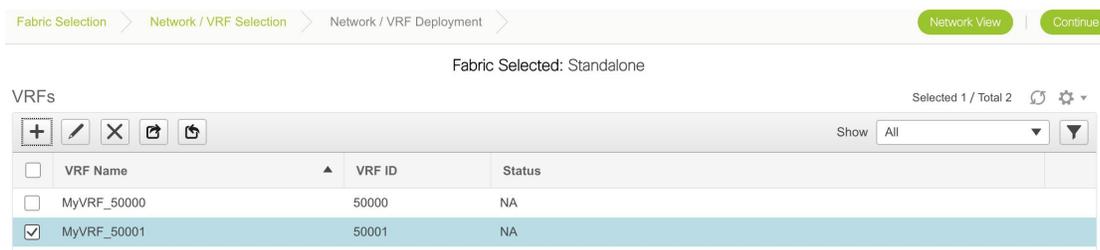
## ▼ VRF Profile

| General | Advanced                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|         | <p>Routing Tag <input type="text" value="12345"/> ? [0-4294967295]</p> <p>Redistribute Direct Route Map <input type="text" value="FABRIC-RMAP-REDIST-SUBNET"/> ?</p> <p>Max BGP Paths <input type="text" value="1"/> ? [1-64]</p> <p>Max iBGP Paths <input type="text" value="2"/> ? [1-64]</p> <p>TRM Enable <input type="checkbox"/> ? Enable Tenant Routed Multicast</p> <p>Is RP External <input type="checkbox"/> ? Is RP external to the fabric?</p> <p>RP Address <input type="text" value="224.0.0.2"/> ? IPv4 Address</p> <p>RP Loopback ID <input type="text" value="3"/> ? 0-1023</p> <p>Underlay Mcast Add... <input type="text" value="224.0.0.10"/> ? IPv4 Multicast Address</p> <p>Overlay Mcast Groups <input type="text" value="224.0.0.0/8"/> ? 224.0.0.0/8 to 239.255.255.255/8</p> <p>Enable IPv6 link-loc... <input type="checkbox"/> ? Enables IPv6 link-local Option under VRF SVI</p> <p>Advertise Host Routes <input type="checkbox"/> ? Flag to Control Advertisement of /32 and /128 Routes to Edge Routers</p> <p>Advertise Default Route <input checked="" type="checkbox"/> ? Flag to Control Advertisement of Default Route Internally</p> |

Create VRF

4. Click **Create VRF**.

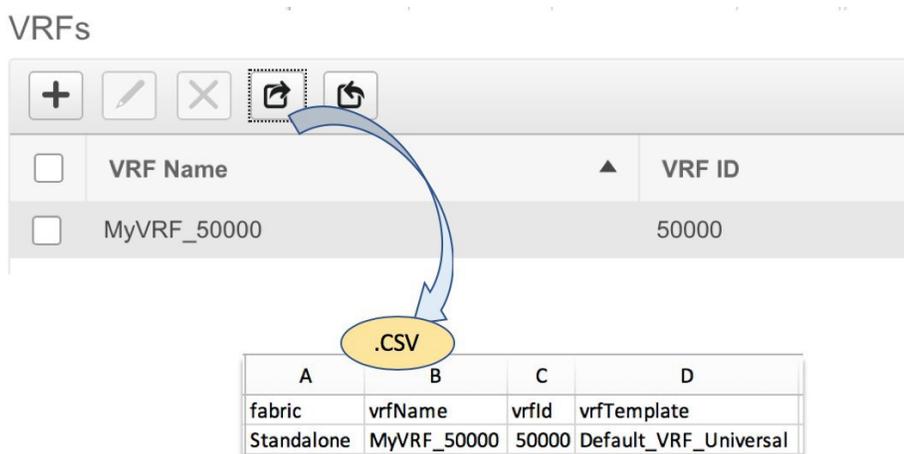
The *MyVRF\_50001* VRF is created and appears on the VRFs page.



## Export and Import VRF Information

You can export VRF information to a .CSV file. The exported file contains information pertaining to each VRF, including the fabric it belongs to, the templates used to create the VRF, and all other configuration details that you saved during VRF creation.

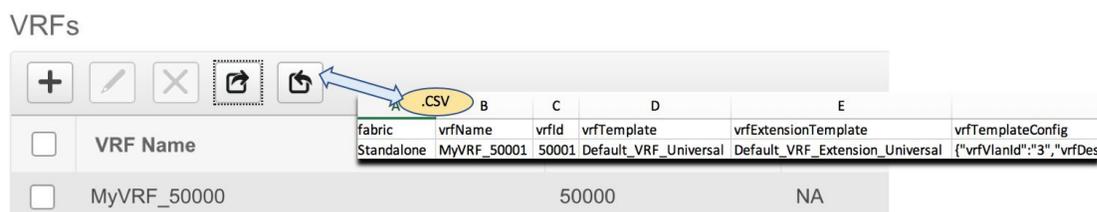
In the VRFs screen, click the Export icon to export VRF information as a .CSV file.



You can use the exported .CSV file for reference or use it as a template for creating new VRFs. To import VRFs, do the following:

1. Update new records in the .CSV file. Ensure that the **vrfTemplateConfig** field contains the JSON Object.
2. In the VRFs screen, click **Import** icon and import the .CSV file into DCNM.

A message at the bottom right part of the screen displays errors and success messages. This screenshot depicts a new VRF being imported.



You can see that the imported VRF is displayed in the VRFs screen.

VRFs Selected 0 / Total 2  

| <input type="checkbox"/> | VRF Name    | VRF ID | Status |
|--------------------------|-------------|--------|--------|
| <input type="checkbox"/> | MyVRF_50000 | 50000  | NA     |
| <input type="checkbox"/> | MyVRF_50001 | 50001  | NA     |

## Editing VRFs for the Standalone Fabric

1. Choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

Data Center Network Manager SCOPE: bgp2  admin 

Network / VRF Selection > Network / VRF Deployment > Network View | Continue

Fabric Selected: bgp2

VRFs Selected 1 / Total 1  

| <input type="checkbox"/>            | VRF Name    | VRF ID | Status |
|-------------------------------------|-------------|--------|--------|
| <input checked="" type="checkbox"/> | MyVRF_50000 | 50000  | NA     |

2. From the **Select a Fabric** drop-down list, select the fabric *Standalone*, and click **Continue** on the top right part of the screen. The Networks page is displayed.
3. Click the **VRF View** at the top right part of the screen. The VRFs page appears.

Fabric Selected: New7200

VRFs Selected 0 / Total 2  

| <input type="checkbox"/> | VRF Name    | VRF ID | Status |
|--------------------------|-------------|--------|--------|
| <input type="checkbox"/> | MyVRF_50000 | 50000  | NA     |
| <input type="checkbox"/> | MyVRF_50001 | 50001  | NA     |

4. Select the **VRF** and click the **Edit** option at the top left part of the screen. The **Edit VRF** screen comes up.

Edit VRF ✕

▼ VRF Information

\* VRF ID

\* VRF Name

\* VRF Template

VRF Extension Template

---

▼ VRF Profile

General

Advanced

VRF Vlan Name  ?

VRF Intf Description  ?

VRF Description  ?

[Save](#) [Cancel](#)

5. Update the fields in the **General** and **Advanced** tabs of the **VRF Profile** section as needed.
6. Click **Save** at the bottom right part of the screen to save the updates.

## Deploying Networks for the Standalone and MSD Fabrics

*Before you begin:* Ensure that you have created networks for the fabric.

1. Click **Control** > **Networks** (under **Fabrics** submenu).

The Networks screen comes up.

2. Choose the correct fabric from **SCOPE**. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

Data Center Network Manager SCOPE: bgp2 admin

Network / VRF Selection > Network / VRF Deployment > [VRF View](#) [Continue](#)

Fabric Selected: bgp2

Networks Selected 1 / Total 1

Show: All

| <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      | NA       |                     |                     | NA     |         |

3. Select networks that you want to deploy. In this case, select the check boxes next to both the networks and click **Continue** at the top right part of the screen.

The Network Deployment page appears. On this page, you can see the network topology of the Standalone fabric.

You can deploy networks simultaneously on multiple switches. The selected devices should have the same role (Leaf, Border Gateway, and so on).



---

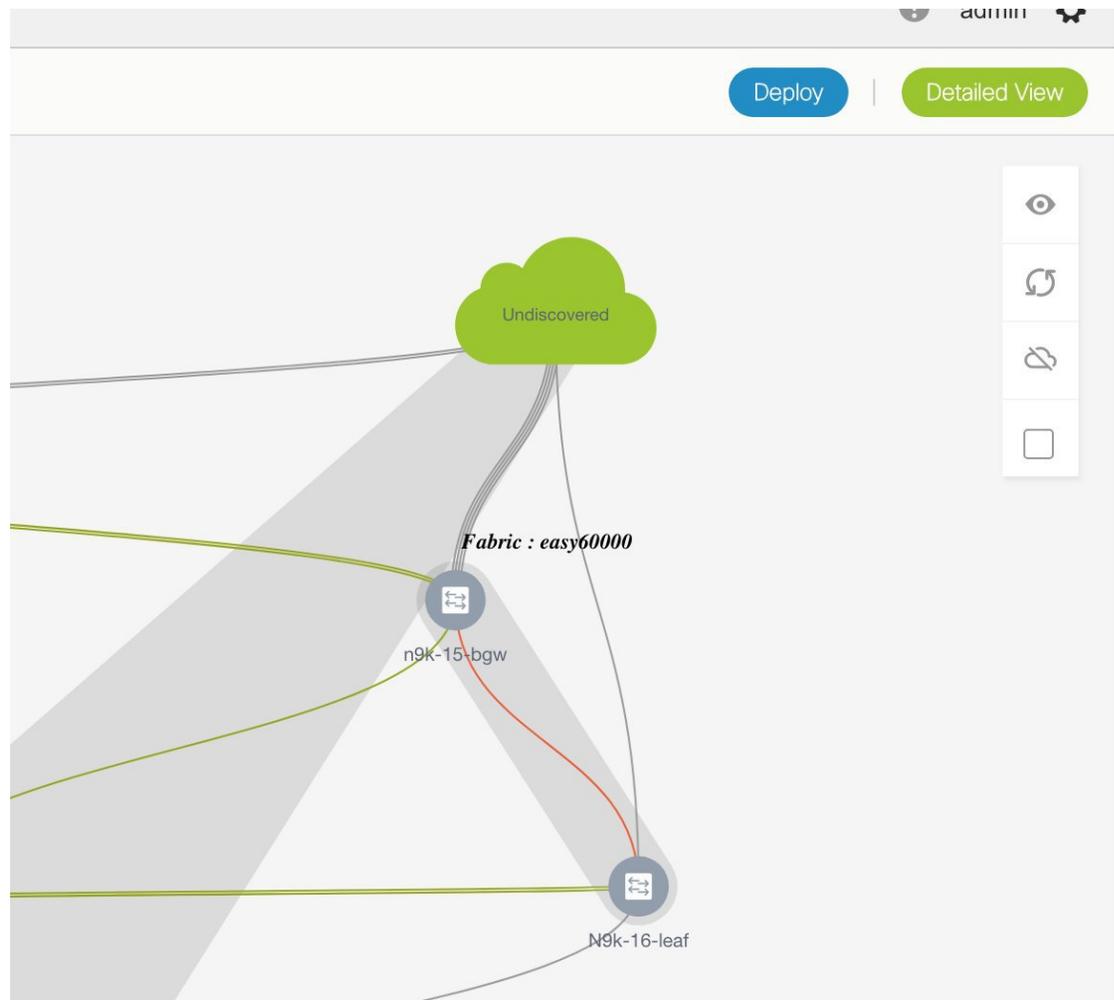
**Note** In an MSD fabric, all member fabrics are visible from this screen.

---

At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly. Blue for *Pending* state, yellow for *In Progress* when the provisioning is in progress, green when successfully deployed, and so on.

The overlay networks (/VRFs) provisioning status is context-specific. It is a combination of networks that you chose for provisioning and the relevant switches in the topology. In this example, it means that the networks *MyNetwork\_30000* and *MyNetwork\_30001* are yet to be deployed on any switch in this fabric.

**Undiscovered cloud** display – To display (or not display) an **Undiscovered** cloud in this screen, click the cloud icon in the vertical panel, at the top-right part of the screen. When you click the icon, the **Undiscovered** cloud and its links to the fabric topology are not displayed. Click the icon again to display the **Undiscovered** cloud.



You can move the topology around the screen by clicking the left mouse button on the screen and moving it in the direction you desire. You can enlarge or shrink the switch icons proportionately by moving the cursor roller. You can also use corresponding alternatives on the touchpad.

4. Click ... in the **Interfaces** column.

The **Interfaces** box opens up. It lists interfaces or port channels. You can select interfaces/port channels to associate them with the selected network. For each interface, port type and description, channel number and connected neighbor interface details are displayed.

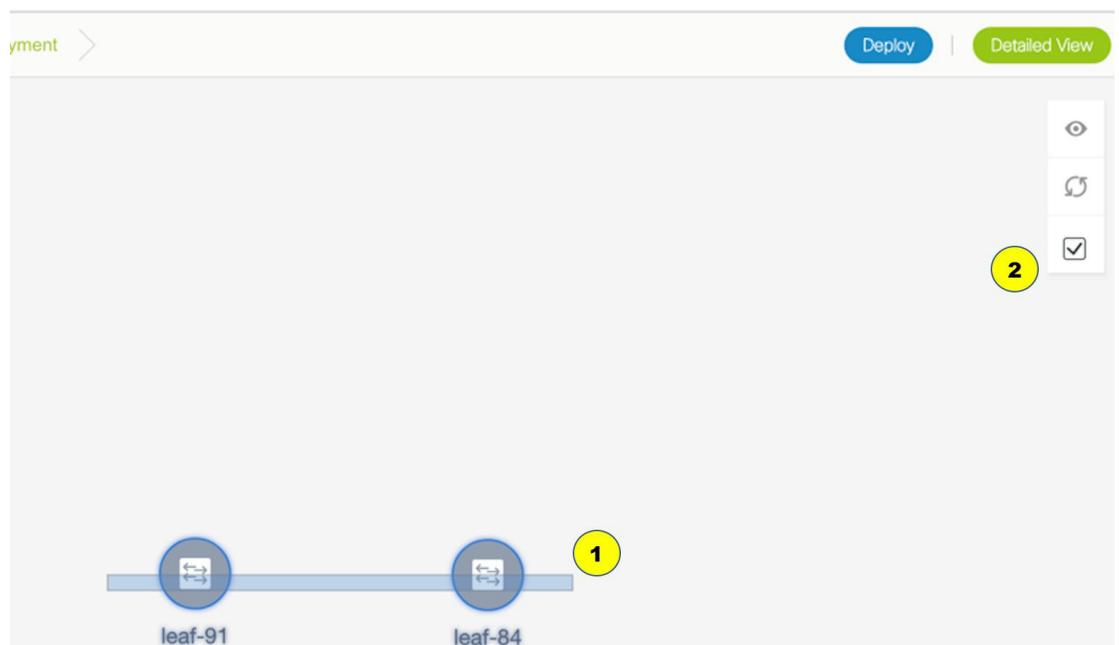
## Interfaces



| <input type="checkbox"/>            | Interface/Ports ▲ | Channel ... | Port Ty... | Port Desc... | Neighbor Info |
|-------------------------------------|-------------------|-------------|------------|--------------|---------------|
| <input type="checkbox"/>            | Ethernet1/1       | NA          | trunk      |              |               |
| <input checked="" type="checkbox"/> | Ethernet1/10      | NA          | trunk      |              |               |
| <input checked="" type="checkbox"/> | Ethernet1/11      | NA          | trunk      |              |               |
| <input type="checkbox"/>            | Ethernet1/12      | NA          | trunk      |              |               |
| <input type="checkbox"/>            | Ethernet1/13      | NA          | trunk      |              |               |

Save

- Double-click a switch to deploy the networks on it. For deployment of networks on multiple switches, click Multi-Select from the panel at the top right part of the screen (the topology freezes to a static state), and drag the cursor across the switches.



Immediately the Network Attachment dialog box appears.

Network Attachment - Attach networks for given switch(es) 

Fabric Name: Standalone

## Deployment Options

 Select the row and click on the cell to edit and save changes

| MyNetwork_30000          |             | MyNetwork_30001 |            |                 |        |  |
|--------------------------|-------------|-----------------|------------|-----------------|--------|--|
| <input type="checkbox"/> | Switch ▲    | VLAN            | Interfaces | CLI Freeform    | Status |  |
| <input type="checkbox"/> | n9k-16-leaf | 2300            | ...        | Freeform config | NA     |  |

Save

A tab represents each network (the first network is displayed by default) that is being deployed. In each network tab, the switches are displayed. Each row represents a switch.

Click the check box next to the **Switch** column to select all switches. The network is ready to be provisioned on the switches.

VLAN - Update the VLAN ID if needed.

When you update a VLAN ID and complete the network deployment process, the old VLAN is not automatically removed. To complete the process, you should go to the fabric topology screen (click **Control > Fabric Builder** and click within the corresponding fabric box to go to the screen) and use the Save and Deploy option.

When updating the VLAN ID for a given network, the original VLAN ID is not automatically removed from the attached trunk interface. In order to remove the old or original VLAN ID, you must perform **Save and Deploy + Config Deploy** operation from within the fabric in Fabric Builder. For this, go to the fabric topology screen (click **Control > Fabric Builder** and click within the corresponding fabric box to go to the screen) and execute the **Save and Deploy** operation. Verify that config compliance is removing the expected config, then execute **Deploy Config** operation to remove the configs.

Interfaces – Click ... in the column to add interfaces associated with the selected network.

VLAN to trunk port mapping – The selected trunk ports include the VLAN as an allowed VLAN on the port.

VLAN to vPC domain mapping - If you want to associate the VLAN to port channels of a vPC domain, add the port channels from the list of interfaces. The vPC port channels include the VLAN as an allowed VLAN.

Freeform configurations – Click Freeform config to enable additional configurations on the switch. After the configurations are saved, the Freeform config button gets highlighted.

6. Select the other network tab and make the same selections.

- Click **Save** (at the bottom right part of your screen) to save the configurations.



**Note** Addition and removal of interfaces are displayed in the **Interfaces** column of the Switches Deploy screen. Though the interface-related updates (like addition or removal of trunk ports) are provisioned on the switches, the correct configurations will not reflect in the preview screen. When you add or remove a trunk or access port, the preview shows the addition or removal of configurations for the interface under that network.

The topology window appears again. Click *Refresh* in the vertical panel at the top right part of the screen. The blue color on the switch icons indicates that the deployment is pending.

- Preview the configurations by clicking *Preview* (the eye icon above the Multi-Select option). Since *MyNetwork\_30000* and *MyNetwork\_30001* are networks of VRF *50000*, the configurations contain VRF configurations followed by the network configurations.

## Preview Configuration

Select a Switch:

n9k-16-leaf ▼

Select a Network

MyNetwork\_30000 ▼

Generated Configuration:

```
configure profile MyVRF_50000
vlan 2000
vn-segment 50000
interface vlan2000
vrf member myvrf_50000
ip forward
ipv6 forward
no ip redirects
no ipv6 redirects
mtu 9216
no shutdown
vrf context myvrf_50000
vni 50000
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
address-family ipv6 unicast
route-target both auto
route-target both auto evpn
router bgp 60000
vrf myvrf_50000
address-family ipv4 unicast
advertise I2vpn evpn
redistribute direct route-map fabric-rmap-redis-subnet
maximum-paths ibgp 2
address-family ipv6 unicast
advertise I2vpn evpn
redistribute direct route-map fabric-rmap-redis-subnet
maximum-paths ibgp 2
interface nve1
member vni 50000 associate-vrf
configure terminal
apply profile MyVRF_50000
```

**MyVRF\_50000  
Configuration**

## Preview Configuration

**Select a Switch:** n9k-16-leaf ▼

**Select a Network:** MyNetwork\_30000 ▼

**Generated Configuration:**

```
vrf myvrf_50000
address-family ipv4 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-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
```

**MyNetwork\_30000 Configuration**

```
configure profile MyNetwork_30000
vlan 2300
vn-segment 30000
interface vlan2300
vrf member myvrf_50000
fabric forwarding mode anycast-gateway
no shutdown
interface nve1
member vni 30000
mcast-group 239.1.1.0
evpn
vni 30000 l2
rd auto
route-target import auto
route-target export auto
configure terminal
apply profile MyNetwork_30000
```

**Interfaces Configuration**

```
interface ethernet1/11
switchport trunk allowed vlan add 2300
interface ethernet1/10
switchport trunk allowed vlan add 2300
```

On the preview screen, you can select from the **Select a switch** and **Select a network** drop-down boxes at the top of the screen to view other network configurations.

After checking the configurations, close the screen. The Topology screen appears again.

- Click **Deploy** on the top right part of the screen. The color of the switch icons changes to yellow and a message appears at the bottom right part of the screen indicating that the deployment is in progress. After the networks' deployment is complete, the color of the switch icons changes to green, indicating successful deployment.



**Note** In case you click **Deploy** and there is no configuration diff that has to be deployed, a pop-up window comes up stating **No switches PENDING for deployment**.



**Note** When you select multiple networks on the *Topology View* screen and proceed to the deployment screen, the switch color reflects the status of the first network in the selected list of networks. In this example, the switch color turns green when *MyNetwork\_30000* is provisioned on the switch.

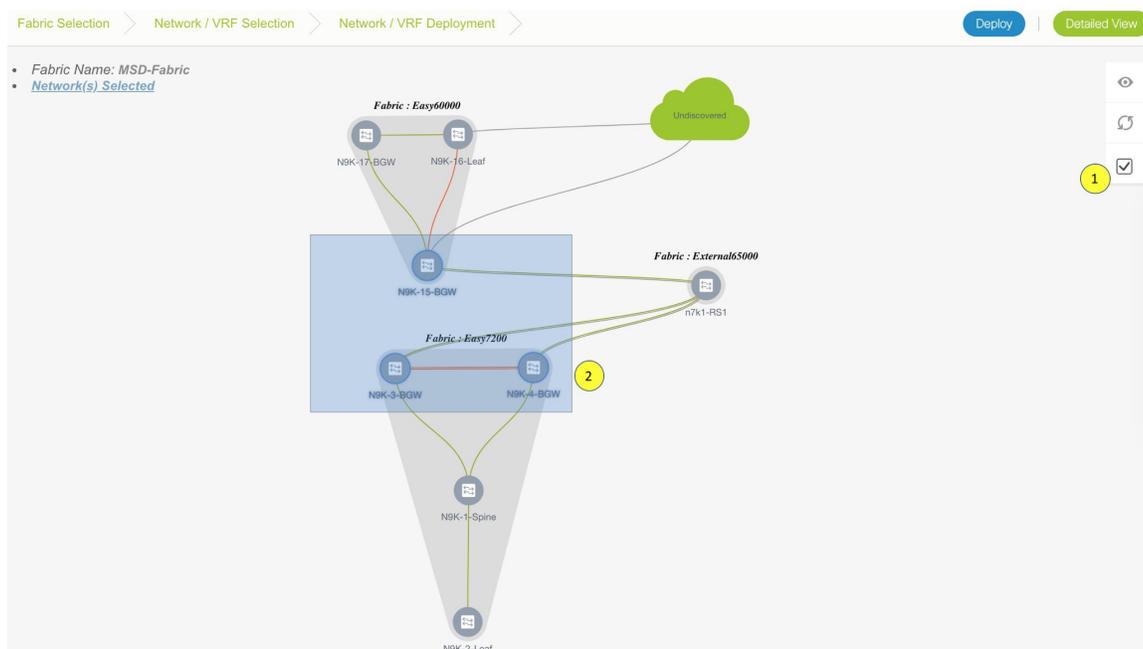
Go to the [Networks](#) page to view the individual status for all networks.

### Network Deployment for an MSD Fabric

Consider a scenario wherein you are deploying the same networks on different member fabric border devices. You can choose one fabric, deploy networks on its border devices, and then choose the second fabric and deploy networks.

Alternatively, you can choose the MSD fabric, and deploy the networks from a single topology view of all member fabric border devices.

This is a topology view of an MSD fabric wherein the two member fabrics topologies and their connections are depicted. You can deploy networks on the BGWs of the fabrics at once.



### Detailed View

You can also use the Detailed View option to deploy networks and VRFs. Click **Detailed View** at the top right part of the screen. The Detailed View window appears. This lists the networks in a tabular view.

| <input type="checkbox"/> | Name            | Switch      | Ports                     | Status   | Fabric Name | Role   |
|--------------------------|-----------------|-------------|---------------------------|----------|-------------|--------|
| <input type="checkbox"/> | MyNetwork_30000 | N9k-15-bgw  |                           | NA       | new60000    | border |
| <input type="checkbox"/> | MyNetwork_30001 | N9k-15-bgw  |                           | NA       | new60000    | border |
| <input type="checkbox"/> | MyNetwork_30001 | n9k-16-leaf | Ethernet1/1               | DEPLOYED | new60000    | leaf   |
| <input type="checkbox"/> | MyNetwork_30000 | n9k-16-leaf | Ethernet1/10,Ethernet1/11 | DEPLOYED | new60000    | leaf   |

The options:

Edit - Select a network and click the Edit icon at the top left part of the screen.



#### Note

If you select one network/switch entry and click on Edit, the Network Attach dialog box appears. To maintain consistency across the Topology View and Detailed View screens, the Network Attach screen displays all networks, and not just the selected network/switch.

Preview – Click Preview to preview configurations before deployment. You can only preview pending configurations, and not uninitiated or deployed configurations.

Deploy – Click Deploy to provision networks onto the switches.

History – Select a row and click History to view the configuration instances and status. Network and VRF-wise configurations are displayed. Click in the Status column of any instance for more details.

The fields in the table contain the configuration instance in each row, the associated switch and fabric names, the switch role, trunk ports (if any), and the deployment status.

Apply/Save – Selecting a network and clicking Apply/Save will select a switch for the network to be deployed on.

On the Detailed View page, the network profile configuration history is displayed. If you have associated specific trunk interfaces to that network, then the interface configuration is displayed as a separate configuration instance.



**Note** When you upgrade from an earlier release (such as DCNM 10.4[2]) to the DCNM 11.0(1) release, overlay networks and VRFs deployment history information from the earlier DCNM release is not retained.

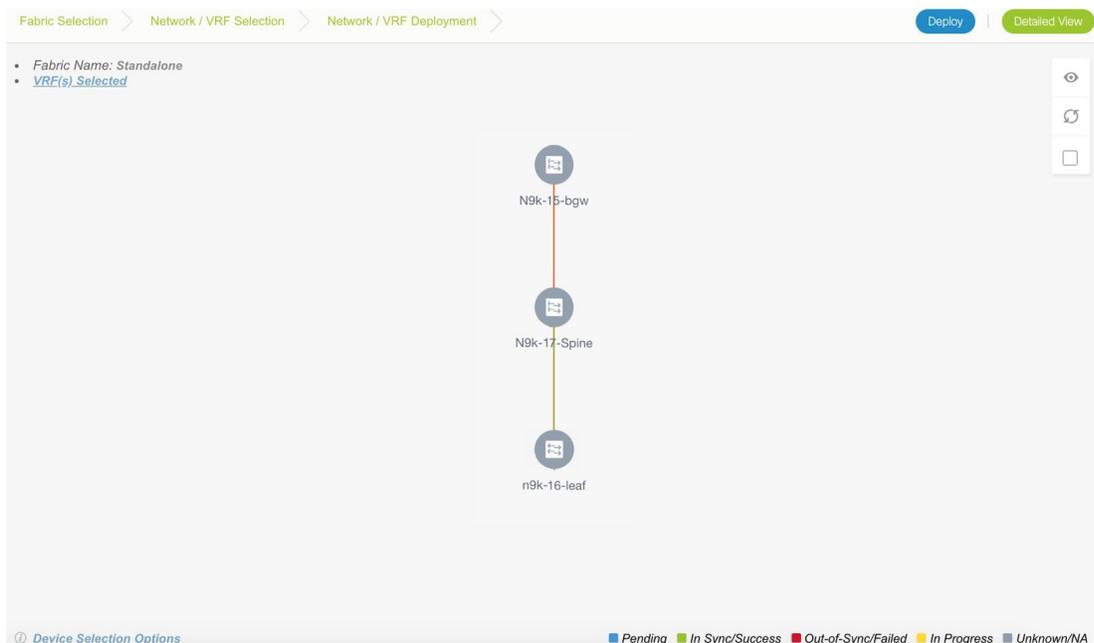
## Deploying VRFs for the Standalone and MSD Fabrics

1. Choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

| Check                               | VRF Name    | VRF ID | Status |
|-------------------------------------|-------------|--------|--------|
| <input checked="" type="checkbox"/> | MyVRF_50000 | 50000  | NA     |

2. Select check boxes next to the VRFs that you want to deploy and click Continue at the top right part of the screen.

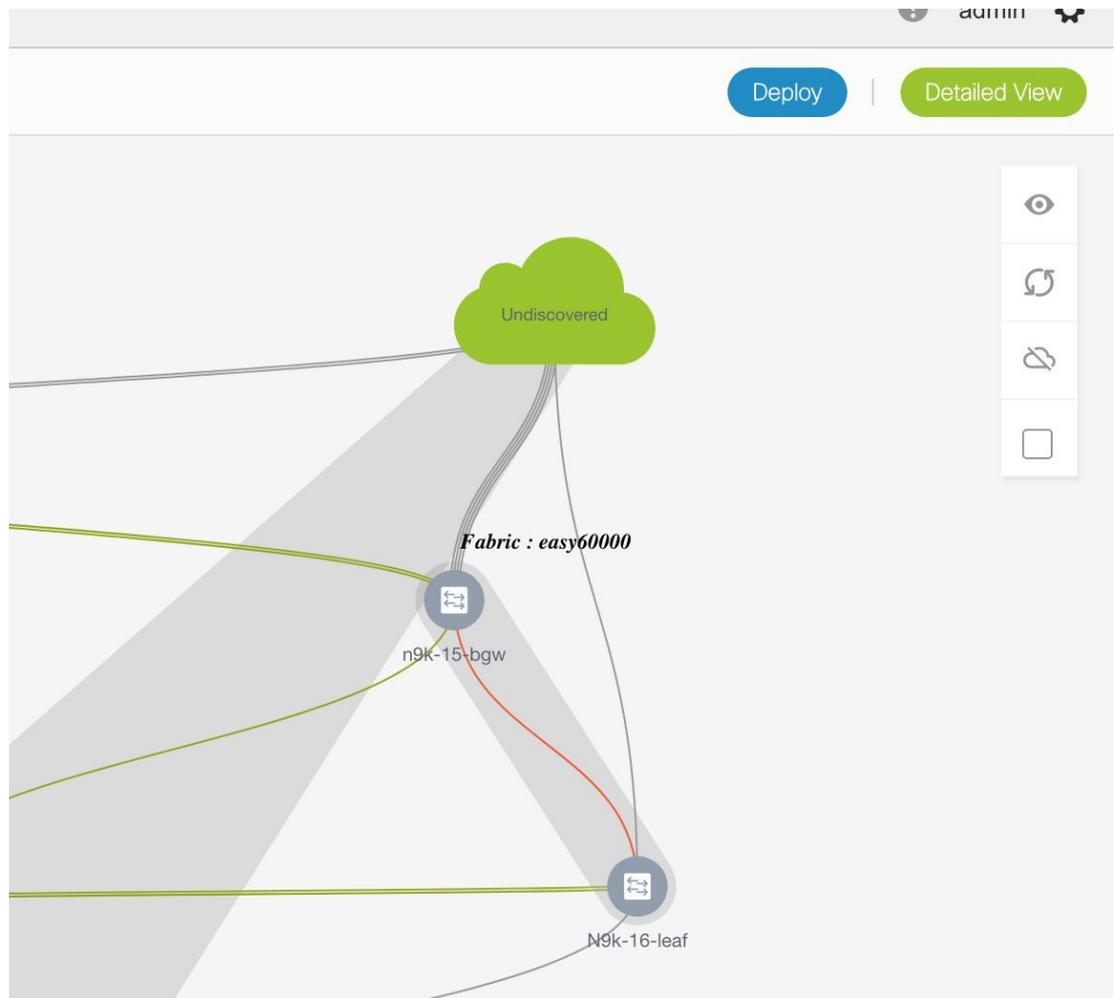
The VRF Deployment screen appears. On this page, you can see the topology of the Standalone fabric. The following example shows you how to deploy the VRFs MyVRF\_50000 and MyVRF\_50001 on the leaf switch. You can deploy VRFs simultaneously on multiple switches but of the same role (Leaf, Border Gateway, and so on).



At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly. Blue for *Pending* state, yellow for *In Progress* state when the provisioning is in progress, red for failure state, green when successfully deployed, and so on.

The overlay networks (or VRFs) provisioning status is context-specific. It is a combination of VRFs that you chose for provisioning and the relevant switches in the topology. In this example, it means that the VRFs are yet to be deployed on any switch in this fabric.

**Undiscovered cloud** display – To display (or not display) an **Undiscovered** cloud in this screen, click the cloud icon in the vertical panel, at the top-right part of the screen. When you click the icon, the **Undiscovered** cloud and its links to the fabric topology are not displayed. Click the icon again to display the **Undiscovered** cloud.



You can move the topology around the screen by clicking the left mouse button on the screen and moving it in the direction you desire. You can enlarge or shrink the switch icons proportionately by moving the cursor roller. You can also use corresponding alternatives on the touchpad.

3. Double-click a switch to deploy VRFs on it. The VRF Attachment screen comes up.



**Note** For deployment of VRFs on multiple switches, click the Multi-Select option from the panel at the top right part of the screen (This freezes the topology to a static state), and drag the cursor across the switches.

## VRF Attachment - Attach VRFs for given switch(es).



Fabric Name: Standalone

## Deployment Options

*Select the row and click on the cell to edit and save changes*

| MyVRF_50000              |             | MyVRF_50001 |      |                 |        |
|--------------------------|-------------|-------------|------|-----------------|--------|
| <input type="checkbox"/> | Switch      | ▲           | VLAN | CLI Freeform    | Status |
| <input type="checkbox"/> | n9k-16-leaf |             | 2000 | Freeform config | NA     |

Save

A tab represents each VRF that is being deployed (the first selected VRF is displayed by default). In each VRF tab, the selected switches are displayed. Each row represents a switch.

VLAN ID - Click within the VLAN column to update the VRF VLAN ID, if needed.

Freeform configurations – Click Freeform config to enable additional configurations on the switch. After you save freeform configurations, the Freeform config button gets highlighted.

Click the checkbox next to the Switch column to select all switches. VRF MyVRF\_50000 is ready to be provisioned on the switch

- 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:

n9k-16-leaf ▼

Select a VRF

MyVRF\_50000 ▼

Generated Configuration:

```
configure profile MyVRF_50000
vlan 2000
vn-segment 50000
interface vlan2000
vrf member myvrf_50000
ip forward
ipv6 forward
no ip redirects
no ipv6 redirects
mtu 9216
no shutdown
vrf context myvrf_50000
vni 50000
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
address-family ipv6 unicast
route-target both auto
route-target both auto evpn
router bgp 60000
vrf myvrf_50000
address-family ipv4 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redirect-subnet
maximum-paths ibgp 2
address-family ipv6 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redirect-subnet
maximum-paths ibgp 2
interface nve1
member vni 50000 associate-vrf
configure terminal
apply profile MyVRF_50000
```

After checking the configurations, close the screen. The *Topology View* screen appears.

- Click the **Deploy** button on the top right part of the screen. The color of the switch icons changes to yellow and a message appears at the bottom right part of the screen indicating that the deployment is in progress. After the VRF deployment is complete, the color of the switch icons changes to green, indicating successful deployment.

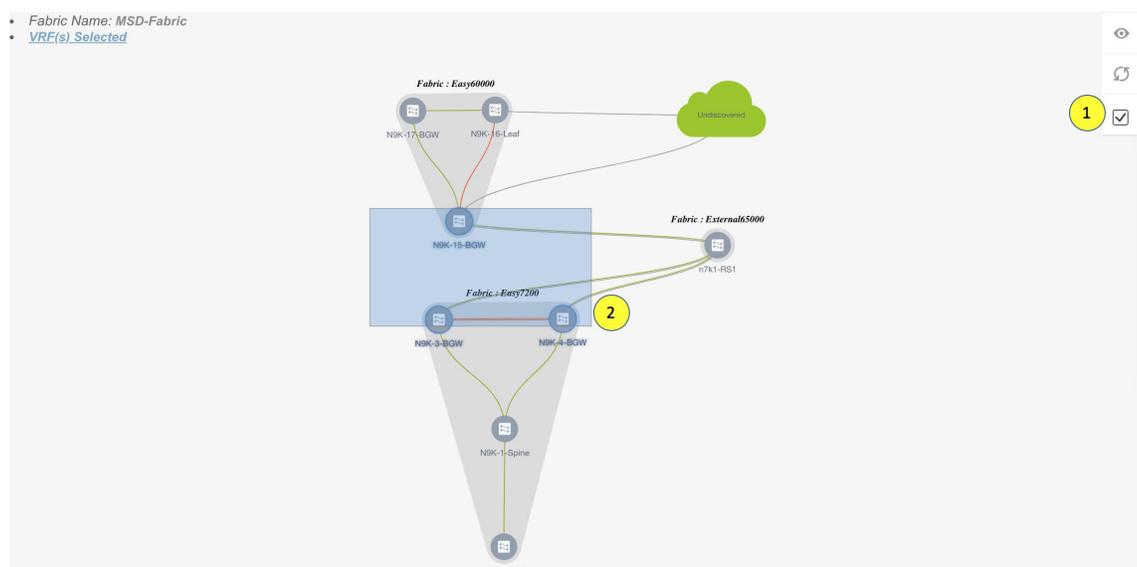


**Note** In case you click **Deploy** and there is no configuration diff that has to be deployed, a pop-up window comes up stating **No switches PENDING for deployment**.

### VRFs Deployment for an MSD Fabric

Consider a scenario wherein you are deploying the same VRFs on different member fabric border devices. You can choose one fabric, deploy VRFs on its border devices, and then choose the second fabric and deploy the VRFs.

Alternatively, you can choose the MSD fabric, and deploy the VRFs from a single topology view of all member fabric border devices at once.



### Detailed View

You can also use the **Detailed View** button to deploy networks and VRFs.

Click **Detailed View** at the top right part of the screen. The Detailed View screen comes up. This lists the VRFs in a tabular view.

Fabric Selection > Network / VRF Selection > Network / VRF Deployment > Topology View

Fabric Name: Standalone VRF(s) Selected Selected 0 / Total 4

| <input type="checkbox"/> | Name        | Switch      | Ports | Status   | Fabric Name | Role |
|--------------------------|-------------|-------------|-------|----------|-------------|------|
| <input type="checkbox"/> | MyVRF_50000 | n9k-15-BL   |       | NA       | Easy60000   | leaf |
| <input type="checkbox"/> | MyVRF_50000 | n9k-16-leaf |       | DEPLOYED | Easy60000   | leaf |
| <input type="checkbox"/> | MyVRF_50001 | n9k-15-BL   |       | NA       | Easy60000   | leaf |
| <input type="checkbox"/> | MyVRF_50001 | n9k-16-leaf |       | DEPLOYED | Easy60000   | leaf |

The options:

Edit - Select a VRF and click the Edit icon at the top left part of the screen.



#### Note

If you select one VRF/switch entry, the VRF Attach screen comes up. To maintain consistency across the Topology View and Detailed View screens, the VRF Attach screen displays all VRFs, and not just the selected VRF/switch entry.

Preview – Click Preview to preview configurations before deployment. You can only preview pending configurations, and not uninitiated or deployed configurations.

Deploy – Click Deploy to provision VRFs onto the switches.

History – Select a row and click History to view the configuration instances and status. Network and VRF-wise configurations are displayed. Click in the Status column of any instance for more details.

The fields in the table contain the configuration instance in each row, the associated switch and fabric names, the switch role, and the deployment status.

Apply/Save – Selecting a VRF and clicking Apply/Save will select a switch for the VRF to be deployed on.



**Note** When you upgrade from an earlier release (such as DCNM 10.4[2]) to the DCNM 11.0(1) release, overlay networks and VRFs deployment history information from the earlier DCNM release is not retained.

## Undeploying Networks for the Standalone Fabric

You can undeploy VRFs and networks from the deployment screen. The DCNM screen flow for undeployment is similar to the deployment process flow. Go to the deployment screen (Topology View) to undeploy networks:

1. Click **Control** > **Networks** (under **Fabrics** submenu).

The Networks screen comes up.

2. Choose the correct fabric from SCOPE. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

The screenshot shows the Cisco Data Center Network Manager interface. The top navigation bar includes the Cisco logo, 'Data Center Network Manager', and a 'SCOPE' dropdown menu set to 'bgp2'. Below the navigation bar, there are breadcrumb links for 'Network / VRF Selection' and 'Network / VRF Deployment', along with 'VRF View' and 'Continue' buttons. The main content area is titled 'Fabric Selected: bgp2' and displays a table of networks. The table has columns for Network Name, Network ID, VRF Name, IPv4 Gateway/Subnet, IPv6 Gateway/Prefix, Status, and VLAN ID. One network is listed: 'MyNetwork\_30000' with Network ID '30000', VRF Name 'NA', and Status 'NA'. The 'MyNetwork\_30000' row is selected, indicated by a checkmark in the first column.

| Network Name                                        | Network ID | VRF Name | IPv4 Gateway/Subnet | IPv6 Gateway/Prefix | Status | VLAN ID |
|-----------------------------------------------------|------------|----------|---------------------|---------------------|--------|---------|
| <input checked="" type="checkbox"/> MyNetwork_30000 | 30000      | NA       |                     |                     | NA     |         |

3. Select the networks that you want to undeploy and click Continue. The topology view comes up.
4. Select the Multi-Select button (if you are undeploying the networks from multiple switches), and drag the cursor across switches with the same role. The Network Attachment screen comes up.
  - (For a single switch, double-click the switch and the Network Attachment screen comes up).
  - (For a single switch, double-click the switch and the Switches Deploy screen comes up).
5. In the Network Attachment screen, the Status column for the deployed networks is displayed as DEPLOYED. Clear the check boxes next to the switches, as needed. Ensure that you repeat this on all tabs since each tab represents a network.
6. Click **Save** (at the bottom right part of the screen) to initiate the undeployment of the networks. The *Topology View* comes up again.



**Note** Alternatively, you can click the **Detailed View** button to undeploy networks.

7. Refresh the screen, preview configurations if needed and click **Deploy** to remove the network configurations on the switches. After the switch icons turn green, it indicates successful undeployment.
8. Go to the Networks page to verify if the networks are undeployed.

## Undeploying VRFs for the Standalone Fabric

You can undeploy VRFs from the deployment screen. The DCNM screen flow for undeployment is similar to the deployment process flow.

1. Choose **Control > Fabrics > VRFs**.
2. Choose the correct fabric from **SCOPE**. When you select a fabric, the **VRFs** screen refreshes and lists networks of the selected fabric.
3. Select the VRFs that you want to undeploy and click **Continue**. The *Topology View* page comes up.
4. Select the Multi-Select option (if you are undeploying the VRFs from multiple switches), and drag the cursor across switches with the same role. The VRF Attachment screen comes up.  
(For a single switch, double-click the switch and the VRF Attachment screen comes up).
5. In the Switches Deploy screen, the **Status** column for the deployed VRFs is displayed as DEPLOYED. Clear the check boxes next to the switches, as needed. Ensure that you repeat this on all tabs since each tab represents a VRF.
6. Click **Save** (at the bottom right part of the screen) to initiate the undeployment of the VRFs. The topology view comes up again.




---

**Note** Alternatively, you can click the **Detailed View** button to undeploy VRFs.

---

7. Refresh the screen, preview configurations if needed and click **Deploy** to remove the VRF configurations on the switches. After the switch icons turn green, it indicates successful undeployment.
8. Go to the VRFs page to verify if the networks are undeployed.

## Deleting Networks and VRFs

If you want to delete networks and corresponding VRFs in the MSD fabric, follow this order:

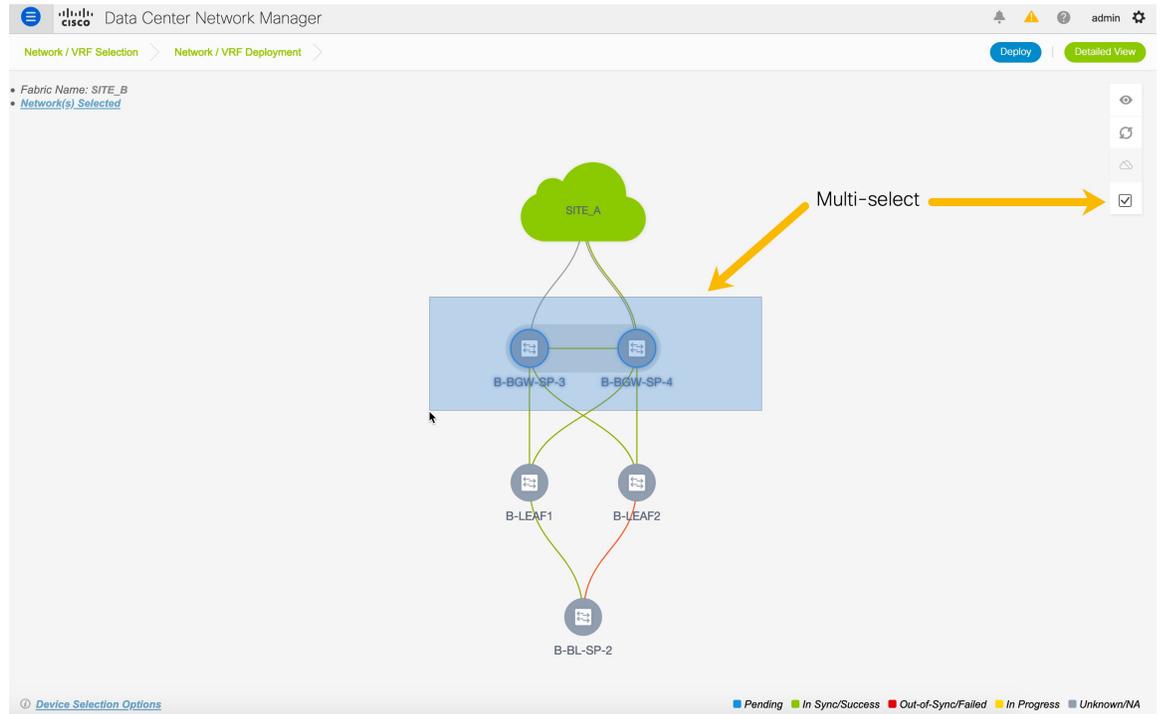
1. Undeploy the networks, if not already done.
2. Delete the networks.
3. Undeploy the VRFs, if not already done.
4. Delete the VRFs.

## Configuring Multiple VLAN IDs to a Single VNI

The following procedure shows how to tag multiple VLAN IDs to a single VNI in DCNM.

## Procedure

- Step 1** Navigate to **Control > Networks**.
- Step 2** Select the fabric from the **Scope** drop-down list and then select the network. Click **Continue**.
- Step 3** Check the **Multi-Select** check box and drag the cursor over the switches that needs to be updated with VLAN IDs.



- Step 4** In the **Network Attachment** window, edit the VLAN ID for the switches and click **Save**.

Network Extension Attachment - Attach extensions for given switch(es) ✕

Fabric Name: SITE\_B

Deployment Options

① Select the row and click on the cell to edit and save changes

MyNetwork\_30000 ← Network VNI

| Switch                              | VLAN | Extend    | Interfaces | CLI Freeform    | Status |
|-------------------------------------|------|-----------|------------|-----------------|--------|
| <input type="checkbox"/> B-BGW-SP-3 | 2300 | MULTISITE |            |                 | NA     |
| <input type="checkbox"/> B-BGW-SP-4 | 2300 | MULTISITE | ...        | Freeform config | NA     |

← Switches

**Save**

**Step 5** Click **Deploy** to deploy the configuration.

---

## Fabric Backup and Restore

This section describes the fabric backup and restore in Cisco DCNM.

### Backing Up Fabrics

You can back up all fabric configurations and intents automatically or manually. You can save configurations in DCNM, which are the intents. The intent may or may not be pushed on to the switches.

DCNM doesn't back up the following fabrics:

- External fabrics in monitor-only mode: Backing up of external fabrics in monitor-only mode isn't supported because you can't restore any configurations or intent.
- Parent MSD fabrics in releases earlier than Cisco DCNM, Release 11.4(1): You can only back up the configurations and intent of member fabrics in an MSD fabric individually.



---

**Note** From Cisco DCNM, Release 11.4(1), you can take backups of MSD fabrics. When you initiate a backup from the parent fabric, the backup process is applicable for the member fabrics as well. However, DCNM stores all the backed-up information of the member fabrics and the MSD fabric together in a single directory.

---

The backup does not capture the intent related to IFC. When you're backing up an external fabric, the checkpoints are copied from the switches to DCNM. The backup configuration files are stored in the following path in DCNM: `/usr/local/cisco/dcm/dcnm/data/archive`

By default, DCNM archives only 50 backups, and removes the older backups.

You can set the number of backup files to be archived in the **Server Properties** window. Search for the **Number of archived files per fabric to be retained:** section in the **Server Properties** window. Enter a value in the **archived.versions.limit** field.

### Backing Up Fabrics Automatically

You can enable an automatic hourly backup or scheduled backup for fabric configurations and intents. There are two types of automatic backup.

The backup has the information related to intent and fabric configurations in addition to associated state of the resource manager in terms of used resources on fabrics. DCNM backs up only when there's a configuration push. DCNM triggers the automatic backup only if you didn't trigger any manual backup after the last configuration push.

There are two types of automatic backup.

- **Hourly Fabric Backup:** You can enable an hourly backup.
- **Scheduled Fabric Backup:** You can schedule a fabric backup for regular intervals.



---

**Note** In external fabrics, DCNM backs up the changes in the running configurations as well. The configuration push happens after a deploy. If you didn't deploy the changes, you can't back up them in an intent.

---

Hourly and scheduled backup processes happen only during the next periodic configuration compliance activity, and there can be a delay of up to an hour.

### *Hourly and Scheduled Backup of Fabrics*

To enable automatic backup of fabric configurations and intents from the Cisco DCNM Web client, perform the following steps:

#### **Procedure**

---

- Step 1** Choose **Control > Fabrics > Fabric Builder**.
- The **Fabric Builder** window appears.
- Step 2** Click the **Edit Fabric** icon for the fabric you want to backup.
- Step 3** Click the **Configuration Backup** tab.
- Step 4** Choose the nature of backup by checking the appropriate check box.

The valid options are **Hourly Fabric Backup** and **Scheduled Fabric Backup**. If you want to enable both the backups, check the **Hourly Fabric Backup** check box and the **Scheduled Fabric Backup** check box.

**Note** If you check the **Scheduled Fabric Backup** check box, specify the scheduled backup time in the **Scheduled Time** field. Enter the value in `HH:MM` format.

- Step 5** Click **Save**.
- Step 6** Click the fabric and go to the fabric topology window.
- Step 7** Click **Save & Deploy**.
- 

### **Backing Up Fabrics Manually**

You can enable a manual backup for fabric configurations and intents. Regardless of the settings you choose under the **Configuration Backup** tab in the **Edit Fabric** dialog box, you can initiate a backup using this option.

To initiate a manual backup of fabric configurations and intents from the Cisco DCNM Web UI, perform the following steps:

#### **Procedure**

---

- Step 1** Choose **Control > Fabrics > Fabric Builder**.
- The **Fabric Builder** window appears.
- Step 2** Click the fabric for which you want to backup immediately.

The fabric topology window appears.

**Step 3** Click **Backup Now** in the **Actions** pane.

The **Backup Now** dialog appears.

**Step 4** Enter a tag name in the **Tag** field.

**Step 5** Click **OK**.

A confirmation message appears that the backup is triggered successfully.

**Note** The confirmation message only states that the backup is triggered and not if the backup is successful.

**Step 6** (Optional) Click **Restore Fabric** from the **Actions** pane to confirm if the manual backup is successful or not.

When you hover over the backup, the name has the tag you mentioned in *Step 4* confirming that it's a manual backup.

---

## Restoring Fabrics

This section describes the fabric restoring for different types of fabrics. Cisco DCNM supports configuration restore at fabric level. Take a backup of the configuration to restore it.

### Restoring Easy Fabrics

To restore an easy fabric in Cisco DCNM, perform the following steps from the Cisco DCNM Web UI:

#### Procedure

---

**Step 1** Choose **Control > Fabrics > Fabric Builder** and select a fabric.

**Step 2** Select **Restore Fabric** from the **Actions** menu.

The **Restore Fabric** window appears.

**Step 3** Choose the time for which you want to restore the configuration.

Valid values are **1m**, **3m**, **6m**, **YTD**, **1y**, and **All**. You can zoom into the graph. By default the backup information for **1m**, which is one month, appears. You can also select a custom date range. The backup information includes the following information:

- Backup date
- Total number of devices
- Number of devices in sync
- Number of devices out of sync

**Step 4** Click **View Backup Summary** to see the selected backup information of the devices in sync.

The switch name, switch serial number, IP address, status, and the configuration details of the devices appear.

**Note** If you add or remove devices from the fabric, the backup isn't valid. You can restore only the valid backups.

- Step 5** Click **Get Config** to preview the configuration details.
- Config Preview** window appears, which has two tabs.
- **Backup Config**: This tab displays the backup configuration for the selected device.
  - **Current Config**: This tab displays the current configuration for the selected device.

**Step 6** Go back to **View Backup Summary** window.

**Step 7** Click **Restore Intent** to proceed with the restoring.

The **Restore Status** window appears. You can view the status of the following:

- **Validating Backup**
- **Restoring fabric intent**
- **Restoring underlay intent**
- **Restoring interface intent**
- **Restoring overlay intent**

The valid values for the status of any action are **In Progress**, **Pending**, or **Failed**.

**Note** If the status of **Validating Backup** is **Failed**, other restoring actions won't be listed in this window.

**Step 8** Click **Next** after the intent is restored.

The **Configuration Preview** window appears. You can view the following details in this window:

- Switch name
- IP address
- Switch serial number
- Preview configuration
- Status
- Progress

**Step 9** Click **Deploy** to deploy the restored configuration.

The **Configuration Deployment Status** window appears. You can view the details of the switch name, IP address, status, status description, and the progress.

**Step 10** Click **Close** after the restoring process is complete.

---

## Restoring External Fabrics

When you restore an external fabric, the backed-up checkpoint is copied from DCNM to switches. To restore an external fabric in Cisco DCNM, perform the following steps from the Cisco DCNM Web UI:

## Procedure

---

- Step 1** Choose **Control > Fabrics > Fabric Builder** and select a fabric.
- Step 2** Select **Restore Fabric** from the Actions menu.  
The **Restore Fabric** window appears.
- Step 3** Select the time for which you want to restore the configuration.  
Valid values are **1m**, **3m**, **6m**, **YTD**, **1y**, and **All**. You can zoom into the graph. By default the backup information for **1m**, which is one month, appears.

When you select a backup version, the vertical bar representing it turns grey, and corresponding information is displayed at the bottom part of the screen. It includes the following information:

- Backup date
- DCNM Version
- Total number of devices
- Number of devices in sync
- Number of devices out of sync

You can select a custom date range either by rearranging the date slide below the vertical bars, or using the **From** and **To** boxes at the top right part of the screen.

**Step 4** Click **View Backup Summary** to see the selected backup information of the devices in sync.

The switch name, switch serial number, IP address, status, Restore Supported (indicating whether the device supports checkpoint rollback or not), the configuration details of the devices, and the VRF appear.

**Note** For information about the support for the checkpoint rollback feature in platforms, refer to the respective platform documentation.

By default, the management VRF is displayed in the VRF column because it is used for the copy operation during the restore process. If you want to use a different VRF for the copy operation, update the VRF column. To update the same VRF for all devices, use the Apply for all devices option at the bottom-left part of the screen. A sample screenshot:

**Note** If you added or removed devices to the fabric, you can't restore a fabric from the present day to a past date.

**Step 5** Click **Get Config** to preview device configuration details.

The **Config Preview** window appears, which has three tabs.

- **Backup Config:** This tab displays the backup configuration for the selected device.
- **Current Config:** This tab displays the current running configuration of the selected device.
- **Side-by-side Comparison:** This tab displays current running configuration on the switch, and the backup configuration (or expected configuration).

**Step 6** Go back to the **View Backup Summary** window.

**Step 7** Click **Restore Intent** to proceed with the restoring.

The **Restore Status** window appears. You can view the status of the following:

- **Validating Backup**
- **Restoring fabric intent**
- **Restoring underlay intent**
- **Restoring interface intent**
- **Restoring overlay intent**
- **Intent Regeneration**

The valid values for the status of any action are **In Progress**, **Pending**, **Completed**, or **Failed**.

**Note** If the status of **Validating Backup** is **Failed**, other restoring actions won't be listed in this window.

**Step 8** Click **Close** after the restore process is complete.

---

## Deleting a VXLAN BGP EVPN Fabric

Choose **Control > Fabric Builder**. On the Fabric Builder page, click **X** on the rectangular box that represents the fabric. Ensure the following before deleting a fabric.

- Fabric devices should not be in transition such as migration into or out of the fabric, ongoing network or VRF provisioning, and so on. Delete a fabric after the transition is complete.
- Remove devices that are still attached to the fabric. Remove non-Cisco Nexus 9000 Series switches first and then remove the 9000 Series switches.

## Post DCNM 11.2(1) Upgrade for VXLAN BGP EVPN, External, and MSD Fabrics

Note the following guidelines after you upgrade to the DCNM Release 11.2(1):

- As part of the upgrade from an earlier DCNM release, the fabric and associated templates are carried over to the DCNM Release 11.2(1).

- You can use the old fabric template, but you will not be able to use the features introduced in the DCNM Release 11.2(1). Perform the following steps to use the new features:

- Edit the settings of each fabric by updating the old fabric template with the equivalent new fabric template and clicking **Save**.

The following table shows the old and new fabric template names in DCNM.

| Old Template Fabric Name | New Template Fabric Name |
|--------------------------|--------------------------|
| Easy_Fabric              | Easy_Fabric_11_1         |
| External_Fabric          | External_Fabric_11_1     |
| MSD_Fabric               | MSD_Fabric_11_1          |

- Under the **Advanced** tab, the **iBGP Peer-Template Config** field displays the value **null**. If this value is displayed, remove it and click **Save**.



**Note** You can skip this step if you are upgrading from the DCNM Release 11.1(1) to the DCNM Release 11.2(1).

- Navigate to each fabric in the Topology view, and click **Save & Deploy** to deploy any changes.

If you encounter any new or unexpected pending configurations after you click **Save & Deploy**, refer [Configuration Compliance in DCNM, on page 214](#).



**Caution** Some configuration changes can be expected as part of this step. Therefore, perform it only during a scheduled maintenance window.

- After a multi-level upgrade from Cisco DCNM 10.4(2) or 11.0(1), you can change the VRF templates to **Default\_VRF\_Universal** or **Default\_VRF\_Extension\_Universal** to enable **ipv6 address use-link-local-only**.

## Changing ISIS Configuration from Level 1 to Level 2

This procedure shows how to change ISIS configuration on switches from Level 1 to Level 2 in a VXLAN fabric deployment.

- Choose **Control > Fabrics > Fabric Builder**.
- Click a fabric in the **Fabric Builder** window.
- Click **Tabular view** under **Actions** menu.
- Search for all the **base\_isis** policies in the **Template** search field.
- Select all the **base\_isis** policies and click the **Delete** icon to delete policies
- Click **Save & Deploy**.

After all the **base\_isis** policies are deleted, DCNM considers the migrated brownfield fabric as a greenfield fabric and creates the **base\_isis\_level2** policies on the switches.

## Configuration Compliance in DCNM

The entire intent or expected configuration defined for a given switch is stored in DCNM. When you want to push this configuration down to one or more switches, the configuration compliance (CC) module is triggered. CC takes the current intent, the current running configuration, and then comes up with the set of configurations that are required to go from the current running configuration to the current expected config so that everything will be In-Sync.

When performing a software or firmware upgrade on the switches, the current running configuration on the switches is not changed. Post upgrade, if CC finds that the current running configuration does not have the current expected configuration or intent, it reports an Out-of-Sync status. There is no auto deployment of any configurations. You can preview the diffs that will get deployed to get one or more devices back In-Sync.

With CC, the sync is always from the DCNM to the switches. There is no reverse sync. So, if you make a change out-of-band on the switches that conflicts with the defined intent in DCNM, CC captures this diff, and indicates that the device is Out-of-Sync. The pending diffs will undo the configs done out-of-band to bring back the device In-Sync. Note that such conflicts due to out-of-band changes are captured by the periodic CC run that occurs every 60 mins by default, or when you click the RESYNC option either on a per fabric or per switch basis. Note that you can also capture the out-of-band changes for the entire switch by using the CC REST API. For more information, see *Cisco DCNM REST API Guide, Release 11.2(1)*.

From Cisco DCNM Release 11.2(1), to improve ease of use and readability of deployed configurations, CC in DCNM has been enhanced with the following:

- All displayed configurations in DCNM are easily readable and understandable.
- Repeated configuration snippets are not displayed.
- Pending configurations precisely show only the diff configuration.
- Side-by-side diffs has greater readability, integrated search or copy, and diff summary functions.

All freeform configurations have to strictly match the **show running configuration** output on the switch and any deviations from the configuration will show up as a diff during **Save & Deploy**. You need to adhere to the leading space indentations.

You can typically enter configuration snippets in DCNM using the following methods:

- User-defined profile and templates
- Switch, interface, overlay, and vPC freeform configurations
- Network and VRF per switch freeform configurations
- Fabric settings for Leaf, Spine, or iBGP configurations



---

**Caution**

The configuration format should be identical to the **show running configuration** of the corresponding switch. Otherwise, any missing or incorrect leading spaces in the configuration can cause unexpected deployment errors and unpredictable pending configurations. If any unexpected diffs or deployment errors are displayed, check the user-provided or custom configuration snippets for incorrect values.

---

If DCNM displays the "Out-of-Sync" status due to unexpected pending configurations, and this configuration is either unable to be deployed or stays consistent even after a deployment, perform the following steps to recover:

1. Check the lines of config highlighted under the **Pending Config** tab in the **Config Preview** window.
2. Check the same lines in the corresponding **Side-by-side Comparison** tab. This tab shows whether the diff exists in "intent", or "show run", or in both with different leading spaces. Leading spaces are highlighted in the **Side-by-side Comparison** tab.
3. If the pending configurations or switch with an out-of-sync status is due to any identifiable configuration with mismatched leading spaces in "intent" and "running configuration", this indicates that the intent has incorrect spacing and needs to be edited.
4. To edit incorrect spacing on any custom or user-defined policies, navigate to the switch and edit the corresponding policy:
  - a. If the source of the policy is **UNDERLAY**, you will need to edit this from the Fabric settings screen and save the updated configuration.
  - b. If the source is blank, it can be edited from the **View/Edit policies** window for that switch.
  - c. If the source of the policy is **OVERLAY**, but it is derived from a switch freeform configuration. In this case, navigate to the appropriate **OVERLAY** switch freeform configuration and update it.
  - d. If the source of the policy is **OVERLAY** or a custom template, perform the following steps:
    1. Navigate to **Administration > DCNM Server > Server Properties**, set the **template.in\_use.check** property to **false**. This allows the profiles or templates to be editable.
    2. Edit the specific profile or template from the **Control > Template Library** edit window, and save the updated profile template with the right spacing.
    3. Click **Save & Deploy** to recompute the diffs for the impacted switches.
    4. After the configurations are updated, set the **template.in\_use.check** property to **true**, as it slows down the performance of the DCNM system, specifically for **Save & Deploy** operations.

To confirm that the diffs have been resolved, click **Save & Deploy** after updating the policy to validate the changes.



---

**Note** DCNM checks only leading spaces, as it implies hierarchy of the command, especially in case of multi-command sequences. DCNM does not check any trailing spaces in command sequences.

---

### Example 1: Configuration Compliance in Switch Freeform Policy

Let us consider an example with an incorrect spacing in the Switch Freeform Config field.

The switch freeform policy is created as shown:

Edit Policy
✕

Policy ID: POLICY-30630

Entity Type: SWITCH

\* Priority (1-1000):

Template Name: switch\_freeform

Entity Name: SWITCH

General

Variables:

\* Switch Freeform Config

```

ip dhcp relay
ip dhcp relay information option
ip dhcp relay information option vpn
ip dhcp snooping
ip domain-lookup
ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25
ip pim ssm range 232.0.0.0/8
ipv6 dhcp relay
ipv6 switch-packets lla
          
```

After deploying this policy successfully to the switch, DCNM persistently reports the following diffs:

### Config Preview - Switch 70.70.70.73

Pending Config

Side-by-side Comparison

```

ip domain-lookup
 ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25
 ip pim ssm range 232.0.0.0/8
 ipv6 dhcp relay
 ipv6 switch-packets lla
 configure terminal
          
```

After clicking the **Side-by-side Comparison** tab, you can see the cause of the diff. As seen below, the **ip pim rp-address** line has 2 leading spaces, while the running configuration has 0 leading spaces.

## Config Preview - Switch 70.70.70.73

Pending Config | Side-by-side Comparison

To re-compute the *running config*, please click the Re-sync button on the previous screen. Lastly, to resolve unexpected diffs, please review the leading spaces and edit the appropriate policies to match show run outputs.

| Running config                                               | Expected config                                        |
|--------------------------------------------------------------|--------------------------------------------------------|
| 281 description "vpc-peer-link"                              | description "vpc-peer-link"                            |
| 282                                                          | no shutdown                                            |
| 283 spanning-tree port type network                          | spanning-tree port type network                        |
| 284 switchport                                               | switchport                                             |
| 285 switchport mode trunk                                    | switchport mode trunk                                  |
| 286 vpc peer-link                                            | vpc peer-link                                          |
| 287 ip dhcp relay                                            | ip dhcp relay                                          |
| 288 ip dhcp relay information option                         | ip dhcp relay information option                       |
| 289 ip dhcp relay information option vpn                     | ip dhcp relay information option vpn                   |
| 290 ip dhcp snooping                                         | ip dhcp snooping                                       |
| 291 ip domain-lookup                                         | ip domain-lookup                                       |
| 292                                                          | ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25 |
| 293                                                          | ip pim ssm range 232.0.0.0/8                           |
| 294                                                          | ipv6 dhcp relay                                        |
| 295                                                          | ipv6 switch-packets lla                                |
| 296 ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25   | ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25 |
| 297 ip pim ssm range 232.0.0.0/8                             | ip pim ssm range 232.0.0.0/8                           |
| 298 ipv6 dhcp relay                                          | ipv6 dhcp relay                                        |
| 299 ipv6 switch-packets lla                                  | ipv6 switch-packets lla                                |
| 300 line console                                             | line console                                           |
| 301 line vty                                                 | line vty                                               |
| 302 ngoam install acl                                        | ngoam install acl                                      |
| 303 nv overlay evpn                                          | nv overlay evpn                                        |
| 304 nxapi http port 80                                       | nxapi http port 80                                     |
| 305 rmon event 1 description FATAL(1) owner PMON@FATAL       |                                                        |
| 306                                                          | power redundancy-mode ps-redundant                     |
| 307 rmon event 2 description CRITICAL(2) owner PMON@CRITICAL |                                                        |
| 308 rmon event 3 description ERROR(3) owner PMON@ERROR       |                                                        |
| 309 rmon event 4 description WARNING(4) owner PMON@WARNING   |                                                        |

To resolve this diff, edit the corresponding Switch Freeform policy so that the spacing is correct.

## Edit Policy

Policy ID: POLICY-30630 | Template Name: switch\_freeform  
 Entity Type: SWITCH | Entity Name: SWITCH

\* Priority (1-1000):

General

Variables:

\* Switch Freeform Config

```
ip dhcp relay
ip dhcp relay information option
ip dhcp relay information option vpn
ip dhcp snooping
ip domain-lookup
ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25
ip pim ssm range 232.0.0.0/8
ipv6 dhcp relay
ipv6 switch-packets lla
```

Save Push Config Cancel

After you save, you can use the **Push Config** or **Save & Deploy** option to re-compute diffs.

As shown below, the diffs are now resolved. The **Side-by-side Comparison** tab confirms that the leading spaces are updated.

Config Preview - Switch 70.70.70.73

Pending Config Side-by-side Comparison

Running config

```

276 interface nve1
277 host-reachability proto
278 no shutdown
279 source-interface loopba
280 interface port-channel500
281 description "vpc-peer-L
282
283 spanning-tree port type
284 switchport
285 switchport mode trunk
286 vpc peer-link
287 ip dhcp relay
288 ip dhcp relay information
289 ip dhcp relay information
290 ip dhcp snooping
291 ip domain-lookup
292 ip pim rp-address 10.254.
293 ip pim ssm range 232.0.0.
294 ipv6 dhcp relay
295 ipv6 switch-packets lla
296 line console
297 line vty
298 ngoam install acl
299 nv overlay evpn
300 nxapi http port 80

```

## Example 2: Resolving a Leading Space Error in Overlay Configurations

Let us consider an example with a leading space error that is displayed in the **Pending Config** tab.

Config Preview - Switch 80.80.80.62

Pending Config Side-by-side Comparison

```

terminal dont-expunge
router bgp 65000
  vrf common-dmz
    redistribute static route-map allow
    default-information originate
configure terminal
terminal dont-expunge
router bgp 65000
  vrf common-dmz
    address-family ipv4 unicast
      no default-information originate
      no redistribute static route-map allow
configure terminal

```

Unexpected Pending configurations after upgrade or after configuration updates.

In the **Side-by-side Comparison** tab, search for diffs line by line to understand context of the deployed configuration.

terminal dont-expunge 0/0

SCOPE: green

Search for the Diffs, line by line in the Side-by-Side to understand context of the deployed configuration.

Matched count of 0, means this is some special configuration that DCNM has evaluated it needs to be pushed to the switch.

Config Preview - Switch 80.80.80.62

Pending Config Side-by-side Comparison

To re-compute the *running config*, please click the Re-sync button on the previous screen. Lastly, to resolve unexpected diffs, please review the leading spaces and edit the appropriate policies to match show run outputs.

| Running config                                                 | Expected config                                                                                   |
|----------------------------------------------------------------|---------------------------------------------------------------------------------------------------|
| 1 !Command: show running-config                                |                                                                                                   |
| 2                                                              | !Command: Intent from DCNM Fabric Builder. Any Intent not captured in Pending Config are defaults |
| 3 !Running configuration last done at: Tue Jun 4 14:19:01 2019 |                                                                                                   |
| 4                                                              | aaa group server radius radius                                                                    |
| 5 !Time: Tue Jun 4 16:03:38 2019                               |                                                                                                   |
| 6                                                              | use-vrf default                                                                                   |
| 7 aaa group server tacacs+ ACS                                 | aaa group server tacacs+ ACS                                                                      |
| 8 server 10.145.249.150                                        | server 10.145.249.150                                                                             |
| 9 server 10.2.98.28                                            | server 10.2.98.28                                                                                 |
| 10 server 10.20.0.201                                          | server 10.20.0.201                                                                                |
| 11 source-interface mgmt0                                      | source-interface mgmt0                                                                            |
| 12 use-vrf management                                          | use-vrf management                                                                                |
| 13 boot nxos bootflash:/nxos.9.2.3.bin                         |                                                                                                   |
| 14 cfs eth distribute                                          | cfs eth distribute                                                                                |
| 15 configure profile Auto_Net_VNI20006_VLAN6                   | configure profile Auto_Net_VNI20006_VLAN6                                                         |
| 16 evpn                                                        | evpn                                                                                              |

A matched count of 0 means that it is a special configuration that DCNM has evaluated to push it to the switch.

redistribute static route-map 1/14

SCOPE: green

Searching for the next line in the pending configuration, shows the problem. The leading spaces are mismatched between running and expected configurations. 6 leading spaces in "Running configurations" and 4 leading spaces in "Expected Configuration". Similar mismatch is seen for "default-information originate" as well. For the VRF common-dmz as shown below.

Config Preview - Switch 80.80.80.62

Pending Config Side-by-side Comparison

To re-compute the *running config*, please click the Re-sync button on the previous screen. Lastly, to resolve unexpected diffs, please review the leading spaces and edit the appropriate policies to match show run outputs.

| Running config                           | Expected config                     |
|------------------------------------------|-------------------------------------|
| 2604 bfd                                 | bfd                                 |
| 2605 remote-as 65000                     | remote-as 65000                     |
| 2606 update-source loopback501           | update-source loopback501           |
| 2607 router-id 192.168.0.4               | router-id 192.168.0.4               |
| 2608 vrf common-dmz                      | vrf common-dmz                      |
| 2609 address-family ipv4 unicast         | address-family ipv4 unicast         |
| 2610 default-information originate       | default-information originate       |
| 2611                                     |                                     |
| 2612 redistribute static route-map allow | redistribute static route-map allow |
| 2613                                     |                                     |
| 2614 vrf common-mgmt                     | vrf common-mgmt                     |
| 2615 address-family ipv4 unicast         | address-family ipv4 unicast         |
| 2616 default-information originate       | default-information originate       |
| 2617 redistribute static route-map allow | redistribute static route-map allow |
| 2618 vrf ecd                             | vrf ecd                             |
| 2619 address-family ipv4 unicast         | address-family ipv4 unicast         |
| 2620 default-information originate       | default-information originate       |
| 2621 redistribute static route-map allow | redistribute static route-map allow |
| 2622 vrf ialab                           | vrf ialab                           |
| 2623 address-family ipv4 unicast         | address-family ipv4 unicast         |
| 2624 default-information originate       | default-information originate       |
| 2625 redistribute static route-map allow | redistribute static route-map allow |
| 2626 vrf lc                              | vrf lc                              |

You can see that the leading spaces are mismatched between running and expected configurations.

Navigate to the respective freeform configs and correct the leading spaces, and save the updated configuration.

The screenshot shows the Cisco Data Center Network Manager (DCNM) interface. The main window displays the 'Network / VRF Deployment' section for a fabric named 'green'. A 'VRF Attachment - Attach VRFs for given switch(es)' dialog is open, showing a list of VRFs with 'COMMON-DMZ' selected. A 'Freeform Config (n9k12\_bp2-lfsw01-l001)' window is also open, displaying the following configuration:

```

vrf context COMMON-DMZ
ip route 0.0.0.0/0 10.9.8.17 name COMMON-DMZ-DG
ip route 10.0.0.0/8 10.9.8.17 name OT-Networks-1-via-Mgmt&Tools-VDOM
ip route 10.9.16.0/20 10.9.8.17 name Mgmt&Tools-Networks
ip route 10.9.32.0/19 10.9.8.25 name RS-VDOM
ip route 10.9.128.0/19 10.9.8.33 name ECD-DG
ip route 10.9.254.0/23 10.9.8.9 name to-RA-LAB-VDOY
ip route 149.235.128.0/16 10.9.8.17 name OT-Networks-4-via-Mgmt&Tools-VDOM
ip route 172.16.0.0/12 10.9.8.17 name OT-Networks-5-2-via-Mgmt&Tools-VDOM
ip route 192.168.0.0/16 10.9.8.17 name OT-Networks-3-via-Mgmt&Tools-VDOM
router bgp 65000
vrf COMMON-DMZ
address-family ipv4 unicast
 redistribute static route-map allow
 default-information originate
  
```

A callout box with an arrow points to the Freeform Config field, containing the text: "Navigate to the corresponding VRF, Common-VRF and edit the freeform attachment to find the incorrectly spaced lines".

Navigate to the **Fabric Builder** window for the fabric and click **Save & Deploy**.

In the **Config Deployment** window, you can see that all the devices are in-sync.

Config Deployment

Step 1. Configuration Preview > Step 2. Configuration Deployment Status >

| Switch Name      | IP Address  | Switch Serial | Preview Config | Status  | Re-sync | Progress |
|------------------|-------------|---------------|----------------|---------|---------|----------|
| n9k12_bp2-lfs... | 80.80.80.62 | SAL18422FX8   | 0 lines        | In-Sync |         | 100%     |
| n9k13_bp2-lfs... | 80.80.80.63 | SAL18422FXE   | 0 lines        | In-Sync |         | 100%     |
| n9k7_bp2-lfsw... | 80.80.80.57 | SAL1833YM64   | 0 lines        | In-Sync |         | 100%     |
| n9k14_bp2-sp...  | 80.80.80.64 | SAL2016NXXB   | 0 lines        | In-sync |         | 100%     |
| n9k8_bp2-sps...  | 80.80.80.58 | SAL1833YM0V   | 0 lines        | In-sync |         | 100%     |

Deploy Config

## Configuration Compliance in External Fabrics

With external fabrics, any Nexus switch can be imported into the fabric, and there is no restriction on the type of deployment. It can be LAN Classic, VXLAN, FabricPath, vPC, HSRP, etc. When switches are imported into an external fabric, the configuration on the switches is retained so that it is non-disruptive. Only basic policies such as the switch username and mgmt0 interface are created after a switch import.

In the external fabric, for any intent that is defined in the DCNM, configuration compliance (CC) ensures that this intent is present on the corresponding switch. If this intent is not present on the switch, CC reports an Out-of-Sync status. Additionally, there will be a Pending Config generated to push this intent to the switch to change the status to In-Sync. Any additional configuration that is on the switch but not in intent defined in DCNM, will be ignored by CC, as long as there is no conflict with anything in the intent.

When there is user-defined intent added on DCNM and the switch has additional configuration under the same top-level command, as mentioned earlier, CC will only ensure that the intent defined in DCNM is present on the switch. When this user defined intent on DCNM is deleted as a whole with the intention of removing it from the switch and the corresponding configuration exists on the switch, CC will report an Out-of-Sync status for the switch and will generate **Pending Config** to remove the config from the switch. This **Pending Config** includes the removal of the top-level command. This action leads to removal of the other out-of-band configurations made on the switch under this top-level command as well. If you choose to override this behavior, the recommendation is that, you create a freeform policy and add the relevant top-level command to the freeform policy.

Let us see this behavior with an example.

1. A **switch\_freeform** policy defined by the user in DCNM and deployed to the switch.

### Edit Policy ✕

**Policy ID:** POLICY-51710      **Template Name:** switch\_freedom  
**Entity Type:** SWITCH      **Entity Name:** SWITCH

\* **Priority (1-1000):**

General

\* **Switch Freeform Config**

```
router bgp 1234
neighbor 10.2.0.1
address-family l2vpn evpn
send-community both
remote-as 1234
update-source loopback0
```

**Variables:**

- Additional configuration exists under **router bgp** in **Running config** that does not exist in user-defined DCNM intent **Expected config**. Note that there is no **Pending Config** to remove the additional config that exists on the switch without a user defined intent on DCNM.

### Config Preview - Switch 172.29.21.130 ✕

Pending Config    Side-by-side Comparison

To re-compute the *running config*, please click the Re-sync button on the previous screen. Lastly, to resolve unexpected diffs, please review the leading spaces and edit the appropriate policies to match show `run outputs`.

| Running config                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Expected config                                                                                                                                                                          |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>593 rmon event 3 description ERROR(3) owner PMON@ERROR 594 rmon event 4 description WARNING(4) owner PMON@WARNING 595 rmon event 5 description INFORMATION(5) owner PMON@INFO 596 route-map fabric-rmap-redirect-subnet permit 10 597 match tag 12345 598 router bgp 1234 599 neighbor 10.2.0.1 600 address-family l2vpn evpn 601 send-community both 602 remote-as 1234 603 update-source loopback0 604 neighbor 20.2.0.2 605 address-family ipv4 unicast 606 send-community both 607 router-id 10.2.0.2 608 router ospf UNDERLAY 609 router-id 10.2.0.2 610 service dhcp 611 snmp-server host 172.28.194.124 traps version 2c public udp-port 2162 612 snmp-server host 172.28.194.126 traps version 2c public udp-port 2162 613 snmp-server host 172.28.194.130 traps version 2c public udp-port 2162 614 tacacs-server host 1.1.1.11 key 7 "cisco123" 615 vdc N9K-z1 id 1 616 limit-resource m4route-mem minimum 58 maximum 58 617 limit-resource m6route-mem minimum 8 maximum 8 618 limit-resource port-channel minimum 0 maximum 511 619 limit-resource u4route-mem minimum 248 maximum 248 620 limit-resource u6route-mem minimum 96 maximum 96 621 limit-resource vlan minimum 16 maximum 4094 622 limit-resource vrf minimum 2 maximum 4096 623 version 7.0(3)I7(3) 624 vlan 1 625 vrf context management 626 ip route 0.0.0.0/0 172.29.21.1</pre> | <pre>router bgp 1234 neighbor 10.2.0.1 address-family l2vpn evpn send-community both remote-as 1234 update-source loopback0  vrf context management ip route 0.0.0.0/0 172.29.21.1</pre> |

## Config Preview - Switch 172.29.21.130



Pending Config

Side-by-side Comparison

3. The **Pending Config** and the **Side-by-side Comparison** when the intent that was pushed earlier via DCNM is deleted from DCNM by deleting the switch\_freeform policy that was created in the Step 1.

## Config Preview - Switch 172.29.21.130



Pending Config

Side-by-side Comparison

To re-compute the *running config*, please click the Re-sync button on the previous screen. Lastly, to resolve unexpected diffs, please review the leading spaces and edit the appropriate policies to match `show run` outputs.

| Running config                                                            | Expected config            |
|---------------------------------------------------------------------------|----------------------------|
| 584 ip domain-lookup                                                      |                            |
| 585 ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25                |                            |
| 586 ip pim ssm range 232.0.0.0/8                                          |                            |
| 587 ipv6 dhcp relay                                                       |                            |
| 588 ipv6 switch-packets lla                                               |                            |
| 589 line console                                                          |                            |
| 590 line vty                                                              |                            |
| 591 ngoam install acl                                                     |                            |
| 592 no password strength-check                                            | no password strength-check |
| 593 nv overlay evpn                                                       |                            |
| 594 rmon event 1 description FATAL(1) owner PMON@FATAL                    |                            |
| 595 rmon event 2 description CRITICAL(2) owner PMON@CRITICAL              |                            |
| 596 rmon event 3 description ERROR(3) owner PMON@ERROR                    |                            |
| 597 rmon event 4 description WARNING(4) owner PMON@WARNING                |                            |
| 598 rmon event 5 description INFORMATION(5) owner PMON@INFO               |                            |
| 599 route-map fabric-rmap-redis-subnet permit 10                          |                            |
| 600 match tag 12345                                                       |                            |
| 601 <del>router tag</del> 1234                                            |                            |
| 602 neighbor 10.2.0.1                                                     |                            |
| 603 address-family l2vpn evpn                                             |                            |
| 604 send-community both                                                   |                            |
| 605 remote-as 1234                                                        |                            |
| 606 update-source loopback0                                               |                            |
| 607 neighbor 20.2.0.2                                                     |                            |
| 608 address-family ipv4 unicast                                           |                            |
| 609 send-community both                                                   |                            |
| 610 router-id 10.2.0.2                                                    |                            |
| 611 router ospf UNDERLAY                                                  |                            |
| 612 router-id 10.2.0.2                                                    |                            |
| 613 service dhcp                                                          |                            |
| 614 snmp-server host 172.28.194.124 traps version 2c public udp-port 2162 |                            |
| 615 snmp-server host 172.28.194.126 traps version 2c public udp-port 2162 |                            |
| 616 snmp-server host 172.28.194.130 traps version 2c public udp-port 2162 |                            |
| 617 tacacs-server host 1.1.1.11 key 7 "cisco123"                          |                            |
| 618 tacacs-server host 172.28.1.203 key 7 "Fewhg12345"                    |                            |

## Config Preview - Switch 172.29.21.130

Pending Config    Side-by-side Comparison

```
no router bgp 1234
configure terminal
```

- A **switch\_freeform** policy with the top-level **router bgp** command needs to be created. This enables CC to generate the configuration needed to remove only the desired sub-config which was pushed from DCNM earlier.

### Edit Policy ✕

Policy ID: POLICY-51770      Template Name: switch\_freeform  
 Entity Type: SWITCH      Entity Name: SWITCH

\* Priority (1-1000):

General

---

Variables:

\* Switch Freeform Config

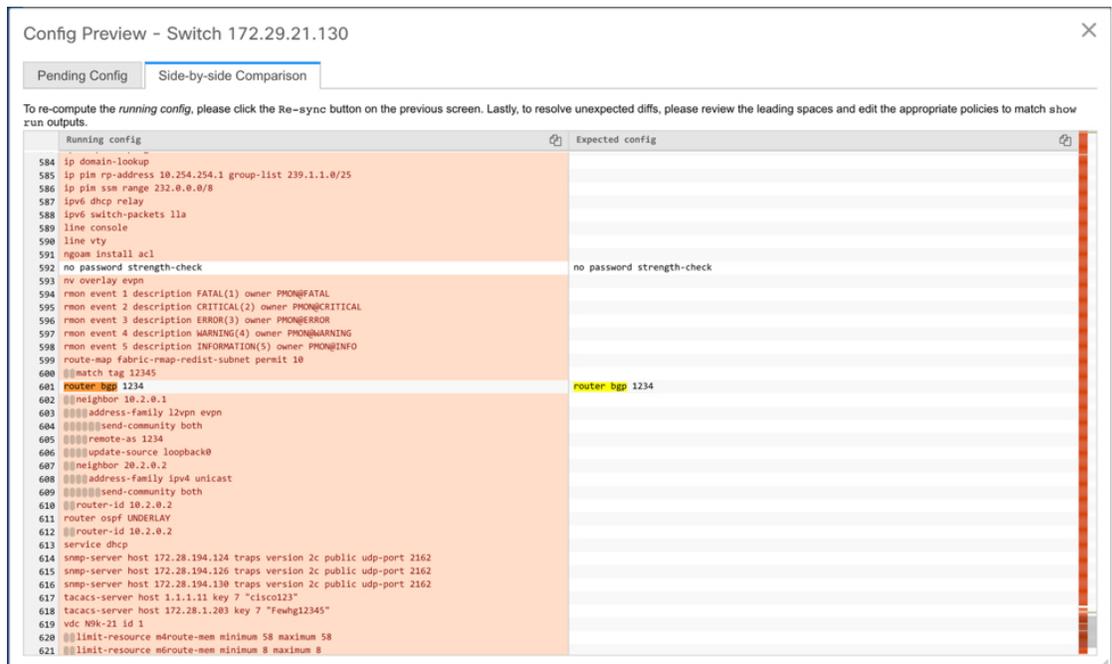
```
router bgp 1234
```

- The removed configuration is only the subset of the configuration that was pushed earlier from DCNM.

## Config Preview - Switch 172.29.21.130

Pending Config    Side-by-side Comparison

```
router bgp 1234
  no neighbor 10.2.0.1
configure terminal
```



For interfaces on the switch in the external fabric, DCNM either manages the entire interface or does not manage it at all. CC checks interfaces in the following ways:

- For any interface, if there is a policy defined and associated with it, then this interface is considered as managed. All configurations associated with this interface must be defined in the associated interface policy. This is applicable for both logical and physical interfaces. Otherwise, CC removes any out-of-band updates made to the interface to change the status to **In-Sync**.
- Interfaces created out-of-band (applies for logical interfaces such as port-channels, sub interfaces, SVIs, loopbacks, etc.), will be discovered by DCNM as part of the regular discovery process. However, since there is no intent for these interfaces, CC will not report an **Out-of-Sync** status for these interfaces.
- For any interface, there can always be a monitor policy associated with it in DCNM. In this case, CC will ignore the interface's configuration when it reports the **In-Sync** or **Out-of-Sync** config compliance status.

## Special Configuration CLIs Ignored for Configuration Compliance

The following configuration CLIs are ignored during configuration compliance checks:

- Any CLI having 'username' along with 'password'
- Any CLI that starts with 'snmp-server user'

Any CLIs that match the above will not show up in pending diffs and clicking **Save & Deploy** in the **Fabric Builder** window will not push such configurations to the switch. These CLIs will not show up in the **Side-by-side Comparison** window also.

To deploy such configuration CLIs, perform the following procedure:

1. Select **Control>Template Library**, and click + to create a new custom template with the required configuration CLIs and the following attributes:
  - Template Type: Policy
  - Template Sub Type: Device
  - Template Content Type: TEMPLATE\_CLI
2. Select **Control>Fabric Builder**, click **Tabular View**, and select a switch in the **Name** column or select **Control>Fabric Builder** and right-click on the device.
3. Click **View/Edit Policies**.
4. Select the custom template created in **Step 1** and click **Push Config** to deploy the configuration to the switch(es).

## Resolving Diffs for Case Insensitive Commands

By default, all diffs generated in DCNM while comparing intent, also known as Expected Configuration, and Running Configuration, are case sensitive. However, the switch has many commands that are case insensitive, and therefore it may not be appropriate to flag these commands as differences. These outlier cases are captured in the **compliance\_case\_insensitive\_clis.txt** text file.

There could be additional commands not included in the existing **compliance\_case\_insensitive\_clis.txt** file that should be treated as case insensitive. If the pending configuration is due to the differences of cases between the Expected Configuration in DCNM and the Running Configuration, you can configure DCNM to ignore these case differences as follows:

1. Modify the following file on the DCNM file system:

```
/usr/local/cisco/dcm/dcm/model-config/compliance_case_insensitive_clis.txt
```

The sample entries in **compliance\_case\_insensitive\_clis.txt** file are displayed as:

```
[root@dcnm98 model-config]# pwd
/usr/local/cisco/dcm/dcm/model-config
[root@dcnm98 model-config]# cat compliance_case_insensitive_clis.txt
"^(no |)interface\s+Port(.)"
"^(no |)interface\s+Loo(.)"
"^(no |)interface\s+Eth(.)"
"^update-source\s+Loo(.)"
"^vrf\s+"
"^hardware profile portmode\s+"
"^(.*)route-map\s+(.)"
"^(.*)neighbor-policy(.)"
"(no |)encapsulation\s+(.)"
"(.*)alert-group\s+(.)"
"^streetaddress\s+(.)"
"^transport email\s+(.)"
"(no |)action\s+(.)"
"(no|)\s+\d*\s+remark.*"
[root@dcnm98 model-config]# █
```

If newer patterns are detected during deployment, and they are triggering pending configurations, you can add these patterns to this file. The patterns need to be valid regex patterns.

This enables DCNM to treat the documented configuration patterns as case insensitive while performing comparisons.

2. Run the following command for each fabric to restart the config compliance container:

```
# docker exec -it `docker ps | grep compliance | grep <fabric name> | awk '{print $1}'`
/usr/bin/pkill python
```

3. Click **Save & Deploy** for fabrics to see the updated comparison outputs.

## Enabling Freeform Configurations on Fabric Switches

In DCNM, you can add custom configurations through freeform policies in the following ways:

1. Fabric-wide
  - On all leaf, border leaf, and border gateway leaf switches in the fabric, at once.
  - On all spine, super spine, border spine, border super spine, border gateway spine and border switches, at once.
2. On a specific switch at the global level.
3. On a specific switch on a per Network or per VRF level.

Leaf switches are identified by the roles Leaf, Border, and Border Gateway. The spine switches are identified by the roles Spine, Border Spine, Border Gateway Spine, Super Spine, Border Super Spine, and Border Gateway Super Spine.



**Note** You can deploy freeform CLIs when you create a fabric or when a fabric is already created. The following examples are for an existing fabric. However, you can use this as a reference for a new fabric.

### Deploying Fabric-Wide Freeform CLIs on Leaf and Spine Switches

1. Click **Control > Fabric Builder**. The Fabric Builder screen comes up. A rectangular box represents each fabric.
2. Click the **Edit Fabric** icon (located on the top right part of the rectangular box) for adding custom configurations to an existing fabric. The **Edit Fabric** screen comes up.  
(If you are creating a fabric for the first time, click **Create Fabric**).
3. Click the **Advanced** tab and update the following fields:

**Leaf Freeform Config** – In this field, add configurations for all leaf, border leaf, and border gateway leaf switches in the fabric.

**Spine Freeform Config** - In this field, add configurations for all Spine, Border Spine, Border Gateway Spine, Super Spine, Border Super Spine, and Border Gateway Super Spine switches in the fabric.



**Note** Copy-paste the intended configuration with correct indentation, as seen in the running configuration on the Nexus switches. For more information, see [Resolving Freeform Config Errors in Switches, on page 230](#).

4. Click **Save**. The fabric topology screen comes up.
5. Click **Save & Deploy** at the top right part of the screen to save and deploy configurations.

Configuration Compliance functionality will ensure that the intended configuration as expressed by those CLIs are present on the switches and if they are removed or there is a mismatch, then it will flag it as a mismatch and indicate that the device is Out-of-Sync.

*Incomplete Configuration Compliance* - On some Cisco Nexus 9000 Series switches, in spite of configuring pending switch configurations using the **Save & Deploy** option, there could be a mismatch between the intended and switch configuration. To resolve the issue, add a **switch\_freeform** policy to the affected switch (as explained in the *Deploy Freeform CLIs on a Specific Switch* section). For example, consider the following persistent pending configurations:

```
line vty
logout-warning 0
```

After adding the above configurations in a policy and saving the updates, click **Save and Deploy** in the topology screen to complete the deployment process.

To bring the switch back in-sync, you can add the above configuration in a **switch\_freeform** policy saved and deployed onto the switch.

### Deploying Freeform CLIs on a Specific Switch

1. Click **Control > Fabric Builder**. The Fabric Builder screen comes up.
2. Click on the rectangular box that represents the fabric. The Fabric Topology screen comes up.



#### Note

To provision freeform CLIs on a new fabric, you have to create a fabric, import switches into it, and then deploy freeform CLIs.

3. Right-click the switch icon and select the **View/edit policies** option.  
The **View/Edit Policies** screen comes up.
4. Click +. The **Add Policy** screen comes up.  
In the **Priority** field, the priority is set to 500 by default. You can choose a higher priority (by specifying a lower number) for CLIs that need to appear higher up during deployment. For example, a command to enable a feature should appear earlier in the list of commands.
5. From the **Policy** field, select **switch\_freeform**.
6. Add or update the CLIs in the **Freeform Config CLI** box.  
Copy-paste the intended configuration with correct indentation, as seen in the running configuration on the Nexus switches. For more information, see [Resolving Freeform Config Errors in Switches, on page 230](#).
7. Click **Save**.  
After the policy is saved, it gets added to the intended configurations for that switch.
8. Close the policy screens. The Fabric Topology screen comes up again.

## 9. Right click the switch and click **Deploy Config**.

The **Save & Deploy** option can also be used for deployment. However, the **Save & Deploy** option will identify mismatch between the intended and running configuration *across all* fabric switches.

### Pointers for `switch_freeform` Policy Configuration:

- You can create multiple instances of the policy.
- For a vPC switch pair, create consistent **switch\_freeform** policies on both the vPC switches.
- When you edit a **switch\_freeform** policy and deploy it onto the switch, you can see the changes being made (in the **Side-by-side** tab of the Preview option).

## Freeform CLI Configuration Examples

### Console line configuration

This example involves deploying some fabric-wide freeform configurations (for all leaf, and spine switches), and individual switch configurations.

Fabric-wide session timeout configuration:

```
line console
  exec-timeout 1
```

Console speed configuration on a specific switch:

```
line console
  speed 115200
```

### ACL configuration

ACL configurations are typically configured on specific switches and not fabric-wide (leaf/spine switches). When you configure ACLs as freeform CLIs on a switch, you should include sequence numbers. Else, there will be a mismatch between the intended and running configuration. A configuration sample with sequence numbers:

```
ip access-list ACL_VTY
  10 deny tcp 172.29.171.67/32 172.29.171.36/32
  20 permit ip any any
ip access-list vlan65-acl
  10 permit ip 69.1.1.201/32 65.1.1.11/32
  20 deny ip any any

interface Vlan65
  ip access-group vlan65-acl in
line vty
  access-class ACL_VTY in
```

If you have configured ACLs without sequence numbers in a **switch\_freeform** policy, update the policy with sequence numbers *as shown in the running configuration of the switch*.

After the policy is updated and saved, right click the device and select the per switch **Deploy Config** option to deploy the configuration. Alternatively, use the **Save and Deploy** option in the fabric topology screen

(within Fabric Builder) so that the fabric triggers Configuration Compliance and resolves the configuration mismatch.

### Resolving Freeform Config Errors in Switches

Copy-paste the running-config to the freeform config with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. Otherwise, configuration compliance in DCNM marks switches as out-of-sync.

Let us see an example of the freeform config of a switch.

```
feature bash-shell
feature telemetry

clock timezone CET 1 0
# Daylight saving time is observed in Metropolitan France from the last Sunday in March
(02:00 CET) to the last Sunday in October (03:00 CEST)
clock summer-time CEST 5 Sunday March 02:00 5 Sunday October 03:00 60
clock protocol ntp

telemetry
  destination-profile
  use-vrf management
```

The highlighted line about the daylight saving time is a comment that is not displayed in the **show running config** command output. Therefore, configuration compliance marks the switch as out-of-sync because the intent does not match the running configuration.

Let us check the running config in the switch for the clock protocol.

```
spinel# show run all | grep "clock protocol"
clock protocol ntp vdc 1
```

You can see that **vdc 1** is missing from the freeform config.

In this example, let us copy-paste the running config to the freeform config.

Here is the updated freeform config:

```
feature bash-shell
feature telemetry

clock timezone CET 1 0
clock summer-time CEST 5 Sunday March 02:00 5 Sunday October 03:00 60
clock protocol ntp vdc 1

telemetry
  destination-profile
  use-vrf management
```

After you copy-paste the running config and deploy, the switch will be in-sync. When you click **Save & Deploy**, the **Side-by-side Comparison** tab in the **Config Preview** window provides you information about the difference between the defined intent and the running config.

# Management

The Management menu includes the following submenus:

## Resources

Cisco DCNM allows you to manage the resources. The following table describes the fields that appear on this page.

| Field              | Description                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Scope Type         | Specifies the scope level at which the resources are managed. The scope types can be <b>Fabric</b> , <b>Device</b> , <b>DeviceInterface</b> , <b>DevicePair</b> , <b>Fabric</b> , and <b>Link</b> .                                                    |
| Scope              | Specifies the resource usage scope. Valid values are the switch serial numbers or fabric names. Resources with serial numbers are unique, and can be used on the serial number of the switch only.                                                     |
| Allocated Resource | Specifies if the resources are managed with device, device interface, or fabric. Valid values are ID type, subnet, or IP addresses.                                                                                                                    |
| Allocated To       | Specifies the entity name for which the resource is allocated.                                                                                                                                                                                         |
| Resource Type      | Specifies the resource type. The valid values are <b>TOP_DOWN_VRF_LAN</b> , <b>TOP_DOWN_NETWORK_VLAN</b> , <b>LOOPBACK_ID</b> , <b>VPC_ID</b> , and so on.                                                                                             |
| Is Allocated?      | Specifies if the resource is allocated or not. The value is set to <b>True</b> if the resource is permanently allocated to the given entity. The value is set to <b>False</b> if the resource is reserved for an entity and not permanently allocated. |
| Allocated On       | Specifies the date and time of the resource allocation.                                                                                                                                                                                                |

## Adding, Editing, Re-Discovering and Removing VMware Servers

This section contains the following:

### Adding a Virtual Center Server

You can add a virtual center server from Cisco DCNM.

#### Procedure

**Step 1** Choose **Control > Management > Virtual Machine Manager**.

You see the list of VMware servers (if any) that are managed by Cisco DCNM-LAN in the table.

- Step 2** Click **Add**.  
You see the **Add VCenter** window.
- Step 3** Enter the **Virtual Center Server** IP address for this VMware server.
- Step 4** Enter the **User Name** and **Password** for this VMware server.
- Step 5** Click **Add** to begin managing this VMware server.
- 

## Deleting a VMware Server

You can remove a VMware server from the Cisco DCNM.

### Procedure

---

- Step 1** Choose **Control > Management > Virtual Machine Manager**.
- Step 2** Select the check box next to the VMware server that you want to remove and click **Delete** to discontinue data collection for that VMware server.
- 

## Editing a VMware Server

You can edit a VMware server from Cisco DCNM Web Client.

### Procedure

---

- Step 1** Choose **Control > Management > Virtual Machine Manager**.
- Step 2** Check the check box next to the VMware server that you want to edit and click **Edit** virtual center icon.  
You see the **Edit VCenter** dialog box.
- Step 3** Enter a the **User Name** and **Password**.
- Step 4** Select managed or unmanaged status.
- Step 5** Click **Apply** to save the changes.
- 

## Rediscovering a VMware Server

You can rediscover a VMware server from Cisco DCNM.

### Procedure

---

- Step 1** Choose **Control > Management > Virtual Machine Manager**.
- Step 2** Select the check box next to the VMware that you want to rediscover.
- Step 3** Click **Rediscover**.  
A dialog box with warning "Please wait for rediscovery operation to complete." appears.

**Step 4** Click **OK** in the dialog box.

## Template Library

You can add, edit, or delete templates that are configured across different Cisco Nexus and Cisco MDS platforms using Cisco DCNM Web client. From Cisco DCNM Web client home page, choose **Control > Template Library > Templates**. The following parameters are displayed for each template that is configured on Cisco DCNM Web client. Templates support JavaScript. You can use the JavaScript function in a template to perform arithmetic operations and string manipulations in the template syntax.

The following table describes the fields that appear on this page.

**Table 1: Templates Operations**

| Field                    | Description                                                                                                                                                                                                   |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Add Template             | Allows you to add a new template.                                                                                                                                                                             |
| Modify/View Template     | Allows you to view the template definition and modify as required.                                                                                                                                            |
| Save Template As         | Allows you to save the selected template in a different name. You can edit the template as required.                                                                                                          |
| Delete Template          | Allows you to delete a template                                                                                                                                                                               |
| Import Template          | Allows you to import a template from your local directory, one at a time.                                                                                                                                     |
| Export template          | Allows you to export the template configuration to a local directory location.                                                                                                                                |
| Import Template Zip File | Allows you to import .zip file, that contains more than one template that is bundled in a .zip format<br><br>All the templates in the ZIP file are extracted and listed in the table as individual templates. |



**Note** Notifications appear next to **Import Template Zip File** if there are issues while loading templates after restarting the server. Click the notifications to see the errors in the **Issues in loading Template** window. Templates with errors are not listed in the **Templates** window. To import these templates, correct the errors, and import them.

**Table 2: Template Properties**

| Field         | Description                                   |
|---------------|-----------------------------------------------|
| Template Name | Displays the name of the configured template. |

| Field                 | Description                                                                                                                                                                                             |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Template Description  | Displays the description that is provided while configuring templates.                                                                                                                                  |
| Tags                  | Displays the tag that is assigned for the template and aids to filter templates based on the tags.                                                                                                      |
| Supported Platforms   | Displays the supported Cisco Nexus platforms compatible with the template. Check the check box of platforms that are supported with the template.<br><br><b>Note</b> You can select multiple platforms. |
| Template Type         | Displays the type of the template.                                                                                                                                                                      |
| Template Sub Type     | Specifies the sub type that is associated with the template.                                                                                                                                            |
| Template Content Type | Specifies if it is Jython or Template CLI.                                                                                                                                                              |

**Table 3: Advanced Template Properties**

| Field        | Description                                       |
|--------------|---------------------------------------------------|
| Implements   | Displays the abstract template to be implemented. |
| Dependencies | Specifies the specific feature of a switch.       |
| Published    | Specifies if the template is published or not.    |
| Imports      | Specifies the base template for importing.        |

In addition, from the menu bar, choose **Control > Template Library > Templates** and you can also:

- Click **Show Filter** to filter the templates that is based on the headers.
- Click **Print** to print the list of templates.
- Click **Export to Excel** to export the list of template to a Microsoft Excel spreadsheet.

This section contains the following:

## Template Structure

The configuration template content mainly consists of four parts. Click the **Help** icon next to the **Template Content** for information about editing the content of the template.

This section contains the following:

## Template Format

This section describes the basic information of the template. The possible fields are as detailed in the table below.

| Property Name      | Description                                                                                            | Valid Values                                                                                                                                                                                                                                                                                       | Optional? |
|--------------------|--------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| name               | The name of the template                                                                               | Text                                                                                                                                                                                                                                                                                               | No        |
| description        | Brief description about the template                                                                   | Text                                                                                                                                                                                                                                                                                               | Yes       |
| userDefined        | Indicates whether the user created the template.<br>Value is 'true' if user created.                   | "true" or "false"                                                                                                                                                                                                                                                                                  | Yes       |
| supportedPlatforms | List of device platforms supports this configuration template. Specify 'All' to support all platforms. | N1K, N3K, N3500, N4K, N5K, N5500, N5600, N6K, N7K, N9K, MDS, VDC, N9K-9000v, IOS-XE, IOS-XR, Others, All list separated by comma.                                                                                                                                                                  | No        |
| templateType       | Specifies the type of Template used.                                                                   | <ul style="list-style-type: none"> <li>• CLI</li> <li>• POAP</li> </ul> <p><b>Note</b> POAP option is not applicable for Cisco DCNM LAN Fabric deployment.</p> <ul style="list-style-type: none"> <li>• POLICY</li> <li>• SHOW</li> <li>• PROFILE</li> <li>• FABRIC</li> <li>• ABSTRACT</li> </ul> | Yes       |

| Property Name   | Description                                          | Valid Values | Optional? |
|-----------------|------------------------------------------------------|--------------|-----------|
| templateSubType | Specifies the sub type associated with the template. |              |           |

| Property Name | Description | Valid Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Optional? |
|---------------|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
|               |             | <ul style="list-style-type: none"> <li>• CLI               <ul style="list-style-type: none"> <li>• N/A</li> </ul> </li> <li>• POAP               <ul style="list-style-type: none"> <li>• N/A</li> <li>• VXLAN</li> <li>• FABRICPATH</li> <li>• VLAN</li> <li>• PMN</li> </ul> </li> <li><b>Note</b> POAP option is not applicable for Cisco DCNM LAN Fabric deployment.</li> <li>• POLICY               <ul style="list-style-type: none"> <li>• VLAN</li> <li>• INTERFACE_VLAN</li> <li>• INTERFACE_VPC</li> <li>• INTERFACE_ETHNET</li> <li>• INTERFACE_BD</li> <li>• INTERFACE_CHANNEL</li> <li>• INTERFACE_FC</li> <li>• INTERFACE_MGMT</li> <li>• INTERFACE_COBACK</li> <li>• INTERFACE_NVE</li> <li>• INTERFACE_VFC</li> <li>• INTERFACE_NFC</li> <li>• DEVICE</li> <li>• FEX</li> <li>• NIRA_FABRIC_LINK</li> <li>• NIER_FABRIC_LINK</li> </ul> </li> </ul> |           |

| Property Name | Description | Valid Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Optional? |
|---------------|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
|               |             | <ul style="list-style-type: none"> <li>• INTERFACE</li> <li>• SHOW               <ul style="list-style-type: none"> <li>• VLAN</li> <li>• INTERFACE_VLAN</li> <li>• INTERFACE_VPC</li> <li>• INTERFACE_ETHNET</li> <li>• INTERFACE_BD</li> <li>• <del>INTERFACE_CHANNEL</del></li> <li>• INTERFACE_FC</li> <li>• INTERFACE_MGMT</li> <li>• INTERFACE_COBACK</li> <li>• INTERFACE_NVE</li> <li>• INTERFACE_VFC</li> <li>• <del>INTERFACE_CHANNEL</del></li> </ul> </li> <li>• DEVICE               <ul style="list-style-type: none"> <li>• FEX</li> <li>• <del>NIRAFABRIC_LINK</del></li> <li>• <del>NIRAFABRIC_LINK</del></li> </ul> </li> <li>• INTERFACE               <ul style="list-style-type: none"> <li>• PROFILE                   <ul style="list-style-type: none"> <li>• VXLAN</li> </ul> </li> </ul> </li> <li>• FABRIC               <ul style="list-style-type: none"> <li>• NA</li> </ul> </li> </ul> |           |

| Property Name | Description | Valid Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Optional? |
|---------------|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
|               |             | <ul style="list-style-type: none"> <li>• ABSTRACT</li> <li>• VLAN</li> <li>• INTERFACE_VLAN</li> <li>• INTERFACE_VPC</li> <li>• INTERFACE_ETH</li> <li>• INTERFACE_HRNET</li> <li>• INTERFACE_BD</li> <li>• INTERFACE_FANL</li> <li>• INTERFACE_FC</li> <li>• INTERFACE_MGMT</li> <li>• INTERFACE_LOOPBACK</li> <li>• INTERFACE_NVE</li> <li>• INTERFACE_VFC</li> <li>• INTERFACE_NFCANL</li> <li>• DEVICE</li> <li>• FEX</li> <li>• NIRA_FABRIC_LINK</li> <li>• NIER_FABRIC_LINK</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> <li><b>Note</b> POAP option is not applicable for Cisco DCNM LAN Fabric deployment</li> <li>• POLICY               <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> <li>• PYTHON</li> </ul> </li> <li>• SHOW               <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> </ul> </li> <li>• PROFILE               <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> <li>• PYTHON</li> </ul> </li> <li>• FABRIC               <ul style="list-style-type: none"> <li>• PYTHON</li> </ul> </li> <li>• ABSTRACT               <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> <li>• PYTHON</li> </ul> </li> </ul> | Yes       |
| implements    | Used to implement the abstract template.                         | Text                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Yes       |
| dependencies  | Used to select the specific feature of a switch.                 | Text                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Yes       |
| published     | Used to Mark the template as read only and avoids changes to it. | “true” or “false”                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Yes       |

## Template Variables

This section contains declared variables, the data type, default values, and valid values conditions for the parameters that are used in the template. These declared variables are used for value substitution in the template content section during the dynamic command generation process. Also these variables are used in decision making and in iteration blocks in the template content section. Variables have predefined data types. You can also add a description about the variable. The following table describes the syntax and usage for the available datatypes.

| Variable Type  | Valid Value                                                                                           | Iterative? |
|----------------|-------------------------------------------------------------------------------------------------------|------------|
| boolean        | true false                                                                                            | No         |
| enum           | Example: running-config,<br>startup-config                                                            | No         |
| float          | Floating number format                                                                                | No         |
| floatRange     | Example: 10.1,50.01                                                                                   | Yes        |
| Integer        | Any number                                                                                            | No         |
| integerRange   | Contiguous numbers separated by<br>“_”<br>Discrete numbers separated by “,”<br>Example: 1-10,15,18,20 | Yes        |
| interface      | Format: <if type><slot>[/<sub<br>slot>]/<port><br>Example: eth1/1, fa10/1/2 etc.                      | No         |
| interfaceRange | Example: eth10/1/20-25,<br>eth11/1-5                                                                  | Yes        |
| ipAddress      | IPv4 OR IPv6 address                                                                                  | No         |

| Variable Type          | Valid Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Iterative? |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| ipAddressList          | <p>You can have a list of IPv4, IPv6, or a combination of both types of addresses.</p> <p>Example 1: 172.22.31.97,<br/>172.22.31.99,<br/>172.22.31.105,<br/>172.22.31.109</p> <p>Example 2:<br/>2001:0cb8:85a3:0000:0000:8a2e:0370:7334,<br/><br/>2001:0cb8:85a3:0000:0000:8a2e:0370:7335,<br/><br/>2001:0cb8:85a3:1230:0000:8a2f:0370:7334</p> <p>Example 3: 172.22.31.97,<br/>172.22.31.99,<br/><br/>2001:0cb8:85a3:0000:0000:8a2e:0370:7334,<br/><br/>172.22.31.254</p> | Yes        |
| ipAddressWithoutPrefix | <p>Example: 192.168.1.1</p> <p>or</p> <p>Example: 1:2:3:4:5:6:7:8</p>                                                                                                                                                                                                                                                                                                                                                                                                      | No         |
| ipV4Address            | IPv4 address                                                                                                                                                                                                                                                                                                                                                                                                                                                               | No         |
| ipV4AddressWithSubnet  | Example: 192.168.1.1/24                                                                                                                                                                                                                                                                                                                                                                                                                                                    | No         |
| ipV6Address            | IPv6 address                                                                                                                                                                                                                                                                                                                                                                                                                                                               | No         |
| ipV6AddressWithPrefix  | <p>Example: 1:2:3:4:5:6:7:8</p> <p>22</p>                                                                                                                                                                                                                                                                                                                                                                                                                                  | No         |
| ipV6AddressWithSubnet  | IPv6 Address with Subnet                                                                                                                                                                                                                                                                                                                                                                                                                                                   | No         |
| ISISNetAddress         | <p>Example:</p> <p>49.0001.00a0.c96b.c490.00</p>                                                                                                                                                                                                                                                                                                                                                                                                                           | No         |
| long                   | Example: 100                                                                                                                                                                                                                                                                                                                                                                                                                                                               | No         |
| macAddress             | 14 or 17 character length MAC address format                                                                                                                                                                                                                                                                                                                                                                                                                               | No         |
| string                 | <p>Free text, for example, used for the description of a variable</p> <p>Example:</p> <pre>string scheduledTime { regularExpr=^([01]\d 2[0-3]):([0-5]\d)\$; }</pre>                                                                                                                                                                                                                                                                                                        | No         |

| Variable Type                                    | Valid Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Iterative?                                                                                              |
|--------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| string[]                                         | Example: {a,b,c,str1,str2}                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Yes                                                                                                     |
| struct                                           | <p>Set of parameters that are bundled under a single variable.</p> <pre> struct &lt;structure name declaration &gt; { &lt;parameter type&gt; &lt;parameter 1&gt;; &lt;parameter type&gt; &lt;parameter 2&gt;; ... } [&lt;structure_inst1&gt;] [, &lt;structure_inst2&gt;] [, &lt;structure_array_inst3 []&gt;;  struct interface_detail { string inf_name; string inf_description; ipAddress inf_host; enum duplex { validValues = auto, full, half; }; }myInterface, myInterfaceArray[]; </pre> | <p>No</p> <p><b>Note</b> If the struct variable is declared as an array, the variable is iterative.</p> |
| wwn<br>(Available only in Cisco DCNM Web Client) | <p>Example:<br/>20:01:00:08:02:11:05:03</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                      | No                                                                                                      |

## Variable Meta Property

Each variable that is defined in the template variable section has a set of meta properties. The meta properties are mainly the validation rules that are defined for the variable.

The following table describes the various meta properties applicable for the available variable types.

| Variable Type | Description                       | Variable Meta Property |              |                |     |     |          |          |          |          |            |            |              |
|---------------|-----------------------------------|------------------------|--------------|----------------|-----|-----|----------|----------|----------|----------|------------|------------|--------------|
|               |                                   | default Value          | valid Values | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| boolean       | A boolean value.<br>Example: true | Yes                    |              |                |     |     |          |          |          |          |            |            |              |
| enum          |                                   |                        | Yes          |                |     |     |          |          |          |          |            |            |              |

| Variable Type  | Description                                                    | Variable Meta Property |              |                |     |     |          |          |          |          |            |            |              |
|----------------|----------------------------------------------------------------|------------------------|--------------|----------------|-----|-----|----------|----------|----------|----------|------------|------------|--------------|
|                |                                                                | default Value          | valid Values | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| float          | signed real number<br>Example:<br>75.56,<br>-8.5               | Yes                    | Yes          | Yes            | Yes | Yes |          |          |          |          |            |            |              |
| floatRange     | range of signed real numbers<br>Example:<br>50.5<br>-<br>54.75 | Yes                    | Yes          | Yes            | Yes | Yes |          |          |          |          |            |            |              |
| integer        | signed number<br>Example:<br>50,<br>-75                        | Yes                    | Yes          |                | Yes | Yes |          |          |          |          |            |            |              |
| integerRange   | Range of signed numbers<br>Example:<br>50-65                   | Yes                    | Yes          |                | Yes | Yes |          |          |          |          |            |            |              |
| interface      | specifies interface<br>Example:<br>Ethernet<br>5/10            | Yes                    | Yes          |                |     |     | Yes      | Yes      | Yes      | Yes      |            |            |              |
| interfaceRange |                                                                | Yes                    | Yes          |                |     |     | Yes      | Yes      | Yes      | Yes      |            |            |              |
| ipAddress      | IP address in IPv4 or IPv6 format                              | Yes                    |              |                |     |     |          |          |          |          |            |            |              |

| Variable Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                              | Variable Meta Property |              |                |     |     |          |          |          |          |            |            |              |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|--------------|----------------|-----|-----|----------|----------|----------|----------|------------|------------|--------------|
|               |                                                                                                                                                                                                                                                                                                                                                                                                                                          | default Value          | valid Values | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| ipAddress     | <p>You can have a list of IPv4, IPv6, or a combination of both types of addresses.</p> <p>Example 1:<br/>172.31.9,<br/>172.31.9,<br/>172.31.15,<br/>172.31.10</p> <p>Example 2:<br/><del>172.31.9,</del><br/><del>172.31.9,</del><br/><del>172.31.15,</del></p> <p>Example 3:<br/>172.31.9,<br/>172.31.9,<br/><del>172.31.9,</del><br/>172.31.24</p> <p><b>Note</b> Separate the addresses in the list using commas and not hyphens.</p> | Yes                    |              |                |     |     |          |          |          |          |            |            |              |

| Variable Type      | Description                                    | Variable Meta Property |              |                |     |     |          |          |          |          |            |            |              |
|--------------------|------------------------------------------------|------------------------|--------------|----------------|-----|-----|----------|----------|----------|----------|------------|------------|--------------|
|                    |                                                | default Value          | valid Values | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| <del>ipAdns</del>  | IPv4 or IPv6 Address (does not require prefix) |                        |              |                |     |     |          |          |          |          |            |            |              |
| <del>ipAdns</del>  | IPv4 address                                   | Yes                    |              |                |     |     |          |          |          |          |            |            |              |
| <del>ipAdns</del>  | IPv4 Address with Subnet                       | Yes                    |              |                |     |     |          |          |          |          |            |            |              |
| <del>ip6Adns</del> | IPv6 address                                   | Yes                    |              |                |     |     |          |          |          |          |            |            |              |
| <del>ip6Adns</del> | IPv6 Address with prefix                       | Yes                    |              |                |     |     |          |          |          |          |            |            |              |
| <del>ip6Adns</del> | IPv6 Address with Subnet                       | Yes                    |              |                |     |     |          |          |          |          |            |            |              |
| <del>ip6Adns</del> | Example:<br>4008:5:50                          |                        |              |                |     |     |          |          |          |          |            |            |              |
| long               | Example:<br>100                                | Yes                    |              |                | Yes | Yes |          |          |          |          |            |            |              |
| <del>macAdns</del> | MAC address                                    |                        |              |                |     |     |          |          |          |          |            |            |              |

| Variable Type | Description                                                                                              | Variable Meta Property |              |                |     |     |          |          |          |          |            |            |              |
|---------------|----------------------------------------------------------------------------------------------------------|------------------------|--------------|----------------|-----|-----|----------|----------|----------|----------|------------|------------|--------------|
|               |                                                                                                          | default Value          | valid Values | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| string        | literal string<br><br>Example for string<br><br>Regular expression string<br><br>statement {<br>...<br>} | Yes                    |              |                |     |     |          |          |          |          | Yes        | Yes        | Yes          |
| string[]      | string literals that are separated by a comma (,)<br><br>Example:<br>{string1,<br>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 <del>params</del> that are bundled under a single variable.<br><br>struct<br><br><structure name declaration><br>> {<br><parameter type><br><br><parameter 1>;<br><parameter type><br><br><parameter 2>;<br>...<br>}<br><struct1><br>[,<br><struct2><br>[,<br><struct3><br>[ ]>; |                        |              |                |     |     |          |          |          |          |            |            |              |
| wnn           | WWN address                                                                                                                                                                                                                                                                             |                        |              |                |     |     |          |          |          |          |            |            |              |

### Example: Meta Property Usage

```
##template variables

integer VLAN_ID {
min = 100;
max= 200;
};

string USER_NAME {
defaultValue = admin123;
minLength = 5;
};

struct interface_a{
```

```

string inf_name;
string inf_description;
ipAddress inf_host;
enum duplex {
    validValues = auto, full, half;
};
}myInterface;

##

```

## Variable Annotation

You can configure the variable properties marking the variables using annotations.



**Note** Variable Annotations are available for POAP only. However, the annotations do not impact on the template type 'CLI'.

The following annotations can be used in the template variable section.

| Annotation Key          | Valid Values                                                         | Description                                       |
|-------------------------|----------------------------------------------------------------------|---------------------------------------------------|
| AutoPopulate            | Text                                                                 | Copies values from one field to another           |
| DataDepend              | Text                                                                 |                                                   |
| Description             | Text                                                                 | Description of the field appearing in the window  |
| DisplayName             | Text<br><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"                                                    |                                                   |

| Annotation Key          | Valid Values                                                                                       | Description                                                                                                                               |
|-------------------------|----------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| IsFEXID                 | “true” or “false”                                                                                  |                                                                                                                                           |
| IsGateway               | “true” or “false”                                                                                  | Validates if the IP address is a gateway                                                                                                  |
| IsInternal              | “true” or “false”                                                                                  | Makes the fields internal and does not display them on the window<br><br><b>Note</b> Use this annotation only for the ipAddress variable. |
| IsManagementIP          | “true” or “false”<br><br><b>Note</b> This annotation must be marked only for variable “ipAddress”. |                                                                                                                                           |
| IsMandatory             | “true” or “false”                                                                                  | Validates if a value should be passed to the field mandatorily                                                                            |
| IsMTU                   | “true” or “false”                                                                                  |                                                                                                                                           |
| IsMultiCastGroupAddress | “true” or “false”                                                                                  |                                                                                                                                           |
| IsMultiLineString       | “true” or “false”                                                                                  | Converts a string field to multiline string text area                                                                                     |
| IsMultiplicity          | “true” or “false”                                                                                  |                                                                                                                                           |
| IsPassword              | “true” or “false”                                                                                  |                                                                                                                                           |
| IsPositive              | “true” or “false”                                                                                  | Checks if the value is positive                                                                                                           |
| IsReplicationMode       | “true” or “false”                                                                                  |                                                                                                                                           |
| IsShow                  | “true” or “false”                                                                                  | Displays or hides a field on the window                                                                                                   |
| IsSiteId                | “true” or “false”                                                                                  |                                                                                                                                           |
| IsSourceDevice          | “true” or “false”                                                                                  |                                                                                                                                           |
| IsSourceFabric          | “true” or “false”                                                                                  |                                                                                                                                           |
| IsSourceInterface       | “true” or “false”                                                                                  |                                                                                                                                           |

| Annotation Key           | Valid Values      | Description                                          |
|--------------------------|-------------------|------------------------------------------------------|
| IsSourceSwitchName       | “true” or “false” |                                                      |
| IsSwitchName             | “true” or “false” |                                                      |
| IsRMID                   | “true” or “false” |                                                      |
| IsVPCDomainID            | “true” or “false” |                                                      |
| IsVPCID                  | “true” or “false” |                                                      |
| IsVPCPeerLinkPort        | “true” or “false” |                                                      |
| IsVPCPeerLinkPortChannel | “true” or “false” |                                                      |
| IsVPCPortChannel         | “true” or “false” |                                                      |
| Password                 | Text              | Validates the password field                         |
| PeerOneFEXID             | “true” or “false” |                                                      |
| PeerTwoFEXID             | “true” or “false” |                                                      |
| PeerOnePCID              | “true” or “false” |                                                      |
| PeerTwoPCID              | “true” or “false” |                                                      |
| PrimaryAssociation       |                   |                                                      |
| ReadOnly                 | “true” or “false” | Makes the field read-only                            |
| ReadOnlyOnEdit           | “true” or “false” |                                                      |
| SecondaryAssociation     | Text              |                                                      |
| Section                  |                   |                                                      |
| UsePool                  | “true” or “false” |                                                      |
| UseDNSReverseLookup      |                   |                                                      |
| Username                 | Text              | Displays the username field on the window            |
| Warning                  | Text              | Provides text to override the Description annotation |

#### Example: AutoPopulate Annotation

```
##template variables
string BGP_AS;
@ (AutoPopulate="BGP_AS")
```

```
    string SITE_ID;
##
```

### Example: DisplayName Annotation

```
##template variables
@(DisplayName="Host Name", Description = "Description of the host")
String hostname;
@(DisplayName="Host Address", Description = " test description" IsManagementIP=true)
ipAddress hostAddress;
##
```

### Example: IsMandatory Annotation

```
##template variables
@(IsMandatory="ipv6!=null")
ipV4Address ipv4;
@(IsMandatory="ipv4!=null")
ipV6Address ipv6;
##
```

### Example: IsMultiLineString Annotation

```
##template variables
@(IsMultiLineString=true)
string EXTRA_CONF_SPINE;
##
```

### IsShow Annotation

```
##template variables
boolean isVlan;
@(IsShow="isVlan==true")
integer vlanNo;
##

##template variables
boolean enableScheduledBackup;
@(IsShow="enableScheduledBackup==true",Description="Server time")
string scheduledTime;
##
The condition "enableScheduledBackup==true" evaluates to true/false

##template variables
@(Enum="Manual,Back2BackOnly,ToExternalOnly,Both")
string VRF_LITE_AUTOCONFIG;
@(IsShow="VRF_LITE_AUTOCONFIG!=Manual", Description="Target Mask")
integer DCI_SUBNET_TARGET_MASK
##
The condition "VRF_LITE_AUTOCONFIG!=Manual" matches string comparison to evaluate to true
or false
```

### Example: Warning Annotation

```
##template variables
@(Warning="This is a warning msg")
    string SITE_ID;
##
```

## Templates Content

This section includes the configuration commands and any parameters that you want to include in the template. These commands can include the variables declared in the template variables section. During the command generation process the variable values are substituted appropriately in the template content.



**Note** You must specify the commands that you include as if you were entering them in the global configuration command mode on any device. You must consider the command mode when you include commands.

Template content is governed by the usage of variables.

- **Scalar variables:** does not take a range or array of values which cannot be used for iteration (In the variable types table those marked iterate-able as 'No'). Scalar variables must be defined inside the template content.

```
Syntax: $$<variable name>$$
Example: $$USER_NAME$$
```

- **Iterative variables:** used for block iteration. These loop variable must be accessed as shown below inside the iteration block.

```
Syntax:@<loop variable>
Example:
foreach val in $$INTEGER_RANGE_VALUE$$ {
@val
}
```

- **Scalar Structure Variable:** Structure member variables can be accessed inside the template content.

```
Syntax: $$<structure instance name>.<member variable name>$$
Example: $$myInterface.inf_name$$
```

- **Array Structure Variable:** Structure member variables can be accessed inside the template content.

```
Syntax: $$<structure instance name>.<member variable name>$$
Example: $$myInterface.inf_name$$
```

In addition to the template variables, you can use the conditional and iterative command generation using the following statements:

- **if-else if-else Statement:** makes a logical decision in inclusion/exclusion of set of configuration command based on the value assigned for the variable in it.

```
Syntax: if(<operand 1> <logical operator> <operand 2>){
command1 ..
command2..
..
}
else if (<operand 3> <logical operator> <operand 4> )
{
Command3 ..
Command4..
..
}
else
{
```

```

Command5 ..
Command6..
..
}
Example: if-else if-else statement
if($$USER_NAME$$ == 'admin'){
Interface2/10
no shut
}
else {
Interface2/10
shut
}

```

- **foreach Statement:** used for iterating a block of commands. The iteration is performed based on the assigned loop variable value.

```

Syntax:
foreach <loop index variable> in $$<loop variable>$$ {
@<loop index variable> ..
}
Example: foreach Statement
foreach ports in $$MY_INF_RANGES$${
interface @ports
no shut
}

```

- **Optional parameters:** By default all parameters are mandatory. To make a parameter optional, you must annotate the parameter.

In the variable section, you can include the following command:

- **@(IsMandatory=false)**
- **Integer frequency;**

In the template content section, a command can be excluded or included without using “if” condition check, by assigning a value to the parameter. The optional command can be framed as below:

- **probe icmp [frequency frequency-value] [timeout seconds] [retry-count retry-count-value]**

## Template Content Editor

The template content editor has the following features:

- **Syntax highlighting:** The editor highlights the syntax, like different types of statements, keywords, and so on, for Python scripting.
- **Autocompletion:** The editor suggests the template datatypes, annotations, or metaproperties when you start typing.
- **Go to line:** You can navigate to the exact line in the template content editor instead of scrolling. Press **Command-L** in Mac or **Ctrl-L** in Windows, and enter the line number to which you want to navigate to in the pop-up window.

If you enter a value greater than the number of lines in the editor, you will be navigated to the last line in the editor window.

- **Template search and replace:** Press **Command-F** in Mac or **Ctrl-F** in Windows, enter the search term in the **Search for** field, and select the type of search in the search window. You can perform the following searches in the editor:
  - **RegExp Search:** You can perform the regular expression search in the editor.
  - **CaseSensitive Search:** You can perform a case-sensitive search in the editor.
  - **Whole Word Search:** You can perform a whole word search to find the exact words in the editor. For example, a regular search for the word "play" returns results where it is part of words like "display," but the whole word search returns results only when there is an exact match for the word "play".
  - **Search In Selection:** You can perform a search in the selected content. Select the content to which you want to limit the search and enter the search term.

Choose the + icon in the search window to use the replace option. Enter the replacing word in the **Replace with** field. You can replace the selected word once by selecting **Replace**. To replace all the occurrences of the selected word, select **All**.

- **Code folding:** You can expand or group code blocks in the editor by clicking the arrow next to their line numbers.
- **Other features:** The editor automatically indents the code, the closing braces, and highlights the matching parenthesis.

## Template Editor Settings

You can edit the following features of a template editor by clicking **Template Editor Settings**.

- **Theme:** Select the required theme for the editor from the drop-down list.
- **KeyBinding:** Select the editor mode from the **KeyBinding** drop-down list to customize the editor. **Vim** and **Ace** modes are supported. The default is **Ace**.
- **Font Size:** Select the required font size for the editor.

## Advanced Features

The following are the advanced features available to configure templates.

- **Assignment Operation**

Config template supports assignment of variable values inside the template content section. The values are validated for the declared data type of the variable. If there is a mismatch, the value is not assigned.

Assignment operation can be used under the following guidelines:

- The operator on the left must be any of the template parameters or a for loop parameter.
- The operator on the right values can be any of the values from template parameters, for loop parameters, literal string values surrounded by quotes or simple string values.

If a statement does not follow these guidelines, or if it does not suit this format, it will not be considered as assignment operation. It is substituted during command generation like other normal lines.

Example: Template with assignment operation

```

##template properties
name =vlan creation;
userDefined= true;
supportedPlatforms = All;
templateType = CLI;
published = false;
##
##template variables
integerRange vlan_range;
@(internal=true)
integer vlanName;
##
##template content
foreach vlanID in $$vlan_range$${
vlan @vlanID
$$vlanName$$=@vlanID
name myvlan$$vlanName$$
}
##

```

- Evaluate methods

Config template uses the Java runtime provided JavaScript 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 JavaScript file.

- Dynamic decision

Config template provides a special internal variable “LAST\_CMD\_RESPONSE”. This variable stores the last command response from the device during the execution of the command. This can be used in the config template content to make dynamic decisions to deliver the commands that are based on the device condition.




---

**Note** The if block must be followed by an else block in a new line, which can be empty.

---

An example use case to create a VLAN, if it does not exist on the device.

Example: Create VLAN

```

##template content
show vlan id $$vlan_id$$
if($$LAST_CMD_RESPONSE$$ contains "not found"){
vlan $$vlan_id$$
}
else{

```

```
}
##
```

This special implicit variable can be used only in the “IF” blocks.

- **Template referencing**

You can have a base template with all the variables defined. This base template can be imported to multiple templates. The base template content is substituted in the appropriate place of the extending template. The imported template parameters and the contents can be accessed inside the extending template.

```
Example: Template Referencing
Base template:
##template properties
  name =a vlan base;
  userDefined= true;
  supportedPlatforms = All;
  templateType = CLI;
  published = false;
  timestamp = 2015-07-14 16:07:52;
  imports = ;
##
##template variables
  integer vlan_id;
##
##template content
  vlan $$vlan_id$$
##

Derived Template:
##template properties
  name =a vlan extended;
  userDefined= true;
  supportedPlatforms = All;
  templateType = CLI;
  published = false;
  timestamp = 2015-07-14 16:07:52;
  imports = a vlan base,template2;
##
##template variables
  interface vlanInterface;
##
##template content
  <substitute a vlan base>
  interface $$vlanInterface$$
  <substitute a vlan base>
##
```

When you launch the extended template, the parameter inputs for the base template are also obtained. In addition, the substituted content is used for complete CLI command generation.

## Adding a Template

To add user-defined templates and schedule jobs from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

**Step 1** Choose **Control > Template Library**.

The **Templates** window is displayed with the name of the template along with its description, supported platforms, and tags.

**Step 2** Click **Add** to add a new template.

The Template Properties window appears.

**Step 3** Specify a template name, description, tags, and supported platforms for the new template.

**Step 4** Specify a **Template Type** for the template.

**Step 5** Select a **Template Sub Type** and **Template Content Type** for the template.

**Step 6** Click the **Advanced** tab to edit other properties like **Implements**, **Dependencies**, **Published**, and **Imports**. Select **Published** to make the template read-only. You cannot edit a published template.

**Step 7** From the **Imports > Template Name** list, check the template check box.

The base template content is displayed in the **Template Content** window. The base template displays the template properties, template variables, and template content. This template can be imported in to another template and the base template content is substituted in the appropriate place of the extending template. When you launch the extended template, the parameter inputs for the base template are also obtained. Also, the substituted content is used for complete CLI command generation.

**Note** The base templates are CLI templates.

**Step 8** Click **OK** to save the template properties, or click the cancel icon at the top-right corner of the window to revert the changes.

**Note** You can edit the template properties by clicking **Template Property**.

**Step 9** Click **Template Content** to edit the template syntax. For information about the structure of the Configuration Template, see the *Template Structure* section.

**Step 10** Click **Validate Template Syntax** to validate the template values.

If an error or a warning message appears, you can check the validation details in **Validation Table** by clicking the error and warnings field.

**Note** You can continue to save the template if there are warnings only. However, if there is an error, you must edit the templates to fix the errors before you proceed. Click the line number under the Start Line column to locate the error in the template content. You will get an error if you validate a template that does not have a template name.

**Step 11** Click **Save** to save the template.

**Step 12** Click **Save and Exit** to save the configuration and go back to the configuring templates screen.

---

## Modifying a Template

You can edit the user-defined templates. However, the predefined templates and templates that are already published cannot be edited.

### Procedure

---

- Step 1** From **Control > Template Library**, select a template.
- Step 2** Click **Modify/View template**.
- Step 3** Edit the template description and tags.  
The edited template content is displayed in a pane on the right.
- Step 4** From the **Imports > Template Name** list, check the template check box.  
The base template content is displayed in the **Template Content** window. You can edit the template content based on your requirement in the **Template Content** window. Click the help icon next to the **Template Content** window for information about editing the content of the template.
- Step 5** Edit the supported platforms for the template.
- Step 6** Click **Validate Template Syntax** to validate the template values.
- Step 7** Click **Save** to save the template.
- Step 8** Click **Save and Exit** to save the configuration and go back to the configuring templates screen.
- 

## Copying a Template

To copy a template from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Template Library**, and select a template.
- Step 2** Click **Save Template As**.
- Step 3** Edit the template name, description, tags, and other parameters.  
The edited template content is displayed in the right-hand pane.
- Step 4** From the **Imports > Template Name** list, check the template check box.  
The base template content is displayed in the **Template Content** window. You can edit the template content that is based on your requirement in the **Template Content** window. Click the help icon next to the **Template Content** window for information about editing the content of the template.
- Step 5** Edit the supported platforms for the template.
- Step 6** Click **Validate Template Syntax** to validate the template values.
- Step 7** Click **Save** to save the template.
- Step 8** Click **Save and Exit** to save the configuration and go back to the configuring templates screen.
-

## Deleting a Template

You can delete the user-defined templates. However, you cannot delete the predefined templates. From Cisco DCNM Release 11.0(1), you can delete multiple templates at once.

To delete a template from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Template Library**.
- Step 2** Use the check box to select a template and click **Remove template** icon.
- The template is deleted without any warning message.
- 

### What to do next

The template is deleted from the list of templates on the DCNM Web UI. When you restart the DCNM services, the deleted templates are displayed on the **Control > Template Library** page.

To delete the template permanently, delete the template that is located in your local directory: `Cisco Systems\dcm\dcnm\data\templates\`.

## Importing a Template

To import a template from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Template Library** and click **Import Template**.
- Step 2** Browse and select the template that is saved on your computer.
- You can edit the template parameters, if necessary. For information, see [Modifying a Template, on page 258](#).
- Note** The “\n” in the template is considered as a new line character when imported and edited, but it works fine when imported as a ZIP file.
- Step 3** Click **Validate Template Syntax** to validate the template.
- Step 4** Click **Save** to save the template or **Save and Exit** to save the template and exit.
- 

## Exporting a Template

To export a template from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Template Library**.
- Step 2** Use the check box to select a template and click **Export Template**.  
The browser requests you to open or save the template to your directory.
- 

## Image Management

Upgrading your devices to the latest software version manually might take a long time and prone to error, which requires a separate maintenance window. To ensure rapid and reliable software upgrades, image management automates the steps associated with upgrade planning, scheduling, downloading, and monitoring. Image management is supported only for Cisco Nexus switches.



**Note** Before you upgrade, ensure that the POAP boot mode is disabled for Cisco Nexus 9000 Series switches and Cisco Nexus 3000 Series switches. To disable POAP, run the `no boot poap enable` command on the switch console. You can however, enable it after the upgrade.

---

The **Image Management** menu includes the following submenu:

This feature allows you to upload or delete images that are used during POAP and switch upgrade. To view the window from the Cisco DCNM Web UI homepage, choose .

You can view the following details in the window.

### Deleting an Image

To delete an image from the repository from the Cisco DCNM Web UI, perform the following steps:

#### Procedure

---

- Step 1** Choose .  
The window appears.
- Step 2** Choose an existing image from the list and click the **Delete Image** icon.  
A confirmation window appears.
- Step 3** Click **Yes** to delete the image.
- 

### Image Upload

To upload different types of images to the server from the Cisco DCNM Web UI, perform the following steps:




---

**Note** Devices use these images during POAP or image upgrade.

Your user role should be **network-admin** to upload an image. You can't perform this operation with the **network-stager** user role.

---

### Procedure

---

- Step 1** Choose .
- The window appears.
- Step 2** Click **Image Upload**.
- The **Select File to Upload** dialog box appears.
- Step 3** Click **Choose file** to choose a file from the local repository of your device.
- Step 4** Choose the file and click **Upload**.
- Step 5** Click **OK**.
- The upload takes some time depending on the file size and network bandwidth.

**Note** You can upload images for all Cisco Nexus Series Switches.

---

## Install & Upgrade

The **Install & Upgrade** menu includes the following submenus:

### Upgrade History

This feature enables you to upgrade the Cisco Nexus Platform Switches using In-Service Software Upgrade (ISSU). This upgrade procedure may be disruptive or non-disruptive based on the device configuration. You can select the Kickstart, System, or NX-OS images from image repository or the file system on the device. To select the images from the repository, the same needs to be uploaded from **Control > Image Management > Image upload** tab.

The following table describes the fields that appear on **Control > Image Management > Upgrade History**.

| Field   | Description                                                                                          |
|---------|------------------------------------------------------------------------------------------------------|
| Task Id | Specifies the serial number of the task. The latest task will be listed in the top.                  |
|         | <b>Note</b> If Failover is triggered in Native HA, the Task Id sequence number is incremented by 32. |

| Field          | Description                                                                                                                                                                                                                                                                                                              |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Task Type      | Specifies the type of task. <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> <p><b>Note</b> If the job fails on a single or multiple devices, the status field shows COMPLETED WITH EXCEPTION indicating a failure.</p> |
| Created Time   | Specifies the time when the task was created.                                                                                                                                                                                                                                                                            |
| Scheduled At   | Specifies the time when the task is specified to be executed. You can also choose to schedule a task to be executed at a later time.                                                                                                                                                                                     |
| Completed Time | Specifies the time when the task was completed.                                                                                                                                                                                                                                                                          |
| Comment        | Shows any comments that the Owner has added while performing the task.                                                                                                                                                                                                                                                   |



**Note** After a fresh Cisco DCNM installation, this page will have no entries.

You can perform the following:

## View

To view the image upgrade history from the Cisco DCNM Web UI, perform the following steps:

### Procedure

**Step 1** Choose **Control > Image Management > Install & Upgrade > Upgrade History**, check the task ID check box.

Select only one task at a time.

**Step 2** Click **View**.

The **Installation Task Details** window appears.

**Step 3** Click **Settings**. Expand the **Columns** menu and choose the details you want to view.

You can view the following information in this window:

- Location of the kickstart and system images
- Compatibility check status
- Installation status
- Descriptions
- Logs

**Step 4** Select the device.

The detailed status of the task appears. For the completed tasks, the response from the device appears.

If the upgrade task is in progress, a live log of the installation process appears.

**Note** • This table autorefreshes every 30 secs for jobs in progress, when you're on this window.

## Delete

To delete a task from the Cisco DCNM Web UI, perform the following steps:

### Procedure

**Step 1** Choose **Control > Image Management > Install & Upgrade > Upgrade History**, and check the **Task ID** check box.

**Step 2** Click **Delete**.

**Step 3** Click **OK** to confirm deletion of the job.

## New Installation

To upgrade the devices that are discovered from the Cisco DCNM, perform the following steps:

### Procedure

**Step 1** Choose **Control > Image Management > Install & Upgrade > Upgrade History**.

**Step 2** Choose **New Installation** to install, or upgrade the kickstart and the system images on the devices.

The devices with default VDCs are displayed in the **Select Switches** window.

**Step 3** Select the check box to the left of the switch name.

You can select more than one switch and move the switches to the right column.

**Step 4** Click **Add** or **Remove** icons to include the appropriate switches for upgrade.

The selected switches appear in a column on the right.

**Step 5** Click **Next**.

The **Specify Software Images** window appears. This tab displays the switches that you selected in the previous screen. You can choose the images for upgrade as well.

- The **Auto File Selection** check box enables you to specify an image version, and a path where you can apply the upgraded image to the selected devices.
- **Select File Server** is disabled, and the default server is used.
- In the **Image Version** field, specify the image version as displayed in the **Image Upload** window.
- The **Path** field is disabled, and the default image path is used.

**Step 6** Click **Select Image** in the **Kickstart image** column.

The **Software Image Browser** dialog box appears.

- Note**
- Cisco Nexus 9000 Series Switches require only the system image to load the Cisco NX-OS operating system. Therefore, the option to select kickstart images for these devices is disabled.
  - If there's an issue in viewing the **Software Image Browser** dialog box, reduce the font size of your browser and retry.

**Step 7** Click **Select Image** in the **System Image** column.

The **Software Image Browser** dialog box appears.

**Step 8** On the **Software Image Browser** dialog box, you can choose the image from **File Server** or **Switch File System**.

If you choose **File Server**:

- a) From the **Select the File server** list, choose the Default\_SCP\_Repository file server on which the image is stored.
- b) From the **Select Image** list, choose the appropriate image. Check the check box to use the same image for all other selected devices of the same platform.

Example: For platform types N9K-C93180YC-EX and N9K-C93108TC-EX, logic matches platform (N9K) and three characters (C93) from subplatform. The same logic is used across all platform switches.

**Note** Only files with BIN extension are listed if you select **File Server**. To view other files, choose **Administration > DCNM Server > Server Properties**, set **FILE\_SELECTION\_FILTER** to **false**, and restart the server. It is set to **true** by default.

**Note** Only image files present in the **Image Upload** window can be selected. You can't select images present in any other paths.

- c) Click **OK** to choose the kickstart image or **Cancel** to revert to the **Specify Software Images** window.

If you choose **Switch File System**:

- a) From the **Select Image** list, choose the appropriate image that is located on the flash memory of the device.

**Note** Only files with BIN extension are listed if you select **Switch File System**. To view other files, choose **Administration > DCNM Server > Server Properties**, set **FILE\_SELECTION\_FILTER** to **false**, and restart the server. It is set to **true** by default.

b) Click **OK** to choose the kickstart image or **Cancel** to revert to the **Specify Software Images** dialog box.

**Step 9** The **Vrf** column indicates the name of the virtual routing and forwarding (VRF).

**Step 10** In the **Available Space** column, specify the available space for the **Primary Supervisor** and **Secondary Supervisor** modules of the switch.

**Available Space** column shows the available memory in MB on the switch (for less than 1 MB, it's shown and marked as KB).

Bootflash browser shows the filename, size, and last modified date for all the files and directories on the switch bootflash. You can delete files by selecting them and clicking **Delete** to increase the available space on the switch.

**Step 11** **Selected Files Size** column shows the size of images that are selected from the server.

If the total size of selected images is greater than available space on a switch, the file size is marked in red. We recommend that you create more space on the switch to copy images to it and install.

**Step 12** Drag and drop the switches to reorder the upgrade task sequence.

**Step 13** Select **Skip Version Compatibility** if you are sure that the version of the Cisco NX-OS software on your device is compatible with the upgraded images that you have selected.

**Step 14** Select **Select Parallel Line Card upgrade** to upgrade all the line cards at the same time.

Upgrading a parallel line card isn't applicable for Cisco MDS devices.

**Step 15** Select **Options** under the **Upgrade Options** column to choose the type of upgrade.

**Upgrade Options** window appears with two upgrade options. The drop-down list for **Upgrade Option 1** has the following options:

- NA
- bios-force
- non-disruptive

NA is the default value.

The drop-down list for **Upgrade Option 2** has the following options:

- NA
- bios-force

When **NA** is selected under **Upgrade Option 1**, **Upgrade Option 2** is disabled.

When **bios-force** is selected under **Upgrade Option 1**, **Upgrade Option 2** is disabled.

When **non-disruptive** is selected under **Upgrade Option 1**, you can choose **NA** or **bios-force** under **Upgrade Option 2**.

Check the **Use this Option for all other selected devices** check box to use the selected option for all the selected devices and click **OK**.

- Note**
- The upgrade options are applicable only for Cisco Nexus 3000 Series and 9000 Series switches.
  - Selecting the non-disruptive option for upgrading does not ensure a non-disruptive upgrade. Perform a compatibility check to ensure that the device supports non-disruptive upgrade.

**Step 16** Click **Next**.

If you didn't select **Skip Version Compatibility**, the Cisco DCNM performs a compatibility check.

You can choose to wait until the check is complete or click **Finish Installation Later**.

The installation wizard is closed and a compatibility task is created in **Control > Image Management > Install & Upgrade > Upgrade History** tasks.

The time that is taken to check the image compatibility depends on the configuration and the load on the device.

The **Version Compatibility Verification** status column displays the status of verification.

If you skip the version compatibility check by choosing **Skip Version Compatibility**, Cisco DCNM displays only the name of the device. The **Current Action** column displays **Completed**, and the **Version Compatibility Verification** column displays **Skipped**.

**Step 17** Click **Finish Installation Later** to perform the upgrade later.

**Step 18** Click **Next**.

**Step 19** Check the check box to save the running configuration to the startup configuration before upgrading the device.

**Step 20** You can schedule the upgrade process to occur immediately or later.

- a. Select **Deploy Now** to upgrade the device immediately.
- b. Select **Choose time to Deploy** and specify the time in MMM/DD/YYYY HH:MM:SS format to perform the upgrade later.

This value is relative to the server time. If the selected time to deploy is in the past, the job is executed immediately.

**Step 21** You can choose the execution mode based on the devices and the line cards you have chosen to upgrade.

- a. Select **Sequential** to upgrade the devices in the order you chose them.
- b. Select **Concurrent** to upgrade all the devices at the same time.

**Step 22** Click **Finish** to begin the upgrade process.

The Installation wizard closes and a task to upgrade is created on the **Control > Image Management > Install & Upgrade > Upgrade History** page.

---

### What to do next

After you complete the ISSU on the switch, ensure that you wait for 20 minutes to allow the switch to reboot, and stabilize the SNMP agent. DCNM discovers polling cycles in order to display the new version of the switch on the Cisco DCNM Web UI.

## Finish Installation

You can choose to complete the installation for tasks which was completed on the **Compatibility Check** page. Perform the following task to complete the upgrade process on the devices.

### Procedure

- 
- Step 1** Choose **Control > Image Management > Install & Upgrade > Upgrade History**, select a task for which the compatibility check is complete.
- Select only one task at a time.
- Step 2** Click **Finish Installation**.
- Software Installation Wizard** appears.
- Step 3** Check the check box to save the running configuration to the startup configuration before upgrading the device.
- Step 4** Check the check box to put a device in maintenance mode before upgrade. This option is valid only for the devices that support maintenance mode.
- Step 5** You can schedule the upgrade process to occur immediately or later.
- Select **Deploy Now** to upgrade the device immediately.
  - 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.
- Select **Sequential** to upgrade the devices in the order in which they were chosen.
  - Select **Concurrent** to upgrade the devices at the same time.
- Step 7** Click **Finish** to complete the upgrade process.
- 

## Switch Level History

You can view the history of the upgrade process at a switch level. You can view the current version of the switch and other details.

The following table describes the fields that appear on **Control > Image Management > Install & Upgrade > Switch Level History**.

| Field           | Description                                          |
|-----------------|------------------------------------------------------|
| Switch Name     | Specifies the name of the switch                     |
| IP Address      | Specifies the IP Address of the switch               |
| Platform        | Specifies the Cisco Nexus switch platform            |
| Current Version | Specifies the current version on the switch software |

Click the radio button next to a switch name to select the switch and view its upgrade history. Click **View** to view the upgrade task history for the selected switch.

The following table describes the fields that appear on **Control > Image Management > Install & Upgrade > Switch Level History > View Device Upgrade Tasks**:

| Field              | Description                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| Owner              | Specifies the owner who initiated the upgrade.                                                                                           |
| Job Status         | Specifies the status of the job. <ul style="list-style-type: none"> <li>• Planned</li> <li>• In Progress</li> <li>• Completed</li> </ul> |
| KickStart Image    | Specifies the kickStart image that is used to upgrade the Switch.                                                                        |
| System Image       | Specifies the system image that is used to upgrade the switch.                                                                           |
| Completed Time     | Specifies the date and time at which the upgrade was successfully completed.                                                             |
| Status Description | Specifies the installation log information of the job.                                                                                   |

## Endpoint Locator

The Endpoint Locator (EPL) feature allows real-time tracking of endpoints within a data center. The tracking includes tracing the network life history of an endpoint and getting insights into the trends that are associated with endpoint additions, removals, moves, and so on.

Information about the Endpoint Locator is displayed on a single landing page or dashboard . The dashboard displays an almost real-time view of data (refreshed every 30 seconds) pertaining to all the active endpoints on a single pane. The data that is displayed on this landing page depends on the scope selected by you from the **SCOPE** drop-down list.

## Endpoint Locator

The Endpoint Locator (EPL) feature allows real-time tracking of endpoints within a data center. The tracking includes tracing the network life history of an endpoint and getting insights into the trends that are associated with endpoint additions, removals, moves, and so on. An endpoint is anything with at least one IP address (IPv4 and/or IPv6) and MAC address. An endpoint can be a virtual machine (VM), container, bare-metal server, service appliance and so on.

**Important**

- EPL is supported for VXLAN BGP EVPN fabric deployments only in the DCNM LAN fabric installation mode. The VXLAN BGP EVPN fabric can be deployed as Easy fabric, Easy eBGP fabric, or an External fabric (managed or monitored mode). EPL is not supported for 3-tier access-aggregation-core based network deployments.
- EPL displays endpoints that have at least one IP address (IPv4 and/or IPv6). Also, these endpoints must be residing in networks where the gateway or SVI is configured on the network switches within the VXLAN EVPN fabric. In other words, EPL cannot determine the identity (IPv4/IPv6 address) of the endpoints for networks that are deployed as Layer-2 Only within the fabric.

EPL relies on BGP updates to track endpoint information. Hence, typically the DCNM needs to peer with the BGP Route-Reflector (RR) to get these updates. For this purpose, IP reachability from the DCNM to the RR is required. This can be achieved over in-band network connection to the DCNM eth2 interface.

Some key highlights of the Endpoint Locator are:

- Support for dual-homed and dual-stacked (IPv4 + IPv6) endpoints
- Support for up to two BGP Route Reflectors or Route Servers
- Support real-time and historical search for all endpoints across various search filters such as VRF, Network, Layer-2 VNI, Layer-3 VNI, Switch, IP, MAC, port, VLAN, and so on.
- Support for real-time and historical dashboards for insights such as endpoint lifetime, network, endpoint, VRF daily views, and operational heat map.
- Support for iBGP and eBGP based VXLAN EVPN fabrics. From Release 11.2(1), the fabrics may be created as Easy Fabrics or External Fabrics. EPL can be enabled with an option to automatically configure the spine or RRs with the appropriate BGP configuration (new in DCNM 11.2).
- Support for high availability
- Support for endpoint data that is stored for up to 180 days, amounting to a maximum of 5 G storage space.
- Supported scale: 10K endpoints

For more information about EPL, refer to the following sections:

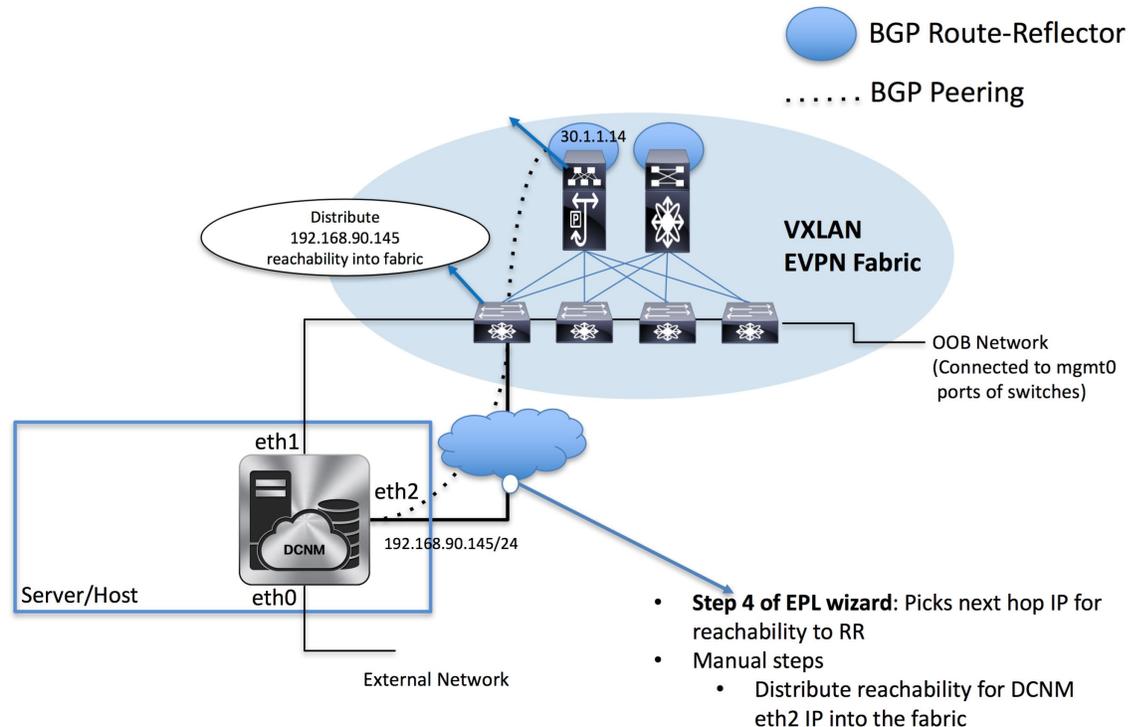
## Configuring Endpoint Locator

The DCNM OVA or the ISO installation comes with three interfaces:

- eth0 interface for external access
- eth1 interface for fabric management (Out-of-band or OOB)
- eth2 interface for in-band network connectivity

# Configuration

The Server Hosting DCNM has IP connectivity to BGP RR(s)



The eth1 interface provides reachability to the devices via the mgmt0 interface either Layer-2 or Layer-3 adjacent. This allows DCNM to manage and monitor these devices including POAP. EPL requires BGP peering between the DCNM and the Route-Reflector. Since the BGP process on Nexus devices typically runs on the default VRF, in-band IP connectivity from the DCNM to the fabric is required. For this purpose, the eth2 interface can be configured using the **appmgr setup inband** command. Optionally, you can configure the eth2 interface during the Cisco DCNM installation.

If you need to modify the already configured in-band network (eth2 interface), execute the **ifconfig eth2 0.0.0.0** command and run the **appmgr setup inband** command again. Refer [Editing Network Properties Post DCNM Installation](#) to run the **appmgr setup inband** command.



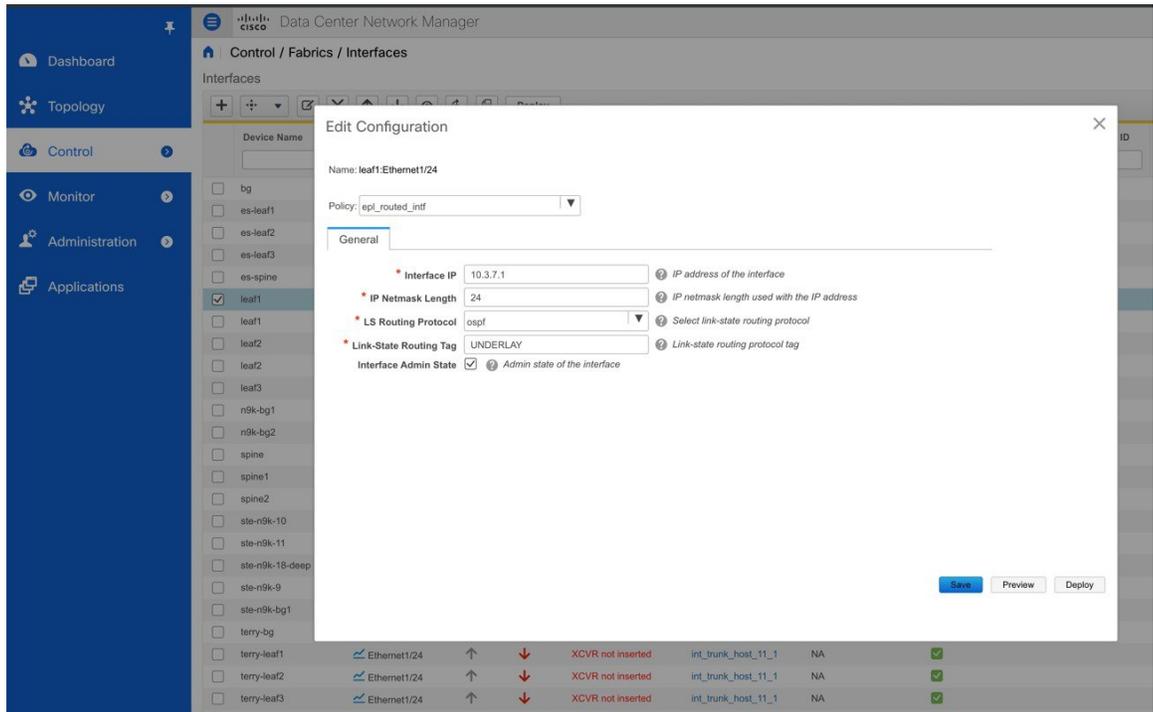
**Note** The setup of eth2 interface on the DCNM is a prerequisite of any application that requires the in-band connectivity to the devices within fabric. This includes EPL and Network Insights Resources (NIR).



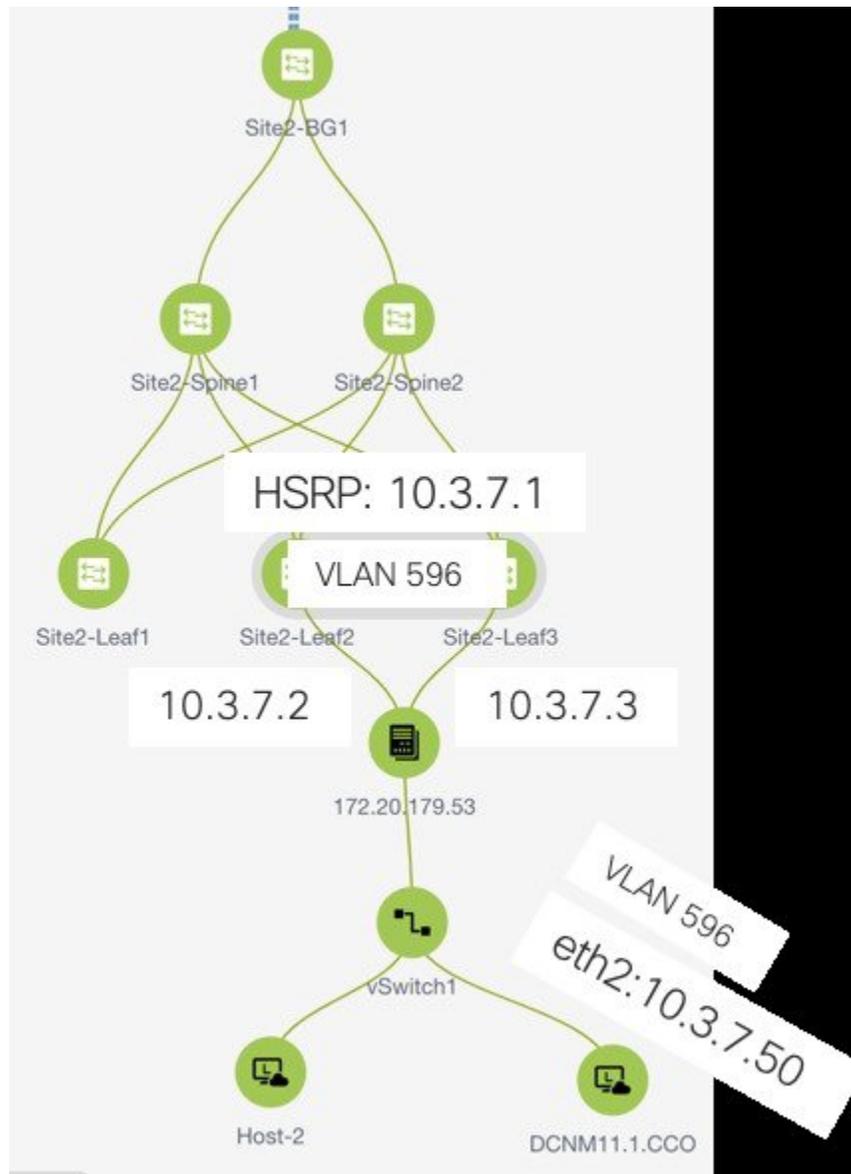
**Note** For configuring EPL in standalone mode, you must add a single neighbor to EPL. DCNM eth2 IP address is EPL IP.

On the fabric side, for a standalone DCNM deployment, if the DCNM eth2 port is directly connected to one of the front-end interfaces on a leaf, then that interface can be configured using the **epl\_routed\_intf** template.

An example scenario of how this can be done when IS-IS or OSPF is employed as the IGP in the fabric, is depicted below:



However, for redundancy purposes, it is always advisable to have the server on which the DCNM is installed to be dual-homed or dual-attached. With the OVA DCNM deployment, the server can be connected to the switches via a port-channel. This provides link-level redundancy. To also have node-level redundancy on the network side, the server may be attached to a vPC pair of Leaf switches. In this scenario, the switches must be configured such that the HSRP VIP serves as the default gateway of the eth2 interface on the DCNM. The following image depicts an example scenario configuration:



In this example, the server with the DCNM VM is dual-attached to a vPC pair of switches that are named Site2-Leaf2 and Site2-Leaf3 respectively. VLAN 596 associated with the IP subnet 10.3.7.0/24 is employed for in-band connectivity. You can configure the vPC host port toward the server using the **interface vpc trunk host** policy as shown in the following image:

For the HSRP configuration on Site2-Leaf2, the **switch\_freemom** policy may be employed as shown in the following image:

You can deploy a similar configuration on Site2-Leaf3 while using IP address 10.3.7.2/24 for SVI 596. This establishes an in-band connectivity from the DCNM to the fabrics over the eth2 interface with the default gateway set to 10.3.7.1.

After you establish the in-band connectivity between the physical or virtual DCNM and the fabric, you can establish BGP peering. There is a simple wizard for enabling Endpoint Locator.

During the EPL configuration using the wizard, the route reflectors (RRs) are configured to accept DCNM as a BGP peer. During the same configuration, the DCNM is also configured by adding routes to the BGP loopback IP on the spines/RRs via the eth2 gateway.



**Note** Cisco DCNM queries the BGP RR to glean information for establishment of the peering, like ASN, RR, IP, and so on.

Endpoint Locator

Endpoint Locator enables real-time tracking of current and past location information about network endpoints.

Please select a fabric to configure endpoint locator feature

Endpoint Locator

Endpoint Locator enables real-time tracking of current and past location information about network endpoints.

Please select a fabric to configure endpoint locator feature

- Data Center
- MSD1
  - epi-ex-site
  - lerry-k2
  - Default\_LAN

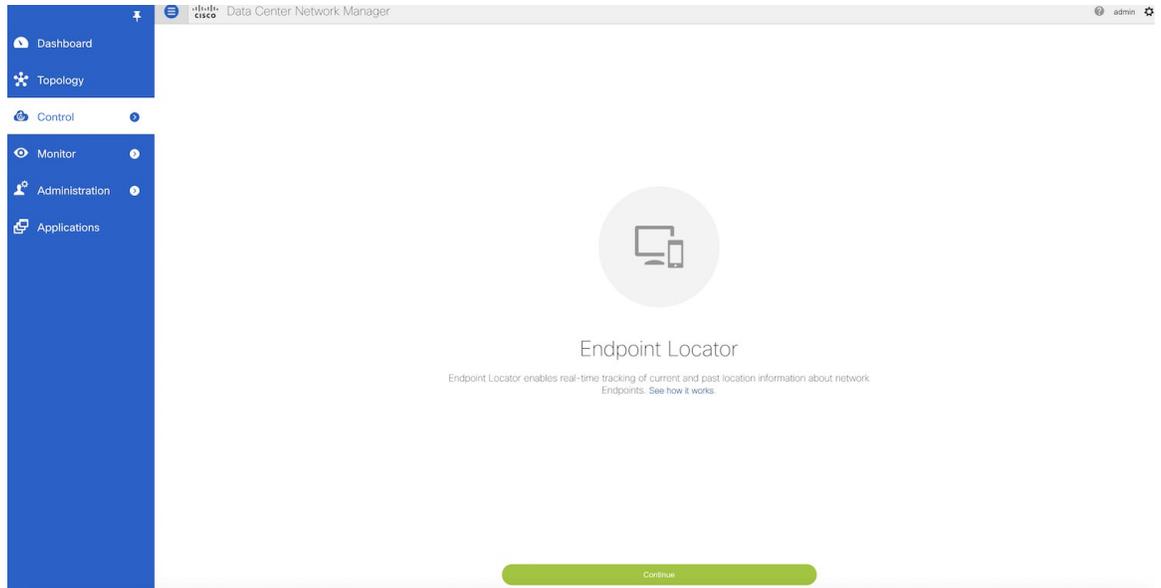
For more information about the EPL dashboard, refer [Monitoring Endpoint Locator](#).

To configure Endpoint Locator from Cisco Web UI, perform the following steps:

## Procedure

**Step 1** Choose **Control > Endpoint Locator > Configure**.

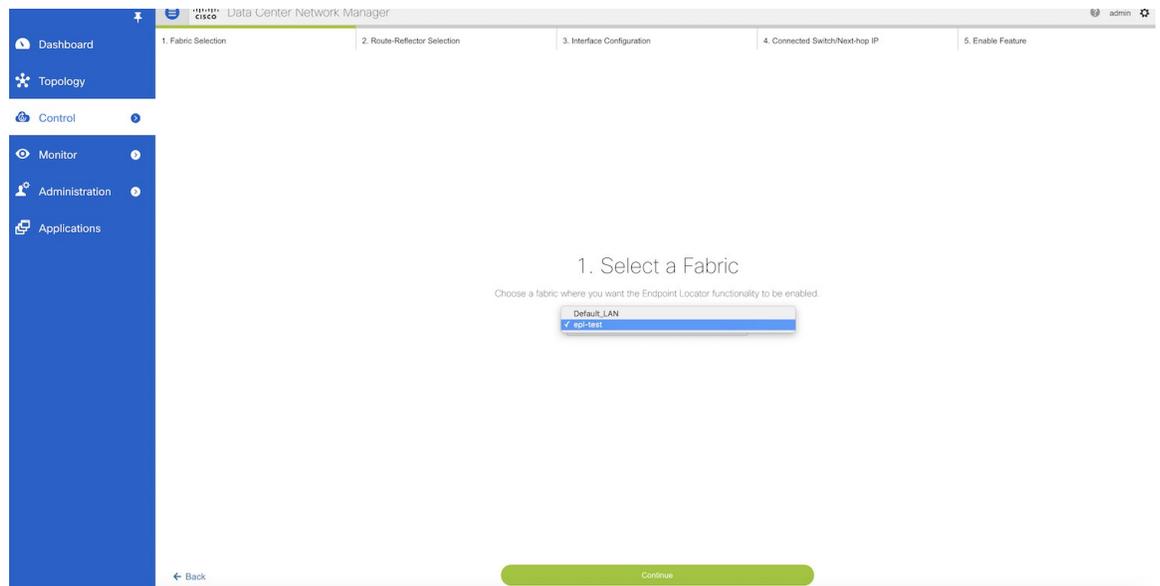
The **Endpoint Locator** window appears, with a **See how it works** help link.



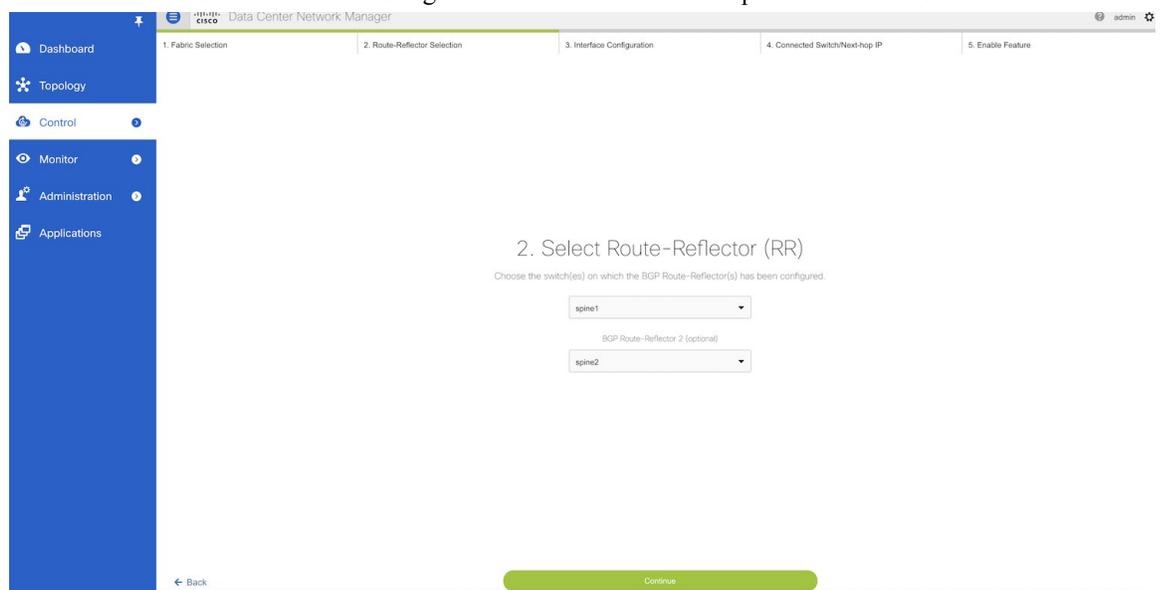
**Step 2** Click **Continue**.

**Step 3** Select the appropriate fabric on which the endpoint locator feature should be enabled to track endpoint activity.

You can enable EPL for one fabric. It can be DFA or EVPN.



**Step 4** Select the switches on the fabric hosting the RRs. Cisco DCNM will peer with the RRs.



**Step 5** Check the Cisco DCNM eth2 configuration for IP reachability to the RR.

3. Verify DCNM In-band Interface

Choose the Ethernet interface on the DCNM that will provide reachability to the BGP Route-Reflector(s) within the fabric.

eth2

Interface IP

192.168.94.124 / 24

← Back Continue

**Step 6** By default, the “Configure my fabric” option is selected. This knob controls whether BGP configuration will be pushed to the selected spines/RRs as part of the enablement of the EPL feature. If the spine/RR needs to be configured manually with a custom policy for the EPL BGP neighborship, then this option should be unchecked. Check Next-hop IP and ensure the eth2 gateway IP is correct. If there is an error go to command line and reconfigure the eth2 port using the **appmgr setup inband** command. To flush the prior eth2 configuration, perform **ifconfig eth2 0.0.0.0** before proceeding with the **appmgr setup inband** command.

4. Next-hop IP

Provide the next-hop IP that provides reachability to the BGP Route-Reflector (RR).

Configure my fabric

Next-hop IP

192.168.94.1

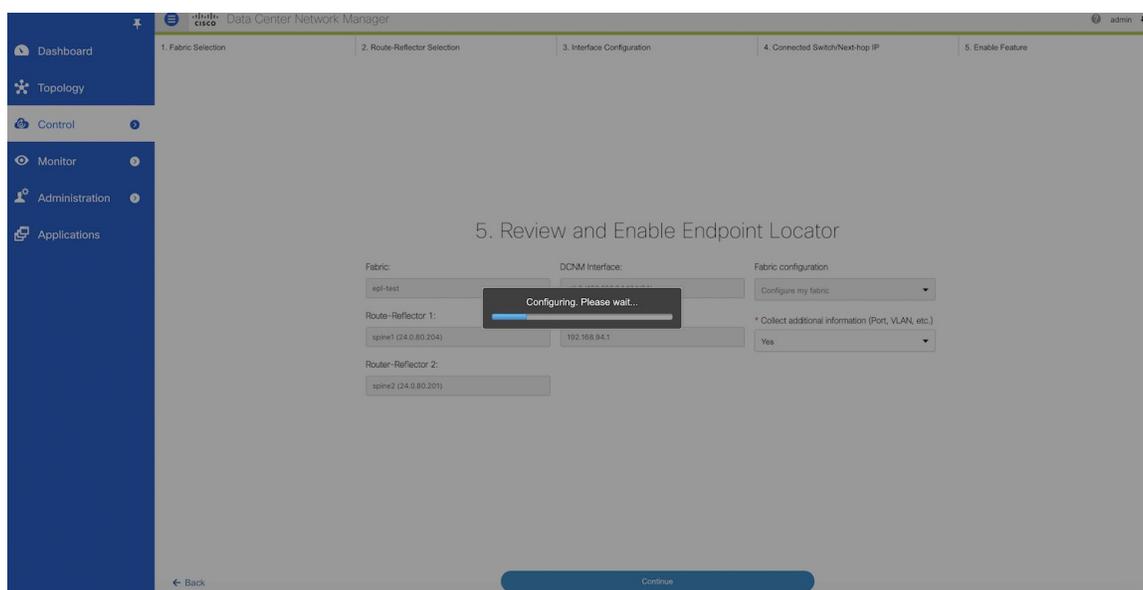
← Back Continue

**Step 7** The last step provides a summary of the information entered in the previous steps. The wizard view allows navigation to any particular step whereby one can make necessary changes/edits. In this step, one must specify whether additional information such as PORT, VLAN, VRF etc. is required when enabling the EPL feature. If the **No** option is selected, then this information will not be collected and reported by EPL.

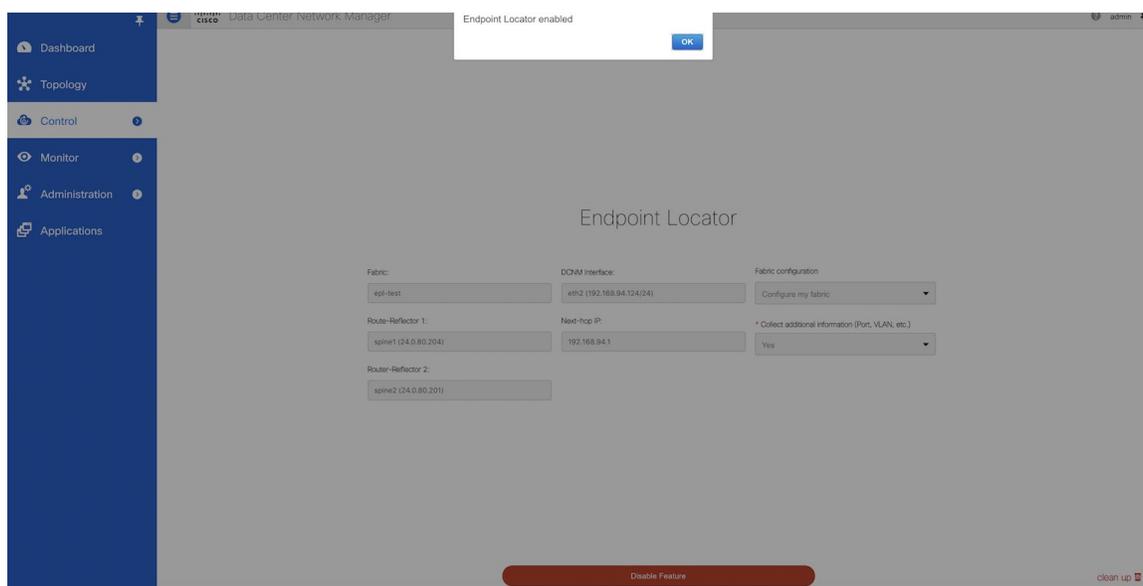
However, if the **Yes** option is selected in the drop-down, a warning pop-up appears that feature NX-API must be supported and enabled on the switches, ToRs, and leafs to gather this information. Otherwise, you cannot fetch or report this additional information.

### Step 8

Once the appropriate selections are made and various inputs have been reviewed, click **Continue** to enable EPL. A progress bar will appear indicating the status of the EPL feature enable process.



If there are any errors while you enable EPL, the enable process aborts and the appropriate error message is displayed. Otherwise, EPL is successfully enabled and on clicking **OK**, the screen is automatically redirected to the EPL dashboard.



When the Endpoint Locator feature is enabled, there are a number of steps that occur in the background. DCNM contacts the selected RRs and determines the ASN. It also determines the interface IP that is bound to the BGP process. Also, appropriate BGP neighbor statements are added on the RRs or spines in case of eBGP underlay, to get them ready to accept the BGP connection that will be initiated from the DCNM. The neighbor address is the same as that of the eth2 interface shown in step 2. For the native HA DCNM deployment, both the primary and secondary DCNM eth2 interface IPs will be added as BGP neighbors but only one of them will be active at any given time. Once EPL is successfully enabled, the user is automatically redirected

to the EPL dashboard that depicts operational and exploratory insights into the endpoints that are present in the fabric.

### What to do next

For more information about monitoring EPL, see [Monitoring Endpoint Locator, on page 335](#).

## Flushing the Endpoint Database

After you enable the Endpoint Locator feature, you can clean up or flush all the Endpoint information. This allows starting from a clean-slate with respect to ensuring no stale information about any endpoint is present in the database. After the database is clean, the BGP client re-populates all the endpoint information learnt from the BGP RR.

To flush all the Endpoint Locator information from the Cisco DCNM Web UI, perform the following steps:

### Procedure

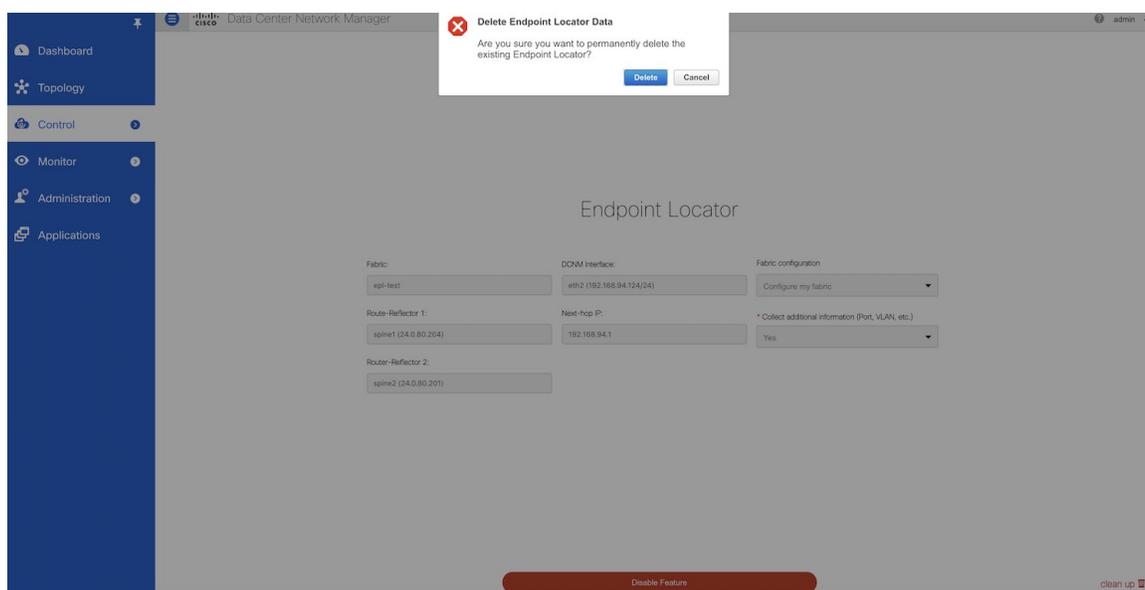
**Step 1** Choose **Control > Endpoint Locator > Configure**, and click **clean up** link.

The screenshot shows the Cisco DCNM Web UI for the Endpoint Locator configuration. The sidebar on the left is blue and contains the following menu items: Dashboard, Topology, Control (highlighted), Monitor, Administration, and Applications. The main content area is titled "Endpoint Locator" and contains the following configuration fields:

- Fabric: epi-test
- DCNM interface: eth2 (192.168.94.124/24)
- Fabric configuration: Configure my fabric
- Route-Reflector 1: spine1 (24.0.80.204)
- Next-hop IP: 192.168.94.1
- \* Collect additional information (Port, VLAN, etc.): Yes
- Route-Reflector 2: spine2 (24.0.80.201)

At the bottom center, there is a red button labeled "Disable Feature". At the bottom right, there is a "clean up" link.

A warning is displayed with a message indicating that all the endpoint information that is stored in the database will be flushed.



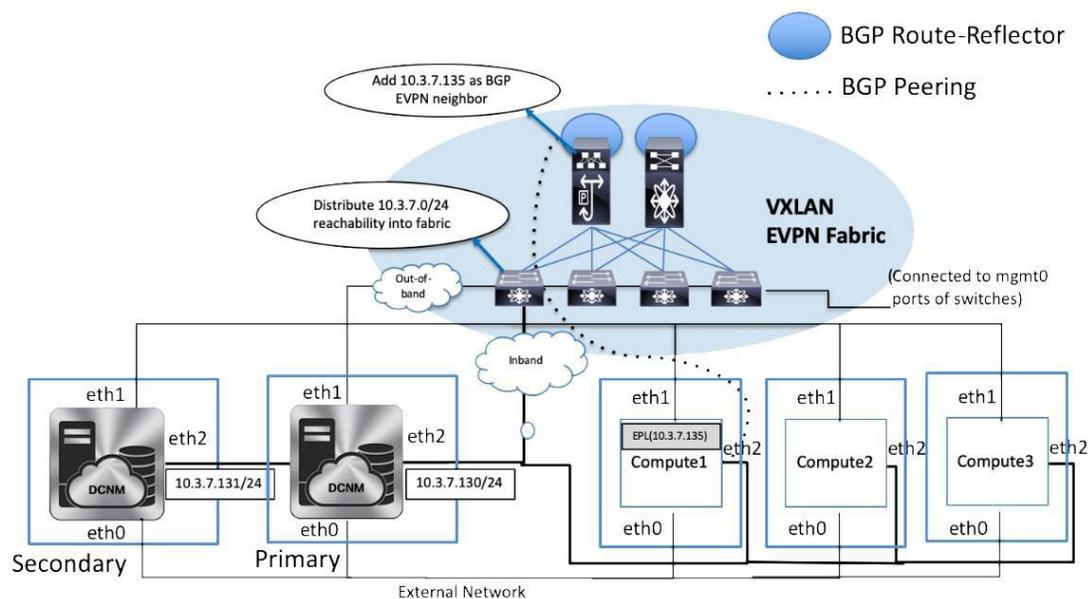
**Step 2** Click **Delete** to continue or **Cancel** to abort.

## Configuring Endpoint Locator in DCNM Cluster Mode



**Note** For configuring EPL in cluster mode, you must add a single neighbor to EPL. DCNM EPL container Inband IP address is EPL IP.

With the DCNM cluster mode deployment, in addition to the DCNM nodes, an additional 3 compute nodes are present in the deployment. For information about deploying applications in cluster mode, see *Cisco DCNM in Clustered Mode*.



In DCNM Cluster mode, all applications including EPL run on the compute nodes. The DCNM application framework takes care of the complete life cycle management of all applications that run on the compute nodes. The EPL instance runs as a container that has its own IP address allocated out of the inband pool assigned to the compute nodes. This IP address will be in the same IP subnet as the one allocated to the eth2 or inband interface. Using this IP address, the EPL instance forms a BGP peering with the spines/RRs when the EPL feature is enabled. If a compute node hosting the EPL instance will go down, the EPL instance will be automatically respawned on one of the remaining 2 compute nodes. All IP addresses and other properties associated with the EPL instance are retained.

The Layer-2 adjacency requirement of the compute nodes dictates that the compute node eth2 interfaces should be part of the same IP subnet as the DCNM nodes. Again, in this case, connecting the compute nodes to the same vPC pair of switches is the recommended deployment option. Note that for cluster mode DCNM OVA setups, ensure that promiscuous mode is enabled in the port group corresponding to eth2 interface in order to establish inband connectivity as depicted below:

## EPL-Inband - Edit Settings

| Properties           |                     |                                              |        |
|----------------------|---------------------|----------------------------------------------|--------|
| Security             | Promiscuous mode    | <input checked="" type="checkbox"/> Override | Accept |
| Traffic shaping      | MAC address changes | <input checked="" type="checkbox"/> Override | Accept |
| Teaming and failover | Forged transmits    | <input checked="" type="checkbox"/> Override | Accept |

CANCEL

OK

The enablement of the EPL feature for DCNM cluster mode is identical to that in the non-cluster mode. The main difference is that on the spine/RRs, only a single BGP neighborship is required that points to the IP address allocated to the EPL instance. Recall that for the DCNM native HA deployment in the non-cluster mode, all spines/RRs always had 2 configured BGP neighbors, one pointing to the DCNM primary eth2 interface and other one pointing to the DCNM secondary eth2 interface. However, only one neighbor would be active at any given time.

## Configuring Endpoint Locator for External Fabrics

In addition to Easy fabrics, DCNM Release 11.2(1) allows you to enable EPL for VXLAN EVPN fabrics comprising of switches that are imported into the external fabric. The external fabric can be in managed mode or monitored mode, based on the selection of **Fabric Monitor Mode** flag in the **External Fabric Settings**. In case the monitor or read-only fabric option is selected for the fabric, while enabling EPL, the **Configure my fabric** option must be unchecked; because, the EPL neighborship is added to the spines or RRs via some other means.

## Configuring Endpoint Locator for eBGP EVPN Fabrics

From Cisco DCNM Release 11.2(1), you can enable EPL for VXLAN EVPN fabrics, where eBGP is employed as the underlay routing protocol. Note that with an eBGP EVPN fabric deployment, there is no traditional RR similar to iBGP. The reachability of the in-band subnet must be advertised to the spines that behave as Route Servers. To configure EPL for eBGP EVPN fabrics from the Cisco DCNM Web UI, perform the following steps:

### Procedure

**Step 1** Choose **Control > Fabric Builder**.

Select the fabric to configure eBGP on or create eBGP fabric with the **Easy\_Fabric\_eBGP** template.

Add Fabric



\* Fabric Name :

\* Fabric Template :

General | EVPN | vPC | Advanced | Manageability | Bootstrap | Configuration Backup

\* BGP ASN for Spines  ⓘ 1-4294967295 | 1-65535[0-65535]

\* BGP AS Mode  ⓘ Multi-AS: Unique ASN per Leaf/Border  
Dual-AS: One ASN for all Leafs/Borders

\* Routing Loopback Id  ⓘ 0-512

\* Underlay Subnet IP Mask  ⓘ Mask for Underlay Subnet IP Range

Manual Underlay IP Address Allocation  ⓘ Checking this will disable Dynamic Underlay IP Address Allocations

\* Underlay Routing Loopback IP Range  ⓘ Typically Loopback0 IP Address Range

\* Underlay Subnet IP Range  ⓘ Address range to assign Numbered and Peer Link SVI IPs

\* Subinterface Dot1q Range  ⓘ Per Border Dot1q Range For VRF Lite Connectivity (Min:2, Max:511)

NX-OS Software Image Version  ⓘ If Set, Image Version Check Enforced On All Switches. Images Can Be Uploaded From Control:Image Upload

**Step 2** Use the **leaf\_bgp\_asn** policy to configure unique ASNs on all leaves.

View/Edit Policies for leaf1 ( FDO23070AC0 )

Add Policy ✕

\* Priority (1-1000):

\* Policy:

General

\* Leaf BGP AS #  ? Leaf BGP Autonomous System number

Variables:

- Step 3** Add the **ebgp\_overlay\_leaf\_all\_neighbor** policy to each leaf.  
 Fill **Spine IP List** with the spines' BGP interface IP addresses, typically the loopback0 IP addresses.  
 Fill **BGP Update-Source Interface** with the leaf's BGP interface, typically loopback0.

View/Edit Policies for leaf1 ( FDO23070AC0 )

Add Policy ✕

\* Priority (1-1000):

\* Policy:

General

\* Spine IP List  ? list of spine IP address for peering list e.g. 10.2.

\* BGP Update-Source Interface  ? Source of BGP session and updates

Enable Tenant Routed Multicast  ? For Overlay Multicast Support In VXLAN Fabrics

Enable BGP Authentication  ? BGP Authentication needs to match the fabric setting

Variables:

- Step 4** Add the **ebgp\_overlay\_spine\_all\_neighbor** policy to each spine.  
 Fill **Leaf IP List** with the leaves' BGP interface IPs, typically the loopback0 IPs.

Fill **Leaf BGP ASN** with the leaves' ASNs in the same order as in **Leaf IP List**.

Fill **BGP Update-Source Interface** with the spine's BGP interface, typically loopback0.

View/Edit Policies for spine ( FDO231003AG )

Add Policy ✕

\* Priority (1-1000):

\* Policy:  ▼

General

\* Leaf IP List  ? list of leaf IP address for peering list e.g. 10.2.0.

\* Leaf BGP ASN  ? BGP ASN of each leaf, separated by ,

\* BGP Update-Source Interface  ? Source of BGP session and updates

Enable Tenant Routed Multicast  ? Tenant Routed Multicast setting needs to match the fabric setting

Enable BGP Authentication  ? BGP Authentication needs to match the fabric setting

Variables:

After the in-band connectivity is established, the enablement of the EPL feature remains identical to what is listed so far. EPL becomes a iBGP neighbor to the Route Servers running on the spines.

## EPL Connectivity Options

Sample topologies for the various EPL connectivity options are as given below.

Cisco DCNM supports the following web browsers:

### DCNM Cluster Mode: Physical Server to VM Mapping

We recommend a minimum of 3 physical servers, or a maximum of 5 physical servers in which each DCNM and compute is located on an individual physical server.

Figure 2: A minimum of 3 physical servers

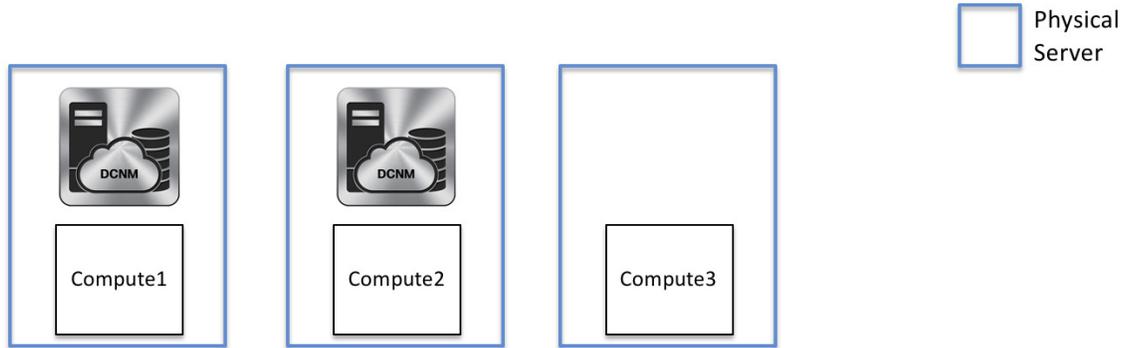
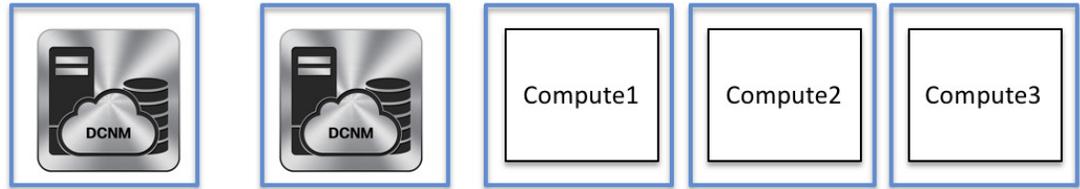
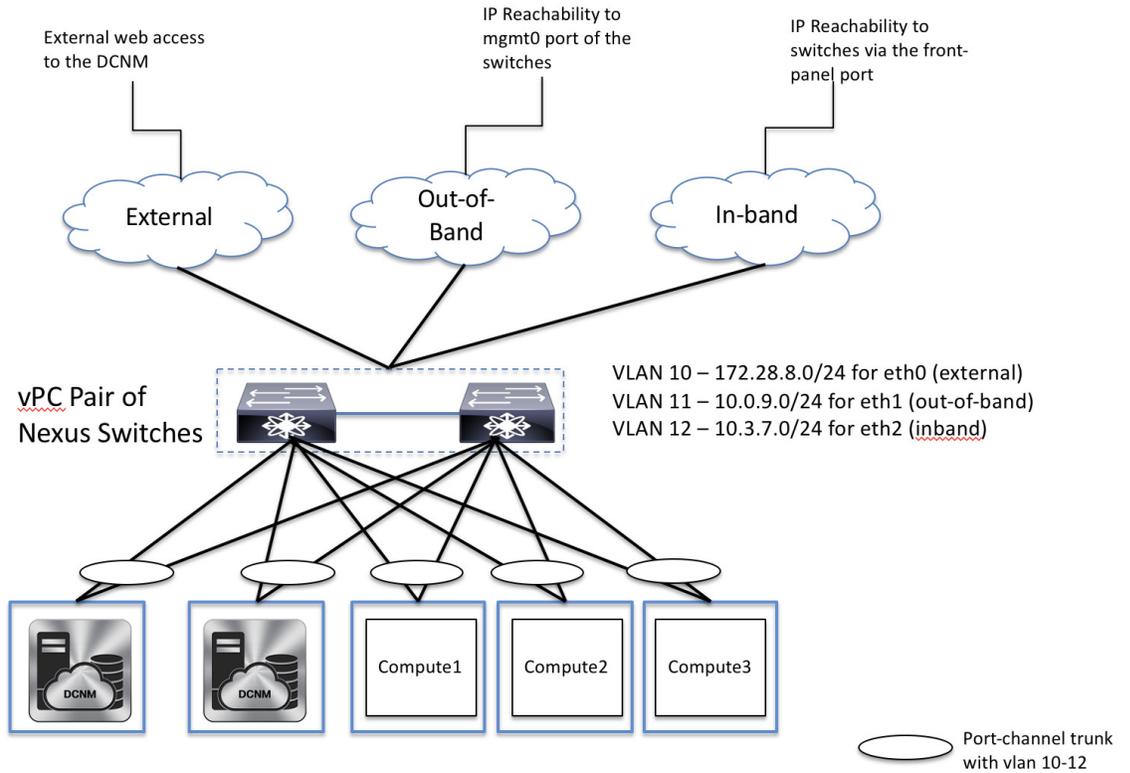


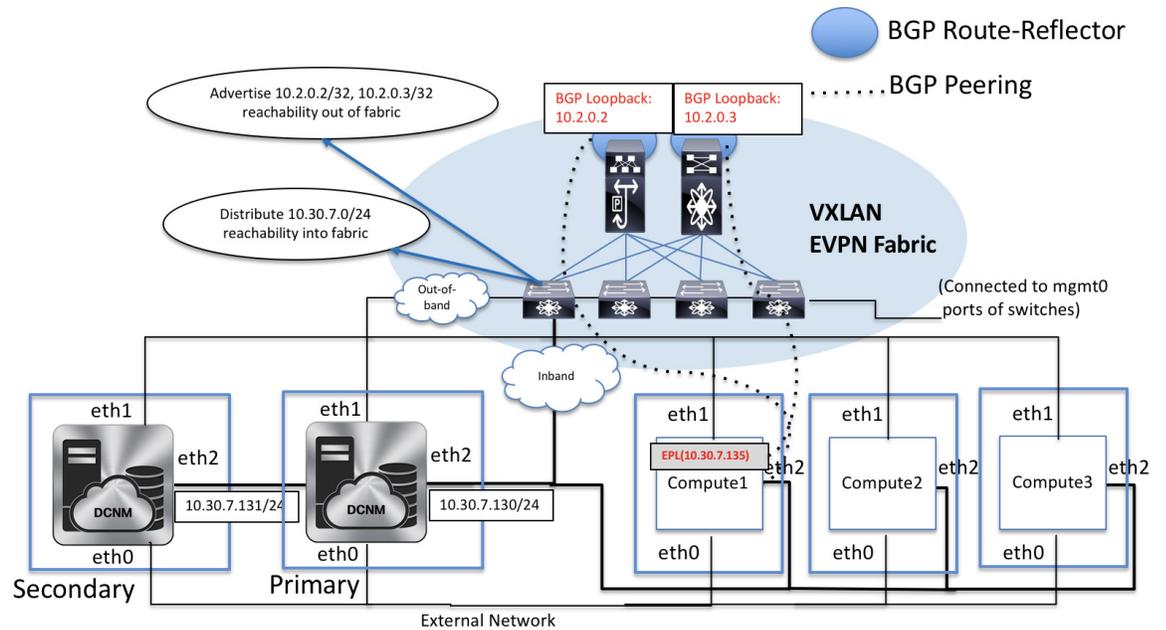
Figure 3: A maximum of 5 physical servers



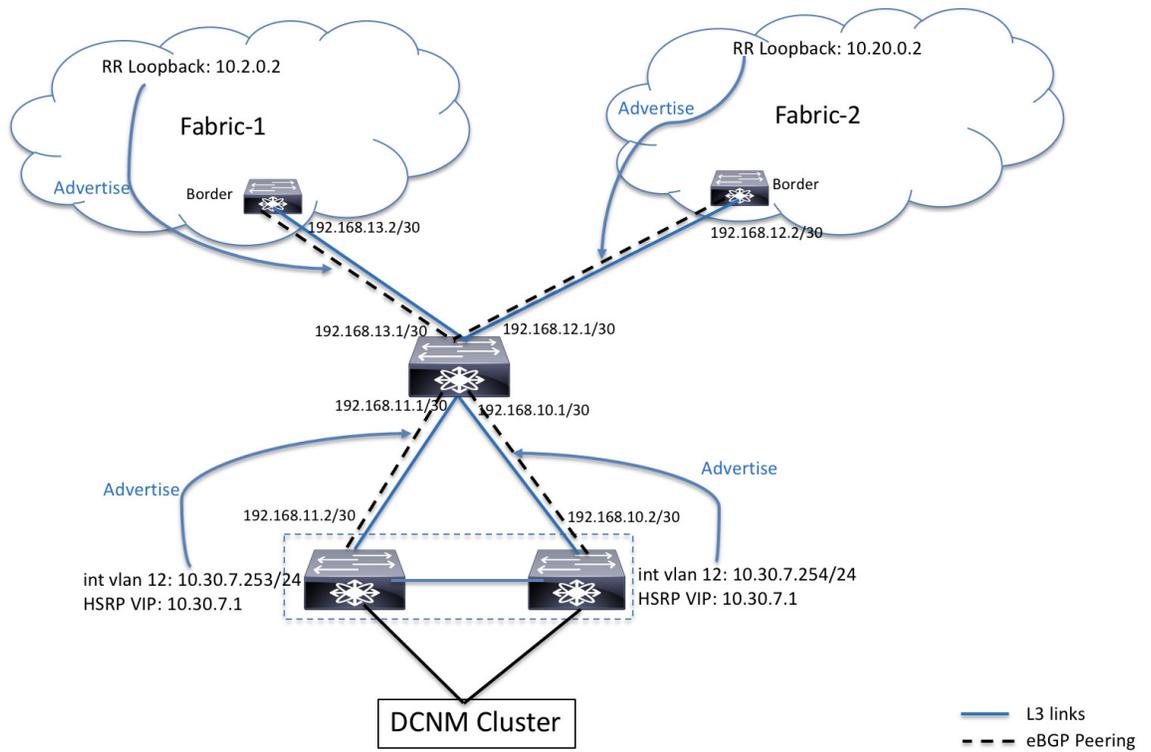
### DCNM/Compute VM Physical Connectivity



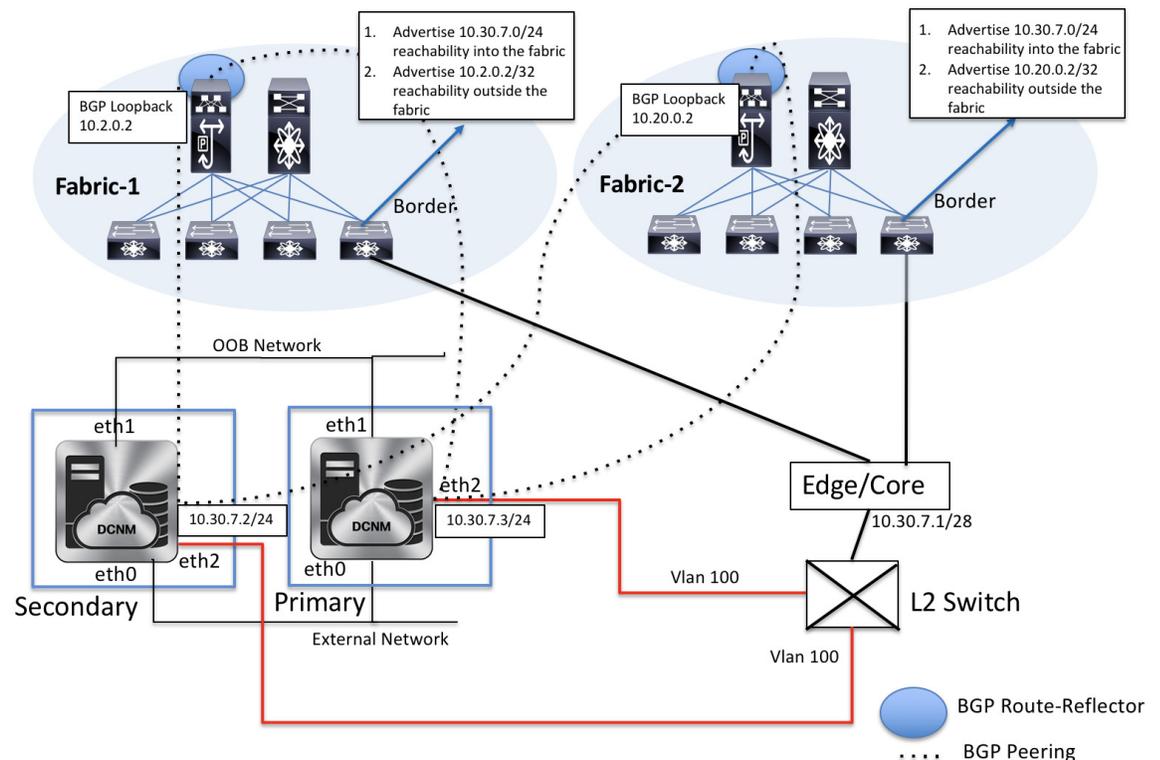
### DCNM Cluster Mode



### DCNM Multi-Fabric Connectivity



## EPL Connectivity for Native HA



## Disabling Endpoint Locator

To disable endpoint locator from the Cisco DCNM Web UI, perform the following steps:

### Procedure

- Step 1** Choose **Control > Endpoint Locator > Configure**.  
The **Endpoint Locator** window appears and the fabric configuration details are displayed.
- Step 2** Click **Disable Feature**.

## Troubleshooting Endpoint Locator

There may be multiple reasons why enabling the Endpoint Locator feature may fail. Typically, if the appropriate devices are selected and the IP addresses to be used are correctly specified, the connectivity of the DCNM to the BGP RR may not be present due to which the feature cannot be enabled. This is a sanity check that is present to ensure that basic IP connectivity is available. The following image shows an example error scenario that was encountered during an attempt to enable the EPL feature.

The logs for EPL are located at the following location: `/usr/local/cisco/dcm/fm/logs`. The log that provides further details on what all occurred when the EPL feature is enabled or disabled, are present in the file `epl.log`.

The following example provides a snapshot of the log that provides the user further information on when EPL enablement failed.

The following example helps you to understand why Endpoint Locator is unable to Connect to a switch.

```
#tail -f epl.log
2017.04.08 07:47:05 INFO [epl] Running script: [/sbin/appmgr, status, epls]
2017.04.08 07:47:05 INFO [epl] Received response:
2017.04.08 07:47:05 INFO [epl]
2017.04.08 07:47:05 INFO [epl] >>> Sat Apr 8 07:47:05 PDT 2017
2017.04.08 07:47:05 INFO [epl] appmgr status epls
2017.04.08 07:47:05 INFO [epl]
2017.04.08 07:47:05 INFO [epl] EPLS is stopped...
2017.04.08 07:47:08 INFO [epl] Running command: ifdown eth2
2017.04.08 07:47:08 INFO [epl] Received response:
2017.04.08 07:47:08 INFO [epl] EPL disabled successfully
2017.04.08 08:00:06 INFO [epl] Enable End Point Locator
2017.04.08 08:00:13 ERROR [epl] Failed to connect to switch 192.169.6.2:java.lang.Exception:
  Authentication failed : Ssh/Telnet failed to connect with the switch
2017.04.08 08:00:13 INFO [epl] Failed to Enabled End Point Locator. Trying to removing
configuration
2017.04.08 08:00:13 INFO [epl] Disable EPL
2017.04.08 08:00:19 ERROR [epl] Failed to connect to switch 192.169.6.2:java.lang.Exception:
  Authentication failed : Ssh/Telnet failed to connect with the switch
2017.04.08 08:00:19 ERROR [epl] Failed to connect to switch: 192.169.6.2
2017.04.08 08:00:21 ERROR [epl] Failed to unconfigure BGP neighbor or failed to connect to
switch or fabric information not provided.
2017.04.08 08:00:21 ERROR [epl] Failed to unconfigure BGP neighbor or failed to connect to
switch or fabric information not provided.
2017.04.08 08:00:21 INFO [epl] Received response: configure terminal
Interface Ethernet1/1
no ip address
switchport
end
Enter configuration commands, one per line. End with CNTL/Z.
(config)# Interface Ethernet1/1
(config-if)# no ip address
(config-if)# switchport
(config-if)# end
# from 192.169.6.45
2017.04.08 08:00:21 ERROR [epl] Failed to disable EndPoint locator:
java.lang.NullPointerException
2017.04.08 08:00:21 INFO [epl] EPL disabled 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 `/var/afw/applogs/` under the directory for the associated fabric. For example, if EPL is enabled for the `test` fabric, the logs will be in `/var/afw/applogs/epl_cisco_test_afw_log/epl/` starting with filename `afw_bgp.log.1`. Depending on the scale of the network and the number of endpoint events, the file size will increase. Therefore, there is a restriction on the maximum number and size of `afw_bgp.log`. Up to 10 such files will be stored with each file size of maximum of 10MB.




---

**Note** EPL creates a symlink in this directory inside the docker container, hence it appears broken when accessed natively.

---

The EPL relies on BGP updates to get endpoint information. In order for this to work, the switch loopback or VTEP interface IP addresses must be discovered on the DCNM for all switches that have endpoints. To validate, navigate to the Cisco DCNM Web UI > **Dashboard** > **Switch** > **Interfaces** tab, and verify if the IP address and the prefix associated with the corresponding Layer-3 interfaces (typically loopbacks) are displayed correctly.

In a Cisco DCNM Cluster deployment, if EPL cannot establish BGP peering and the active DCNM is able to ping the loopback IP address of the spine, while the EPL container cannot, it implies that the eth2 port group for Cisco DCNM and its computes does not have Promiscuous mode set to **Accept**. After changing this setting, the container can ping the spine and EPL will establish BGP.

In a large-scale setup, it may take more than 30 seconds (default timer set in Cisco DCNM) to get this information from the switch. If this occurs, the `ssh.read-wait-timeout` property (in the **Administration** > **DCNM Server** > **Server Properties**) must be changed from 30000 (default) to 60000 or a higher value.

## Monitoring Endpoint Locator

Information about the Endpoint Locator is displayed on a single landing page or dashboard. The dashboard displays an almost real-time view of data (refreshed every 30 seconds) pertaining to all the active endpoints on a single pane. The data that is displayed on this dashboard depends on the scope selected by you from the **SCOPE** drop-down list. The DCNM scope hierarchy starts with the fabrics. Fabrics can be grouped into a Multi-Site Domain (MSD). A group of MSDs constitute a Data Center. The data that is displayed on the Endpoint Locator dashboard is aggregated based on the selected scope. From this dashboard, you can access Endpoint History, Endpoint Search, and Endpoint Life.

### Endpoint Locator Dashboard

To explore endpoint locator details from the Cisco DCNM Web UI, choose **Monitor** > **Endpoint Locator** > **Explore**. The **Endpoint Locator** dashboard is displayed.



**Note** Due to an increase in scale from Cisco DCNM Release 11.3(1), the system may take some time to collect endpoint data and display it on the dashboard. Also, on bulk addition or removal of endpoints, the endpoint information displayed on the EPL dashboard takes a few minutes to refresh and display the latest endpoint data.

You can also filter and view the endpoint locator details for a specific **Switch**, **VRF**, **Network**, and **Type**, by using the respective drop-down lists. Starting from Cisco DCNM Release 11.3(1), you can select MAC type of endpoints as a filter attribute. By default, the selected option is **All** for these fields.

You can reset the filters to the default options by clicking the **Reset Filters** icon.

The 'top pane' of the window displays the number of active endpoints, active VRFs, active networks, dual attached endpoints, single attached endpoints and dual stacked endpoints, for the selected scope. Support for displaying the number of dual attached endpoints, single attached endpoints and dual stacked endpoints has been added from Cisco DCNM Release 11.3(1). A dual attached endpoint is an endpoint that is behind at least two switches. A dual stacked endpoint is an endpoint that has at least one IPv4 address and one IPv6 address.

Historical analysis of data is performed and a statement mentioning if any deviation has occurred or not over the previous day is displayed at the bottom of each tile.

Click any tile in the top pane of the EPL dashboard to go to the [Endpoint History](#) window.

The 'middle pane' of the window displays the following information:

- **Top 10 Networks by Endpoints** - A pie chart is displayed depicting the top ten networks that have the most number of endpoints. Hover over the pie chart to display more information. Click on the required section to view the number of IPv4, IPv6, and MAC addresses.
- **Top 10 Switches by Endpoints** - A pie chart is displayed depicting the top ten switches that are connected to the most number of endpoints. Hover over the pie chart to display more information. Click on the required section to view the number of IPv4, IPv6, and MAC addresses.
- **Top Switches by Networks** - Bar graphs are displayed depicting the number of switches that are associated with a particular network. For example, if a vPC pair of switches is associated with a network, the number of switches associated with the network is 2.

The 'bottom pane' of the window displays the list of active endpoints.

Click the search icon in the **Endpoint Identifier** column to search for specific IP addresses.

In certain scenarios, the datapoint database may go out-of-sync and information, such as the number of endpoints, may not be displayed correctly due to network issues such as -

- Endpoint moves under the same switch between ports and the port information needs some time to be updated.
- An orphan endpoint is attached to the second VPC switch and is no longer an orphan endpoint.
- NX-API not enabled initially and then enabled at a later point in time.
- NX-API failing initially due to misconfiguration.
- Change in Route Reflector (RR).
- Management IPs of the switches are updated.

In such cases, clicking the **Resync**  icon leads to the dashboard syncing to the data currently in the RR. However, historical data is preserved. We recommend not clicking **Resync** multiple times as this is a compute-intensive activity.

Click the **Pause**  icon to temporarily stop the near real-time collection and display of data.

Consider a scenario in which EPL is first enabled and the **Process MAC-Only Advertisements** checkbox is selected. Then, EPL is disabled and enabled again without selecting the **Process MAC-Only Advertisements** checkbox. As the cache data in elasticsearch is not deleted on disabling of EPL, the MAC endpoint information is still displayed in the EPL dashboard. The same behavior is observed when a Route-Reflector is disconnected. Depending on the scale, the endpoints are deleted from the EPL dashboard after some time. In certain cases, it may take up to 30 minutes to remove the older MAC-only endpoints. However, to display the latest endpoint data, you can click the **Resync** icon at the top right of the EPL dashboard.

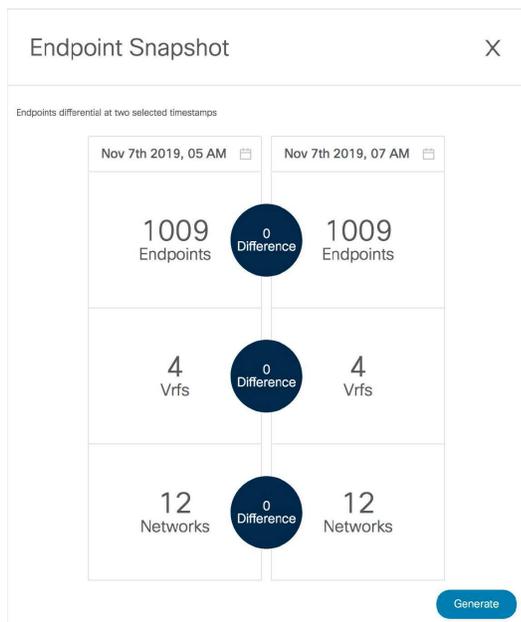
## Endpoint History

Click any tile in the top pane of the EPL dashboard to go to the **Endpoint History** window. A graph depicting the number of active endpoints, VRFs and networks, dual attached endpoints and dual stacked MAC endpoints at various points in time is displayed. The graphs that are displayed here depict all the endpoints and not only the endpoints that are present in the selected fabric. Endpoint history information is available for the last 180 days amounting to a maximum of 100 GB storage space.

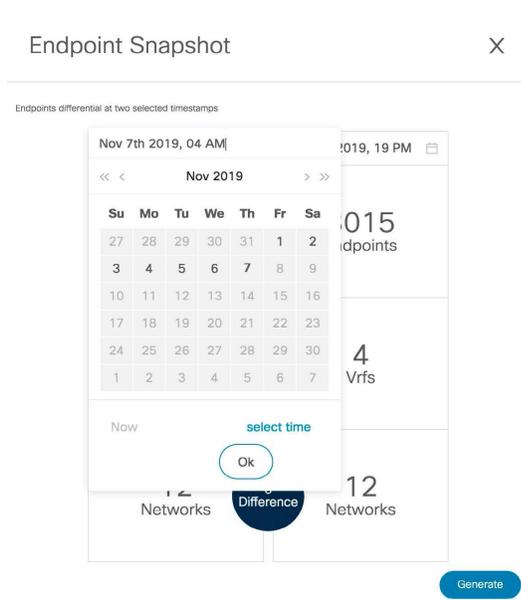
## Endpoint Snapshots

Starting from Cisco DCNM Release 11.3(1), you can compare endpoint data at two specific points in time. To display the **Endpoint Snapshot** window, click the **Endpoint Snapshot** icon at the top right of the **Active Endpoints** graph in the **Endpoint History** window.

By default, endpoint snapshot comparison data for the previous hour is displayed.

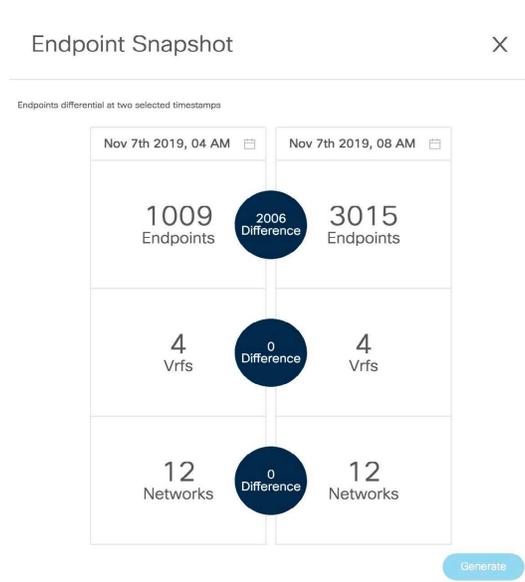


To compare endpoint snapshots at specific points in time, select two points in time, say T1 and T2, and click **Generate**.



A comparison of the endpoints, VRFs, and networks at the selected points in time are displayed. Click each tile to download more information about the endpoints, VRFs, or networks. Click the **Difference** icon to

download details about the differences in data for the specified time interval. Snapshots are stored for a maximum of three months and then discarded.



## Endpoint Search

Click the **Endpoint Search** icon at the top right of the Endpoint Locator landing page to view a real-time plot displaying endpoint events for the period specified in a date range.

## Endpoint Life

Click the **Endpoint Life** icon at the top right of the Endpoint Locator landing page to display a time line of a particular endpoint in its entire existence within the fabric.

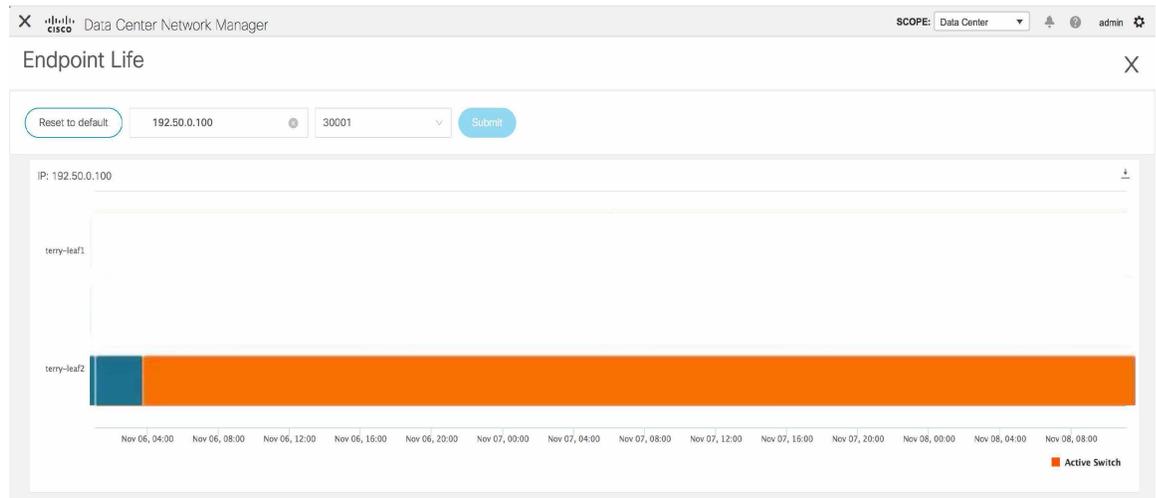
Specify the IP or MAC address of an endpoint and the VXLAN Network Identifier (VNI) to display the list of switches that an endpoint was present under, including the associated start and end dates. Click **Submit**.

The screenshot shows a browser window titled "Data Center Network Manager" with a "SCOPE: lerry-f2" dropdown and "admin" user. The main content is a window titled "Endpoint Life" with a close button (X). Below the title bar, there is a search form with the following elements:

- "Reset to default" button
- "Enter IP or MAC" input field
- "Select VNI" dropdown menu
- "Submit" button

Below the form is a message: "Please enter IP & VNI to see the graph".

The window that is displayed is essentially the endpoint life of a specific endpoint. The bar that is orange in color represents the active endpoint on that switch. If the endpoint is viewed as active by the network, it will have a band here. If an endpoint is dual-homed, then there will be two horizontal bands reporting the endpoint existence, one band for each switch (typically the vPC pair of switches). In case the endpoints are deleted or moved, you can also see the historical endpoint deletions and moves on this window.



## LAN Telemetry Health

Starting from DCNM 11.2(1), Streaming LAN Telemetry preview feature in DCNM is obsolete and is replaced by Network Insights Resources (NIR) application. NIR can be deployed using Cisco DCNM Applications Framework on **Web UI > Applications**. After the NIR is enabled on a fabric, you can monitor the status on the window in the Cisco DCNM Web UI.

When the connection status is shown as **Disconnected** the port configuration may not be accepted by the switch correctly. On the switch image 7.0(3)I7(6), if a switch already had **nxapi** configuration, and later it was managed by DCNM and telemetry was enabled on that fabric, DCNM pushes **http port 80** configuration so that it could query some NXAPI commands such as **show telemetry transport** and **show telemetry data collector details**, to monitor telemetry connection statistics. In this case, the switch does not update **http port 80** in its configuration even though the command was executed correctly. In such a scenario, issue the following commands on the switch:

```
switch# configure
switch(config)# no feature nxapi
switch(config)# feature nxapi
switch(config)# http port 80
```

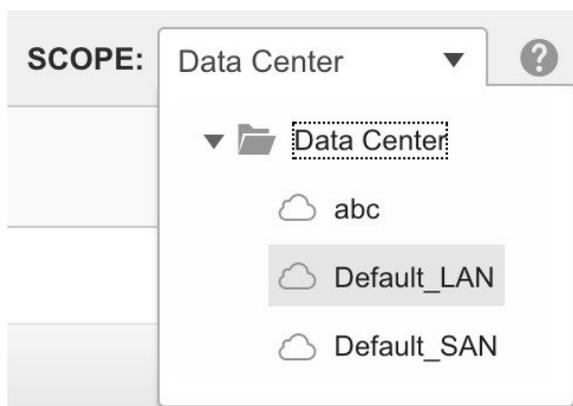


**Note** You cannot configure ICAM on the Cisco Nexus 9000 Series Switches Release 7.0(3)I7(6), and therefore, the telemetry will fail until the switch issue is resolved.

LAN Telemetry has the following topics:

## Health

Cisco DCNM allows you to monitor the health attributes for each fabric. The attributes are displayed for a particular fabric or all fabrics based on the selected **SCOPE**. **Default\_LAN** displays all fabrics.



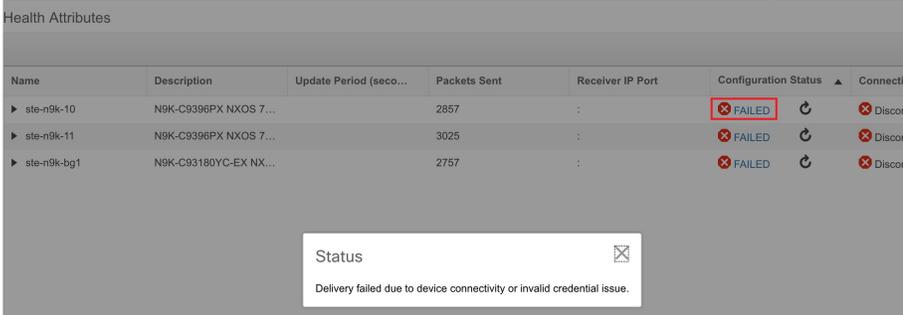
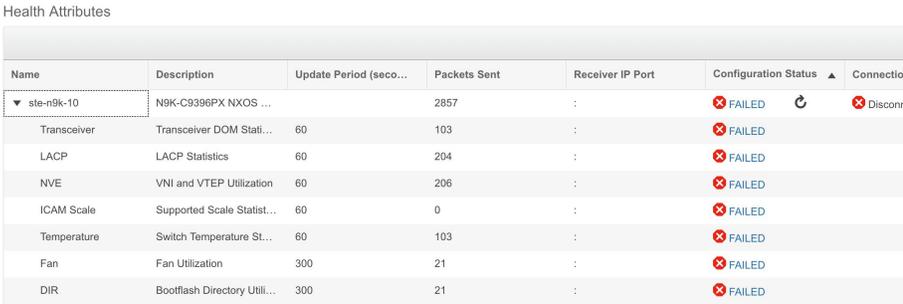
| Health Attributes |                              |                        |              |                  |
|-------------------|------------------------------|------------------------|--------------|------------------|
| Name              | Description                  | Update Period (seco... | Packets Sent | Receiver IP Port |
| ▶ ste-n9k-18-deep | N9K-C9396PX NXOS ...         |                        | 320189       | 24.0.0.4:57500   |
| ▶ ste-n9k-9       | N9K-C9396PX NXOS ...         |                        | 347061       | 24.0.0.2:57500   |
| ▼ ste-n9k-10      | N9K-C9396PX NXOS ...         |                        | 0            | 24.0.0.3:57500   |
| Module            | Modules                      | 60                     | 0            | 24.0.0.3:57500   |
| Fan               | Fan Utilization              | 300                    | 0            | 24.0.0.3:57500   |
| Clock             | System Clock                 | 300                    | 0            | 24.0.0.3:57500   |
| Routing Mode      | System Routing Mode          | 3600                   | 0            | 24.0.0.3:57500   |
| CPU               | Per Process CPU Utiliz...    | 60                     | 0            | 24.0.0.3:57500   |
| Resources         | Overall System Resour...     | 60                     | 0            | 24.0.0.3:57500   |
| MAC               | MAC Address Utilization      | 60                     | 0            | 24.0.0.3:57500   |
| Memory            | Memory Utilization           | 60                     | 0            | 24.0.0.3:57500   |
| VRF               | VRF Utilization              | 60                     | 0            | 24.0.0.3:57500   |
| LACP              | LACP Statistics              | 60                     | 0            | 24.0.0.3:57500   |
| ACL/CAM           | ACL TCAM Utilization         | 60                     | 0            | 24.0.0.3:57500   |
| DIR               | Bootflash Directory Utili... | 300                    | 0            | 24.0.0.3:57500   |
| Interface         | Interface Statistics         | 60                     | 0            | 24.0.0.3:57500   |
| NVE               | VNI and VTEP Utilization     | 60                     | 0            | 24.0.0.3:57500   |
| IP                | IPv4/v6 Unicast/Multica...   | 60                     | 0            | 24.0.0.3:57500   |
| FWD CAM           | Forwarding TCAM Utili...     | 60                     | 0            | 24.0.0.3:57500   |

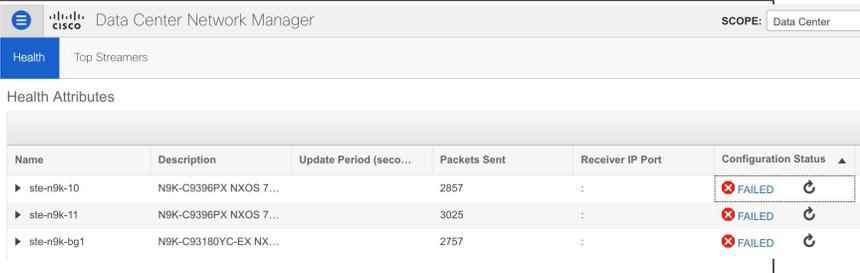
The following table describes the columns in the **LAN Telemetry > Health** tab.

**Table 4: Fields and Description on Health tab**

| Field                   | Description                                                                                                                                                                                            |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name                    | Displays the switch name at the top. The drop-down list displays all the metric names such as CPU, Memory, and so on, that are configured on the switch by the telemetry manager.                      |
| Description             | At the switch level, it displays the switch model and switch image version.<br>At the metric level, it displays the metric description.                                                                |
| Update Period (seconds) | At the switch level, nothing is displayed. At the metric level, it displays the <b>metric collection interval in seconds</b> . Ex: 60 means that the switch shall stream that metric every 60 seconds. |
| Packets Sent            | At the switch level, nothing is displayed.<br>At the metric level, it displays the number of metric samples is collected till time.                                                                    |
| Receiver IP Port        | Displays the IP address of the UTR micro-service and the port that runs as a part of the Network Insights application, which collects the streamed telemetry data from the switches.                   |

| Field                | Description |
|----------------------|-------------|
| Configuration Status |             |

| Field | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|       | <p>Displays the telemetry configuration status on the switches. The following statuses are displayed:</p> <ul style="list-style-type: none"> <li>• <b>MONITOR</b> implies that the switch in the fabric was configured as "Monitored" in the NIR app. In this case, it is the user's responsibility to configure these switches with the required telemetry configurations and stream it to the right destination UTR IP address and Port displayed in the "Receiver IP Port" column.</li> <li>• <b>PROCESSING</b>: This means that the switch belonging to the fabric was configured as "Managed" in the NIR app. In this case, the telemetry manager will configure the switches and when configuration is in progress, it is displayed as "PROCESSING".</li> <li>• <b>SUCCESS</b>: This means that the switches were successfully configured.</li> <li>• <b>FAILED</b>: This means that the switches could not be configured successfully. It could be a partial failure i.e. some of the metrics could not be configured or a full failure i.e. none of configurations went through successfully. Click on the FAILED link to the failure reason.</li> </ul>  <p>Click on the <b>Name</b> drop-down list to check which metrics failed.</p>  <p><b>Note</b> FWD TCAM and ACLCAM metric configurations are not supported on Cisco Nexus C9504, C9508, C9516 series platforms. Check the release notes for any limitations or caveats that could cause a failure.</p> <ul style="list-style-type: none"> <li>• <b>Failure Retry</b>: Click on the retry button to the right of FAILED to reconfigure the switches again.</li> </ul> |

| Field                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Connection Status      | <p>This indicates the status of the connection used to transport telemetry data between the switch and the UTR micro-service running in the NIR app. The value can either be <b>Connected</b> or <b>Disconnected</b>.</p> <p>When the connection status is <b>Disconnected</b>, and if <b>Configuration Status</b> shows either <b>MONITOR</b> or <b>SUCCESS</b>, login to the switch and check the <b>nxapi</b> configuration. When a Cisco DCNM managed fabric is enabled with telemetry, the telemetry manager pushes the <b>http port 80</b> configuration. If the switch does not have <b>http port 80</b> configuration, run the following commands on the switch:</p> <pre>switch# configure switch(config)# no feature nxapi switch(config)# feature nxapi switch(config)# http port 80</pre> |
| Additional Information | Displays the switch serial number, fabric name, and the switch management IP address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



**Note** The **Health** table data gets refreshed every 2 minutes automatically. It can be manually refreshed by clicking the refresh button on the top right corner.

### Failure Troubleshooting

To get more details on the failures, choose **Control > Fabric Builder**. Navigate to **Jobs** tab. Click on the fabric that you want to debug for failures.



## Fabric Builder

Fabric Builder creates a managed and controlled SDN fabric. Select an existing fabric below or define a new *VXLAN* fabric, add switches, set the roles of the switches and deploy settings to devices.

Create Fabric

Fabrics (3)

ABC ⚙️ ✕

Type: Switch Fabric  
ASN: 65000  
Replication Mode: Multicast  
Technology: VXLAN Fabric

DEF\_HW ⚙️ ✕

Type: Switch Fabric  
ASN: 65001  
Replication Mode: Multicast  
Technology: VXLAN Fabric

GHI ⚙️ ✕

Type: Switch Fabric  
ASN: 65003  
Replication Mode: Multicast  
Technology: VXLAN Fabric

Click on the Tabular View.

Actions —

+ — 🔄 🗑️

☰ Tabular view

🔄 Refresh topology

💾 Save layout

✕ Delete saved layout

Hierarchical ▾

🔄 Restore Fabric

🔄 Re-sync Fabric

+ Add switches

⚙️ Fabric Settings

Select a switch. Click **History**.

← Fabric Builder: GHI

Switches Links

+ ↺ ✎ ⏻ ✕ View/Edit Policies Manage Interfaces **History** Deploy

|   | <input type="checkbox"/>            | Name              | IP Address  | Role  | Serial Number | Fabric Name | Fabric Status | Discovery Status |
|---|-------------------------------------|-------------------|-------------|-------|---------------|-------------|---------------|------------------|
| 1 | <input checked="" type="checkbox"/> | gmurthy-spine3    | 15.15.15.25 | spine | FDO223615XK   | GHI         | In-Sync       | ok               |
| 2 | <input type="checkbox"/>            | gmurthy-n9k-leaf6 | 15.15.15.23 | leaf  | FDO22480V9W   | GHI         | In-Sync       | ok               |
| 3 | <input type="checkbox"/>            | gmurthy-n9k-leaf7 | 15.15.15.26 | leaf  | FDO22480VAE   | GHI         | In-Sync       | ok               |

This will list all the Policies that were deployed on that switch and their Status. Check the Status Description column to see the failure reason.

Policy Deployment History for gmurthy-spine3 ( FDO223615XK )

Show Quick Filter

| Entity Name | Entity Type | Source    | Status       | Status Description                                                      | User  | Time of Completion      |
|-------------|-------------|-----------|--------------|-------------------------------------------------------------------------|-------|-------------------------|
| FDO223615XK | SWITCH      | TELEMETRY | NOT_EXECUTED | Delivery failed due to device connectivity or invalid credential issue. | admin | 2019-06-29 21:09:16.681 |
| FDO223615XK | SWITCH      | TELEMETRY | NOT_EXECUTED | Delivery failed due to device connectivity or invalid credential issue. | admin | 2019-06-29 21:09:16.662 |
| FDO223615XK | SWITCH      | TELEMETRY | NOT_EXECUTED | Delivery failed due to device connectivity or invalid credential issue. | admin | 2019-06-29 21:09:16.66  |
| FDO223615XK | SWITCH      | TELEMETRY | NOT_EXECUTED | Delivery failed due to device connectivity or invalid credential issue. | admin | 2019-06-29 21:08:38.703 |
| FDO223615XK | SWITCH      | TELEMETRY | NOT_EXECUTED | Delivery failed due to device connectivity or invalid credential issue. | admin | 2019-06-29 21:08:38.684 |
| FDO223615XK | SWITCH      | TELEMETRY | NOT_EXECUTED | Delivery failed due to device connectivity or invalid credential issue. | admin | 2019-06-29 21:08:38.681 |
| FDO223615XK | SWITCH      | DCNM      | SUCCESS      | Successfully deployed                                                   | admin | 2019-06-19 11:58:35.885 |
| FDO223615XK | SWITCH      | TELEMETRY | SUCCESS      | Successfully deployed                                                   | admin | 2019-06-19 11:36:28.729 |
| FDO223615XK | SWITCH      | TELEMETRY | SUCCESS      | Successfully deployed                                                   | admin | 2019-06-19 11:36:27.526 |
| FDO223615XK | SWITCH      | TELEMETRY | SUCCESS      | Successfully deployed                                                   | admin | 2019-06-19 11:36:26.725 |
| FDO223615XK | SWITCH      | TELEMETRY | SUCCESS      | Successfully deployed                                                   | admin | 2019-06-19 11:36:26.524 |
| FDO223615XK | SWITCH      | DCNM      | SUCCESS      | Successfully deployed                                                   | admin | 2019-06-04 12:58:42.147 |

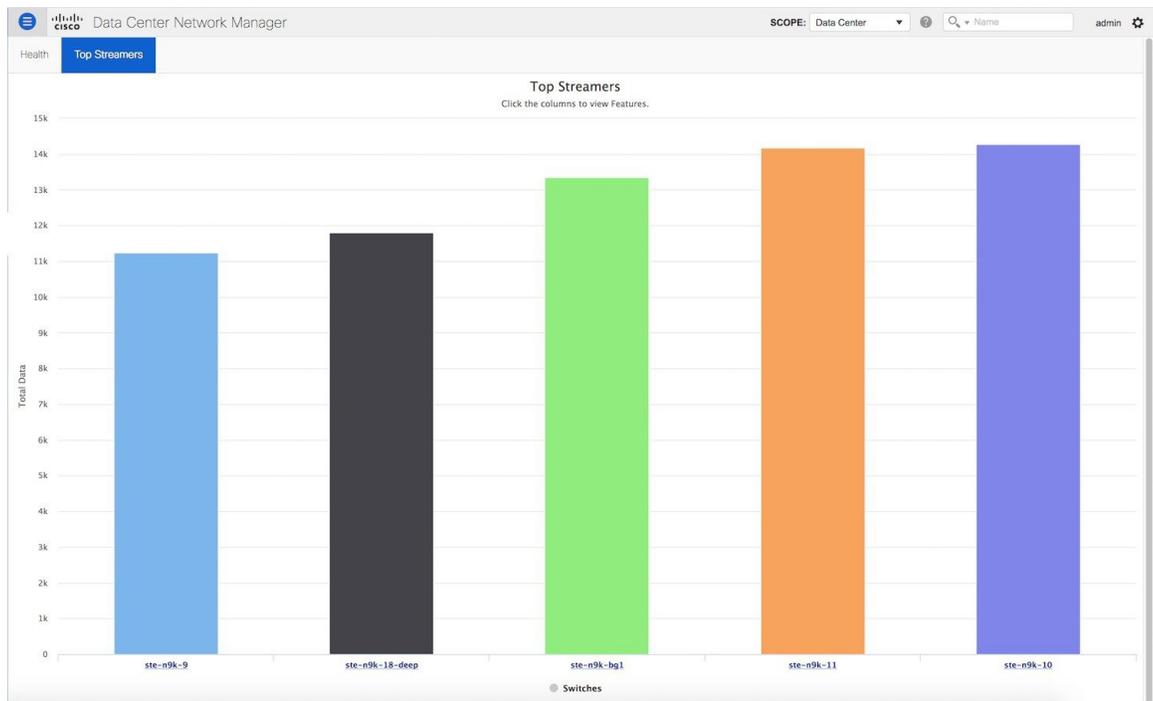
Click on the NOT\_EXECUTED in the Status column to check the commands that were not executed or failed.

## Command Execution Details for gmurthy-spine3 ( FDO223615XK )

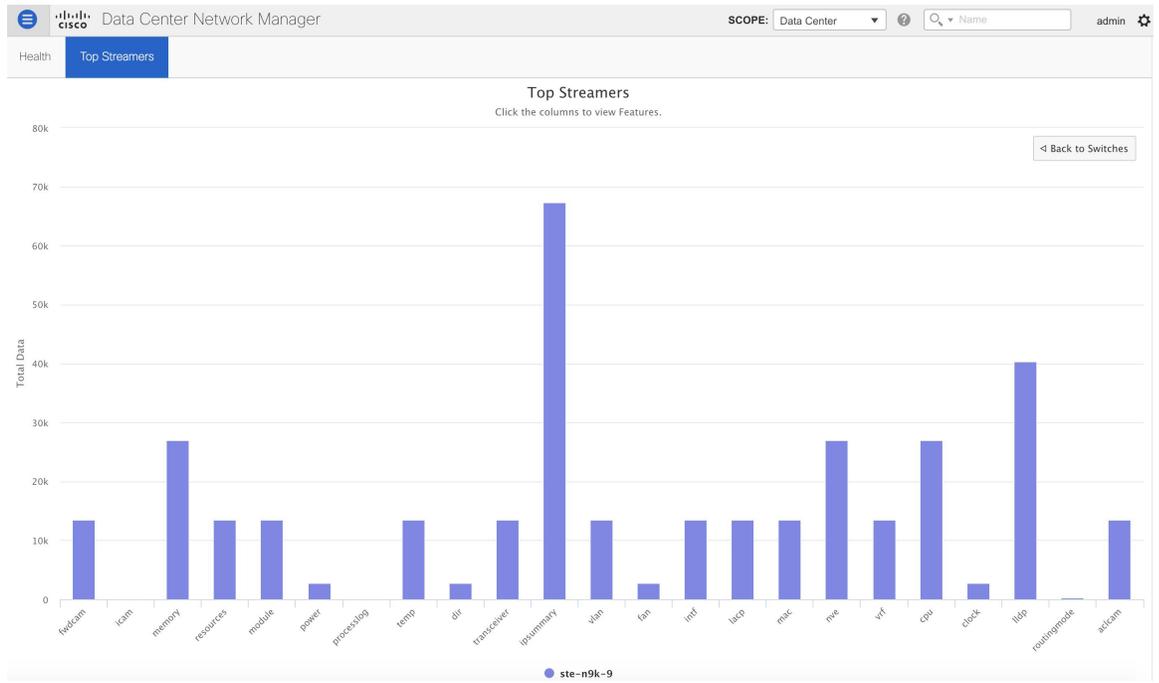
| Config                                       | Status       | CLI Response |
|----------------------------------------------|--------------|--------------|
| telemetry                                    | NOT_EXECUTED |              |
| destination-profile                          | NOT_EXECUTED |              |
| use-vrf pepsi                                | NOT_EXECUTED |              |
| destination-group 500                        | NOT_EXECUTED |              |
| ip address 17.17.17.242 port 57500 p...      | NOT_EXECUTED |              |
| sensor-group 500                             | NOT_EXECUTED |              |
| data-source NX-API                           | NOT_EXECUTED |              |
| path "show interface" depth unbounded        | NOT_EXECUTED |              |
| path "show lacp counters detail" dept...     | NOT_EXECUTED |              |
| path "show lacp interface" depth unb...      | NOT_EXECUTED |              |
| path "show interface transceiver detai...    | NOT_EXECUTED |              |
| path "show mac address-table count"...       | NOT_EXECUTED |              |
| path "show nve vrf" depth unbounded          | NOT_EXECUTED |              |
| path "show ip route summary vrf all" ...     | NOT_EXECUTED |              |
| path "show ipv6 route summary vrf all...     | NOT_EXECUTED |              |
| path "show ip mroute summary vrf all...      | NOT_EXECUTED |              |
| path "show ipv6 mroute summary vrf ...       | NOT_EXECUTED |              |
| path "show lldp traffic interface all" de... | NOT_EXECUTED |              |
| path "show lldp all" depth unbounded         | NOT_EXECUTED |              |

## Top Streamers

Choose **LAN Telemetry > Health > Top Streamers** to view the graphs that depicts the top five streaming switches.



Click on the switch level bar chart to visualize a feature-wise break-down.







## CHAPTER 5

# Monitor

---

This chapter contains the following topics:

- [Inventory, on page 307](#)
- [Monitoring Switch, on page 328](#)
- [Monitoring LAN, on page 331](#)
- [Monitoring Endpoint Locator, on page 335](#)
- [Alarms, on page 343](#)

## 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 DCNM Web UI 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 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 375](#) 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.
- 

## Hosts

You can view host details of switch.

To view the **Hosts** tab, choose **Monitor > Inventory > Switches**, click a switch name in the **Device Name** column, and navigate to the **Hosts** tab.

The following table describes the fields that are displayed:

**Table 5: The Hosts Tab**

| Field            | Description                                                              |
|------------------|--------------------------------------------------------------------------|
| VRF              | Displays VRF details of switch.                                          |
| Host IP          | Displays host IP address of switch.                                      |
| Host MAC Address | Displays host MAC address of switch.                                     |
| VLAN             | Displays configured VLAN on switch.                                      |
| Port             |                                                                          |
| L2 VNI           | Displays layer 2 VXLAN network identifier (L2 VNI) configured on switch. |
| L3 VNI           | Displays layer 3 VXLAN network identifier (L3 VNI) configured on switch. |

## Capacity

You can view the physical capacity of switch.

Capacity tab shows information about the physical ports that are present on the switch.

To view the **Capacity** tab, choose **Monitor > Inventory > Switches**, click a switch name in the **Device Name** column, and navigate to the **Capacity** tab.

The following table describes the fields that are displayed:

**Table 6: The Capacity Tab**

| Field       | Description                                      |
|-------------|--------------------------------------------------|
| Tier        | Displays physical ports available on the switch. |
| Used Ports  | Displays number of used ports on switch.         |
| Total Ports | Displays total number of ports on switch.        |
| Days Left   | Displays total days left.                        |

## Features

You can view features enabled on the switch.

To view the **Features** tab, choose **Monitor > Inventory > Switches**, click a switch name in the **Device Name** column, and navigate to the **Features** tab.

## VXLAN

You can view VXLANs and their details under the **VXLAN** tab.

To view VXLANs, choose **Monitor > Inventory > View > Switches**, and then click a switch name in the **Device Name** column.

The following table describes the fields that are displayed:

**Table 7: 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.                                 |

## VLAN

You can view VLANs and their details under the **VLAN** tab.

To view VLANs, choose **Monitor > Inventory > View > Switches**, and then click a switch name in the **Device Name** column.

The following table describes the fields that are displayed:

**Table 8: The VLAN Tab**

| Field  | Description                                                                                        |
|--------|----------------------------------------------------------------------------------------------------|
| VLAN   | Displays the VLAN configured on the switch.                                                        |
| Name   | Displays the name of VLAN.                                                                         |
| Type   | Displays whether the VLAN is associated with a network.                                            |
| Policy | Displays the name of associated policy. If a policy is not associated, by default it is Undefined. |
| Mode   | Displays the VLAN modes.                                                                           |
| Status | Displays the status of VLAN.                                                                       |
| Ports  | Specifies the port number to which the VLAN is physically connected to the Switch.                 |

## Switch Modules

You can view the switch modules and their details under the **Modules** tab.

To view the **Modules** tab, choose **Monitor** > **Inventory** > **Switches**, click a switch name in the **Device Name** column, and navigate to the **Modules** tab.

The following table describes the fields that are displayed:

**Table 9: The Modules Tab**

| Field        | Description                                                                                                        |
|--------------|--------------------------------------------------------------------------------------------------------------------|
| Name         | Specifies the name of the module.                                                                                  |
| ModelName    | Specifies the model name of the module.                                                                            |
| SerialNum    | Specifies the serial number of the module.                                                                         |
| Type         | Specifies the module type. Valid values are <b>chassis</b> , <b>module</b> , <b>fan</b> , and <b>powerSupply</b> . |
| OperStatus   | Specifies the operational status of the module.                                                                    |
| Slot         | Specifies the slot number of the module.                                                                           |
| H/W Revision | Specifies the NX-OS hardware version.                                                                              |
| S/W Revision | Specifies the NX-OS software version.                                                                              |
| AssetID      | Specifies the asset ID of the module.                                                                              |

## 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 10: FEX Operations**

| Field | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Show  | <p>Allows you to view various configuration details for the selected FEX ID. You can select the following from the drop-down list.</p> <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> <p>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.</p> <p>You can create a show template for FEX. Select template type as SHOW and sub type as FEX.</p> |

**Table 11: 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.                                                    |
| 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.                                                                            |

| Field        | Description                                                                                                                                                                                        |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Serial No.   | Specifies the configured serial number.<br><br><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.                                                                                                                                                           |

## 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.
- ```
flex 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 12: 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 delete the VDC. Select any active VDC radio button and click Delete to remove the VDC associated with the device.   |
| Resume     | Allows you to resume a suspended VDC.   |
| Suspend    | <p>Allows you to suspend an active non-default VDC.</p> <p>Save the VDC running configuration to the startup configuration before suspending the VDC. Otherwise, you will lose the changes to the running configuration.</p> <p><b>Note</b> You cannot suspend the default VDC.</p> <p><b>Caution</b> Suspending a VDC disrupts all traffic on the VDC.</p> |
| 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       | <p>Allows you to view the Interfaces and Resources that are allocated to the selected VDC.</p> <p>In the Interface tab, you can view the mode, admin-status, and operational status for each interface associated with the VDC.</p> <p>In the Resource tab, you can view the allocation of resources and current usage of these resources.</p>              |

Table 13: Vdc Table Field and Description

| Field                      | Description  |
|----------------------------|--|
| Name                       | Displays the unique name for the VDC   |
| Type                       | <p>Species the type of VDC. The two types of VDCs are:</p> <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.   |

| Field  | Description   |
|--|---|
| 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> |
| 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> | Species the IP Address of the VDC Management interface. The status shows if the interface if up or down.  |
| SSH  | Specifies the SSH status  |



---

**Note** If you change the VDC hostname of a neighbor device after initial configuration, the link to the old VDC hostname is not replaced with the new hostname automatically. As a workaround, we recommend manually deleting the link to the old VDC hostname.

---

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.
- [Configuring Ethernet VDCs](#)
  - [Configuring Storage VDCs](#)
- 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 14: 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                             |         |                                    |
| IPv6 multicast route memory                 |         |                                    |
| IPv4 multicast route memory                 |         |                                    |
| IPv6 unicast route memory                   |         |                                    |
| IPv4 unicast route memory                   |         |                                    |

| Resource | Minimum | Maximum |
|----------|---------|---------|
| 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
  - 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.

---

- 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.
  - **IO FPGA** column displays the IO field programmable gate arrays (FPGA) version.
  - **MI FPGA** column displays the MI field programmable gate arrays (FPGA) version.
- 

## 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.

- **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 isplayed.

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 15: vPC Performance, on page 333](#) displays the following vPC configuration details in the data grid view.

**Table 15: 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.                    |

| Column                           | Description   |
|----------------------------------|---|
| Primary vPC Peer - Avg. Tx/sec   | Displays the average sending speed of primary vPC peer.         |
| Primary vPC Peer - Peak Util%    | Displays the peak utilization percentage of primary vPC peer.   |
| Secondary vPC Peer - Device Name | Displays the vPC secondary device name.                         |
| Secondary vPC Interface          | Displays the secondary vPC interface.                           |
| Secondary vPC Peer - Capacity    | Displays the capacity for the secondary vPC peer.               |
| Secondary vPC Peer - Avg. Rx/sec | Displays the average receiving speed of secondary vPC peer.     |
| Secondary vPC Peer - Avg. Tx/sec | Displays the average sending speed of secondary vPC peer.       |
| Secondary vPC Peer - Peak Util%  | 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** are displayed.

The **vPC Consistency**, **Peer-link Consistency**, and **vPC Type2 Consistency** for the vPC are 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:

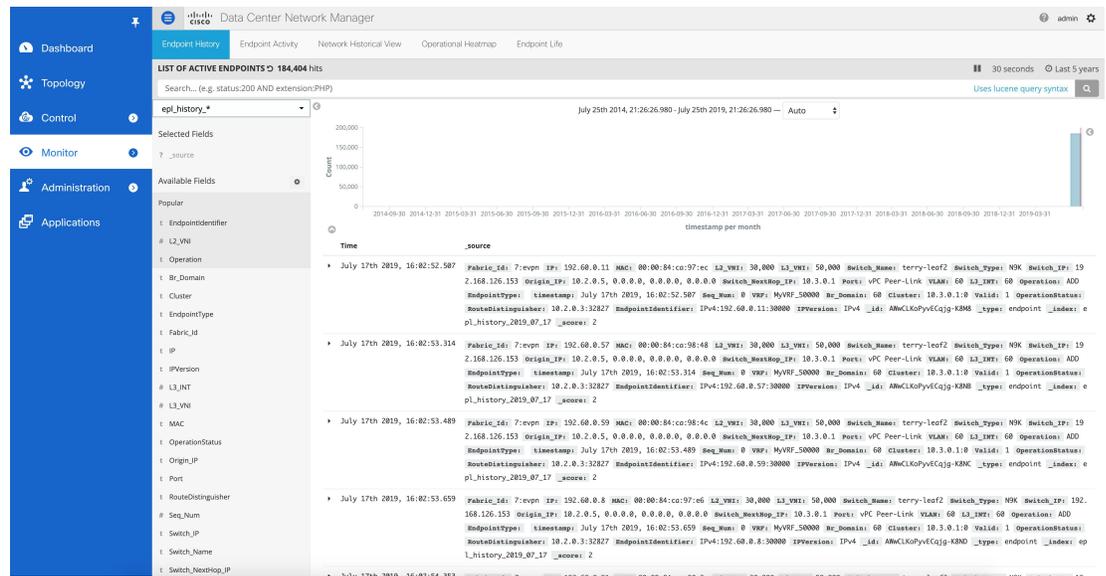
#### Procedure

---

Choose **Monitor > Endpoint Locator > Explore**. The Endpoint Locator dashboard appears.

The Endpoint Locator Dashboard displays the following information:

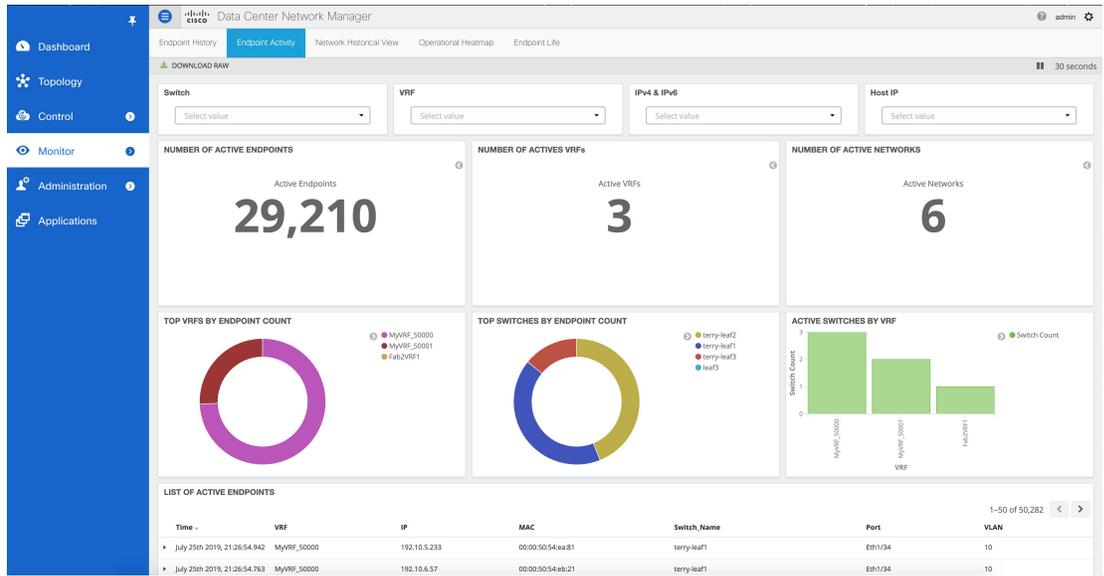
- **Endpoint History**—Real time plot displaying Endpoint events for the period specified in the relative or absolute date range. A user can search for a specific metric value in the search bar. Search is supported on any of the fields as specified under the “Available Fields” column on the menu on the left. A sample screenshot of the endpoint history based on an IP address specified in the search field is depicted below.



- **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:

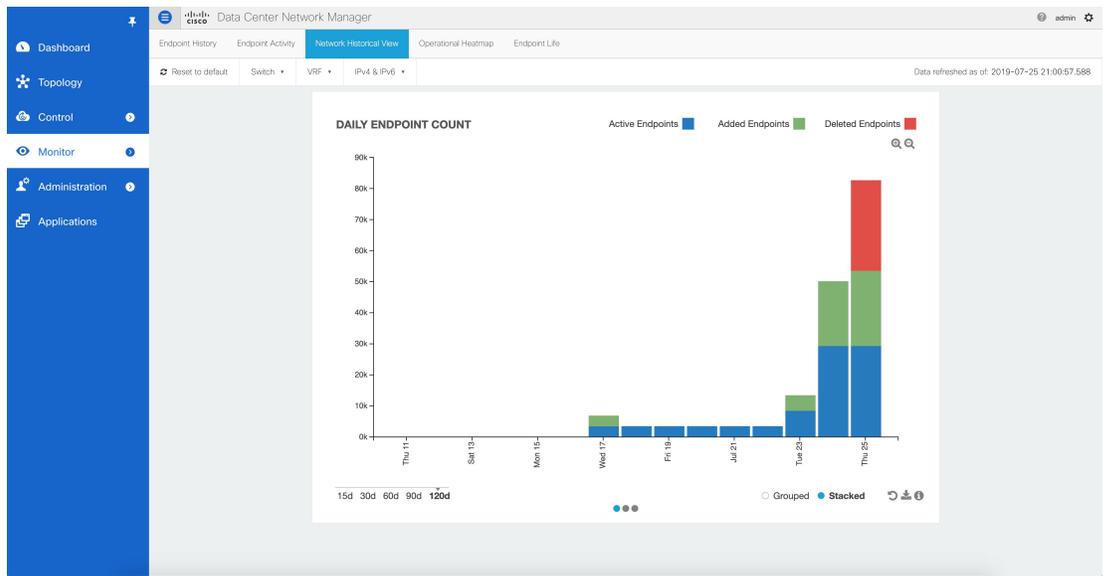
| 1  | Fabric_id | IP         | MAC           | L2_VNI | L3_VNI | Switch_Nam | Switch_Type | Switch_IP | Origin_IP   | Switch_Next | Port        | VLAN | L3_INT | Operation | EndpointType | timestamp        | Seq_Num | VRF   | Br_Domain | Cluster    | Valid | Op |
|----|-----------|------------|---------------|--------|--------|------------|-------------|-----------|-------------|-------------|-------------|------|--------|-----------|--------------|------------------|---------|-------|-----------|------------|-------|----|
| 2  | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 3  | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 4  | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 5  | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 6  | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 7  | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 8  | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 9  | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 10 | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 11 | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 12 | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 13 | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 14 | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 15 | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 16 | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 1  |
| 17 | 3sevpn    | 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      | ACTIVE    |              | 2018-06-30 12:31 | 0       | 50004 | 0         | 10.2.0.1.0 | 1     | 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:

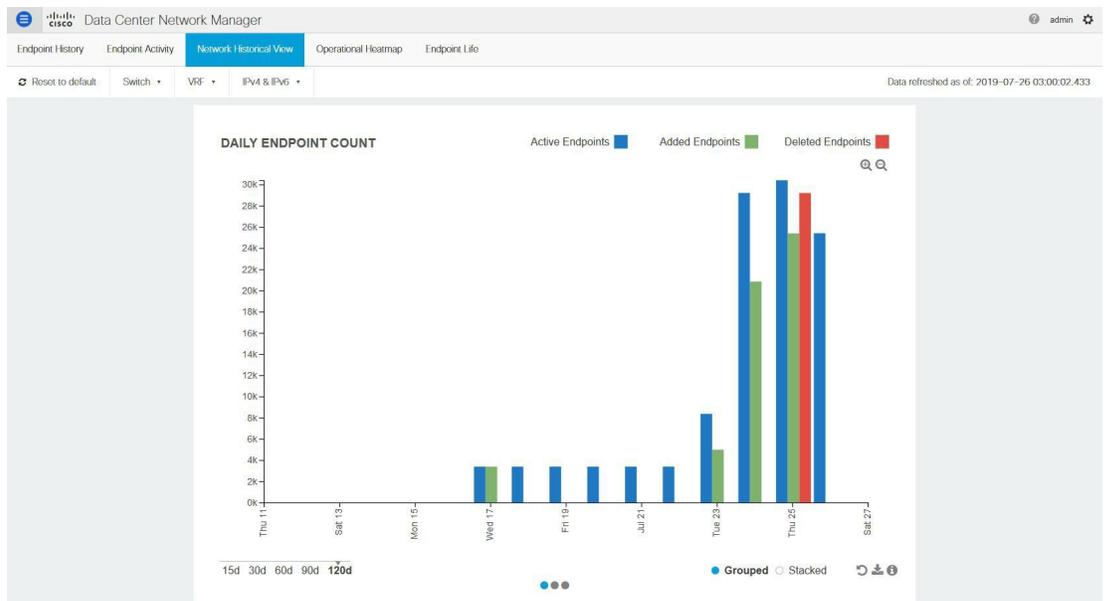
- Go to the LIST OF ACTIVE ENDPOINTS data table and click on any row. This will expand the row as shown below:

| 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:09:9f | 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:09:9f | leaf1, leaf2 |      | 0    |
| ▶ June 30th 2018, 12:31:11.429 | 50004 | 1.0.14.112 | 00:00:00:2f:09:9d | leaf1, leaf2 |      | 0    |
| ▶ June 30th 2018, 12:31:11.409 | 50004 | 1.0.14.112 | 00:00:00:2f:09:9d | leaf1, leaf2 |      | 0    |

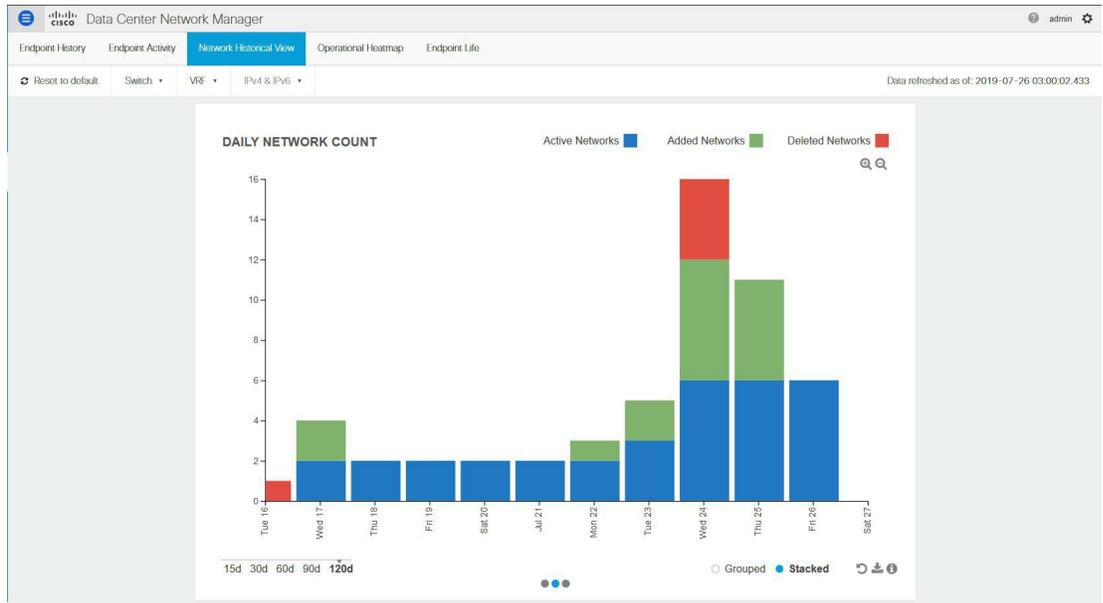




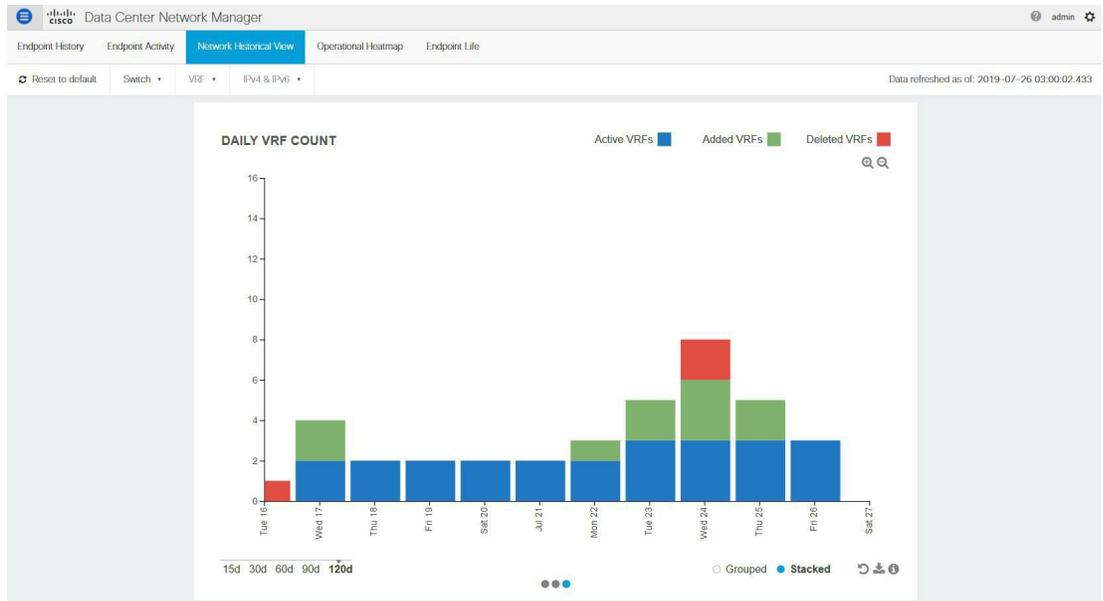
The same representation with the Grouped visualization selection is shown below:



Similarly, the figure below depicts the historic network count in stacked format:



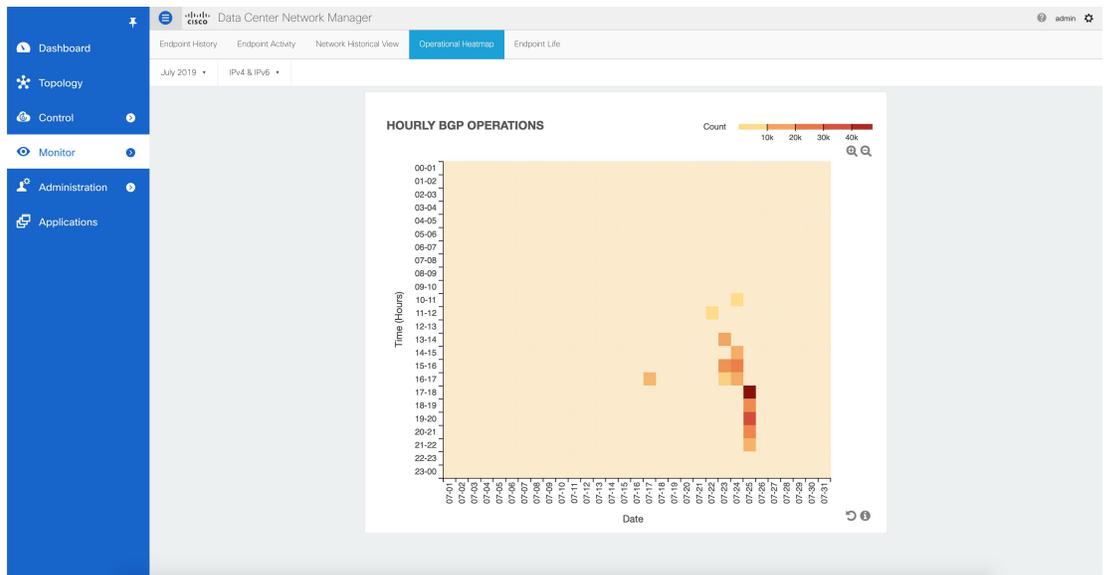
Along the same lines, the figure below depicts the historic vrf count:



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.

| Date       | VRF         | Switch | Operation |
|------------|-------------|--------|-----------|
| 07-25-2019 | Fab2VRF1    | All    | ACTIVE    |
| 07-25-2019 | MyVRF_50001 | All    | ACTIVE    |
| 07-25-2019 | MyVRF_50000 | All    | ACTIVE    |

- **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.

Cisco Data Center Network Manager

Endpoint History | Endpoint Activity | Network Historical View | **Operational Heatmap** | Endpoint Life

July 2019 | IPv4 & IPv6

Complete data set will be available in the downloaded csv. [Download](#)

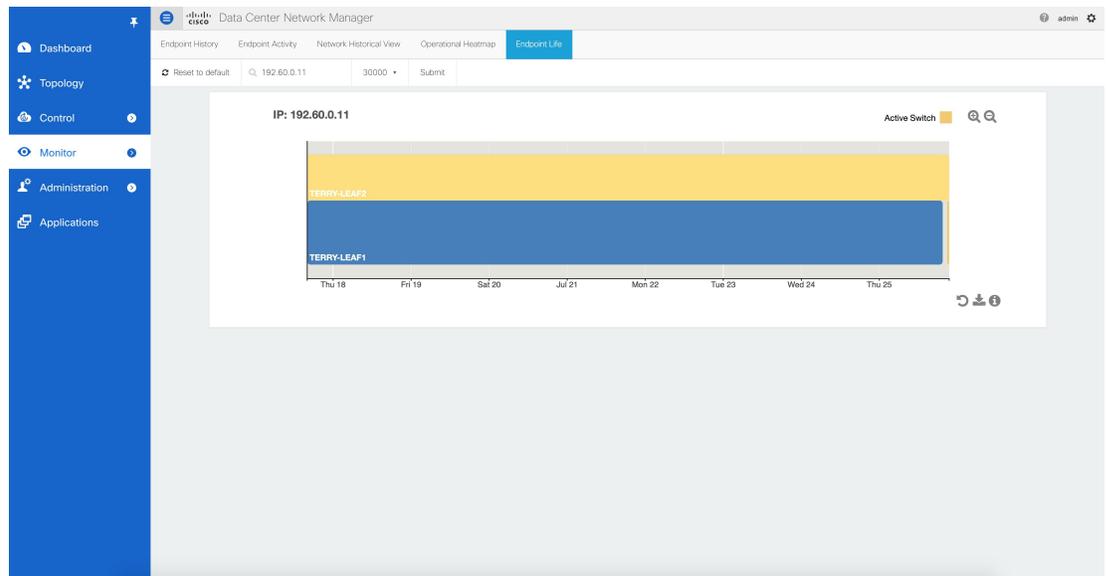
OPERATIONS: 07-25-2019 6:00PM - 7:00PM

| Time                | VRF         | IP            | MAC               | Switch Name | Operation | VLAN |
|---------------------|-------------|---------------|-------------------|-------------|-----------|------|
| 2019-07-25 18:03:51 | MYVRF_50000 | 192.60.21.147 | 00:00:84:ca:c2fc  | terry-leaf1 | ADD       | 60   |
| 2019-07-25 18:03:53 | MYVRF_50000 | 192.60.17.213 | 00:00:84:ca:bb:80 | terry-leaf1 | ADD       | 60   |
| 2019-07-25 18:03:53 | MYVRF_50000 | 192.60.21.235 | 00:00:84:ca:c3ac  | terry-leaf1 | ADD       | 60   |
| 2019-07-25 18:03:53 | MYVRF_50000 | 192.60.19.79  | 00:00:84:ca:be:74 | terry-leaf1 | ADD       | 60   |
| 2019-07-25 18:03:54 | MYVRF_50000 | 192.60.23.41  | 00:00:84:ca:c5:28 | terry-leaf1 | ADD       | 60   |
| 2019-07-25 18:03:55 | MYVRF_50000 | 192.60.22.122 | 00:00:84:ca:c4:ca | terry-leaf1 | ADD       | 60   |
| 2019-07-25 18:03:57 | MYVRF_50000 | 192.60.19.19  | 00:00:84:ca:bd:fc | terry-leaf1 | ADD       | 60   |
| 2019-07-25 18:03:59 | MYVRF_50000 | 192.60.22.195 | 00:00:84:ca:c5:5c | terry-leaf1 | ADD       | 60   |
| 2019-07-25 18:03:59 | MYVRF_50000 | 192.60.20.217 | 00:00:84:ca:c1:88 | terry-leaf1 | ADD       | 60   |
| 2019-07-25 18:03:59 | MYVRF_50000 | 192.60.24.187 | 00:00:84:ca:c9:4c | terry-leaf1 | ADD       | 60   |
| 2019-07-25 18:04:00 | MYVRF_50000 | 192.60.23.21  | 00:00:84:ca:c8:00 | terry-leaf1 | ADD       | 60   |
| 2019-07-25 18:03:45 | MYVRF_50000 | 192.60.6.58   | 00:00:84:ca:a4:4a | terry-leaf1 | ADD       | 60   |
| 2019-07-25 18:03:46 | MYVRF_50000 | 192.60.7.84   | 00:00:84:ca:a6:7e | terry-leaf1 | ADD       | 60   |
| 2019-07-25 18:03:49 | MYVRF_50000 | 192.60.8.9    | 00:00:84:ca:a7:a8 | terry-leaf1 | ADD       | 60   |
| 2019-07-25 18:03:50 | MYVRF_50000 | 192.60.26.97  | 00:00:84:ca:cc:98 | terry-leaf1 | ADD       | 60   |

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:

| 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_NextPort | VLAN | L3_INT | Operation | EndpointType | Timestamp     | Seq_Num | VRF   | Br_Domain | Clust    |
|----|-----------|------|-----------|--------------|--------|------------|--------------|-----------|-------------|-------------|-------------|-------------|-----------------|------|--------|-----------|--------------|---------------|---------|-------|-----------|----------|
| 2  | 3         | evpn | 51.1.1.33 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 3  | 3         | evpn | 51.1.1.53 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 4  | 3         | evpn | 51.1.1.93 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 5  | 3         | evpn | 51.1.1.12 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 6  | 3         | evpn | 51.1.1.35 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 7  | 3         | evpn | 51.1.1.88 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 8  | 3         | evpn | 51.1.1.50 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 9  | 3         | evpn | 51.1.1.79 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 10 | 3         | evpn | 51.1.1.45 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 11 | 3         | evpn | 51.1.1.71 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 12 | 3         | evpn | 51.1.1.67 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 13 | 3         | evpn | 51.1.1.38 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 14 | 3         | evpn | 51.1.1.27 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 15 | 3         | evpn | 51.1.1.94 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 16 | 3         | evpn | 51.1.1.96 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 17 | 3         | evpn | 51.1.1.47 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 18 | 3         | evpn | 51.1.1.56 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 19 | 3         | evpn | 51.1.1.60 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 20 | 3         | evpn | 51.1.1.83 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 21 | 3         | evpn | 51.1.1.18 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 22 | 3         | evpn | 51.1.1.57 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 23 | 3         | evpn | 51.1.1.61 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 24 | 3         | evpn | 51.1.1.12 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 25 | 3         | evpn | 51.1.1.19 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 26 | 3         | evpn | 51.1.1.65 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 0         | ADD          | Sun Jul 01 2C | 0       | 50002 |           | 0.10.2.1 |
| 27 | 3         | evpn | 51.1.1.75 | 00:00:48:69: | 30009  | 50002      | leaf1, leaf2 | NSK       | 24.0.80.203 | 10.1.0.7    | 0.0.0.0     | 0.0.0.0     | 10.2.0.1        |      |        | 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   | End Timestamp   | Active |
| 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) |   | TRUE   |
| 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   |

## 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 **Monitor > Alarms > View**.

**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



### Note

- Alarm policies are stored in compute nodes. Therefore, run the **appmgr backup** command on each compute node in addition to taking a backup of DCNM.

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 `fmsserver.jks` located at

```
/usr/local/cisco/dcm/wildfly-10.1.1.0.Final/standalone/configuration to
/etc/elasticsearch. If you do not copy the fmsserver.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. Under **Device Features**, you can select the BFD, BGP, and HSRP protocols. When these checkboxes are selected, alarms are triggered for the following traps: **BFD**- ciscoBfdSessDown, ciscoBfdSessUp, **BGP**- bgpEstablishedNotification, bgpBackwardTransNotification, cbgpPeer2BackwardTransition (), cbgpPeer2EstablishedNotification, and **HSRP**- cHsrpStateChange. Please refer <https://snmp.cloudapps.cisco.com/Support/SNMP/do/BrowseOID.do?local=en> for detailed trap OID definition.
- **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 16: Example 1**

| Identifier  | ID1-ID2   |
|-------------|---|
| 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 17: Example 2**

| Identifier  | ID1-ID2  |
|-------------|--|
| 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 18: Example 3

| Identifier  | ID1-ID2  |
|-------------|--|
| 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
__pfm_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
__pfm_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.
- 

## Health Monitor Alarms

Starting from Cisco DCNM Release 11.4(1), alarms are registered and created under the External alarm category by the Health Monitor.

### Health Monitor: Alarm Policy

The Health Monitor external alarm category policy is automatically activated and enabled on all the devices in a fabric. The severity level of this alarm policy can be MINOR, MAJOR, or CRITICAL.

Alarms are raised and categorized as CRITICAL for the following events:

- Elasticsearch (ES) Cluster Status is Red: Critical (For Cluster/HA mode only)
- CPU/Memory/Disk Utilization/ES JVM Heap Used Percentage  $\geq 90\%$

Alarms are raised and categorized as MAJOR for the following events:

- ES Cluster Status is Yellow (For Cluster/HA mode only)
- ES has unassigned shards (For Cluster/HA mode only)

- CPU/Memory/Disk Utilization/ES JVM Heap Used Percentage  $\geq 80\%$  and  $<90\%$

Alarms are raised and categorized as MINOR for the following events:

- CPU/Memory/Disk Utilization/ES JVM Heap Used Percentage  $\geq 65\%$  and  $<80\%$
- Kafka: Number of partitions without active leader  $> 0$
- Kafka: Qualified partition leader not found. Unclear leaders  $> 0$

Choose **Monitor>Alarms>Policies** to display the Health Monitor alarm policies. These alarm policies are not editable on the web UI. Click **Activate** or **Deactivate** to activate or deactivate the selected policy.

The screenshot shows the Cisco Data Center Network Manager interface for the 'Monitor / Alarms / Policies' section. It features a table of alarm policies with the following columns: Name, Description, Status, Policy Type, Devices, Interfaces, and Details. The table contains six rows of policies, all with a status of 'Active' and 'External' policy type. The policies are: EPL: Terry-FX2: MINOR, Config-Compliance: Terry-F..., EPL: Terry-FX2: CRITICAL, Health-Monitor: Critical, Health-Monitor: Major, and Health-Monitor: Minor. Above the table, there are buttons for 'Add', 'Edit', 'Delete', 'Activate', 'Deactivate', 'Import', and 'Export'. A 'Show' dropdown menu is set to 'Quick Filter'.

| Name                          | Description                    | Status | Policy Type | Devices     | Interfaces | Details   |
|-------------------------------|--------------------------------|--------|-------------|-------------|------------|---|
| EPL: Terry-FX2: MINOR         | MINOR EPL alarms               | Active | External    | All Devices |            | MINOR alarms auto generated by EPL                    |
| Config-Compliance: Terry-F... | Device level Config-Compla...  | Active | External    | All Devices |            | Alarm created when device status is Out-of-Sync, clea |
| EPL: Terry-FX2: CRITICAL      | CRITICAL EPL alarms            | Active | External    | All Devices |            | CRITICAL alarms auto generated by EPL                 |
| Health-Monitor: Critical      | Critical Health Monitor alarms | Active | External    | All Devices |            | Critical alarms auto generated by Health Monitor      |
| Health-Monitor: Major         | Major Health Monitor alarms    | Active | External    | All Devices |            | Major alarms auto generated by Health Monitor         |
| Health-Monitor: Minor         | Minor Health Monitor alarms    | Active | External    | All Devices |            | Minor alarms auto generated by Health Monitor         |

In case an alarm policy is deactivated using the GUI, any alarms created or cleared for that policy will not be displayed in the **Monitor>Alarms>View** tab. To delete a policy, select the checkbox next to the policy and click **Delete**. However, we recommend not deleting a policy from the GUI. When a fabric is deleted, the alarm policy along with all the active alarms for the devices in that fabric are deleted.

### Health Monitor: Active Alarms

Choose **Monitor>Alarms>View** to display the active alarms.

To clear active alarms, select the checkbox next to the alarm, click **Change Status** and select **Clear**.

To delete active alarms, select the checkbox next to the alarm and click **Delete**.

### Health Monitor: Cleared Alarms

To view the cleared alarms, select **Monitor>Alarms>View>Cleared Alarms**.

Click the arrow icon  to display detailed information about the required alarm.

To delete a cleared alarm from the list of cleared alarms, select the checkbox next to the alarm and click **Delete**.

For more information on Alarms and Policies, refer [Alarms](#).





## CHAPTER 6

# Administration

---

This chapter contains the following topics:

- [DCNM Server, on page 351](#)
- [Management Users, on page 367](#)
- [Performance Setup, on page 374](#)
- [Event Setup, on page 375](#)
- [Credentials Management, on page 379](#)

## DCNM Server

The DCNM Server menu includes the following submenus:

### Starting, Restarting, and Stopping Services

By default, the ICMP connectivity between DCNM and its switches validates the connectivity during Performance Management. If you disable ICMP, Performance Management data will not be fetched from the switches. You can configure this parameter in the **server properties**. To disable ICMP connectivity check from Cisco DCNM Web UI, choose **Administration > DCNM Server > Server Properties**, and set `skip.checkPingAndManageable` parameter value to `true`.

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 action you want to perform. You can perform the following actions:
- Start or restart a service.
  - Stop a service.
  - Clean up the stale PM DB entries.

- Reinitialize the Elasticsearch DB schema.

**Step 3** View the status in the **Status** column.

---

### What to do next

See the latest status in the **Status** column.

### Using the Commands Table

The commands table contains links to commands that launch new dialog boxes to provide information about the server status and server administrative utility scripts. You can execute these commands directly on the server CLI.

- **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.
- **appmgr show vmware-info**: click this link to view information about the CPU and Memory of Virtual Machine.
- **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, 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.

Beginning with Release 11.2(1), for DCNM OVA and DCNM ISO installations, all log files with .log extension are also listed.




---

**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** (Optional) Click **Generate Techsupport** to generate and download files required for technical support.

This file contains more information in addition to log files.

**Note** A TAR.GZ file will be downloaded for OVA and ISO deployments, and a ZIP file will be downloaded for all other deployments.

- Step 5** (Optional) 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/dcnm/data/archive
```

The number of archived files that can be retained is set in the **# Number of archived files per device to be retained:** field. In the 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**
- **Smart License**
- **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.  |
| Expiration Date  | Displays the expiry date of the license.<br><b>Note</b> Text under the <b>Expiration Date</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.

After you register smart license, if you click **Assign License** for a switch that does not have a permanent license, a smart license is assigned to the switch. The priority of licenses that are assigned are in the following order:

1. **Permanent**
2. **Smart**
3. **Eval**

Disabling smart licensing unassigns licenses of switches that were smart-licensed.

The evaluation license is assigned for switches that do not support smart licensing. The license state is **Eval** and the license type is **DCNM-Server**. See *Cisco DCNM Licensing Guide, Release 11.x* to view the list of switches that support smart licensing.

## Smart License

From Cisco DCNM Release 11.1(1), you can use the smart licensing feature to manage licenses at device-level and renew them if required. From Cisco DCNM Web UI, choose **Administration > DCNM Server > License > Smart License**. You will see a brief introduction on Cisco smart licensing, a menu bar, and the **Switch Licenses** area.

In the introduction, click **Click Here** to view the information on smart software licensing.

The menu bar has the following icons:

- **Registration Status:** Displays details of the current registration in a pop-up window when clicked. The value is **UNCONFIGURED** if the smart licensing is not enabled. After you enable the smart licensing without registering, the value is set to **DEREGISTERED**. The value is set to **REGISTERED** after you register. Click the registration status to view the last action, account details, and other registration details in the **Registration Details** pop-up window.
- **License Status:** Specifies the status of the license. The value is **UNCONFIGURED** if the smart licensing is not enabled. After you enable the smart licensing without registering, the value is set to **NO LICENSES IN USE**. The value is set to **AUTHORIZED** or **OUT-OF-COMPLIANCE** after registering and assigning licenses. Click the license status to view the last action, last authorization attempt, next authorization attempt, and the authorization expiry in the **License Authorization Details** pop-up window.
- **Control:** Allows you to enable or disable smart licensing, register tokens, and renew the authorization.

The following table describes the fields that appear in the **Switch Licenses** section.

| Field        | Description  |
|--------------|--|
| Name         | Specifies the license name.  |
| Count        | Specifies the number of licenses used.   |
| Status       | Specifies the status of the licenses used. Valid values are <b>Authorized</b> and <b>Out of Compliance</b> . |
| Description  | Specifies the type and details of the license.   |
| Last Updated | Specifies the timestamp when switch licenses were last updated.  |

| Field  | Description   |
|--------|---|
| Print  | Allows you to print the details of switch licenses. |
| Export | Allows you to export the license details.           |

After you remove a product license from your account in Cisco Smart Software Manager, disable the smart licensing and register it again.

## Enabling Smart Licensing

To enable smart licensing from Cisco DCNM Web UI, perform the following steps:

### Procedure

- 
- Step 1** Choose **Administration > DCNM Server > License > Smart License**.
- Step 2** Click **Control** and choose **Enable** in the drop-down list to enable the smart licensing.  
A confirmation window appears.
- Step 3** Click **Yes**.  
Instructions to register the DCNM instance appear.  
The registration status changes from **UNCONFIGURED** to **DEREGISTERED**, and the license status changes from **UNCONFIGURED** to **No Licenses in Use**.
- 

## Registering a Cisco DCNM Instance

### Before you begin

Create a token in Cisco Smart Software Manager.

### Procedure

- 
- Step 1** Choose **Administration > DCNM Server > License > Smart License**.
- Step 2** Click **Control** and choose **Register** in the drop-down list.  
The **Register** window appears.
- Step 3** Select the transport option to register the smart licensing agent.  
The options are:
- **Default - DCNM communicates directly with Cisco's licensing servers**  
This option uses the following URL: <https://tools.cisco.com/its/service/odcce/services/DDCCEService>
  - **Transport Gateway - Proxy via Gateway or Satellite**  
Enter the URL if you select this option.

- **Proxy - Proxy via intermediate HTTP or HTTPS proxy**

Enter the URL and the port if you select this option.

**Step 4** Enter the registration token in the **Token** field.

**Step 5** Click **Submit** to register the license.

The registration status changes from **DEREGISTERED** to **REGISTERED**. The name, count, and status of switch licenses appear.

Click **Registration Status: REGISTERED** to see the details of the registered token.

The switch details are updated under the **Switches/VDCs** section of the **License Assignments** tab. The license type and the license state of switches that are licensed using the smart license option are **Smart**.

### What to do next

Troubleshoot communication errors, if any, that you encounter after the registration.

### Troubleshooting Communication Errors

To resolve the communication errors during registration, perform the following steps:

#### Procedure

**Step 1** Stop the DCNM service.

**Step 2** Open the server properties file from the following path: `/usr/local/cisco/dcm/fm/conf/server.properties`

**Note** The server properties file for Windows will be in the following location: `C:/Program Files/Cisco/dcm/fm/conf/server.properties`

**Step 3** Include the following property in the server properties file: `#cisco.smart.license.production=false`  
`#smartlicense.url.transport=https://CiscoSatellite_Server_IP/Transportgateway/services/DeviceRequestHandler`

**Step 4** Update the Cisco satellite details in Host Database in the `/etc/hosts` file in the following syntax:  
`Satellite_Server_IP CiscoSatellite`

**Step 5** Start the DCNM service.

### Renew Authorization

You can manually renew the authorization only if you have registered. Automatic reauthorization happens periodically. Click **License Status** to view details about the next automatic reauthorization. To renew authorization from Cisco DCNM Web UI, perform the following steps:

#### Procedure

**Step 1** Choose **Administration > DCNM Server > License > Smart License**.

**Step 2** Click **Control** and choose **Renew Authorization** in the drop-down list to renew any licensing authorizations.

A request is sent to Cisco Smart Software Manager to fetch updates, if any. The **Smart Licenses** window is refreshed after the update.

## Disabling Smart Licensing

To disable smart licensing from Cisco DCNM Web UI, perform the following steps:

### Procedure

- Step 1** Choose **Administration > DCNM Server > License > Smart License**.
- Step 2** Select **Control** and select **Disable** to disable smart licensing.  
A confirmation window appears.
- Step 3** Click **Yes**.  
The license status of the switches using this token, under the **License Assignments** tab, changes to **Unlicensed**. This token is removed from the list under the **Product Instances** tab in the Cisco Smart Software Manager.  
If a smart license is not available and you disable smart licensing, release the license manually from the **License Assignments** tab.

## Server License Files

From Cisco DCNM Web UI, choose **Administration > DCNM Server > License > 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.  |
| Expiration Date  | Displays the expiry date of the license.<br><b>Note</b> Text in the <b>Expiration Date</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.

a. SSH into the DCNM active host.

b. 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/ pgsq-standby-backup.tgz. If the file exists, restore database data files:

```
rm -rf data/*
tar -zxf replication/ pgsq-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.

---

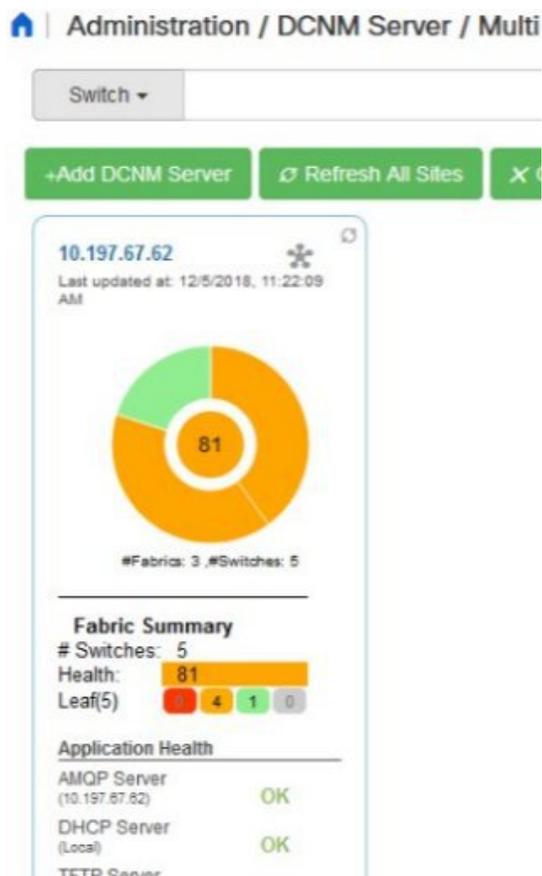
## Multi Site Manager

Using Multi Site Manager, you can view the health of a DCNM server application and retrieve switch information for switches in local and remote sites. To access switch information for remote DCNM servers, you must register the server in Multi Site Manager. The procedures to access remote DCNM servers and search for switch information are explained:

### Add Remote DCNM Server Information

This procedure allows you to access a DCNM server in a remote site from the DCNM server that you are currently logged on to. For the remote site to access the current DCNM server, registration is required on the remote site.

1. Choose **Administration > DCNM Server > Multi Site Manager**. The Multi Site Manager screen comes up.



The currently logged on DCNM application health status is displayed on the screen.



**Note** The **Application Health** function is only available for the DCNM ISO/OVA installation type and not for the Windows/RHEL installation type.

- Click **+Add DCNM Server**. The **Enter Remote DCNM Server Information** screen comes up.

Enter the remote DCNM server name, its IP address or URL, the user credentials of the remote DCNM server, and optionally, the port number.



**Note** Do not disable the **Use HTTPS** check box. If you disable, DCNM will not be accessible.

## Enter Remote DCNM Server Information

|               |   |
|---------------|---|
| * DCNM Name   | <input type="text" value="remote-DCNM"/>  |
| * IP/DNS Name | <input type="text" value="172.28.8.125"/> |
| * User        | <input type="text" value="admin"/>        |
| * Password    | <input type="password" value="....."/>    |
| Use HTTPS     | <input checked="" type="checkbox"/>       |
| Port Number   | <input type="text" value="1099"/>         |

- Click **OK**. After validation, the remote DCNM server is represented in the screen, next to the local DCNM server.

The screenshot shows the Multi Site Manager interface. At the top, there is a 'Switch' dropdown, 'Search', and 'Clear' buttons. Below these are three green buttons: '+Add DCNM Server', 'Refresh All Sites', and 'Clear All Search Result'. The main area displays two server cards. The left card is for IP 10.197.67.62, last updated at 12/5/2018, 11:22:09 AM. It features a donut chart with the number 81 in the center and the text '#Fabrics: 3, #Switches: 5'. Below the chart is a 'Fabric Summary' section with '# Switches: 5'. The right card is for 'remote-DCNM', last updated at 10/10/2018, 5:39:59 PM. It features a donut chart with the number 32 in the center and the text '#Fabrics: 1, #Switches: 7'. Below the chart is an 'Application Health' section with 'AMQP Server (172.28.8.125)' and a green 'OK' status. A red arrow points to the 'remote-DCNM' card.

You can click **Refresh All Sites** to display updated information.

### Retrieve Switch Information

- Choose **Administration > DCNM Server > Multi Site Manager**. The Multi Site Manager screen comes up

- From the search box at the top of the screen, search for a switch based on one of the following parameters:
  - VM information (**VM IP** and **VM Name** fields) - A connected VM's IP address or name.
  - Switch information (**Switch** and **MAC** fields) – A switch's name or MAC address.
  - Segment (**Segment ID** field) that has presence on the switch.

If there is a match, the switch name appears as a hyperlink below the search box, in the appropriate local or remote DCNM server depiction.

In this example, the switch **leaf3** is available in the remote site managed by a DCNM server. A link to **leaf3** is available in the **remote-DCNM** panel.

- Click **leaf3** to view detailed switch information in an adjacent browser tab.

At any point in time, you can click the **Launch Topology View** icon to view the fabric's topology.

## Device Connector

The Device Connector is an embedded management controller that enables the capabilities of Cisco Intersight, a cloud-based management platform.

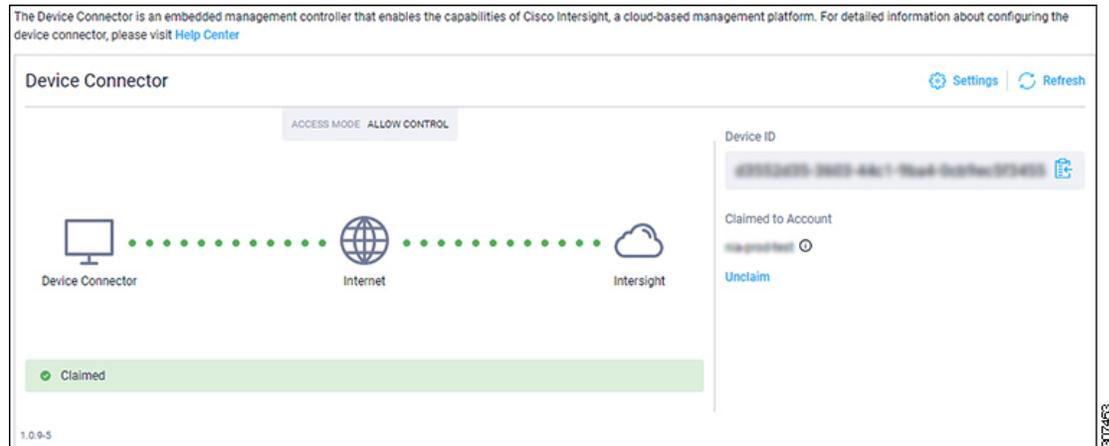
Networks Insights applications are connected to the Cisco Intersight cloud portal through a Device Connector which is embedded in the management controller of the Cisco DCNM platform. Cisco Intersight is a virtual appliance that helps manage and monitor devices through the Network Insights application. The Device Connector provides a secure way for connected DCNM to send information and receive control instructions from the Cisco Intersight portal, using a secure Internet connection.

## Configuring Device Connector

To configure the Device Connector from the Cisco DCNM Web UI, perform the following steps:

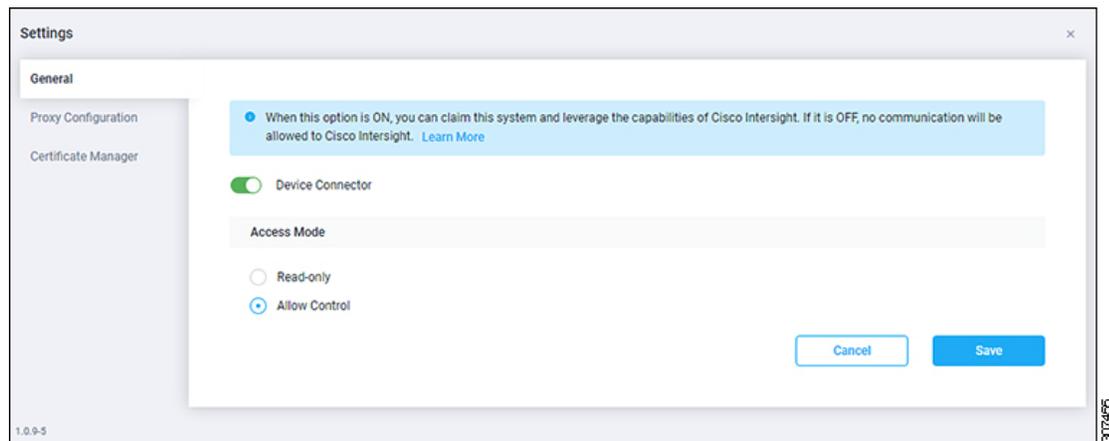
1. Choose **Administration > DCNM Server > Device Connector**.

The Device Connector work pane appears.



2. Click **Settings**.

The **Settings - General** window appears.



- **Device Connector (switch)**

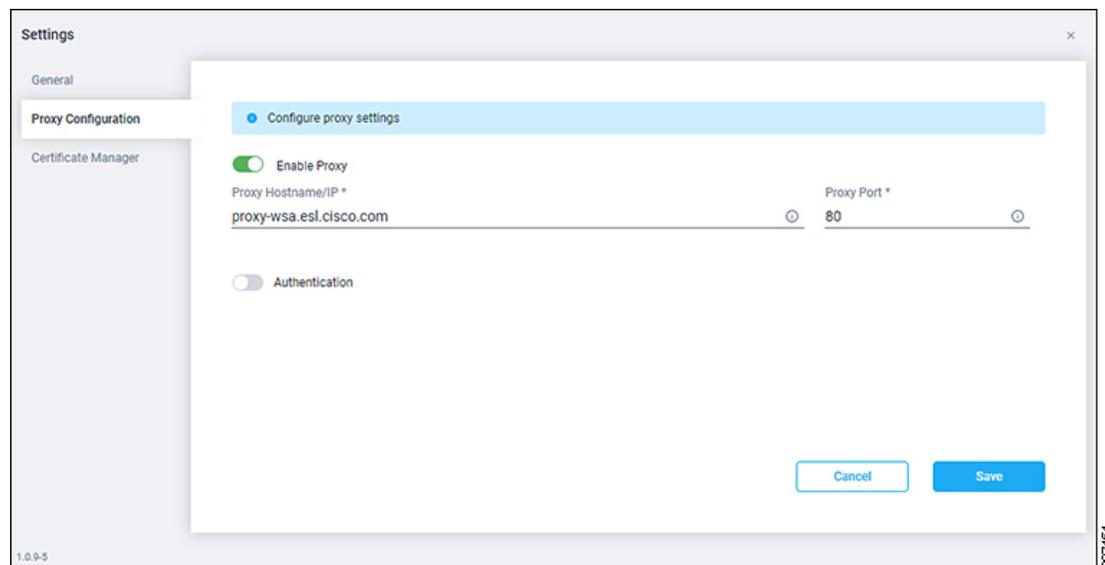
This is the main switch for the Device Connector communication with Cisco Intersight. When the switch is on (green highlight), the Device Connector claims the system and leverages the capabilities of the Cisco Intersight. If the switch is off (gray highlight), no communication can occur between Cisco DCNM and Cisco Intersight.

- **Access Mode**

- **Read-only:** This option ensures that there are no changes to this device from Intersight. For example, actions such as upgrading firmware or a profile deployment is not allowed in the Read-Only mode. However, the actions depend on the features available for a particular system.
- **Allow Control:** This option (selected by default) enables you to perform full read/write operations from the appliance, based on the features available in Cisco Intersight.

3. Set the Device Connector to on (green highlight) and choose **Allow Control**.
4. Click **Proxy Configuration**.

The **Settings - Proxy Configuration** window appears.



- **Enable Proxy (switch)**

Enable HTTPS Proxy to configure the proxy settings.



**Note** Network Insights requires Proxy settings.

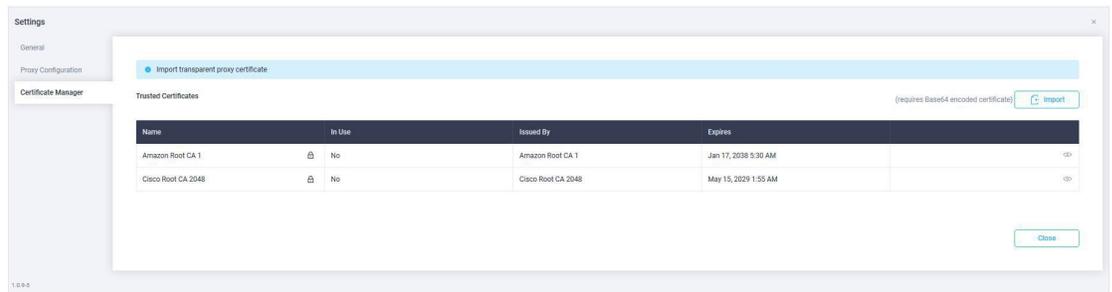
- **Proxy Hostname/IP\* and Proxy Port\*:** Enter a proxy hostname or IP address, and a proxy port number.
- **Authentication (switch)**

Enable proxy access through authentication. When the switch is on (green highlight), authentication to the proxy server is required. If the switch is off (gray highlight), it does not require authentication.

**Username\* and Password:** Enter a user name and password for authentication.

The device connector does not mandate the format of the login credentials, they are passed as-is to the configured HTTP proxy server. The username must be a qualified domain name depending on the configuration of the HTTP proxy server.

5. Enable the proxy (green highlight) and enter a hostname and port number.
6. (Optional) If proxy authentication is required, enable it (green highlight) and enter a username and password.
7. Click **Save**.
8. Click **Certificate Manager**.



The trusted certificates appear in the table.

A list of trusted certificates appears. You can import a valid trusted certificate.

- **Import**

Browse the directory, choose, and import a CA signed certificate.




---

**Note** The imported certificate must be in the **\*.pem (base64encoded)** format.

---

- You can view the list of certificates with the following information:

- **Name**—Common name of the CA certificate.
- **In Use**—Whether the certificate in the trust store is used to successfully verify the remote server.
- **Issued By**—The issuing authority for the certificate.
- **Expires**—The expiry date of the certificate.




---

**Note** You cannot delete bundled certificates.

---

## Management Users




---

**Note** Every time you login to DCNM, the DCNM server fetches information from the ISE server for AAA authentication. The ISE server will not authenticate again, after the first login.

---

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 TACACS 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.

**Note** For IPv6 transport, enter Physical and VIP address for AAA authentication as the order of addresses changes during failover situation.

- 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.

The screenshot shows the Cisco Data Center Network Manager interface for configuring Remote AAA. The left sidebar contains navigation options: Dashboard, Topology, Inventory, Monitor, Configure, and Administration. The main content area is titled 'Administration / Management Users / Remote AAA'. It features several configuration fields and options:

- Auth Mode:** Radio buttons for Local, Radius, TACACS+, Switch, and LDAP (selected).
- Host:** Text input field containing 'ds.cisco.com' and a 'Test...' button.
- Port:** Text input field containing '389'.
- SSL Enabled:** A checkbox that is currently unchecked.
- Base DN:** Text input field containing 'DC=cisco,DC=com'.
- Filter:** Text input field containing '\$userid@cisco.com'.
- Auth Non-Restricted:** A checkbox that is currently unchecked.
- Determine Role By:** Radio buttons for Attribute and Admin Group Map (selected).
- Role Admin Group:** Text input field containing 'dcnm-admins'.
- Map TO DCNM Role:** Text input field containing 'network-admin'.
- Access Map:** An empty text input field.

**Step 2** In the **Host** field, enter either the IPv4 or IPv6 address.

If DNS service is enabled, you can enter DNS address (hostname) of the LDAP server.

**Step 3** In the **Port** field, enter a port number.

Enter 389 for non-SSL; enter 636 for SSL. By default, the port is configured for non-SSL.

**Step 4** Select the **SSL Enabled** check box, if SSL is enabled on the AAA server.

**Note** You must enter **636** in the Port field, and select **SSL Enabled** check box to use LDAP over SSL.

This ensures the integrity and confidentiality of the transferred data by causing the LDAP client to establish a SSL session, before sending the bind or search request.

**Step 5** In the **Base DN** field, enter the base domain name.

The LDAP server searches this domain. You can find the base DN by using the **dsquery.exe user -name <display\_name>** command on the LDAP server.

For example:

```
ldapsrvr# dsquery.exe users -name "John Smith"
```

```
CN=john smith,CN=Users,DC=cisco,DC=com
```

The Base DN is DC=cisco,DC=com.

**Note** Ensure that you enter the elements within the Base DN in the correct order. This specifies the navigation of the application when querying Active Directory.

**Step 6** In the **Filter** field, specify the filter parameters.

These values are used to send a search query to the Active Directory. The LDAP search filter string is limited to a maximum of 128 characters.

For example:

- \$userid@cisco.com

This matches the user principal name.

- CN=\$userid,OU=Employees,OU=Cisco Users

This matches the exact user DN.

**Step 7** Choose an option to determine a role. Select either **Attribute** or **Admin Group Map**.

- **Admin Group Map**: In this mode, DCNM queries LDAP server for a user based on the Base DN and filter. If the user is a part of any user group, the DCNM role will be mapped to that user group.
- **Attribute**: In this mode, DCNM queries for a user attribute. You can select any attribute. When you choose **Attribute**, the **Role Admin Group** field changes to **Role Attributes**.

**Step 8** Enter value for either **Roles Attributes** or **Role Admin Group** field, based on the selection in the previous step.

- If you chose **Admin Group Map**, enter the name of the admin group in the **Role Admin Group** field.
- If you chose **Attribute**, enter the appropriate attribute in the **Attributes** field.

**Step 9** In the **Map to DCNM Role** field, enter the name of the DCNM role that will be mapped to the user.

Generally, **network-admin** or **network-operator** are the most typical roles.

For example:

```
Role Admin Group: dcnm-admins
Map to DCNM Role: network-admin
```

This example maps the Active Directory User Group **dcnm-admins** to the **network-admin** role.

To map multiple Active Directory User Groups to multiple roles, use the following format:

```
Role Admin Group:
Map To DCNM Role: dcnm-admins:network-admin;dcnm-operators:network-operator
```

Note that **Role Admin Group** is blank, and **Map To DCNM Role** contains two entries delimited by a semicolon.

**Step 10** In the **Access Map** field, enter the Role Based Access Control (RBAC) device group to be mapped to the user.

**Step 11** Click **Test** to verify the configuration. The Test AAA Server window appears.

**Step 12** Enter a valid **Username** and **Password** in the Test AAA Server window.

If the configuration is correct, the following message is displayed.

```
Authentication succeeded.
The cisco-av-pair should return 'role=network-admin' if this user needs to
see the DCNM Admin pages. 'SME' roles will allow SME page access. All other
roles - even if defined on the switches - will be treated
as network operator.
```

This message is displayed regardless of 'Role Admin Group' or 'Attribute' mode. It implies that Cisco DCNM can query your Active Directory, the groups, and the roles are configured correctly.

If the test fails, the LDAP Authentication Failed message is displayed.

**Warning** Don't save the configuration unless the test is successful. You cannot access DCNM if you save incorrect configurations.

**Step 13** Click **Apply Changes** icon (located in the right top corner of the screen) to save the configuration.

- Step 14** Restart the DCNM SAN service.
- For Windows – On your system navigate to **Computer Management > Services and Applications > Services**. Locate and right click on the DCNM application. Select **Stop**. After a minute, right click on the DCNM application and select **Start** to restart the DCNM SAN service.
  - For Linux – Go to `/etc/init.d/FMServer.restart` and hit return key to restart DCNM SAN service.
- 

## 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.
- Note** All special characters, except SPACE is allowed in 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

You can select specific groups or fabrics that local users can access. This restricts local users from accessing specific groups or fabrics for which they have not been provided access. To do this, 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 specific groups or fabrics that the user can access and click **Apply**.

The screenshot shows the Cisco Data Center Network Manager interface. The breadcrumb navigation is Administration / Management Users / Local. The 'Local Users' section contains a table with the following data:

| User Name                                | Role          | Access      | Password Expiration Status |
|--|---------------|-------------|----------------------------|
| <input type="checkbox"/> admin           | network-admin | Data Center | Password never expires.    |
| <input type="checkbox"/> poap            | network-admin | Data Center | Password never expires.    |
| <input type="checkbox"/> root            | network-admin | Data Center | Password never expires.    |
| <input checked="" type="checkbox"/> john | network-admin | Data Center | Password never expires.    |

The 'User Access' dialog box is open, showing a list of access groups with checkboxes:

- Cloud-Connect
  - CSR-Azure
  - CSR-OnPrem
  - ext-fabric5
  - site2
- ext
- s1
- services-setup
- john-fx2
- fx2
- Default\_LAN

The 'Apply' button is highlighted.

## 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.



**Note** To collect Performance Manager data, ICMP ping must be enabled between the switch and DCNM server. Set `pm.skip.checkPingAndManageable` server property to true and then restart the DCNM. Choose Web UI > **Administration** > **DCNM Server** > **Server Properties** to set the server property.

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 Manager. The Performance Manager has to be restarted for any new setting to take effect.

## 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:
- From the **Type** drop-down list, choose an event type.
  - Check the **Storage Ports Only** check box to select only the storage ports.
  - From the **Minimum Severity** drop-down list, select the severity level of the messages to receive.
  - Click **Add** to add the notification.
- If you select **Syslog**, then:
- In the **Facility** list, select the syslog facility.
  - Specify the syslog **Type**.
  - In the **Description Regex** field, specify a description that matches with the event description.
  - From the **Minimum Severity** drop-down list, select the severity level of the messages to receive.
  - 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.  
  
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 use selects **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 381](#)
- [Validate Credentials, on page 381](#)

- [Clear Switch Credentials, on page 381](#)

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

# Applications

Cisco Data Center Network Manager (DCNM) uses the application framework to host various plugins and microservices to support operations and related features in Cisco DCNM.

The Applications Framework provides the following features:

- An infrastructure for hosting applications that require more system resources as the scale of the network increases.
- An independent application development-deployment-management lifecycle for applications.

Cisco DCNM Applications Framework supports two modes namely clustered mode and unclustered mode. In clustered mode, the compute nodes are clustered together whereas in the latter only the DCNM server nodes namely the active/standby exist. Most of the applications for ex: Network Insights require clustered setup to be ready before they can be uploaded and deployed using DCNM Applications Framework.

- [Cisco DCNM in Unclustered Mode, on page 383](#)
- [Cisco DCNM in Clustered Mode, on page 384](#)
- [Installing and Deploying Applications, on page 401](#)
- [Application Framework User Interface, on page 406](#)
- [Catalog, on page 407](#)
- [Compute, on page 410](#)
- [Preferences, on page 411](#)
- [Failure Scenario, on page 418](#)

## Cisco DCNM in Unclustered Mode

From Cisco DCNM Release 11.0(1), the unclustered mode is the default deployment mode in both Standalone and Native HA environment. In this mode, the Cisco DCNM runs some of its internal services as containers, also.

- Endpoint Locator is running as a container application, from Cisco DCNM Release 11.1(1).
- Configuration Compliance service is a container application, from Cisco DCNM Release 11.0(1).
- Virtual Machine Manager (VMM) is also a container application, from Cisco DCNM Release 11.0(1)

Cisco DCNM leverages resources from the Standby node for running some containers applications. The Cisco DCNM Active and Standby nodes work together to extend resources to the overall functionality and deployment

of DCNM and its applications. However, it has limited resources to run some of the advanced applications and to extend the system to deploy more applications delivered through the Cisco AppCenter. For example, you cannot deploy the Network Insights applications that are downloaded from the Cisco AppCenter, for production, in unclustered mode.

To install and deploy applications, see [Installing and Deploying Applications, on page 401](#).

For best practices and recommended deployments for IP address configurations of all interfaces of the Cisco DCNM and Compute nodes, see *Best Practices for Deploying Cisco DCNM and Computes* in *Cisco DCNM Installation Guide*, for your deployment type.

## Cisco DCNM in Clustered Mode

By default, the clustered mode is not enabled on the Cisco DCNM deployments. Enable the cluster mode after you deploy the Cisco DCNM Server. In a clustered mode, the Cisco DCNM Server with more compute nodes provides an architecture to expand resources, as you deploy more applications.

Compute nodes are scale out application hosting nodes that run resource-intensive services to provide services to the larger Fabric. When compute nodes are added, all services that are containers, run only on these nodes. This includes Config Compliance, Endpoint Locator, and Virtual Machine Manager. The Elasticsearch time series database for these features runs on compute nodes in clustered mode. In the clustered mode deployment, the DCNM Servers do not run containerized applications. All applications that work in unclustered mode works in the clustered mode, also.



---

**Note** The clustered mode is not supported on Cisco DCNM for Media Controller deployment.

---

From Cisco DCNM Release 11.1(1), in a Native HA setup, 80 switches with Endpoint Locator, Virtual Machine Manager, config compliance are validated in the unclustered mode. For a network exceeding 80 switches, with these features in a given Cisco DCNM instance (maximum qualified scale is 256 switches), we recommend that you enable clustered mode.

While the Cisco DCNM core functionalities only run on the Native HA nodes, addition of compute nodes beyond 80 switches is to build a scale-out model for Cisco DCNM and related services.

From Release 11.2(1), you can configure IPv6 address for Network Management for compute clusters. However, DCNM does not support IPv6-address for containers, and must connect to DCNM using only IPv4 address only.

For best practices and recommended deployments for IP address configurations of all interfaces of the Cisco DCNM and Compute nodes, see *Best Practices for Deploying Cisco DCNM and Computes* in *Cisco DCNM Installation Guide*, for your deployment type.

## Requirements for Cisco DCNM Clustered Mode



---

**Note** We recommend that you install the Cisco DCNM in the Native HA mode.

---

## Cisco DCNM LAN Fabric Deployment Without Network Insights (NI)



**Note** For information about various system requirements for proper functioning of Cisco DCNM LAN Fabric deployment, see [System Requirements](#).

Refer to [Network Insights User guide](#) for sizing information for Cisco DCNM LAN Deployment with Network Insights (NI).

To see the verified scale limits for Cisco DCNM 11.4(1) for managing LAN Fabric deployments, see [Verified Scale Limits for Cisco DCNM LAN Fabric Deployment](#).

**Table 19: Upto 80 Switches**

| Node     | CPU Deployment Mode | CPU      | Memory | Storage  | Network |
|----------|---------------------|----------|--------|----------|---------|
| DCNM     | OVA/ISO             | 16 vCPUs | 32G    | 500G HDD | 3xNIC   |
| Computes | NA                  | —        | —      | —        | —       |

**Table 20: 81–250 Switches**

| Node                      | CPU Deployment Mode | CPU      | Memory | Storage  | Network |
|---------------------------|---------------------|----------|--------|----------|---------|
| DCNM                      | OVA/ISO             | 16 vCPUs | 32G    | 500G HDD | 3xNIC   |
| Computes x 3 <sup>1</sup> | OVA/ISO             | 16 vCPUs | 64G    | 500G HDD | 3xNIC   |

<sup>1</sup> Cisco DCNM must be deployed with Compute cluster nodes to use NI applications.

## Cisco DCNM LAN Deployment with NIA and NIR Software Telemetry



**Note** We recommend that you install the Cisco DCNM in the Native HA mode.

**Table 21: Upto 80 Switches**

| Node         | CPU Deployment Mode | CPU      | Memory | Storage  | Network |
|--------------|---------------------|----------|--------|----------|---------|
| DCNM         | OVA/ISO             | 16 vCPUs | 32G    | 500G HDD | 3xNIC   |
| Computes x 3 | OVA/ISO             | 16 vCPUs | 64G    | 500G HDD | 3xNIC   |

Table 22: 81–250 Switches

| Node         | CPU Deployment Mode | CPU      | Memory | Storage    | Network            |
|--------------|---------------------|----------|--------|------------|--------------------|
| DCNM         | OVA/ISO             | 16 vCPUs | 32G    | 500G HDD   | 3xNIC              |
| Computes x 3 | ISO                 | 32 vCPUs | 256G   | 2.4-TB HDD | 3xNIC <sup>2</sup> |

<sup>2</sup> Network card: Quad-port 10/25G

### Subnet Requirements

In general, Eth0 of the Cisco DCNM server is used for Management, Eth1 is used to connect Cisco DCNM Out-Of-Band with switch management, and eth2 is used for In-Band front panel connectivity of Cisco DCNM. The same concept extends into compute nodes as well. Some services in clustered mode have other requirements. Some services require a switch to reach into Cisco DCNM. For example, Route Reflector to Endpoint Locator connection or switch streaming telemetry into the Telemetry receiver service of the application require a switch to reach DCNM. This IP address needs to remain sticky during all failure scenarios. For this purpose, an IP pool must be provided to Cisco DCNM at the time of cluster configuration for both out-of-band and In-Band subnets.

### Telemetry NTP Requirements

For telemetry to work correctly, the Cisco Nexus 9000 switches and Cisco DCNM must be time that is synchronized. Cisco DCNM telemetry manager does the required NTP configuration as part of enablement. If there is a use-case to change the NTP server configuration manually on the switches ensure that the DCNM and the switches are time synchronized, always. To set up telemetry network configuration, see [In-Band Telemetry Network and NTP Configuration, on page 394](#).

## Installing a Cisco DCNM Compute



**Note** With Native HA installations, ensure that the HA status is **OK** before DCNM is converted to cluster mode.

A Cisco DCNM Compute can be installed using an ISO or OVA of a regular Cisco DCNM image. It can be deployed directly on a bare metal using an ISO or a VM using the OVA. After you deploy Cisco DCNM, using the DCNM web installer, choose **Compute** as the install mode for Cisco DCNM Compute nodes. On a Compute VM, you will not find DCNM processes or postgres database; it runs a minimum set of services that are required to provision and monitor applications.

## Networking Policies for OVA Installation

For each compute OVA installation, ensure that the following networking policies are applied for the corresponding vSwitches of host:

- Log on to the vCenter.
- Click on the Host on which the computes OVA is running.

- Click **Configuration > Networking**.
- Right click on the port groups corresponding to the eth1 and eth2, and select **Edit Settings**.  
The **VM Network - Edit Settings** is displayed.
- In Security settings, for **Promiscuous** mode, select **Accepted**.
- If a DVS Port-group is attached to the compute VM, configure these settings on the **vCenter > Networking > Port-Group**. If a normal vSwitch port-group is used, configure these settings on **Configuration > Networking > port-group** on each of the Compute's hosts.

**Figure 4: Security Settings for vSwitch Port-Group**

VM Network - Edit Settings

|                      |                     |  |        |
|----------------------|---------------------|--|--------|
| Properties           |                     |  |        |
| <b>Security</b>      | Promiscuous mode    | <input checked="" type="checkbox"/> Override | Accept |
| Traffic shaping      | MAC address changes | <input checked="" type="checkbox"/> Override | Accept |
| Teaming and failover | Forged transmits    | <input checked="" type="checkbox"/> Override | Accept |

CANCEL OK

Figure 5: Security Settings for DVSwitch Port-Group

OobFabric - Edit Settings

|                      |                     |        |
|----------------------|---------------------|--------|
| General              |                     |        |
| Advanced             | Promiscuous mode    | Accept |
| VLAN                 | MAC address changes | Accept |
| <b>Security</b>      | Forged transmits    | Accept |
| Teaming and failover |                     |        |
| Traffic shaping      |                     |        |
| Monitoring           |                     |        |
| Miscellaneous        |                     |        |

CANCEL OK



**Note** Ensure that you repeat this procedure on all the hosts, where a Compute OVA is running.

## Enabling the Compute Cluster

From Cisco DCNM Release 11.2(1), you must enable the compute cluster mode using the **appmgr afw** command.



**Note** Ensure that you enable Compute Cluster before you install applications. The applications that are installed via the AppCenter will not work if you enable the compute cluster after installing the applications.



**Note** The services are down until the configuration is complete. Ensure that the session is active while configuration is in progress.

Use the following command to enable the compute cluster.

```
appmgr afw config-cluster
[--ewpool<InterApp-Subnet>]--oobpool<OutOfBand-Subnet>--ibpool<Inband-Subnet>--computeip<compute-IP>
```

Where:

- **ewpool**: specifies the east-west pool subnet; for inter-service connectivity.

This field is optional, if the inter-application subnet is specified during Cisco DCNM installation for your deployment type. These addresses are not used directly between the computes, or to communicate with another node. These are used by containers to communicate with each other. This subnet must be minimum of /24 (256 addresses) and a maximum of a /20 (4096 addresses).

This field is optional if the Inter-app subnet is specified during Cisco DCNM deployment installation.

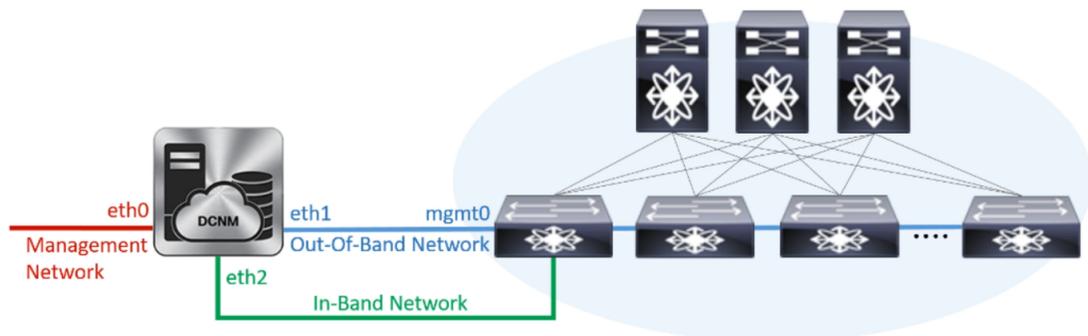
- **oobpool**: specifies the out-of-band pool; a smaller prefix of available IP addresses from the eth1 subnet. For example: Use 10.1.1.240/28 if the eth1 subnet was configured as 10.1.1.0/24 during installation.

This subnet must be a minimum of /28 (16 addresses) and maximum of /24 (256 addresses). It should also be longer than the east-west pool. This subnet is assigned to containers, to communicate with the switches.

- **ibpool**: specifies the in-band pool; a smaller prefix of available IP addresses eth2 subnet. For example: Use 11.1.1.240/28 if the eth2 subnet was configured as 11.1.1.0/24 during installation.

This subnet must be a minimum of /28 (16 addresses) and maximum of /24 (256 addresses). It should also be longer than the east-west pool. This subnet is assigned to containers, to communicate with the switches.

- **computeip**: specifies the dcnm-mgmt network (eth0) interface IP address of the first compute node added to the cluster. This compute is added into the cluster as part of this command process and is used to migrate application data from DCNM servers to computes.



| Compute IP Address                  | In-Band Interface | Out-Band Interface | Status     | Memory | Disk | Uptime                  |
|-------------------------------------|-------------------|--------------------|------------|--------|------|-------------------------|
| <input type="radio"/> 172.28.12.205 | eth2              | eth1               | Joined     | 80%    | 80%  | -- Hrs : 4 Min : 17 Sec |
| <input type="radio"/> 172.28.12.210 | NA                | NA                 | Discovered |        |      |                         |
| <input type="radio"/> 172.28.12.206 | NA                | NA                 | Discovered |        |      |                         |

The other two computes are Discovered automatically, and is displayed on the Cisco DCNM Web UI > Applications > Compute.

The In-Band or out-of-band pools are used by services to connect with switches as required. The IP addresses from these pools must be available for configuration.



**Note** To add computes to the cluster mode, see [Adding Computes into the Cluster Mode, on page 390](#).

## Managing Application Network Pools

When you alter the eth1 or eth2 interface subnets, the corresponding oob pool and inband pool must be modified to match the new configuration. Network Insights and Endpoint Locator applications use the IP addresses from the Out-of-Band and In-Band pools.

To modify the IP addresses that are assigned to services running in the compute cluster, use the following command:



**Note** The inband or out-of-band pools are used by applications to connect with Cisco Nexus Switches. Hence, the IP addresses from these pools must be available and free.

```
appmgr afw config-pool [--ewpool <InterApp-Subnet>] --oobpool <OutOfBand-Subnet> --ibpool <Inband-Subnet> --computeip <compute-IP>
```

Where:

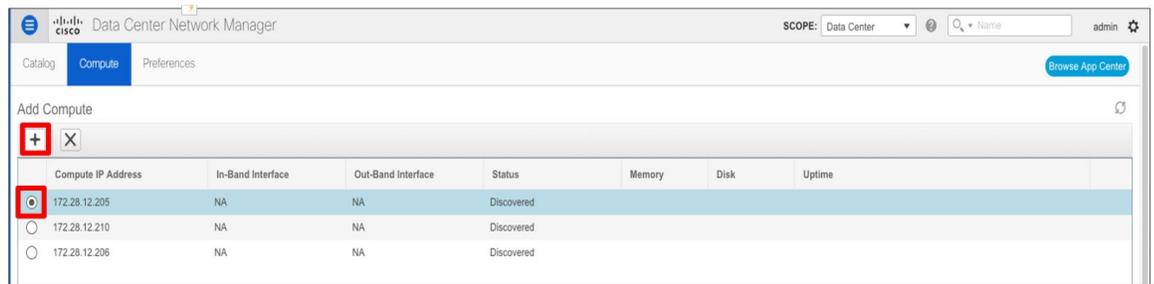
- **ewpool**: specifies the east west pool subnet; for inter-service connectivity.  
The network mask ranges from 20 to 24. These addresses aren't used directly between the computes, or to communicate with another node. These are used by containers to communicate with each other.
- **oobpool**: specifies the out-of-band pool; a smaller prefix of available IP Addresses from eth1 subnet.  
The network mask ranges from 24 to 28.
- **ibpool**: specifies the inband pool; a smaller prefix of available IP addresses from eth2 subnet.  
The network mask ranges from 24 to 28.
- **ipv6oobpool**: specifies the out-of-band IPv6 pool; a smaller prefix of available IPv6 addresses from eth1 subnet.  
If IPv6 is enabled, these addresses are required on both inband and out-of-band subnet.  
The network mask ranges from 112 to 124.
- **ipv6ibpool**: specifies the inband IPv6 pool; a smaller prefix of available IPv6 addresses from eth2 subnet.  
If IPv6 is enabled, these addresses are required on both inband and out-of-band subnet.  
The network mask ranges from 112 to 124.

## Adding Computes into the Cluster Mode

To add computes into the cluster mode from Cisco DCNM Web UI, perform the following steps:

### Procedure

- Step 1** Choose **Applications > Compute**.  
The Compute tab displays the computes enabled on the Cisco DCNM.
- Step 2** Select a Compute node which is in **Discovered** status. Click the **Add Compute (+)** icon.

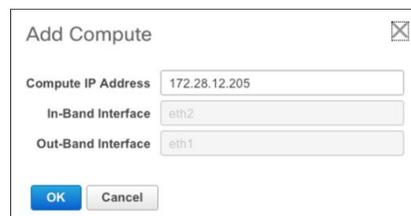


- While using Compute, ensure that Cisco DCNM GUI shows nodes as Joined.
- Offline indicates some connectivity issues, therefore no applications are running on Offline Computes.
- Failed indicates that the compute node could not join the cluster.
- Health indicates the amount of free memory and disk on the Compute node. The Watch Tower application provides more detailed statistics.
- Cisco DCNM 3 node cluster is resilient to single node failure only.
- If the Performance Manager was stopped during or after the inline upgrade and after all the computes have changed to Joined, you must restart the Performance Manager.

The Compute window allows you to monitor the health of computes. The health essentially indicates the amount of memory that is left in the compute, this is based on applications that are enabled. If a Compute is not properly communicating with the DCNM Server, the status of the Compute appears as Offline, and no applications are running on Offline Computes.

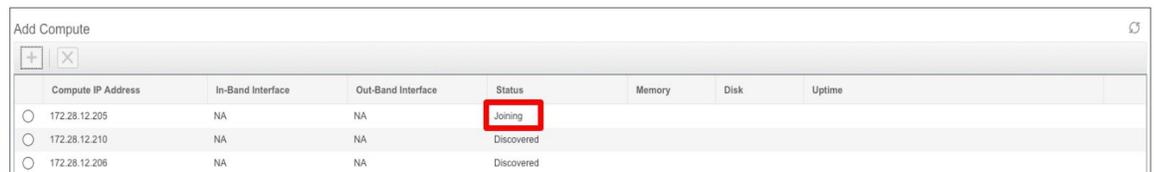
**Step 3** In the **Add Compute** dialog box, view the **Compute IP Address**, **In-Band Interface**, and the **Out-Band Interface** values.

**Note** The interface value for each compute node is configured by using the `appmgr afw config-cluster` command.



**Step 4** Click **OK**.

The Status for that Compute IP changes to **Joining**.



Wait until the Compute IP status shows **Joined**.

| Add Compute                         |                   |                    |            |        |      |                         |  |
|-------------------------------------|-------------------|--------------------|------------|--------|------|-------------------------|--|
| Compute IP Address                  | In-Band Interface | Out-Band Interface | Status     | Memory | Disk | Uptime                  |  |
| <input type="radio"/> 172.28.12.205 | eth2              | eth1               | Joined     | 60%    | 99%  | -- Hrs : 4 Min : 17 Sec |  |
| <input type="radio"/> 172.28.12.210 | NA                | NA                 | Discovered |        |      |                         |  |
| <input type="radio"/> 172.28.12.206 | NA                | NA                 | Discovered |        |      |                         |  |

**Step 5** Repeat the above steps to add the remaining compute node.

All the Computes appear as **Joined**.

| Add Compute                         |                   |                    |        |        |      |                           |  |
|-------------------------------------|-------------------|--------------------|--------|--------|------|---------------------------|--|
| Compute IP Address                  | In-Band Interface | Out-Band Interface | Status | Memory | Disk | Uptime                    |  |
| <input type="radio"/> 172.28.12.205 | eth2              | eth1               | Joined | 40%    | 99%  | 183 Hrs : 15 Min : 41 Sec |  |
| <input type="radio"/> 172.28.12.210 | eth2              | eth1               | Joined | 57%    | 99%  | -- Hrs : 4 Min : 9 Sec    |  |
| <input type="radio"/> 172.28.12.206 | eth2              | eth1               | Joined | 58%    | 99%  | -- Hrs : 2 Min : 18 Sec   |  |

**Note** When you install compute as a virtual machine on the VMware platform, vSwitch or DV switch port groups associated eth1 and eth2 must allow for packets that are associated with Mac address other than eth1 and eth2 to be forwarded.

## Transitioning Compute Nodes

### Transitioning Compute nodes from VM to Service Engine

To transition Cisco DCNM Compute Nodes from VMs to Applications Services Engine using the Cisco DCNM Web Client, perform the following steps:

#### Before you begin

- Ensure that Cisco DCNM Web Client is functioning.
- On the Cisco DCNM Web Client > **Applications** > **Compute**, all the Compute nodes must be in **Joined** state.

#### Procedure

**Step 1** Choose **Applications** > **Compute**.

For example, let us indicate the three Compute nodes as **compute1**, **compute2**, and **compute3**.

**Step 2** Open the vCenter Server application and connect to the vCenter Server with your vCenter user credentials.

**Step 3** Navigate to **Home** > **Inventory** > **Hosts and Clusters** and identify the VM on which the DCNM Compute nodes are deployed.

**Step 4** For **compute1**, make a note of the configurations and setup details provided during installation.

**Step 5** Turn off **compute1**. Right click on the VM, select **Power off**.

On the **Web UI > Applications > Compute**, the status of **compute1** shows **Offline**.

- Step 6** Using the configuration details of the compute node VM, install the compute node on Cisco Applications Services Engine.  
For instructions, refer to [Installing DCNM Compute Node on Cisco ASE](#).
- Step 7** Launch the Web UI, and choose **Applications > Compute**.  
The newly added compute automatically joins the cluster. The status of **compute1** changes from **Offline** → **Joining** → **Joined**.
- Step 8** Repeat Steps [Step 4, on page 392](#) to [Step 7, on page 393](#) on **compute2** and **compute3** compute nodes.  
After completion, all the Compute nodes on **Web UI > Applications > Compute** are in the **Joined** state.  
All are Compute nodes are successfully hosted on the Cisco Applications Services Engine.

---

## Transitioning Compute nodes from Service Engine to VM

To transition Cisco DCNM Compute Nodes from Applications Services Engine to VMs using the Cisco DCNM Web Client, perform the following steps:

### Before you begin

- Ensure that Cisco DCNM Web Client is functioning.
- On the Cisco DCNM **Web Client > Applications > Compute**, all the Compute nodes must be in **Joined** state.

### Procedure

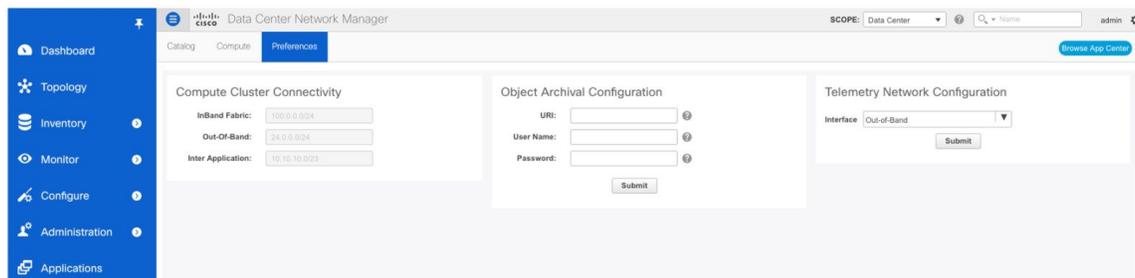
---

- Step 1** Choose **Applications > Compute**.  
For example, let us indicate the three Compute nodes as **compute1**, **compute2**, and **compute3**.
- Step 2** On the Cisco Applications Server console, for **compute1**, make a note of the configurations and setup details provided during installation.
- Step 3** Power off the Applications Service Engine to turn off **compute1**.  
On the Cisco DCNM **Web UI > Applications > Compute**, the status of **compute1** shows **Offline**.
- Step 4** Using the configuration details of the compute node on Applications Service Engine, install the compute node on the VM.  
For instructions, refer to [Installing DCNM on ISO Virtual Appliance](#).
- Step 5** Launch the Web UI, and choose **Applications > Compute**.  
The newly added compute automatically joins the cluster. The status of **compute1** changes from **Offline** → **Joining** → **Joined**.
- Step 6** Repeat Steps [#unique\\_290 unique\\_290\\_Connect\\_42\\_make-note-setup](#) to [#unique\\_290 unique\\_290\\_Connect\\_42\\_join-cluster](#) on **compute2** and **compute3** compute nodes.

After completion, all the Compute nodes on **Web UI > Applications > Compute** are in the **Joined** state. All are Compute nodes are successfully hosted on the VMs.

## Preferences

This tab is relevant to the cluster mode of deployment, where the application instances are placed. This tab enables you to compute cluster connectivity and configure the Cluster Connectivity preferences.



### Object Archival Configuration

The NIA application collects tech support logs for all switches in Fabric, and determines the advisory, based on the data. The logs are saved on the Cisco DCNM server for further analysis or troubleshooting. If you need to download these logs before their life span ends or to create some space on the DCNM server, you can move the logs to a remote server.

In the **URI** field, enter the relative path to the archive folder, in the format `host[:port]/[path to archive]`. Enter the username and password to access the URI, in the **username** and **Password** field. Click **Submit** to configure the remote server.

## In-Band Telemetry Network and NTP Configuration

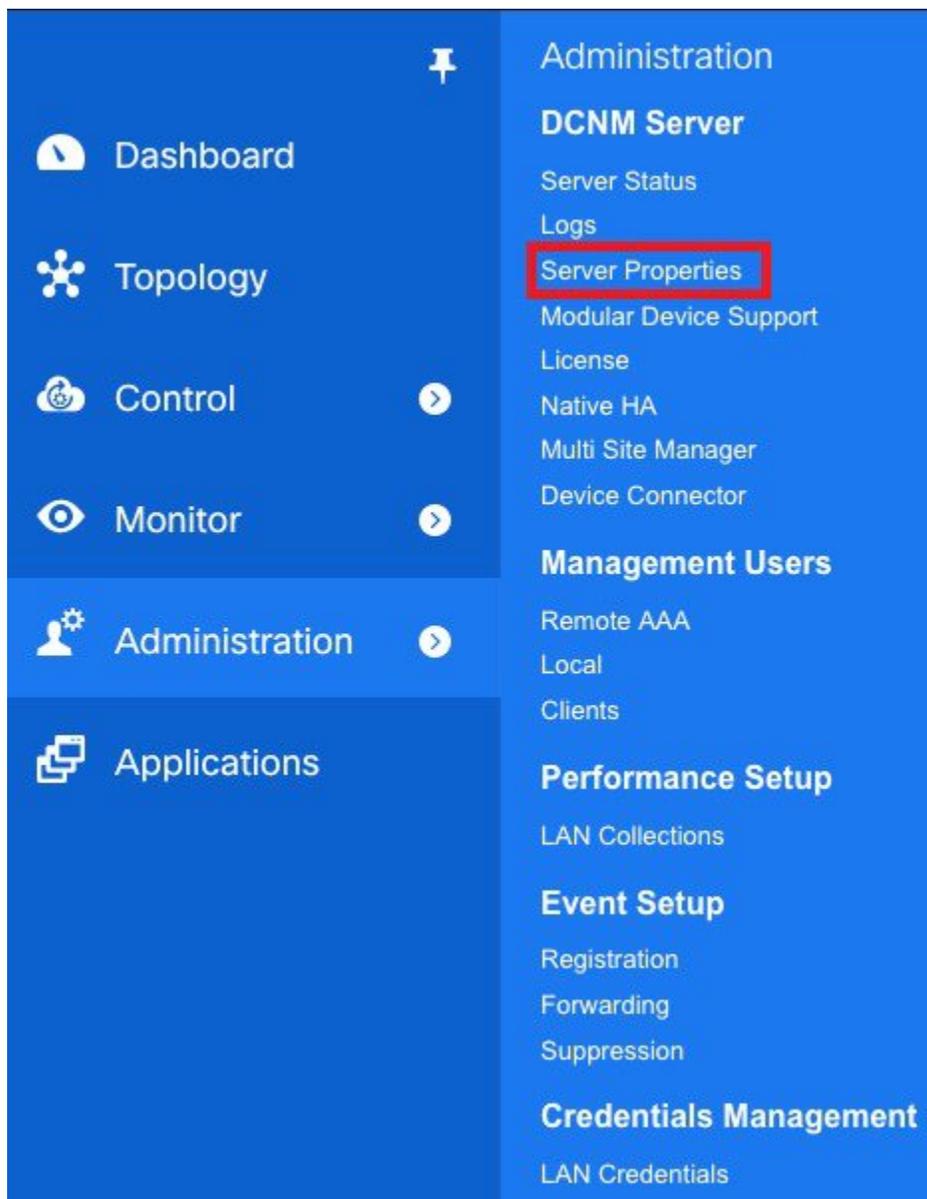
### Before you begin

If Network Insights Resources (NIR) application is running, you must disable it on all the fabrics before you begin the procedure.

### Procedure

#### Step 1

Choose **Administration > DCNM Server > Server Properties**.



- Step 2** In the # **Template Properties** area, locate the **template.in\_use.check** field.
- Step 3** Set the **template.in\_use.check** field to **true**. Click **Apply Changes**.

The screenshot shows the Cisco Data Center Network Manager (DCNM) interface. The left sidebar contains navigation options: Dashboard, Topology, Control, Monitor, Administration, and Applications. The main content area is titled "Administration / DCNM Server / Server Properties". It displays several configuration fields:

- `# periodically restart ECT baseline training after number of days`  
`# (Default is 28)`  
`san.telemetry.train.reset` 28
- `# Template Properties`
  - `template.in_use.check` true (highlighted with a red box)
  - `template.use_cache` true
  - `template.server_validation_check` false
  - `template.server_validation_continue_on_error` false
- `#Image Management Property`  
`FILE_SELECTION_FILTER` true

An "Apply Changes" button is located in the top right corner, also highlighted with a red box.

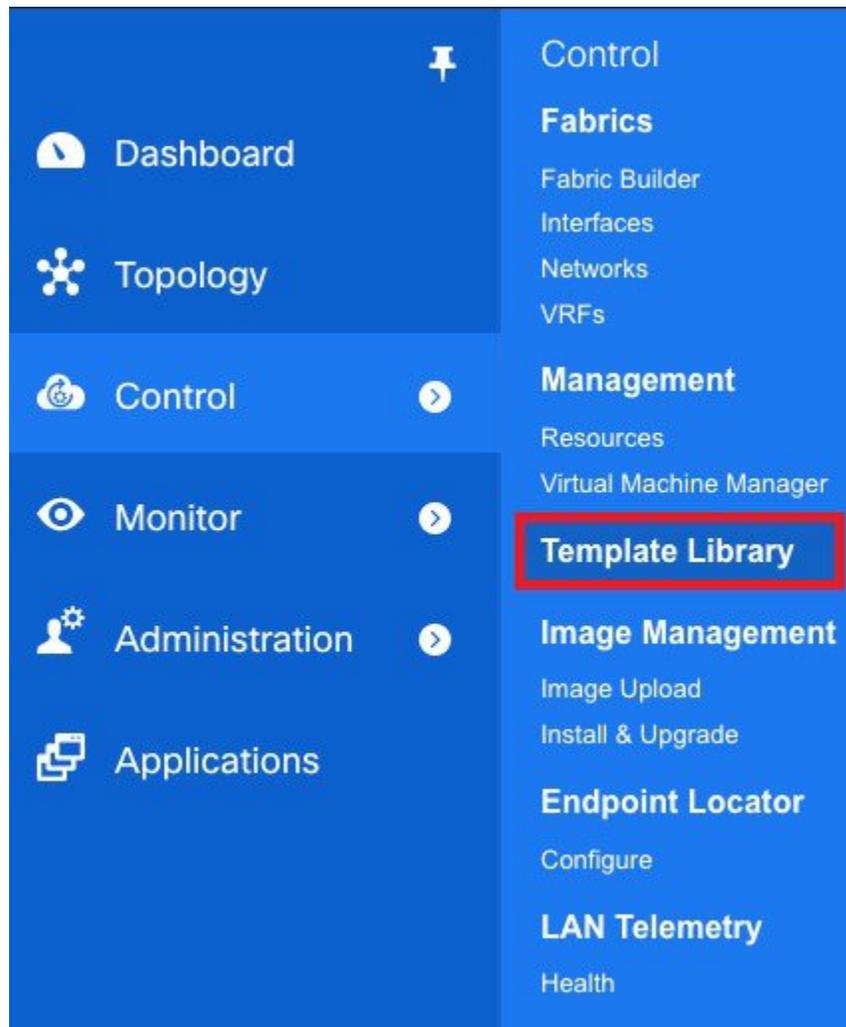
**Step 4** Click **OK** to confirm that this property change does not require restarting the Cisco DCNM server.

The dialog box contains the following text:

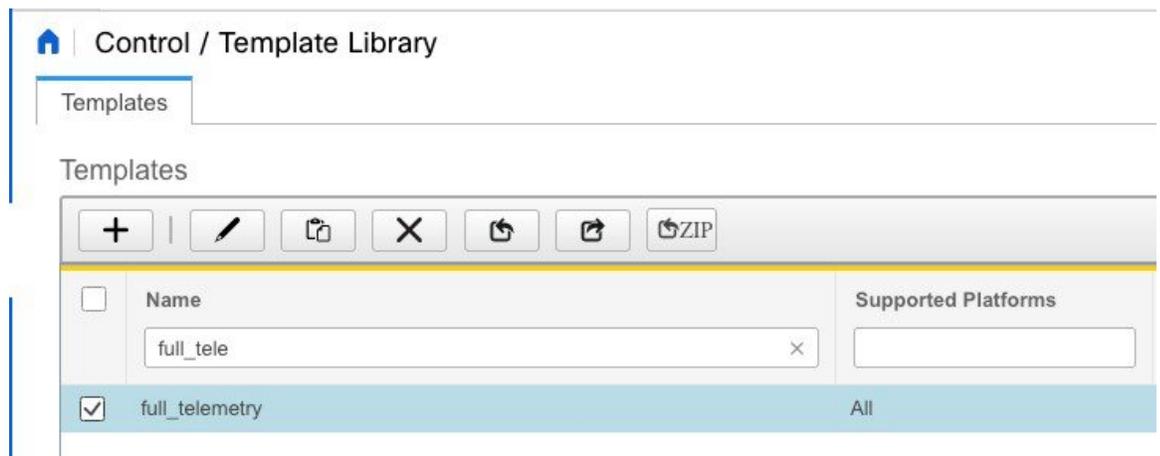
Please restart DCNM SAN service if you update properties other than EMC Callhome properties(`server.callhome.enable`, `server.callhome.xmlDir`), Event Forwarding properties (`server.forward.event.enable`), Template properties (`template.in_use.check` property), Event Registration properties(`syslog.disable`) or `fabric.enableNpvDiscovery` properties. (Note: please restart all instances if federation is deployed). Please resync vmm if you updated `vmm.resync.timer`

An "Ok" button is located at the bottom right of the dialog.

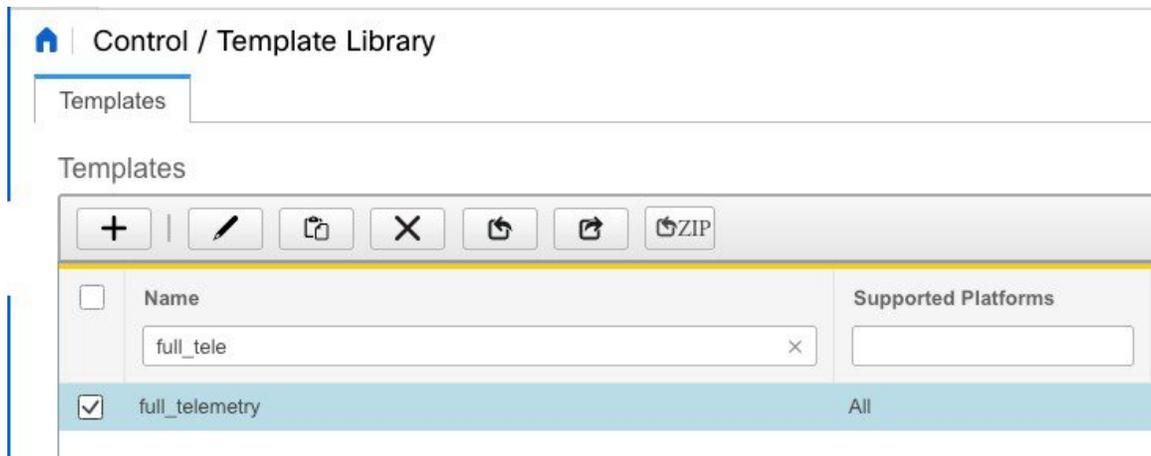
**Step 5** Choose **Control > Template Library**.



**Step 6** Locate and select **full\_telemetry** template, and click **Edit** icon.



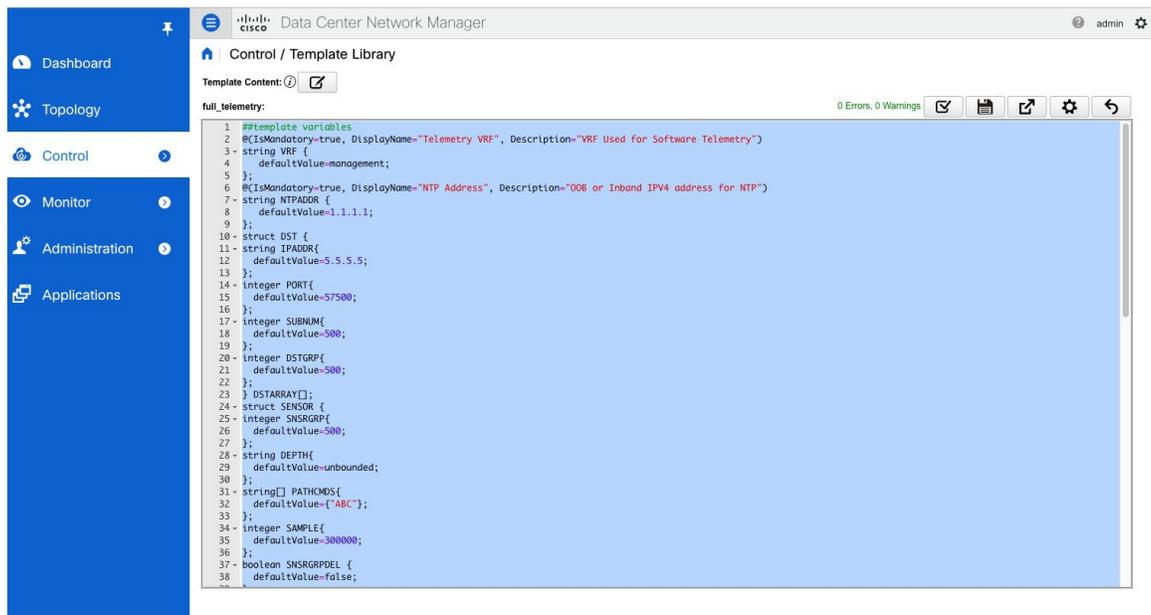
**Step 7** Click **Yes** to cancel any pending scheduled template jobs associated with this template.



The **Template Content** window is displayed.

### Step 8

Locate the **##template content** in the text section.



### Step 9

Change the **use-vrf \$\$VRFS\$** to **use-vrf management**.

```

##template content
if ( $$NTPADDR$$ != "" ) {
  if ( $$SVRF$$ != "management" ) {
ntp server $$NTPADDR$$ prefer use-vrf $$SVRF$$
  }
  else {
ntp server $$NTPADDR$$ prefer use-vrf management
  }
}
telemetry

##template content
if ( $$NTPADDR$$ != "" ) {
  if ( $$SVRF$$ != "management" ) {
ntp server $$NTPADDR$$ prefer use-vrf management
  }
  else {
ntp server $$NTPADDR$$ prefer use-vrf management
  }
}

```

Click **Save** icon. Click **OK** on the **Validation Result** window.

- Step 10** Choose **Applications > Preferences**.  
The **Preferences** tab is displayed.

The screenshot shows the Cisco Data Center Network Manager interface. The 'Preferences' tab is active. Under the 'Telemetry Network Configuration' section, the 'Interface' dropdown menu is open, showing 'Out-of-Band' as the selected option. Other sections like 'Compute Cluster Connectivity' and 'Object Archival Configuration' are also visible.

- Step 11** From the **Interface** drop-down list, select **In-Band**.

This is a close-up of the 'Telemetry Network Configuration' form. The 'Interface' dropdown menu is open, and 'In-Band' is selected. The 'Out-of-Band' option is also visible in the list.

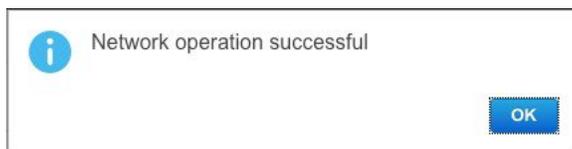
- Step 12** In the **VRF** field, enter the VRF name, which has connectivity to the Cisco DCNM In-Band interface.  
The VRF value is set to default. You can also change this to any other VRF assuming that the switches are configured with the VRF and there is availability to the in-band interface of the Cisco DCNM through that VRF.

This is a close-up of the 'Telemetry Network Configuration' form. The 'Interface' dropdown menu is set to 'In-Band' and the 'VRF' field is set to 'default'. A 'Submit' button is visible at the bottom.

**Note** If the Cisco NIR telemetry is already enabled for some fabrics, disable the Telemetry on all the enabled fabrics before you modify the Telemetry Network Configuration. You can enable NIR telemetry after you modify the telemetry network to begin streaming of data through the in-band interface.

Click **Submit**.

**Step 13** Click **OK** to confirm the network operation.



The telemetry network configuration is now complete.

### What to do next

To download and install the Network Insights applications, see [Catalog, on page 407](#).

After the telemetry is enabled via Network Insights Resources (NIR) application, the telemetry manager in Cisco DCNM will push the necessary NTP server configurations to the switches.

## Telemetry Network and NTP Requirements

For the Network Insights Resource (NIR) application, a UTR micro-services running inside the NIR receives the telemetry traffic from the switches either through Out-Of-Band (Eth1) or In-Band (Eth2) interface. By default, the telemetry is configured, and is streaming via the Out-Of-Band interface. You can choose to change it to In-Band interface as well.

For the Cisco Network Insights for Resources (NIR) Release 2.1 and later, and flow telemetry, **feature lldp** command is one of the required configuration.

Cisco DCNM pushes **feature lldp** on the switches only for the Easy Fabric deployments, that is, for the eBGP routed fabric or VXLAN EVPN fabric. Therefore, NIR users need to enable **feature lldp** on all the switches in the following scenarios:

- External fabric in Monitored or Managed Mode

### Telemetry Using Out-of-band (OOB) Network

By default, the telemetry data is streamed through the management interface of the switches to the Cisco DCNM OOB network eth1 interface. This is a global configuration for all fabrics in Cisco DCNM LAN Fabric Deployment, or switch-groups in Cisco DCNM Classic LAN Deployment. After the telemetry is enabled via NIR application, the telemetry manager in Cisco DCNM will push the necessary NTP server configurations to the switches by using the DCNM OOB IP address as the NTP server IP address. The following example is sample output for **show run ntp** command.

```
switch# show run ntp

!Command: show running-config ntp
!Running configuration last done at: Thu Jun 27 18:03:07 2019
!Time: Thu Jun 27 20:32:18 2019
```

```
version 7.0(3)I7(6) Bios:version 07.65
ntp server 192.168.126.117 prefer use-vrf management
```

### Telemetry Using In-Band (IB) Network:

The switches stream telemetry data through their front panel ports to Cisco DCNM assuming the connectivity from the switches to the Cisco DCNM In-Band network eth2 interface.

To set up the In-Band network for the Cisco DCNM LAN Fabric Deployment, see [In-Band Telemetry Network and NTP Configuration, on page 394](#).

## Installing and Deploying Applications

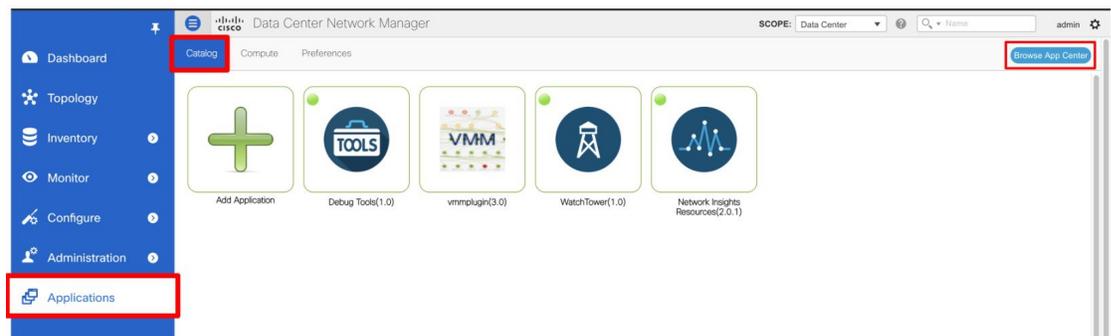
The following sections describes how to download, add, start, stop, and delete applications from the Cisco DCNM Web UI.

### Download App from the App Store

To download new applications from the Cisco DCNM Web UI, perform the following steps:

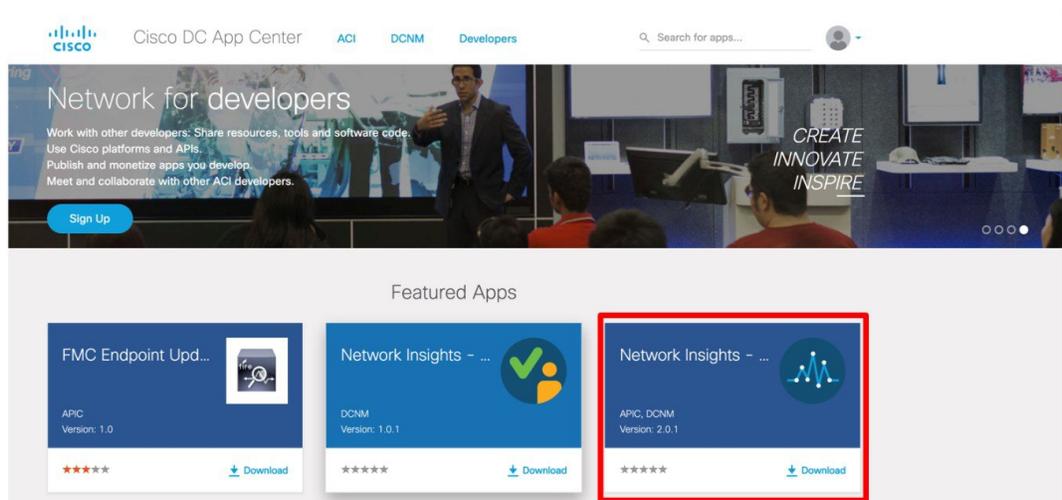
1. Choose **Applications**.

By default, the **Catalog** tab displays.



2. Click **Browse App Center** on the top-right corner on the window.

On the Cisco ACI App Center, locate the required application and click the download icon.



3. Save the application executable file on your local directory.

### Add a New Application to DCNM

To add new applications from the Cisco DCNM Web UI, perform the following steps:

1. Choose **Applications**.

By default, the **Catalog** tab displays.

2. Click **Add Application (+)** icon.



On the Application Upload window, from the Type drop-down field, choose one of the following to upload the application.

From the Type drop-down list, select one of the following:

- If the file is located in a local directory, select **Local-file**.

In the Upload field, click **Select files...** Navigate to the directory where you have stored the application file.

Select the application file and click **Open**.

Click **Upload**.

- If the application is located on a remote server, select **Secure copy**.



**Note** Ensure that the remote server must be capable of serving Secure-copy (SCP).

In the URI field, provide the path to the application file. The path must be in `{host-ip}:{filepath}` format.

In the Username field, enter the username to access the URI.

In the Password field, enter the appropriate password for accessing the URI.

Click **Upload**.

After the application successfully uploaded, it is displayed in the Catalog window.



The green icon on the left-top corner indicates that the application is launched successfully and is operational. If there is no green icon on the application, it indicates that the application is not running. Click the application to Launch it.



**Note** Ensure that the Compute Cluster is enabled before you install applications. A few applications may not work if the compute cluster is configured after installing the applications.

Click the gear icon on the left-bottom on the application icon to view the Application Specifications. The Info tab displays the running container information. The Specs tab displays the configuration.

### Starting Application

After the application is installed on the Cisco DCNM server, you must deploy the application. Click on the Application to begin deployment. Cisco DCNM starts all the services in the backend that are required for the application.



Network Insights  
Resources(2.0.1)

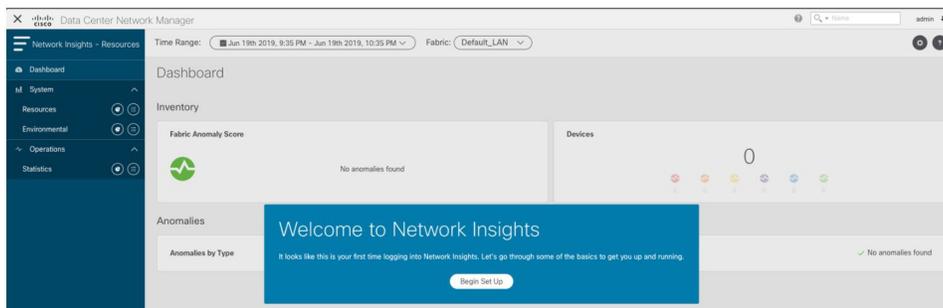
The green icon on the left-top corner indicates that the application is launched successfully and is operational.



Network Insights  
Resources(2.0.1)

The applications utilizing the Kafka infrastructure services require three actively joined compute nodes, when you begin the application. For example, NIR and NIA applications. If the application has a user interface, after the application is successfully started the UI redirects to the index page served by the application.

If the application has a user interface, after the application is successfully started the UI redirects to the index page served by the application.



To check the services that are running go back to **Applications > Catalog**. Click the gear icon on the left-bottom on the application icon to view the Application Specifications. The Info tab displays the running container information and the Specs tab displays the configuration as shown in the figures below.



Network Insights  
Resources(2.0.1)

Application Specifications

Info Spec

Running Instance Info

| Container Name        | Compute             | East-West IP | Fabric IP |
|-----------------------|---------------------|--------------|-----------|
| scheduler_Cisco_...   | nilesh-vm210.cis... | 10.10.10.10  |           |
| predictor_Cisco_af... | nilesh-vm208.cis... | 10.10.10.12  |           |
| correlator_Cisco_a... | nilesh-vm208.cis... | 10.10.10.26  |           |
| eventcollector_Cis... | nilesh-vm208.cis... | 10.10.10.30  |           |
| eventcollector_Cis... | nilesh-vm205.cis... | 10.10.10.28  |           |
| eventcollector_Cis... | nilesh-vm210.cis... | 10.10.10.29  |           |
| postprocessor_Cis...  | nilesh-vm210.cis... | 10.10.10.32  |           |
| postprocessor_Cis...  | nilesh-vm208.cis... | 10.10.10.33  |           |
| postprocessor_Cis...  | nilesh-vm205.cis... | 10.10.10.34  |           |
| utr_Cisco_afw.1       | nilesh-vm208.cis... | 10.10.10.38  | 24.0.0.4  |
| utr_Cisco_afw.3       | nilesh-vm205.cis... | 10.10.10.37  | 24.0.0.3  |
| utr_Cisco_afw.2       | nilesh-vm210.cis... | 10.10.10.36  | 24.0.0.2  |
| apiserver_Cisco_a...  | nilesh-vm208.cis... | 10.10.10.42  |           |
| apiserver_Cisco_a...  | nilesh-vm205.cis... | 10.10.10.40  |           |
| apiserver_Cisco_a...  | nilesh-vm210.cis... | 10.10.10.41  |           |

For information on how to remove computes from the cluster, stopping or deleting the applications, see [Application Framework User Interface, on page 406](#).

### Stop and Delete Applications

To delete the applications from the Catalog on the Cisco DCNM Web UI, perform the following steps:

1. Choose **Applications**.

By default, the **Catalog** tab displays, showing all the installed applications.

2. Click the red icon on the right-bottom corner to stop the application.



3. Check the **Wipe Volumes** check box to erase all the data that is related to the application.



WatchTower(1.0)

4. Click **Stop** to stop the application from streaming data from Cisco DCNM.  
The Green icon disappears after the application is successfully stopped.



WatchTower(1.0)

5. After you stop the application, click the **Waste Basket** icon to remove the application from the Catalog.



WatchTower(1.0)

## Application Framework User Interface

To use the Applications Framework feature, in the Cisco DCNM home page's left pane, click **Applications**.  
The Applications window displays the following tabs:

- **Catalog**—This tab lists the applications that are used by Cisco DCNM. These applications are used for performing various functions within Cisco DCNM. For more information, see *Catalog*.
- **Compute**—This tab displays the existing compute nodes. The tab shows nodes that are part of the hosting infrastructure. The uptime indicates how long they have been part of the infrastructure. In a High Availability (HA) setup, both the active and the standby nodes appear as joined. For more information, see [Compute, on page 410](#).



---

**Note** In the cluster mode, the Cisco DCNM servers will not appear under the Compute tab.

---

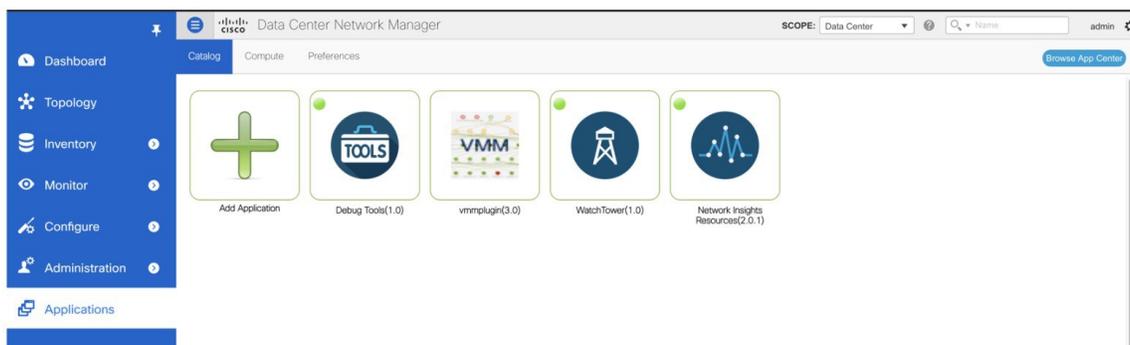
- **Preferences**—This tab is relevant to the cluster mode of deployment, where the application instances are placed. This tab enables you to compute the cluster connectivity and configure the Cluster Connectivity preferences. For more information, see [Preferences, on page 394](#).

Cisco DCNM uses the following applications:

- **Compliance**: This application helps in building fabrics for the Easy Fabric installation. The Compliance application runs as one instance per fabric. It is enabled when fabric is created. Similarly, it is disabled when fabric is deleted.
- **Kibana**: This is an open-source data-visualization plug-in for Elasticsearch, which provides visualization capabilities. Cisco DCNM uses the Kibana application for the Media Controller, and Endpoint Locator.
- **vmmplugin**: The Virtual Machine Manager (VMM) plug-in stores all the computes and the virtual machine information that connects to the fabric or the switch groups that are loaded into Cisco DCNM. VMM gathers compute repository information and displays the VMs, VSwitches/DVS, hosts in the topology view.
- **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 an IP and MAC address. In that sense, an endpoint can be a virtual machine (VM), container, bare-metal server, service appliance and so on.

## Catalog

The Catalog allows you to view all the applications that you have installed or enabled on the Cisco DCNM. Few applications are installed and are operational by default, when you install the Cisco DCNM.



The following applications appears based on the Cisco DCNM Deployments:

- Compliance (2.0)
- Debug plug-in (1.0)
- Endpoint Locator (1.1)
- Kibana (1.1)
- Vmmplugin (3.0)
- WatchTower (1.0)



**Note** The applications started by default, or also installed on the DCNM utilizes infrastructure services are operational, by default.

You can install more applications from the App Center, via the Web UI.

For instructions about downloading, adding, starting, stopping, and deleting applications from the Cisco DCNM Web UI, see [Installing and Deploying Applications, on page 401](#).

## Watch Tower

The Watch Tower helps you to monitor the infrastructure health and status. You can monitor the Alerts, Service Utilization, and Compute Utilization using the Watch Tower application. When you install or upgrade to 11.2(1), the Watch Tower application is installed and operational, by default.

To launch the Watch Tower app, on the Cisco DCNM Web UI, choose **Applications**. On the Catalog tab, click on **Watch Tower** to launch the application.



**Note** Watch Tower application is installed by default in Cisco DCNM cluster mode.

Health Monitor app broadly monitors and alerts on the following metrics for Services, Computes and DCNM server:

- CPU utilization
- Memory utilization

- Network I/O (eth0)
- Disk I/O

You can monitor the following using the Watch Tower application:

## Alerts

The Alerts window provides information about the number of alerts that have occurred, from the specified date and time. You can view the alerts, based on the following categories, in the graphical view and the list view.

In the graphical view, the categories are:

- **Severity** displays the alerts, based on the severity: Critical/Major/Minor/Info.
- **Type** displays the alerts, based on the cluster type.
- **Compute** displays the alerts, for each compute node.
- **Service** displays the alerts, for all the services running on Cisco DCNM.

Click on the Refresh icon to refresh the alerts. Click on the list view icon to view the alerts in list format.

In the List View, alerts are displayed in tabular format with the following categories:

- **Timestamp** displays the time when the alert triggers. Format is MM/DD HH:MM AM/PM.
- **Alert Severity** displays the severity of alert.
- **Alert Type** displays the cluster alert type.
- **Node Name** displays the node name where the alert triggers.
- **Container Name** specifies the service name which triggered the alert.
- **Alert Description** displays the summary of the alert.

Click on the right or left navigation arrows to move to the next or the previous page.

You can also choose to set the number of items to view on page. Select a suitable number from the **Objects Per Page** drop-down list.

Click on the **Graphical representation** icon to go to the graphical view. Click on **Download Data** icon to download alerts information for troubleshooting purposes.

## Service Utilization

You can monitor all the services running on the Cisco DCNM on this window. Based on the time range and the service, the graphical view shows the CPU and Memory utilization for service. Click on the **Service Utilization** icon on the top-right corner to launch the CPU utilization graphical view.

From the **Time Range** drop-down list, choose the time range for which you want to view the utilization. You can select a specific time interval to view the metrics during that time interval. Click the fields showing the date and time to select the required date and time interval. You can also click the date on the calendar to set range. Click **Apply** to confirm the time range.

From the **Services** drop-down list, choose the service to view its Service utilization. This list comprises of all the services that are currently running on the Cisco DCNM.

Select the Time Range to view the **Service**, the **Cpu Utilization**, and **Memory Utilization** graphs. You can hover over specific points on the respective graphs for more information on CPU and Memory utilization at specific time.

The memory utilization graphical view depicts the actual memory consumption (RAM) in Gigabytes (GB).

Click [X] icon on the top-right corner to close the Service Utilization window and revert to the Alerts window.

## Compute Utilization

You can monitor all the computes installed with the Cisco DCNM. Based on the time range and the service, the graphical view shows the CPU and Memory utilization for service. Click on the **Compute Utilization** icon on the top-right corner to launch the CPU utilization graphical view.

From the **Time Range** drop-down list, choose the time range for which you want to view the utilization. You can select a specific time interval to view the metrics during that time interval. Click the fields showing the date and time to select the required date and time interval. You can also click the date on the calendar to set range. Click **Apply** to confirm the time range.

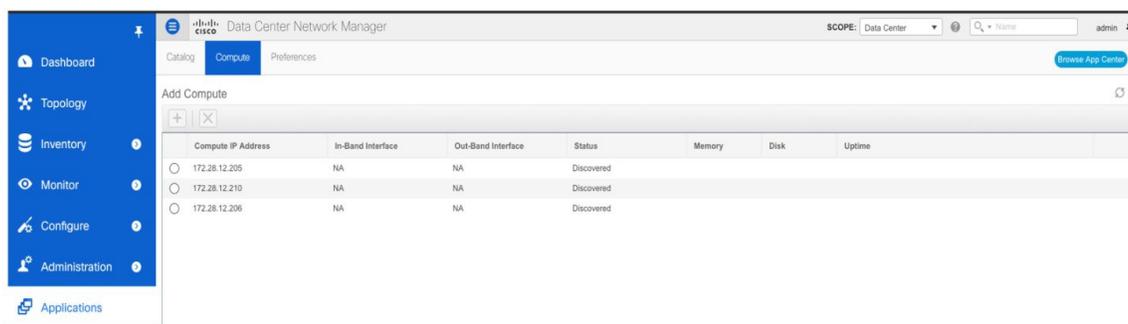
Select the Time Range to view the **Service**, the **Cpu Utilization**, and **Memory Utilization** graphs. You can hover over specific points on the respective graphs for more information on CPU and Memory utilization at specific time.

The memory utilization graphical view depicts the actual memory consumption (RAM) in Gigabytes (GB).

Click [X] icon on the top-right corner to close the Service Utilization window and revert to the Alerts window.

## Compute

This tab displays the existing compute nodes. The tab shows nodes that are part of the hosting infrastructure. The uptime indicates how long they have been part of the infrastructure. In a High Availability (HA) setup, both the active and the standby nodes appear as joined. In clustered mode, the compute nodes status indicate if the nodes are joined or discovered.



| Compute IP Address | In-Band Interface | Out-Band Interface | Status     | Memory | Disk | Uptime |
|--------------------|-------------------|--------------------|------------|--------|------|--------|
| 172.28.12.205      | NA                | NA                 | Discovered |        |      |        |
| 172.28.12.210      | NA                | NA                 | Discovered |        |      |        |
| 172.28.12.206      | NA                | NA                 | Discovered |        |      |        |



**Note** If the NTP server for compute nodes is not synchronized with the NTP server for DCNM Servers (Active and Standby) and Computes, you cannot configure a cluster.

The certificates are generated with a timestamp. If you configure the Compute nodes using a different NTP server, the mismatch in timestamp will not allow to validate the certificates. Therefore, if the compute cluster is configured despite of a mismatch of NTP server, the applications will not function properly.



**Note** In clustered mode, the Cisco DCNM servers will not appear under the Compute tab.

The following table describes the fields that appear on **Applications > Compute**.

**Table 23: Field and Description on Compute Tab**

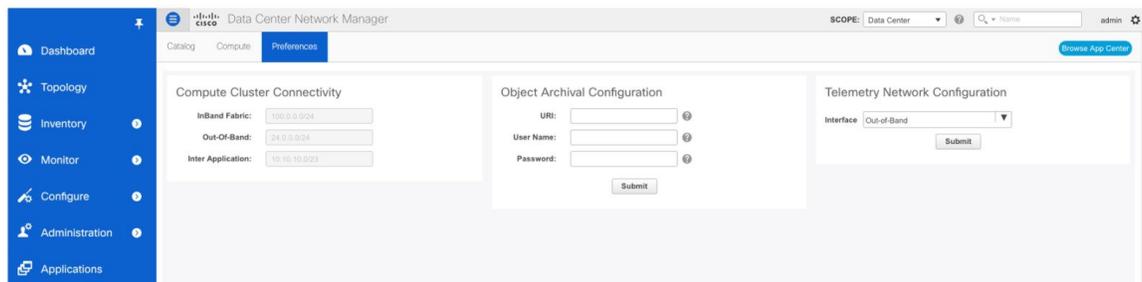
| Field              | Description   |
|--------------------|---|
| Compute IP Address | Specifies the IP Address of the Compute node.   |
| In-Band Interface  | Specifies the in-band management interface.   |
| Out-Band Interface | Specifies the out-band management interface.  |
| Status             | Specifies the status of the Compute node. <ul style="list-style-type: none"> <li>• Joined</li> <li>• Discovered</li> <li>• Failed</li> <li>• Offline</li> </ul> |
| Memory             | Specifies the memory that is consumed by the node.  |
| Disk               | Specifies the disk space that is consumed on the compute node.  |
| Uptime             | Specifies the duration of the uptime for a compute node.  |

When you install a compute node with correct parameters, it appears as **Joined** in the Status column. However, the other two computes appears as Discovered. To add computes to the cluster mode from Cisco DCNM Web UI, see [Adding Computes into the Cluster Mode, on page 390](#).

To configure or modify the Cluster Connectivity preferences, see [Preferences, on page 394](#).

## Preferences

This tab is relevant to the cluster mode of deployment, where the application instances are placed. This tab enables you to compute cluster connectivity and configure the Cluster Connectivity preferences.



### Object Archival Configuration

The NIA application collects tech support logs for all switches in Fabric, and determines the advisory, based on the data. The logs are saved on the Cisco DCNM server for further analysis or troubleshooting. If you need to download these logs before their life span ends or to create some space on the DCNM server, you can move the logs to a remote server.

In the **URI** field, enter the relative path to the archive folder, in the format `host[:port]/[path to archive]`. Enter the username and password to access the URI, in the **username** and **Password** field. Click **Submit** to configure the remote server.

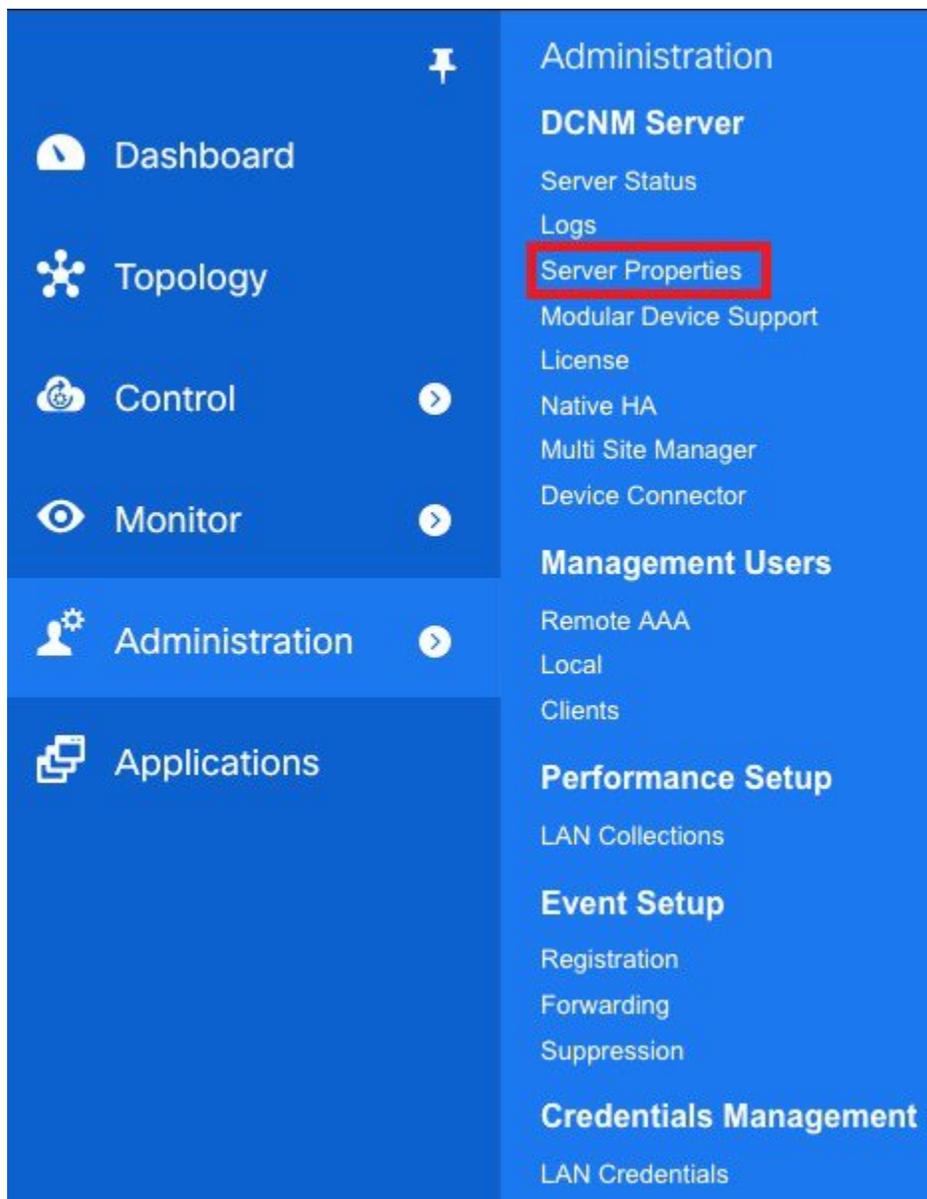
## In-Band Telemetry Network and NTP Configuration

### Before you begin

If Network Insights Resources (NIR) application is running, you must disable it on all the fabrics before you begin the procedure.

### Procedure

**Step 1** Choose **Administration > DCNM Server > Server Properties**.



- Step 2** In the # **Template Properties** area, locate the **template.in\_use.check** field.
- Step 3** Set the **template.in\_use.check** field to **true**. Click **Apply Changes**.

The screenshot shows the Cisco Data Center Network Manager (DCNM) Administration / DCNM Server / Server Properties page. The left sidebar contains navigation options: Dashboard, Topology, Control, Monitor, Administration, and Applications. The main content area displays configuration fields for various properties:

- # periodically restart ECT baseline training after number of days**
  - # (Default is 28)
  - san.telemetry.train.reset: 28
- # Template Properties**
  - template.in\_use.check: true (highlighted with a red box)
  - template.use\_cache: true
  - template.server\_validation\_check: false
  - template.server\_validation\_continue\_on\_error: false
- #Image Management Property**
  - FILE\_SELECTION\_FILTER: true

An **Apply Changes** button is located in the top right corner, also highlighted with a red box.

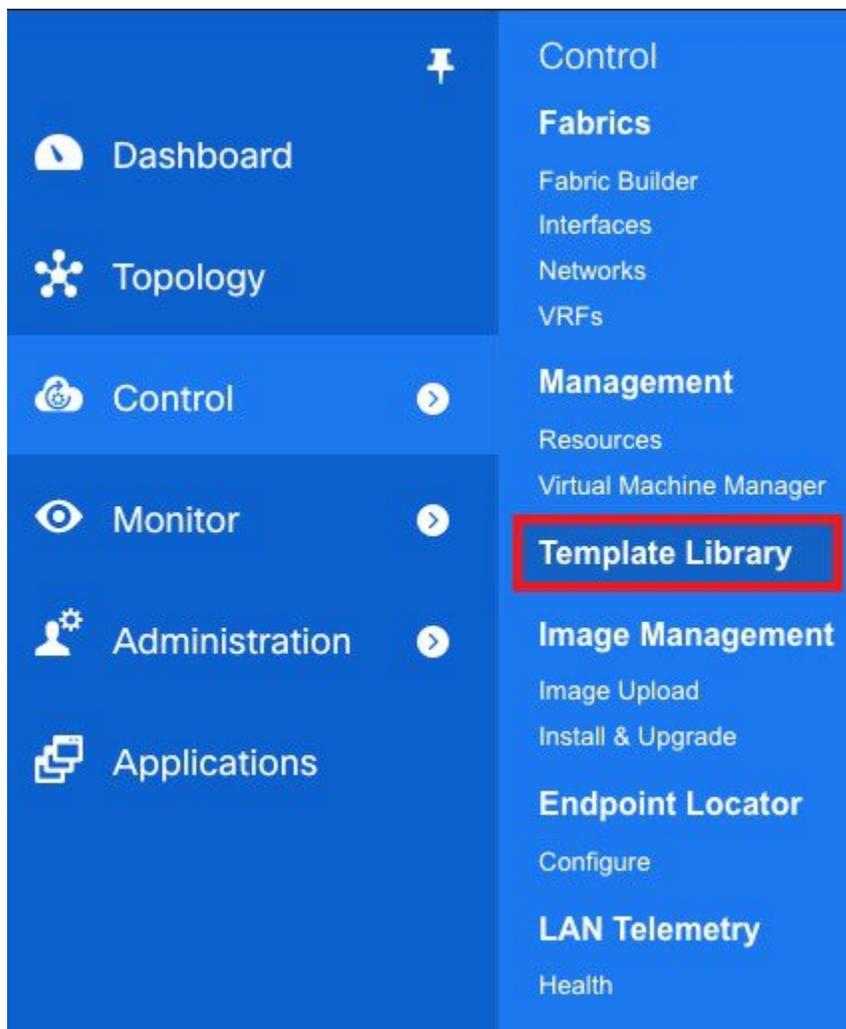
**Step 4** Click **OK** to confirm that this property change does not require restarting the Cisco DCNM server.

The dialog box contains the following text:

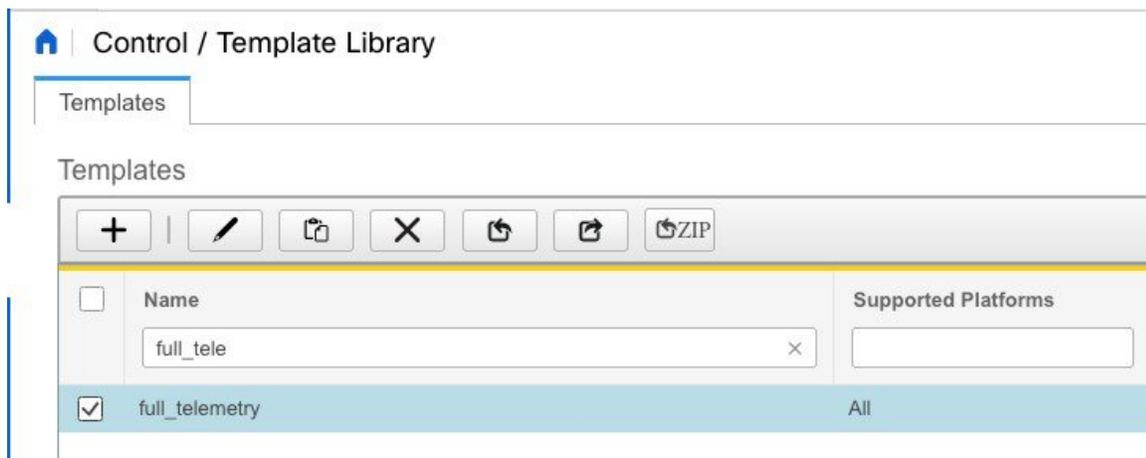
Please restart DCNM SAN service if you update properties other than EMC Callhome properties(server.callhome.enable, server.callhome.xmlDir), Event Forwarding properties (server.forward.event.enable), Template properties (template.in\_use.check property), Event Registration properties(syslog.disable) or fabric.enableNpvDiscovery properties. (Note: please restart all instances if federation is deployed). Please resync vmm if you updated vmm.resync.timer

An **Ok** button is located at the bottom right of the dialog.

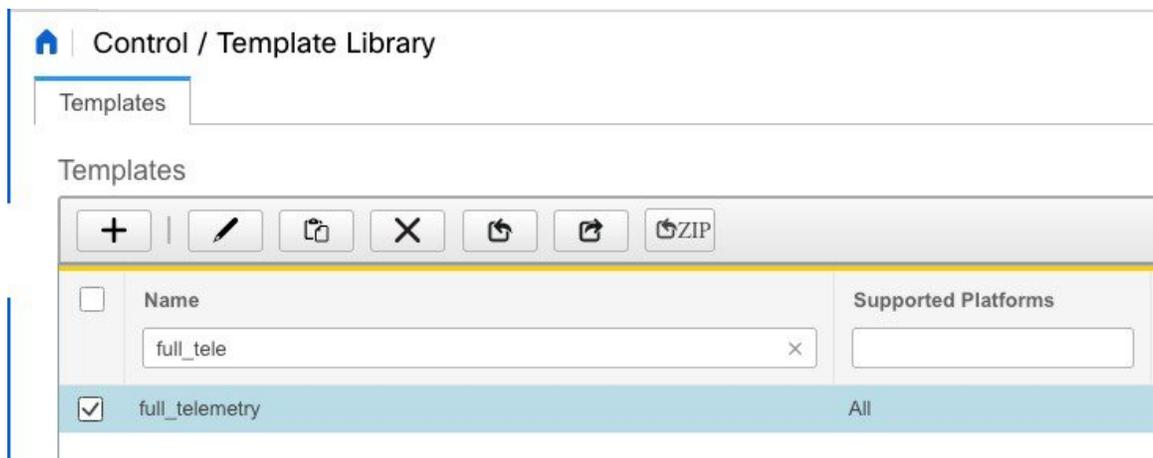
**Step 5** Choose **Control > Template Library**.



**Step 6** Locate and select **full\_telemetry** template, and click **Edit** icon.



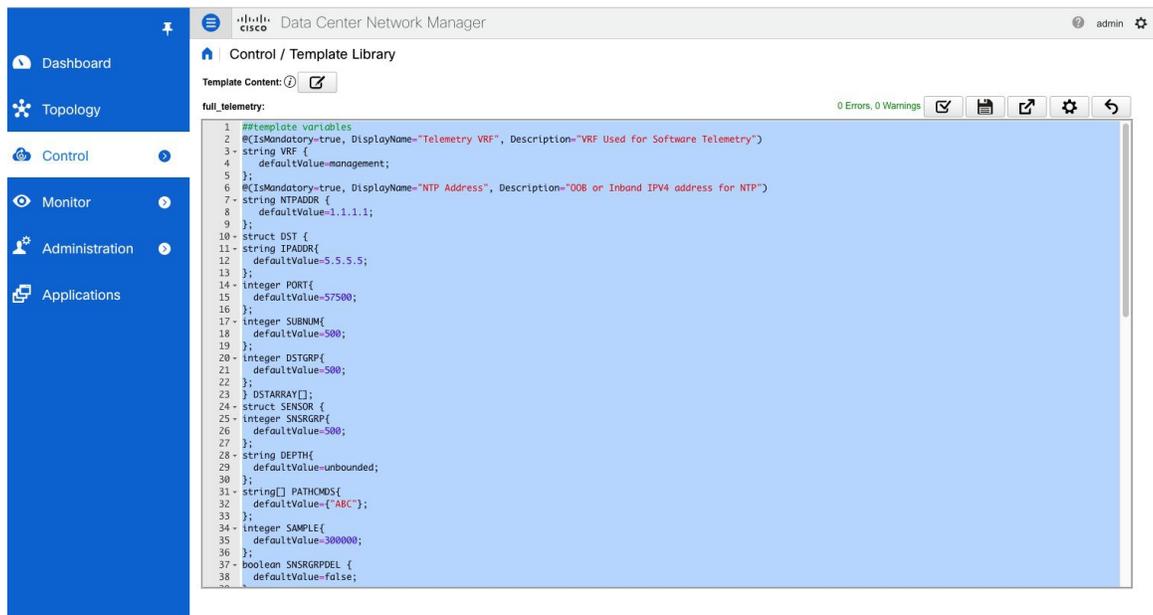
**Step 7** Click **Yes** to cancel any pending scheduled template jobs associated with this template.



The **Template Content** window is displayed.

### Step 8

Locate the **##template content** in the text section.



### Step 9

Change the **use-vrf \$VRFS\$** to **use-vrf management**.

```

##template content
if ( $$NTPADDR$$ != "" ) {
  if ( $$SVRF$$ != "management" ) {
ntp server $$NTPADDR$$ prefer use-vrf $$SVRF$$
  }
  else {
ntp server $$NTPADDR$$ prefer use-vrf management
  }
}
telemetry

##template content
if ( $$NTPADDR$$ != "" ) {
  if ( $$SVRF$$ != "management" ) {
ntp server $$NTPADDR$$ prefer use-vrf management
  }
  else {
ntp server $$NTPADDR$$ prefer use-vrf management
  }
}

```

Click **Save** icon. Click **OK** on the **Validation Result** window.

- Step 10** Choose **Applications > Preferences**.  
The **Preferences** tab is displayed.

The screenshot shows the Cisco Data Center Network Manager interface. The 'Preferences' tab is active. Under the 'Telemetry Network Configuration' section, the 'Interface' dropdown menu is open, showing 'Out-of-Band' as the selected option. Other sections like 'Compute Cluster Connectivity' and 'Object Archival Configuration' are also visible.

- Step 11** From the **Interface** drop-down list, select **In-Band**.

This is a close-up of the 'Interface' dropdown menu in the 'Telemetry Network Configuration' section. The 'In-Band' option is selected and highlighted in grey.

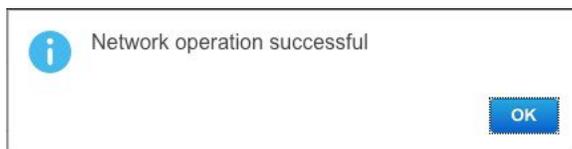
- Step 12** In the **VRF** field, enter the VRF name, which has connectivity to the Cisco DCNM In-Band interface.  
The VRF value is set to default. You can also change this to any other VRF assuming that the switches are configured with the VRF and there is availability to the in-band interface of the Cisco DCNM through that VRF.

This is a close-up of the 'Telemetry Network Configuration' form. The 'Interface' dropdown is set to 'In-Band' and the 'VRF' text input field contains the word 'default'. A 'Submit' button is visible below the fields.

**Note** If the Cisco NIR telemetry is already enabled for some fabrics, disable the Telemetry on all the enabled fabrics before you modify the Telemetry Network Configuration. You can enable NIR telemetry after you modify the telemetry network to begin streaming of data through the in-band interface.

Click **Submit**.

**Step 13** Click **OK** to confirm the network operation.



The telemetry network configuration is now complete.

---

### What to do next

To download and install the Network Insights applications, see [Catalog, on page 407](#).

After the telemetry is enabled via Network Insights Resources (NIR) application, the telemetry manager in Cisco DCNM will push the necessary NTP server configurations to the switches.

## Failure Scenario

Recommendation for minimum redundancy configuration with a DCNM OVA install is as follows:

- DCNM Active Node(Active) and compute node 1 in server1.
- DCNM Standby Node and compute node 2 in server2.
- Compute node 3 in server3.

When DCNM Active node is down, the Standby node takes full responsibility of running the core functionality.

When a compute node is down, the applications may continue to function with limited functionality. If this situation persists for a longer duration, it affects the performance and reliability of the applications. When more than one node is down, it affects the applications functionality and most of the applications fail to function.

You must maintain 3 compute nodes at any time. If a compute node goes down, rectify the issue as soon as possible, for the services to function as expected.

## Compute Node Disaster Recovery

When a compute node is lost due to a disaster and is irrecoverable, you must install another compute node with the same parameters. This will essentially appear as a reboot of the compute with lost data and it tries to join the cluster automatically. After it joins the cluster, all the data will synchronize from the other two compute nodes.



## CHAPTER 8

# Connecting Cisco Data Center and a Public Cloud

---

- [Connecting Cisco Data Center and a Public Cloud, on page 419](#)

## Connecting Cisco Data Center and a Public Cloud

This section explains the preview functionality that allows public cloud connectivity from a Cisco DCNM provisioned VXLAN EVPN fabric to the Microsoft Azure public cloud. The layer-3 connectivity ensures a seamless and secure communication between the workloads on premise and the Microsoft Azure cloud. The connectivity is provisioned through the Cisco Cloud Services Router 1000v (Cisco CSR 1000v) that is managed by Cisco DCNM. BGP EVPN is employed for the control plane and VXLAN is employed for the data plane. A secure IPsec tunnel is established between the Cisco CSR 1000v in the premise and the Cisco CSR 1000v in the public cloud.



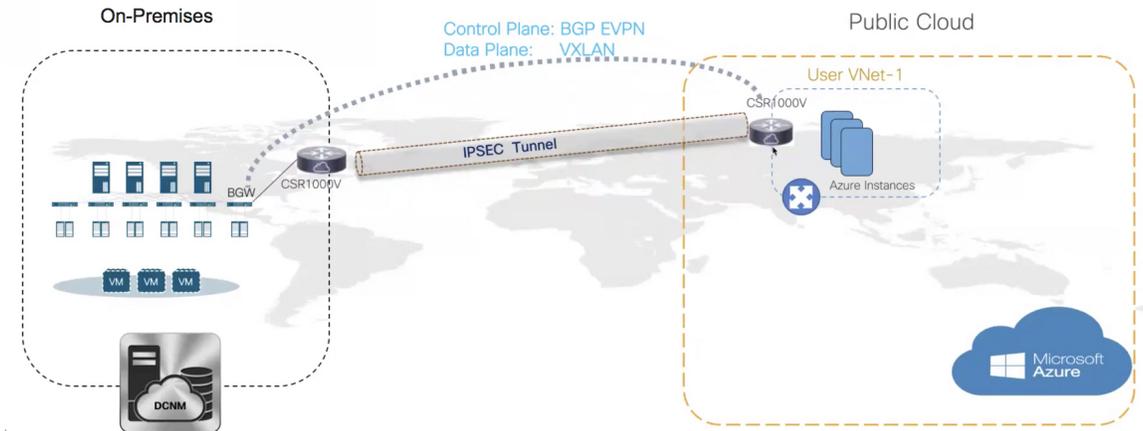
---

**Note** Cisco DCNM supports discovery and management of IOS-XE, specifically Cisco CSR 1000v.

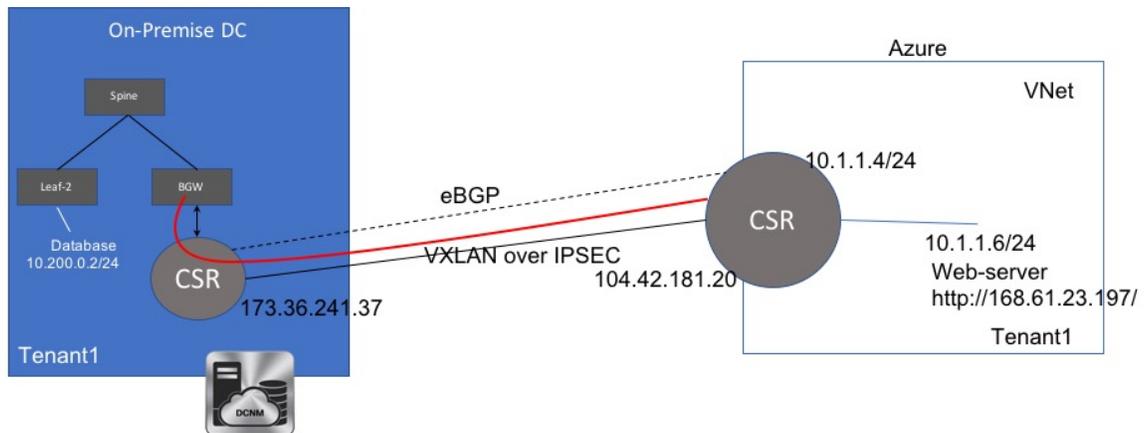
---

## Topology Overview

Figure 6: Topology Overview



The on-premise data center has the required switches. One of these switches is a border gateway (BGW) that interfaces with an core router for WAN connectivity to the public cloud. The Cisco CSR 1000v is the core router in this use case. You can import this core router into an external fabric in Cisco DCNM. The following figure depicts the sample topology that is employed.



In this example, we list the tasks that are required to provide a layer-3 connectivity between a VM behind standalone leaf and a VM in the Microsoft Azure cloud in a specific user VNET.

The public cloud has a Cisco CSR 1000v, Microsoft Azure instances, Azure Virtual Networks (Azure VNets), and a VM. The Cisco CSR 1000v in the cloud has an interface with the VM.

We are using eBGP between the two core routers for exchanging underlay routing and reachability. The VXLAN connects the on-premises BGW and the core router on Microsoft Azure, over the IPsec tunnel.

In this use-case, we are going to configure the setup as follows:

## Guidelines and Limitations

The following are the guidelines and limitations for connecting an on-premises data center and a public cloud:

- This is a preview-only feature. We recommend that you use this feature only in lab setups, and not in production environments.
- This preview functionality is supported only for Cisco CSR 1000v Series Routers.
- Cisco CSR 1000v Series Routers support route-based IP Security (IPsec) tunnel interface.
- Use Cisco Nexus 9000 Series Switches or Cisco Nexus 3000 Series Switches in the VXLAN EVPN Easy fabric in Cisco DCNM.
- The IP addresses specified in this document are sample addresses. Ensure that your setup reflects the IP addresses used in the production network.

## Prerequisites

- Create an account with Microsoft Azure.
- Create VNets for the public-cloud core router in Microsoft Azure.
- Deploy a Cisco CSR 1000v in Microsoft Azure. This Cisco CSR 1000v is the public-cloud core router. See the [Deploying Cisco CSR 1000v on Microsoft Azure, on page 439](#) section for more information.
- Use switches that support Cisco NX-OS Release 7.0(3)I7(x) or higher versions as border gateways are required.
- Set up the Cisco DCNM, switches, Cisco CSR 1000v, and other devices in a DMZ or equivalent zone to have access to the public internet.
- Familiarity with VXLAN BGP EVPN data center fabric architecture and configuration through DCNM.
- Familiarity with MSD fabrics.



---

**Note** Refer to the *Control* chapter in the *Cisco DCNM LAN Fabric Configuration Guide*, for information on various tasks that are required in setting up.

---

## Task Summary

The following sections list the task summary to establish a connection between the on-premises data center and the public cloud.

### On-premises Data Center

1. Enable the preview functionality.
2. Create a fabric with switches for the on-premises data center, and configure one of the switches with BGW role.
3. Create an external fabric for the on-premises core router. Discover a Cisco CSR 1000v as the core router.

4. Simulate an IP address as on-premises host on the BGW.

### Public Cloud

1. Create an external fabric for the public cloud core router.
2. Discover a Cisco CSR 1000v for the public cloud, which is the core router.

### Connectivity

1. Create an MSD fabric and import the fabrics that were created previously.
2. Connect the BGW and the on-premises core router.
3. Create an IPsec tunnel between the on-premises core router and the public-cloud core router.
4. Create an eBGP underlay connection between the core routers that runs over the IPsec Tunnel.
5. Connect the BGW and the public cloud core router using VXLAN EVPN.
6. Extend the VRFs in fabrics.

The procedure that is involved in each task in this section is explained in the following sections.

## Enabling the Preview Functionality

To enable the preview functionality of a public cloud connectivity from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Administration > DCNM Server > Server Properties**.
- The **Server Properties** window appears.
- Step 2** Locate the **Private and public cloud connectivity** properties.
- Step 3** Set the value of the **private\_public\_cloud\_connectivity.enable** field to **true**.
- Note** The value of this field is set to **false** by default.
- Step 4** (Optional) Set the polling time in the **private\_public\_cloud\_connectivity.stats.polling\_time** field to **15000**.
- The value is in milliseconds. Here, Cisco DCNM queries the on-premises core router and updates the state of the routing table every 15 seconds.

The screenshot shows the Cisco Data Center Network Manager (DCNM) web interface. The breadcrumb navigation is Administration / DCNM Server / Server Properties. Below the navigation, there is a configuration section for private and public cloud connectivity. The configuration includes two input fields: one for 'private\_public\_cloud\_connectivity.enable' set to 'true' and another for 'private\_public\_cloud\_connectivity.stats.polling\_time' set to '15000'.

```
#
# Private and public cloud connectivity
#
private_public_cloud_connectivity.enable true
private_public_cloud_connectivity.stats.polling_time 15000
```

**Step 5** Click **Apply Changes**.

**Step 6** Restart Cisco DCNM using the **appmgr restart dcnm** command.

A warning about the preview features enabled appears after you log in to the Cisco DCNM Web UI.

**Note** This is a preview only feature. We recommend that you use this feature only in lab setups, and not in production environments.

## Setting Up the On-premise External Fabric with CSR 1000v

Create an external fabric for the on-premises edge router.

### Creating an External Fabric

To create an external fabric from Cisco DCNM Web UI, perform the following steps:

#### Procedure

**Step 1** Choose **Control > Fabrics > Fabric Builder**.

The **Fabric Builder** window appears.

**Step 2** Click **Create Fabric**.

The **Add Fabric** dialog box appears.

**Step 3** Enter the fabric name as **CSR-OnPrem** in the **Fabric Name** field.

**Step 4** Choose **External\_Fabric\_11\_1** from the Fabric Template drop-down list.

**Step 5** Enter the BGP AS number in the **BGP AS #** field.

**Step 6** Uncheck the **Fabric Monitor Mode** check box.

**Step 7** Click **Save**.

A fabric is created and the fabric topology window appears.

**What to do next**

Discover the on-premises core router.

**Discovering the On-Premises Core Router**

Cisco CSR 1000v is used for on-premises core routing. To discover the core router in the fabric topology window, perform the following steps:

**Before you begin**

Ensure that you know the credentials of the core router.

**Procedure**

**Step 1** Click **Add switches** in the Actions pane.

The **Inventory Management** dialog box appears.

**Step 2** Enter values for the following fields under the **Discover Existing Switches** tab:

| Field       | Description   |
|-------------|---|
| Seed IP     | Enter the IP address of the core router.  |
| Device Type | Choose <b>IOS XE</b> from the drop-down list.<br><b>Note</b> You can see the <b>IOS XE</b> option only after you enable the preview functionality and restart the DCNM. |
| Username    | Enter the username of the core router for SSH access.   |
| Password    | Enter the password of the core router for SSH access.   |

**Note** An error appears if you try to discover a switch that is already discovered.

**Step 3** Click **Start Discovery**.

The fabric topology window appears, and a pop-up message appears at the bottom-right about the discovery.

For example: <ip-address> added for discovery.

**Note** Discovering switches might take some time.

**Step 4** Click **Tabular view** in the Actions pane.

The switches and links window appears, where you can view the scan details. The discovery status is discovering in red with a warning icon next to it if the discovery is in progress.

**Step 5** View the details of the core router.

After the router is discovered:

- The discovery status changes to **ok** in green with a check box checked next to it.
- The value of the router under the **Fabric Status** column will be **In-Sync**.

**Step 6** Go back to the fabric topology window and refresh the topology.

---

#### What to do next

Set up a VXLAN EVPN fabric for the on-premises data center, which has a BGW.

## Setting Up the VXLAN EVPN Fabric

Create a fabric for the BGW.

### Creating a VXLAN EVPN Fabric

To create a VXLAN EVPN fabric from Cisco DCNM Web UI, perform the following steps:

#### Procedure

---

- Step 1** Choose **Control > Fabrics > Fabric Builder**.
- The **Fabric Builder** window appears.
- Step 2** Click **Create Fabric**.
- The **Add Fabric** dialog box appears.
- Step 3** Enter the fabric name as **site2** in the **Fabric Name** field.
- Step 4** Choose **Easy\_Fabric\_11\_1** from the **Fabric Template** drop-down list.
- Step 5** Enter values in all the mandatory fields.
- Step 6** Click **Save**.
- A fabric is created and the fabric topology window appears.
- 

#### What to do next

Add switches in this fabric and assign the BGW role for one of the switches.

### Assigning the BGW Role

To assign a switch with the BGW role, perform the following steps:

#### Before you begin

Add switches to the **site2** fabric.

#### Procedure

---

- Step 1** Right-click the switch for which you need to set the BGW role.

A list of actions that you can perform on the switch appears.

**Step 2** Choose **Set role > Border Gateway**.

---

#### What to do next

Set up a fabric for the public cloud.

## Setting Up the External Fabric with CSR in Azure

Create an external fabric for the public cloud core router.

### Creating an External Fabric

To create an external fabric from Cisco DCNM Web UI, perform the following steps:

#### Procedure

---

- Step 1** Choose **Control > Fabrics > Fabric Builder**.  
The **Fabric Builder** window appears.
- Step 2** Click **Create Fabric**.  
The **Add Fabric** dialog box appears.
- Step 3** Enter the fabric name as **CSR-Azure** in the **Fabric Name** field.
- Step 4** Choose **External\_Fabric\_11\_1** from the **Fabric Template** drop-down list.
- Step 5** Enter the BGP AS number in the **BGP AS # field**.
- Step 6** Uncheck the **Fabric Monitor Mode** check box.
- Step 7** Click **Save**.  
A fabric is created and the fabric topology window appears.
- 

#### What to do next

Discover the public-cloud core router in this fabric.

### Discovering the Core Router

Cisco CSR 1000v Series router is used for the public-cloud core routing as well. To discover the core router in the fabric topology window, perform the following steps:

#### Before you begin

Ensure that you know the credentials of the core router.

## Procedure

**Step 1** Click **Add switches** in the **Actions** pane.

The **Inventory Management** dialog box appears.

**Step 2** Enter values for the following fields under the **Discover Existing Switches** tab:

| Field       | Description   |
|-------------|---|
| Seed IP     | Enter the IP address of the core router.  |
| Device Type | Choose <b>IOS XE</b> from the drop-down list.<br><b>Note</b> You can see the <b>IOS XE</b> option only after you enable the preview functionality and restart the DCNM. |
| Username    | Enter the username of the core router for SSH access.   |
| Password    | Enter the password of the core router for SSH access.   |

**Note** An error message appears if you try to discover a switch that is already discovered.

**Step 3** Click **Start Discovery**.

The fabric topology window appears, and a pop-up message appears at the bottom-right about the switch discovery. For example: **<ip-address> added for discovery**

**Note** Discovering switches takes some time.

**Step 4** Click **Tabular view** in the **Actions** pane.

The switches and links window appears, where you can view the scan details. The discovery status is **discovering** in red with a warning icon next to it if the discovery is in progress.

**Step 5** View the details of the core router.

After the discovery of the router:

- The discovery status changes to **ok** in green with a check box checked next to it.
- The value of the router under the **Fabric Status** column changes to **In-Sync**.

**Step 6** Go back to the fabric topology window and refresh the topology.

### What to do next

Create an MSD fabric and import other fabrics, created previously, into it.

## Setting Up the MSD Fabric for Connectivity

Create an MSD fabric to bring all the standalone fabrics together for connectivity.

## Creating an MSD Fabric

To create an MSD fabric from Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Fabrics > Fabric Builder**.  
The **Fabric Builder** window appears.
- Step 2** Click **Create Fabric**.  
The **Add Fabric** dialog box appears.
- Step 3** Enter the fabric name as **Cloud-Connect** in the **Fabric Name** field.
- Step 4** Choose **MSD\_Fabric\_11\_1** from the **Fabric Template** drop-down list.
- Step 5** Enter values in all the mandatory fields.
- Step 6** Click **Save**.  
A fabric is created and the fabric topology window appears.
- 

### What to do next

Move other fabrics into this MSD fabric.

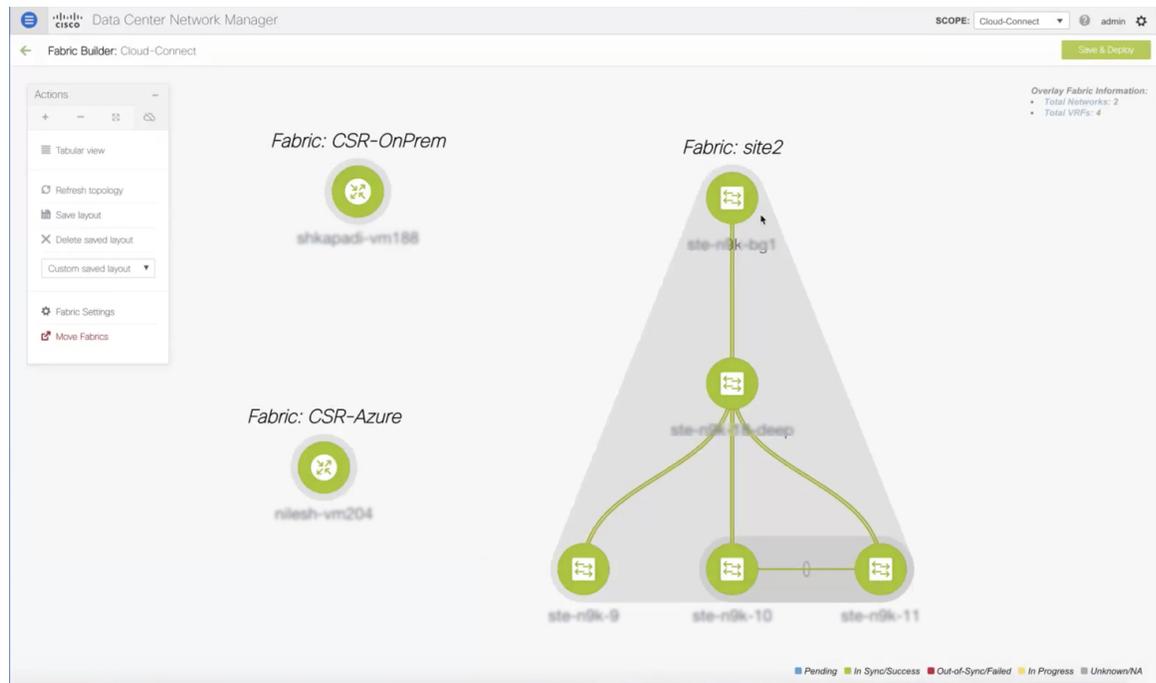
## Moving Other Fabrics into the MSD Fabric

To move other fabrics into the **Cloud-Connect** fabric from the fabric topology window, perform the following steps:

### Procedure

---

- Step 1** Click **Move Fabric** in the **Actions** pane.  
The **Move Fabric** dialog box appears. It contains a list of fabrics.
- Step 2** Choose **CSR-OnPREm, site2**, and **CSR-Azure** fabrics.
- Step 3** Click **Add**.
- Step 4** Close the dialog box and refresh the fabric topology.  
All the member fabrics appear in the **Cloud-Connect** fabric.



### What to do next

Set up the connections between fabrics.

## Setting Up Connections

Connect the fabrics that you created previously using different links.

### Connecting the On-Premises BGW and the On-Premises Core Router

To add a link between the on-premises BGW and the on-premises core router, perform the following steps:

#### Procedure

- Step 1** Right-click anywhere in the **Cloud-Connect** topology window.  
The actions that you can perform in the fabric appears in a list. Alternatively, from the fabric topology window, choose **Tabular view** in the **Actions** pane, and click the **Links** tab.
- Step 2** Choose **Add Link**.  
The **Link Management - Add Link** dialog box appears.
- Step 3** Enter values for the following fields:

| Field                 | Description   |
|-----------------------|---|
| Link Type             | Choose the <b>Inter-Fabric</b> link type from the drop-down list.   |
| Link Sub-Type         | Choose the <b>MULTISITE_UNDERLAY</b> link sub-type from the drop-down list.   |
| Link Template         | Choose the <b>csr_ext_multisite_underlay_setup</b> link template from the drop-down list.<br><br><b>Note</b> This template is available only after you enable the preview functionality and restart the DCNM. |
| Source Fabric         | Choose <b>site2</b> as the source fabric from the drop-down list.   |
| Destination Fabric    | Choose <b>CSR-OnPrem</b> as the destination fabric from the drop-down list.   |
| Source Device         | Choose the BGW from the drop-down list.   |
| Source Interface      | Choose the BGW's interface.   |
| Destination Device    | Choose the on-premises core router from the drop-down list.   |
| Destination Interface | Choose the on-premises core router's interface from the drop-down list.   |

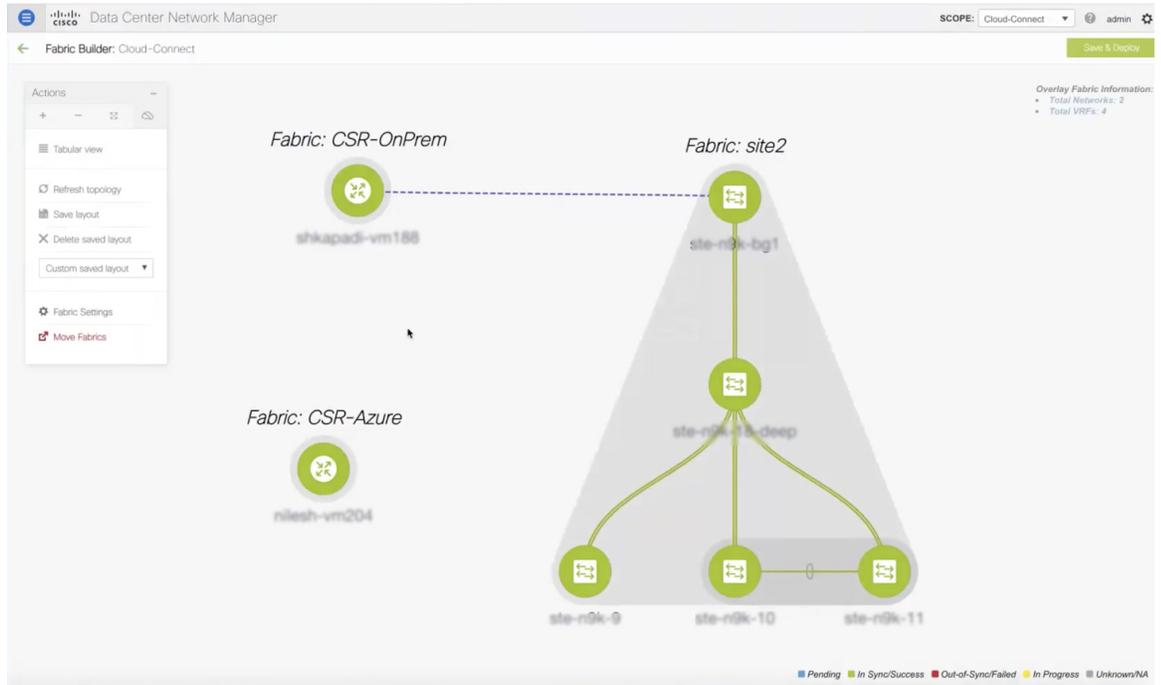
**Step 4** Enter values for the following fields under the **Link Profile** area in the **General** tab:

| Field       | Description   |
|-------------|---|
| IP_MASK     | Enter the IPv4 address of the source interface with a subnet. |
| NEIGHBOR_IP | Enter the IPv4 address of the destination interface.          |

To verify the IP address from the Cisco DCNM Web UI, choose **Control > Fabrics > Interfaces**. Choose the fabric from the **Scope** drop-down list, and search the device. The IP address of the device will be listed in the **IP/Prefix** column.

**Step 5** Click **Save**.

The fabric topology window refreshes. A link is added between the on-premises BGW in the **site2** fabric and the on-premises core router in the **CSR-OnPrem** fabric.



**What to do next**

Connect the on-premises core router and the public-cloud core router.

**Connecting the On-prem Core Router and the Public-cloud Core Router with IPsec Tunnel**

To add a link between the on-prem core router and the public-cloud core router, perform the following steps:

**Procedure**

- Step 1** Right-click anywhere in the **Cloud-Connect** topology window.  
The actions that you can perform in the fabric appears in a list. Alternatively, from the fabric topology window, choose **Tabular view** in the **Actions** pane, and click the **Links** tab.
- Step 2** Choose **Add Link**.  
The **Link Management - Add Link** dialog box appears.
- Step 3** Enter values for the following fields:

| Field         | Description  |
|---------------|--|
| Link Type     | Choose the <b>Inter-Fabric</b> link type from the drop-down list.          |
| Link Sub-Type | Choose the <b>BGP_OVER_IPSEC</b> link sub-type from the drop-down list.    |
| Link Template | Choose the <b>csr_link_template</b> link template from the drop-down list. |

| Field                 | Description  |
|-----------------------|--|
| Source Fabric         | Choose <b>CSR-OnPrem</b> as the source fabric from the drop-down list.     |
| Destination Fabric    | Choose <b>CSR-Azure</b> as the destination fabric from the drop-down list. |
| Source Device         | Choose the on-prem core router from the drop-down list.                    |
| Source Interface      | Choose the on-prem core router's interface.                                |
| Destination Device    | Choose the public-cloud core router from the drop-down list.               |
| Destination Interface | Choose the public-cloud core router's interface from the drop-down list.   |

**Step 4** In the **Link Profile** area under the **General** tab, enter the the pass key used for IPsec tunnel in the **SHARED\_KEY** field.

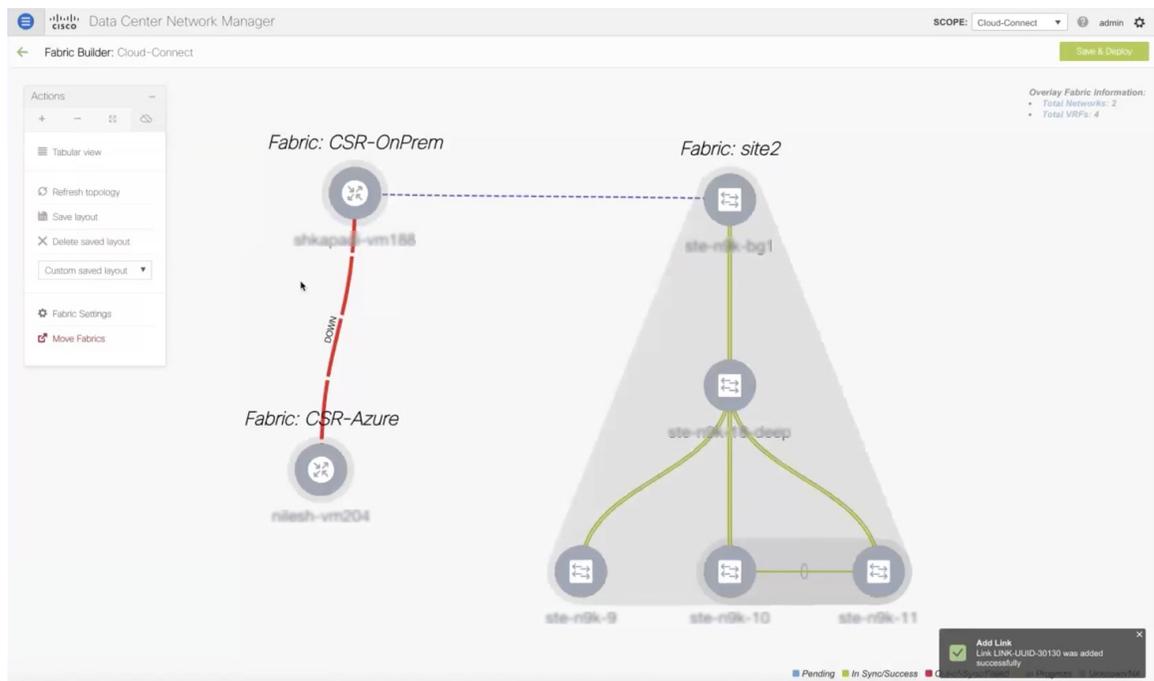
**Step 5** (Optional) In the Link Profile area, choose the **Advanced** tab.

The fields under this tab have default values populated. Change the values if needed. This will create a loopback for which the eBGP peering is configured between the two core routers.

**Step 6** Click **Save**.

The fabric topology window refreshes, and a link is added between the core routers in the **CSR-OnPrem** fabric and the **CSR-Azure** fabric.

**Note** The link will be down till you push it into the configuration.



**What to do next**

Connect the on-prem BGW and the public-cloud core router.

**Connecting the On-prem BGW and the Public-cloud Core Router using EVPN Peering**

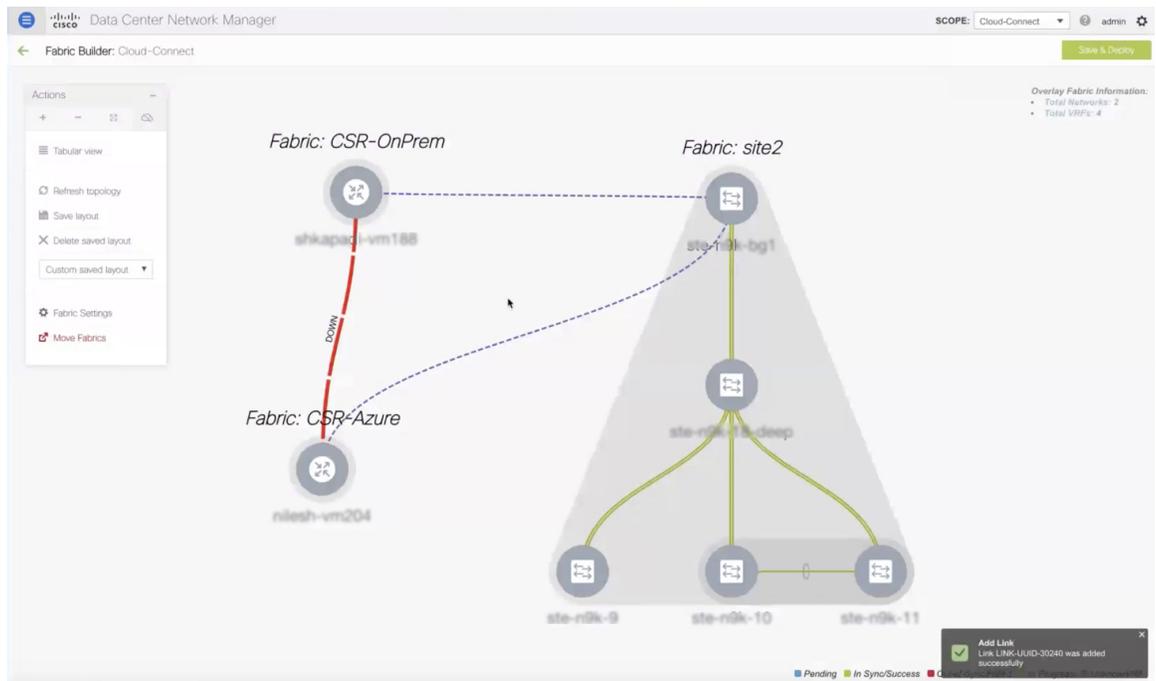
To add a link between the on-prem core router and the public-cloud core router, perform the following steps:

**Procedure**

- Step 1** Right-click anywhere in the **Cloud-Connect** topology window.
- The actions that you can perform in the fabric appears in a list. Alternatively, from the fabric topology window, choose **Tabular view** in the **Actions** pane, and click the **Links** tab.
- Step 2** Choose **Add Link**.
- The **Link Management - Add Link** dialog box appears.
- Step 3** Enter values for the following fields:

| Field                 | Description  |
|-----------------------|--|
| Link Type             | Choose the <b>Inter-Fabric</b> link type from the drop-down list.  |
| Link Sub-Type         | Choose the <b>MULTISITE_OVERLAY</b> link sub-type from the drop-down list.   |
| Link Template         | Choose the <b>csr_ext_evpn_multisite_overlay_setup</b> link template from the drop-down list.  |
| Source Fabric         | Choose <b>site2</b> as the source fabric from the drop-down list.  |
| Destination Fabric    | Choose <b>CSR-Azure</b> as the destination fabric from the drop-down list.   |
| Source Device         | Choose the on-prem BGW from the drop-down list.  |
| Source Interface      | Choose the on-prem BGW's loopback interface.   |
| Destination Device    | Choose the public-cloud core router from the drop-down list.   |
| Destination Interface | Choose the public-cloud core router's interface from the drop-down list.<br><br><b>Note</b> If you did not create an interface, the destination interface will not appear in the drop-down list and you have to enter the destination interface. |

- Step 4** Click **Save**.
- The fabric topology window refreshes, and a link is added between the BGW in the **site2** fabric and the core router in the **CSR-Azure** fabric.
- Note** The link will be down till you push it into the configuration.



### What to do next

Save and deploy the configurations.

## Saving and Deploying Configurations

To save and deploy the configurations in the fabric topology window, perform the following steps:

### Procedure

- Step 1** Click **Save & Deploy**.  
The **Config Deployment** dialog box appears, and you will see the **Configuration Preview** step. The intents for the links created among the BGW, on-prem data center, and the public cloud are generated.
- Step 2** (Optional) Click the field against the BGW in the **Preview Config** column.  
The **Config Preview** dialog box appears for the BGW.
- Step 3** (Optional) View the configuration details in the **Pending Config** column.  
It includes details about the underlay peering and overlay peering.
- Step 4** (Optional) Click the field against the on-prem core router in the **Preview Config** column.  
The **Config Preview** dialog box appears for the on-prem core router.
- Step 5** (Optional) View the configuration details in the **Pending Config** column.

It includes details about the interfaces, the IPsec tunnel, shared key, BGP peering between the core routers, and EVPN peering. Route maps are added indicating that all the BGP traffic and the data traffic should go through the tunnel.

**Step 6** (Optional) Click the field against the public cloud core router in the **Preview Config** column.

The **Config Preview** dialog box appears for the on-prem core router.

**Step 7** (Optional) View the configuration details in the **Pending Config** column.

It includes the details about VTEPs in addition to the details mentioned for the on-prem core router.

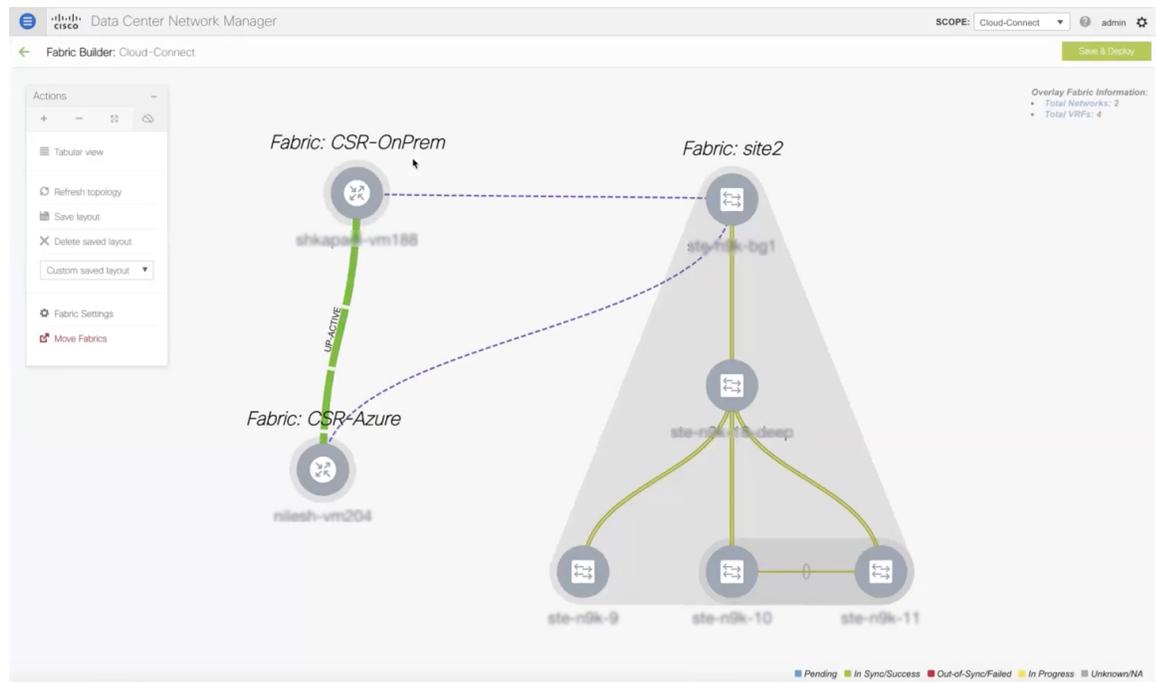
**Step 8** Click **Deploy Config**.

The **Configuration Deployment Status** step appears, where you can see the deployment status of the configurations.

**Step 9** Click **Close** after the successful deployment.

The fabric topology window appears. The IPsec tunnel will be up and active.

**Note** The deployment might take some time.



### What to do next

Extend VRFs and deploy them.

## Extending VRFs

VRFs are extended so that the workloads can be shared between the data center and the public cloud.

## Deploying and Extending the VRF On-prem Core Router

To extend a VRF and deploy it on the on-prem core router from the fabric topology window of the MSD fabric, perform the following steps:

### Procedure

- Step 1** Click the **Total VRF** link in the **Overlay Fabric Information** area, which is below the **Save & Deploy** icon. The **Network / VRF Selection** area of the VRFs window appears for the fabric.
- Step 2** Choose the VRF for the on-prem core router and click **Continue**. The **Network / VRF Deployment** area of the VRFs window appears. The network topology of the fabric appears. You can hide the undiscovered cloud.
- Step 3** Double-click the BGW. The **VRF Extension Attachment** dialog box appears.
- Step 4** Choose the BGW and click the edit icon under the **Extend** column, to enable multi-site on it. A drop-down list appears under the **Extend** column.
- Step 5** Choose **MULTISITE** from the drop-down list.
- Step 6** Enter the loopback ID and the loopback IPv4 address under the respective columns to simulate the host on BGW.

VRF Extension Attachment - Attach extensions for given switch(es) ✕

Fabric Name: Cloud-Connect  
Deployment Options

① Select the row and click on the cell to edit and save changes

| MyVRF_50000 | CLI Freeform | Status | Loopback Id | Loopback IPv4 Address | Loopback IPv6 Address |
|-------------|--------------|--------|-------------|-----------------------|-----------------------|
| ▼           |              | NA     | 101         | 14.14.14.14           |                       |

**Save**

- Step 7** Click **Save**. The network topology of the fabric appears and the BGW will turn blue indicating that the deployment is pending.
- Step 8** Click the preview option.

The **Preview Configuration** dialog box appears. The EVPN configurations are pushed and the loopback interface is created.

**Step 9** Click **Deploy**.

---

#### What to do next

Create a VRF and deploy it on the public cloud.

## Creating and Deploying VRF on Public Cloud

To extend a VRF and deploy it on the public cloud core router from the fabric topology window, perform the following steps:

#### Before you begin

Ensure the VM is up and running. The VM should be attached to the public-cloud core router.

#### Procedure

---

- Step 1** Choose the **CSR-Azure** fabric from the **Fabric Builder** window.  
The fabric topology window appears.
- Step 2** Right-click the public cloud core router.  
A list of actions that you can perform on the router appears.
- Step 3** Choose **View/edit policies** from the list.  
The **View/Edit Policies** dialog box appears.
- Step 4** Click the **Add Policy** icon.  
The **Add Policy** dialog box appears.
- Step 5** Choose the **csr\_vrf\_evpn** policy from the **Policy** drop-down list.
- Step 6** Enter values in mandatory fields in the **General** tab.
- Step 7** Click **Save**.  
The **View/Edit Policies** dialog box appears.
- Step 8** Click **View All** to view the networks and interfaces created.  
The **Generated Config** dialog box appears. Details about the VRF, bridge domain, and the mapped VNI can also be viewed in this dialog box.
- 

#### What to do next

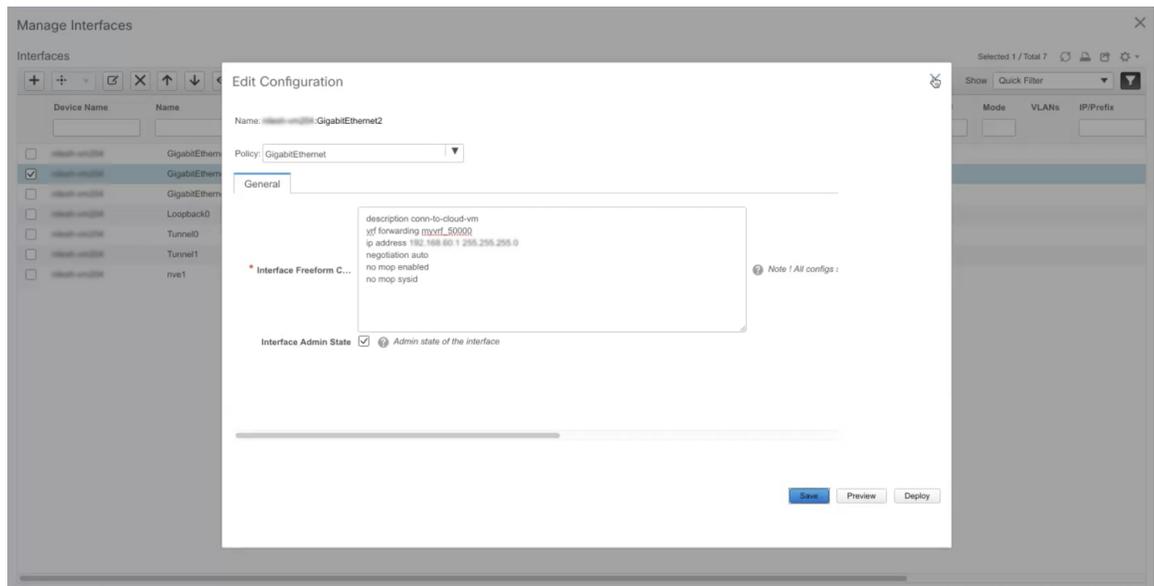
Configure a default gateway on the public-cloud core router for the VM in the public cloud.

## Configuring Default Gateway for the VM

To configure a default gateway on the public-cloud core router from the fabric topology window, perform the following steps:

### Procedure

- Step 1** Choose the **CSR-Azure** fabric from the Fabric Builder window.  
The fabric topology window appears.
- Step 2** Right-click the public-cloud core router.  
A list of actions that you can perform on the router appears.
- Step 3** Choose **Manage Interfaces** from the list.  
The **Manage Interfaces** dialog box appears.
- Step 4** Click **Edit Configuration** to edit the interface for which the policy is created.  
The **Edit Configuration** dialog box appears.
- Step 5** Edit the freeform config, click **Save**, and close the **Manage Interfaces** dialog box.



The fabric topology window appears.

- Step 6** Right-click the public-cloud core router and choose **Deploy Config** from the list.  
The **Config Deployment** dialog box appears.
- Step 7** Click the value under the **Preview Config** column to check the preview configuration.
- Step 8** Click **Deploy Config** to deploy the configuration.  
The configuration will be pushed and deployed.

- Step 9** Click **Close**.
- Step 10** Log on to the CLI to view the traffic flow.  
The traffic flows between the core routers and through the VRF.
- 

## Verifying the Connectivity

To verify the connectivity between the on-prem data center and the public cloud from Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Fabrics > VRFs**.  
The **VRFs** window appears.
- Step 2** Choose the **Cloud-Connect** fabric.  
VRFs in this fabric are listed.
- Step 3** Choose the VRF and click **Continue**.
- Step 4** Right-click the BGW.  
The **VRF Extension Attachment** dialog box appears.
- Step 5** Uncheck the check box and click **Save**.  
The network topology window appears.
- Step 6** Click **Deploy** to push the configurations.  
The VRF is disabled on the BGW.
- Step 7** Check the CLI.  
The traffic will stop.
- Step 8** Enable the VRF again on BGW.
- Step 9** Check the CLI.  
The traffic will flow. Alternatively, access the HTTP address of the web server in the public cloud. You will get a **Database Reachable** message.
- 

## Deploying Cisco CSR 1000v on Microsoft Azure

To deploy a Cisco CSR 1000v in Microsoft Azure, perform the following steps:

## Procedure

- Step 1** From the **Microsoft Azure** UI, choose **Virtual Machines**.  
The **Virtual Machines** window appears.
- Step 2** Click **Add**.  
The **Create a virtual machine** window appears.
- Step 3** Click the **Create VM from Azure Marketplace** hyperlink.  
The **Marketplace** window appears, where you can search for the standard classic VMs.
- Step 4** Search for the CSR deployments in the marketplace.
- Step 5** Choose **Cisco Cloud Services Router (CSR) 1000V** from the search results.
- Step 6** Choose **Cisco CSR 1000V Bring Your Own License – XE 16.9** or higher versions from the **Select a software plan** drop-down list.
- Step 7** Click **Create**.
- Step 8** Enter the project details and instance details in the **Create a virtual machine** window.
- Step 9** Choose the **Password** authentication type in the administrator account section.  
Cisco DCNM does not support the SSH public key.
- Step 10** Create a username and password.

The screenshot shows the 'Create a virtual machine' page in the Microsoft Azure portal. The left sidebar contains navigation options like 'Home', 'Dashboard', and 'All services'. The main content area is titled 'Create a virtual machine' and includes a breadcrumb trail: Home > Virtual machines > Create a virtual machine > Marketplace > Cisco Cloud Services Router (CSR) 1000V > Create a virtual machine. Below the breadcrumb, there's a section for 'Subscription' (Pay-As-You-Go) and 'Resource group' (demo-csr2). The 'INSTANCE DETAILS' section includes 'Virtual machine name' (csr3), 'Region' ((US) West US), 'Availability options' (No infrastructure redundancy required), 'Image' (Cisco CSR 1000V Bring Your Own License - XE 16.9), and 'Size' (Standard DS2 v2, 2 vcpus, 7 GiB memory). The 'ADMINISTRATOR ACCOUNT' section shows 'Authentication type' set to Password, with fields for 'Username' (cisco), 'Password', and 'Confirm password'. A 'Review + create' button is at the bottom, along with navigation arrows for '< Previous' and 'Next : Disks >'. A tooltip at the bottom right states 'Password and confirm password must match.'

- Step 11** Click **Next : Disks >**.
- Step 12** Choose the **Standard HDD** option from the OS disk type drop-down list.

- Step 13** Click **Next : Networking** >.
- Step 14** Enter values in the required fields.
- Step 15** Choose a public IP for the network.

Home > Virtual machines > Create a virtual machine > Marketplace > Cisco Cloud Services Router (CSR) 1000V > Create a virtual machine

## Create a virtual machine

Basics Disks **Networking** Management Advanced Tags Review + create

Define network connectivity for your virtual machine by configuring network interface card (NIC) settings. You can control ports, inbound and outbound connectivity with security group rules, or place behind an existing load balancing solution. [Learn more](#)

**NETWORK INTERFACE**

When creating a virtual machine, a network interface will be created for you.

\* Virtual network ⓘ demo-csr2    
[Create new](#)

\* Subnet ⓘ subnet1 (10.1.0.0/24)    
[Manage subnet configuration](#)

Public IP ⓘ (new) csr3-ip    
[Create new](#)

NIC network security group ⓘ  None  Basic  Advanced   
 This VM image has preconfigured NSG rules

The selected subnet 'subnet1 (10.1.0.0/24)' is already associated to a network security group 'demo-csr2-SSH-SecurityGroup'. We recommend managing connectivity to this virtual machine via the existing network security group instead of creating a new one here.

\* Configure network security group (new) csr3-nsg    
[Create new](#)

Accelerated networking ⓘ  On  Off  
 The selected image does not support accelerated networking.

- Step 16** Use the default values in other fields.
- Step 17** Click **Review + create**.
- A VM will be created for Cisco CSR 1000v in Microsoft Azure with a public IP address.

## What to do next

- Attach network interfaces:

1. Choose the **Networking** setting of the VM.
2. Choose **Attach network interface** to add a Nic.

Attach one Nic each for both the subnets. IP addresses are automatically assigned.

3. Add an SSH rule using the port 22 to enable the SSH access of the core router.  
Cisco DCNM discovers the core router using this SSH access.



**Note** Two UDP rules using the ports 500 and 4500 to enable the IPsec tunnel are added automatically.

Home > Virtual machines > demo-csr2 - Networking

demo-csr2 - Networking

Virtual machine

Search (Ctrl+F)

Attach network interface Detach network interface

demo-csr2-Nic0-newVnet demo-csr2-Nic1-newVnet

**Network Interface: demo-csr2-Nic0-newVnet** Effective security rules Topology  
Virtual network/subnet: demo-csr2/subnet1 NIC Public IP: 104.42.181.20 NIC Private IP: 10.1.0.4 Accelerated networking: Disabled

Inbound port rules Outbound port rules Application security groups Load balancing

Network security group demo-csr2-SSH-SecurityGroup (attached to subnet: subnet1)  
Impacts 1 subnets, 2 network interfaces [Add inbound port rule](#)

| PRIORITY | NAME                          | PORT | PROTOCOL | SOURCE            | DESTINATION    | ACTION |
|----------|-------------------------------|------|----------|-------------------|----------------|--------|
| 100      | SSH-Rule                      | 22   | TCP      | Internet          | Any            | Allow  |
| 101      | UDP-Rule1                     | 500  | UDP      | Internet          | Any            | Allow  |
| 102      | UDP-Rule2                     | 4500 | UDP      | Internet          | Any            | Allow  |
| 65000    | AllowVnetInBound              | Any  | Any      | VirtualNetwork    | VirtualNetwork | Allow  |
| 65001    | AllowAzureLoadBalancerInBound | Any  | Any      | AzureLoadBalancer | Any            | Allow  |
| 65500    | DenyAllInBound                | Any  | Any      | Any               | Any            | Deny   |

Network security group demo-csr2-SSH-SecurityGroup (attached to network interface: demo-csr2-Nic0-newVnet)  
Impacts 1 subnets, 2 network interfaces [Add inbound port rule](#)

| PRIORITY | NAME      | PORT | PROTOCOL | SOURCE   | DESTINATION | ACTION |
|----------|-----------|------|----------|----------|-------------|--------|
| 100      | SSH-Rule  | 22   | TCP      | Internet | Any         | Allow  |
| 101      | UDP-Rule1 | 500  | UDP      | Internet | Any         | Allow  |

- Create routes in the **Routes** setting of the VM to create traffic routes between the on-prem data center and Microsoft Azure. You can use the default route to redirect traffic from the VNet to Cisco CSR 1000v.

Home > subnet2-CSR-RouteTable

**subnet2-CSR-RouteTable**  
Route table

Search (Ctrl+F)

Move Delete Refresh

Resource group (change) : demo-csr2 Associations : 1 subnet associations

Location : West US

Subscription (change) : Pay-As-You-Go

Subscription ID : 1cda121a-974e-4166-9625-a1e5f69bec73

Tags (change) : Click here to add tags

**Routes**

Search routes

| NAME                       | ADDRESS PREFIX   | NEXT HOP |     |
|----------------------------|------------------|----------|-----|
| Route-to-192.168.202.0-AWS | 192.168.202.0/24 | 10.1.1.4 | ... |
| Route-to-Onprem            | 10.200.0.0/24    | 10.1.1.4 | ... |

**Subnets**

Search subnets

| NAME    | ADDRESS RANGE | VIRTUAL NETWORK | SECURITY GROUP |     |
|---------|---------------|-----------------|----------------|-----|
| subnet2 | 10.1.1.0/24   | demo-csr2       | -              | ... |

See *Cisco CSR 1000v Deployment Guide for Microsoft Azure* for more information.

## Viewing Links and Core Routers Details

To view the details of links and core routers from the fabric topology window, perform the following steps:

### Procedure

- Step 1** From the **Actions** pane, choose **Tabular view > Links**.  
The **Links** window appears.
- Step 2** Refresh the window.  
The three links that you created will appear in the list.
- Step 3** (Optional) Double-click the on-prem core router to view the IP route information.  
The **IP Route Information** dialog box appears.
- Step 4** (Optional) Click the **Crypto Session** tab to view the details about the IPsec tunnel.
- Step 5** (Optional) Click the **BGP Session** tab to view the details about the BGP session.
- Step 6** (Optional) Click the **Packet Counter** tab to view the packet counter details.

You can reset the counter value you see in the **Packet Counter** tab. See the [Resetting Packet Counter Using API, on page 443](#) section more information.

## Resetting Packet Counter Using API

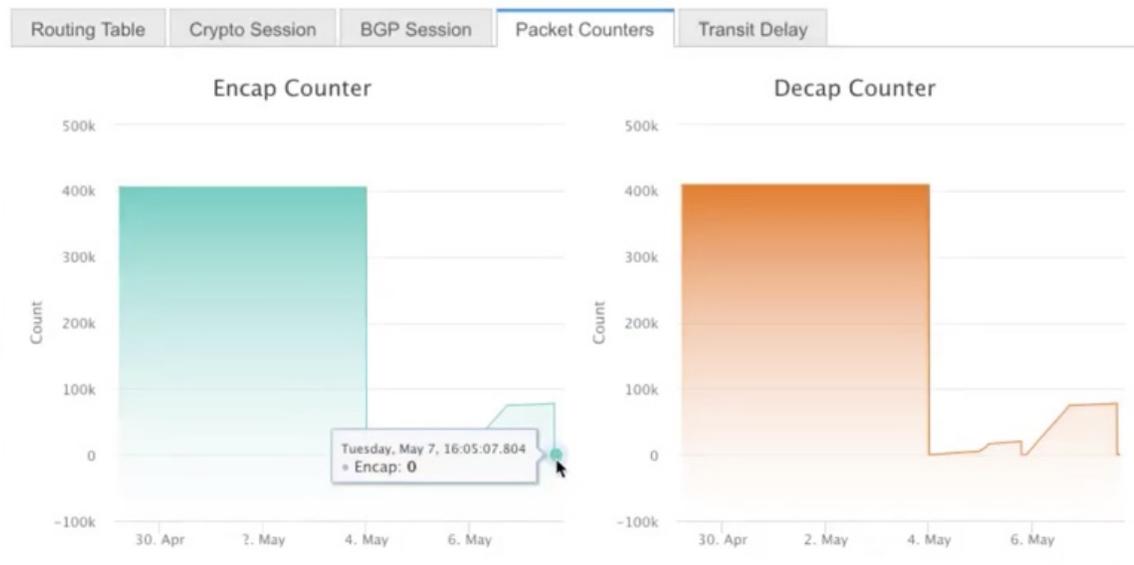
To reset the packet counter, perform the following steps:

## Procedure

- Step 1** Log into Cisco DCNM.
- Step 2** Navigate to the `https://DCNM-IP/api-docs` URL.
- Step 3** Expand the `GET /cloud-extension/status/{ipAddress}` API under cloud extension.
- Step 4** Enter the IP address of the on-prem core router.
- Step 5** Set the `fetchLatestFromSwitch` value to `true`.
- Step 6** Click **Try it out**.

The packet counter is cleared and the count drops to zero.

### IP Route Information





## CHAPTER 9

# Managing a Brownfield VXLAN BGP EVPN Fabric

This chapter explains how to migrate a Brownfield fabric into Cisco DCNM.

- [Overview, on page 445](#)
- [Prerequisites, on page 446](#)
- [Guidelines and Limitations, on page 446](#)
- [Fabric Topology Overview, on page 448](#)
- [DCNM Brownfield Deployment Tasks, on page 449](#)
- [Verifying the Existing VXLAN BGP EVPN Fabric, on page 449](#)
- [Creating a VXLAN BGP EVPN Fabric, on page 452](#)
- [Adding Switches and Transitioning VXLAN Fabric Management to DCNM, on page 460](#)
- [Verifying the Import of the VXLAN BGP EVPN Fabric, on page 473](#)
- [Migrating a Bottom-Up VXLAN Fabric to DCNM, on page 480](#)
- [Resolving Config Compliance Error on Switches with Cisco NX-OS Release 7.0\(3\)I4\(8b\) and 7.0\(4\)I4\(x\) Images, on page 488](#)
- [Modifying VLAN Names in a Switch with Cisco NX-OS Release 7.0\(3\)I4\(8b\) and 7.0\(4\)I4\(x\) Images, on page 492](#)
- [Changing a Brownfield Imported BIDIR Configuration, on page 494](#)
- [Manually Adding PIM-BIDIR Configuration for Leaf or Spine Post Brownfield Migration , on page 495](#)
- [Migrating an MSD Fabric with Border Gateway Switches , on page 495](#)

## Overview

This use case shows how to migrate an existing VXLAN BGP EVPN fabric to Cisco DCNM. The transition involves migrating existing network configurations to DCNM.

Typically, your fabric would be created and managed through manual CLI configuration or custom automation scripts. Now, you can start managing the fabric through DCNM. After the migration, the fabric underlay and overlay networks will be managed by DCNM.

For information about MSD fabric migration, see *Migrating an MSD Fabric with Border Gateway Switches*.

## Prerequisites

- DCNM-supported NX-OS software versions. For details, refer *Cisco DCNM Release Notes, Release 11.2(1)*.
- Underlay routing protocol is OSPF or IS-IS.
- The supported underlay is based on the DCNM 10.2(1) POAP template's best practices for the VXLAN fabric (dcnm\_ip\_vxlan\_fabric\_templates.10.2.1.ST.1.zip) available on Cisco.com.
- The following fabric-wide loopback interface IDs must not overlap:
  - Routing loopback interface for IGP/BGP.
  - VTEP loopback ID
  - Underlay rendezvous point loopback ID if ASM is used for multicast replication.
- BGP configuration uses the 'router-id', which is the IP address of the routing loopback interface.
- If the iBGP peer template is configured, then it must be configured on the leaf switches and route reflectors. The template name that needs to be used between leaf and route reflector should be identical.
- The BGP route reflector and multicast rendezvous point (if applicable) functions are implemented on spine switches. Leaf switches do not support the functions.
- Familiarity with VXLAN BGP EVPN fabric concepts and functioning of the fabric from the DCNM perspective.
- Fabric switch nodes are operationally stable and functional and all fabric links are up.
- vPC switches and the peer links are up before the migration. Ensure that no configuration updates are in progress or changes pending.
- Create an inventory list of the switches in the fabric with their IP addresses and credentials. DCNM uses this information to connect to the switches.
- Shut down any other controller software you are using presently so that no further configuration changes are made to the VXLAN fabric. Alternatively, disconnect the network interfaces from the controller software (if any) so that no changes are allowed on the switches.
- The switch overlay configurations must have the mandatory configurations defined in the shipping DCNM Universal Overlay profiles. Additional network or VRF overlay related configurations found on the switches are preserved in the freeform configuration associated with the network or VRF DCNM entries.
- All the overlay network and VRF profile parameters such as VLAN name and route map name should be consistent across all devices in the fabric for the brownfield migration to be successful.

## Guidelines and Limitations

- Fabric interfaces can be numbered or unnumbered.
- Various other interface types are supported.

- The following features are unsupported.
  - eBGP underlay
  - Layer 3 port channel
  - Configuration profiles present in the brownfield configurations (the expectation is that the overlays should be configured through regular CLIs).
  - vPC Fabric Peering
- Take a backup of the switch configurations and save them before the migration.
- No configuration changes (unless instructed to do so in this document) must be made to the switches until the migration is completed. Else, significant network issues can occur.
- Migration to Cisco DCNM is only supported for Cisco Nexus 9000 switches.
- Multi-line banner configuration on the switch is preserved in the `switch_freeform` configuration, along with other configurations captured in the `switch_freeform` configuration, if any.
- From DCNM Release 11.2(1), the Border Spine and Border Gateway Spine roles are supported for the brownfield migration.
- Fabrics with IS-IS Level-1 and Level-2 are supported for the Brownfield migration.
- Switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images support the Brownfield migration. For information about feature compatibility, refer your respective platform documentation and for information about the supported software images, see *Compatibility Matrix for Cisco DCNM*.

Note the following guidelines and limitations:

- The VLAN name for the network or VRF is not captured in the overlay profile if at least one of the non-spine switches have the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images. The VLAN name is captured in the freeform config associated with the overlay network or VRF. The VLAN name can be changed by updating the freeform config. For more information, see *Modifying VLAN Names in a Switch with Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) Images*.
- Config compliance difference for TCAM CLIs on Cisco Nexus 9300 Series switches and Cisco Nexus 9500 Series switches with X9500 line cards. For more information, see *Resolving Config Compliance Error on Switches with Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) Images*.
- The overlay profile refresh feature is unsupported for the brownfield migration of switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images.
- Cisco Nexus 9500 Series Switches are supported as VTEPs with border spine, BGW spine, or leaf roles for Cisco NX-OS Release 7.0.3.I7(3) or later.
- During the brownfield migration in the Cisco DCNM Release 11.1(1), the overlay configuration profiles are deployed to switches and all the overlay related configurations are captured in the respective network or VRF freeform configs. Post migration, switches have both the original configuration CLIs and the config-profiles.

From Cisco DCNM Release 11.2(1), during the brownfield migration, the overlay config-profiles are deployed to the switches, and the original configuration CLIs are removed. Post migration, the switches only have the configuration profiles and any extra configuration that is not part of the configuration profile if the switches in the brownfield migration have the following Cisco NX-OS images:

- Cisco NX-OS Release 7.0(3)I7(6) or newer
- Cisco NX-OS Release 9.2(3) or newer

If the switches do not meet these requirements, the brownfield migration behavior is the same as described for the Cisco DCNM Release 11.1(1).

- First, guidelines for updating the settings are noted. Then each VXLAN fabric settings tab is explained:
  - Some values (BGP AS Number, OSPF, etc) are considered as reference points to your existing fabric, and the values you enter must match the existing fabric values.
  - For some fields (such as IP address range, VXLAN ID range), the values that are auto-populated or entered in the settings are only used for future allocation. The existing fabric values are honored during migration.
  - Some fields relate to new functions that may not exist in your existing fabric (such as advertise-pip). Enable or disable it as per your need.
  - At a later point in time, after the fabric transition is complete, you can update settings if needed.

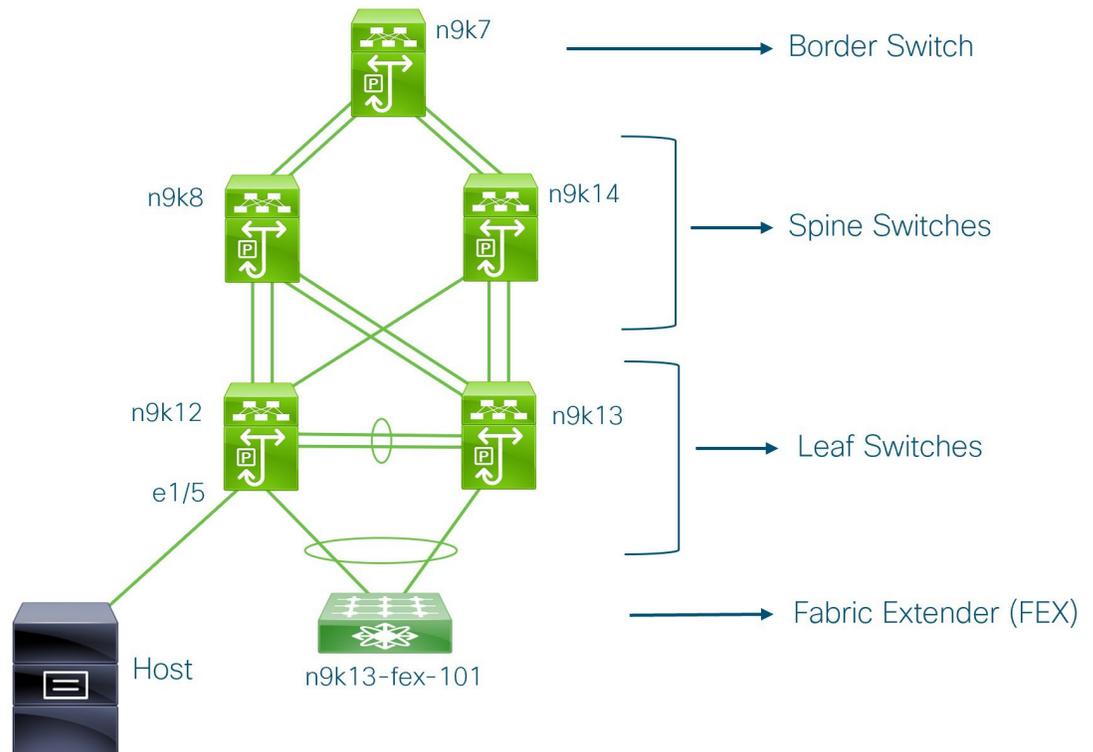
## Fabric Topology Overview

This example use case uses the following hardware and software components:

- Five Cisco Nexus 9000 Series Switches running NX-OS Release 7.0(3)I7(6)
- One Fabric Extender or FEX
- One host

For information about the supported software images, see [Compatibility Matrix for Cisco DCNM](#).

Before we start the transition of the existing fabric, let us see its topology.



You can see that there is a border switch, two spine switches, two leaf switches, and a Fabric Extender or FEX.

A host is connected to the n9k12 leaf switch through the interface Ethernet 1/5.

## DCNM Brownfield Deployment Tasks

The following tasks are involved in a Brownfield migration:

1. [Verifying the Existing VXLAN BGP EVPN Fabric, on page 449](#)
2. [Creating a VXLAN BGP EVPN Fabric, on page 452](#)
3. [Adding Switches and Transitioning VXLAN Fabric Management to DCNM, on page 460](#)
4. [Verifying the Import of the VXLAN BGP EVPN Fabric, on page 473](#)

## Verifying the Existing VXLAN BGP EVPN Fabric

Let us check the network connectivity of the **n9k12** switch from the console terminal.

## Procedure

### Step 1 Verify the Network Virtual Interface or NVE in the fabric.

```
n9k12# show nve vni summary
Codes: CP - Control Plane      DP - Data Plane
       UC - Unconfigured
```

```
Total CP VNIs: 84    [Up: 84, Down: 0]
Total DP VNIs: 0     [Up: 0, Down: 0]
```

There are 84 VNIs in the control plane and they are up. Before the Brownfield migration, make sure that all the VNIs are up.

### Step 2 Check consistency and failures of vPC.

```
n9k12# show vpc
Legend:
          (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id           : 2
Peer status              : peer adjacency formed ok
vPC keep-alive status   : peer is alive
Configuration consistency status : success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role                 : secondary
Number of vPCs configured : 40
Peer Gateway             : Enabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status    : Enabled, timer is off.(timeout = 300s)
Delay-restore status    : Timer is off.(timeout = 60s)
Delay-restore SVI status : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router : Disabled
.
.
.
```

### Step 3 Check the EVPN neighbors of the n9k-12 switch.

```
n9k12# show bgp 12vpn evpn summary
BGP summary information for VRF default, address family L2VPN EVPN
BGP router identifier 192.168.0.4, local AS number 65000
BGP table version is 637, L2VPN EVPN config peers 2, capable peers 2
243 network entries and 318 paths using 57348 bytes of memory
BGP attribute entries [234/37440], BGP AS path entries [0/0]
BGP community entries [0/0], BGP clusterlist entries [2/8]

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.0.0   4 65000   250    91      637   0   0 01:26:59 75
192.168.0.1   4 65000   221    63      637   0   0 00:57:22 75
```

You can see that there are two neighbors corresponding to the spine switches.

Note that the ASN is 65000.

### Step 4 Verify the VRF information.

```
n9k12# show run vrf internet

!Command: show running-config vrf Internet
!Running configuration last done at: Fri Aug  9 01:38:02 2019
```

```

!Time: Fri Aug  9 02:48:03 2019

version 7.0(3)I7(6) Bios:version 07.59

interface Vlan347
  vrf member Internet

interface Vlan349
  vrf member Internet

interface Vlan3962
  vrf member Internet

interface Ethernet1/25
  vrf member Internet

interface Ethernet1/26
  vrf member Internet
vrf context Internet
  description Internet
  vni 16777210
  ip route 204.90.141.0/24 204.90.140.129 name LC-Networks
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
router ospf 300
  vrf Internet
    router-id 204.90.140.3
    redistribute direct route-map allow
    redistribute static route-map static-to-ospf
router bgp 65000
  vrf Internet
    address-family ipv4 unicast
      advertise l2vpn evpn

```

The VRF **Internet** is configured on this switch.

The host connected to the **n9k-12** switch is part of the VRF **Internet**.

You can see the VLANs associated with this VRF.

Specifically, the host is part of **Vlan349**.

**Step 5** Verify the layer 3 interface information.

```
n9k12# show run interface vlan349
```

```

!Command: show running-config interface Vlan349
!Running configuration last done at: Fri Aug  9 01:38:02 2019
!Time: Fri Aug  9 02:49:27 2019

version 7.0(3)I7(6) Bios:version 07.59

interface Vlan349
  no shutdown
  vrf member Internet
  no ip redirects
  ip address 204.90.140.134/29
  no ipv6 redirects
  fabric forwarding mode anycast-gateway

```

Note that the IP address is **204.90.140.134**. This IP address is configured as the anycast gateway IP.

- Step 6** Verify the physical interface information. This switch is connected to the Host through the interface Ethernet 1/5.

```
n9k12# show run interface ethernet1/5

!Command: show running-config interface Ethernet1/5
!Running configuration last done at: Fri Aug  9 01:38:02 2019
!Time: Fri Aug  9 02:50:05 2019

version 7.0(3)I7(6) Bios:version 07.59

interface Ethernet1/5
  description to host
  switchport mode trunk
  switchport trunk native vlan 349
  switchport trunk allowed vlan 349,800,815
  spanning-tree bpduguard enable
  mtu 9050
```

You can see that this interface is connected to the host and is configured with VLAN 349.

- Step 7** Verify the connectivity from the host to the anycast gateway IP address.

```
host# ping 204.90.140.134 count unlimited interval 1
PING 204.90.140.134 (204.90.140.134): 56 data bytes
64 bytes from 204.90.140.134: icmp_seq=0 ttl=254 time=1.078 ms
64 bytes from 204.90.140.134: icmp_seq=1 ttl=254 time=1.129 ms
64 bytes from 204.90.140.134: icmp_seq=2 ttl=254 time=1.151 ms
64 bytes from 204.90.140.134: icmp_seq=3 ttl=254 time=1.162 ms
64 bytes from 204.90.140.134: icmp_seq=4 ttl=254 time=1.84 ms
64 bytes from 204.90.140.134: icmp_seq=5 ttl=254 time=1.258 ms
64 bytes from 204.90.140.134: icmp_seq=6 ttl=254 time=1.273 ms
64 bytes from 204.90.140.134: icmp_seq=7 ttl=254 time=1.143 ms
```

We let the ping command run in the background while we transition the existing brownfield fabric into DCNM.

## Creating a VXLAN BGP EVPN Fabric

This procedure describes how to create a VXLAN BGP EVPN fabric in DCNM.

### Procedure

- Step 1** Choose **Control > Fabric Builder**.

The **Fabric Builder** screen appears. When you log in for the first time, the **Fabrics** section has no entries. After you create a fabric, it is displayed on the **Fabric Builder** screen, wherein a rectangular box represents each fabric.

A standalone or member fabric contains Switch\_Fabric (in the Type field), the AS number (in the ASN field), and mode of replication (in the Replication Mode field).

- Step 2** Click **Create Fabric**. The **Add Fabric** window appears.

**Fabric Template** - From the drop-down menu, choose the **Easy\_Fabric\_11\_1** fabric template. The fabric settings for creating a standalone fabric comes up.

**Fabric Name** - Enter the name of the fabric.

The tabs and their fields in the screen are explained in the subsequent points. The overlay and underlay network parameters are included in these tabs.

**Note** If you are creating a standalone fabric as a potential member fabric of an MSD fabric (used for provisioning overlay networks for fabrics that are connected through EVPN Multi-Site technology), then browse through the Multi-Site Domain for VXLAN BGP EVPN Fabrics topic before member fabric creation.

### Step 3

The **General** tab is displayed by default. The fields in this tab are:

**BGP ASN:** Enter the BGP AS number the fabric is associated with.

**Enable IPv6 Underlay:** Select this check box to enable the IPv6 underlay feature.

Brownfield migration is supported for the VXLANv6 fabrics. Note that L3 vPC keep-alive using IPv6 address is not supported for brownfield migration. This vPC configuration is deleted after the migration. However, L3 vPC keep-alive using IPv4 address is supported.

**Fabric Interface Numbering:** Specify whether you are using a point-to-point (p2p) or unnumbered network in your existing setup.

**Underlay Subnet IP Mask** - Specify the subnet mask you are using for the fabric underlay IP address subnets in your existing setup.

**Link-State Routing Protocol:** The IGP used in the existing fabric, OSPF, or IS-IS.

**Route-Reflectors** – The Route Reflector count is only applicable post-migration. The existing route reflector configuration is honored when importing into the DCNM setup.

The number of spine switches that are used as route reflectors for transporting BGP traffic. Choose 2 or 4 from the drop-down box. The default value is 2.

To deploy spine devices as route reflectors, DCNM sorts the spine devices based on their serial numbers, and designates two or four spine devices as route reflectors. If you add more spine devices, existing route reflector configuration will not change.

*Increasing the count* - You can increase the route reflectors from two to four at any point in time. Configurations are automatically generated on the other 2 spine devices designated as route reflectors.

*Decreasing the count*

When you reduce four route reflectors to two, you must remove the unneeded route reflector devices from the fabric. Follow these steps to reduce the count from 4 to 2.

a. Change the value in the drop-down box to 2.

b. Identify the spine switches designated as route reflectors.

An instance of the **rr\_state** policy is applied on the spine switch if it is a route reflector. To find out if the policy is applied on the switch, right-click the switch, and choose **View/edit policies**. In the View/Edit Policies screen, search **rr\_state** in the **Template** field. It is displayed on the screen.

c. Delete the unneeded spine devices from the fabric (right-click the spine switch icon and choose **Discovery > Remove from fabric**).

If you delete existing route reflector devices, the next available spine switch is selected as the replacement route reflector.

d. Click Save and Deploy at the top right part of the fabric topology screen.

You can preselect RRs and RPs before performing the first **Save & Deploy** operation. For more information, see *Preselecting Switches as Route-Reflectors and Rendezvous-Points*.

**Anycast Gateway MAC:** Enter the Anycast gateway MAC address of the existing fabric.

**NX-OS Software Image Version:** Leave this field blank. You can update this post-transition, as desired.

#### Step 4

Click the **Replication** tab. Most of the fields are auto generated.

**Replication Mode:** The mode of replication that is used in the existing fabric, Ingress Replication, or Multicast.

When you choose Ingress replication, the multicast replication fields get disabled.

**Multicast Group Subnet** - The IP address prefix for multicast communication is used for post-migration allocation. The IP address prefix used in your existing fabric is honored during the transition.

A unique IP address is allocated from this group for each overlay network.

**Enable Tenant Routed Multicast** – Select the check box to enable Tenant Routed Multicast (TRM) as the fabric overlay multicast protocol.

If you enable TRM, the Multicast address for TRM must be entered. All the TRM specific tenant configuration is captured in the switch freeform policy linked to the tenant network and VRF profile.

Note that the TRM feature is unsupported on switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images.

**Default MDT address for TRM VRFs** – Enter the default multicast distribution tree (MDT) IPv4 address for TRM VRFs.

**Rendezvous-Points** - Enter the number of spine switches acting as rendezvous points.

**RP mode** – Select **asm** (Any-Source Multicast) or **bidir** (Bidirectional PIM) mode.

When you choose ASM, the BiDir related fields are not enabled.

The **asm** RP mode supports up to 4 RPs.

The **bidir** mode supports up to 2 RPs. An error message is displayed if the BIDIR configuration indicates that more than 2 RPs are used.

After brownfield migration, only 2 RPs are supported in the migrated fabric. An error message is displayed when you click **Save & Deploy** after changing the RP count to 4.

If an RP is down or deleted from the fabric, this RP cannot be replaced by another spine as Easy Fabric does not remember the configuration of a removed switch. Easy Fabric uses a specific scheme to generate RP configuration for Bidir. Therefore, the generated Bidir configuration will not work with the brownfield imported configuration. After brownfield migration, if you change the RP count or add new spine or leaf switches, you should manually configure the PIM-Bidir feature. If a manual configuration is required, a warning message is displayed after you click **Save & Deploy**. For more information, see *Manually Adding PIM-BIDIR Configuration for Leaf or Spine Post Brownfield Migration*.

You can also modify a brownfield imported bidir configuration to use the configuration generated by **Fabric Builder**. For more information, see *Changing a Brownfield Imported BIDIR Configuration*.

**Underlay RP Loopback ID** – The loopback ID has to match your existing setup's loopback ID. This is the loopback ID used for the rendezvous point (RP), for multicast protocol peering purposes in the fabric underlay.

The next two fields are enabled if you choose BIDIR-PIM as the multicast mode of replication.

**Underlay Primary RP Loopback ID** – The primary loopback ID used for the phantom RP, for multicast protocol peering purposes in the fabric underlay.

**Underlay Backup RP Loopback ID** – The secondary loopback ID used for the phantom RP, for multicast protocol peering purposes in the fabric underlay.

The next two fields are enabled if **Rendezvous-Points** is set to 4. However, the fabric can have only 2 RPs for the brownfield migration.

**Underlay Second Backup RP Loopback Id** – The second fallback loopback ID for Phantom RP, for multicast protocol peering purposes in the fabric underlay.

**Underlay Third Backup RP Loopback Id** – The third fallback loopback ID for Phantom RP, for multicast protocol peering purposes in the fabric underlay.

### Step 5

Click the **vPC** tab. Most of the fields are auto generated.

**vPC Peer Link VLAN** - Enter the VLAN ID used for the vPC peer link SVI in the existing fabric.

**vPC Peer Keep Alive option** – Choose the management or loopback option, as used in the existing fabric. If you want to use IP addresses assigned to the management port and the management VRF, choose management. If you use IP addresses assigned to loopback interfaces (and a non-management VRF), choose loopback.

If you only use IPv6 addresses on the management interface, you must use the loopback option.

During the transition, the switch configuration is not checked for the following fields in the vPC tab. The switch configurations will get updated if they are different.

**vPC Auto Recovery Time** - Specify the vPC auto recovery time-out period in seconds, as needed.

**vPC Delay Restore Time** - Specify the vPC delay restore period in seconds, as needed.

**vPC Peer Link Port Channel Number** - Specifies the Port Channel ID for a vPC Peer Link. By default, the value in this field is 500. Change the value based on your existing settings.

**vPC IPv6 ND Synchronize** – Enables IPv6 Neighbor Discovery synchronization between vPC switches. The check box is enabled by default. Clear the check box to disable the function as needed.

**vPC advertise-pip** - Select the check box to enable the Advertise PIP feature.

Note that the Advertise PIP feature is unsupported on switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images.

### Step 6

Click the **Protocols** tab. Most of the fields are auto generated. You can update the fields if needed.

**Underlay Routing Loopback Id** - The loopback interface ID is populated as 0 since loopback0 is usually used for fabric underlay IGP peering purposes. This must match the existing configuration on the switches. This must be the same across all the switches.

**Underlay VTEP Loopback Id** - The loopback interface ID is populated as 1 since loopback1 is usually used for the VTEP peering purposes. This must match the existing configuration on the switches. This must be the same across all the switches where VTEPs are present.

**Link-State Routing Protocol Tag** - Enter the existing fabric's routing protocol tag in this field to define the type of network.

**OSPF Area ID** – The OSPF area ID of the existing fabric, if OSPF is used as the IGP within the fabric.

**Note** The OSPF or IS-IS authentication fields are enabled based on your selection in the **Link-State Routing Protocol** field in the **General** tab.

**Enable OSPF Authentication** – Select the check box to enable the OSPF authentication. Deselect the check box to disable it. If you enable this field, the **OSPF Authentication Key ID** and **OSPF Authentication Key** fields are enabled.

**OSPF Authentication Key ID** – Enter the OSPF authentication key ID.

**OSPF Authentication Key** - The OSPF authentication key must be the 3DES key from the switch.

**Note** Plain text passwords are not supported. Login to the switch, retrieve the OSPF authentication details.

You can obtain the OSPF authentication details by using the **show run ospf** command on your switch.

```
# show run ospf | grep message-digest-key
ip ospf message-digest-key 127 md5 3 c7c83ec78f38f32f3d477519630faf7b
```

In this example, the OSPF authentication key ID is **127** and the authentication key is **c7c83ec78f38f32f3d477519630faf7b**.

For information about how to configure a new key and retrieve it, see *Retrieving the Authentication Key*.

**IS-IS Level** - Select the IS-IS level from this drop-down list.

**Enable IS-IS Authentication** - Select the check box to enable IS-IS authentication. Deselect the check box to disable it. If you enable this field, the IS-IS authentication fields are enabled.

**IS-IS Authentication Keychain Name** - Enter the keychain name.

**IS-IS Authentication Key ID** - Enter the IS-IS authentication key ID.

**IS-IS Authentication Key** - Enter the Cisco Type 7 encrypted key.

**Note** Plain text passwords are not supported. Login to the switch, retrieve the IS-IS authentication details.

You can obtain the IS-IS authentication details by using the **show run | section "key chain"** command on your switch.

```
# show run | section "key chain"
key chain CiscoIsisAuth
  key 127
    key-string 7 075e731f
```

In this example, the keychain name is **CiscoIsisAuth**, the key ID is **127**, and the type 7 authentication key is **075e731f**.

**Enable BGP Authentication** - Select the check box to enable BGP authentication. Deselect the check box to disable it. If you enable this field, the **BGP Authentication Key Encryption Type** and **BGP Authentication Key** fields are enabled.

**BGP Authentication Key Encryption Type** – Choose the 3 for 3DES encryption type, and 7 for Cisco encryption type.

**BGP Authentication Key** - Enter the encrypted key based on the encryption type.

**Note** Plain text passwords are not supported. Login to the switch, retrieve the BGP authentication details.

You can obtain the BGP authentication details by using the **show run bgp** command on your switch.

```
# show run bgp
```

```
neighbor 10.2.0.2
remote-as 65000
password 3 sd8478fswerdfw3434fsw4f4w34sdsd8478fswerdfw3434fsw4f4w3
```

In this example, the BGP authentication key is displayed after the encryption type **3**.

**Enable BFD feature** – Select the check box to enable the BFD feature.

The BFD feature is disabled by default.

Make sure that the BFD feature setting matches with the switch configuration. If the switch configuration contains **feature bfd** but the BFD feature is not enabled in the fabric settings, config compliance generates diff to remove the BFD feature after brownfield migration. That is, **no feature bfd** is generated after migration.

The following config is pushed after you select the **Enable BFD** check box:

```
feature bfd
```

For information about BFD feature compatibility, refer your respective platform documentation and for information about the supported software images, see *Compatibility Matrix for Cisco DCNM*.

**Enable BFD for iBGP:** Select the check box to enable BFD for the iBGP neighbor. This option is disabled by default.

**Enable BFD for OSPF:** Select the check box to enable BFD for the OSPF underlay instance. This option is disabled by default, and it is grayed out if the link state protocol is ISIS.

**Enable BFD for ISIS:** Select the check box to enable BFD for the ISIS underlay instance. This option is disabled by default, and it is grayed out if the link state protocol is OSPF.

**Enable BFD for PIM:** Select the check box to enable BFD for PIM. This option is disabled by default, and it is grayed out if the replication mode is Ingress.

Here are the examples of the BFD global policies:

```
router ospf <ospf tag>
  bfd

router isis <isis tag>
  address-family ipv4 unicast
  bfd

ip pim bfd

router bgp <bgp asn>
  neighbor <neighbor ip>
  bfd
```

**Enable BFD Authentication:** Select the check box to enable BFD authentication. If you enable this field, the **BFD Authentication Key ID** and **BFD Authentication Key** fields are editable.

**Note**

- BFD Authentication is not supported when the **Fabric Interface Numbering** field under the **General** tab is set to **unnumbered**. The BFD authentication fields will be grayed out automatically.

**BFD Authentication Key ID:** Specifies the BFD authentication key ID for the interface authentication. The default value is 100.

**BFD Authentication Key:** Specifies the BFD authentication key.

For information about how to retrieve the BFD authentication parameters, see *Retrieving the Authentication Key*.

**iBGP Peer-Template Config** – Add iBGP peer template configurations on the leaf switches and route reflectors to establish an iBGP session between the leaf switch and route reflector. Set this field based on switch configuration. If this field is blank, it implies that the iBGP peer template is not used. If the iBGP peer template is used, enter the peer template definition as defined on the switch. The peer template name on devices configured with BGP should be the same as defined here.

**Note** If you use the iBGP peer template, include the BGP authentication configuration in this template config field. Additionally, uncheck the Enable BGP Authentication check box to avoid duplicating the BGP configuration.

### Step 7

Click the **Advanced** tab. Most of the fields are auto generated.

**VRF Template** and **VRF Extension Template**: Specifies the VRF template for creating VRFs, and the VRF extension template for enabling VRF extension to other fabrics.

**Network Template** and **Network Extension Template**: Specifies the network template for creating networks, and the network extension template for extending networks to other fabrics.

You must not change the templates when migrating. Only the Universal templates are supported for overlay migration.

**Site ID** - The ID for this fabric if you are moving this fabric within an MSD. You can update this field post-migration.

**Intra Fabric Interface MTU** - Specifies the MTU for the intra fabric interface. This value should be an even number.

**Layer 2 Host Interface MTU** - Specifies the MTU for the layer 2 host interface. This value should be an even number.

**Power Supply Mode** - Choose the appropriate power supply mode.

**CoPP Profile** - Choose the Control Plane Policing (CoPP) profile policy used in the existing fabric. By default, the strict option is populated.

**VTEP HoldDown Time** - Specifies the NVE source interface hold down time.

**Enable VXLAN OAM** - Enables the VXLAN OAM function for existing switches.

This is enabled by default. Clear the check box to disable VXLAN OAM function.

If you want to enable the VXLAN OAM function on specific switches and disable on other switches in the fabric, you can use freeform configurations to enable OAM and disable OAM in the fabric settings.

**Note** The VXLAN OAM feature in Cisco DCNM is only supported on a single fabric or site.

Note that the NGOAM feature is unsupported on switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images.

**Enable Tenant DHCP** – Select the check box to enable the tenant DHCP support.

**Note** Ensure that **Enable Tenant DHCP** is enabled before enabling DHCP related parameters in the overlay profiles.

**Enable NX-API** - Specifies enabling of NX-API.

**Enable NX-API on HTTP** - Specifies enabling of NX-API on HTTP.

**Enable Policy-Based Routing (PBR)** - Select this check box to enable routing of packets based on the specified policy. For information on Layer 4-Layer 7 service, refer [Layer 4-Layer 7 Service](#).

**Enable Strict Config Compliance** - Enable the Strict Config Compliance feature by selecting this check box. By default, this feature is disabled. For more information, refer *Strict Configuration Compliance*.

**Note** If Strict Config Compliance is enabled in a fabric, you cannot deploy Network Insights for Resources on Cisco DCNM.

**Enable AAA IP Authorization** - Enables AAA IP authorization, when IP Authorization is enabled in the AAA Server.

**Greenfield Cleanup Option** – Enable or disable the switch cleanup option for Greenfield switches. This is applicable post-migration when new switches are added.

**Enable MPLS Handoff**: Select the check box to enable the MPLS Handoff feature. For more information, see [Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - MPLS SR and LDP Handoff](#).

**Note**: For the brownfield import, you need to select the **Enable MPLS Handoff** feature. Most of the IFC configuration will be captured in **switch\_freeform**.

**Underlay MPLS Loopback Id**: Specifies the underlay MPLS loopback ID. The default value is 101.

**Leaf Freeform Config** and **Spine Freeform Config** - You can enter these fields after fabric transitioning is complete, as needed.

**Intra-fabric Links Additional Config** - You can enter this field after fabric transitioning is complete, as needed.

## Step 8

Click the **Resources** tab.

**Manual Underlay IP Address Allocation** – *Do not* select this check box if you are transitioning your VXLAN fabric management to DCNM.

Review the ranges and ensure they are consistent with the existing fabric. The migration will honor the existing resources as found on the fabric. The range settings apply to post migration allocation.

**Underlay Routing Loopback IP Range** - Specifies loopback IP addresses for the protocol peering.

**Underlay VTEP Loopback IP Range** - Specifies loopback IP addresses for VTEPs.

**Underlay RP Loopback IP Range** - Specifies the anycast or phantom RP IP address range.

**Underlay Subnet IP Range** - IP addresses for underlay P2P routing traffic between interfaces.

**Layer 2 VXLAN VNI Range** and **Layer 3 VXLAN VNI Range** - Specifies the VXLAN VNI IDs for the fabric.

**Network VLAN Range** and **VRF VLAN Range** - VLAN ranges for the Layer 3 VRF and overlay network.

**Subinterface Dot1q Range** - Specifies the subinterface range when L3 sub interfaces are used.

**VRF Lite Deployment** - Specify the VRF Lite method for extending inter fabric connections.

If you select Manual, the VRF Lite subnet details are required so that the resource manager can reserve the address space.

If you select Back2BackOnly, ToExternalOnly, or Both, then the VRF Lite subnet fields are enabled.

**VRF Lite Subnet IP Range** and **VRF Lite Subnet Mask** – These fields are populated with the DCI subnet details. Update the fields as needed.

The values shown in your screen are automatically generated. If you want to update the IP address ranges, VXLAN Layer 2/Layer 3 network ID ranges or the VRF/Network VLAN ranges, ensure the following:

**Note** When you update a range of values, ensure that it does not overlap with other ranges. You should only update one range of values at a time. If you want to update more than one range of values, do it in separate instances. For example, if you want to update L2 and L3 ranges, you should do the following.

- a. Update the L2 range and click **Save**.
- b. Click the **Edit Fabric** option again, update the L3 range and click **Save**.

The remaining tabs do not require updates. However, their purpose is mentioned.

**Step 9** Click the **Manageability** tab.

Enter the DNS, NTP, AAA, or syslog servers' IP address, VRF, and other applicable information matching the switch configuration. If there are more than two servers for these features, add the configurations of the additional servers to the **Leaf Freeform Config** and **Spine Freeform Config** fields in the **Advanced** tab.

**Step 10** Click the **Bootstrap** tab. Update the fields in this tab post transition, when new switches are added to the fabric.

**Step 11** Click the **Configuration Backup** tab. Leave the fields in this tab blank. You can update post transition.

**Step 12** Click **Save** after filling and updating relevant information. A note appears briefly at the bottom right part of the screen, indicating that the fabric is created. When a fabric is created, the fabric page comes up. The fabric name appears at the top left part of the screen.

The **Actions** panel at the left part of the screen allows you to perform various functions. One of them is the **Add switches** option to add switches to the fabric. After you create a fabric, you should add fabric devices. The process is explained next:

---

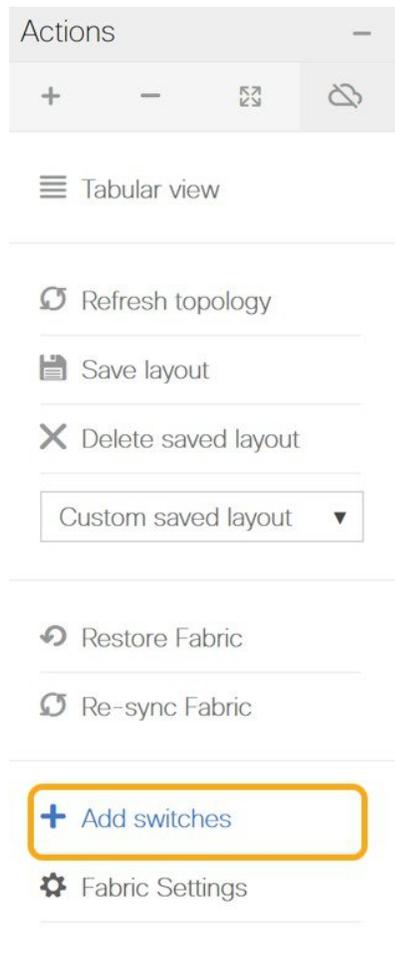
## Adding Switches and Transitioning VXLAN Fabric Management to DCNM

Let us discover and add switches to the newly created fabric.

### Procedure

---

**Step 1** Click **Add Switches** in the **Actions** menu.



**Step 2** Under the **Discover Existing Switches** tab, enter the IP address of the switch in the **Seed IP** field. Enter the username and password of the switches that you want to discover.

### Inventory Management ✕

Discover Existing Switches
PowerOn Auto Provisioning (POAP)

Discovery Information > Scan Details >

Seed IP   
Ex: "2.2.2.20"; "10.10.10.40-60"; "2.2.2.20, 2.2.2.21"

Authentication Protocol

Username

Password

Max Hops    hop(s)

Preserve Config  no  yes  
Selecting 'no' will clean up the configuration on switch(es)

By default, the value in the **Max Hops** field is **2**. The switch with the specified IP address and the switches that are 2 hops from it will be populated after the discovery is complete.

Make sure that the **Preserve Config** toggle button is set to **yes**.

The **yes** setting means that the current configuration of the switches will be retained.

**Important** - Ensure that the Preserve Config field remains set to **yes**. Selecting **no** can cause significant configuration loss and fabric disruption.

The POAP tab is only used for adding new switches to the fabric. Use the tab only after migrating your existing fabric to DCNM.

**Step 3** Click **Start discovery**.

✕

### Inventory Management

Discover Existing Switches
PowerOn Auto Provisioning (POAP)

Discovery Information > Scan Details >

Seed IP   
Ex: \*2.2.2.20\*; \*10.10.10.40-60\*; \*2.2.2.20, 2.2.2.21\*

Authentication Protocol

Username

Password

Max Hops   hop(s)

Preserve Config  no  yes  
Selecting 'no' will clean up the configuration on switch(es)

Start discovery

The switch with the specified IP address and switches up to two hops away (depending on the setting of Max Hops) from it are populated in the Scan Details section.

#### Step 4

Check the check box next to the switches that have to be imported into the fabric and click **Import into fabric**.

It is best practice to discover multiple switches at the same time in a single attempt. The switches must be cabled and connected to the DCNM server and the switch status must be manageable.

If switches are imported in multiple attempts, then all the switches must be added to the fabric before you make any changes to the fabric, that is, before you click **Save & Deploy**.

Inventory Management ✕

Discover Existing Switches | PowerOn Auto Provisioning (POAP)

Discovery Information > Scan Details >

← Back Note: Preserve Config selection is 'yes'. Import into fabric

|                                     | Name  | IP Address  | Model       | Version     | Status     | Progress |
|-------------------------------------|-------|-------------|-------------|-------------|------------|----------|
| <input checked="" type="checkbox"/> | n9k13 | 80.80.80.63 | N9K-C939... | 7.0(3)I7(6) | manageable |          |
| <input checked="" type="checkbox"/> | n9k8  | 80.80.80.58 | N9K-C939... | 7.0(3)I7(6) | manageable |          |
| <input checked="" type="checkbox"/> | n9k12 | 80.80.80.62 | N9K-C939... | 7.0(3)I7(6) | manageable |          |
| <input checked="" type="checkbox"/> | n9k7  | 80.80.80.57 | N9K-C939... | 7.0(3)I7(6) | manageable |          |
| <input checked="" type="checkbox"/> | n9k14 | 80.80.80.64 | N9K-C921... | 7.0(3)I7(6) | manageable |          |

Show All ▾

Close

**Step 5** Click **Import into fabric**.

The switch discovery process is initiated. The **Progress** column displays progress for all the selected switches. It displays **done** for each switch after completion.

**Note** You should not close the screen and try to import switches again until all selected switches are imported or an error message comes up.

If an error message comes up, close the screen. The fabric topology screen comes up. The error messages are displayed at the top-right part of the screen. Resolve the errors and initiate the import process again by clicking **Add Switches** in the **Actions** panel.

**Step 6** After a successful import, the progress bar shows **Done** for all the switches. Click **Close**.

### Inventory Management ✕

Discover Existing Switches
PowerOn Auto Provisioning (POAP)

Discovery Information >
Scan Details >

← Back
Note: Preserve Config selection is 'yes'.
Import into fabric

| <input checked="" type="checkbox"/> | Name  | IP Address  | Model       | Version     | Status     | Progress |
|-------------------------------------|-------|-------------|-------------|-------------|------------|----------|
| <input checked="" type="checkbox"/> | n9k13 | 80.80.80.63 | N9K-C939... | 7.0(3)I7(6) | manageable | done     |
| <input checked="" type="checkbox"/> | n9k8  | 80.80.80.58 | N9K-C939... | 7.0(3)I7(6) | manageable | done     |
| <input checked="" type="checkbox"/> | n9k12 | 80.80.80.62 | N9K-C939... | 7.0(3)I7(6) | manageable | done     |
| <input checked="" type="checkbox"/> | n9k7  | 80.80.80.57 | N9K-C939... | 7.0(3)I7(6) | manageable | done     |
| <input checked="" type="checkbox"/> | n9k14 | 80.80.80.64 | N9K-C921... | 7.0(3)I7(6) | manageable | done     |

Show All ▼

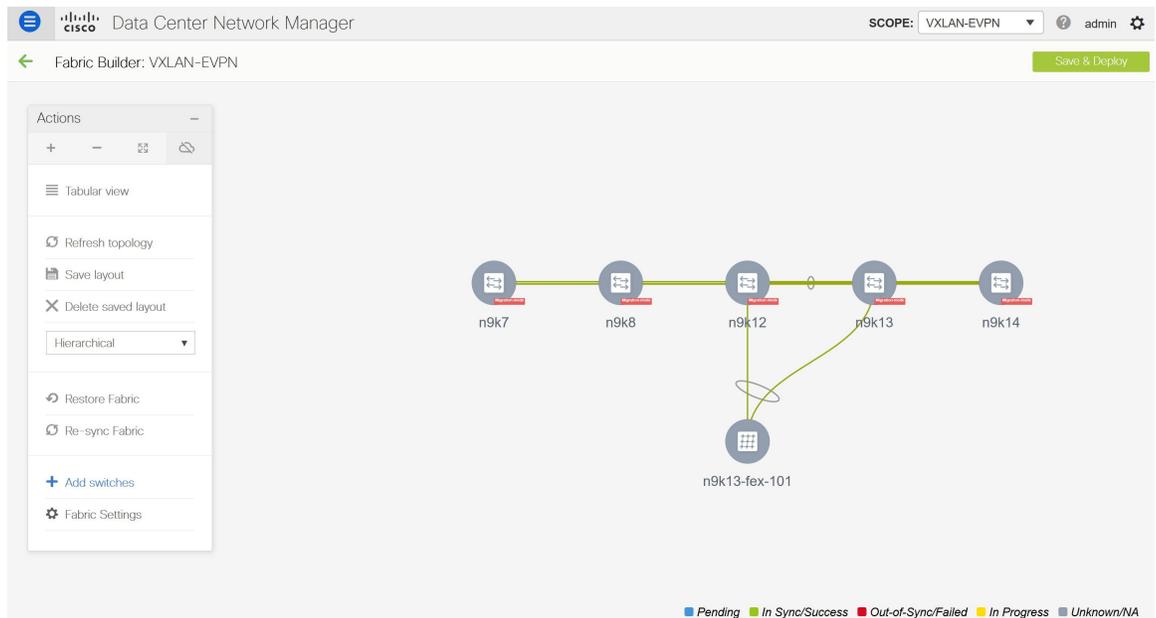
Close

After closing the window, the fabric topology window comes up again. The switch is in Migration Mode now, and the Migration mode label is displayed on the switch icons.

At this point, you must not try to add Greenfield or *new* switches. Support is not available for adding new switches during the migration process. It might lead to undesirable consequences for your network. However, you can add a new switch after the migration process is complete.

**Step 7**

After all the network elements are discovered, they are displayed in the **Fabric Builder** window in a connected topology. Each switch is assigned the **Leaf** role by default.



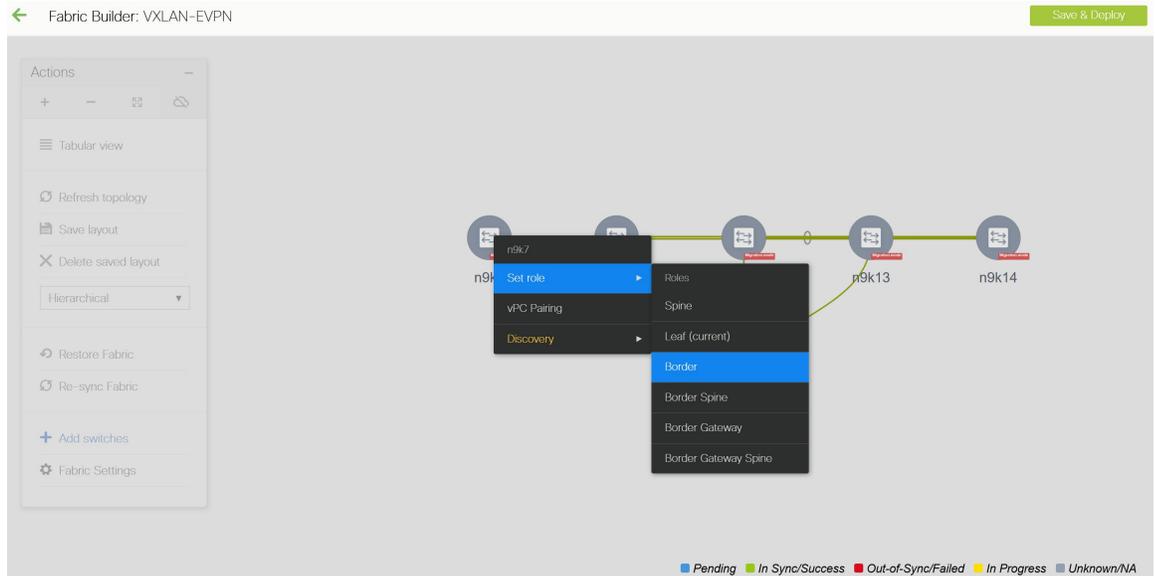
The switch discovery process might fail for a few switches, and the Discovery Error message is displayed. However, such switches are still displayed in the fabric topology. You should remove such switches from the fabric (Right-click the switch icon and click **Discovery > Remove** from fabric), and import them again.

You should not proceed to the next step until all switches in the existing fabric are discovered in DCNM.

If you choose the Hierarchical layout for display (in the Actions panel), the topology automatically gets aligned as per role assignment, with the leaf switches at the bottom, the spine switches connected on top of them, and the border switches at the top.

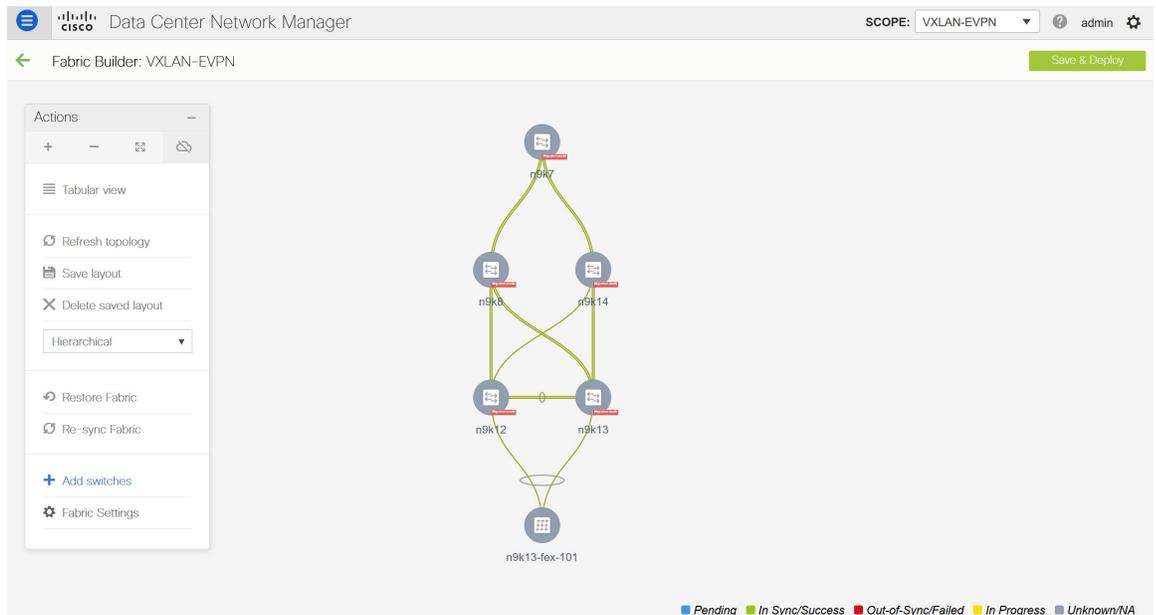
**Note** The supported roles for switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images are Border Leaf, Border Spine, Leaf, and Spine

**Step 8** Right-click the **n9k-7** switch, select **Set Role**, and choose **Border** from the **Roles** drop-down list.



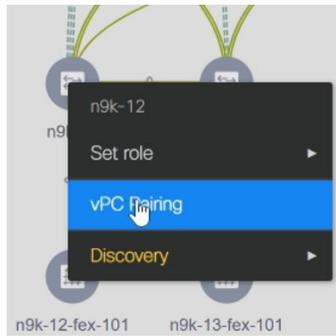
Similarly, set the **Spine** role for the **n9k-14** and **n9k-8** spine switches.

**Note** You need to manually create a vPC pairing when the L3 keep alive is configured on the switches. Otherwise, the vPC configuration is automatically picked up from the switches. For more information, see [Adding a vPC L3 Peer Keep-Alive Link](#), on page 63.



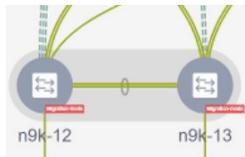
**vPC Pairing** - The vPC pairing must be done for switches where the Layer 3 vpc peer-keep alive is used. The vPC configuration is automatically picked up from the switches when the vpc peer keep alive is established through the management option. This pairing reflects in the GUI only after the migration is complete.

- a. Right-click the switch icon and click vPC Pairing to set a vPC switch pair.



The Select vPC peer screen comes up. It lists potential vPC peer switches.

- b. Select the appropriate switch and click OK. The fabric topology comes up again. The vPC pair is formed now.



**Note** Check if you have added all switches from the current fabric. If you have missed adding switches, add them now. Once you are certain that you have imported all existing switches, move to the next step, the Save and Deploy option.

### Step 9 Click **Save & Deploy**.

When you click **Save & Deploy**, DCNM obtains switch configurations and populates the state of every switch from the current running config to the current expected config, which is the intended state maintained in DCNM.

The Saving Fabric Configuration message comes up immediately. This indicates that overlay and underlay network migration, and switch and port channel settings migration to DCNM is initiated.

If there are configuration mismatches, error messages are displayed. Update changes in the fabric settings or the switch configuration as needed, and click Save and Deploy again.

After the migration of underlay and overlay networks, the Configuration Deployment screen comes up.

- Note**
- The brownfield migration requires best practices to be followed on the existing fabric such as maintain consistency of the overlay configurations. For more information, see [Guidelines and Limitations](#).
  - Any errors or inconsistencies that are found during the migration is reported in fabric errors. The switches continue to remain in the Migration mode. You should fix these errors and complete the migration again by clicking **Save & Deploy** until no errors are reported.

### Step 10 After the configurations are generated, review them by clicking the links in the **Preview Config** column.

## Config Deployment



Step 1. Configuration Preview &gt; Step 2. Configuration Deployment Status &gt;

| Switch Name | IP Address  | Switch Serial | Preview Config             | Status      | Re-sync | Progress  |
|-------------|-------------|---------------|----------------------------|-------------|---------|---|
| n9k12       | 80.80.80.62 | SAL18422FX8   | <a href="#">2405 lines</a> | Out-of-sync |         | <div style="width: 100%;"><div style="width: 100%;"></div></div> 100% |
| n9k13       | 80.80.80.63 | SAL18422FXE   | <a href="#">2405 lines</a> | Out-of-sync |         | <div style="width: 100%;"><div style="width: 100%;"></div></div> 100% |
| n9k7        | 80.80.80.57 | SAL1833YM64   | <a href="#">2200 lines</a> | Out-of-sync |         | <div style="width: 100%;"><div style="width: 100%;"></div></div> 100% |
| n9k14       | 80.80.80.64 | SAL2016NXXB   | <a href="#">2 lines</a>    | Out-of-sync |         | <div style="width: 100%;"><div style="width: 100%;"></div></div> 100% |
| n9k8        | 80.80.80.58 | SAL1833YM0V   | <a href="#">3 lines</a>    | Out-of-sync |         | <div style="width: 100%;"><div style="width: 100%;"></div></div> 100% |

Deploy Config

We strongly recommend that you preview the configuration before proceeding to deploy it on the switches. Click the Preview Config column entry. The Config Preview screen comes up. It lists the pending configurations on the switch.

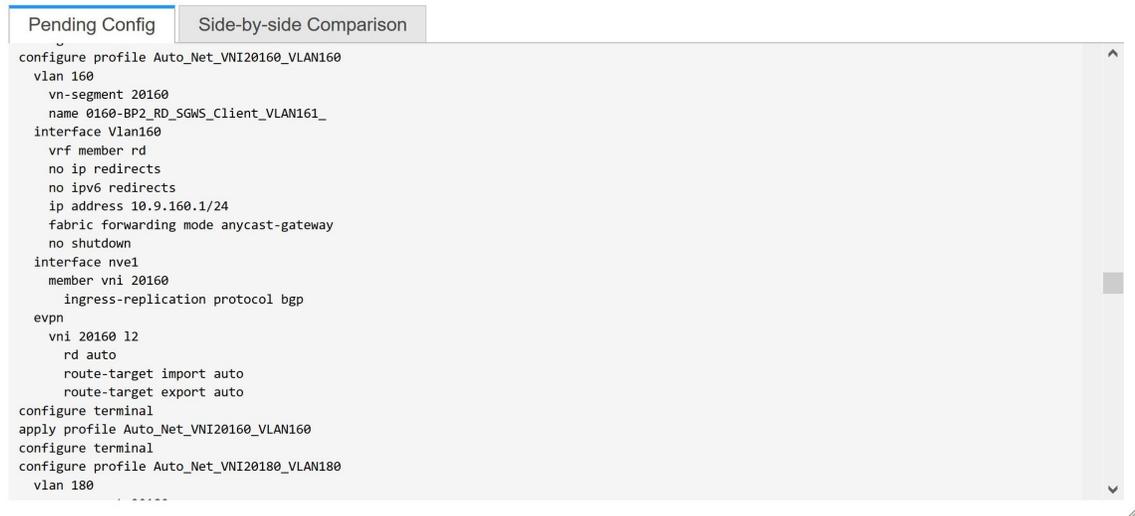
The Side-by-side Comparison tab displays the running configuration and expected configuration side-by-side.

The **Pending Config** tab displays the set of configurations that need to be deployed on a switch in order to go from the current running configuration to the current expected or intended configuration.

The **Pending Config** tab may show many config lines that will be deployed to the switches. Typically, on a successful brownfield import, these lines correspond to the configuration profiles pushed to the switches for an overlay network configuration. Note that the existing network and VRF-related overlay configurations are not removed from the switches.

The configuration profiles are DCNM required constructs for managing the VXLAN configurations on the switches. During the Brownfield import process, they capture the same information as the original VXLAN configurations already present on the switches. In the following image, the configuration profile with **vlan 160** is applied.

## Config Preview - Switch 80.80.80.62



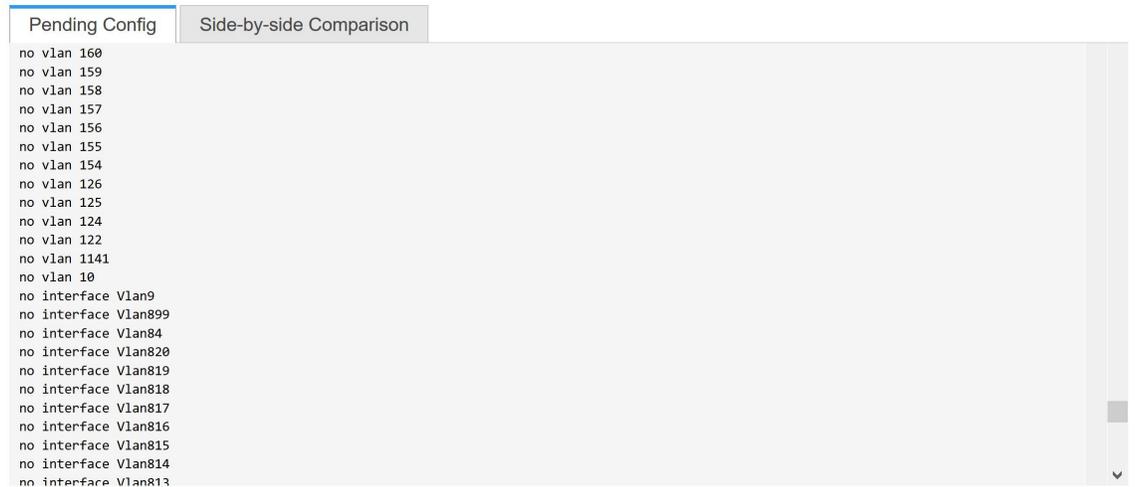
```

Pending Config | Side-by-side Comparison
-----
configure profile Auto_Net_VNI20160_VLAN160
vlan 160
  vn-segment 20160
  name 0160-BP2_RD_SGWS_Client_VLAN161_
interface Vlan160
  vrf member rd
  no ip redirects
  no ipv6 redirects
  ip address 10.9.160.1/24
  fabric forwarding mode anycast-gateway
  no shutdown
interface nve1
  member vni 20160
  ingress-replication protocol bgp
evpn
  vni 20160 12
  rd auto
  route-target import auto
  route-target export auto
configure terminal
apply profile Auto_Net_VNI20160_VLAN160
configure terminal
configure profile Auto_Net_VNI20180_VLAN180
vlan 180

```

As part of the import process, after the configuration profiles are applied, the original CLI based configuration references will be removed from the switches. These are the ‘no’ CLIs that will be seen towards the end of the diffs. The VXLAN configurations on the switches will be persisted in the configuration profiles. In the following image, you can see that the configurations will be removed, specifically, **no vlan 160**.

## Config Preview - Switch 80.80.80.62



```

Pending Config | Side-by-side Comparison
-----
no vlan 160
no vlan 159
no vlan 158
no vlan 157
no vlan 156
no vlan 155
no vlan 154
no vlan 126
no vlan 125
no vlan 124
no vlan 122
no vlan 1141
no vlan 10
no interface Vlan9
no interface Vlan899
no interface Vlan84
no interface Vlan820
no interface Vlan819
no interface Vlan818
no interface Vlan817
no interface Vlan816
no interface Vlan815
no interface Vlan814
no interface Vlan813

```

The **Side-by-side Comparison** tab displays the Running Config and Expected Config on the switch.

- Step 11** Close the **Config Preview Switch** window after reviewing the configurations.
- Step 12** Click **Deploy Config** to deploy the pending configuration onto the switches.

## Config Deployment



Step 1. Configuration Preview > Step 2. Configuration Deployment Status >

| Switch Name | IP Address  | Status    | Status Description    | Progress |
|-------------|-------------|-----------|-----------------------|----------|
| n9k14       | 80.80.80.64 | COMPLETED | Deployed successfully | 100%     |
| n9k8        | 80.80.80.58 | COMPLETED | Deployed successfully | 100%     |
| n9k12       | 80.80.80.62 | COMPLETED | Deployed successfully | 100%     |
| n9k7        | 80.80.80.57 | COMPLETED | Deployed successfully | 100%     |
| n9k13       | 80.80.80.63 | COMPLETED | Deployed successfully | 100%     |

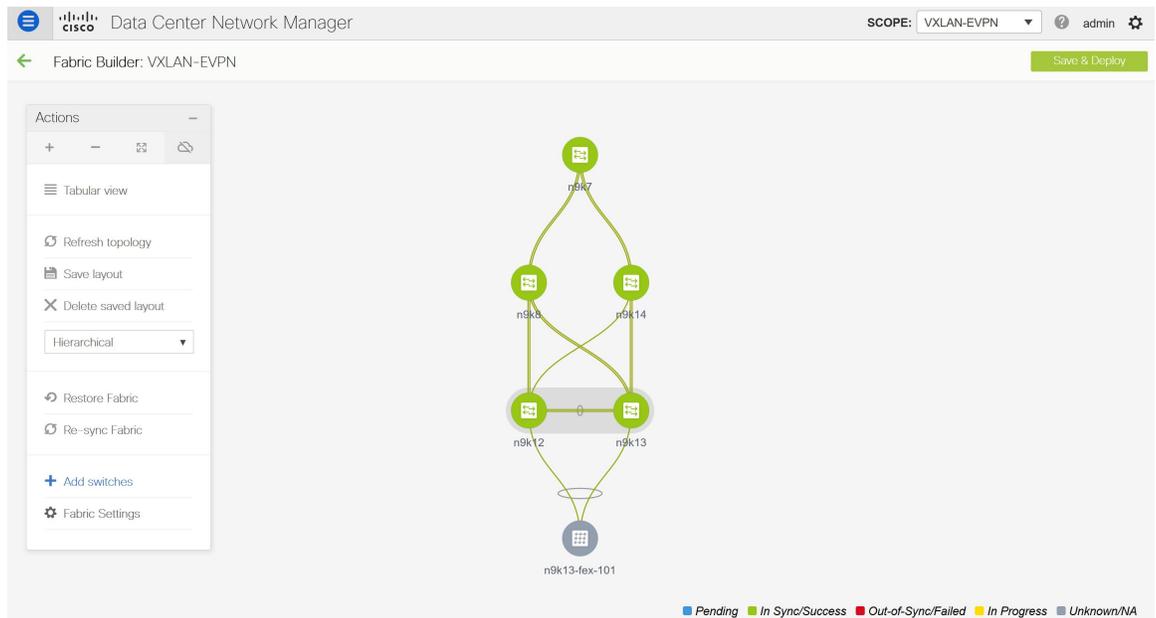
Close

If the **Status** column displays **FAILED**, investigate the reason for failure to address the issue.

The progress bar shows **100%** for each switch. After correct provisioning and successful configuration compliance, close the screen.

In the fabric topology screen that comes up, all imported switch instances are displayed in green color, indicating successful configuration. Also, the **Migration Mode** label is not displayed on any switch icon.

DCNM has successfully imported a VXLAN-EVPN fabric.



**Post-transitioning of VXLAN fabric management to DCNM** - This completes the transitioning process of VXLAN fabric management to DCNM. Now, you can add new switches and provision overlay networks for your fabric. For details, refer the respective section in the Fabrics topic in the configuration guide.

### Fabric Options

- **Tabular View** - By default, the switches are displayed in the topology view. Use this option to view switches in the tabular view.
- **Refresh topology** - Allows you to refresh the topology.
- **Save Layout** – Saves a custom view of the topology. You can create a specific view in the topology and save it for ease of use.
- **Delete saved layout** – Deletes the custom view of the topology
- **Topology views** - You can choose between Hierarchical, Random and Custom saved layout display options.
  - **Hierarchical** - Provides an architectural view of your topology. Various Switch Roles can be defined that draws the nodes on how you configure your CLOS topology.
  - **Random** - Nodes are placed randomly on the screen. DCNM tries to make a guess and intelligently place nodes that belong together in close proximity.
  - **Custom saved layout** - You can drag nodes around to your liking. Once you have the positions as how you like, you can click Save Layout to remember the positions. Next time you come to the topology, DCNM will draw the nodes based on your last saved layout positions.
- **Restore Fabric** – Allows you to restore the fabric to a prior DCNM configuration state (one month back, two months back, and so on). For more information, see the *Restore Fabric* section.
- **Resync Fabric** - Use this option to resynchronize DCNM state when there is a large scale out-of-band change, or if configuration changes do not register in the DCNM properly. The resync operation does a full CC run for the fabric switches and recollects “show run” and “show run all” commands from the

switches. When you initiate the re-sync process, a progress message is displayed on the screen. During the re-sync, the running configuration is taken from the switches. The Out-of-Sync/In-Sync status for the switches is recalculated based on the intent defined in DCNM.

- **Add Switches** – Allows you to add switch instances to the fabric.
- **Fabric Settings** – Allows you to view or edit fabric settings.

## Verifying the Import of the VXLAN BGP EVPN Fabric

Let us verify whether the Brownfield migration was successful.

### Verifying VXLANs and Commands on Switches

#### Procedure

- Step 1** To verify the VXLANs in this fabric, double click a switch and click **Show more details** in the switch pane.

The screenshot displays a network diagram on the left and a detailed switch configuration pane on the right. The network diagram shows five switches: n9k7 at the top, n9k8 and n9k14 in the middle, n9k12 and n9k13 at the bottom, and n9k13-fex-101 at the very bottom. A legend at the bottom of the diagram indicates that green circles represent 'In Sync/Success' and blue circles represent 'Pending'. The switch details pane on the right is titled 'Summary' and shows the following information:

|                |             |
|----------------|-------------|
| Status:        | ✓ ok        |
| Serial number: | SAL18422FX8 |
| CPU:           | 22%         |
| Memory:        | 30%         |

Below the summary, the VPC Domain ID is 2. The Role is Secondary, Peer is n9k13, Peerlink State is Peer is OK, Keep Alive State is Peer is alive, Consistency State is Consistent, Send Interface is mgmt0, and Receive Interface is mgmt0. There is a 'Tags' section with a plus sign and a 'System Tags' section with a 'VTEP' tag. At the bottom of the pane, there is a green button with a left arrow and the text 'Show more details'.

- Step 2** Click the **VXLAN** tab.

n9k12  
80.80.80.62  
N9K-C9396PX

System Info Modules FEX License Features **VXLAN** Port Capacity

VXLAN Total 84

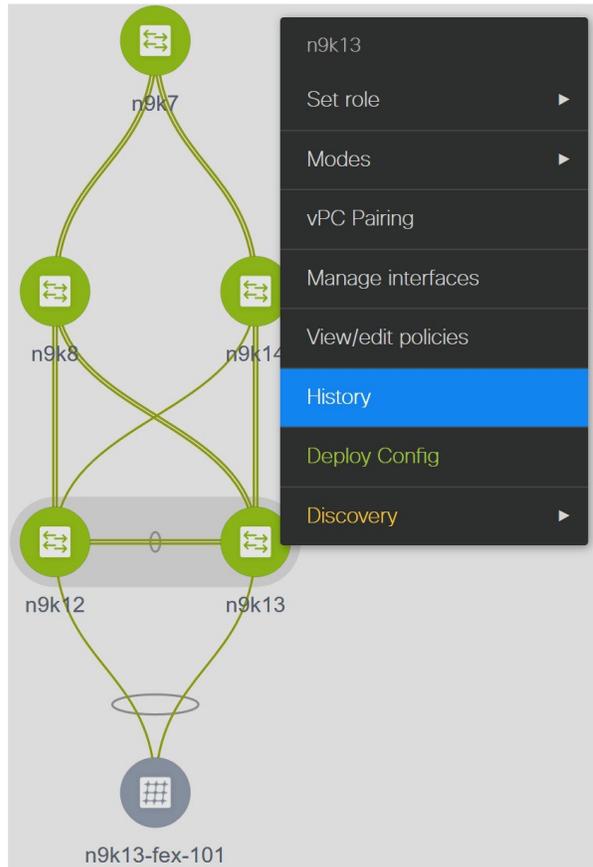
Show Quick Filter

| NVE Interface | VNI   | Multicast Address | VNI Status | Mode          | Type    | VRF | Mapped VLAN |
|---------------|-------|-------------------|------------|---------------|---------|-----|-------------|
| nve1          | 20006 | UnicastBGP        | Up         | Control-Plane | Layer-2 | -   | 6           |
| nve1          | 20009 | UnicastBGP        | Up         | Control-Plane | Layer-2 | -   | 9           |
| nve1          | 20010 | UnicastBGP        | Up         | Control-Plane | Layer-2 | -   | 10          |
| nve1          | 20017 | UnicastBGP        | Up         | Control-Plane | Layer-2 | -   | 17          |
| nve1          | 20018 | UnicastBGP        | Up         | Control-Plane | Layer-2 | -   | 18          |
| nve1          | 20027 | UnicastBGP        | Up         | Control-Plane | Layer-2 | -   | 27          |
| nve1          | 20028 | UnicastBGP        | Up         | Control-Plane | Layer-2 | -   | 28          |
| nve1          | 20029 | UnicastBGP        | Up         | Control-Plane | Layer-2 | -   | 29          |
| nve1          | 20030 | UnicastBGP        | Up         | Control-Plane | Layer-2 | -   | 30          |
| nve1          | 20031 | UnicastBGP        | Up         | Control-Plane | Layer-2 | -   | 31          |
| nve1          | 20036 | UnicastBGP        | Up         | Control-Plane | Layer-2 | -   | 36          |
| nve1          | 20040 | UnicastBGP        | Up         | Control-Plane | Layer-2 | -   | 40          |

You can see that all the VXLANs have been migrated successfully.

**Note** You can verify remaining information by clicking the different tabs in this window.

**Step 3** Right-click a switch and select **History** to see the commands pushed by DCNM.



**Step 4** Click the **Success** hyperlink under the **Status** column to view the commands pushed by DCNM.

## Policy Deployment History for n9k13 ( SAL18422FXE )

| Entity Name | Entity Type | Source   | Status  | Status Description    | User  | Time of Completion      |
|-------------|-------------|----------|---------|-----------------------|-------|-------------------------|
| SAL18422FXE | SWITCH      | DCNM     | SUCCESS | Successfully deployed | admin | 2019-08-08 22:47:13.353 |
| SWITCH      | SWITCH      | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-08-08 22:36:32.101 |
| SWITCH      | SWITCH      | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-08-08 22:36:14.783 |
| SWITCH      | SWITCH      | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-08-08 22:36:07.129 |
| SWITCH      | SWITCH      | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-08-08 22:36:06.122 |
| SWITCH      | SWITCH      | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-08-08 22:36:05.116 |
| SWITCH      | SWITCH      | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-08-08 22:36:04.109 |
| SWITCH      | SWITCH      | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-08-08 22:36:03.102 |
| SWITCH      | SWITCH      | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-08-08 22:36:02.095 |
| SWITCH      | SWITCH      | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-08-08 22:36:01.089 |
| SWITCH      | SWITCH      | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-08-08 22:36:00.081 |
| SWITCH      | SWITCH      | UNDERLAY | SUCCESS | Successfully deployed | admin | 2019-08-08 22:35:59.275 |

## Verifying Resources

DCNM has a resource manager that tracks all the resources used in a fabric. Navigate to **Control > Management > Resources** in the left menu.

SCOPE: VXLAN-EVPN admin

Control / Management / Resources

Resource Allocation Selected 0 / Total 429

| Scope Type | Scope       | Device Name | Device IP   | Allocated Resource | Allocated To            | Resource Type   | Is Allocated? | Allocated On   |
|------------|-------------|-------------|-------------|--------------------|-------------------------|-----------------|---------------|----------------|
| Device     | SAL18422FX8 | n9k12       | 80.80.80.62 | 80                 | Auto_Net_VNI20080_VL... | TOP_DOWN_NE...  | Yes           | 09/08/2019,... |
| Device     | SAL18422FX8 | n9k12       | 80.80.80.62 | 500                | loopback500             | LOOPBACK_ID     | Yes           | 09/08/2019,... |
| Device     | SAL18422FX8 | n9k12       | 80.80.80.62 | 501                | loopback501             | LOOPBACK_ID     | Yes           | 09/08/2019,... |
| Device     | SAL1833YM64 | n9k7        | 80.80.80.57 | 101                | port-channel101         | PORT_CHANNEL... | Yes           | 09/08/2019,... |
| Device     | SAL1833YM64 | n9k7        | 80.80.80.57 | 3957               | ECD                     | TOP_DOWN_VR...  | Yes           | 09/08/2019,... |
| Device     | SAL1833YM64 | n9k7        | 80.80.80.57 | 3959               | LC-DMZ                  | TOP_DOWN_VR...  | Yes           | 09/08/2019,... |
| Device     | SAL1833YM64 | n9k7        | 80.80.80.57 | 3958               | RD                      | TOP_DOWN_VR...  | Yes           | 09/08/2019,... |
| Device     | SAL1833YM64 | n9k7        | 80.80.80.57 | 3965               | COMMON-MGMT             | TOP_DOWN_VR...  | Yes           | 09/08/2019,... |
| Device     | SAL1833YM64 | n9k7        | 80.80.80.57 | 3961               | DCI                     | TOP_DOWN_VR...  | Yes           | 09/08/2019,... |
| Device     | SAL1833YM64 | n9k7        | 80.80.80.57 | 58                 | Auto_Net_VNI20058_VL... | TOP_DOWN_NE...  | Yes           | 09/08/2019,... |
| Device     | SAL1833YM64 | n9k7        | 80.80.80.57 | 57                 | Auto_Net_VNI20057_VL... | TOP_DOWN_NE...  | Yes           | 09/08/2019,... |
| Device     | SAL1833YM64 | n9k7        | 80.80.80.57 | 3964               | COMMON-DMZ              | TOP_DOWN_VR...  | Yes           | 09/08/2019,... |
| Device     | SAL1833YM64 | n9k7        | 80.80.80.57 | 3963               | LC                      | TOP_DOWN_VR...  | Yes           | 09/08/2019,... |
| Device     | SAL1833YM64 | n9k7        | 80.80.80.57 | 3967               | switchpool-default      | TOP_DOWN_VR...  | Yes           | 09/08/2019,... |
| Device     | SAL1833YM64 | n9k7        | 80.80.80.57 | 3960               | IALAB                   | TOP_DOWN_VR...  | Yes           | 09/08/2019,... |
| Device     | SAL1833YM64 | n9k7        | 80.80.80.57 | 3962               | Internet                | TOP_DOWN_VR...  | Yes           | 09/08/2019,... |

The resources that are being utilized by the VXLAN EVPN fabric such as VLAN IDs, port channel IDs, point to point IP addresses, and loopback IDs are displayed in this window.

## Verifying Networks

### Procedure

**Step 1** From the menu, choose **Control > Fabrics > Networks**.

**Step 2** Choose **VXLAN-EVPN** from the **Scope** drop-down list.

All the networks that are displayed in this window were learned and populated by DCNM as part of the brownfield migration.

**Step 3** From the **Show** drop-down list, choose **Quick Filter** and enter **349** in the VLAN ID field.

Network / VRF Selection > Network / VRF Deployment

Fabric Selected: VXLAN-EVPN

| Network Name              | Network ID | VRF Name | IPv4 Gateway/Subnet | IPv6 Gateway/Prefix | Status   | VLAN ID |
|---------------------------|------------|----------|---------------------|---------------------|----------|---------|
| Auto_Net_VNI20349_VLAN... | 20349      | Internet | 204.90.140.134/29   |                     | DEPLOYED | 349     |

This network is associated with the VLAN ID 349 and is configured with the anycast IP 204.90.140.134.

You can see that this network has been deployed.

Select this network and click **Continue**.

**Step 4** Click **Detailed View**.

This network has been deployed on the leaf switches and the border switch.

Note that **Ethernet 1/5** is one of the ports on the leaf switch.

| Name                      | Network ID | VLAN ID | Switch | Ports   | Status   | Role   |
|---------------------------|------------|---------|--------|---|----------|--------|
| Auto_Net_VNI20349_VLAN... | 20349      | 349     | n9k12  | Ethernet1/5, Port-channel500, Port-channel502 | DEPLOYED | leaf   |
| Auto_Net_VNI20349_VLAN... | 20349      | 349     | n9k13  | Port-channel503, Port-channel505              | DEPLOYED | leaf   |
| Auto_Net_VNI20349_VLAN... | 20349      | 349     | n9k7   |   | DEPLOYED | border |

Let us verify the overlay network associated with this interface.

**Step 5** From the menu, click **Control > Fabrics > Interfaces**.

All the imported interfaces, including port channels, vPC, and mgmt0 interfaces are displayed in the **Interfaces** window.

**Step 6** In the name field, enter **Ethernet 1/5**.

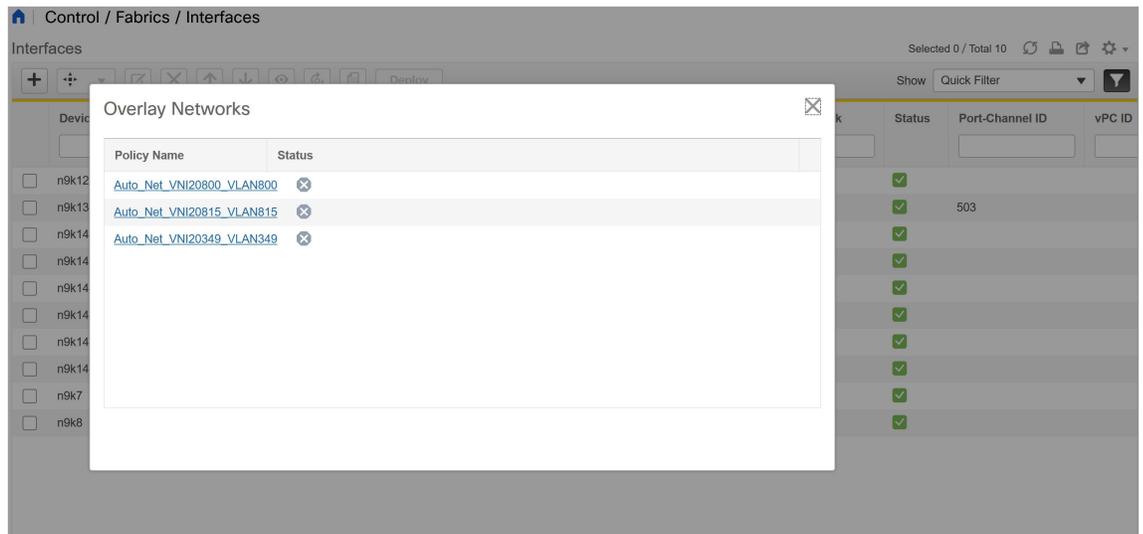
Control / Fabrics / Interfaces

Interfaces

| Device Name | Name         | Admin | Oper | Reason             | Policy                | Overlay Network | Status | Port-Channel ID | vPC ID |
|-------------|--------------|-------|------|--------------------|-----------------------|-----------------|--------|-----------------|--------|
| n9k12       | Ethernet1/5  | ↑     | ↑    | ok                 | int_trunk_host_11_1   | Networks        | ✓      |                 |        |
| n9k13       | Ethernet1/5  | ↑     | ↓    | XCVR not inserted  | int_vpc_trunk_po_memt | NA              | ✓      | 503             |        |
| n9k14       | Ethernet1/5  | ↑     | ↓    | XCVR not inserted  | int_trunk_host_11_1   | NA              | ✓      |                 |        |
| n9k14       | Ethernet1/50 | ↑     | ↓    | Link not connected | int_trunk_host_11_1   | NA              | ✓      |                 |        |
| n9k14       | Ethernet1/51 | ↑     | ↓    | XCVR not inserted  | int_trunk_host_11_1   | NA              | ✓      |                 |        |
| n9k14       | Ethernet1/52 | ↑     | ↓    | XCVR not inserted  | int_trunk_host_11_1   | NA              | ✓      |                 |        |
| n9k14       | Ethernet1/53 | ↑     | ↓    | XCVR not inserted  | int_trunk_host_11_1   | NA              | ✓      |                 |        |
| n9k14       | Ethernet1/54 | ↑     | ↓    | XCVR not inserted  | int_trunk_host_11_1   | NA              | ✓      |                 |        |
| n9k7        | Ethernet1/5  | ↑     | ↓    | XCVR not inserted  | int_trunk_host_11_1   | Networks        | ✓      |                 |        |
| n9k8        | Ethernet1/5  | ↑     | ↓    | XCVR not inserted  | int_trunk_host_11_1   | NA              | ✓      |                 |        |

This interface is attached to the host through the **n9k-12** switch.

**Step 7** In the **Overlay Networks** column, click **Networks** that corresponds to the n9k-12 switch and the Ethernet 1/5 interface.

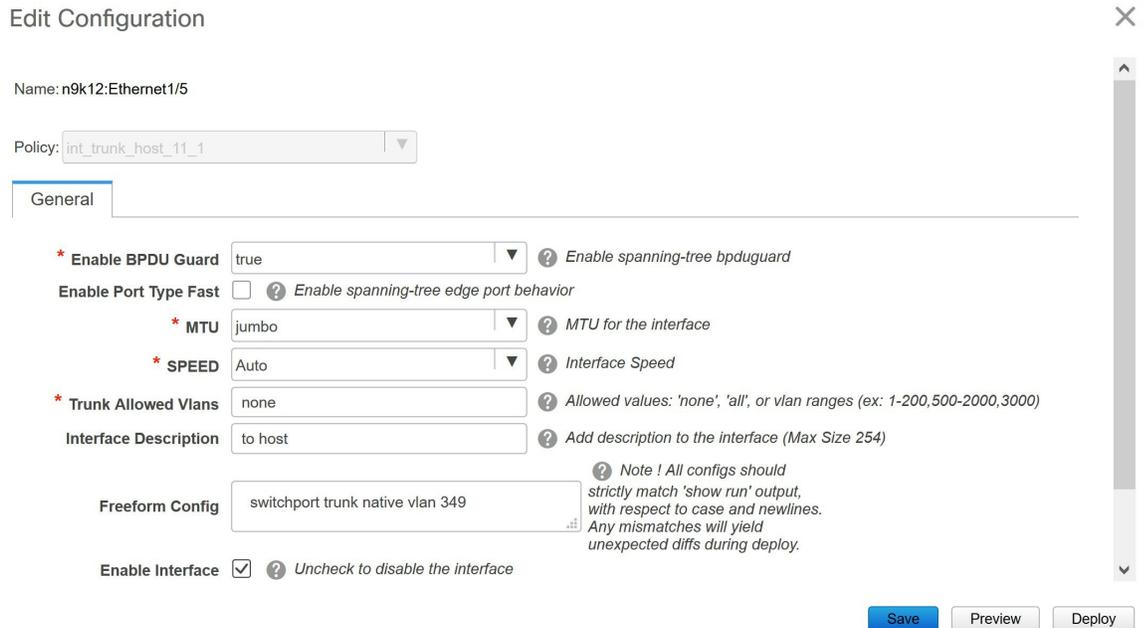


These are the networks that are attached to the **Ethernet 1/5** interface.

**VLAN 349** is also one among them.

You can click this network to see the expected config.

**Step 8** Select the **n9k-12** switch corresponding to the **Ethernet1/5** interface and click the **Edit** icon.



You can see that all the settings for this interface have been successfully imported, including the BPDU guard settings and the interface description.

Let us go back to the host.

The ping command is still running.

**Step 9** End the **ping** command.

```
64 bytes from 204.90.140.134: icmp_seq=4100 ttl=254 time=1.188 ms
64 bytes from 204.90.140.134: icmp_seq=4101 ttl=254 time=1.122 ms
64 bytes from 204.90.140.134: icmp_seq=4102 ttl=254 time=1.224 ms
64 bytes from 204.90.140.134: icmp_seq=4103 ttl=254 time=1.09 ms
64 bytes from 204.90.140.134: icmp_seq=4104 ttl=254 time=1.054 ms
64 bytes from 204.90.140.134: icmp_seq=4105 ttl=254 time=1.079 ms
64 bytes from 204.90.140.134: icmp_seq=4106 ttl=254 time=1.172 ms
64 bytes from 204.90.140.134: icmp_seq=4107 ttl=254 time=1.226 ms
--- 204.90.140.134 ping statistics ---
4108 packets transmitted, 4108 packets received, 0.00% packet loss
round-trip min/avg/max = 1.003/1.264/3.412 ms
```

You can see that 4108 packets are transmitted and received during the migration, and there was zero percent packet loss.

The Brownfield fabric is successfully migrated in to DCNM.

---

## Migrating a Bottom-Up VXLAN Fabric to DCNM

This procedure shows how to migrate a bottom-up VXLAN fabric to DCNM.

Typically, your fabric is created and managed through manual CLI configuration or custom automation scripts. After the migration, the fabric underlay and overlay networks can be managed by using DCNM.

The guidelines and limitations, and prerequisites for bottom-up VXLAN migration are the same as the Brownfield migration. For more information, see *Brownfield Deployment-Transitioning VXLAN Fabric Management to DCNM*.

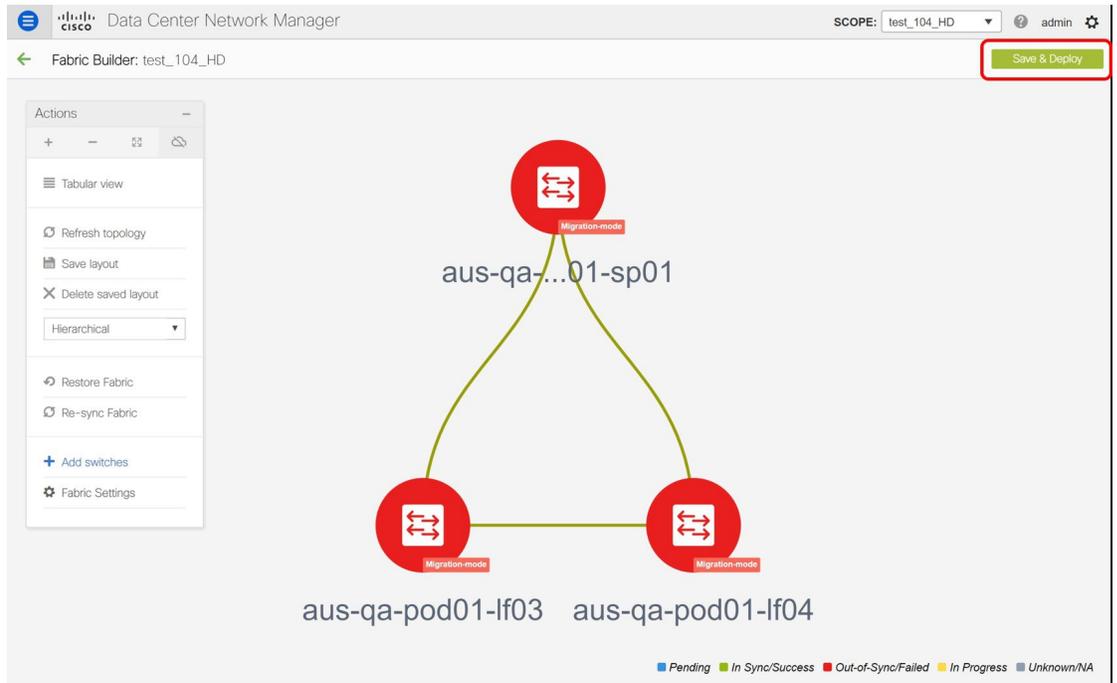
1. Create a VXLAN BGP EVPN fabric.

For more information, see the *Creating a New VXLAN BGP EVPN Fabric* section in *Brownfield Deployment-Transitioning VXLAN Fabric Management to DCNM*.

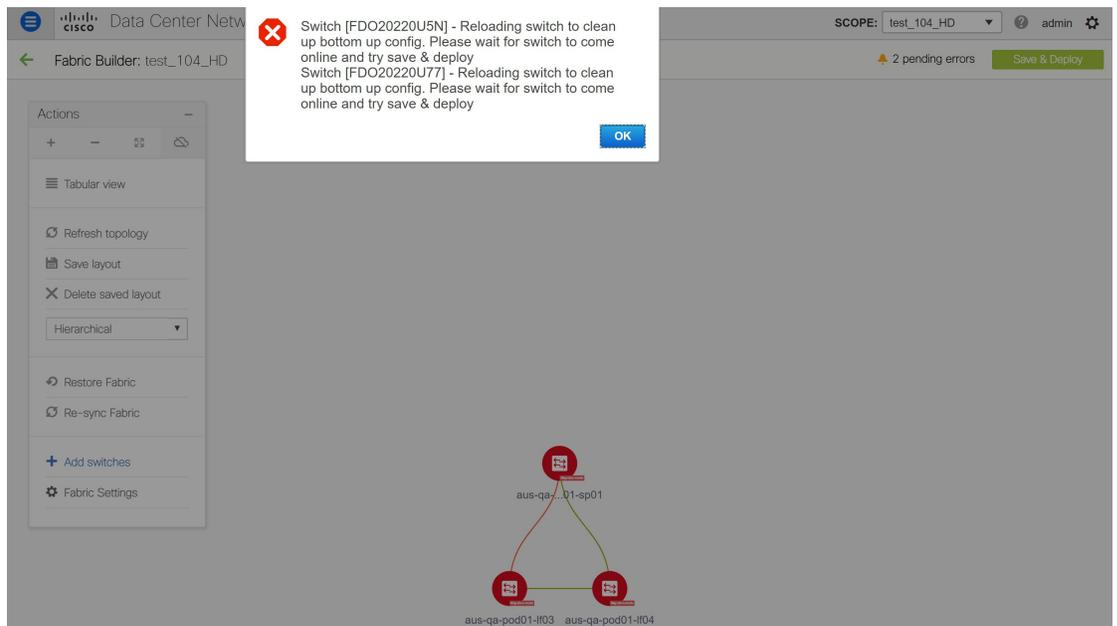
2. Add switch instances to the fabric.

For more information, follow the Step 1 to Step 5 in the *Adding Switch Instances and Transitioning VXLAN Fabric Management* section in *Brownfield Deployment-Transitioning VXLAN Fabric Management to DCNM*.

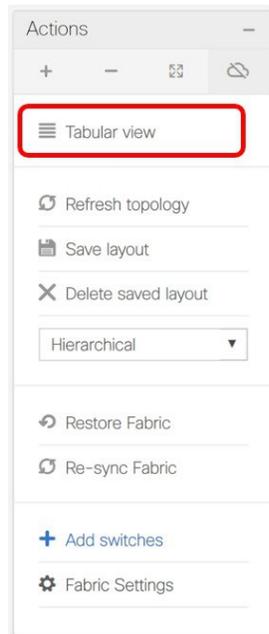
3. Click **Save & Deploy** to sync configurations between the switches and DCNM.



If the added switches contain bottom-up configurations, an error is displayed saying – Reloading switch to clean up bottom up config. Please wait for switch to come online and try **Save & Deploy**.



4. Wait for the switches to complete the reload operation. Click **Tabular view** under the **Actions** menu to view the status of the switches.



- (Optional) Rediscovery of the reloaded switches occurs every 5 minutes. If you want to manually rediscover switches, select the switches and click the **Rediscover switch** icon.

|   | <input checked="" type="checkbox"/> | Name              | IP Address  | Role  | Serial Number | Fabric Name | Fabric Status | Discovery Status | Model          | Software Vers |
|---|-------------------------------------|-------------------|-------------|-------|---------------|-------------|---------------|------------------|----------------|---------------|
| 1 | <input checked="" type="checkbox"/> | aus-qa-pod01-1f03 | 80.80.80.68 | leaf  | FDO20220U5N   | test_104_HD | Out-of-sync   | Discovery timec  | N9K-C9236C     | 7.0(3)17(6)   |
| 2 | <input checked="" type="checkbox"/> | aus-qa-pod01-1f04 | 80.80.80.69 | leaf  | FDO20220U77   | test_104_HD | Out-of-sync   | ok               | N9K-C9236C     | 7.0(3)17(6)   |
| 3 | <input checked="" type="checkbox"/> | aus-qa-pod01-s... | 80.80.80.65 | spine | SAL2016NXX2   | test_104_HD | Out-of-sync   | ok               | N9K-C92160YC-X | 7.0(3)17(6)   |



**Note** Click the **Refresh** icon to refresh the **Fabric Builder** window and see the updated discovery status of switches.

- Check the **Discovery Status** of the switches after the reloading and rediscovering operations are completed. Make sure that the status for all the switches is **ok**.

The screenshot shows the Cisco Data Center Network Manager interface. At the top, it says 'Data Center Network Manager' and 'Fabric Builder: test\_104\_HD'. Below that, there are tabs for 'Switches' and 'Links'. A toolbar contains icons for adding, refreshing, editing, deleting, and buttons for 'View/Edit Policies', 'Manage Interfaces', 'History', and 'Deploy'. Below the toolbar is a table with the following data:

|   | <input type="checkbox"/> | Name              | IP Address  | Role  | Serial Number | Fabric Name | Fabric Status | Disco                                  |
|---|--------------------------|-------------------|-------------|-------|---------------|-------------|---------------|--|
| 1 | <input type="checkbox"/> | aus-qa-pod01-lf03 | 80.80.80.68 | leaf  | FDO20220U5N   | test_104_HD | Out-of-sync   | <input checked="" type="checkbox"/> ok |
| 2 | <input type="checkbox"/> | aus-qa-pod01-lf04 | 80.80.80.69 | leaf  | FDO20220U77   | test_104_HD | Out-of-sync   | <input checked="" type="checkbox"/> ok |
| 3 | <input type="checkbox"/> | aus-qa-pod01-s... | 80.80.80.65 | spine | SAL2016NXX2   | test_104_HD | Out-of-sync   | <input checked="" type="checkbox"/> ok |

7. Click **Save & Deploy** again to sync configurations between the switches and DCNM.

The **Saving Fabric Configuration** message comes up immediately. This indicates that overlay and underlay network migration, and switch and port channel settings migration to DCNM is initiated.

After the migration of underlay and overlay networks, the **Config Deployment** window is displayed.

## Config Deployment



Step 1. Configuration Preview &gt; Step 2. Configuration Deployment Status &gt;

| Switch Name      | IP Address  | Switch Serial | Preview Config | Status      | Re-sync | Progress |
|------------------|-------------|---------------|----------------|-------------|---------|----------|
| aus-qa-pod01-... | 80.80.80.68 | FDO20220U5N   | 498 lines      | Out-of-sync |         | 100%     |
| aus-qa-pod01-... | 80.80.80.65 | SAL2016NXX2   | 0 lines        | In-sync     |         | 100%     |
| aus-qa-pod01-... | 80.80.80.69 | FDO20220U77   | 534 lines      | Out-of-sync |         | 100%     |

Deploy Config

The **Preview Config** column is updated with entries denoting a specific number of lines.

We strongly recommend that you preview the configuration before proceeding to deploy it on the switches. Click a **Preview Config** column entry. The **Config Preview** window is displayed. This window lists the pending configurations on the switch. The **Side-by-side Comparison** tab displays the running configuration and expected configuration side-by-side.

## Config Preview - Switch 80.80.80.68



Pending Config

Side-by-side Comparison

```

router bgp 65500
  no neighbor 10.96.32.2
  nxapi http port 80
  vpc domain 998
  auto-recovery reload-delay 360
  configure profile Auto_Net_VNI30113_VLAN113
  vlan 113
  vn-segment 30113
  name aus-qa-sf1-prim
  interface vlan113
    description aus-qa-sf1-prim
  vrf member qa:common
  no ip redirects
  no ipv6 redirects
  ip address 172.18.113.1/24 tag 12345
  ip dhcp relay address 172.20.16.79
  fabric forwarding mode anycast-gateway
  no shutdown
  interface nve1
    member vni 30113
  mcast-group 239.1.1.20
  suppress-arp
  evpn
  
```

Close the **Config Preview** window.

- Click **Deploy Config** at the bottom part of the **Config Deployment** window to initiate pending configuration onto the switch. The **Status** column displays the completion state. For a failed state, investigate the reason for failure to address the issue.

Config Deployment ✕

Step 1. Configuration Preview > Step 2. Configuration Deployment Status >

| Switch Name      | IP Address  | Status    | Status Description      | Progress |
|------------------|-------------|-----------|-------------------------|----------|
| aus-qa-pod01-... | 80.80.80.65 | COMPLETED | No Commands to execute. | 100%     |
| aus-qa-pod01-... | 80.80.80.69 | COMPLETED | Deployed successfully   | 100%     |
| aus-qa-pod01-... | 80.80.80.68 | COMPLETED | Deployed successfully   | 100%     |

[Close](#)

The progress bar shows 100% for each switch. After correct provisioning and successful configuration compliance, close the **Config Deployment** window.

In the fabric topology window, all imported switch instances are displayed in green color, indicating successful configuration. Also, the **Migration Mode** label is not displayed on any switch icon.

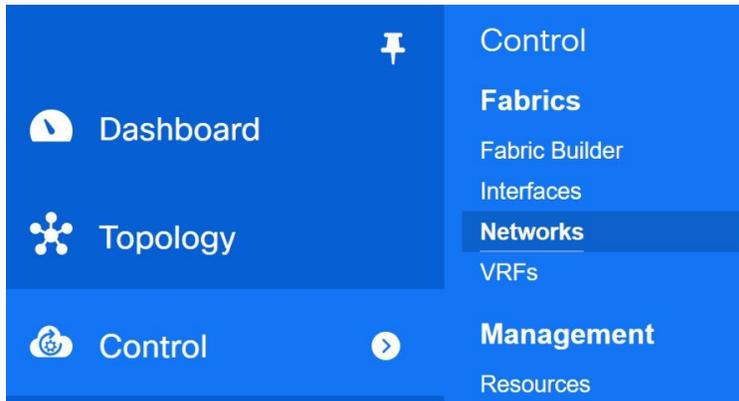
The screenshot displays the Cisco Data Center Network Manager (DCNM) interface, specifically the Fabric Builder for a test environment named 'test\_104\_HD'. The interface shows a network topology with three green circular nodes connected by lines. The top node is labeled 'aus-qa-...01-sp'. The bottom-left node is labeled 'aus-qa-pod01-lf03'. The bottom-right node is partially labeled 'aus-o'. An 'Actions' menu is open on the left, listing various options: Tabular view, Refresh topology, Save layout, Delete saved layout, Hierarchical (dropdown), Restore Fabric, Re-sync Fabric, Add switches, and Fabric Settings.

This completes the migration process of bottom-up VXLAN fabric to DCNM.

Now, you can add new switches and provision overlay networks for your fabric. For details, refer the respective section in the Fabrics topic in the configuration guide.

You can also verify the migrated networks by following the below steps.

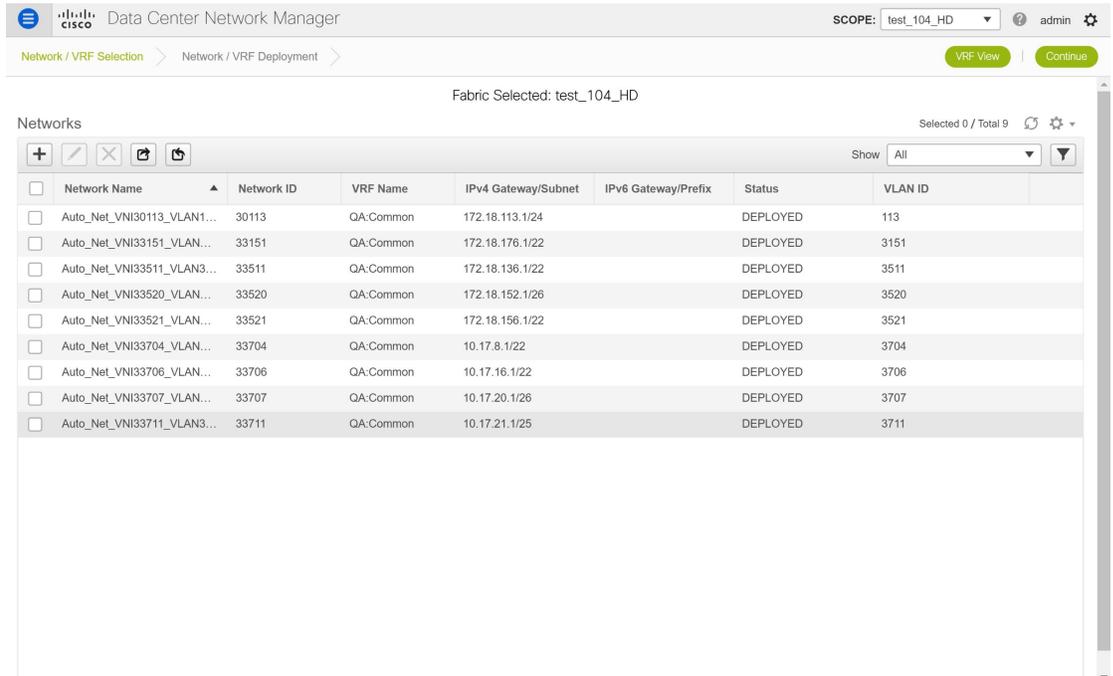
1. Choose **Control > Fabrics > Networks**.



2. Select the fabric from the **SCOPE** drop-down list in the **Networks** window.



3. Check the networks that are migrated from the bottom-up VXLAN fabric and their deployment status.



# Resolving Config Compliance Error on Switches with Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) Images

After brownfield deployment of Cisco Nexus 9300 Series switches and Cisco Nexus 9500 Series switches with X9500 line cards with Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images, config compliance difference is displayed. You need to remove the `tcam_pre_config_vxlan` policy from these switches to resolve the config compliance error.

## Resolving Config Compliance Error on Switches Post Brownfield Deployment

The following procedure shows how to remove the `tcam_pre_config_vxlan` policy from switches after brownfield deployment.

1. Choose **Control > Fabrics > Fabric Builder**.
2. Click the brownfield fabric that contains a Cisco Nexus 9300 Series switch or Cisco Nexus 9500 Series switches with X9500 line cards in the **Fabric Builder** window.
3. (Optional) Click **Save & Deploy** to see the Config Compliance error.

### Config Deployment ✕

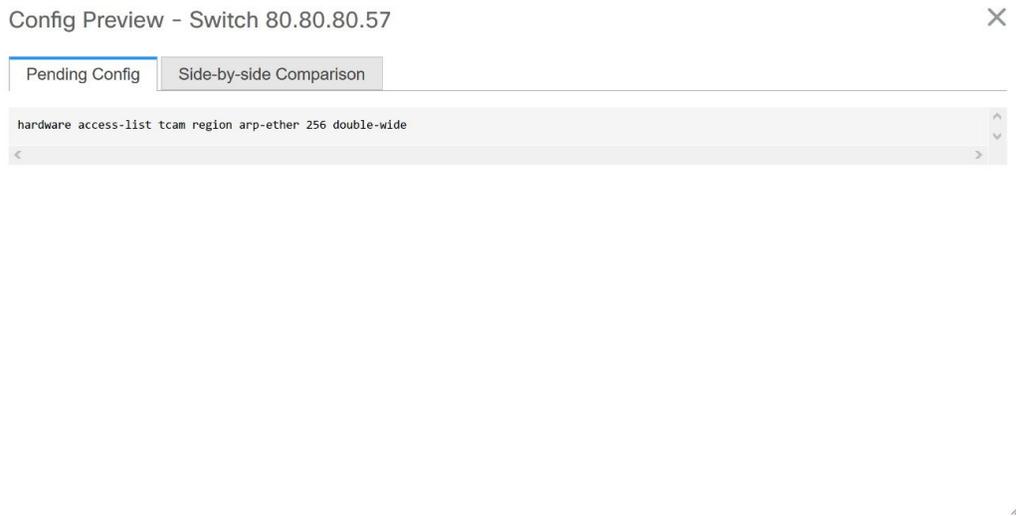
Step 1. Configuration Preview >
Step 2. Configuration Deployment Status >

| Switch Name      | IP Address  | Switch Serial | Preview Config | Status      | Re-sync | Progress   |
|------------------|-------------|---------------|----------------|-------------|---------|--|
| n9k7_bp2-lfsw... | 80.80.80.57 | SAL1833YM64   | 1 lines        | Out-of-sync |         | <div style="width: 100%; height: 10px; background-color: #76923c; border: 1px solid #ccc;"></div> 100% |
| n9k8_bp2-sps...  | 80.80.80.58 | SAL1833YM0V   | 0 lines        | In-sync     |         | <div style="width: 100%; height: 10px; background-color: #76923c; border: 1px solid #ccc;"></div> 100% |

[Deploy Config](#)

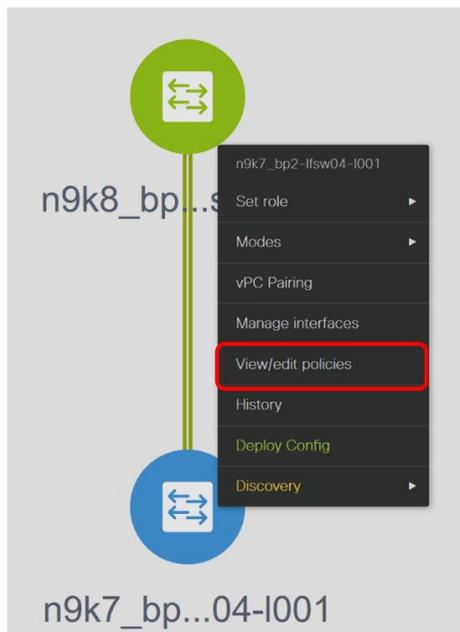
4. (Optional) Click the entry showing **1 lines** under the **Preview Config** column.

You can see the TCAM command under the **Pending Config** tab in the **Config Preview** window.



Close the **Config Preview** window.

5. Right-click a switch and click **View/Edit Policies**.



6. Search for the **tcam\_pre\_config\_vxlan** policy in the **Template** search field.
7. Select the **tcam\_pre\_config\_vxlan** policy and click the **Delete** icon to delete the policy.

View/Edit Policies for n9k7\_bp2-lfsw04-I001 ( SAL1833YM64 )

| Template              | Priority | Fabric Name | Serial Number | Editable | Entity Type | Entity Name | Source |
|-----------------------|----------|-------------|---------------|----------|-------------|-------------|--------|
| 1_pre_config_vxlan x  |          |             |               |          |             |             |        |
| tcam_pre_config_vxlan | 151      | test        | SAL1833YM64   | true     | SWITCH      | SWITCH      |        |

Close the **View/Edit Policies** window.

- (Optional) Click **Save & Deploy** to verify whether there are any pending configs.

Config Deployment

| Switch Name      | IP Address  | Switch Serial | Preview Config | Status  | Re-sync | Progress |
|------------------|-------------|---------------|----------------|---------|---------|----------|
| n9k7_bp2-lfsw... | 80.80.80.57 | SAL1833YM64   | 0 lines        | In-sync |         | 100%     |
| n9k8_bp2-sps...  | 80.80.80.58 | SAL1833YM0V   | 0 lines        | In-sync |         | 100%     |

Deploy Config

### Resolving Config Compliance Error on Switches for RMA, and Write Erase and Reload Operations

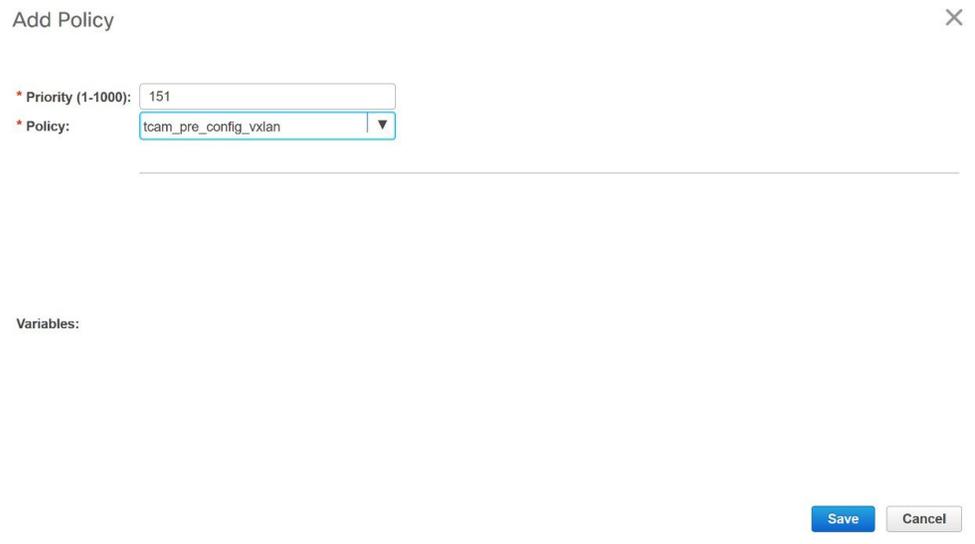
Perform the following procedure before you perform RMA or Write Erase and Reload operation on Cisco Nexus 9300 Series switches and Cisco Nexus 9500 Series switches with X9500 line cards with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images.

- Choose **Control > Fabrics > Fabric Builder**.

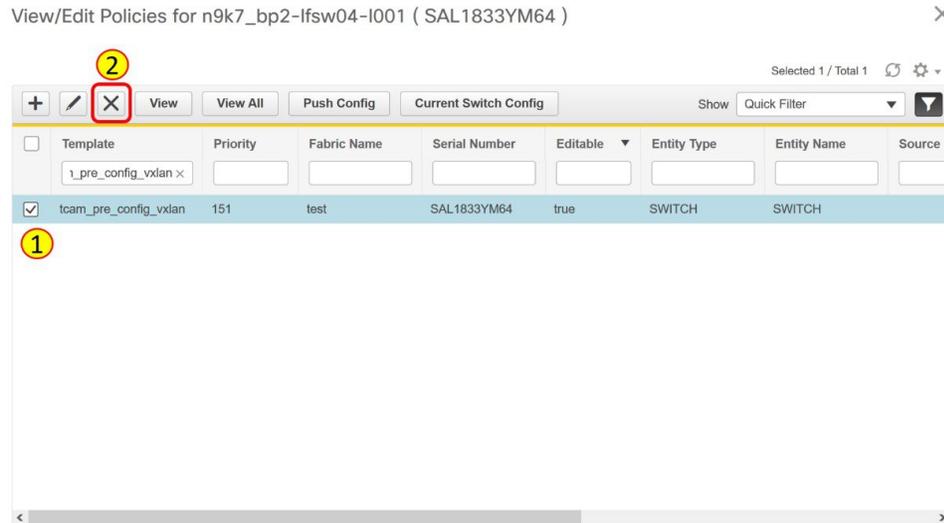
2. Click the brownfield fabric that contains the specified switches with Cisco images.
3. Right-click the switch and click **View/Edit Policies**.
4. Click the **Add** icon.



5. Enter 151 in the Priority (1-1000) field and select **tcam\_pre\_config\_vxlan** from the **Policy** drop-down list.



6. Click **Save**.
7. Complete the RMA or Write Erase and Reload operation.  
After the switch is online, it will be Out-of-Sync.
8. Right-click a switch and click **View/Edit Policies**.
9. Search for the **tcam\_pre\_config\_vxlan** policy in the **Template** search field.
10. Select the **tcam\_pre\_config\_vxlan** policy and click the **Delete** icon to delete the policy.



Close the **View/Edit Policies** window.

## Modifying VLAN Names in a Switch with Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) Images

Post brownfield migration, the VLAN name for the network or VRF is not captured in the overlay profile if at least one of the non-spine switches have the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images.

This procedure shows how to check the VLAN name and modify it.

### Procedure

- 
- Step 1** Choose **Control > Fabrics > Networks**.
  - Step 2** From the **SCOPE** drop-down list, select a fabric containing the non-spine switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images.
  - Step 3** Select a check box for a network in the **Networks** window and click the **Edit Network** icon.

The screenshot shows the 'Edit Network' window with the 'Network Profile' section expanded. The 'General' tab is selected. The 'Vlan Name' field is empty and highlighted with a red box. Other fields include IPv4 Gateway/NetMask (172.16.6.1/24), IPv6 Gateway/Prefix (1111::2222/48), Interface Description, MTU for L3 interface (1500), IPv4 Secondary GW1 (2.2.2.2/24), and IPv4 Secondary GW2 (3.3.3.3/24). There are 'Save' and 'Cancel' buttons at the bottom right.

In the **Edit Network** window, the **Vlan Name** field is empty because DCNM has not captured this info in the overlay profile. Instead, the VLAN name is captured in the freeform config associated with the overlay network or VRF.

**Note** If a VLAN did not have a name before the brownfield migration, you can add the name in the **Vlan Name** field in the **Edit Network** window.

Close the **Edit Network** window.

**Step 4** Click **Continue** in the **Networks** window.

**Step 5** Double-click a switch in the **Topology View** window.

**Step 6** In the **Network Attachment** window for a switch, click the **Freeform config** button under the **CLI Freeform** column.

The screenshot shows the 'Network Attachment - Attach networks for given switch(es)' window. The 'Fabric Name' is 'test'. Under 'Deployment Options', there is a table with columns: Switch, VLAN, Interfaces, CLI Freeform, and Status. The first row is selected, and the 'Freeform config' button in the 'CLI Freeform' column is highlighted with a red box. A 'Save' button is at the bottom right.

| Switch                              | VLAN           | Interfaces | CLI Freeform    | Status   |
|-------------------------------------|----------------|------------|-----------------|----------|
| <input checked="" type="checkbox"/> | n9k7_bp2-if... | 6          | Freeform config | DEPLOYED |

**Step 7** Verify the VLAN name in the **Free Form Config** window.

Free Form Config -n9k7\_bp2-lfsw04-l001 (Auto\_Net\_VNI20006\_VLAN6) ✕

```
vlan 6
name 0006-BP2-IALAB-IP-Storage_172_16
vn-segment 20006
interface Vlan6
no shutdown
vrf member IALAB
no ip redirects
ip address 172.16.6.1/24
ip address 2.2.2.2/24 secondary
ip address 3.3.3.3/24 secondary
ipv6 address 1111::2222/48
no ipv6 redirects
fabric forwarding mode anycast-gateway
ip dhcp relay address 10.1.1.1
ip dhcp relay address 10.3.3.3
```

[Save Config](#)

**Step 8** Modify the VLAN name in the **Free Form Config** window and click **Save Config**.

Here is an example:

```
vlan 6
name Storage_172_16_Deb
vn-segment 20006
interface Vlan6
.
.
.
```

Free Form Config -n9k7\_bp2-lfsw04-l001 (Auto\_Net\_VNI20006\_VLAN6) ✕

```
vlan 6
name Storage_172_16_Deb
vn-segment 20006
interface Vlan6
no shutdown
vrf member IALAB
no ip redirects
ip address 172.16.6.1/24
ip address 2.2.2.2/24 secondary
ip address 3.3.3.3/24 secondary
ipv6 address 1111::2222/48
no ipv6 redirects
fabric forwarding mode anycast-gateway
ip dhcp relay address 10.1.1.1
ip dhcp relay address 10.3.3.3
```

[Save Config](#)

**Step 9** Click **Save** in the **Network Attachment** window.

**Step 10** Click **Deploy** in the **Networks** window.

The modified VLAN name in the selected network is deployed on the switch.

## Changing a Brownfield Imported BIDIR Configuration

This procedure shows how to change a brownfield imported BIDIR configuration to use the configuration generated by **Fabric Builder**.

### Procedure

---

- Step 1** Choose **Control > Fabrics > Networks**.
- Step 2** Click the brownfield fabric.
- Step 3** Click **Tabular View** under the **Actions Panel** in the **Fabric Builder** window.
- Step 4** Select all the devices and click the **View/Edit Policies** icon.
- Step 5** Delete the following policies for all the devices in the **View/Edit Policies** window
- **base\_pim\_bidir\_11\_1**
  - If there is 1 RP in the fabric, delete the **rp\_lb\_id** policy.  
If there are 2 RPs in the fabric, delete the **phantom\_rp\_lb\_id1** and **phantom\_rp\_lb\_id2** policies.
- Step 6** Close the **View/Edit Policies** window.
- Step 7** Click the **Manage Interfaces** button in the **Fabric Builder** window.
- Step 8** Delete all the RP loopback interfaces in the **Interfaces** window and close this window.
- Step 9** Click **Save & Deploy** in the **Fabric Builder** window.
- This action generates a new set of BIDIR-related configuration based on the fabric settings for the devices.
- 

## Manually Adding PIM-BIDIR Configuration for Leaf or Spine Post Brownfield Migration

After brownfield migration, if you add new spine or leaf switches, you should manually configure the PIM-BIDIR feature.

The following procedure shows how to manually configure the PIM-BIDIR feature for a new Leaf or Spine:

### Procedure

---

- Step 1** Check the **base\_pim\_bidir\_11\_1** policies that are created for an RP added through the brownfield migration. Check the RP IP and Multicast Group used in each **ip pim rp-address RP\_IP group-list MULTICAST\_GROUP bidir** command.
- Step 2** Add respective **base\_pim\_bidir\_11\_1** policies from the **View/Edit Policies** window for the new Leaf or Spine, push the config for each **base\_pim\_bidir\_11\_1** policy.
- 

## Migrating an MSD Fabric with Border Gateway Switches

When you migrate an existing MSD fabric with a border gateway switch into DCNM, make sure to note the following guidelines:

- Underlay Multisite peering: The eBGP peering and corresponding routed interfaces for underlay extensions between sites are captured in **switch\_freeform** and **routed\_interfaces**, and optionally in the **interface\_freeform** configs. This configuration includes all the global configs for multisite. Loopbacks for EVPN multisite are also captured via the appropriate interface templates.
  - Overlay Multisite peering: The eBGP peering is captured as part of **switch\_freeform** as the only relevant config is under **router bgp**.
  - Overlays containing Networks or VRFs: The corresponding intent is captured with the profiles on the Border Gateways with **extension\_type = MULTISITE**.
1. Create all the required fabrics including the Easy\_Fabric\_11\_1 and External\_Fabric\_11\_1 fabrics with the required fabric settings. Disable the Auto VRF-Lite options as mentioned above. For more information, refer to *Creating VXLAN EVPN Fabric* and *External Fabric* sections.
  2. Import all the switches into all the required fabrics and set roles accordingly.
  3. Click **Save & Deploy** in each of the fabrics and ensure that the Brownfield Migration process reaches the 'Deployment' phase. Now, do not click **Deploy Config**.
  4. Create an **MSD\_Fabric\_11\_1** fabric with the required fabric settings and disable the **Auto MultiSite IFC** options as shown in Guidelines. For more information, see *Creating an MSD Fabric in Cisco DCNM LAN Fabric Configuration Guide*.
  5. Move all the member fabrics into the MSD. Do not proceed further till this step is completed successfully. For more information, see *Moving the Member1 Fabric Under MSD-Parent-Fabric in Cisco DCNM LAN Fabric Configuration Guide*.

**Note**

The Overlay Networks and VRFs definitions in each of the Easy Fabrics must be symmetric for them to get added successfully to the MSD. Errors will be reported if any mismatches are found. These must be fixed by updating the overlay information in the fabric(s) and added to the MSD.

6. Create all the Multisite Underlay IFCs such that they match the IP address and settings of the deployed configuration. Navigate to **Tabular View** and edit the IFC links.

Below is an example IFC Edit Link window.

## Link Management - Edit Link



|                         |                               |
|-------------------------|-------------------------------|
| * Link Type             | Inter-Fabric                  |
| * Link Sub-Type         | MULTISITE_UNDERLAY            |
| * Link Template         | multisite_underlay_setup_11_1 |
| * Source Fabric         | fab1-11-2                     |
| * Destination Fabric    | CORE-BGW                      |
| * Source Device         | Leaf4-BL2                     |
| * Source Interface      | Ethernet1/9                   |
| * Destination Device    | N7k1-CORE                     |
| * Destination Interface | Ethernet1/1                   |

## Link Profile

## General

## Advanced

|                     |              |                               |
|---------------------|--------------|-------------------------------|
| * BGP Local ASN     | 65000        | ? Local BGP Autonomous S      |
| * IP Address/Mask   | 10.10.1.1/30 | ? IP address with mask (e.g.  |
| * BGP Neighbor IP   | 10.10.1.2    | ? Neighbor IP address         |
| * BGP Neighbor ASN  | 4000         | ? Neighbor BGP Autonomou      |
| * BGP Maximum Paths | 1            | ? Maximum number of iBGP,     |
| * Routing TAG       | 54321        | ? Routing tag associated with |

Save



**Note** Additional interface configurations must be added to the Source/Destination interface freeform fields in the Advanced section as needed.

For more information, see *Configuring Multi-Site Overlay IFCs*.

7. Create all the Multisite Overlay IFCs such that they match the IP address and settings of the deployed configuration. You will need to add the IFC links. For more information, see *Configuring Multi-Site Overlay IFCs*.
8. If there are VRF-Lite IFCs also, create them as well.



**Note** If the Brownfield Migration is for the case where Configuration Profiles already exist on the switches, the VRF-Lite IFCs will be created automatically in Step #3.

9. If Tenant Routed Multicast (TRM) is enabled in the MSD fabric, edit all the TRM related VRFs and Network entries in MSD and enable the TRM parameters.

This step needs to be performed if TRM is enabled in the fabric. If TRM is not enabled, you still need to edit each Network entry and save it.

10. Now click **Save & Deploy** in the MSD fabric, but, do not click **Deploy Config**.

11. Navigate to each member fabric, click **Save & Deploy**, and then click **Deploy Config**.

This completes the Brownfield Migration. You can now manage all the networks or VRFs for BGWs by using the regular DCNM Overlay workflows.



## CHAPTER 10

# Template Usage in Cisco DCNM LAN Fabric Deployment

|              |                                      |  |
|--------------|--------------------------------------|--|
| templateType | Specifies the type of Template used. | <ul style="list-style-type: none"><li>• CLI</li><li>• POLICY</li><li>• SHOW</li><li>• PROFILE</li><li>• ABSTRACT</li></ul> |
|--------------|--------------------------------------|--|

- [Policy Template, on page 499](#)
- [Fabric Template, on page 503](#)
- [Profile Template, on page 503](#)
- [Viewing, Editing, and Adding Policies, on page 504](#)
- [Deploying New Configurations, on page 508](#)
- [switch\\_freeform Template Usage, on page 509](#)
- [Changing the Contents of a Template in Use, on page 512](#)

## Policy Template

For the policy template, there are 2 template content types: CLI and PYTHON. With CLI content type, the policy templates are basically parameterized CLI templates. They can have a lot of variables and CLIs. Typically, CLI policy templates are small and do not have any if-else-for etc. like constructs. An example CLI policy template for AAA server configuration is shown below:

The screenshot shows the Cisco Data Center Network Manager interface. The breadcrumb navigation is 'Control / Template Library'. The selected template is 'aaa\_radius'. The template content is displayed in a code editor with the following text:

```

1  ##template variables
2
3  #   Copyright (c) 2018 by Cisco Systems, Inc.
4  #   All rights reserved.
5
6  @(DisplayName="AAA Server Name/IP", Description="Name or IPv4/IPv6 Address of an AAA Server")
7  ipAddressWithoutPrefix AAA_SERVER;
8
9  @(DisplayName="AAA group", Description="Name of AAA Group")
10 - string AAA_GROUP {
11     minLength = 1;
12     maxLength = 127;
13 };
14
15 ##
16 ##template content
17
18 aaa group server radius $AAAA_GROUP$$
19 server $AAAA_SERVER$$
20
21 ##

```

But you can also have policy templates of template content type PYTHON. Essentially, this allows multiple CLI policy templates to be combined together with a common “source” so that they get all applied/un-applied at one go. For example, when you want to create a vPC host port, it has to be created symmetrically on both peers that are part of the vPC pair. In addition, you have to create port-channel, member interfaces, channel-group, etc. This is why a python vPC host policy template has been added. An example interface PYTHON template for setting up a routed interface is shown below:

Control / Template Library

Template Content: 

int\_routed\_host\_11\_1 0 Errors, 0 Warnings     

```

1  ##template variables
2
3  # Copyright (c) 2018 by Cisco Systems, Inc.
4  # All rights reserved.
5
6  @(IsInternal=true)
7  string SERIAL_NUMBER;
8
9  @(PrimaryAssociation=true, IsInternal=true)
10 interface INTF_NAME;
11
12 @(IsMandatory=false, DisplayName="Interface VRF", Description="Interface VRF name, default VRF if not specified")
13 string INTF_VRF {
14     minLength = 1;
15     maxLength = 32;
16 };
17
18 @(IsMandatory=false, DisplayName="Interface IP", Description="IP address of the interface")
19 ipv4Address IP;
20
21 @(IsMandatory="IP!=null", DisplayName="IP Netmask Length", Description="IP netmask length used with the IP address (Min:1, Max:31)")
22 integer PREFIX {
23     min = 1;
24     max = 31;
25 };
26
27 @(IsMandatory=false, DisplayName="Routing TAG", Description="Routing tag associated with interface IP")
28 string ROUTING_TAG;
29
30 @(DisplayName="MTU", IsMTU=true, Description="MTU for the interface (Min:576, Max:9216)")
31 integer MTU {
32     min = 576;
33     max = 9216;
34     defaultValue=9216;
35 };
36
37 @(DisplayName="SPEED", Description="Interface Speed")
38 enum SPEED {
39     validValues=Auto,100Mb,1Gb,10Gb,25Gb,40Gb,100Gb;
40     defaultValue=Auto;
41 };
42
43 @(IsMandatory=false, DisplayName="Interface Description", Description="Add description to the interface (Max Size 254)")
44 string DESC {
45     minLength = 1;
46     maxLength = 254;
47 };
48
49 @(IsMandatory=false, IsMultiLineString=true, DisplayName="Freeform Config", Description="Additional CLI for the interface")
50 string CONF;
51
52 @(DisplayName="Enable Interface", Description="Uncheck to disable the interface")
53 boolean ADMIN_STATE {
54     defaultValue=true;
55 };
56
57 ##
58 ##template content
59
60 from com.cisco.dcbu.vincil.rest.services.jython import PTIWrapper
61 from com.cisco.dcbu.vincil.rest.services.jython import Wrapper
62 from com.cisco.dcbu.vincil.rest.services.jython import WrappersResp
63 from utility import *
64
65 def add():
66     try:
67         if CONF != "":
68             respObj, conf = Util.adjustIntfFreeformConfig(SERIAL_NUMBER, INTF_NAME, CONF)
69             if respObj.isRetCodeFailure():
70                 return respObj
71
72     # modify to be done, calling delete now to clean up PTIs before add
73     delete()
74
75     intfVrf = "default"
76     try:
77         if INTF_VRF != "":
78             intfVrf = INTF_VRF
79     except:
80         Wrapper.print("Switch/Intf = [%s/%s] - Template[int_routed_host_11_1]: INTF_VRF not defined" %
81             (SERIAL_NUMBER, INTF_NAME))
82         pass
83
84     routingTag = ""
85     try:
86         if ROUTING_TAG != "":
87             routingTag = ROUTING_TAG
88     except:
89         Wrapper.print("Switch/Intf = [%s/%s] - Template[int_routed_host_11_1]: ROUTING_TAG not defined" %
90             (SERIAL_NUMBER, INTF_NAME))
91         pass
92
93     # routed_interface has only one CLI command: no switchport
94     # It must be configured before interface_vrf
95     # p2p_routed_interface that configures the IP address must come after interface_vrf
96     Util.exe(PTIWrapper.createOrUpdate(SERIAL_NUMBER, "INTERFACE",
97         INTF_NAME, INTF_NAME,
98         ConfigPriority.CONFIG_PRIO_INTF,
99         "routed_interface",
100         {"INTF_NAME": INTF_NAME}))
101
102     if intfVrf != "default":
103         # Create/Update PTI for interface VRF
104         Util.exe(PTIWrapper.createOrUpdate(SERIAL_NUMBER, "INTERFACE",
105             INTF_NAME, INTF_NAME,
106             ConfigPriority.CONFIG_PRIO_INTF_SUB_LVL1,
107             "interface_vrf",
108             {"INTF_NAME": INTF_NAME, "INTF_VRF": intfVrf}))
109
110     if IP != "":
111         if routingTag == "":
112             Util.exe(PTIWrapper.createOrUpdate(SERIAL_NUMBER, "INTERFACE",
113                 INTF_NAME, INTF_NAME,
114                 ConfigPriority.CONFIG_PRIO_INTF_SUB_LVL2,
115                 "p2p_routed_interface",
116                 {"INTF_NAME": INTF_NAME, "IP": IP, "PREFIX": PREFIX}))

```

Each policy template has a template subtype like DEVICE, INTERFACE, etc. This allows the right policy template to appear at the right selection point. For example, in the Interface window, you will only see the interface policy templates.

| Name                              | Supported Platforms | Tags            | Template ... | Template ... | Published | Modified T... | D... |
|-----------------------------------|---------------------|-----------------|--------------|--------------|-----------|---------------|------|
| csr1kv_loopback                   | CSR1KV              | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| epi_routed_intf                   | N9K                 | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| GigabitEthernet                   | CSR1KV              | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| GigabitEthernet_freem             | CSR1KV              | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| int_access_host_11_1              | All                 | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| int_loopback_11_1                 | All                 | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| int_mgmt_11_1                     | N9K                 | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| int_monitor_ethernet_11_1         | N9K                 | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| int_monitor_port_channel_11_1     | N9K                 | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| int_port_channel_access_host_11_1 | All                 | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| int_port_channel_trunk_host_11_1  | All                 | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| int_routed_host_11_1              | All                 | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| int_subif_11_1                    | All                 | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| int_trunk_host_11_1               | All                 | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| int_vpc_access_host_11_1          | All                 | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |
| int_vpc_trunk_host_11_1           | All                 | [interface_...] | POLICY       | INTERFAC...  | false     | 2019-06-03... |      |

In the View/Edit Policies window on the Fabric Builder, you will only see device policy templates.

| Name                     | Supported Platforms | Tags | Template ... | Template ... | Published | Modified T... | D... |
|--------------------------|---------------------|------|--------------|--------------|-----------|---------------|------|
| aaa_radius               | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| aaa_radius_deadtime      | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| aaa_radius_key           | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| aaa_radius_src_interface | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| aaa_radius_use_vrf       | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| aaa_tacacs               | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| aaa_tacacs_key           | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| aaa_tacacs_src_interface | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| aaa_tacacs_use_vrf       | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| anycast_gateway          | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| anycast_rp               | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| azure_network_selector   | CSR1KV              |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| banner                   | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| base_aaa                 | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| base_bgp                 | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| base_bgp_external        | N9K,N7K             |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |
| base_dhcp                | N9K                 |      | POLICY       | DEVICE       | false     | 2019-06-03... |      |

You can make a copy of any of these templates and customize them as per their needs. That is the typical use-case for customization. **Do not** modify existing policies but make a copy, and then customize as per the requirements. Otherwise, after a DCNM upgrade, the changes may be lost.

In general, a template already in use, meaning one that is already applied to some switch within any fabric, cannot be edited.



**Note** No Type-CLI templates are used in the LAN fabric installation mode. They are all replaced with more powerful Policy templates which are a super set.

## Fabric Template

A fabric template is basically a python template, specifically jython, which is java + python. A fabric template is quite comprehensive, and in that it embeds the rules that are required for deploying a fabric, including all the logic required to generate intended configuration of all switches within the entire fabric. Configuration is generated based on published Cisco best practice guidelines. In addition to the embedded rules, the fabric template also integrates with other entities such as resource manager, topology database, device roles, configuration compliance, etc. and generates the configuration accordingly for all the devices in the fabric. This is the inherent part of the DCNM fabric builder.

The expectation is that users will not create their own fabric templates. DCNM provides a few fabric templates out of the box such as Easy Fabric, External Fabric, MSD Fabric, eBGP Fabric (introduced in DCNM 11.2).

| Name                 | Supported Platforms | Tags | Template ... | Template ... | Published | Modified T... | D... |
|----------------------|---------------------|------|--------------|--------------|-----------|---------------|------|
| Easy_Fabric_11_1     | All                 |      | FABRIC       | NA           | false     | 2019-06-03... | F... |
| Easy_Fabric_eBGP     | All                 |      | FABRIC       | NA           | false     | 2019-06-03... | F... |
| External_Fabric_11_1 | All                 |      | FABRIC       | NA           | false     | 2019-06-03... | F... |
| MSD_Fabric_11_1      | All                 |      | FABRIC       | NA           | false     | 2019-06-03... | F... |

## Profile Template

A profile template is used for provisioning of overlays (networks or VRFs). The idea is that when you apply some overlay configuration, there are multiple pieces of configurations that should go together. For example, valid layer-3 network configuration in a VXLAN EVPN fabric requires VLAN, SVI, int nve config, EVPN route-target, etc. All of these pieces are put together into what is called a configuration profile (NX-OS construct) and then effectively applied at one go. Either the whole configuration profile gets applied or nothing gets applied, on the switch. In this way, you are not left with any dangling or stray configurations on the switches. For any kind of overlay configurations, whether it is on the leaf or on the borders, DCNM employs profile templates.

There are four kinds of profile templates that are distinguished with tags as depicted below:

- Network Profile (applied to all devices with role leaf)
- Network Extension Profile (applied to all devices with role 'border\*')
- VRF Profile (applied to all devices with role leaf)
- VRF Extension Profile (applied to all devices with role 'border\*')

The screenshot shows the Cisco Data Center Network Manager (DCNM) interface. The breadcrumb navigation is 'Control / Template Library'. Below this, there is a 'Templates' section with a search bar and a table of templates. The table has columns for Name, Supported Platforms, Tags, Template Type, Published, Modified T..., and D... (Details). The table lists several templates, including 'base\_external\_router', 'Default\_Network\_Extension\_Universal', 'Default\_Network\_Universal', 'Default\_VRF\_Extension\_Universal', 'Default\_VRF\_Universal', 'ext\_base\_setup', 'ext\_fabric\_intf', 'ext\_fabric\_multisite\_intf\_11\_1', 'ext\_multisite\_overlay\_setup\_11\_1', 'ext\_multisite\_rs\_base\_feature', and 'ext\_multisite\_rs\_base\_setup'.

| Name                                | Supported Platforms | Tags           | Template ... | Template ... | Published | Modified T... | D... |
|-------------------------------------|---------------------|----------------|--------------|--------------|-----------|---------------|------|
| base_external_router                | N9K                 |                | PROFILE      | NA           | false     | 2019-06-03... | s... |
| Default_Network_Extension_Universal | All                 | [networkEx...  | PROFILE      | VXLAN        | false     | 2019-06-03... | D... |
| Default_Network_Universal           | All                 | [network]      | PROFILE      | VXLAN        | false     | 2019-06-03... | D... |
| Default_VRF_Extension_Universal     | All                 | [vrfExtension] | PROFILE      | VXLAN        | false     | 2019-06-03... | D... |
| Default_VRF_Universal               | All                 | [vrf]          | PROFILE      | VXLAN        | false     | 2019-06-03... | D... |
| ext_base_setup                      | All                 | [borderBase]   | PROFILE      | VXLAN        | false     | 2019-06-03... |      |
| ext_fabric_intf                     | All                 |                | PROFILE      | VXLAN        | false     | 2019-06-03... |      |
| ext_fabric_multisite_intf_11_1      | All                 |                | PROFILE      | VXLAN        | false     | 2019-06-03... |      |
| ext_multisite_overlay_setup_11_1    | All                 | [multiSiteO... | PROFILE      | VXLAN        | false     | 2019-06-03... |      |
| ext_multisite_rs_base_feature       | N9K,N7K             | [multiSiteO... | PROFILE      | VXLAN        | false     | 2019-06-03... | s... |
| ext_multisite_rs_base_setup         | N9K                 | [multiSiteO... | PROFILE      | VXLAN        | false     | 2019-06-03... | s... |

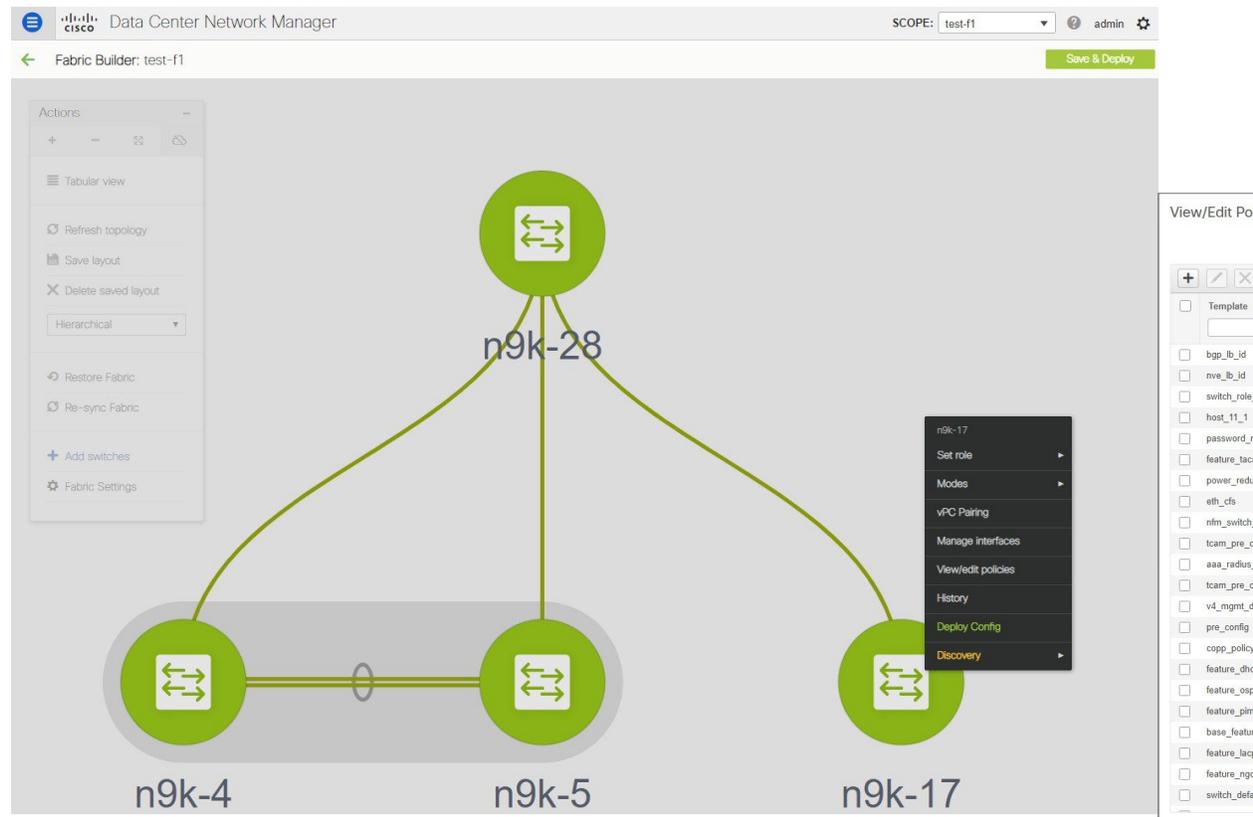
For more information about how to apply overlay configuration via the Networks & VRFs workflow in DCNM, see [Creating and Deploying Networks and VRFs](#).

### Additional Notes

When a policy or profile template is applied, an instance is created for each application of the template. The common terminology used for this is Policy Template Instance or PTI. A PTI is effectively a policy or profile template + the Name-value pairs that give it a specific instance, post substitution. PTIs created for a device can be viewed under the View/Edit policies option for that device in Fabric Builder. In the tabular view, the View/Edit policies button allows selection and bulk creation/deletion of policies across a subset of devices in the entire fabric. For more information, see [Viewing and Editing Policies](#).

## Viewing, Editing, and Adding Policies

To navigate to the View/Edit Policies window, right-click a device in the Fabric Builder window and select View/edit policies.



The View/Edit Policies window can be used to view, edit, or create a policy for a device. Note that Interface policies can only be viewed but cannot be edited/created from the View/Edit Policies window. Interfaces can only be edited, created, or deleted from the Interfaces window.

## Viewing Policies

To view certain policies for a device, you can use filters by specifying the search criteria in the empty boxes under each field. After the policies are found, you can view the content by selecting multiple policies and clicking on the “View” button. Below are examples that show how to use filters and how to view the configuration associated with a policy instance.

### Example: Viewing Policies for a Device

Enter `tcam` in the search field to filter the templates, select the template that you want to view, and click the View button to view TCAM policies created for the device.

| Template                                       | Policy ID   | Fabric Name | Serial Number | Editable | Entity Type | Entity Name |
|--|-------------|-------------|---------------|----------|-------------|-------------|
| tcam   |             |             |               |          |             |             |
| <input type="checkbox"/> tcam_pre_config_9300  | POLICY-9300 | test-f1     | SAL18432P6M   | true     | SWITCH      | SWITCH      |
| <input type="checkbox"/> tcam_pre_config_vxlan | POLICY-9330 | test-f1     | SAL18432P6M   | true     | SWITCH      | SWITCH      |

**Example: Viewing Policies for an Interface**

Enter the interface name in the search field under Entity Name to filter interfaces. Select an interface, and click the View button to view policies created for the interface.

| Template                                     | Policy ID   | Fabric Name | Serial Number | Editable | Entity Type | Entity Name  | Source       | Priority | Content Type | Mark Deleted |
|--|-------------|-------------|---------------|----------|-------------|--------------|--------------|----------|--------------|--------------|
|  |             |             |               |          |             | Ethernet1/29 |              |          |              |              |
| <input type="checkbox"/> trunk_interface     | POLICY-9420 | test-f1     | SAL18432P6M   | false    | INTERFACE   | Ethernet1/29 | Ethernet1/29 | 350      | TEMPLATE_CLI | false        |
| <input type="checkbox"/> int_trunk_host_11_1 | POLICY-9390 | test-f1     | SAL18432P6M   | false    | INTERFACE   | Ethernet1/29 | Ethernet1/29 | 350      | PYTHON       | false        |
| <input type="checkbox"/> interface_mtu       | POLICY-9450 | test-f1     | SAL18432P6M   | false    | INTERFACE   | Ethernet1/29 | Ethernet1/29 | 352      | TEMPLATE_CLI | false        |
| <input type="checkbox"/> porttype_fast_trunk | POLICY-9520 | test-f1     | SAL18432P6M   | false    | INTERFACE   | Ethernet1/29 | Ethernet1/29 | 352      | TEMPLATE_CLI | false        |
| <input type="checkbox"/> no_shut_interface   | POLICY-9530 | test-f1     | SAL18432P6M   | false    | INTERFACE   | Ethernet1/29 | Ethernet1/29 | 352      | TEMPLATE_CLI | false        |



**Note**

- Each interface should be associated with one interface jython policy template.
- An interface jython policy template does not have CLI in its content but rather creates PTIs of CLI policy templates. All these PTIs are combined to generate a complete configuration associated with an interface.

## Editing Policies

Not all device policies can be edited from the View/Edit policies window. Only the policies that are created with an empty Source and have the flag Editable = true, can be edited.

### Procedure

- 
- Step 1** To edit a device policy, select an existing policy and click on the edit or ‘Pencil’ button. The ‘Edit Policy’ window opens.
- Step 2** After changing 1 or more Name-value pairs, press the ‘Save’ button to save the changes on the Edit Policy window.
- Step 3** To deploy the changed config, go back to the Fabric Builder window, right-click on the device and select ‘Deploy Config’.
- This will invoke Configuration Compliance to generate the pending config for the device. Pending config is the diff between the current config on the switch and the new intent config.
- Step 4** If the pending config is correct, click ‘Deploy Config’ to push the pending config onto the switch.
- 

### Example: Editing a Policy

This example shows how to change the IPv4 management default gateway.

The screenshot displays the 'Edit Policy' dialog box for policy POLICY-9140. The dialog is titled 'Edit Policy' and shows the following details:

- Policy ID: POLICY-9140
- Entity Type: SWITCH
- Template Name: v4\_mgmt\_default\_gateway
- Entity Name: SWITCH
- Priority (1-1000): 910
- General tab: Shows the 'Default Gateway' field set to 22.0.0.88. A note indicates: 'Default Gateway IP address to use with mgmt0'.
- Variables: Empty section.
- Buttons: Save, Push Config, Cancel.

In the background, the 'View/Edit Policies for n9k-17 (SAL18432P6M)' window shows a list of policies. The policy 'v4\_mgmt\_default\_gat...' (POLICY-9140) is selected. To the right, the 'Config Deployment' panel shows a table with columns for 'Switch Name' and 'IP Address', with one entry for 'n9k-17' and '22.0.0.17'.

## Adding Policies

### Procedure

**Step 1** To add a policy to a device, click the '+' button on the View/Edit Policies page.  
The 'Add Policy' windows opens.

**Step 2** From the Policy drop-down list, select a policy to be added to the device.

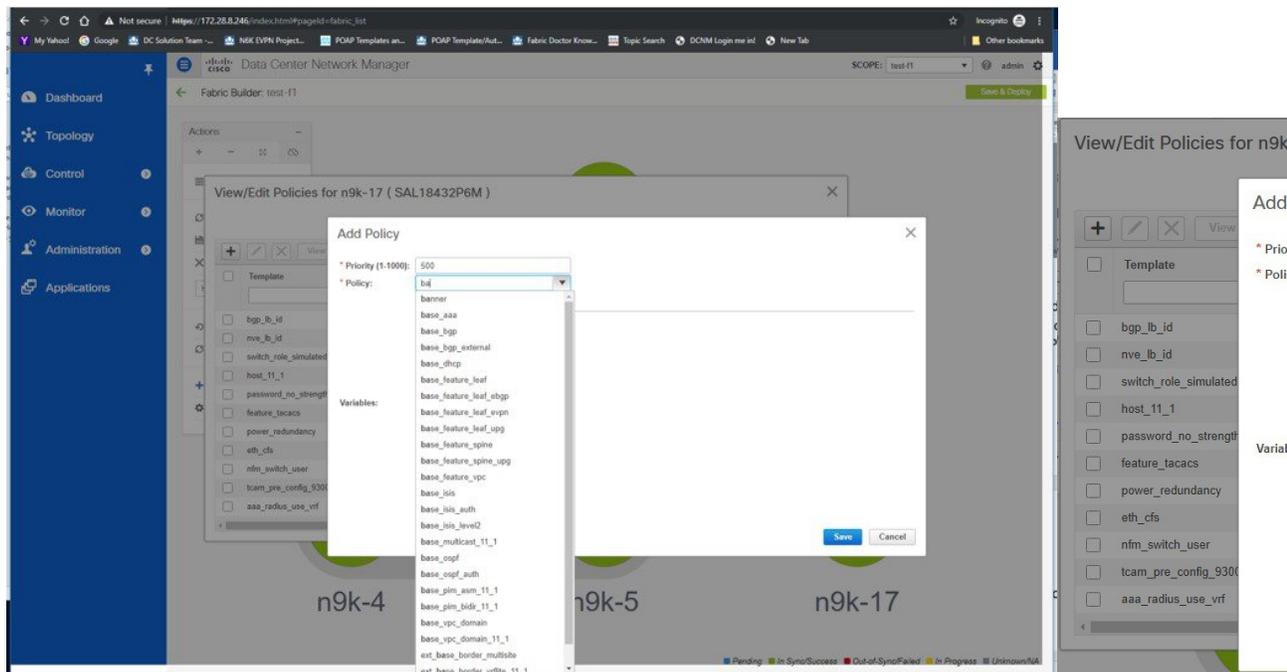
**Step 3** Set the policy priority and input the mandatory fields.

**Step 4** Click the 'Save' button to save and complete adding the policy.

**Note** Policy Priority is used to determine the order in which the configuration will be applied to the switch. Lower priority PTIs are placed before the higher priority PTIs in the expected configuration or intent and this in turn is the order to which the configuration will be pushed via the deployer module. Default priority is 500.

### Adding a Banner Policy

This example shows how to add a banner policy to a device.



## Deploying New Configurations

There are two ways to deploy the new configurations:

1. Navigate to the Fabric Builder window, right-click on the device and select 'Deploy Config' (this is the recommended way).
2. From the View/Edit Policies window, select the newly added policy, click 'View' to verify the config. If the new config looks good, click the 'Push Config' button to push the new config to the device. Note that 'Push Config' will bypass Configuration Compliance. This option should only be used for exception scenarios such as the case where a new user or SNMP user needs to be added to the switch.

## switch\_freeform Template Usage

The **switch\_freeform** is a special policy template that allows users to specify any freeform config for a device. Usage of the template is as follows:

- Specify switch-level config in the **Switch Freeform Config** parameter.
- The specified config must match the **show run** output with respect to case and newlines. Any mismatch will yield unexpected diffs during deploy.
- An internal **switch\_freeform\_config** CLI policy is created for the specified config.
- Should not use this template for interface configuration except for the SVI interface, as SVI interfaces cannot be configured on the Interfaces page currently.
- Users can create many **switch\_freeform** policies for different configs.
- **switch\_freeform** PTIs are sorted together with the other PTIs based on their policy priorities from low to high.
- A **switch\_freeform** policy can be edited before or after the config is deployed.
- If there is any change in the config content, the previously created internal **switch\_freeform\_config** policy will have its priority changed from a positive to a negative number, and a new internal policy is created for the new config.
- A **negative** priority PTI means that CLIs in the PTI need to be deleted; **Configuration Compliance** will generate the **no** commands accordingly.
- Deleting a **switch\_freeform** policy will change the PTI priority of its internal policy to a negative number.

The following section shows how to create a **switch\_freeform** policy, deploy the policy, and subsequently edit and redeploy the updated policy.

### Example: Create a switch\_freeform policy

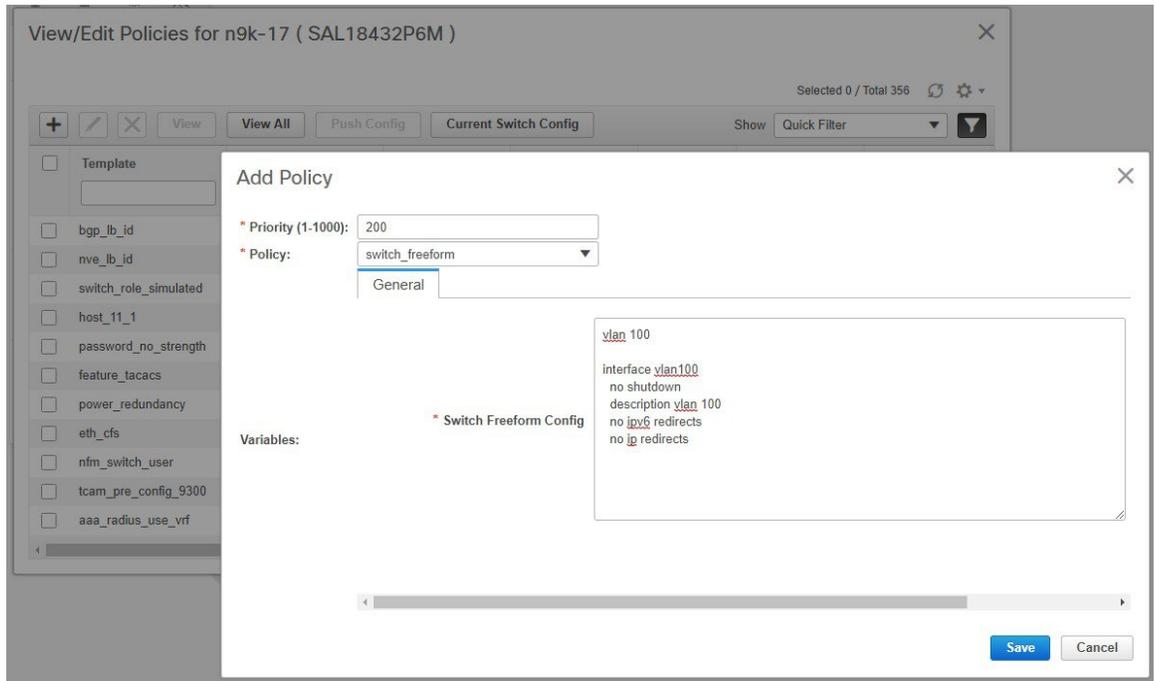
To create a **switch\_freeform** policy, perform the following steps:

#### Procedure

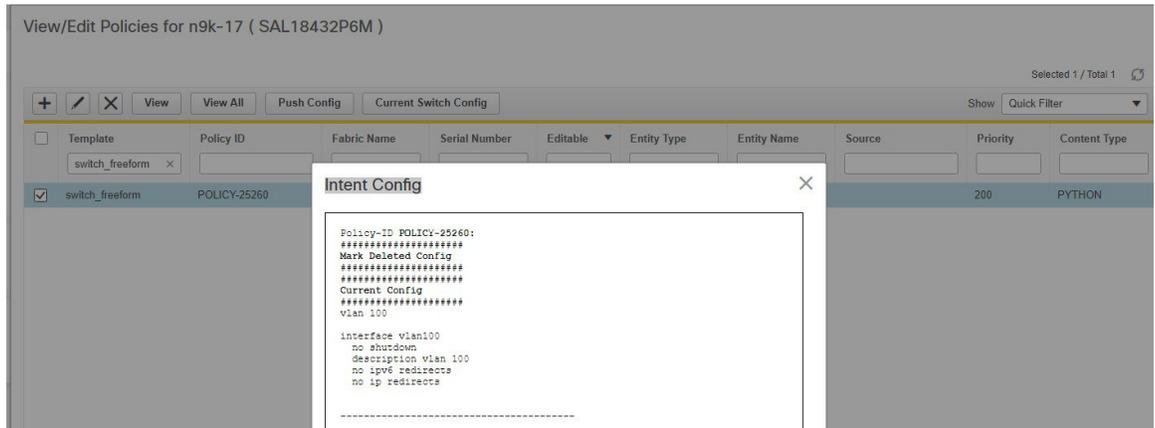
---

- Step 1** Select the **switch\_freeform** template from the policy list in the **Add Policy** screen.  
Set the priority and switch freeform config. Save the policy.

Example: Create a switch\_freeform policy



**Step 2** View the intent config of the **switch\_freeform** policy.



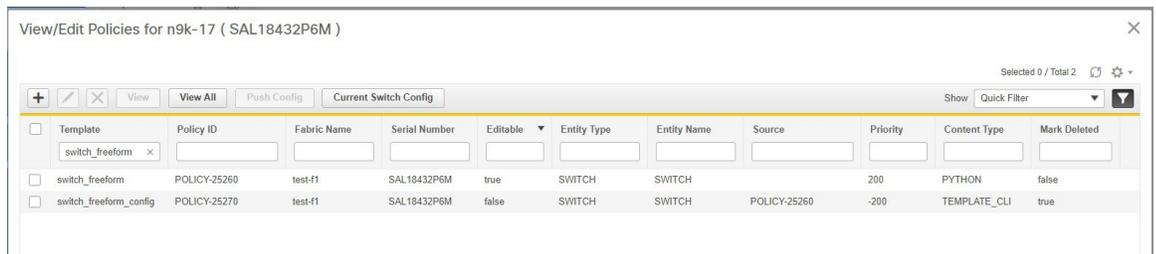
**Step 3** Deploy the switch\_freeform policy from Fabric Builder.

**Step 4** Edit the switch\_freeform policy from the View/Edit Policies window.  
Change the config.

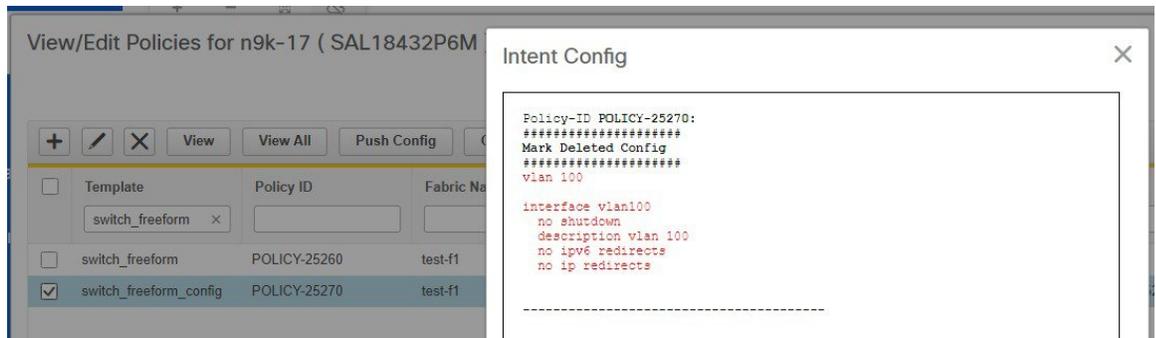


**Step 5** Save the change.

As shown below, the previously created internal **switch\_freeform\_config** policy has its priority changed to a negative number (-200), and the **Mark Deleted** flag is set to true. However, by design, the newly created internal **switch\_freeform\_config** policy is NOT shown.

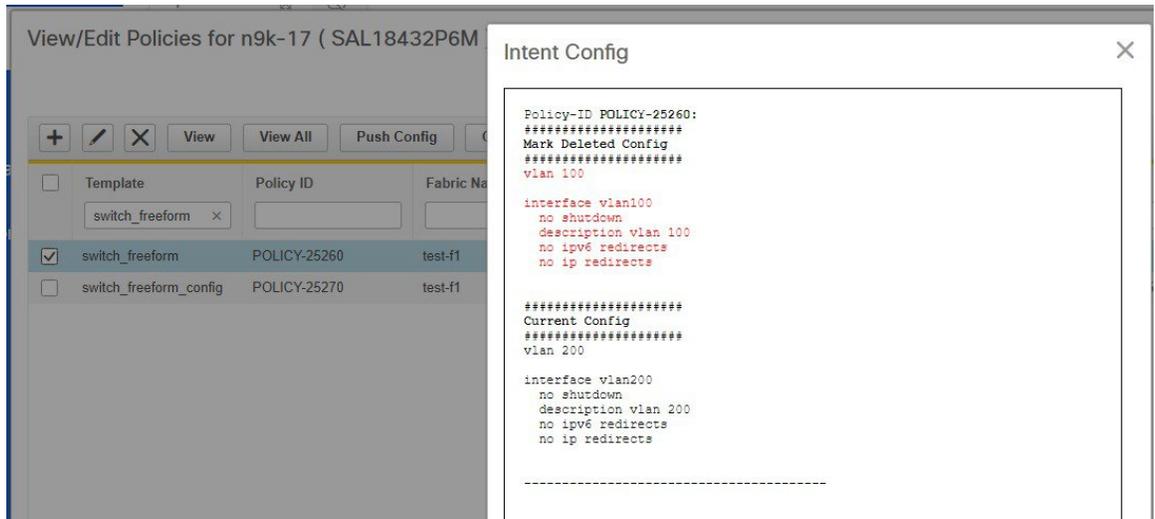


**Step 6** View the intent config of the **mark deleted** internal policy.

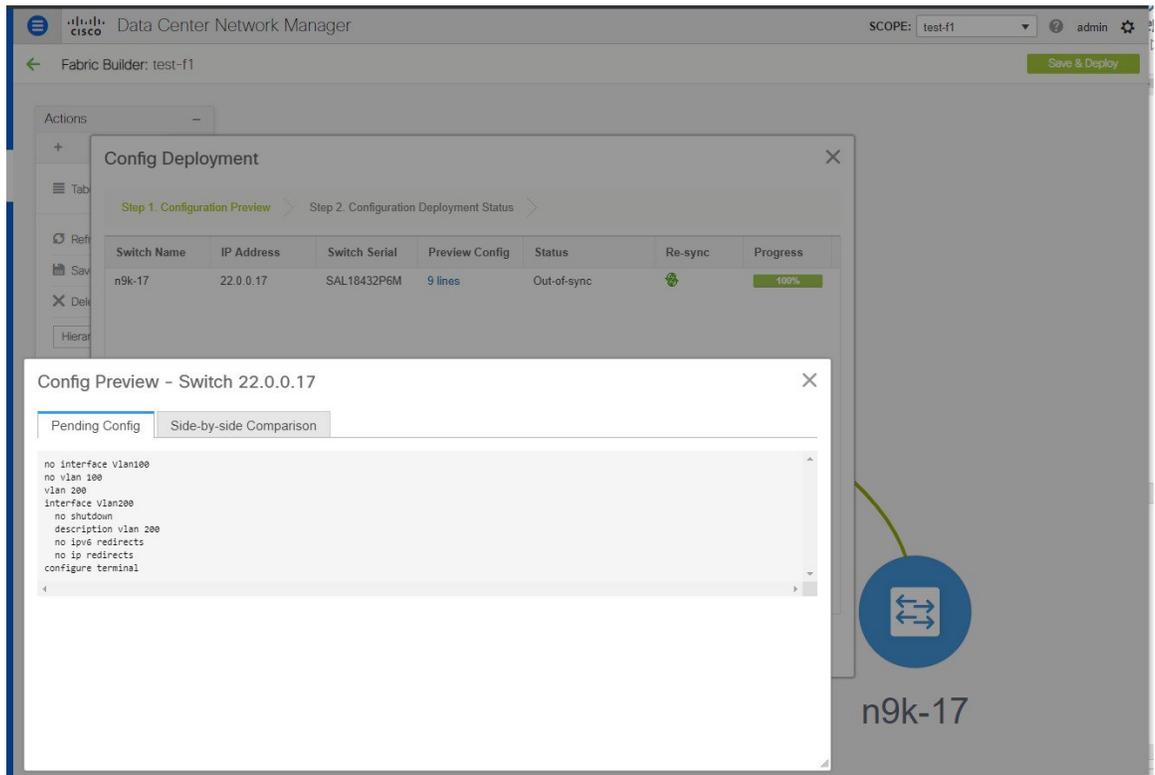


**Step 7** View the intent config of the **changed** switch\_freeform policy before deployment. Note that both the **mark-deleted** and **current configs** are shown.

Changing the Contents of a Template in Use



**Step 8** Deploy the changed config from Fabric Builder.



## Changing the Contents of a Template in Use

A template in general, whether it is a policy, fabric or profile template, cannot be modified once it has been instantiated. However, there could be cases where you want to edit the content of a template, like fixing a bug

in the template or changing an already deployed config. This can be achieved by toggling the `template.in_use.check` option in the **Administration > Server Properties** tab.

## Procedure

- Step 1** Change the `template.in_use.check` from **true (default)** to **false**.
- Step 2** Click 'Apply Changes' at the upper righthand corner.
- A warning will be popped up indicating that a restart of DCNM is needed.
- Ignore this warning as no restart is needed for the `in_use` flag to take effect.
- Step 3** Edit the desired template(s).
- Step 4** Go to the Fabric Builder page and click 'Save & Deploy' for the entire fabric.
- This will regenerate PTIs and the updated content will be picked up and used for the expected configuration (or intent).
- Step 5** Once the contents are re-generated and deployed, change the `template.in_use.check` back to **true** to avoid performance issues.

The screenshot shows the Cisco Data Center Network Manager (DCNM) Administration / DCNM Server / Server Properties page. The page displays various configuration fields for SAN Telemetry. The 'template.in\_use.check' field is highlighted in red and set to 'false'. A warning dialog box is displayed over the page, stating: "Please restart DCNM SAN service if you update properties other than EMC Callhome properties(server.callhome.enable, server.callhome.xmlDir), Event Forwarding properties (server.forward.event.enable), Template properties (template.in\_use.check property), Event Registration properties(sylog.disable) or fabric.enableNpVDiscovery properties. (Note: please restart all instances if federation is deployed). Please resync vmm if you updated vmm.resync.timer". The dialog has an "OK" button.





## CHAPTER 11

# 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 shows you how to extend a VRF to an external fabric. In DNCM, configuration parameters are enhanced as follows:

*Configuration methods* - You can configure VRF Lite through automatic configuration and through the DCNM GUI.

*Supported destination devices* - You can extend VRFs from a VXLAN fabric to Cisco Nexus and non-Nexus devices. A connected non-Cisco device can also be represented in the topology.

- [Prerequisites, on page 515](#)
- [Sample Scenarios, on page 518](#)
- [VRF Lite Through the DCNM GUI – From a BGW Device to a Nexus 7000 Series Edge Router , on page 519](#)
- [VRF Lite Through the DCNM GUI – From a BGW Device To a Non-Nexus Device , on page 531](#)
- [Automatic VRF Lite \(IFC\) Configuration, on page 538](#)
- [Deleting VRF Lite IFCs, on page 541](#)
- [Additional References, on page 543](#)
- [Appendix , on page 543](#)

## Prerequisites

### 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 Configuration Guide, Release 11.2(1).
- Ensure that the role of the designated border device is Border, Border Spine, Border Gateway, or Border Gateway Spine (a switch on which Multi-Site and VRF Lite functions co-exist). 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 role is inappropriate for a border device, set the appropriate role.
- Create an external fabric. If you connect the VLXAN fabric border device to a Nexus 7000 Series switch (or other Nexus device) for external connectivity, add the Nexus 7000 series switch to the external fabric and set its role to **Edge Router**. In DCNM, you can import switches to an external fabric, and update selected configurations. For details, refer the Creating an External Fabric section in the Control chapter.
- To allow inter-subnet communication between end hosts in different VXLAN fabrics, where the subnets are present in both fabrics, you must disable the **Advertise Default Route** feature for the associated VRF. This will result in /32 routes for hosts being seen in both fabrics. For example, Host1 (VNI 30000, VRF 50001) in Fabric1 can send traffic to Host2 (VNI 30001, VRF 50001) in Fabric2 only if the host route is present in both fabrics. When a subnet is present in only one fabric, then default route is sufficient for inter-subnet communication. Steps:
  1. Go to the fabric's **VRFs** screen and select the VRF.
  2. Click the **Edit** option at the top left part of the screen.
  3. In the **Edit VRF** screen, click **Advanced** in the VRF Profile section.
  4. Clear the **Advertise Default Route** checkbox and click **Save**.

▼ VRF Information

\* VRF ID

\* VRF Name

\* VRF Template

\* VRF Extension Template

VLAN ID  Propose VLAN ?

---

▼ VRF Profile

General

Advanced

RP Loopback ID  ? 0-1023

Underlay Mcast Add...  ? IPv4 Multicast Address

Overlay Mcast Groups  ? 224.0.0.0/4 to 239.255.255.255/4

Enable IPv6 link-loc...  ? Enables IPv6 link-local Option under VRF SVI

Enable TRM BGW MSite  ? Enable TRM on Border Gateway Multisite

Advertise Host Routes  ? Flag to Control Advertisement of /32 and /128 Routes to Edge Routers

Advertise Default Route  ? Flag to Control Advertisement of Default Route Internally

Config Static 0/0 Route  ? Flag to Control Static Default Route Configuration

Save
Cancel

The following options apply only when VRF Lite connectivity is enabled on the border devices. By default, following Cisco best practices, DCNM uses eBGP over sub-interfaces for VRF Lite, Option-A peering. In other words, for each VRF Lite Inter-fabric connection (IFC), there is a per VRF per peer eBGP peering session established over IPv4/IPv6 respectively from the border device to the edge/WAN router. As applicable to this VRF Lite peering, there are 3 fields:

- **Advertise Host Routes** – By default, over the VRF Lite peering session, only non-host (/32 or /128) prefixes are advertised. But if host routes (/32 or /128) need to be enabled and advertised from the border device to the edge/WAN router, then the “**Advertise Host Routes**” check box can be enabled. Route-map does outbound filtering. By default, this check box is disabled.
- **Advertise Default Route** – This field controls whether a network statement 0/0 will be enabled under the vrf. This in turn will advertise a 0/0 route in BGP. By default, this field is enabled. When the check box is enabled, this will ensure that a 0/0 route is advertised inside the fabric over EVPN Route-type 5 to the leafs thereby providing a default route out of the leafs toward the border devices.
- **Config Static 0/0 Route** –The field controls whether a static 0/0 route to the edge/WAN router, should be configured under the VRF, on the border device. By default, this field is enabled. If WAN/edge routers are advertising a default route over the VRF Lite peering, to the border device in the fabric, then this field should be disabled. In addition, the “Advertise Default Route” field should also be disabled. This is because the 0/0 route advertised over eBGP will be sent over EVPN to the leafs without the need for any additional configuration. The clean iBGP EVPN separation inside the fabric with eBGP for external out-of-fabric peering, provides for this desired behavior.

Note that all of the options listed are per fabric fields. Hence, in Multi-Site deployments with MSD, these fields can be controlled at a per member fabric level.

5. Follow this procedure for all VRFs deployed on the VXLAN fabrics' border devices connected through VRF Lite.



---

**Note** If you create a new VRF, ensure that you clear the **Advertise Default Route** checkbox.

---



---

**Note** For an explanation on the VRF Lite feature, see the [Cisco Programmable Fabric with VXLAN BGP EVPN Configuration Guide](#) document.

---

## Sample Scenarios

Scenarios explained in this document:

- VRF Lite through the DCNM GUI – From a BGW device to a Nexus 7000 Series edge router.
- VRF Lite through the DCNM GUI – From a BGW device to a non-Nexus device.
- Automatic VRF Lite (IFC) Configuration.



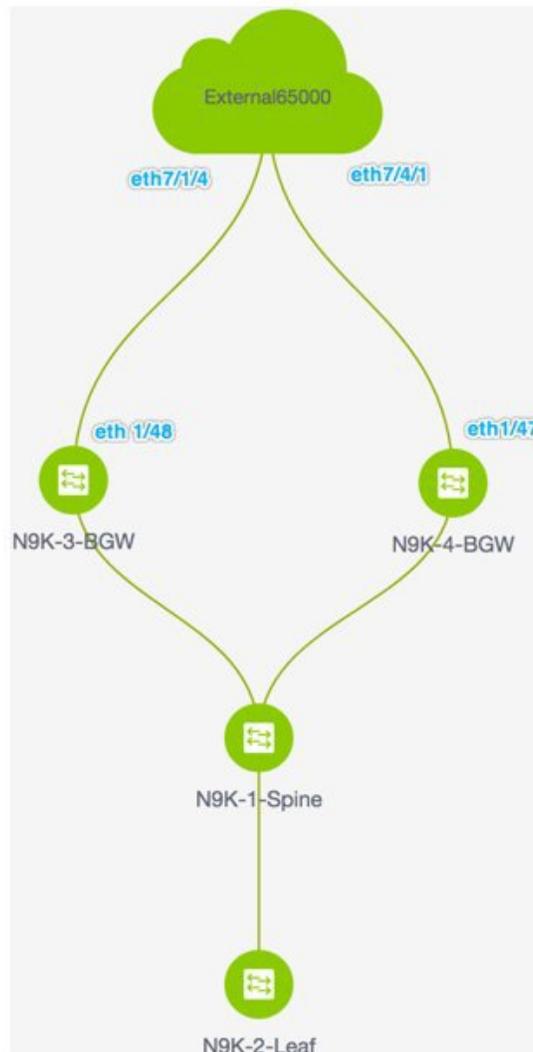
---

**Note**

- The sample scenarios are shown using a Border Gateway role but are equally applicable to the Border nodes as well.
- Anything that applies to Border or Border Gateway roles also applies to Border Spine and Border Gateway Spine roles.

---

## VRF Lite Through the DCNM GUI – From a BGW Device to a Nexus 7000 Series Edge Router



- The topology displays the VXLAN BGP EVPN fabric **Easy7200** connected to the external fabric **External65000** (the cloud icon). The BGWs of the VXLAN fabric are connected to the edge router **n7k1-Edge1** (not visible in the image) in the external fabric.
- The BGWs 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. Network configurations for the VXLAN fabric are 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. This document shows you how to enable Layer 3 north-south traffic between the VXLAN fabric and the external fabric.

- Before VRF Lite configuration, end hosts associated with a specific VRF can send traffic to each other, but only within the fabric. After VRF Lite configuration, end hosts can send traffic outside the VXLAN fabric, towards other (VXLAN or classic LAN) fabrics

### Enabling the VRF Lite feature

For this example, we will enable connectivity between Easy7200 and External65000. The steps:

**Step 1 - Deploy IFC prototypes on physical interfaces, on N9K-3-BGW and N9K-4-BGW.**

**Step 2 - Deploy the individual VRF extensions on the BGWs N9K-3-BGW and N9K-4-BGW.**

**Step 3 - Deploy VRF extensions on the edge router n7k1-Edge1.**

The third step completes the configuration between **Easy7200** and **External65000**.

### Step 1 – Deploying IFC prototypes on physical interfaces on N9K-3-BGW and N9K-4-BGW

For VRF Lite configuration, you should enable eBGP peering between the fabric’s BGW interfaces and the edge router’s interfaces, through point-to-point connections. The BGW physical interfaces are:

- **eth 1/48** on **N9K-3-BGW**, towards **eth 7/1/4** on **n7k1-Edge1**.
- **eth 1/47** on **N9K-4-BGW**, towards **eth 7/4/1** on **n7k1-Edge1**.




---

**Note** You can also enable VRF Lite in a back-to-back topology wherein Border/Border Gateways are directly connected to each other.

---

1. Click **Control > Fabric Builder**. The Fabric Builder screen comes up.
2. Click the **Easy7200** box. The fabric topology comes up.
3. Click **Tabular view**. The **Switches | Links** screen comes up.

The **Links** tab lists fabric links. Each row either represents a link between two devices within **Easy7200** or a link from a device in **Easy7200** to an external fabric.




---

**Note** An inter-fabric link is a physical connection between two Ethernet interfaces or a virtual connection (such as a fabric overlay between two loopback interfaces). When you add a physical connection between devices, the new link appears in the **Links** tab by default.

---

4. Select the link checkbox (that represents the connection between **eth 1/48** on **N9K-3-BGW**, towards **eth 7/1/4** on **n7k1-Edge1**) and click the Edit icon at the top left part of the screen.

| Scope                    | Name   | Policy                         | Info         | Admin State | Oper State |
|--------------------------|--|--------------------------------|--------------|-------------|------------|
| Easy7200                 | N9K-2-Leaf-Ethernet1/47---N9K-1-Spine-Ethernet1/47 | int_intra_fabric_num_link_11_1 | Link Present | Up:Up       | Up:Up      |
| Easy7200<->External65000 | N9K-3-BGW-Ethernet1/48---n7k1-Edge1-Ethernet7/1/4  | int_intra_fabric_num_link_11_1 | Link Present | Up:Up       | Up:Up      |
| Easy7200                 | N9K-3-BGW-Ethernet1/47---N9K-1-Spine-Ethernet1/43  | int_intra_fabric_num_link_11_1 | Link Present | Up:Up       | Up:Up      |
| Easy7200<->External65000 | N9K-4-BGW-Ethernet1/47---n7k1-Edge1-Ethernet7/4/1  | Link Present                   | Up:Up        | Up:Up       | Up:Up      |
| Easy7200<->Easy60000     | N9K-4-BGW-Ethernet1/2---N9K-15-BGW-Ethernet1/8     | Link Present                   | Up:Up        | Up:Up       | Up:Up      |
| Easy7200                 | N9K-4-BGW-Ethernet1/48---N9K-1-Spine-Ethernet1/42  | int_intra_fabric_num_link_11_1 | Link Present | Up:Up       | Up:Up      |

The fields are:

**Scope** – The source and destination fabrics are displayed. For an intra-fabric link, only one fabric name (**Easy7200**) is displayed since the source and destination interfaces are part of the same fabric. An inter-fabric link is displayed as **Easy7200 <->External65000**.

**Name** – The name is formed with the following syntax:

*source device ~ source interface --- destination device ~ destination interface.*

So, the entry is **N9K-4-BGW ~ Ethernet1/47 --- n7k1-Edge1 ~ Ethernet7/4/1**.

**Policy** – The policy used for creating VRF Lite, `ext_fabric_setup_11_1` is displayed.

**Info** – This displays the status of the link (Link Present, Neighbor Present, Neighbor Missing, etc).

**Admin State** – This displays the administrative state of the link (Up, Down, etc).

**Oper State** – This displays the operational state of the link (Up, Down, etc).

The **Link Management – Edit Link** comes up.

#### Link Management - Edit Link

|                         |                       |
|-------------------------|-----------------------|
| * Link Type             | Inter-Fabric          |
| * Link Sub-Type         | VRF_LITE              |
| * Link Template         | ext_fabric_setup_11_1 |
| * Source Fabric         | Easy7200              |
| * Destination Fabric    | External65000         |
| * Source Device         | n9k-2-leaf            |
| * Source Interface      | Ethernet2/1           |
| * Destination Device    | n9k-18-RS             |
| * Destination Interface | Ethernet1/2           |

#### Link Profile

- General
- Advanced

|                    |            |  |
|--------------------|------------|--|
| * BGP Local ASN    | 7200       | ? Local BGP Autonomous System Number       |
| * IP Address/Mask  | 2.2.2.2/24 | ? IP address for sub-interface in each VRF |
| * BGP Neighbor IP  | 2.2.2.1    | ? Neighbor IP address in each VRF          |
| * BGP Neighbor ASN | 65000      | ? Neighbor BGP Autonomous System Number    |

Some fields are explained:

**Link Sub-Type** - By default, the **VRF\_LITE** option is displayed.

**Link Template** – The default template for a VRF Lite IFC, **ext\_fabric\_setup\_11\_1**, is displayed. The template enables the source and destination interfaces as Layer 3 interfaces, configures the **no shutdown** command, and sets their MTU to 9216.

You can edit the **ext\_fabric\_setup\_11\_1** template or create a new one with custom configurations.

In the **General** tab, the BGP AS numbers of **Easy7200** and **External65000** are displayed. Fill in the other fields as explained.

The screenshot shows the 'Link Profile' configuration page with the 'General' tab active. The following fields are visible:

- BGP Local ASN:** 7200 (Local BGP Autonomous System Number)
- IP Address/Mask:** 2.2.2.2/24 (IP address for sub-interface in each VRF)
- BGP Neighbor IP:** 2.2.2.1 (Neighbor IP address in each VRF)
- BGP Neighbor ASN:** 65000 (Neighbor BGP Autonomous System Number)

**IP Address/Mask** – Enter the IP address prefix to assign an IP address for the **Ethernet 1/48** sub interfaces, the source interface of the IFC. A sub-interface is created for each VRF extended over this IFC, and a unique 802.1Q ID is assigned to it. The IP address/Mask entered here, along with the BGP Neighbor IP field (explained below) will be used as the default values for the sub-interface created at VRF extension and can be overwritten.

For example, an 802.1Q ID of 2 is associated with subinterface Eth 1/48.2 for VRF 50000 traffic, and 802.1Q ID of 3 is associated with Eth 1/48.3 and VRF 50001, and so on.

(The VRF extension deployment is explained in a subsequent section).

The IP prefix is reserved with the DCNM resource manager. Ensure that you use a unique IP address prefix for each IFC you create in the topology.

**BGP Neighbor IP** – Enter the IP address of the eBGP neighbor for each VRF extension deployed on this IFC, on the **N9K-3\_BGW** end.

Inter-fabric traffic from VRFs for an IFC will have the same source IP address (**2.2.2.2/24**) and destination IP address (**2.2.2.1**).

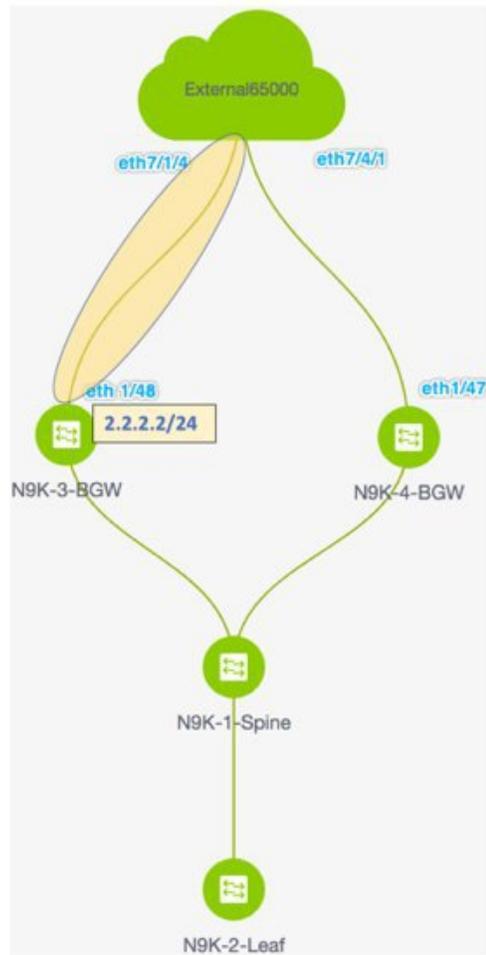
The **Advanced** tab has been added in the **Link Profile** section.

This tab contains the following fields:

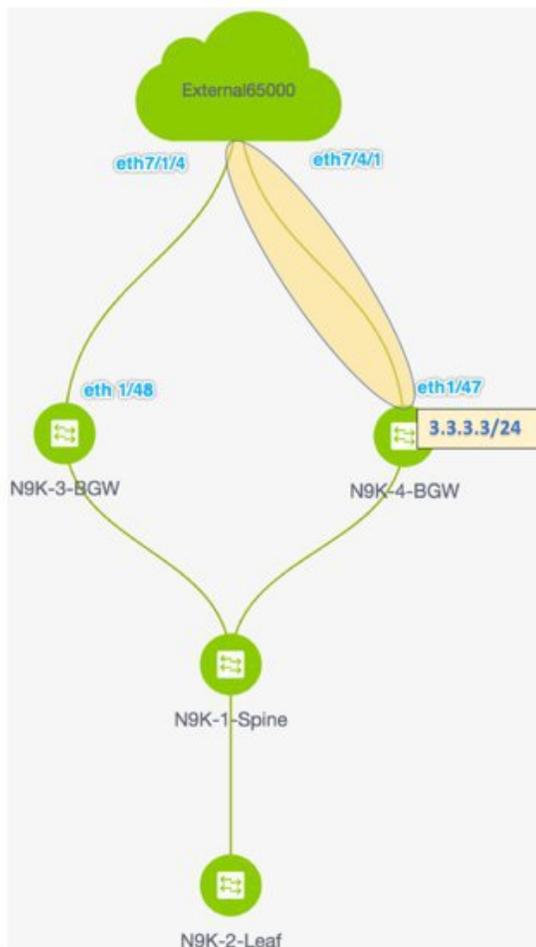
- **Source Interface Description**
- **Destination Interface Description**
- **Source Interface Freeform Config**
- **Destination Interface Freeform Config**

5. Click **Save** at the bottom right part of the screen.

The **Switches|Links** screen comes up again. You can see that the IFC entry is updated with the VRF Lite policy template used for creating the IFC, **ext\_fabric\_setup\_11\_1**. A representation of the topology is shown below.



6. Similarly, create an IFC from **eth 1/47** on **N9K-4-BGW** towards **eth 7/4/1** on **n7k1-Edge1**. An entry is seen in the **Links** screen. A representation of the topology is shown below.



7. Click **Save and Deploy** at the top right part of the screen.

The **Links** tab after executing **Save and Deploy** looks like this. The links on which IFC has deployed have the relevant policy configured in the **Policy** column.

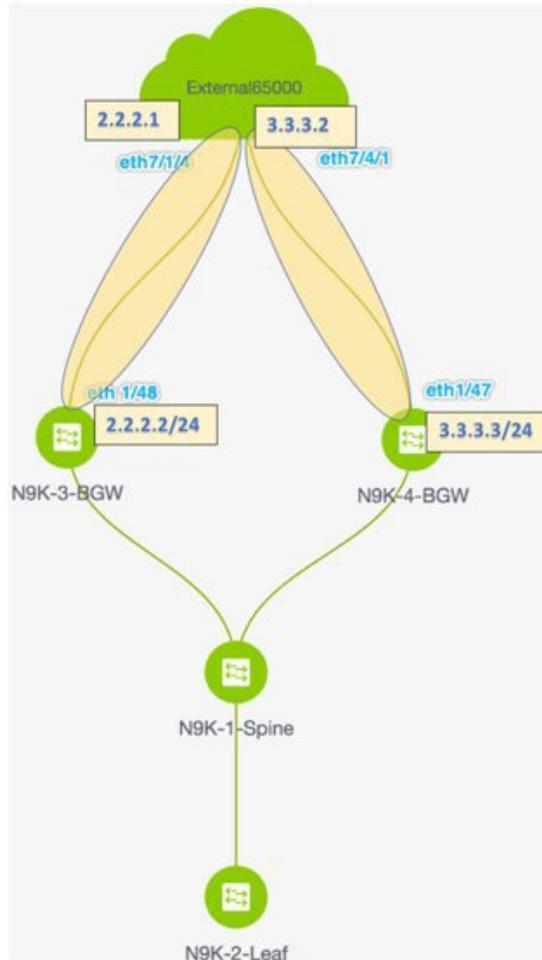
| Scope | Name                    | Policy  | Info                           | Admin State  | Oper St |
|-------|-------------------------|---|--------------------------------|--------------|---------|
| 1     | Easy7200->External65000 | N9K-3-BGW-Ethernet1/48--n7k1-Edge1-Ethernet7/1/4  | ext_fabric_setup_11_1          | Link Present | Up:Up   |
| 2     | Easy7200->External65000 | N9K-4-BGW-Ethernet1/47--n7k1-Edge1-Ethernet7/4/1  | ext_fabric_setup_11_1          | Link Present | Up:Up   |
| 3     | Easy7200                | N9K-3-BGW-Ethernet1/47--N9K-1-Spine-Ethernet1/43  | int_intra_fabric_num_link_11_1 | Link Present | Up:Up   |
| 4     | Easy7200                | N9K-4-BGW-Ethernet1/48--N9K-1-Spine-Ethernet1/42  | int_intra_fabric_num_link_11_1 | Link Present | Up:Up   |
| 5     | Easy7200                | N9K-2-Leaf-Ethernet1/47--N9K-1-Spine-Ethernet1/47 | int_intra_fabric_num_link_11_1 | Link Present | Up:Up   |

8. Go to the **Scope** drop down box at the top right part of the screen and choose **External65000**. The external fabric **Links** screen is displayed. You can see that the two IFCs created from **Easy7200** to **External65000** is displayed here.



**Note** When you create an IFC or edit its setting in the VXLAN fabric, the corresponding entry is automatically created in the connected external fabric.

9. Click **Save and Deploy** to save the IFCs creation on **External65000**.



**Base configurations** – For VRF Lite to function, appropriate route maps and policies that apply to VRFs have to be deployed on the border devices **N9K-3-BGW** and **N9K-4-BGW**. You do not need to manually enable the base configurations. They are automatically deployed via a default template **ext\_base\_border\_vrflite\_11\_1**.

For a device with a Border Leaf or Border Spine role, the base configurations are deployed when you execute the **Save and Deploy** operation (available in the fabric topology screen [via the **Fabric Builder** screen > Fabric Box]) for the first time in a fabric.

For a Border Gateway or Border Gateway Spine role, the base configurations are deployed when you deploy the first VRF Lite IFC on the device.

You need to modify the **ext\_base\_border\_vrflite\_11\_1** template for specific needs before deployment or its policy should be deleted, template modified, and then deploy the template again. The configurations are noted in the **Appendix** section.

The first step in the VRF Lite configuration scenario, creating IFCs on the border devices and edge router, is complete. Next, the VRF extensions are deployed on the switches.

**Step 1** - Deploy IFC prototypes on physical interfaces, on **N9K-3-BGW** and **N9K-4-BGW**.

**Step 2** - Deploy the individual VRF extensions on the BGWs **N9K-3-BGW** and **N9K-4-BGW**.

**Step 3** - Deploy VRF extensions on the edge router **n7k1-Edge1**.

The third step completes the configuration between **Easy7200** and **External65000**.

**Step 2** - Deploy the individual VRF extensions on the BGWs **N9K-3-BGW** and **N9K-4-BGW**

During the IFC creation process, base configurations are created, and IP addresses are reserved for the interfaces that transport the inter-fabric traffic on **N9K-3-BGW** and **N9K-4-BGW**. In this step, the VRF and VRF extension configuration is deployed on the interfaces.

To extend VRFs beyond the fabric, the VRFs should have been created and deployed on relevant fabric devices, except the border devices.

The steps are:

1. Click **Control > Networks and VRFs**. The **Networks & VRFs** screen comes up.
2. Click **Continue**. The **Select a Fabric** screen comes up.
3. Select **Easy7200** and click **Continue** at the top right part of the screen.



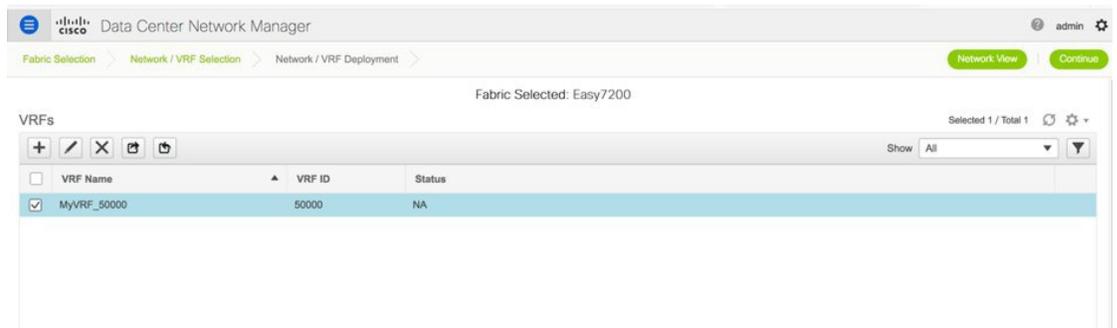
### Select a Fabric

Choose a fabric with appropriate switches where you want the Top Down functionality to be enabled

Easy7200

The **Networks** screen comes up.

4. Click **VRFs** at the top right part of the screen. The **VRFs** screen comes up.
5. Select the VRF that you want to deploy (**MyVRF\_5000** in this case) and click **Continue** at the top right part of the screen.



The **Easy7200** fabric topology comes up.

6. Select the **Multi-Select** checkbox at the top right part of the screen and drag the cursor across the BGWs on which you want to deploy the VRF and VRF extension configuration.



The **VRF Extension Attachment** screen comes up. Each row represents a switch and each tab a VRF. Update settings for each tab as explained.

VRF Extension Attachment - Attach extensions for given switch(es)

Fabric Name: Easy7200

Deployment Options

Select the row and click on the cell to edit and save changes

MyVRF\_50000

| <input type="checkbox"/> | Switch    | VLAN | Extend | CLI Freeform    | Status |
|--------------------------|-----------|------|--------|-----------------|--------|
| <input type="checkbox"/> | N9K-3-BGW | 2000 | NONE   | Freeform config | NA     |
| <input type="checkbox"/> | N9K-4-BGW | 2000 | NONE   | Freeform config | NA     |

[Save](#)

In the **Extend** column, click on **NONE** and choose the **VRF\_LITE** option from the drop down box. Do this for the second row too.

Select the checkboxes in both rows.

The **Extension Details** section comes up at the bottom of the screen. It displays the IFCs created on the selected switches, wherein each row represents an IFC.

Select the IFC check boxes in both rows.

After selecting the IFCs, the screen looks like this.

VRF Extension Attachment - Attach extensions for given switch(es)

Fabric Name: Easy7200

Deployment Options

Select the row and click on the cell to edit and save changes

MyVRF\_50000

| <input checked="" type="checkbox"/> | Switch    | VLAN | Extend   | CLI Freeform    | Status |
|-------------------------------------|-----------|------|----------|-----------------|--------|
| <input checked="" type="checkbox"/> | N9K-3-BGW | 2000 | VRF_LITE | Freeform config | NA     |
| <input checked="" type="checkbox"/> | N9K-4-BGW | 2000 | VRF_LITE | Freeform config | NA     |

Extension Details

| <input checked="" type="checkbox"/> | Source Switch | Type     | IF_NAME      | Dest. Switch | Dest. Interface | DOT1Q_ID | IP_MASK  | NEIGHBOR_IP | NEIGHBOR_ASN | IPV6_MASK |
|-------------------------------------|---------------|----------|--------------|--------------|-----------------|----------|----------|-------------|--------------|-----------|
| <input checked="" type="checkbox"/> | N9K-3-BGW     | VRF_LITE | Ethernet1/48 | Edge1        | Ethernet7/1/4   | 2        | 2.2.2/24 | 2.2.2.1     | 65000        |           |
| <input checked="" type="checkbox"/> | N9K-4-BGW     | VRF_LITE | Ethernet1/47 | Edge1        | Ethernet7/4/1   | 2        | 3.3.3/24 | 3.3.3.1     | 65000        |           |

Click **Save** at the bottom right part of the screen.

The fabric topology screen comes up.

7. Click the **Preview** option at the top right part of the screen to preview VRF and VRF extension configuration.
8. Click **Deploy** at the top right part of the screen.

At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly (Blue for Pending state, yellow for In Progress state when the provisioning is in progress, red for failure state, green when successfully deployed).

When the switch icons turn green, it means that the VRFs are successfully deployed.

The second step in the VRF Lite configuration scenario, deploying VRF extensions on the border devices is complete. Next, the VRF extensions are deployed on the edge router **n7k1-Edge1**.

**Step 1** - Deploy IFC prototypes on physical interfaces, on **N9K-3-BGW** and **N9K-4-BGW**.

**Step 2** - Deploy the individual VRF extensions on the BGWs **N9K-3-BGW** and **N9K-4-BGW**.

**Step 3 - Deploy VRF extensions on the edge router n7k1-Edge1.**

The third step completes the configuration between **Easy7200** and **External65000**.

**Step 3 - Deploy VRF extensions on the edge router n7k1-Edge1**

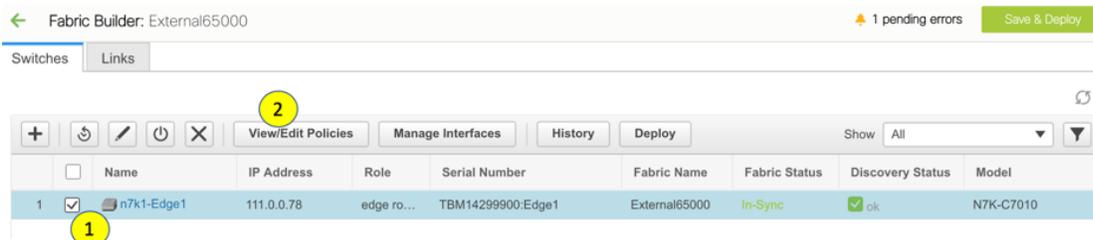
In order to extend VRFs on the edge router, keep a note of the following fields. VRF extension on the border device is on a per interface basis.

- **IP\_MASK** - This will become the neighbor address at the edge router end and mask will be the local mask on the edge router. This is derived from the IFC prototype created in the earlier step.
- **Easy Fabric ASN** - This will become neighbor ASN from the edge router end. This is derived from the IFC prototype created in the earlier step.
- **Dot1Q tag** - This will be same on the edge router. This is derived from the VRF extension table.
- **Neighbor ASN** - This will become LOCAL ASN on the edge router. IFC prototype.
- **Neighbor IP** - This will become Local IP for sub-interface on the edge router. IFC prototype.
- **Destination port** - Will be local port on edge router upon which extension will be deployed.

You have deployed VRF extensions for **MyVRF\_50000** from the BGWs **N9K-3-BGW** and **N9K-4-BGW**. Now, you should deploy the VRF extensions on the other end of the links, on **n7k1-Edge1**. In DCNM, the CLI template used for this is **External\_VRF\_Lite\_eBGP**.

**eBGP configuration on the edge router**

1. In the **External65000** fabric topology screen, click **Tabular view**.  
The **Switches | Links** screen comes up.
2. Select the switch checkbox and click the **View/Edit Policies** button.



The **View/Edit Policies** screen comes up.

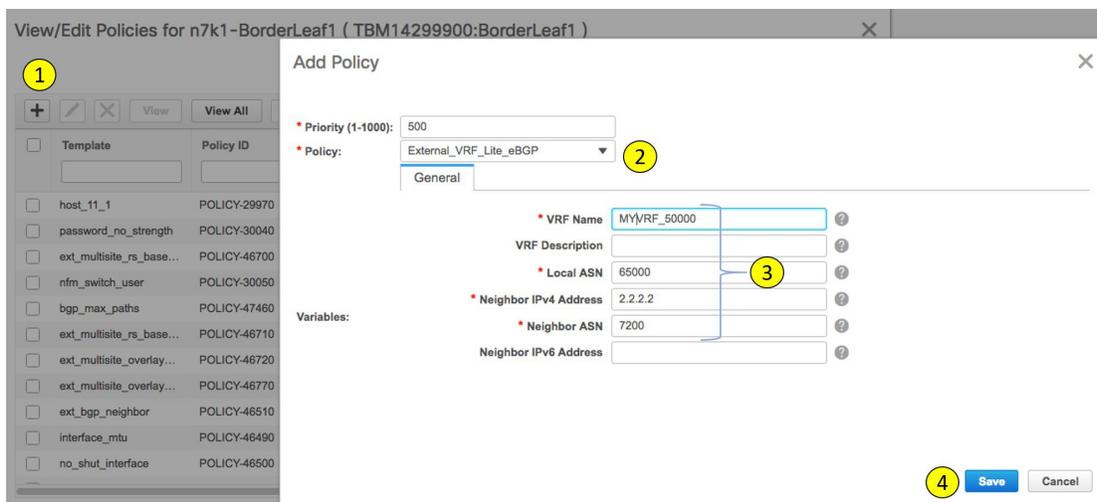
- Click + at the top left part of the screen to add a policy, and fill in the **Add Policy** screen as shown in the image.

You can use a user defined template too in the **Policy** field.



**Note** Note the policy ID for this VRF extension. It is useful when deleting the policy to remove the extension, when applicable.

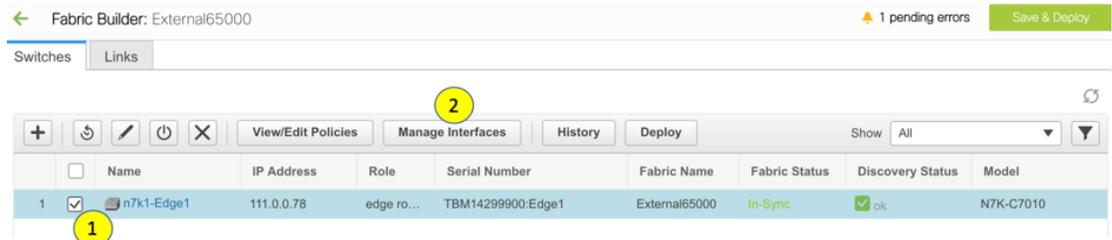
This defines a policy from the edge router towards **N9K-3-BGW**.



- As per the earlier steps, create a policy for the VRF extension towards **N9K-4-BGW**. The **Neighbor IPv4 Address** field for the second extension is updated with 3.3.3.3.

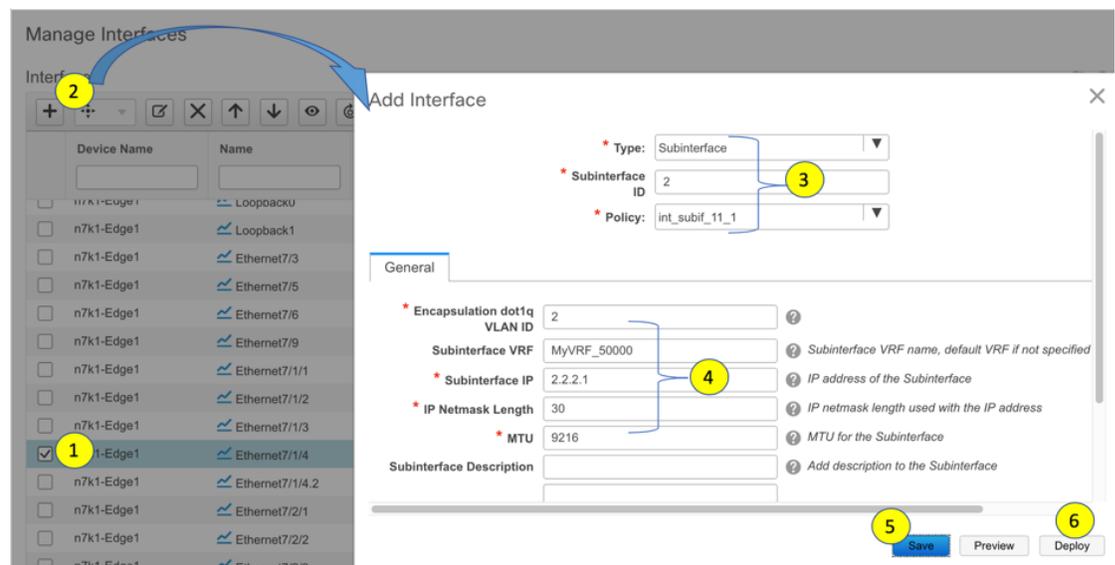
### Sub interface policy on Edge Router

- In the **External65000** fabric topology screen, click **Tabular view**.  
The **Switches | Links** screen comes up.
- Select the switch checkbox and click the **Manage Interfaces** button.



The **Manage Interfaces** screen comes up.

- As shown in the image, select the interface connected to the border device (in this case **Eth7/1/4**), and click + at the top left part of the screen. Then, fill the **Add Interface** screen from corresponding IFC and VRF extensions on the border device.



The example shows a breakout port on the Cisco Nexus 7000 Series switch. This breakout must be performed using the DCNM breakout policy (the template name is **breakout\_interface**). If this is not done, the subinterface deletion is blocked by DCNM.

- Click **Save** to save the settings, and **Deploy** to deploy the settings onto the switch.
- As explained in the earlier steps, create another subinterface policy for the VRF extension towards **N9K-4-BGW**. The **Subinterface IP** field for the second extension is updated with 3.3.3.1.

The third step in the VRF Lite configuration scenario, deploying VRF extensions on the edge router **N7k1-Edge1** is complete. This step completes the configuration between **Easy7200** and **External65000**.

## VRF Lite Through the DCNM GUI – From a BGW Device To a Non-Nexus Device

In this case, the non-Nexus device is an ASR 9000 Series router, **ASR9K-1-Edge** which is connected to the BGW **N9K-3-BGW** in the **Easy7200** fabric. The router is not imported through DCNM nor discovered via

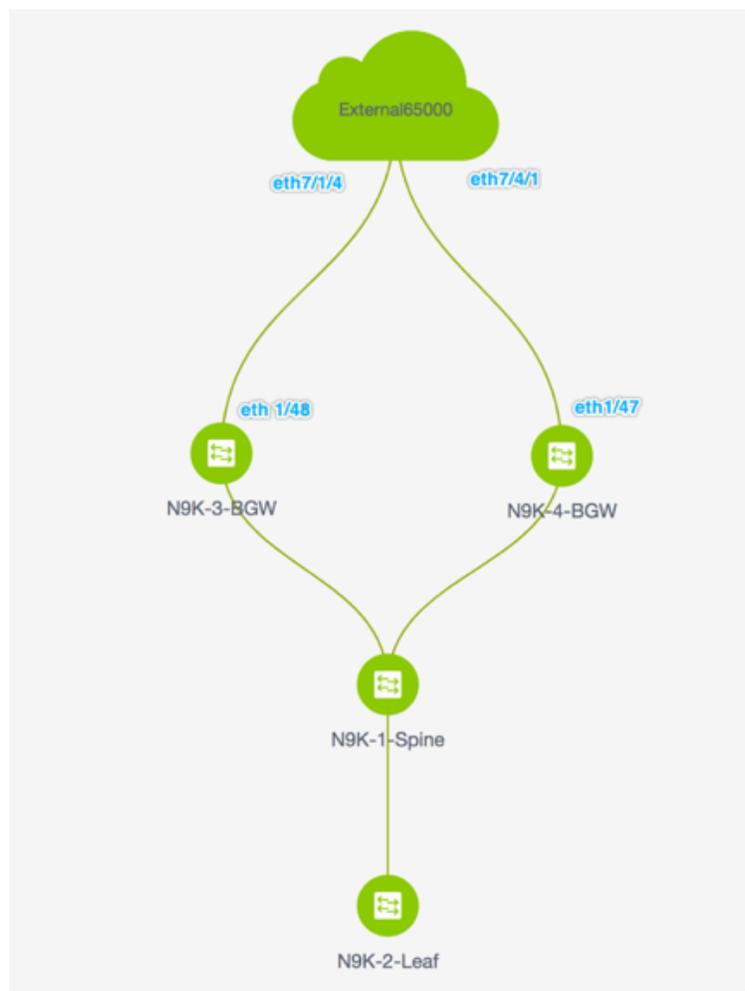
CDP or LLDP. To represent the non-Nexus device, you must create an external fabric. Refer the **Creating an External Fabric** topic to know how to create an external fabric. For this example, the external fabric **External65000** is created.

The device and connection are displayed in the DCNM topology after the IFC creation between **ASR9K-1-Edge** and **N9K-3-BGW**.



**Note** A connected non-Cisco device can also be represented in the topology.

The topology:



The steps are:

**Step 1 - Deploy an IFC prototype on the N9K-3-BGW physical interface that connects to ASR9K-1-Edge.**

**Step 2 - Deploy the individual VRF extensions on N9K-3-BGW.**

This step completes the configuration between **Easy7200** and the non-Nexus device.

**Step 1 - Deploy an IFC prototype on the N9K-3-BGW physical interface that connects to ASR9K-1-Edge**

For VRF Lite configuration, you should enable eBGP peering between the fabric's BGW interface and the **ASR9K-1-Edge** interface, through a point-to-point link.

1. Click **Control > Fabric Builder**. The **Fabric Builder** screen comes up.
2. Click the rectangular box that represents the **Easy7200** fabric. The fabric topology screen comes up.
3. Click **Tabular view**. The **Switches | Links** screen comes up.

The **Links** tab lists fabric links. Each row either represents a link between two devices within **Easy7200** or a link from a device in **Easy7200** to an external fabric.

4. Click + to add a new link. The **Link Management – Add Link** screen comes up.

Link Management - Add Link

\* Link Type: Intra-Fabric

\* Link Sub-Type: Fabric

\* Link Template: int\_intra\_fabric\_num\_link\_11\_1

\* Source Fabric: Easy7200

\* Destination Fabric:

\* Source Device:

\* Source Interface:

\* Destination Device:

\* Destination Interface:

Link Profile

General

\* Source IP: ? IP address of the source interface

\* Destination IP: ? IP address of the destination interface

Interface Admin State:  ? Admin state of the interface

Save

Fill or choose the fields as noted:

**Link Type** – Choose **Inter-Fabric**.

**Link Sub-Type** – **VRF\_Lite** is displayed by default.

**Link Template** - By default, the **ext\_fabric\_setup\_11\_1** template is populated.



**Note** You can add, edit, or delete user-defined templates. See **Template Library** section in the **Control** chapter for more details.

**Source Fabric** - **Easy7200** is selected by default.

**Destination Fabric** – Select **External65000**.

**Source Device** and **Source Interface** - Choose the BGW and the interface that connects to the ASR device.

**Destination Device** and **Destination Interface**— Destination device and interface do not appear in the drop down box. Type any string here that will help identify the device. This name appears in the external fabric topology screen in the **Fabric builder** screen.

General tab in the Link Profile section.

**BGP Local ASN** - In this field, the AS number of the source fabric Easy7200 is autopopulated.

**IP Address/Mask** - Enter the IP address and mask that is used in the VRF Extension Sub-interfaces.

**BGP Neighbor IP** - Enter the IP address that is used on the External box as local interface address for the VRF Extensions.

**BGP Neighbor ASN** - In this field, the AS number of the external fabric External65000 is autopopulated since we selected it as the external fabric.

After filling up the **Add Link** screen, it looks like this:

Link Management - Add Link

\* Link Type: Inter-Fabric  
 \* Link Sub-Type: VRF\_LITE  
 \* Link Template: ext\_fabric\_setup\_11\_1  
 \* Source Fabric: Easy7200  
 \* Destination Fabric: External65000  
 \* Source Device: N9K-3-BGW  
 \* Source Interface: Ethernet1/5  
 \* Destination Device: ASR9K-1-Edge  
 \* Destination Interface: Ethernet1/5

Link Profile  
 General  
 Advanced

\* BGP Local ASN: 7200 ? Local BGP Autonomous System Number  
 \* IP Address/Mask: 5.5.5.2/24 ? IP address for sub-interface in each VRF  
 \* BGP Neighbor IP: 5.5.5.1 ? Neighbor IP address in each VRF  
 \* BGP Neighbor ASN: 65000 ? Neighbor BGP Autonomous System Number

Save

5. Click **Save** at the bottom right part of the screen.

The **Switches|Links** screen comes up again. You can see that the IFC entry is updated.

6. Click **Save and Deploy** at the top right part of the screen.

The links on which the IFC is deployed has the relevant policy (**ext\_fabric\_setup\_11\_1**) configured in the **Policy** column.

7. Go to the **Scope** drop down box at the top right part of the screen and choose **External65000**. The external fabric **Links** screen is displayed. You can see that the IFC created from **Easy7200** to the ASR device is displayed here.

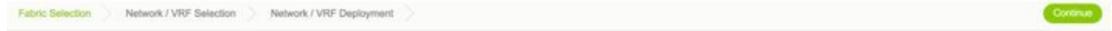
8. Click **Save and Deploy**.

The first step in the VRF Lite configuration scenario from a BGW to a non-Nexus device is complete. Next, the VRF extensions are deployed on the BGW towards the ASR device.

## Step 2 - Deploy the individual VRF extensions on N9K-3-BGW

To extend VRFs beyond the fabric, the VRFs should have been created and deployed on relevant fabric devices, excepting the border devices.

1. Click **Control** > **Networks and VRFs**. The **Networks & VRFs** screen comes up.
2. Click **Continue**. The **Select a Fabric** screen comes up.
3. Select **Easy7200** and click **Continue** at the top right part of the screen.



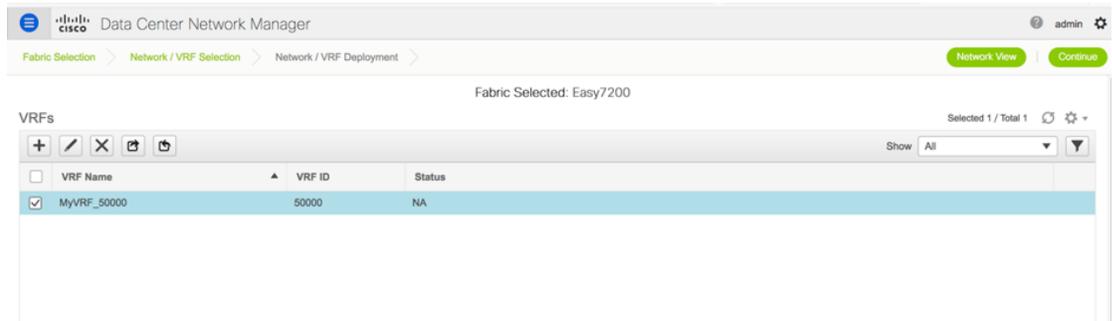
## Select a Fabric

Choose a fabric with appropriate switches where you want the Top Down functionality to be enabled

Easy7200

The **Networks** screen comes up.

4. Click **VRFs** at the top right part of the screen. The **VRFs** screen comes up.
5. Select the VRF that you want to deploy (**MyVRF\_5000** in this case) and click **Continue** at the top right part of the screen.



The Easy7200 fabric topology comes up.

6. Double-click the **N9K-3-BGW** icon on which you want to deploy the VRF and VRF extension configuration.

The **VRF Extension Attachment** screen comes up. Each row represents a switch and each tab a VRF. Only one VRF is extended in this example.

## VRF Extension Attachment - Attach extensions for given switch(es)

Fabric Name: Easy7200

## Deployment Options

① Select the row and click on the cell to edit and save changes

MyVRF\_50000

| <input type="checkbox"/> | Switch ▲  | VLAN | Extend | CLI Freeform    | Status |
|--------------------------|-----------|------|--------|-----------------|--------|
| <input type="checkbox"/> | N9K-3-BGW | 2000 | NONE   | Freeform config | NA     |

[Save](#)

In the **Extend** column, click on **NONE**. A drop down box appears. Choose the **VRF\_LITE** option, and click outside the row.

Select the checkbox next to the switch.

The **Extension Details** section comes up at the bottom of the screen. It displays the IFCs created on the selected switches, wherein each row represents an IFC.

Select the IFC check box. After selecting the IFCs, the screen looks like this.

VRF Extension Attachment - Attach extensions for given switch(es) ✕

Fabric Name: Easy7200

Deployment Options

① Select the row and click on the cell to edit and save changes

| <input type="checkbox"/>            | Switch   | VLAN | Extend   | CLI Freeform    | Loopback Id | Loopback IPv4 Address | Lo |
|-------------------------------------|----------|------|----------|-----------------|-------------|-----------------------|----|
| <input checked="" type="checkbox"/> | N9K-3... | 2000 | VRF_LITE | Freeform config |             |                       |    |

Extension Details

| <input checked="" type="checkbox"/> | Sourc... | type     | IF_NAME      | Dest. Switch | Dest. Interface | DOT1Q_I |
|-------------------------------------|----------|----------|--------------|--------------|-----------------|---------|
| <input checked="" type="checkbox"/> | N9K-3... | VRF_LITE | Ethernet1/48 | Edge1        | Ethernet7/1/4   | 2       |

[Save](#)

Click **Save** at the bottom right part of the screen.

The fabric topology screen comes up.

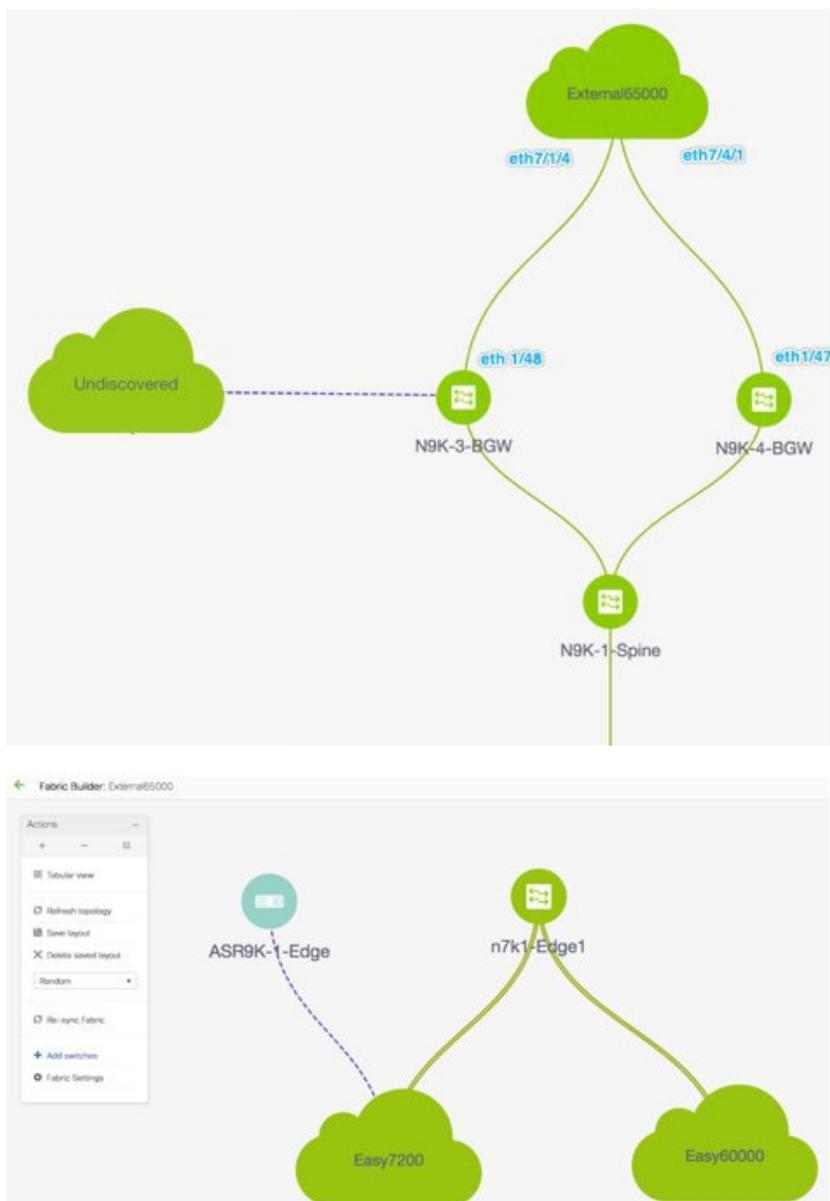
7. Click the **Preview** option at the top right part of the screen to preview VRF and VRF extension configuration.
8. Click **Deploy** at the top right part of the screen.

At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly (Blue for Pending state, yellow for In Progress state when the provisioning is in progress, red for failure state, green when successfully deployed, and so on).

When the switch icons turn green, it means that the VRF is successfully deployed.

The second step in the VRF Lite configuration scenario, deploying VRF extensions on the border device towards the non-Nexus ASR device is complete.

The device and connection will display in the **Easy7200** and **External65000** fabrics.

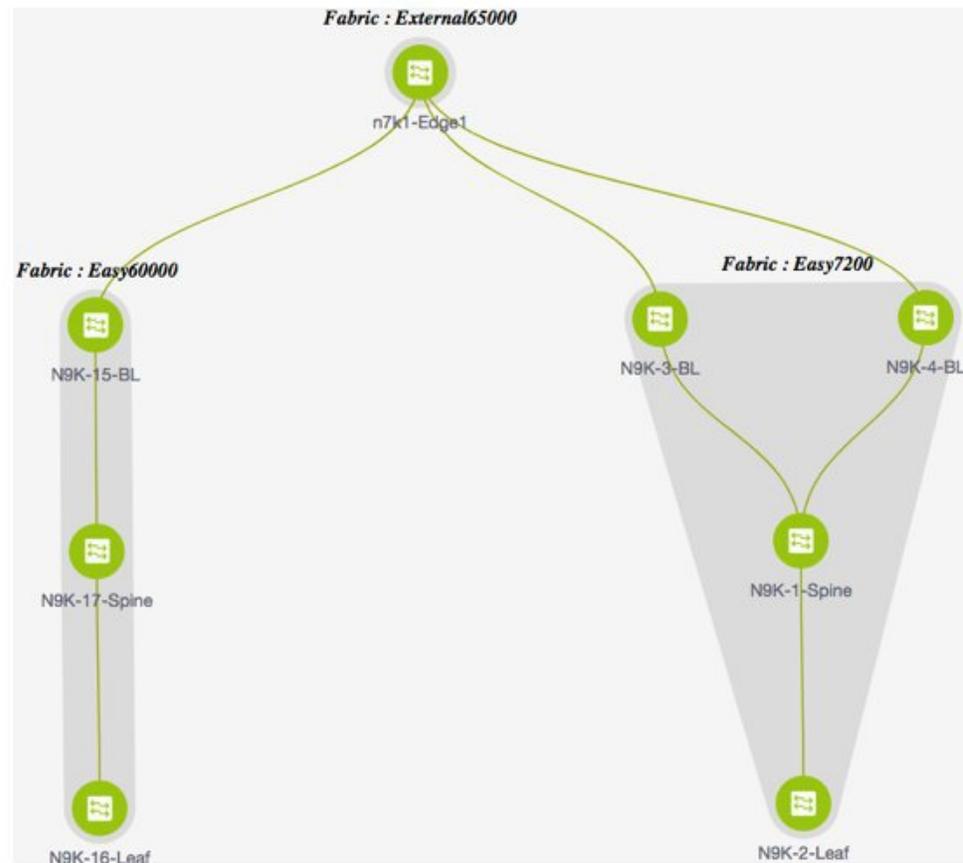


## Automatic VRF Lite (IFC) Configuration

You can enable VRF Lite auto-configuration by changing the fabric settings of the **VRF Lite Deployment** field under the **Resources** tab from **Manual** to any of the auto-configuration settings.



**Note** In the fabric topology screen within **Fabric Builder**, you can view only the individual fabric and the external fabric connected.



- The topology displays VXLAN BGP EVPN fabrics **Easy60000** (at the left) and **Easy7200** (at the right) and external fabric **External65000** (at the top). The border leaf of one VXLAN fabric is connected to the border leaf of the other through the edge router **n7k1-Edge1** in the external fabric.
- The border leafs are special devices that allow clear control and data plane segregation from the fabric to the external Layer 3 domain while allowing for policy enforcement points for any inter-fabric traffic. Multiple border devices in the fabric ensure redundancy in the case of failures and effective load distribution. This document shows you how to enable Layer 3 north-south traffic between the VXLAN fabrics and the external fabric.
- Before VRF Lite configuration, end hosts associated with a specific VRF can send traffic to each other, but only within the fabric. After VRF Lite configuration, end hosts can send traffic across fabrics.
- Network configurations for the VXLAN fabric are provisioned through DCNM.

The template used for VRF Lite IFC auto configuration is **ext\_fabric\_setup\_11\_1**. You can edit the **ext\_fabric\_setup\_11\_1** template or create a new one with custom configurations.

#### Automatic VRF Lite Creation Rules

- Ensure that no user policy is enabled on the interface that connects to the edge router. If a policy exists, then the interface will not be configured.
- Auto configuration is provided for the following cases:

- **Border** role in the VXLAN fabric and **Edge Router** role in the connected external fabric device
- **Border Gateway** role in the VXLAN fabric and **Edge Router** role in the connected external fabric device
- **Border** role to another **Border** role directly

Note that auto configuration is not provided between two BGWs.

If you need a VRF Lite between any other roles, then you have to deploy it manually through the DCNM GUI.

- To deploy configurations in the external fabric, ensure that the **Fabric Monitor Mode** check box is cleared in the external fabric settings of the **External65000** fabric. When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on its switches.

There are four modes available for VRF Lite IFC creation.

1. **Manual** - Use the GUI to deploy the VRF Lite IFCs as shown in the earlier section.
2. **To External Only** - Configure a VRF Lite IFC on each physical interface of a border leaf (Spine) device in the VXLAN fabric that is connected to a device with the **Edge Router** role in the external fabric .
3. **Back to Back Only** - Configure VRF Lite IFCs between directly connected border leaf (Spine) device interfaces of different VXLAN fabrics.
4. **Both** - Use this option to configure IFCs for the modes **To External Only** and **Back to Back Only**.




---

**Note** DCI subnet is required, even if the VRF Lite mode is **Manual**. This helps with the DCNM resource handling.

---

The default mode in fabric settings is Manual Mode. In order to change the mode to any of the others, edit the fabric settings. Under the Resources Tab, modify the VRF Lite Deployment field to one of the above mentioned auto config modes. In this example, ToExternalOnly option is chosen.

**VRF Lite Subnet IP Range:** The IP address for VRF Lite IFC deployment is chosen from this range. The default value is 10.33.0.0/16. Best practice is to ensure that each fabric has its own unique range and distinct from any underlay range in order to avoid possible duplication. These addresses are reserved with the Resource Manager.

**VRF Lite Subnet Mask:** By default its set to /30 which is best practice for P2P links.

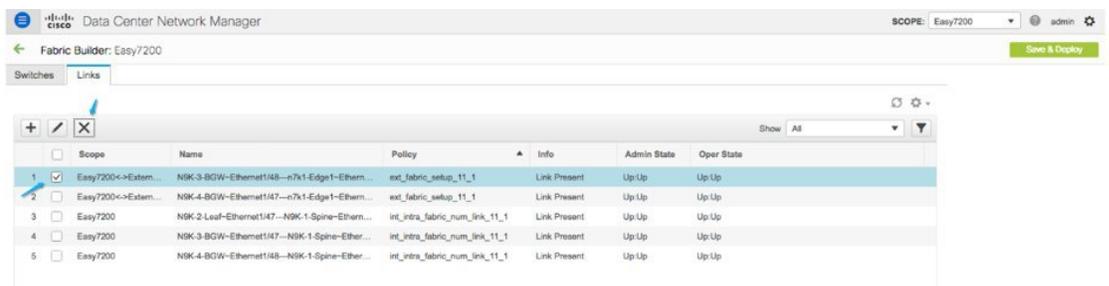
Similarly, update the settings for the Easy60000 fabric too.

Once the fields are set, execute the **Save and Deploy** option in the VXLAN and external fabrics.

## Deleting VRF Lite IFCs

Before deleting the IFC, remove all VRF extensions enabled on the IFC. Else, an error message is reported.

1. Go to the Links tab of the fabric.
2. Select the links with VRF Lite policy configured and click the delete button.



3. Click OK to confirm deletion.
4. Execute the Save and Deploy option in the fabric to reset the VRF Lite policy.

### Deleting VRF Extensions deployed in External Fabric

This is a two part process:

1. Delete the sub interface created using interface TAB.



**Note** Skip this step if the VRF extension is to a non-Nexus device.

2. Delete the policy created for eBGP external connection.

### Deleting the sub-interface

Navigate to the Control->Interfaces page as shown below, select the sub-interface(s) to be deleted and click the delete button.

Control / Fabrics / Interfaces

Interfaces

|                                     | Device Name | Name            | Admin | Oper | Reason                | Policy                  | Overlay N |
|-------------------------------------|-------------|-----------------|-------|------|-----------------------|-------------------------|-----------|
| <input type="checkbox"/>            | n7k1-Edge1  | mgmt0           | ↑     | ↑    | ok                    | NA                      | NA        |
| <input type="checkbox"/>            | n7k1-Edge1  | Vlan1           | ↓     | ↓    | Administratively down | NA                      | NA        |
| <input type="checkbox"/>            | n7k1-Edge1  | Loopback0       | ↑     | ↑    | ok                    | NA                      | NA        |
| <input type="checkbox"/>            | n7k1-Edge1  | Loopback1       | ↑     | ↓    |                       | NA                      | NA        |
| <input type="checkbox"/>            | n7k1-Edge1  | Ethernet7/3     | ↓     | ↓    | Administratively down | NA                      | NA        |
| <input type="checkbox"/>            | n7k1-Edge1  | Ethernet7/5     | ↑     | ↓    | Link not connected    | int_trunk_host_11_1     | NA        |
| <input type="checkbox"/>            | n7k1-Edge1  | Ethernet7/6     | ↑     | ↑    | ok                    | routed_host             | NA        |
| <input type="checkbox"/>            | n7k1-Edge1  | Ethernet7/9     | ↓     | ↓    | Administratively down | NA                      | NA        |
| <input type="checkbox"/>            | n7k1-Edge1  | Ethernet7/1/1   | ↓     | ↓    | Administratively down | NA                      | NA        |
| <input type="checkbox"/>            | n7k1-Edge1  | Ethernet7/1/2   | ↑     | ↓    | Link not connected    | NA                      | NA        |
| <input type="checkbox"/>            | n7k1-Edge1  | Ethernet7/1/3   | ↓     | ↓    | Administratively down | NA                      | NA        |
| <input type="checkbox"/>            | n7k1-Edge1  | Ethernet7/1/4   | ↑     | ↑    | ok                    | ext_int_routed_host_11_ | NA        |
| <input checked="" type="checkbox"/> | n7k1-Edge1  | Ethernet7/1/4.2 | ↑     | ↑    | ok                    | int_subif_11_1          | NA        |

### Deleting the eBGP policy

Navigate to fabric builder page and select the relevant external fabric (External65000 in this example). Select the device and using the second mouse button select view edit policy.

Select the row for the policy ID used in eBGP policy create. Click the “X” as shown below to delete the policy.

Issue a save and deploy in external fabric to deploy the policy change.

View/Edit Policies for n7k1-Edge1 ( TBM14299900:Edge1 )

Selected 1 / Total 30

+ / ✖ / View / View All / Deploy

Show Quick Filter

| Template   | Priority | Fabric Name   | Serial Number     | Editable | Entity Type | Entity Name     | Source          | Policy ID    |
|--|----------|---------------|-------------------|----------|-------------|-----------------|-----------------|--------------|
| <input checked="" type="checkbox"/> External_VRF_Lite_eBGP | 500      | ExternalE5000 | TBM14299900:Edge1 | true     | SWITCH      | SWITCH          |                 | POLICY-33350 |
| <input type="checkbox"/> base_external_router              | 500      | ExternalE5000 | TBM14299900:Edge1 | true     | SWITCH      | SWITCH          |                 | POLICY-33360 |
| <input type="checkbox"/> breakout_interface                | 500      | ExternalE5000 | TBM14299900:Edge1 | true     | SWITCH      | SWITCH          |                 | POLICY-33960 |
| <input type="checkbox"/> routed_interface                  | 350      | ExternalE5000 | TBM14299900:Edge1 | false    | INTERFACE   | Ethernet7/1/4   | LINK-UUID-4770  | POLICY-32770 |
| <input type="checkbox"/> routed_interface                  | 350      | ExternalE5000 | TBM14299900:Edge1 | false    | INTERFACE   | Ethernet7/4/1   | LINK-UUID-4810  | POLICY-32870 |
| <input type="checkbox"/> interface_vrf                     | 350      | ExternalE5000 | TBM14299900:Edge1 | false    | INTERFACE   | Ethernet7/4/1.2 | Ethernet7/4/1.2 | POLICY-33370 |
| <input type="checkbox"/> interface_vrf                     | 350      | ExternalE5000 | TBM14299900:Edge1 | false    | INTERFACE   | Ethernet7/1/4.2 | Ethernet7/1/4.2 | POLICY-33410 |
| <input type="checkbox"/> routed_host                       | 350      | ExternalE5000 | TBM14299900:Edge1 | false    | INTERFACE   | Ethernet7/6     |                 | POLICY-33900 |
| <input type="checkbox"/> trunk_interface                   | 350      | ExternalE5000 | TBM14299900:Edge1 | false    | INTERFACE   | Ethernet7/5     | Ethernet7/5     | POLICY-34170 |
| <input type="checkbox"/> interface_mtu                     | 352      | ExternalE5000 | TBM14299900:Edge1 | false    | INTERFACE   | Ethernet7/1/4   | LINK-UUID-4770  | POLICY-32780 |
| <input type="checkbox"/> no shut interface                 | 362      | ExternalE5000 | TBM14299900:Edge1 | false    | INTERFACE   | Ethernet7/1/4   | LINK-UUID-4770  | POLICY-32800 |

### Deleting IFCs Created By Automatic VRF Lite creation

Editing and deleting IFCs are done through the Link tab in the VXLAN fabric. The extra consideration for auto configured IFCs is that, in order to prevent the regeneration of IFC on next save and deploy, the mode should be changed back to manual mode, or Save config should be done only on the relevant devices.

- In a consecutive scenario, if you delete the VRF lite IFC on one of the fabrics, the VRF lite is deleted from the peer fabric as well.
- When you want to delete a VRF lite between an easy fabric and an external fabric, delete the extension in the easy fabric using the top-down approach. The extension will be automatically deleted from the external fabric.
- Deploy the configurations in both the fabrics.

## 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

### N9K-3-BGW Configurations

N9K-3-BGW (base border configurations) generated by template ext\_base\_border\_vrflite\_11\_1



**Note** *switch(config)#* refers to the global configuration mode. To access this mode, type the following on your switch: *switch# configure terminal*.

```
(config) #
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
route-map extcon-rmap-filter-allow-host deny 10
    match ip address prefix-list default-route
route-map extcon-rmap-filter-allow-host 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
route-map extcon-rmap-filter-v6-allow-host deny 10
    match ipv6 address prefix-list default-route-v6
route-map extcon-rmap-filter-v6-allow-host permit 1000
```

### N9K-3-BGW VRF extension configuration

```
(config) #
configure profile MyVRF_50000
    vlan 2000
        vn-segment 50000
    interface vlan2000
        vrf member myvrf_50000
            ip forward
            ipv6 forward
            no ip redirects
            no ipv6 redirects
            mtu 9216
            no shutdown

(config) #

vrf context myvrf_50000
    vni 50000
    rd auto
    address-family ipv4 unicast
        route-target both auto
        route-target both auto evpn

    ip route 0.0.0.0/0 2.2.2.1
    address-family ipv6 unicast
        route-target both auto
        route-target both auto evpn

router bgp 7200
    vrf myvrf_50000
        address-family ipv4 unicast
            advertise l2vpn evpn
            redistribute direct route-map fabric-rmap-redis-subnet
            maximum-paths ibgp 2
            network 0.0.0.0/0
        address-family ipv6 unicast
            advertise l2vpn evpn
            redistribute direct route-map fabric-rmap-redis-subnet
            maximum-paths ibgp 2
```

```
neighbor 2.2.2.1 remote-as 65000
  address-family ipv4 unicast
  send-community both
  route-map extcon-rmap-filter out

(config) #

interface ethernet1/48.2
  encapsulation dot1q 2
  vrf member myvrf_50000
  ip address 2.2.2.2/24
  no shutdown
interface nve1
  member vni 50000 associate-vrf
configure terminal
  apply profile MyVRF_50000
```





## CHAPTER 12

# Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - Multi-Site

---

This chapter explains LAN Fabric border provisioning using EVPN Multi-Site feature.

- [Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - Multi-Site](#) , on page 547
- [Prerequisites](#) , on page 548
- [Limitations](#) , on page 549
- [Save & Deploy Operation in the MSD Fabric](#) , on page 549
- [EVPN Multi-Site Configuration](#) , on page 551
- [Viewing, Editing and Deleting Multi-Site Overlays](#) , on page 561
- [Deleting Multi-Site IFCs](#) , on page 561
- [Creating and Deploying Networks and VRFs in the MSD Fabric](#) , on page 562
- [Deploying a Legacy Site BGW \(vPC-BGWs\)](#) , on page 566
- [Additional References](#) , on page 570
- [Appendix](#) , on page 570

## 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. Also, you can connect two member fabrics of a Multi-Site Domain (MSD).

Introduced in DCNM 11.0(1) release, MSD 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.

For a detailed explanation on the EVPN Multi-Site feature, see the [VXLAN BGP EVPN Multi-Site Design and Deployment](#) document.

*Configuration methods* - You can create underlay and overlay Inter-Fabric Connections (IFCs) between member fabrics through auto-configuration and through the DCNM GUI.

vPC configuration is supported for BGWs with the role **Border Gateway** from Cisco DCNM Release 11.1(1).

*Supported destination devices* - You can connect a VXLAN fabric to Cisco Nexus and non-Nexus devices. A connected non-Cisco device can also be represented in the topology.

## 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 fabric.
- 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.
- Execute the **Save & Deploy** operation in the member fabrics and external fabrics, and then in the MSD fabric.




---

**Note** The **Save & Deploy** button appears at the top right part of the fabric topology screen (accessible through the **Fabric Builder** window and clicking the fabric).

---

- Ensure that the role of the designated BGW is **Border Gateway** (or **Border Gateway Spine** for spine switches). To verify, right-click the BGW and click **Set role**. You can see that (**current**) is added to the current role of the switch.
- 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. Ensure that each fabric has a unique IP address pool to avoid duplicates.
- 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).



**Note** **Case 1** - During network creation, if a VLAN is specified, then for every switch, when you attach the network to the switch, automatically the VLAN will be autopopulated with the same VLAN that was specified during network creation. The network listing screen shows the VLAN a network level which applies for all the switch (has to be the same). The other thing to keep in mind is that even if one specified a VLAN during network creation, this can still be overwritten on a per switch basis.

**Case 2** - During network creation, if a VLAN is not specified, then for every switch, when you attach the network to the switch, the next free VLAN from the per-switch VLAN pool is autopopulated. This means that on a per-switch basis, the VLAN may be different. The user can always overwrite the autopopulated VLAN and DCNM will honor it. For this case, it is possible that VNI 10000 may use VLAN 10 on leaf1 and VLAN 11 on leaf2. Hence, in the network listing, in this case, the VLAN will not be showcased.

DCNM always keeps track of VLANs on a per switch basis in its resource manager. This is true for either of the 2 cases mentioned above.

## Limitations

- vPC configuration is not supported for the **Border Gateway Spine** role.
- The VXLAN OAM feature in Cisco DCNM is only supported on a single fabric or site.
- FEX is not supported on a Border Gateway or a Border Leaf with vPC or anycast.

## Save & Deploy Operation in the MSD Fabric

These are some operations performed when you execute **Save & Deploy**:

- **Duplicate IP address check:** The MSD fabric checks if any BGW has a duplicate IP address. If so, an error message is displayed.



Change the BGP peering loopback IP address of the BGW(s).

After duplicate IP address issues are resolved, execute the **Save & Deploy** operation again in the MSD fabric.

- **BGW base configuration:** When you execute Save and Deploy for the first time in the MSD fabric (assuming there are currently no IFCs or overlays to deploy), appropriate base configurations are deployed on the BGWs. They are given below:

| Configuration  | Description  |
|--|--|
| <pre>evpn multisite border-gateway 7200 delay-restore time 300</pre>   | <p>7200 is the site ID of the member fabric Easy7200.</p> <p>BGP ASN value is used to auto populate the site ID field. You can override this value. Even if you change the BGP ASN value, the site ID is still set to the first BGP ASN value.</p>   |
| <pre>interface nve1 multisite border-gateway interface loopback100</pre>   | <p>The loopback interface 100 is the configuration set in the MSD fabric settings. Once a loopback ID is chosen and <b>Save and Deploy</b> is executed, the loopback ID cannot be changed.</p> <p>To modify the role of the BGW in the MSD fabric, perform the following steps:</p> <ol style="list-style-type: none"> <li>1. In the easy fabric, modify the role of the BGW to leaf or border.</li> <li>2. Save and deploy the changes.<br/>This will remove the loopback 100 from the switch</li> <li>3. Change role back to BGW, and do a save and deploy.</li> <li>4. In the MSD fabric, change the loopback ID setting to a desired value, and do a save and deploy.</li> </ol> |
| <pre>interface ethernet1/47 evpn multisite fabric-tracking</pre>   | <p>The <b>evpn multisite fabric-tracking</b> command is configured on all ports on a Border Gateway that have a connection to a switch with a Spine role.</p> <p>In case of a Border Gateway Spine role, all ports facing the leaf switch have this command configured</p>   |
| <pre>interface loopback100 ip address 10.10.0.1/32 tag 54321 ip router ospf UNDERLAY area 0.0.0.0 ip pim sparse-mode no shutdown</pre> | <p>The Multi-Site loopback interface. This is configured on all Border Gateway (Spines).</p> <p>All BGWs in the same fabric get the same IP address. Each fabric gets its own unique IP address.</p> <p>It is not possible to change this address or ID without first changing role of the BGW.</p>  |

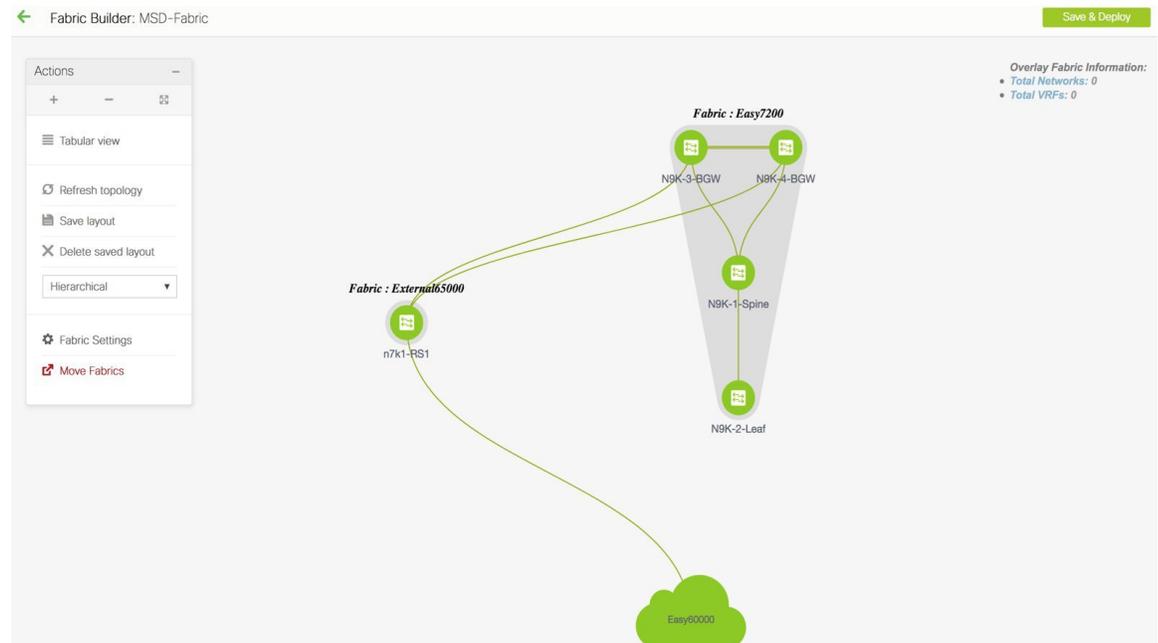
| Configuration   | Description  |
|---|--|
| <pre>route-map rmap-redirect-direct permit 10   match tag 54321</pre> | This is the configuration to redistribute the BGP peering loopback IP address (commonly loopback0), the VTEP primary (in case of vPC, the loopback secondary IP address), commonly loopback1, and the Multi-Site loopback IP address into the Multi-Site eBGP underlay sessions. |

- When you execute the **Save & Deploy** operation in the MSD fabric, it works on all the BGW (or BGW Spine) devices in the member fabrics of an MSD.

After completing the EVPN Multi-Site specific prerequisites, start EVPN Multi-Site configuration. A sample scenario is explained.

## EVPN Multi-Site Configuration

The EVPN Multi-Site feature is explained through an example scenario. Consider two VXLAN BGP EVPN fabrics, **Easy60000** and **Easy7200**, and an external fabric, **External65000**. The three fabrics are member fabrics of the MSD fabric **MSD-Fabric** and identified by a unique AS number. Easy60000 and Easy7200 are connected to a route server in External65000 (each fabric is). This document shows you how to enable end-to-end Layer 3 and Layer 2 traffic between hosts in Easy60000 and Easy7200, through the route server.



VXLAN BGP EVPN intra-fabric configurations, including network and VRF configurations are provisioned on the switches through DCNM software, 11.1(1) release. However, server traffic between the fabrics is only possible through the following configurations:

- A Data Center Interconnect (DCI) function like the Multi-Site feature is configured on the BGWs of both the fabrics (N9K-3-BGW and N9K-4-BGW in Easy7200, and the BGW in Easy60000). As part of

the configuration, since the BGWs of the fabrics are connected to the route server N7k1-RS1 in the external fabric External65000, appropriate eBGP peering configurations are enabled on the BGWs.

- As of now, overlay networks and VRFs are enabled on the non-BGW leaf and spine switches. For a fabric's traffic to go beyond the BGW, networks and VRFs should be deployed on all the BGWs too.

In a nutshell, the EVPN Multi-Site feature configuration comprises of setting up the BGW base configuration (enabled during the Save & Deploy operation), the eBGP underlay and overlay peering from the three BGWs to the route server N7k1-RS1. Both the underlay and overlay peering are established over eBGP through DCNM 11.1(1).

You can create Multi-Site Inter-Fabric Connections (IFCs) between the fabrics through the DCNM GUI or through automatic configuration. First, underlay IFC creation is explained, followed by the overlay IFC creation.

## Configuring Multi-Site Underlay IFCs - DCNM GUI

The end-to-end configurations can be split into these 2 high-level steps.

### Step 1 - EVPN Multi-Site configurations on the BGWs in Easy7200

### Step 2 - EVPN Multi-Site configurations on the BGW in Easy60000



#### Note

An inter-fabric link is a physical connection between two Ethernet interfaces (an underlay connection) or a virtual connection (a fabric overlay connection between two loopback interfaces). When you add a physical connection between devices, the new link appears in the Links tab by default.

### Step 1 - EVPN Multi-Site configurations on the BGWs in Easy7200

For Multi-Site connectivity from Easy7200 to the external fabric, N9K-3-BGW and N9K-4-BGW are connected to the route server N7k1-RS1 in the external fabric. Follow these steps:

#### Deploying underlay IFCs between Easy7200 and External65000

- Deploying Underlay IFC from N9K-3-BGW to N7k1-RS1.
- Deploying Underlay IFC from N9K-4-BGW to N7k1-RS1.

#### Deploying Underlay IFC from N9K-3-BGW to N7k1-RS1

For the Multi-Site DCNM GUI configuration option, the **Deploy Border Gateway Method** field in the MSD fabric's settings (**DCI** tab) is set to **Manual**.

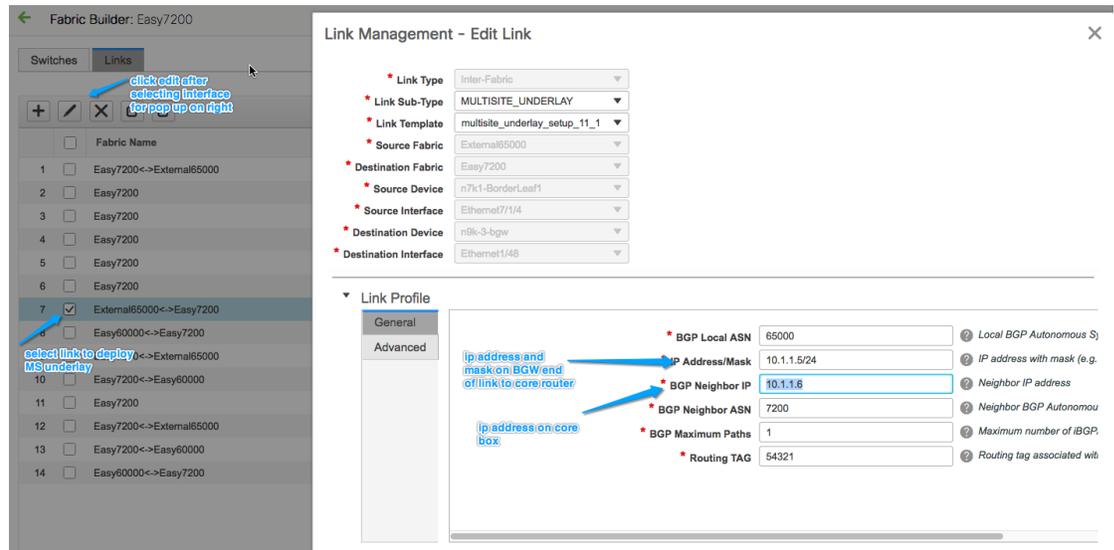
1. Navigate to the **Links** tab and select the physical link connecting N9K-3-BGW to N7k1-RS1.
2. Click the link edit icon as shown in the figure below to bring up the pop up.
3. Select the MS underlay IFC sub type and fill in the required fields.



#### Note

Enter the value as 1 in the **BGP Maximum Paths** field to allow DCNM to pick maximum path value. Enter a value between 2 and 64 to decide the maximum path value.

4. Save and deploy in the MSD will deploy the configuration on the N9K-3-BGW and N7k1-RS1. Similar steps can be used to edit already created IFCs via the Links tab.



5. Similarly, create the underlay IFC from N9K-4-BGW to N7k1-RS1.

This completes Step 1 of the following.

**Step 1** - EVPN Multi-Site configurations on the BGWs in Easy7200.

**Step 2** - EVPN Multi-Site configurations on the BGW in Easy60000.

Next, configurations are enabled on the BGW in Easy60000.

**Step 2** - EVPN Multi-Site configurations on the BGW in Easy60000

For Multi-Site connectivity between the Easy6000 fabric and the external fabric, EVPN Multi-Site configurations are enabled on the BGW interfaces in Easy60000 that are connected to the route server (N7k1-RS1) in the external fabric. Follow the steps as per the explanation for the connections between Easy7200 and External65000.

## Configuring Multi-Site Underlay IFCs - Autoconfiguration

An underlay IFC is a physical link between the devices' interfaces.

- For underlay connectivity from Easy7200 to the external fabric, N9K-3-BGW and N9K-4-BGW are connected to the route server N7k1-RS1 in the external fabric.
- For underlay connectivity from Easy60000 to the external fabric, its BGW is connected to the route server N7k1-RS1.

### Deploying Multi-Site Underlay IFCs Through Autoconfiguration

The underlay generated by DCNM is an eBGP session in the default IPv4 unicast routing table, in order to distribute the three loopback addresses needed for the Multi-Site control plane and data plane to function correctly.

For the Multi-Site autoconfiguration option, the underlay IFCs are automatically deployed by the MSD fabric.

The following rules apply to Multi-Site underlay IFC creation:

1. Check the **Multi-Site Underlay IFC Auto Deployment Flag** check box to enable the multi-site underlay autoconfiguration. Uncheck the check box to disable autoconfiguration. The check box is unchecked by default.

The screenshot shows the 'Edit Fabric' window in the DCNM GUI, specifically the 'Resources' tab. The 'Multi-Site Underlay IFC Auto Deployment Flag' checkbox is unchecked. A blue arrow points to this checkbox with the text 'check to enable, clear for manual'. Other fields include Fabric Name (MSD), Fabric Template (MSD\_Fabric\_11\_1), DCI Subnet IP Range (10.10.1.0/24), Subnet Target Mask (30), and Multi-Site Overlap IFC Deployment Method (Manual). The 'Save' and 'Cancel' buttons are visible at the bottom right.

2. An IFC is deployed on every physical connection between the BGWs of different member fabrics that are physically connected.
3. An IFC is deployed on every physical connection between a BGW and a router with the role Core Router imported into an external fabric which is a member of the MSD fabric.

If you do not want an IFC to be auto generated on a connection, then shut the link, execute the Save & Deploy operation, and delete the undesired IFCs. Also, ensure that there is no existing policy or pre-configured IP address on the interface. Else, use the Manual mode.

4. The IP addresses used to deploy the underlay are derived from the IP address range in the DCI Subnet IP Range field (DCI tab) of the MSD fabric.

Just like overlay IFCs, Multi-Site underlay IFCs can be viewed via the MSD, external and member fabrics. Also, the underlay IFCs can be edited and deleted via the VXLAN or MSD fabrics.

## Configuring Multi-Site Underlay IFCs Towards a Non-Nexus Device - DCNM GUI

In this case, the non-Nexus device is not imported into DCNM, or discovered through Cisco Discovery Protocol or Link Layer Discovery Protocol (LLDP). For example, a Cisco ASR 9000 Series router or even a non-Cisco device.

The steps are similar to the **Configuring Multi-Site Underlay IFCs - DCNM GUI** task.

1. In the **Fabric Builder** window, choose the **Easy7200** fabric.

The **Easy7200** topology window appears.

- From the **Actions** panel at the left, click **Tabular view**.

The **Switches | Links** window appears.

- Click the **Links** tab and click +.

The **Add Link** window appears.

- Fill in the fields.

**Link Type** – Choose **Inter-Fabric**.

**Link Sub-Type** – Choose **MULTISITE\_UNDERLAY**.

**Link Template** - By default, the **ext\_multisite\_underlay\_setup\_11\_1** template is populated.

**Source Fabric** - **Easy7200** is selected by default since the IFC is created from **Easy7200** to the ASR device.

**Destination Fabric** – Select the external fabric. In this case, **External65000** is selected.

**Source Device** and **Source Interface** - Choose the border device and the interface that is connected to the ASR device.

**Destination Device** - Type any string that identifies the device. The destination device **ASR9K-RS2** does not appear in the drop-down list when you create an IFC for the first time. Once you create an IFC towards **ASR9K-RS2** and associate it with the external fabric **External65000**, **ASR9K-RS2** appears in the list of devices displayed in the **Destination Device** field.

Also, after the first IFC creation, **ASR9K-RS2** is displayed in the **External65000** external fabric topology, within Fabric Builder.

**Destination Interface** - Type any string that identifies the interface.

You have to manually enter the destination interface name each time.

**General** tab in the **Link Profile** section.

**Source BGP ASN** - In this field, the AS number of the source fabric **Easy7200** is autopopulated.

**Source IP Address/Mask** - Enter the IP address and mask that is used as the local interface for the Multi-Site underlay IFC.

**Destination IP** - Enter the IP address of the **ASR9K-RS2** interface used as the eBGP neighbor.

**Destination BGP ASN** - In this field, the AS number of the external fabric **External65000** is autopopulated since it is chosen as the external fabric.

5. Click **Save** at the bottom right of the window.  
The **Switches|Links** window appears again. You can see that the IFC entry is updated.
6. Click **Save and Deploy** at the top right of the window.  
The link on which the IFC is deployed has the relevant policy configured in the **Policy** column.
7. Go to the **Scope** drop-down list at the top right of the window and choose **External65000**. The external fabric **Links** window is displayed. You can see that the IFC created from **Easy7200** to the ASR device is displayed here.

## Configuring Multi-Site Overlay IFCs

An overlay IFC is a link between the devices' loopback0 interfaces.

Deploying Overlay IFCs in Easy7200 and Easy60000 comprises of these steps:

- Deploying Overlay IFC from N9K-3-BGW to N7k1-RS1.
- Deploying Overlay IFC from N9K-4-BGW to N7k1-RS1.
- Deploying the Overlay IFC from the BGW in Easy60000 to N7k1-RS1.

### Deploying Overlay IFCs between Easy7200 and External65000

- Deploying Overlay IFC from N9K-3-BGW to N7k1-RS1.
- Deploying Overlay IFC from N9K-4-BGW to N7k1-RS1.

### Deploying Overlay IFCs - from N9K-3-BGW to N7k1-RS1

1. Click **Control > Fabric Builder**. The Fabric Builder window appears.
2. Choose the MSD fabric, **MSD-Fabric**. The fabric topology comes up.
3. Click **Tabular view**. The **Switches | Links** screen comes up.
4. Click the **Links** tab. It lists links within the MSD fabric. Each row either represents an intra-fabric link within Easy7200 or Easy60000, or a link between border devices of member fabrics, including External65000.
5. Click the **Add Link** icon at the top left part of the screen.  
The **Link Management – Add Link** screen comes up.  
Some fields are explained:  
Link Type – Inter-Fabric is autopopulated.  
Link Sub-Type – Choose **MULTISITE\_OVERLAY**.

Link Template – The default template for creating an overlay is displayed.

You can edit the template or create a new one with custom configurations.

In the General tab, the BGP AS numbers of Easy7200 and External65000 are displayed. Fill in the other fields as explained. The BGP AS numbers are derived based on fabric values.

Fabric Builder: MSD-Fabric

Switches Links [Add link icon](#)

Link Management - Add Link

1 \* Link Type Inter-Fabric

2 \* Link Sub-Type MULTISITE\_OVERLAY

3 \* Link Template ext\_evpn\_multisite\_overlay\_se

4 \* Source Fabric Easy7200

5 \* Destination Fabric External65000

6 \* Source Device N9K-3-BGW

\* Source Interface Loopback0

\* Destination Device n7k1-RS1

\* Destination Interface Loopback0

Link Profile

General

\* Local BGP AS # 7200 ? Local BGP Autonomous System Number

\* SOURCE\_IP 10.1.0.1 ?

\* NEIGHBOR\_IP 2.2.2.2 ?

Save

Legend: Link type: Inter-Fabric, Link sub type: multisite overlay

6. Click Save at the bottom right part of the screen.

The Switches|Links screen comes up again. You can see that the IFC entry is updated.

7. Click Save & Deploy at the top right part of the screen.

8. Go to the **Scope** drop-down list at the top right of the window and choose External65000. The external fabric Links screen is displayed. You can see that the two IFCs created from Easy7200 to External65000 is displayed here.



**Note** When you create an IFC or edit its setting in the VXLAN fabric, the corresponding entry is automatically created in the connected external fabric.

9. Click Save and Deploy to save the IFCs creation on External65000.

10. Similarly, create an overlay IFC from N9K-4-BGW to N7k1-RS1.

After the overlay IFCs from N9K-3-BGW and N9K-4-BGW to N7k1-RS1 are deployed, the fabric overlay traffic can flow between Easy7200 and External65000.

11. Similarly, deploy the overlay IFC from the BGW in the Easy60000 fabric to N7k1-RS1.

## Configuring Multi-Site Overlay IFCs - Autoconfiguration

An overlay IFC is a link between the devices' loopback0 interfaces. For overlay connectivity from the Easy7200 and Easy60000 fabrics to the route server N7k1-RS1 in External65000, a link is enabled between the BGW devices and the N7k1-RS1's loopback0 interfaces.

### Deploying Overlay IFCs in Easy7200 and Easy60000

- Deploying Overlay IFC from N9K-3-BGW to N7k1-RS1.
- Deploying Overlay IFC from N9K-4-BGW to N7k1-RS1.
- Deploying the Overlay IFC from the BGW in Easy60000 to N7k1-RS1.

### Deploying Multi-Site Overlay IFCs Through Autoconfiguration

You can automatically configure the Multi-Site overlay through one of these options:

1. Centralized to Route Server - The BGW forms an overlay to the route server. This option is explained in the example.
2. Direct to BGW: A full mesh of Multi-Site Overlay IFC from every BGW in a fabric to every BGW in other member fabrics.

To choose one of the above options, go to the MSD fabric's settings, select the DCI tab, and set the Deploy Border Gateway Method field to Centralized to Route\_Server (such as for this example) or Direct to BGW. By default, the Manual option is selected.

#### Edit Fabric

\* Fabric Name : MSD

\* Fabric Template : MSD\_Fabric\_11\_1

General DCI Resources

DCI Subnet IP Range 10.10.1.0/24 Address range to assign P2P DCI Links

Subnet Target Mask 30 Target Mask for Subnet Range (Min:8, Max:31)

\* Multi-Site Overlay IFC Deployment Method Centralized\_To\_Route\_Server Manual, Auto Overlay EVPN Peering to Route Servers, Auto Overlay EVPN Direct Peering to Border Gateways

\* Multi-Site Route Server List 1.1.1.1 Multi-Site Router-Server peer list, e.g. 128.89.0.1, 128.89.0.2

\* Multi-Site Route Server BGP ASN List 65000 1-4294967295 | 1-65535[0-65535], e.g. 65000, 65001

Multi-Site Underlay IFC Auto Deployment Flag

comma seperated list of route server addresses. Device must be imported into DCNM

comma seperated list if corresponding BGP AS numbers for the router servers list

The IFCs needed for deployment of Networks and VRFs at the BGW nodes can be auto configured via the MSD fabric template. The settings to enable that are in MSD fabric template.

The default mode for the **Deploy Border Gateway Method** field is **Manual**, which implies that the IFCs have to be created via the link tab in MSD fabric. It must be changed to the Centralized to Route\_Server or Direct to BGW mode for autoconfiguration.

The IFCs created via auto configuration can only be edited or deleted via the link tab in MSD or member fabrics (except external fabric). As long as an IFC exists, or there is any user defined policy on the physical or logical link, auto configuration will not touch the IFC configuration.

You can see that Centralized to Route\_Server is selected in the Deploy Border Gateway Method field in the above image.

### Centralized to Route Server

This implies that all BGW devices in all member fabrics will create a Multi-Site overlay BGP connection to one or more route servers in one or more external fabrics which are members of the MSD fabric.

In this topology, there is one route server n7k1-RS1, and its BGP peering address (1.1.1.1) is shown in the route server list. This peering address can be configured out of band or with create interface tab in DCNM. N7k1-RS1 must be imported into the DCNM (in the external fabric, in this example) and the peering address configured before executing the Save & Deploy option.

You can edit the route server peering IP address list at any time, but you can delete a configured Multi-Site overlay only through the Links tab.

The BGP AS number of each route server should be specified in the MSD fabric settings. Note that the route server AS number can be different than the fabric AS number of the external fabric.

## Configuring Multi-Site Overlay IFCs Towards a Non-Nexus Device - DCNM GUI

In this case, the non-Nexus device is not imported into DCNM, or discovered through Cisco Discovery Protocol or Link Layer Discovery Protocol (LLDP). For example, a Cisco ASR 9000 Series router or even a non-Cisco device.

The steps are similar to the **Configuring Multi-Site Overlay IFCs - DCNM GUI** task.

1. In the **Fabric Builder** window, choose the **Easy7200** fabric.

The **Easy7200** topology window appears.

2. From the **Actions** panel, click **Tabular view**.

The **Switches | Links** window appears.

3. Click the **Links** tab and click +.

The **Add Link** screen comes up.

4. Fill in the fields.

The screenshot shows the 'Link Management - Add Link' dialog in the DCNM GUI. The dialog is divided into two main sections: 'Link Management' and 'Link Profile'.

**Link Management Section:**

- Link Type:** Inter-Fabric
- Link Sub-Type:** MULTISITE\_OVERLAY
- Link Template:** ext\_evpn\_multisite\_overlay\_setup
- Source Fabric:** Easy7200
- Destination Fabric:** External65000
- Source Device:** N9K-3-BGW
- Source Interface:** Loopback0
- Destination Device:** RS1
- Destination Interface:** loopback0

**Link Profile Section (General tab):**

- Local BGP AS #:** 7200
- SOURCE\_IP:** 4.4.4.4
- NEIGHBOR\_IP:** 5.5.5.5

A 'Save' button is located at the bottom right of the dialog.

**Link Type** – Choose **Inter-Fabric**.

**Link Sub-Type** – Choose **MULTISITE\_OVERLAY**.

**Link Template** – By default, the **ext\_evpn\_multisite\_overlay\_setup** template is populated.

**Source Fabric** – **Easy7200** is selected by default since the IFC is created from **Easy7200** to the ASR device.

**Destination Fabric** – Select the external fabric. In this case, **External65000** is selected.

**Source Device** and **Source Interface** - Choose the border device and the loopback0 interface that is the source interface of the overlay.

**Destination Device:** Type any string that identifies the device. The destination device **ASR9K-RS1** does not appear in the drop-down list when you create an IFC for the first time. Once you create an IFC towards **ASR9K-RS1** and associate it with the external fabric **External65000**, **ASR9K-RS1** appears in the list of devices displayed in the **Destination Device** field.

Also, after the first IFC creation, **ASR9K-RS1** is displayed in the **External65000** topology screen, within Fabric Builder.

**Destination Interface:** Type any string that identifies the interface.

You have to manually enter the destination interface name each time.

**General** tab in the **Link Profile** section.

**Source BGP ASN:** In this field, the AS number of the source fabric **Easy7200** is autopopulated.

**Source IP Address/Mask:** Enter the IP address of the loopback0 interface for the Multi-Site overlay IFC.

**Destination IP:** Enter the IP address of the **ASR9K-RS1** loopback interface used for this Multi-Site overlay IFC.

**Destination BGP ASN:** In this field, the AS number of the external fabric **External65000** is autopopulated since it is chosen as the external fabric.

5. Click **Save** at the bottom right part of the screen.

The **Switches|Links** screen comes up again. You can see that the IFC entry is updated.

6. Click **Save and Deploy** at the top right part of the screen.

The link on which the IFC is deployed has the relevant policy configured in the **Policy** column.

7. Go to the **Scope** drop-down list at the top right of the window and choose **External65000**. The external fabric **Links** screen is displayed. You can see that the overlay IFC is displayed here.

## Overlay and Underlay Peering Configurations on the Route Server N7k1-RS1

When you execute the Save and Deploy operation in the MSD fabric during the IFCs creation, peering configurations are enabled on the router server N7k1-RS1 towards the BGWs in the VXLAN fabrics.

## Viewing, Editing and Deleting Multi-Site Overlays

The overlay IFCs can be viewed via the MSD and member fabrics Links tab as shown below.

The IFCs can be edited and deleted in the member fabric or in the MSD fabric.

Multi-Site overlay IFCs can also be created by the links tab in MSD fabric.

Once the IFC is deleted, you should execute the Save & Deploy operation in the external and VXLAN fabric (or MSD fabric) to undeploy the IFC on the switches and remove the intent from DCNM.



### Note

Until a particular IFC is completely deleted from DCNM, auto configuration will not regenerate it on a Save & Deploy operation in the MSD fabric.

|    | Scope                     | Name  | Policy                            | Info             | Admin State | Oper State |
|----|---------------------------|---|-----------------------------------|------------------|-------------|------------|
| 1  | Easy7200<->External65000  | N9K-4-BGW-loopback0--n7k1-RS1-Loopback0           | ext_evpr_multisite_overlay_setup  | Neighbor Missing | --          | --         |
| 2  | Easy7200<->External65000  | N9K-3-BGW-loopback0--n7k1-RS1-Loopback0           | ext_evpr_multisite_overlay_setup  | Neighbor Missing | --          | --         |
| 3  | Easy60000<->External65000 | N9K-15-BGW-Ethernet1/3--n7k1-RS1-Ethernet7/10/1   |                                   | Link Present     | Up:Up       | Up:Up      |
| 4  | Easy7200                  | N9K-2-Leaf-Ethernet1/47--N9K-1-Spine-Ethernet1/47 | int_intra_fabric_num_link_11_1    | Link Present     | Up:Up       | Up:Up      |
| 5  | Easy7200<->External65000  | N9K-3-BGW-Ethernet1/48--n7k1-RS1-Ethernet7/1/4    | ext_multisite_underlay_setup_11_1 | Link Present     | Up:Up       | Up:Up      |
| 6  | Easy7200                  | N9K-3-BGW-Ethernet1/47--N9K-1-Spine-Ethernet1/43  | int_intra_fabric_num_link_11_1    | Link Present     | Up:Up       | Up:Up      |
| 7  | Easy7200<->External65000  | N9K-4-BGW-Ethernet1/47--n7k1-RS1-Ethernet7/4/1    | ext_multisite_underlay_setup_11_1 | Link Present     | Up:Up       | Up:Up      |
| 8  | Easy7200                  | N9K-4-BGW-Ethernet1/22--N9K-3-BGW-Ethernet1/22    | int_intra_fabric_num_link_11_1    | Link Present     | Up:Up       | Up:Up      |
| 9  | Easy7200                  | N9K-4-BGW-Ethernet1/21--N9K-3-BGW-Ethernet1/21    | int_intra_fabric_num_link_11_1    | Link Present     | Up:Up       | Up:Up      |
| 10 | Easy7200                  | N9K-4-BGW-Ethernet1/48--N9K-1-Spine-Ethernet1/42  | int_intra_fabric_num_link_11_1    | Link Present     | Up:Up       | Up:Up      |

## Deleting Multi-Site IFCs

1. Navigate to the Links tab, select the IFCs to be deleted and click the Delete icon as shown below.
2. Perform a Save & Deploy in the MSD fabric to complete deletion.



**Note** If auto configuration of IFCs is enabled in the MSD fabric settings, then performing a Save & Deploy may regenerate the IFC intent.

If all or large number of IFCs are to be deleted, then temporarily change the BGW deploy mode to Manual setting before performing Save & Deploy.

- **Deleting IFC in a non-Nexus Switch:** If you delete the last IFC in a non-Nexus switch, the switch is removed from the topology. From Cisco DCNM Release 11.2(1), you can remove non-Nexus switches and neighbor switches like a physical switch from the **Tabular view** window or from the fabric topology window by right-clicking the switch and choosing **Discovery > Remove from fabric** in the drop-down menu.
- **Removing a fabric from an MSD fabric:** Before removing a fabric from an MSD fabric, remove all the multisite overlays from all BGWs in that fabric. Otherwise, you will not be able to remove the fabric. After the following save and deploy in the easy fabric, all the multisite configurations, such as IFC, multisite loopback address configured in MSD are removed from BGWs.
- **Device role change:** You can change the device role from Border to Border Gateway, but the role change from Border Gateway to Border is allowed only if there are no multisite IFCs or overlays deployed on the device.

## Creating and Deploying Networks and VRFs in the MSD Fabric

Networks and VRFs can be created from the MSD context in the Networks and VRF page, these can be deployed on BGW nodes for all member fabrics of that MSD.

The following screenshots show how to select networks and deploy them. From the MSD fabric context, any device can be selected for network or VRF deployment. However, networks or VRFs can be deployed only on BGWs from the MSD context in the network deployment screen. The leaf deployment can be done from the fabric context or from the Fabric Builder context.

Fabric Selection > Network / VRF Selection > Network / VRF Deployment > VRF View | Continue

Step 1: navigate to the network deployment page of the relevant MSD fabric

Fabric Selected: MSD-Fabric

Step 3: press the continue button to go the deployment page

Selected 1 / Total 1

| Network Name  | Network ID | VRF Name    | IPv4 Gateway/Subnet | IPv6 Gateway/Prefix | Status | VLAN ID |
|---|------------|-------------|---------------------|---------------------|--------|---------|
| <input checked="" type="checkbox"/> MyNetwork_30000 | 30000      | MyVRF_50000 | 11.0.0.1/24         |                     | NA     | 111     |

Step 2: select NW(s) to be deployed

Network Extension Attachment - Attach extensions for given switch(es)

Fabric Name: MSD-Fabric

Deployment Options

Select the row and click on the cell to edit and save changes

| Switch  | VLAN | Extend    | Interfaces                        | CLI Freeform    | Status      |
|---|------|-----------|-----------------------------------|-----------------|-------------|
| <input checked="" type="checkbox"/> NSK-3-BGW | 111  | MULTISITE | Applicable to BGW Leaf - VPC only | Freeform config | DEPLOYED    |
| <input checked="" type="checkbox"/> NSK-4-BGW | 111  | MULTISITE | Applicable to BGW Leaf - VPC only | Freeform config | OUT-OF-SYNC |

Step 1: Check this box to multiple BGWs, then use GUI to select one or more BGWs, then this pop up will appear

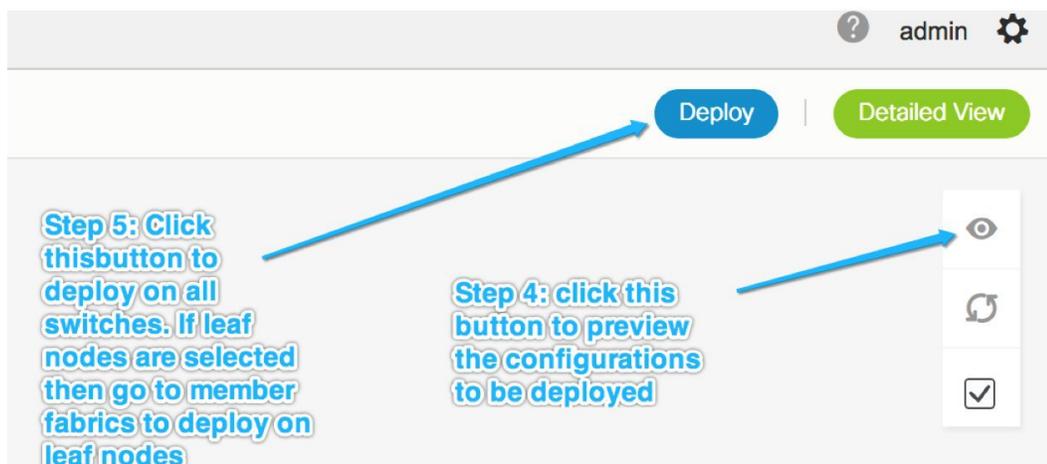
Step 2: Check this boxes to select BGWs on which to deploy the NW(s)

Step 3: click this to move to deployment screen, repeat till all nodes on which NW(s) are to be deployed

Save

Deploy | Detailed View

Undiscovered



### Deploying Networks with a Layer 3 Gateway on a BGW

Perform the following steps:



**Note** Selecting an interface to deploy SVI is only available on vPC BGW setups. This is a device limitation not a DCNM limitation.

1. In order to deploy a network with a Layer 3 gateway on a Border device (Border, Border spine, Border Gateway, Border Gateway spine), perform these steps.

When creating the network, check the **Enable L3 gateway on Border** check box, as shown in the figure below. Note that this is a network wide setting, so whenever this network is deployed on the Border device, the Layer 3 gateway will be deployed. If this is required on only a subset of the Borders, then a custom template is required.

When deploying the network on the Border device, select the interface(s) in the **Interface** column in case of vPC BGW.

Just like the leaf switch, the candidate ports should have **int\_trunk\_host\_policy\_11\_1**, otherwise they will not be in the interface list.

The interface policy can be modified through the **Control > Interfaces** tab.

**Networks**

- Network Name
- MyNetwork\_30000
- MyNetwork\_30001

**Edit Network**

**Network Information**

- \* Network ID: 30001
- \* Network Name: MyNetwork\_30001
- \* VRF Name: MyVRF\_50000
- Layer 2 Only:
- \* Network Template: Default\_Network\_Universal
- \* Network Extension Template: Default\_Network\_Extension\_Univer
- VLAN ID:

**Network Profile**

**General**

**Advanced**

- setting in advanced tab
- DHCPv4 Server 2:  ? DHCP Relay IP
- DHCPv4 Server VRF:  ?
- Loopback ID for DHCP Relay interface:  ?
- Routing Tag: 12345 ? [0-4294967295]
- TRM Enable:  ? Enable Tenant Routed Multicast
- L2 VNI Route-Target Both Enable:  ?
- Enable L3 Gateway on Border:  ? check this box when creating a network, this is a per network setting

**Save** **Cancel**

- When deploying the network on the vPC pair of BGWs, select the interface(s) in the Interfaces column. Only vPC port channel interfaces are the candidate interfaces.

## Network Extension Attachment - Attach extensions for given switch(es)



Fabric Name: MSD

## Deployment Options

*Select the row and click on the cell to edit and save changes*

| MyNetwork_30001                     |          |      |           |                     |                 |          |  |
|-------------------------------------|----------|------|-----------|---------------------|-----------------|----------|--|
| <input type="checkbox"/>            | Switch ▲ | VLAN | Extend    | Interfaces          | CLI Freeform    | Status   |  |
| <input checked="" type="checkbox"/> | BL-1     | 2300 | MULTISITE | ... Port-channel500 | Freeform config | DEPLOYED |  |
| <input checked="" type="checkbox"/> | BL-2     | 2300 | MULTISITE | ... Port-channel500 | Freeform config | DEPLOYED |  |

Save

## Interfaces



| <input type="checkbox"/>            | Interface/Ports ▲ | Port Type |
|-------------------------------------|-------------------|-----------|
| <input checked="" type="checkbox"/> | Port-channel500   | trunk     |

Save

## Deploying a Legacy Site BGW (vPC-BGWs)

The recommended way of integrating non-VXLAN BGP EVPN (legacy) and VXLAN BGP EVPN fabrics is by using a pair of VPC BGWs. For more information about this method, see [NextGen DCI with VXLAN EVPN Multi-Site Using vPC Border Gateways White Paper](#).

The vPC BGW method replaces the Pseudo-Border Gateway method recommended in the DCNM release 11.1(1).

In this section, the tasks from the white paper that can be accomplished by DCNM are explained with an example topology.

### Prerequisites

- Legacy network is already setup by a method. This is out of the scope for this document.
- Familiarity with fabric creation and Multi-Site use case.

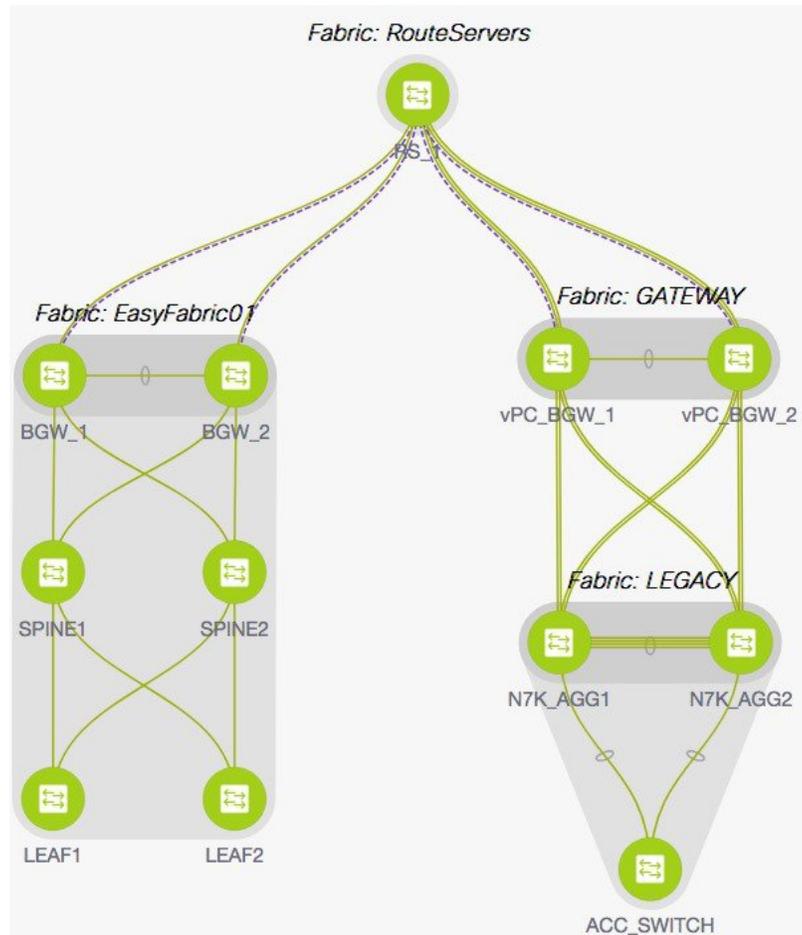
### Tasks Overview

The following information is covered as part of this section:

1. Fabrics to be created using DCNM:
  - a. VXLAN fabric with vPC border gateways.
  - b. Easy Fabric for VXLAN.
  - c. External fabric for Route Server. Note that this fabric is optional if you are using Direct to BGWs topology.
  - d. External fabric to monitor the legacy devices.
  - e. MSD fabric as a container for all fabrics.
2. vPC connection from vPC BGWs towards the legacy site. It is expected that vPC from legacy towards BGWs is done out of band.
3. Multi-Site underlay eBGP inter-fabric connection (IFC) creation.
4. Multi-Site overlay eBGP IFC creation.

### Topology Overview

Let us look at the example topology.



This topology contains the following five fabrics:

### 1. GATEWAY

This fabric is created for the vPC border gateways.

This fabric is an Easy fabric without any Spine nodes and it is set up as a regular Easy fabric with the following characteristics:

- Under the **Replication** tab, the **Replication Mode** is set to **Ingress**.
- The vPC Border gateways role are set as BGW.
- The IFC create method will be set to Manual or auto configuration as per user preference.
- Gateway fabric has a vPC interface configuration towards the Legacy Fabric.
- A member fabric of MSD.
- Save and deploy operation is performed in Easy fabric and MSD fabric.

### 2. LEGACY

This fabric is created for the Legacy network. The fabric type is External and could be kept in the monitor mode. Fully configured devices are imported into this fabric as shown in External Fabric procedure.

### 3. EasyFabric01

This represents a fully functional VXLAN fabric. The Border Gateway switches of this fabric are connected via IFC's to Route Servers or Direct to BGWs of Legacy fabric as per your topology. Both models are supported as shown in the Multi-Site use case.

### 4. RouteServers

In this topology, Centralized to Route Server topology is used. Typically, there would be more than one Route Server for redundancy reasons. This fabric is of type External as shown in the Multi-Site use case.

### 5. MSD

The MSD fabric is created to configure the base multi-site for the member fabrics. All the above four fabrics are imported into the MSD fabric for the BGW base. Optionally, you can enable auto-configuration of all underlay and overlay IFCs.

## Configuring vPC from vPC Border Gateways to Legacy Network

In the **Manage Interfaces** window for the **GATEWAY** fabric, click the **Add (+)** icon and enter the information for the fields as shown in the following image. From the **Policy** drop-down list, select the vPC policy and fill in the fields for your topology.

Edit Configuration
✕

Name: vPC\_BGW\_2~vPC\_BGW\_1:vPC1

Policy:

Note : PeerOne = vPC\_BGW\_2 & PeerTwo = vPC\_BGW\_1

General

|                           |                                       |  |
|---------------------------|---------------------------------------|--|
| Peer-1 Port-Channel ID    | <input type="text" value="1"/>        | <small>Peer-1 VPC port-channel number (Min:1, Max:4096)</small>                        |
| Peer-2 Port-Channel ID    | <input type="text" value="1"/>        | <small>Peer-2 VPC port-channel number (Min:1, Max:4096)</small>                        |
| Peer-1 Member Interfaces  | <input type="text" value="E1/21-24"/> | <small>A list of member interfaces for Peer-1 [e.g. e1/5,eth1/7-9]</small>             |
| Peer-2 Member Interfaces  | <input type="text" value="E1/21-24"/> | <small>A list of member interfaces for Peer-2 [e.g. e1/5,eth1/7-9]</small>             |
| * Port Channel Mode       | <input type="text" value="active"/>   | <small>Channel mode options: on, active and passive</small>                            |
| * Enable BPDU Guard       | <input type="text" value="no"/>       | <small>Enable spanning-tree bpduguard</small>  |
| Enable Port Type Fast     | <input checked="" type="checkbox"/>   | <small>Enable spanning-tree edge port behavior</small>                                 |
| * MTU                     | <input type="text" value="jumbo"/>    | <small>MTU for the Port Channel</small>  |
| * Peer-1 Trunk Allowed... | <input type="text" value="all"/>      | <small>Allowed values: 'none', 'all', or vlan ranges (ex: 1-200,500-2000,3000)</small> |
| * Peer-2 Trunk Allowed... | <input type="text" value="all"/>      | <small>Allowed values: 'none', 'all', or vlan ranges (ex: 1-200,500-2000,3000)</small> |
| Peer-1 PO Description     | <input type="text"/>                  | <small>Add description to Peer-1 VPC port-channel (Max Size 254)</small>               |

After entering all the information, click **Preview** to preview the configurations that are deployed, and then click **Deploy**.

### Multi-Site Underlay eBGP IFC Creation

The Multi-Site underlay configuration is same as MSD shown in the Multi-Site use case. Choose GUI or autoconfiguration based method to create IFCs to the Core router or directly to BGW of other fabric, as per your topology.

In this topology, vPC Border Gateways are physically connected to Route Server (RS1), one MS underlay IFC is configured from each BGW (in GATEWAY and EasyFabric01) to RS1. Both methods are detailed in the Multi-Site use case.

### Configuring Multi-Site Overlay IFCs

Multi-Site overlay IFCs need to be created between vPC BGWs to either a centralized route server or Direct to each BGW in **EasyFabric01**. In the example topology, there is one Overlay IFC from each BGW to RS1.

The summary of the IFCs for this topology are shown in the following image.

The screenshot shows the 'Fabric Builder: MSD' interface with the 'Links' tab selected. A table lists 8 IFCs with columns for Fabric Name, Name, Policy, Info, Admin State, and Oper State.

|   | <input type="checkbox"/> | Fabric Name                 | Name                                      | Policy                            | Info             | Admin State | Oper State |
|---|--------------------------|-----------------------------|---|-----------------------------------|------------------|-------------|------------|
| 1 | <input type="checkbox"/> | EasyFabric01<->RouteServers | BGW_1-loopback0—RS_1-Loopback0            | ext_evprn_multisite_overlay_setup | Neighbor Missing | --          | --         |
| 2 | <input type="checkbox"/> | EasyFabric01<->RouteServers | BGW_2-loopback0—RS_1-Loopback0            | ext_evprn_multisite_overlay_setup | Neighbor Missing | --          | --         |
| 3 | <input type="checkbox"/> | GATEWAY<->RouteServers      | vPC_BGW_1-loopback0—RS_1-Loopback0        | ext_evprn_multisite_overlay_setup | Neighbor Missing | --          | --         |
| 4 | <input type="checkbox"/> | GATEWAY<->RouteServers      | vPC_BGW_2-loopback0—RS_1-Loopback0        | ext_evprn_multisite_overlay_setup | Neighbor Missing | --          | --         |
| 5 | <input type="checkbox"/> | EasyFabric01<->RouteServers | BGW_1-Ethernet4/3—RS_1-Ethernet5/5        | ext_multisite_underlay_setup_11_1 | Link Present     | Up:Up       | Up:Up      |
| 6 | <input type="checkbox"/> | EasyFabric01<->RouteServers | BGW_2-Ethernet1/51—RS_1-Ethernet5/6       | ext_multisite_underlay_setup_11_1 | Link Present     | Up:Up       | Up:Up      |
| 7 | <input type="checkbox"/> | GATEWAY<->RouteServers      | vPC_BGW_1-Ethernet1/14—RS_1-Ethernet5/7/2 | ext_multisite_underlay_setup_11_1 | Link Present     | Up:Up       | Up:Up      |
| 8 | <input type="checkbox"/> | GATEWAY<->RouteServers      | vPC_BGW_2-Ethernet1/13—RS_1-Ethernet5/7/3 | ext_multisite_underlay_setup_11_1 | Link Present     | Up:Up       | Up:Up      |

## 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

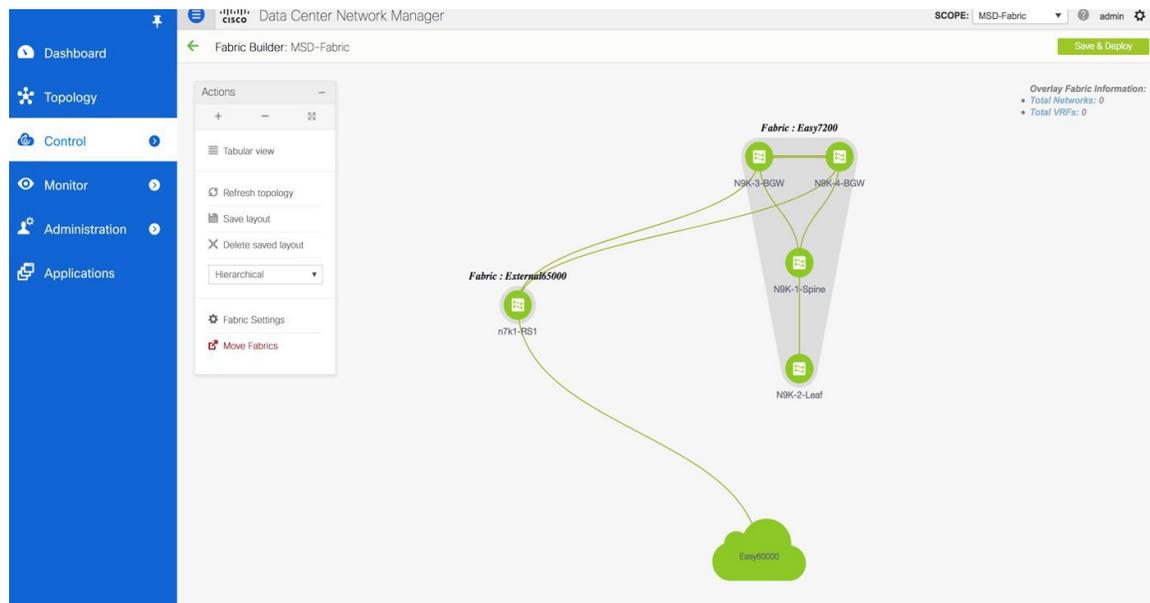
### Multi-Site Fabric Base Configurations – Box Topology

In the Easy7200 fabric, N9K-3-BGW and N9K-4-BGW are connected to each other over two physical interfaces, and the BGWs do not form a vPC pair. Such a topology is called a Box topology. An IBGP session

is configured on each physical connection. One connection is between the Eth1/21 interfaces, and the other is between the Eth1/22 interfaces.

## IBGP Configuration for the Box Topology in the Easy7200 Fabric

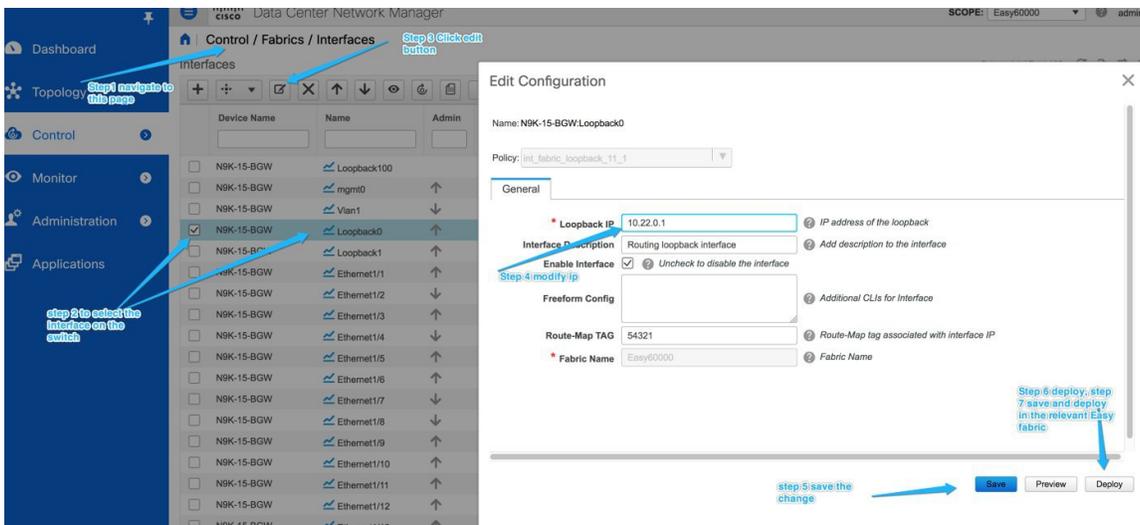
The following configuration is generated on each of the nodes if the fabric has numbered interfaces. In case the fabric interfaces are unnumbered, then the IBGP session is formed via the loopback0 address.



| N9K-BGW-3  | N9K-BGW-4   |
|--|---|
| <pre>router bgp 7200   neighbor 10.4.0.17   remote-as 7200   update-source ethernet1/22   address-family ipv4 unicast   next-hop-self</pre>                    | <pre>router bgp 7200   neighbor 10.4.0.18   remote-as 7200   update-source Ethernet1/22   address-family ipv4 unicast   next-hop-self</pre>                   |
| <pre>router bgp 7200   neighbor 10.4.0.13   remote-as 7200   update-source ethernet1/21   address-family ipv4 unicast   next-hop-self</pre>                    | <pre>router bgp 7200   neighbor 10.4.0.14   remote-as 7200   update-source Ethernet1/21   address-family ipv4 unicast   next-hop-self</pre>                   |
| <pre>interface ethernet1/22   evpn multisite dci-tracking   no switchport   ip address 10.4.0.18/30   description   connected-to-N9K-4-BGW--Ethernet1/22</pre> | <pre>interface Ethernet1/22   evpn multisite dci-tracking   no switchport   ip address 10.4.0.17/30   description   connected-to-N9K-3-BGW-Ethernet1/22</pre> |

| N9K-BGW-3   | N9K-BGW-4   |
|---|---|
| <pre>interface ethernet1/21   evpn multisite dci-tracking   no switchport   ip address 10.4.0.14/30   description   connected-to-N9K-4-BGW-Ethernet1/21</pre> | <pre>interface Ethernet1/21   evpn multisite dci-tracking   no switchport   ip address 10.4.0.13/30   description   connected-to-N9K-3-BGW-Ethernet1/21</pre> |

### Changing loopback0 Policy to Modify IP Address



## Route Server Configuration

The route server overlay and base configurations are only deployed if the external fabric is not in Monitor mode.



**Note** When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on its switches. Refer the *Creating an External Fabric topic* in the *Control* chapter for details.

**Route Server Base Configuration** - These are one time deployed on the route server and may be edited or deleted via the corresponding policy. The router server overlay and base configurations are only deployed if the external fabric is not in Monitor mode.

| Configuration  | Description |
|--|-------------|
| <pre>route-map unchanged permit 10   set ip next-hop unchanged</pre> | —           |

| Configuration   | Description  |
|---|--|
| <pre>router bgp 65000   address-family ipv4 unicast     network 1.1.1.1/32</pre>  | <p>The network command to redistribute the BGP peering address of RS1 to the eBGP underlay sessions so that BGWs know how to reach RS.</p> <p>If operator is using a different method to distribute the route server peering address to BGW, then this is not needed</p>   |
| <pre>interface ethernet1/22   evpn multisite dci-tracking   no switchport   ip address 10.4.0.18/30   description connected-to-N9K-4-BGW--Ethernet1/22</pre>  | <pre>interface Ethernet1/22   evpn multisite dci-tracking   no switchport   ip address 10.4.0.17/30   description connected-to-N9K-3-BGW-Ethernet1/22</pre>  |
| <pre>template peer OVERLAY-PEERING   update-source loopback0   ebgp-multihop 5   address-family l2vpn evpn     route-map unchanged out   address-family l2vpn evpn     retain route-target all     send-community     send-community extended</pre> | <p>The knob in the external fabric controls if send community is sent in the form shown here, or as send-community both.</p> <p>If this form causes a persistent CC difference, then edit the policy on the device in the external fabric as shown in the Deploying the Send-Community Both Attribute section below.</p> |

## Multi-Site Overlay IFC Configuration

In the reference topology, there are two BGWs in the Easy7200 fabric. Each BGW forms a BGP overlay connection with the route server.

| BGW  | Route Server   |
|--|--|
| <pre>router bgp 7200   neighbor 1.1.1.1     remote-as 65000   update-source loopback0   ebgp-multihop 5   peer-type fabric-external   address-family l2vpn evpn     send-community     send-community extended   rewrite-evpn-rt-asn</pre> | <pre>router bgp 65000   neighbor 10.2.0.1     remote-as 7200   inherit peer OVERLAY-PEERING   address-family l2vpn evpn     rewrite-evpn-rt-asn router bgp 65000   neighbor 10.2.0.2     remote-as 7200   inherit peer OVERLAY-PEERING   address-family l2vpn evpn     rewrite-evpn-rt-asn</pre> |

See below for the configurations generated on the BGW and the route server.

## Multi-Site Underlay IFC Configuration – Out-of-Box Profiles

The following table shows the Multi-Site IFC configuration deployed by DCNM with the out-of-the box profiles. If the IFC is between two VXLAN fabrics, then both sides have the BGW configurations shown below.

| BGW Configuration  | Core Router Configuration   |
|--|---|
| <pre>router bgp 7200   neighbor 10.10.1.6   remote-as 65000   update-source ethernet1/47   address-family ipv4 unicast     next-hop-self</pre> | <pre>router bgp 65000   neighbor 10.10.1.5   remote-as 7200   update-source ethernet7/4/1   address-family ipv4 unicast     next-hop-self</pre> |
| <pre>interface ethernet1/47   mtu 9216   no shutdown   no switchport   ip address 10.10.1.5/30 tag 54321   evpn multisite dci-tracking</pre>   | <pre>interface ethernet7/4/1   mtu 9216   no shutdown   no switchport   ip address 10.10.1.6/30 tag 54321</pre>                                 |

The tag 54321 attached to the IP address is not required for correct functioning and will be removed in subsequent releases. It is benign.