Cisco DCNM LAN Fabric Configuration Guide, Release 11.5(3a)

First Published: 2021-11-24

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

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

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

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

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

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

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

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

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/c/en/us/about/legal/trademarks.html. 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)

© 2021 Cisco Systems, Inc. All rights reserved.

CONTENTS

Full Cisco Trademarks with Software License ?

CHAPTER 1 **Overview** 1

Cisco Data Center Network Manager 1

REST API Tool 2

CHAPTER 2 **New and Changed Information** 7

New and Changed Information in Cisco DCNM, Release 11.5(1) 7

CHAPTER 3 **Dashboard** 11

Dashboard 11

Dashlets 12

CHAPTER 4 **Topology** 17

Topology 17

Status 17

Scope 18

Searching 19

Quick Search 19

Host name (vCenter) 19

VM name (OpenStack) 19

Host IP 19

Host MAC 20

Multicast Group 20

Redirected Flows 20

VXLAN ID (VNI) 22

VLAN 22
VXLAN OAM 23
Show Panel 24
Layouts 25
Zooming, Panning, and Dragging 25
Switch Slide-Out Panel 27
Beacon 27
Tagging 27
More Details 27
Link Slide-Out Panel 29
24-Hour Traffic 29
vCenter Compute Visualization 29
Support for Cisco UCS B-Series Blade Servers 30
Enabling vCenter Compute Visualization 32
Using vCenter Compute Visualization 34
Troubleshooting vCenter Compute Visualization 39
Container Orchestrator 40
Using the UI Controls on Container Orchestrator Visualization 42
OpenStack Workload Visibility 46
OpenStack Topology Scale 46
Notifications and Triggers for OpenStack 47
Using OpenStack Visualizer 47
Viewing VMs in OpenStack Clusters 49

CHAPTER 5 **Control** 51
Fabrics 51
VXLAN BGP EVPN Fabrics Provisioning 51
Creating a New VXLAN BGP EVPN Fabric 55
Adding Switches to a Fabric 73
Pre-provisioning Support in DCNM 11 86
Precision Time Protocol for Easy Fabric 100
Support for Super Spine Role in DCNM 101
Changing the TCAM Configuration on a Device 107
Preselecting Switches as Route-Reflectors and Rendezvous-Points 108

Adding a vPC L3 Peer Keep-Alive Link 109
Changing the Local Authentication to AAA Authentication for Switches in a Fabric 112
IPv6 Underlay Support for Easy Fabric 114
Brownfield Deployment-Transitioning VXLAN Fabric Management to DCNM 114
Configuring Fabrics with eBGP Underlay 114
Creating an External Fabric 114
Discovering New Switches 126
Adding non-Nexus Devices to External Fabrics 131
Pre-provisioning a Device 135
Pre-provisioning an Ethernet Interface 139
Creating a vPC Setup 141
Undeploying a vPC Setup 146
Multi-Site Domain for VXLAN BGP EVPN Fabrics 146
Support for CloudSec in Multi-Site Deployment 174
Removing a Fabric From an MSD 178
Moving a Standalone Fabric (With Existing Networks and VRFs) to an MSD Fabric 179
Managing Switches Using LAN Classic Templates 180
Creating a LAN Classic Fabric 180
Adding Switches to LAN Classic Fabric 184
Creating a Fabric Group and Associating Member Fabrics 185
Support for Inter-Fabric Connection in LAN Classic Fabric Template 186
Inband Management in External Fabrics and LAN Classic Fabrics 186
Inband POAP Management in External Fabrics and LAN Classic Fabrics 187
Precision Time Protocol for External Fabrics and LAN Classic Fabrics 188
Sync up Out-of-Band Switch Interface Configurations with DCNM 190
Syncing up Switch Interface Configurations to DCNM 191
MACsec Support in Easy Fabric and eBGP Fabric 194
Enabling MACsec 194
Disabling MACsec 195
Overview of Tenant Routed Multicast 196
Overview of Tenant Routed Multicast with VXLAN EVPN Multi-Site 196
Tenant Routed Multicast with VXLAN EVPN Multi-Site Operations 197
Configuring TRM for Single Site Using Cisco DCNM 197
Configuring TRM for Multi-Site Using Cisco DCNM 201

SSH Key RSA Handling 204
Switch Operations 205
Running EXEC Mode Commands in DCNM 208
Fabric Multi Switch Operations 209
Tabular View - Switches 209
Tabular View - Links 212
Creating Intra-Fabric Links 212
Creating Inter-Fabric Links 217
Exporting Links 221
Importing Links 222
Viewing Details of Fabric Links 222
Viewing the Traffic Details of Fabric Links 223
Symmetric Automatic VRF Lite 224
Layer 3 Port Channels 225
Configuring Layer 3 Port Channel on Interfaces 226
Configuring Layer 3 Port Channel on Interfaces for IOS XE Devices 226
Deploying Policies on Physical Interfaces for non-Nexus Devices 227
Configuring Layer 3 Port Channel on Subinterfaces 228
Configuring Layer 3 Port Channel for Inter-fabric Connectivity 229
Tabular View - Operational View 230
Viewing the Operational Status 231
Viewing Logical Links 232
Viewing Alerts and Event Notifications 232
Support for ToR Switches 232
vPC Fabric Peering 233
Creating a Virtual Peer Link 235
Converting a Physical Peer Link to a Virtual Peer Link 238
Converting a Virtual Peer Link to a Physical Peer Link 239
Advertising PIP on vPC 240
Viewing and Editing Policies 241
Viewing Policies 242
Adding a Policy 244
Deploying Policies 245
Editing a Policy 246

Current Switch Configuration 247
Retrieving the Authentication Key 247
Custom Maintenance Mode Profile Policy 249
Creating and Deploying a Custom Maintenance Mode Profile Policy 250
Deleting a Custom Maintenance Mode Profile Policy 251
Return Material Authorization (RMA) 253
Prerequisites 253
Guidelines and Limitations 253
POAP RMA Flow 253
Manual RMA Flow 256
RMA for User with Local Authentication 258
Interfaces 258
Adding Interfaces 263
Breakout 265
Editing Interfaces 265
Deleting Interfaces 267
Shutting Down and Bringing Up Interfaces 268
Viewing Interface Configuration 268
Rediscovering Interfaces 268
Viewing Interface History 269
Deploying Interface Configurations 269
Creating External Fabric Interfaces 270
Interface Groups 270
Creating and Deploying Networks and VRFs 276
Viewing Networks and VRFs for a Fabric 277
Creating Networks for the Standalone Fabric 277
Editing Networks for the Standalone Fabric 282
Creating VRFs for the Standalone Fabric 283
Editing VRFs for the Standalone Fabric 288
Deploying Networks for the Standalone and MSD Fabrics 289
Deploying VRFs for the Standalone and MSD Fabrics 298
Undeploying Networks for the Standalone Fabric 304
Undeploying VRFs for the Standalone Fabric 305
Deleting Networks and VRFs 305

Configuring Multiple VLAN IDs to a Single VNI 306
Enhanced Role-based Access Control in Cisco DCNM 307
Device-upg-admin Role 307
Access-admin Role 307
Network-Operator Role 308
Network-Stager Role 308
Viewing Policy Change History 309
Freezing Fabrics in Cisco DCNM 310
Fabric Backup and Restore 311
Backing Up Fabrics 311
Restoring Fabrics 315
Restoring a Switch 321
Deleting a VXLAN BGP EVPN Fabric 322
Post DCNM 11.5(1) Upgrade for VXLAN BGP EVPN, External, and MSD Fabrics 323
Changing ISIS Configuration from Level 1 to Level 2 323
Configuration Compliance in DCNM 324
Configuration Compliance in External Fabrics 333
Resolving Diffs for Case Insensitive Commands 338
Enabling Freeform Configurations on Fabric Switches 343
VMM Workload Automation 349
Overview of Network Objects in vCenter 349
How VMM Workload Automation Works 351
Configuration Files for VMM Workload Automation 352
Installing and Starting the VMM Workload Automation Module 355
Additional Functionalities Using REST APIs 357
Events in vCenter 358
Management 359
Resources 359
Allocating a Resource 360
Releasing a Resource 362
Adding, Editing, Re-Discovering and Removing VMware Servers 362
Adding a Virtual Center Server 362
Deleting a VMware Server 363
Editing a VMware Server 363

Rediscovering a VMware Server 363
Container Orchestrator 364
Adding Container Orchestrator 365
Deleting Container Orchestrator 368
Editing Container Orchestrator 368
Rediscover Kubernetes Cluster 368
OpenStack Visualizer 369
Adding OpenStack Cluster 370
Editing OpenStack Cluster 371
Deleting OpenStack Cluster 372
Rediscovering OpenStack Cluster 372
Template Library 372
Template Structure 374
Template Format 374
Template Variables 381
Variable Meta Property 383
Variable Annotation 389
Templates Content 393
Advanced Features 395
Report Template 397
Adding a Template 410
Modifying a Template 411
Copying a Template 411
Deleting a Template 412
Importing a Template 412
Exporting a Template 413
Image Management 413
Smart Image Management 414
Image Upload 415
Deleting an Image 417
Install & Upgrade 417
Upgrade History 417
Switch Level History 427
Packages 428

Installing Packages and Patches 428
Uninstalling Packages and Patches 429
Activating Packages and Patches 430
Deactivate 430
Image Management Policies 430
Adding an Image Management Policy 431
Deleting an Image Management policy 433
Endpoint Locator 434
Layer 4-Layer 7 Service 434
Cross Site Scripting (XSS) threat and mitigation 434
Cross Site Scripting (XSS) threat and Handling of special Characters in Policy Fields 434

CHAPTER 6 **Monitor** 437
Inventory 437
Viewing Inventory Information for Switches 437
Viewing System Information 442
Hosts 442
Capacity 443
Features 443
VXLAN 444
VLAN 444
Switch Modules 445
FEX 445
VDCs 447
Viewing Inventory Information for Modules 454
Viewing Inventory Information for Licenses 455
Monitoring Switch 456
Viewing Switch CPU Information 456
Viewing Switch Memory Information 456
Viewing Switch Traffic and Errors Information 456
Viewing Switch Temperature 457
Enabling Temperature Monitoring 458
Viewing Accounting Information 458
Viewing Events Information 458

Monitoring LAN 459
Monitoring Performance Information for Ethernet 459
Monitoring ISL Traffic and Errors 460
Monitoring a vPC 461
Monitoring vPC Performance 462
Endpoint Locator 463
Alarms 463
Viewing Alarms and Events 464
Monitoring and Adding Alarm Policies 464
Activating Policies 467
Deactivating Policies 468
Importing Policies 468
Exporting Policies 468
Editing Policies 468
Deleting Policies 469
Enabling External Alarms 469
Configuration Compliance Alarms 469
Endpoint Locator Alarms 472
Health Monitor Alarms 475

CHAPTER 7 **Administration** 479
DCNM Server 479
Starting, Restarting, and Stopping Services 479
Customization 481
Network Preferences 482
Viewing Log Information 483
Server Properties 484
Modular Device Support 484
Native HA 485
Multi Site Manager 486
Device Connector 489
NX-API Certificate Management for Switches 492
Uploading the certificates on DCNM 494
Installing Certificates on Switches 494

Unlinking and Deleting certificates 495
Troubleshooting NX API Certificate Management 496
Backing up DCNM 496
Creating a Backup 497
Modifying a Backup 498
Deleting a Backup 499
Job Execution Details 499
Manage Licensing 500
Managing Licenses 500
License Assignments 500
Smart License 506
Switch Smart License 509
Server License Files 510
Switch Features—Bulk Install 511
Application Licenses 513
Management Users 515
Remote AAA 515
Local 515
Radius 516
TACACS+ 516
Switch 516
LDAP 517
Managing Local Users 519
Adding Local Users 519
Deleting Local Users 520
Editing a User 520
User Access 520
Managing Clients 521
Performance Setup 522
Performance Setup LAN Collections 522
Event Setup 523
Viewing Events Registration 523
Notification Forwarding 524
Adding Notification Forwarding 524

Removing Notification Forwarding 525
Event Suppression 526
Add Event Suppression Rules 526
Delete Event Suppression Rule 527
Modify Event Suppression Rule 527
Credentials Management 527
LAN Credentials 528
Credentials Management with Remote Access 530

PART I **Applications** 537

CHAPTER 8 **Applications Framework** 539
Cisco DCNM in Unclustered Mode 539
Cisco DCNM in Clustered Mode 540
Requirements for Cisco DCNM Clustered Mode 541
Installing a Cisco DCNM Compute 542
Networking Policies for OVA Installation 542
Enabling the Compute Cluster 544
Managing Application Network Pools 545
Adding Computes into the Cluster Mode 546
Transitioning Compute Nodes 548
Transitioning Compute nodes from VM to Service Engine 548
Transitioning Compute nodes from Service Engine to VM 549
Preferences 549
Telemetry Network and NTP Requirements 550
Installing and Deploying Applications 551
Application Framework User Interface 554
Catalog 555
Compute 555
Preferences 557
Failure Scenario 557
Compute Node Disaster Recovery 558

CHAPTER 9 **Endpoint Locator** 559

Endpoint Locator 559
Configuring Endpoint Locator 560
Configuring Endpoint Locator in DCNM High Availability Mode 569
Configuring Endpoint Locator in DCNM Cluster Mode 570
Configuring Endpoint Locator for External Fabrics 572
Configuring Endpoint Locator for eBGP EVPN Fabrics 572
EPL Connectivity Options 575
Disabling Endpoint Locator 579
Troubleshooting Endpoint Locator 579
Monitoring Endpoint Locator 583
Endpoint Locator Dashboard 583
Endpoint History 588
Endpoint Search 591
Endpoint Life 592

CHAPTER 10 **IPAM Integrator** 595
Catalog 595
IPAM Integrator 596
Accessing IPAM Integrator 596
Viewing Network IP Scope 597
Viewing Statistics for the Subnet Utilization 598
Viewing IP Allocation for Hosts 599
Viewing Conflicting Networks 600

CHAPTER 11 **Health Monitor** 601
Catalog 601
Health Monitor 602
Alerts 602
Service Utilization 604
Compute Utilization 607

CHAPTER 12 **PTP Monitoring** 609
Catalog 609
PTP Monitoring 610

CHAPTER 13 **Programmable Reports** 613

Catalog 613

Programmable Report 614

Creating a Report Job 616

Viewing a Report Job 618

Downloading Report Information 620

Deleting a Report 621

Comparing Reports 621

Deleting a Report Job 623

Editing a Report Job 623

Rerunning a Report Job 624

Displaying Report Job History 625

Downloading Report Job Information 626

Report Purging 626

CHAPTER 14 **ServiceNow Integration** 629

DCNM Integration with ServiceNow 629

Guidelines and Limitations of DCNM Integration with ServiceNow 630

Installing and Configuring the Cisco DCNM Application on ServiceNow 631

Viewing the Dashboard 634

Contact Us 638

Troubleshooting DCNM Integration with ServiceNow 638

PART II **Easy Provisioning of VXLAN BGP EVPN Fabrics** 641

CHAPTER 15 **Managing a Greenfield VXLAN BGP EVPN Fabric** 643

VXLAN BGP EVPN Fabrics Provisioning 643

Creating a New VXLAN BGP EVPN Fabric 646

Adding Switches to a Fabric 665

Discovering New Switches 665

Discovering Existing Switches 673

VXLAN EVPN Deployment with eBGP EVPN 678

Creating a eBGP New VXLAN EVPN with eBGP-based Underlay 678

Deploying Fabric Underlay eBGP Policies 691
Deploying Fabric Overlay eBGP Policies 692
Deploying Spine Switch Overlay Policies 692
Deploying Leaf Switch Overlay Policies 693

CHAPTER 16 **Managing a Brownfield VXLAN BGP EVPN Fabric** 695
Overview 695
Prerequisites 696
Guidelines and Limitations 696
Fabric Topology Overview 698
DCNM Brownfield Deployment Tasks 699
Verifying the Existing VXLAN BGP EVPN Fabric 699
Creating a VXLAN BGP EVPN Fabric 702
Adding Switches and Transitioning VXLAN Fabric Management to DCNM 716
Verifying the Import of the VXLAN BGP EVPN Fabric 729
Verifying VXLANs and Commands on Switches 729
Verifying Resources 733
Verifying Networks 734
Configuration Profiles Support for Brownfield Migration 737
Migrating a Bottom-Up VXLAN Fabric to DCNM 737
Resolving Config Compliance Error on Switches with Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) Images 746
Modifying VLAN Names in a Switch with Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) Images 750
Changing a Brownfield Imported BIDIR Configuration 752
Manually Adding PIM-BIDIR Configuration for Leaf or Spine Post Brownfield Migration 753
Migrating an MSD Fabric with Border Gateway Switches 753

CHAPTER 17 **Configuring a VXLANv6 Fabric** 757
Overview 757
Creating a VXLAN Fabric with IPv6 Underlay 758

CHAPTER 18 **Auto-Provisioning ToR Switches Attached to VXLAN VTEPs** 763
Overview 763

Supported Topologies for ToR Switches 763
Configuring ToR Switches 769
Deploying Networks on ToR Switches 775

PART III External/WAN Layer 3 Connectivity for VXLAN BGP EVPN Fabrics 779

CHAPTER 19 VRF Lite 781
Prerequisites and Guidelines 781
Sample Scenarios 784
VRF Lite Through the DCNM GUI – From a BGW Device to a Nexus 7000 Series Edge Router 785
VRF Lite Through the DCNM GUI – From a BGW Device To a Non-Nexus Device 797
Automatic VRF Lite (IFC) Configuration 804
Deleting VRF Lite IFCs 808
Additional References 810
Appendix 810
N9K-3-BGW Configurations 810

CHAPTER 20 MPLS SR and LDP Handoff 813
Overview of VXLAN EVPN to SR-MPLS and MPLS LDP Interconnection 813
VXLAN MPLS Topology 815
Configuration Tasks for VXLAN MPLS Handoff 817
Editing Fabric Settings for MPLS Handoff 817
Editing Easy Fabric Settings 817
Editing External Fabric Settings 819
Creating an Underlay Inter-Fabric Connection 820
Creating an Overlay Inter-Fabric Connection 823
Deploying VRFs 825
Changing the Routing Protocol and MPLS Settings 827

PART IV Layer-2/Layer-3 DCI with VXLAN EVPN Multi-Site 829

CHAPTER 21 Auto-Provisioning Border Gateways with Multi-Site Domains 831
Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - Multi-Site 831
Prerequisites 832

Limitations 833
Save & Deploy Operation in the MSD Fabric 833
EVPN Multi-Site Configuration 835
Configuring Multi-Site Underlay IFCs - DCNM GUI 836
Configuring Multi-Site Underlay IFCs - Autoconfiguration 837
Configuring Multi-Site Underlay IFCs Towards a Non-Nexus Device - DCNM GUI 838
Configuring Multi-Site Overlay IFCs 840
Configuring Multi-Site Overlay IFCs - Autoconfiguration 842
Configuring Multi-Site Overlay IFCs Towards a Non-Nexus Device - DCNM GUI 843
Overlay and Underlay Peering Configurations on the Route Server N7k1-RS1 845
Viewing, Editing and Deleting Multi-Site Overlays 845
Deleting Multi-Site IFCs 846
Creating and Deploying Networks and VRFs in the MSD Fabric 847
Deploying a Legacy Site BGW (vPC-BGWs) 850
Additional References 854
Appendix 854
Multi-Site Fabric Base Configurations – Box Topology 854
IBGP Configuration for the Box Topology in the Easy7200 Fabric 855
Route Server Configuration 856

PART V **Network Provisioning for L4-Layer7 Services** 859

CHAPTER 22 **L4-L7 Service Basic Workflow** 861
Layer 4-Layer 7 Service 861
Guidelines and Limitations for Layer 4-Layer 7 Service 864
Types of Layer 4–Layer 7 Service Devices 865
Configuring Fabric Settings for Layer 4-Layer 7 Service 865
Configuring Layer 4-Layer 7 Service 868
Create Service Node 869
Create Route Peering 872
Create Service Policy 881
Templates 882
Adding a Route Peering 884
Adding a Service Policy 885

Deleting a Service Node 886
Editing a Service Node 887
Refreshing the Service Policy and Route Peering List 888
Refreshing a Specific Service Policy or Route Peering 888
Attaching a Service Policy or a Route Peering 888
Detaching a Service Policy or a Route Peering 888
Preview a Service Policy or a Route Peering 889
Deploying a Service Policy or a Route Peering 890
Viewing Deployment History 890
Exporting a Service Policy or a Route Peering Table 893
Importing a Service Policy or a Route Peering Table 894
Deleting a Service Policy 894
Deleting a Route Peering 895
Viewing Service Policy Information 895
Viewing Route Peering Information 898
Service Node Backup and Restore 899
Fabric Backup and Restore 900
Brownfield Migration 900
Audit History 900

CHAPTER 23 **L4-L7 Service Use Cases** 903
Use Case: Intra-tenant Firewall with Policy-based Routing 903
1. Create Service Node 904
2. Create Route Peering 906
3. Create Service Policy 908
4. Deploy Route Peering 910
5. Deploy Service Policy 912
6. View Stats 915
7. View Traffic Flow in Fabric Builder 915
8. Visualize Redirected Flows to Destination in the Topology window 918
Use Case: Inter-tenant Firewall with eBGP Peering 921
1. Create Service Node 923
2. Create Route Peering 924
3. Deploy Route Peering 926

Use Case: One-arm Load Balancer 928
1. Create Service Node 928
2. Create Route Peering 930
3. Create Service Policy 931
4. Deploy Route Peering 931
5. Deploy Service Policy 931
6. View Stats 932
7. View Traffic Flow in Fabric Builder 932
8. Visualize Redirected Flows to Destination in the Topology window 932

PART VI **Public Cloud Connectivity** 933

CHAPTER 24 **Connecting Cisco Data Center and a Public Cloud** 935
Connecting Cisco Data Center and a Public Cloud 935
Topology Overview 936
Guidelines and Limitations 937
Prerequisites 937
Task Summary 937
Setting the Polling Time 938
Setting Up the On-premise External Fabric with CSR 1000v 939
Creating an External Fabric 939
Discovering the On-Premises Core Router 939
Setting Up the VXLAN EVPN Fabric 940
Creating a VXLAN EVPN Fabric 940
Assigning the BGW Role 941
Setting Up the External Fabric with CSR in Azure 941
Creating an External Fabric 941
Discovering the Core Router 942
Setting Up the MSD Fabric for Connectivity 943
Creating an MSD Fabric 943
Moving Other Fabrics into the MSD Fabric 944
Setting Up Connections 945
Connecting the On-Premises BGW and the On-Premises Core Router 945
Connecting the On-prem Core Router and the Public-cloud Core Router with IPsec Tunnel 946

Connecting the On-prem BGW and the Public-cloud Core Router using EVPN Peering 948
Saving and Deploying Configurations 950
Extending VRFs 951
Deploying and Extending the VRF On-prem Core Router 952
Creating and Deploying VRF on Public Cloud 953
Configuring Default Gateway for the VM 954
Verifying the Connectivity 955
Deploying Cisco CSR 1000v on Microsoft Azure 955
Viewing Links and Core Routers Details 959
Resetting Packet Counter Using API 959

PART VII **Easy Provisioning of MSDC Deployments** 961

CHAPTER 25 **Managing BGP-Based Routed Fabrics** 963
Creating an eBGP-based Fabric 963
Adding Switches to a Fabric 973
Discovering Existing Switches 974
Discovering New Switches 979
Deploying Fabric Underlay eBGP Policies 987
Deploying Networks in eBGP-based Fabrics 988
Overview of Networks in a Routed Fabric 988
Creating and Deploying a Network in a Routed Fabric 989
Creating Inter-Fabric Links Between a Routed Fabric and an External Fabric 993

PART VIII **Templates Usage** 997

CHAPTER 26 **Template Usage in Cisco DCNM LAN Fabric Deployment** 999
Policy Template 999
Fabric Template 1003
Profile Template 1003
Viewing, Editing, and Adding Policies 1004
Viewing Policies 1005
Editing Policies 1007
Adding Policies 1008

Deploying New Configurations 1008
switch_freeform Template Usage 1009
Example: Create a switch_freeform policy 1009
Changing the Contents of a Template in Use 1012

CHAPTER 27 Guidelines for Programmable Reports 1015
Prerequisites 1015
CLI Output Processing 1016
Report Template 1017
Template Content 1018

CHAPTER 28 Cisco DCNM Programmable Report APIs 1021
Template 1021
UPGRADE 1021
GENERIC 1021
Template Structure 1021
Template Functions 1022
Context Parameter 1023
Report Layout 1023
Summary View 1024
Detail View 1024
Command Log 1025
Report Python Library 1025
Report APIs 1025
Create Report Object 1025
Add Summary 1025
Add Section 1026
Formatters 1027
Chart 1028
Run CLIs on Device 1030
Get Job Context Information 1031
Analyze with Historical Reports 1031
XML Utilities 1032
WrapperResp 1033

Logger 1033

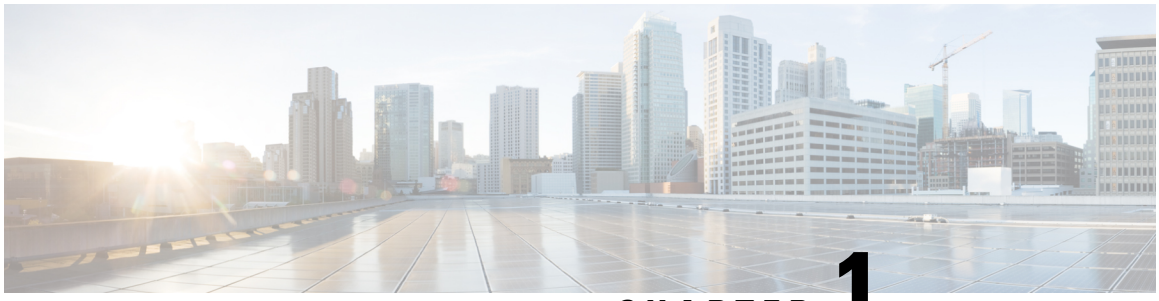

CHAPTER 1

Overview

- Cisco Data Center Network Manager, on page 1
- REST API Tool, on page 2

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.

Note The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Configuring the Device Connector is mandatory if you've deployed Cisco DCNM in LAN Fabric mode. If you did not configure Device Connector during installation, a message appears asking you to configure Device Connector everytime you login. If you check the **Do not show again**, the message will not appear. However, an alarm notification will be added under the **Alarms** icon.

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:

- **Alerts and Notifications**: You can view the alerts and event notifications by clicking the **Alerts and Notifications** icon, next to the **Help** icon, in the top pane of Cisco DCNM.
- **Help**: Launches the context-sensitive online help.
- **Alarms**: The **Alarms** icon flashes when there is an Alarm or when thresholds exceed for your Cisco DCNM Deployment. Click on the flashing **Alarms** icon to view the messages. The following alarms are displayed.

- **Interfaces Limit Exceeded for DCNM** – If the maximum number of endpoints across all fabrics exceeds 100K, the **Alarms** icon flashes and displays a message.
- **Device Connector Disconnected** – If the Device Connector was not configured during installation, this alarm appears to imply that the Device connector is not connected to Intersight. Choose **Administration > Device Connector** to configure the Device Connector and remove the alarm.
- **High Availability (HA) State of DCNM** – If the Native HA setup is not synchronized. One of the nodes or both nodes may have stopped, failed, or not ready if the **Alarms** icon flashes. When the HA setup is synchronized, the notification clears out in 30 minutes (during the polling cycle) or when you logout and login to the Cisco DCNM Web UI.
- **Application down** – If one or more applications are down, an error appears. An alarm message appears when the applications goes online or offline. Click on the respective alarm to navigate to the application on **Web UI > Applications > Catalog**.
- **Compute Node disconnected** – An alarm message appears when one or more Compute node goes down.

- *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.
  - **REST API Tool**: Allows you to examine the APIs invoked for every operation. See the *REST API Tool* section for more information about the API inspection.
  - **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.

REST API Tool

Operations like discovery, fabric management, monitoring, and so on, which are performed in Cisco DCNM Web UI, invoke HTTP calls to fetch and commit the information accessed. The REST API tool enables you to examine the API call by viewing the structure of an API call. This tool also provides a corresponding CURL request to help with building quick prototypes and testing APIs.

The **REST API Tool** dialog box has the following fields.

Table 1: Fields and Description for the REST API Tool Dialog Box

Field	Description
Filter	Enter any keyword to search the log.

Field	Description
scroll to new items	Check this check box to scroll to the new entries when you navigate back to the **REST API Tool** dialog box after you perform an operation in the Web UI. This check box is checked by default.
clear log	Click **clear log** to clear the log in the dialog box.
API-docs	Click API-docs to view the Cisco DCNM REST API documentation in the Web UI. Clicking this option takes you to the following URL: https://*DCNM-IP*/api-docs

All actions you perform in the Cisco DCNM Web UI appear in the API inspector tool. The following information appears in the APIs invoked for every operation:

- HTTP method
- URI
- Payload
- HTTP status code
- Time taken for the operation

The following image displays how the log appears in the **REST API Tool** dialog box.

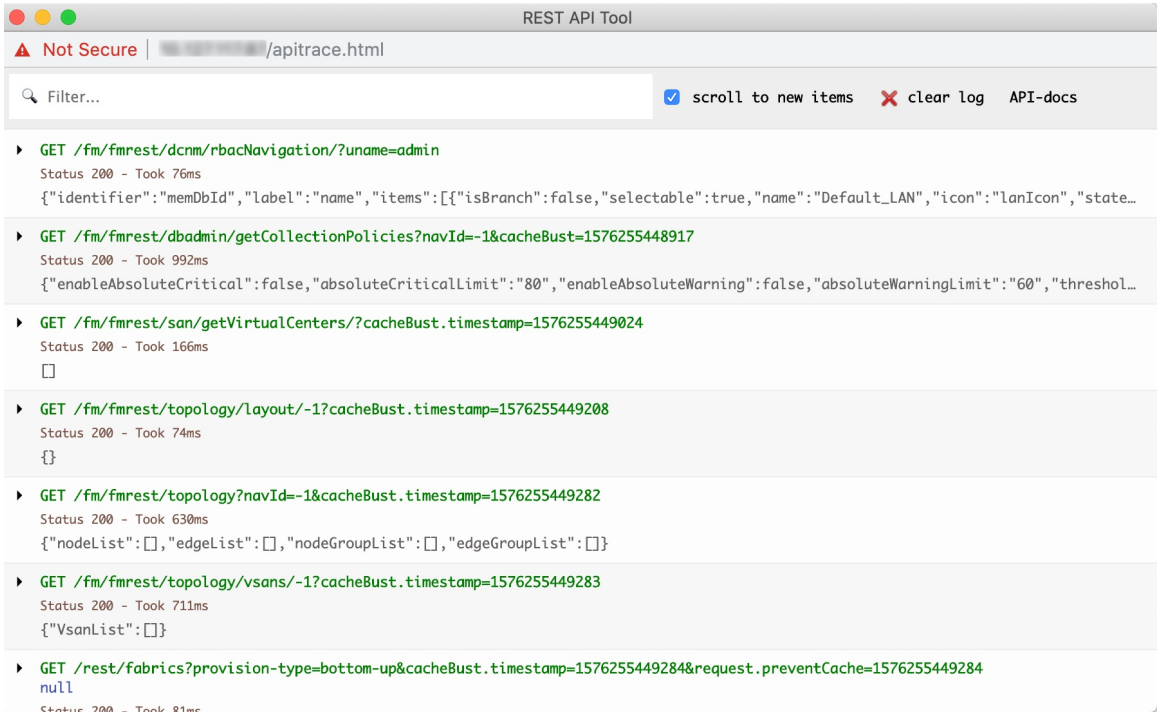

Click the URI to expand or collapse each REST method. You can perform the following actions after expanding a REST method:

- **Prettify output**: Click this option to arrange the response code in a more presentable way, which otherwise appears in a single line. Scroll through the response to view it completely.
- **Copy response**: Click this option to copy the response code to your clipboard.
- **Copy CURL request**: Click this option to copy the CURL request to your clipboard.

```
curl -k -XGET --header 'Dcnm-Token: <DCNM_TOKEN>' --header 'Content-Type:
application/x-www-form-urlencoded'
https://<ip-address>/fm/fmrest/dcnm/rbacNavigation/?uname=admin
```

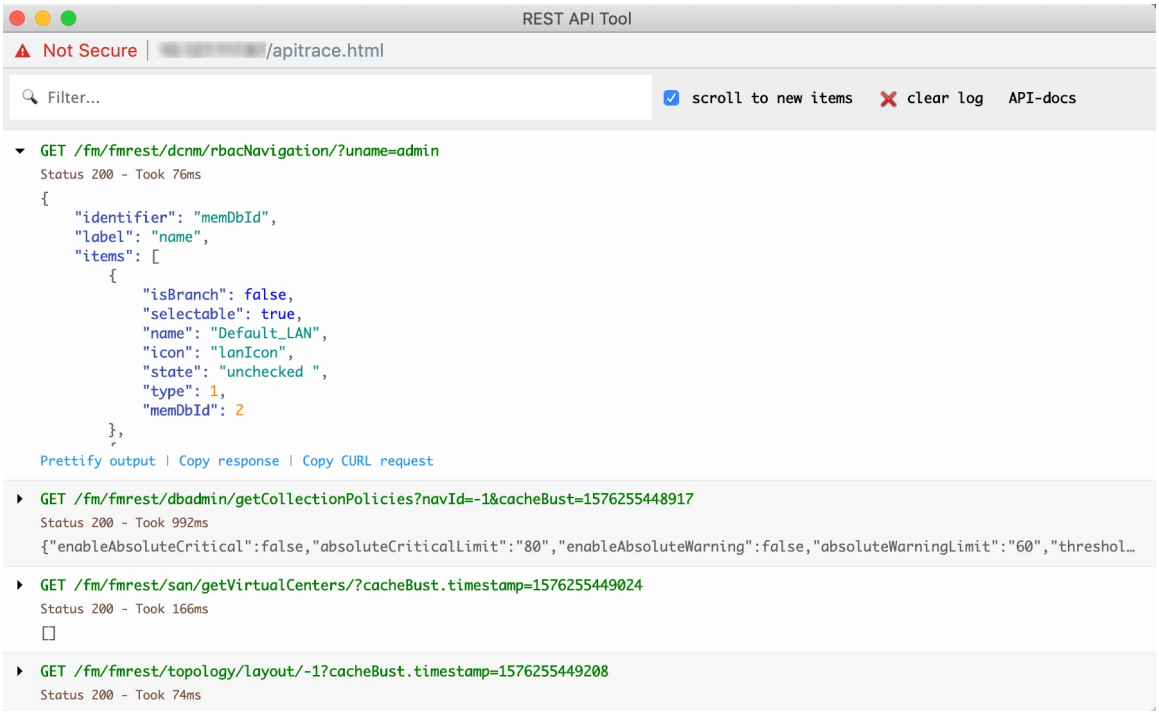

The **REST API Tool** dialog box updates every time the Cisco DCNM Web UI updates.

To use the API inspector from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Click the **Gear** icon in the top pane.

Step 2 Choose **REST API Tool** from the drop-down list.

The **REST API Tool** dialog box appears and the log is empty before you perform any operation in the Cisco DCNM Web UI.

Step 3 Minimize the **REST API Tool** dialog box.

Note You can also keep the dialog box open, but not close it.

Step 4 Perform an operation in the Cisco DCNM Web UI.

Note You can perform any operation in the Cisco DCNM Web UI like viewing any options, adding, deleting, and so on.

Step 5 Navigate back to the **REST API Tool** dialog box.

The log is populated with the REST APIs fetched depending on the operations you performed.

Note Closing the **REST API Tool** dialog box, instead of minimizing it before performing any operations, clears the log.

For a demo on some of the operations that can be performed using the REST API tool, see the Using REST API Tool in Cisco DCNM video.

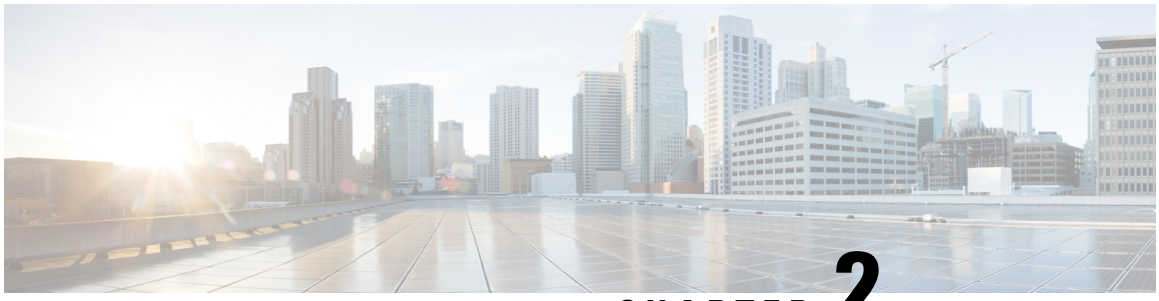

CHAPTER 2

New and Changed Information

This chapter contains the following section:

- New and Changed Information in Cisco DCNM, Release 11.5(1), on page 7

New and Changed Information in Cisco DCNM, Release 11.5(1)

The following table provides an overview of the significant changes to this guide for this current release. The table does not provide an exhaustive list of all changes made to the guide or of the new features in this release.

Table 2: New and Changed Behavior in Cisco DCNM, Release 11.5(1)

Feature	Description	Where Documented
Single-switch Configuration Restore	You can restore configuration for a Cisco Nexus switch in external and LAN classic fabrics from the Cisco DCNM Web UI. The information you restore at switch-level is extracted from the fabric-level backups. The switch-level restoration does not restore fabric-level intents and other configurations applied using the fabric settings. Only switch-level intents are restored.	Restoring a Switch, on page 321
EPLD Golden Upgrade	From Cisco DCNM Release 11.5(1), DCNM supports EPLD golden upgrade as well. When you perform the EPLD upgrade, you have an option to choose the golden or primary region of the Nexus 9000 Series switches. You can view the EPLD golden upgrade notifications in the **Events** window. From the homepage of the Cisco DCNM Web UI, choose **Monitor > Switch >Events**.	EPLD Installation, on page 424

PTP Monitoring Application	The Precision Time Protocol (PTP) is a time synchronization protocol for nodes that are distributed across a network. On a local area network, it achieves clock accuracy in the sub-microsecond range, making it suitable for measurement and control systems. In DCNM, PTP Monitoring can be installed as an application. This PTP monitoring application, which can be previously installed in Media Controller deployment can now be installed in LAN Fabric deployment as a preview feature. We recommend that you do not deploy this feature in production environments.	PTP Monitoring, on page 610
Support for Simplified CLI Configuration for Brownfield Deployment	The Brownfield import in DCNM supports the simplified NX-OS VXLAN EVPN configuration CLIs.	Guidelines and Limitations, on page 696
CloudSec Operational View	You can use the **CloudSec Operational View** tab in DCNM to check the operational status of the CloudSec sessions if CloudSec is enabled on the MSD fabric.	Viewing CloudSec Operational State, on page 177
Sync up Out-of-Band Switch Interface Configurations with DCNM	You can use the **host_port_resync** policy to sync up any out-of-band switch interface level configurations (via CLI) with Cisco DCNM and subsequently manage it. Additionally, the vPC pair configurations are automatically detected and paired.	Sync up Out-of-Band Switch Interface Configurations with DCNM, on page 190
Support for MACsec in Easy Fabric and eBGP Fabric	MACsec is supported in the Easy Fabric and eBGP Fabric on intra-fabric links. You need to enable MACsec on the fabric and on each required intra-fabric link to configure MACsec. Note that this is a preview feature in the Cisco DCNM Release 11.5(1).	MACsec Support in Easy Fabric and eBGP Fabric, on page 194
Interface Group	You can create an interface group that allows grouping of host-facing interfaces at a fabric level. Specifically, you can create an interface group for physical Ethernet interfaces, L2 port-channels, and vPCs. You can attach or unattach multiple overlay networks to the interfaces in an interface group.	Interface Groups, on page 270

L4-7 Services Enhancements	The following enhancements are introduced in DCNM Release 11.5(1): • You can specify an arbitrary network, that has not been defined in the top-down configuration, as a source or destination network in the service policy. This helps in streamlining policy enforcement for north-south traffic. • Layer 4-Layer 7 Service pushes static routes on all VTEPs, including service leaf switches, where the VRF being referenced in the static route is attached. This expedites service node failover with static routes. • The one-arm Virtual Network Function is supported. • Layer 4-Layer 7 Service REST APIs are accessible via DCNM packaged REST API documentation. • Bulk attachment, detachment, preview, and deployment of route peering and service policies is supported and they are limited up to 10 route peerings or 10 service policies only. • **Audit History** feature displays the logs for changes made to service nodes, route peering, and service policies.	Layer 4-Layer 7 Service, on page 861
OpenStack Workload Visibility	OpenStack plugin application is provided by DCNM that helps you to monitor OpenStack Clusters. You can get visibility with respect to the physical network connectivity and virtualized workloads, and debug VM networking specific issues within the context of the data center. Note that this is a preview feature in the Cisco DCNM Release 11.5(1).	• OpenStack Visualizer, on page 369 • OpenStack Workload Visibility, on page 46
Support for L3 Gateway on Border for fabrics	From Cisco DCNM Release 11.5(1), the **Enable L3 Gateway on Border** field is not available as part of the MSD network settings. You can enable a Layer 3 gateway on the border switches at a fabric level.	Multi-Site Domain for VXLAN BGP EVPN Fabrics , on page 146
Periodic report generation frequency	• When you are creating a Periodic NVE VNI Counters report, the report generation frequency has to be set to 60 minutes or more. If the frequency is less than 60 minutes, an error message is displayed. • The generateReport method is invoked while generating a report and contains the report implementation logic. This method accepts any context object.	• Creating a Report Job, on page 616 • Report Template Functions, on page 402

Pre-provisioning a device	From Cisco DCNM Release 11.5(1), extended configuration support to pre-provisioned devices.	Pre-provisioning a Device
Enhanced Role-based Access control	New user roles, **device-upg-admin**, and **access-admin** are added. • A user with the **device-upg-admin** role can perform operations only in **Image Management** window. • A user with the **access-admin** role can perform operations only in **Interface Manager** window for all fabrics.	• Enhanced Role-based Access Control in Cisco DCNM • Interfaces
Switch-smart License	From Cisco DCNM Release 11.5(1), new license type is added for switches.	Switch Smart License
Inband Management in External Fabrics and LAN Classic Fabrics	Cisco DCNM allows you to import or discover switches with inband connectivity for External and LAN Classic fabrics in Brownfield deployments only. Enable inband management, per fabric, while configuring or editing the Fabric settings. You cannot import or discover switches with inband connectivity using POAP.	Inband Management in External Fabrics and LAN Classic Fabrics, on page 186
Precision Time Protocol for External Fabrics or LAN Classic Fabrics	From Release 11.5(1), in the fabric settings for the **External_Fabric_11_1** or **LAN_Classic** template, select the Enable Precision Time Protocol (PTP) check box to enable PTP across a fabric.	Precision Time Protocol for External Fabrics and LAN Classic Fabrics, on page 188
Ability to Edit DNS, NTP Servers from the GUI	Cisco DCNM allows you to modify few network parameters from the Web UI. Modifying these will overwrites the previously configured parameters.	Network Preferences, on page 482

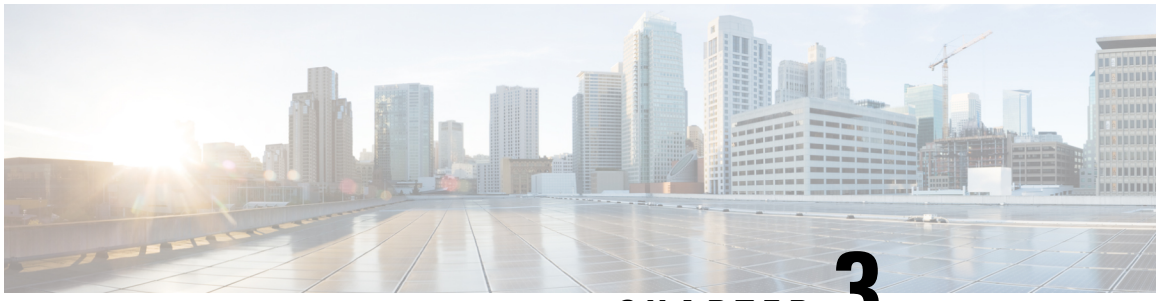

CHAPTER 3

Dashboard

This chapter contains the following topics:

- Dashboard, on page 11

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
- Alarms

- 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 **Critical**, **Error**, and **Warning** severity. In this dashlet, click the **Show Acknowledged Events** link to go to the **Monitor** > **Switch** > **Events**.
Alarms	Displays alarms with **Critical, Major, Minor**, and **Warning** severity. In this dashlet, click the **Show Acknowledged Alarms** link to go to the **Monitor** > **Alarms** > **View** window. Hover the mouse cursor over the blue **i** icon for more information about a specific alarm. Click **ACK** to acknowledge a specific alarm.
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

Dashlet	Description
	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. • The network map dialog box has properties that are different from the Summary dashboard view: • You can click and drag nodes to move them around the map. The map saves their new positions. • You can double click a node to trigger a slider that contains the summary inventory information pertaining to a specific switch group. • You can upload an image of your choice as the background to the network map. **Note** 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.
Server Status	Displays the status of DCNM and federation servers, and the health check status for the components. The following services, server, and status details are displayed under the **DCNM** tab. • Database Server • Search Indexer • Performance Collector • NTPD Server • DHCP Server • SNMP Traps • Syslog Server The following component status and details are displayed under the **Health Check** tab. • AMQP Server • DHCP Server • TFTP Server • EPLS

Dashlet	Description
	• EPLC
Top ISLs/Trunks	Displays the performance data for the top ten performing ISLs, trunk ports or both. Each entry shows the current average receive and transmit percentage, with a graph depicting the percentage of time each trunk spent exceeding the currently configured thresholds.
Top SAN End Ports (SAN only)	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. **Note** This dashlet is only for SAN.
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	Displays the module temperature sensor details of switches. **Note** This dashlet is only for LAN.
Health	Displays the health summary that contains two columns displaying the summary of problems and summary of events for the past 24 hours. 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. 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. From Release 11.4(1), if you have deployed Cisco DCNM in HA mode, the Health Dashlet displays the status of the HA setup. Along with the HA State, it also displays the IP Addresses for the Active, Standby HA nodes and VIP.
Errors	Displays the error packets for the selected interface. This information is retrieved from the **Errors > In-Peak** and **Errors > Out-Peak** columns of the **Monitor > LAN / Ethernet** page.

Dashlet	Description
Discards	Displays the error packets that are discarded for the selected interface. **Note** The Discards dashlet is only for LAN.
Inventory (Ports)	Displays the ports inventory summary information.
Inventory (Modules)	Displays the switches on which the modules are discovered, the models name and the count.
Inventory (ISLs)	Displays the ISLs inventory summary information, such as the category and count of ISLs.
Inventory (Logical)	Displays the logical inventory summary information, such as the category and count of logical links.
Inventory (Switches)	Displays the switches inventory summary information such as the switch models and the corresponding count.
Inventory (Port Capacity)	Displays the port capacity inventory summary information such as the tiers, the number and percentage of the available ports, and the remaining days.

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

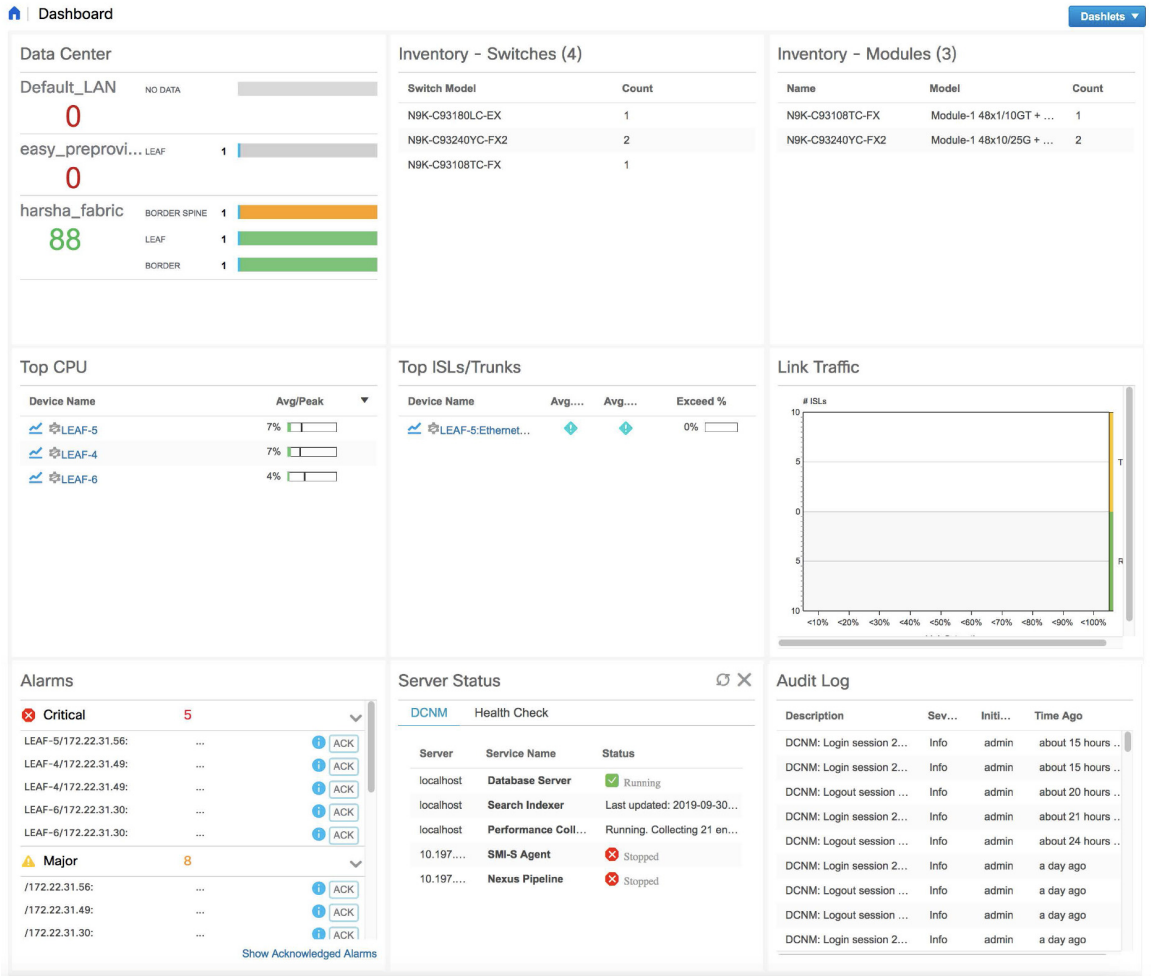
Dashboard
Dashlets
Data Center
Default_LAN
NO DATA
0
easy_preprovi...
LEAF
1
0
harsha_fabric
BORDER SPINE
1
88
LEAF
1
BORDER
1
Inventory - Switches (4)
Switch Model
Count
N9K-C93180LC-EX
1
N9K-C93240YC-FX2
2
N9K-C93108TC-FX
1
Inventory - Modules (3)
Name
Model
Count
N9K-C93108TC-FX
Module-1 48x1/10GT + ...
1
N9K-C93240YC-FX2
Module-1 48x10/25G + ...
2
Top CPU
Device Name
Avg/Peak
LEAF-5
7%
LEAF-4
7%
LEAF-6
4%
Top ISLs/Trunks
Device Name
Avg....
Avg....
Exceed %
LEAF-5:Ethernet...
0%
Link Traffic
ISLs
10
5
0
5
10
<10%
<20%
<30%
<40%
<50%
<60%
<70%
<80%
<90%
<100%
Alarms
Critical
5
LEAF-5/172.22.31.56:
LEAF-4/172.22.31.49:
LEAF-4/172.22.31.49:
LEAF-6/172.22.31.30:
LEAF-6/172.22.31.30:
ACK
Major
8
/172.22.31.56:
/172.22.31.49:
/172.22.31.30:
Show Acknowledged Alarms
Server Status
DCNM
Health Check
Server
Service Name
Status
localhost
Database Server
Running
localhost
Search Indexer
Last updated: 2019-09-30...
localhost
Performance Coll...
Running. Collecting 21 en...
10.197....
SMI-S Agent
Stopped
10.197....
Nexus Pipeline
Stopped
Audit Log
Description
Sev...
Initi...
Time Ago
DCNM: Login session 2...
Info
admin
about 15 hours ..
DCNM: Login session 2...
Info
admin
about 15 hours ..
DCNM: Logout session ...
Info
admin
about 20 hours ..
DCNM: Login session 2...
Info
admin
about 21 hours ..
DCNM: Logout session ...
Info
admin
about 24 hours ..
DCNM: Login session 2...
Info
admin
a day ago
DCNM: Logout session ...
Info
admin
a day ago
DCNM: Logout session ...
Info
admin
a day ago
DCNM: Login session 2...
Info
admin
a day ago

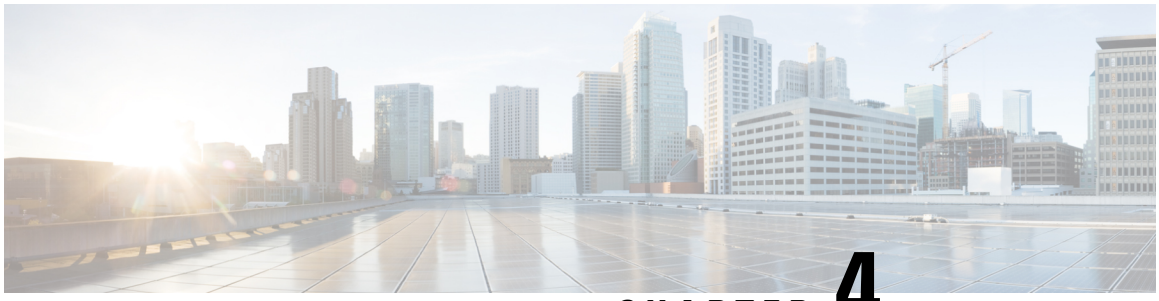

CHAPTER 4

Topology

- Topology, on page 17

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.

Starting from Cisco DCNM Release 11.4(1), if a switch is in maintenance mode, a **Maintenance Mode** badge is displayed next to the switch. If a switch is in migration mode, a **Migration Mode** badge is displayed next to the switch.

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.

The **Quick Search** provides options to search OpenStack resource based on its type, that is, IP address or name. You can search by Host IP, and the corresponding host is highlighted. You can also select a specific OpenStack cluster based on its IP address from the **OpenStack** drop-down list and search within it.

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.

VM name (OpenStack)

Select **VM name (OpenStack)** from the search field and enter a VM name. The path to the fabric switch from the VM is highlighted. You should have selected **All** under OpenStack in the **Show** panel for this search option. Otherwise, this search option is disabled.

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 address. 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.

Redirected Flows

After the physical attachment of a service node to the fabric is defined, select **Redirected Flow** from the **Quick Search** drop-down list in the **Topology** window.

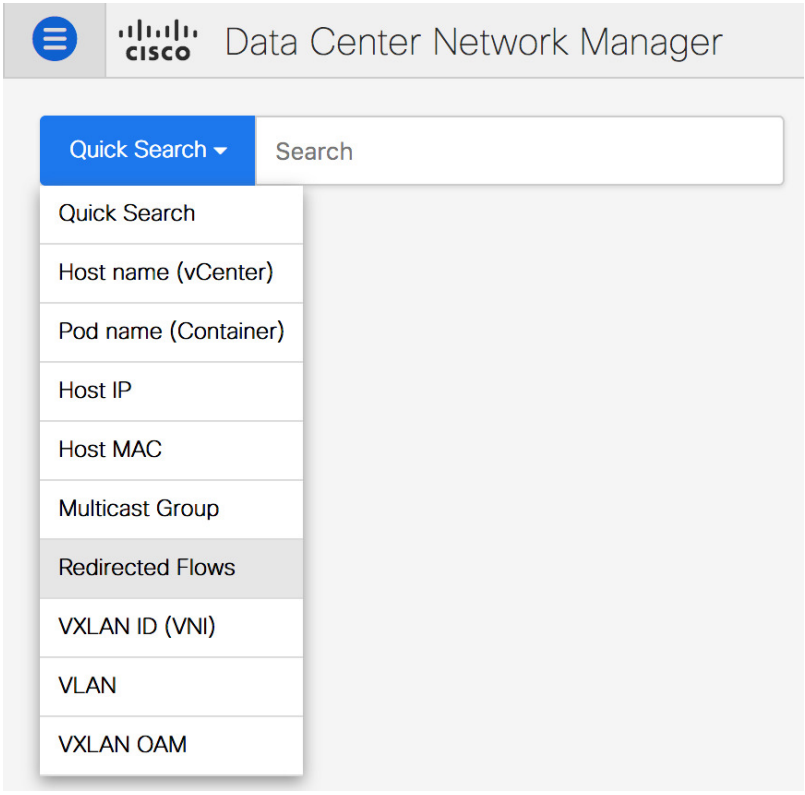

You can select a policy from the drop-down list or initiate a search by entering a policy name, source network and destination network in the search field. The search field is auto-populated based on your input.

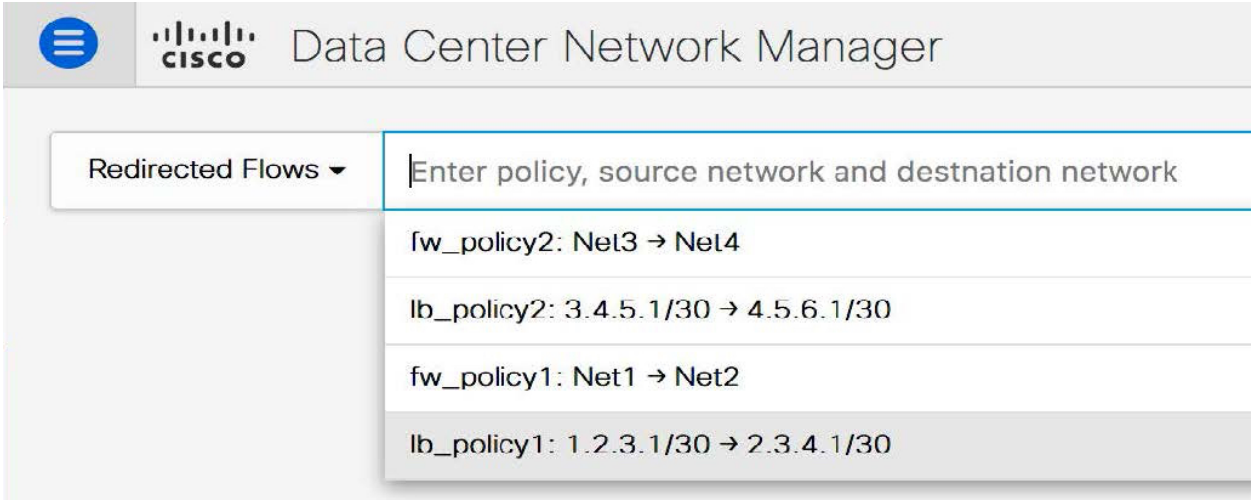

Based on the input in the search field, the switches are highlighted on the topology window. The switches, on which the source and destination network have been attached and the flow has been redirected, are highlighted on the topology window. The service node is shown as connected by a dotted line to the leaf switch on the topology window. Hover over the dotted line to get more information about the interface. Click a switch to display the redirected flows which are initiated, redirected to, or terminated on that switch. Click **Show more flows** to display the **Service Flows** window that has information about all the redirected flows.

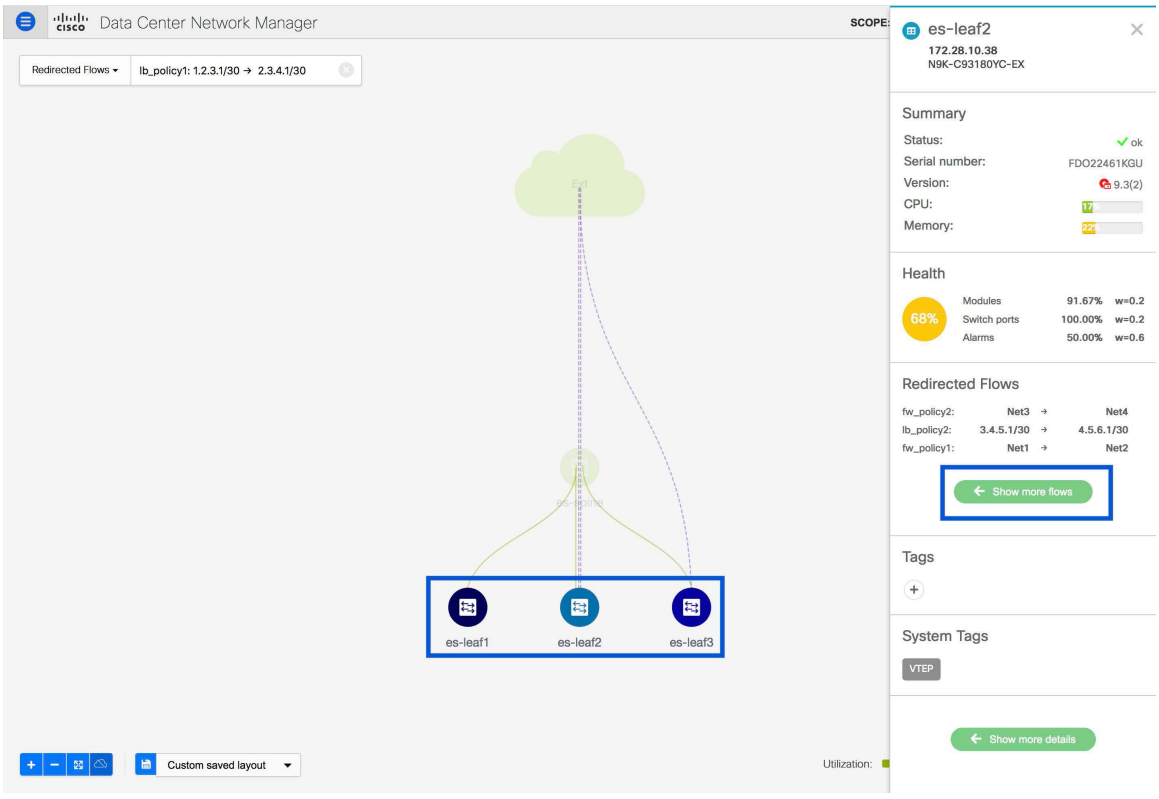

Click **Details** in the **Service Flows** window to display attachment details.

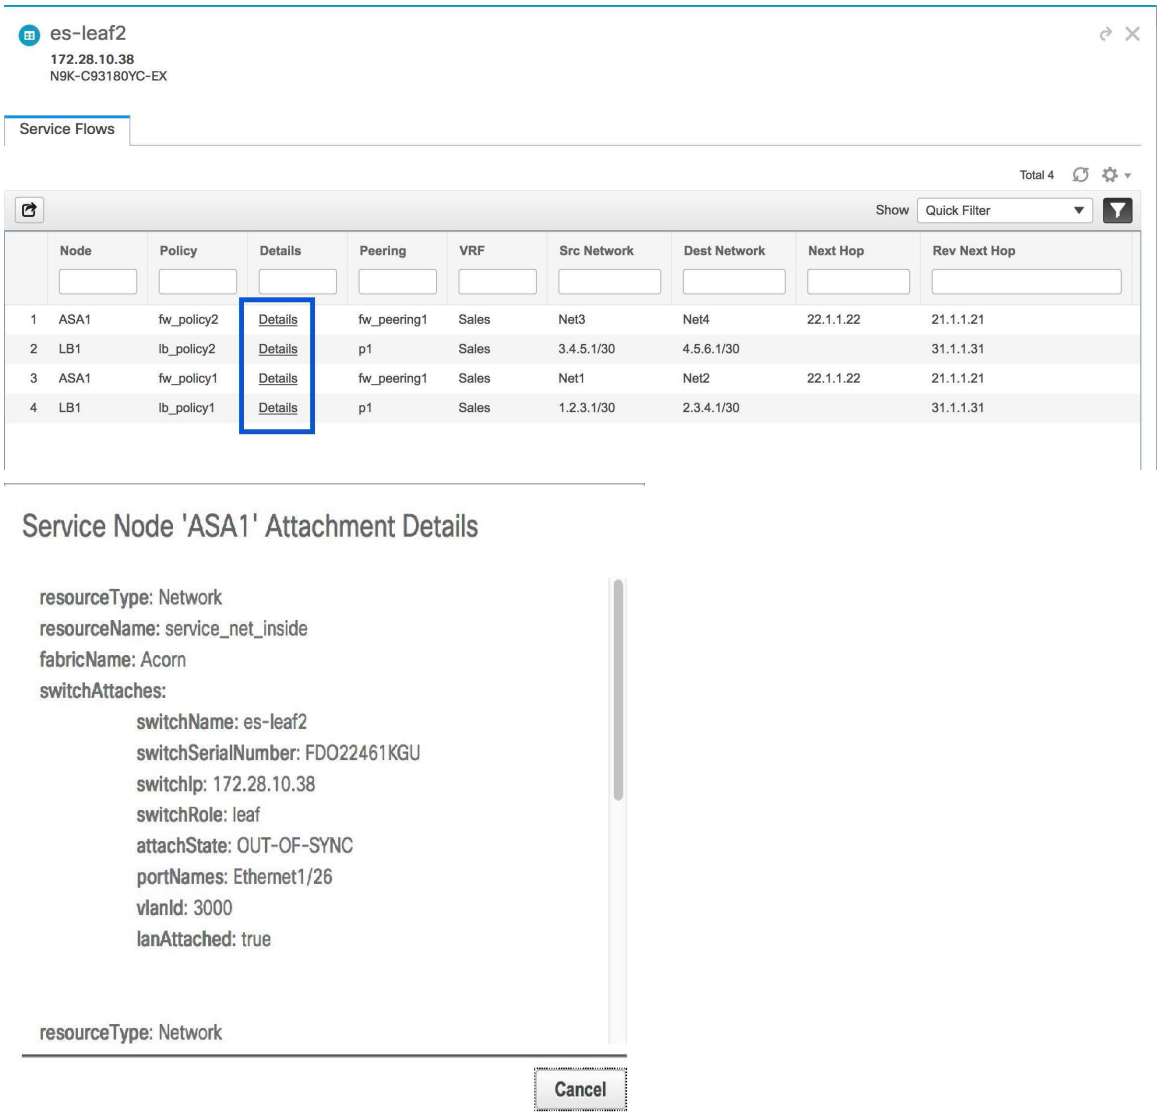

VXLAN ID (VNI)

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.

  From Cisco NX-OS Release 11.4(1), if you uncheck this check box, FEX devices are hidden in the **Fabric Builder** topology window as well. To view FEX in **Fabric Builder**, you need to check this check box. This option is applicable for all fabrics and it is saved per session or until you log out of DCNM. If you log out and log in to DCNM, the FEX option is reset to default, that is, enabled by default. For more information, see Creating a New VXLAN BGP EVPN Fabric, on page 55.

  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).

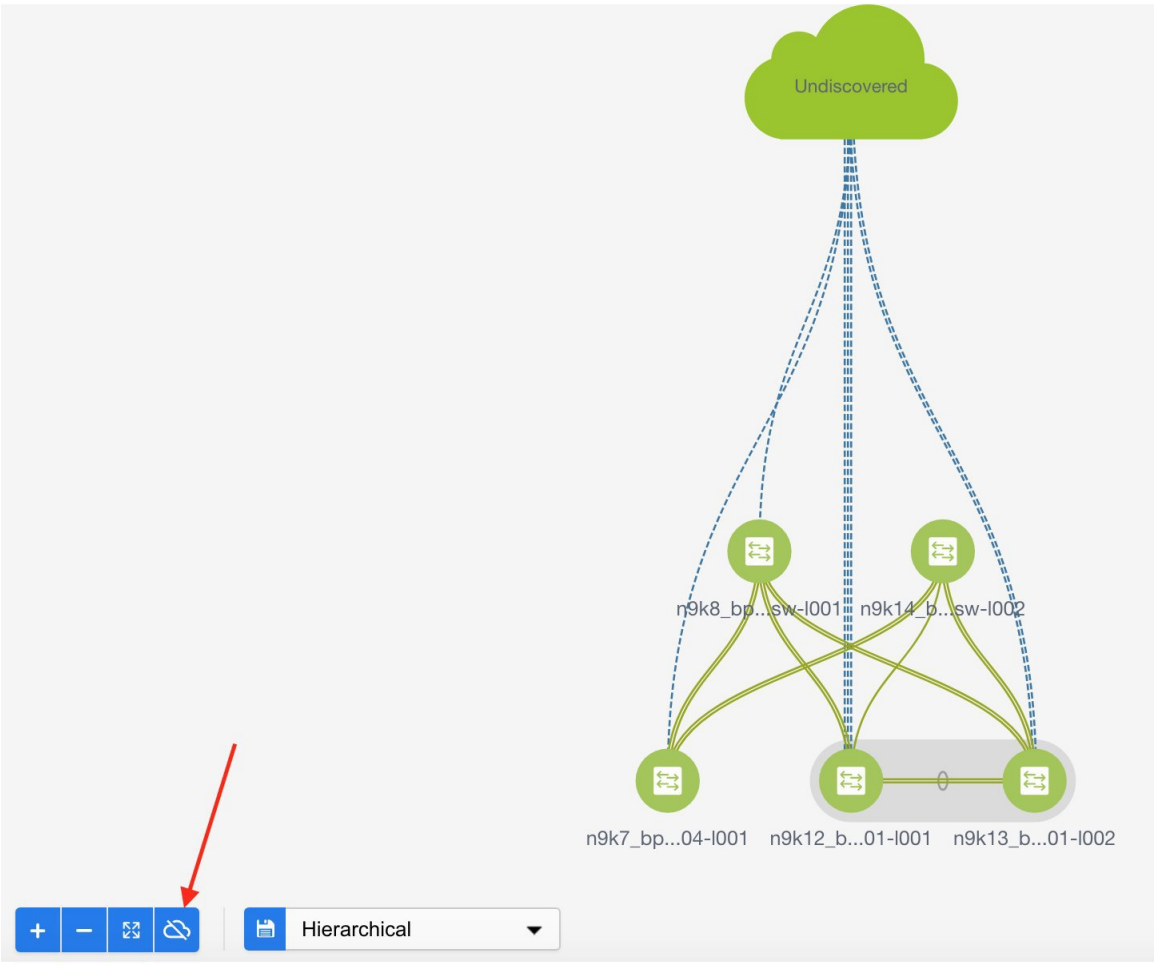

Click again to stop the **Undiscovered** cloud from being displayed. You can see that the **Undiscovered** cloud and its links to the fabric devices are not displayed.

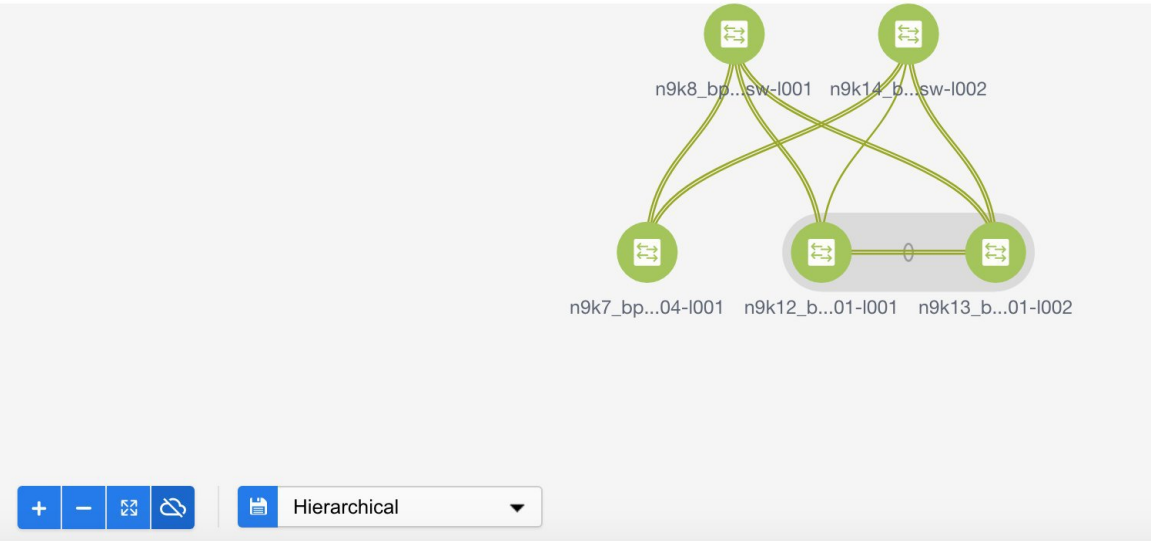

Click the **Cloud** icon again to display the **Undiscovered** cloud.

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-polled CPU utilization, and last-polled memory utilization.

Beacon

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

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

Tagging

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

Use the search functionality to perform searches based on tags.

More Details

Click **Show more details** to display more information under the following tabs: **System Info, Modules, FEX, License, Features, VXLAN, VLAN, Capacity** and **Hosts**.

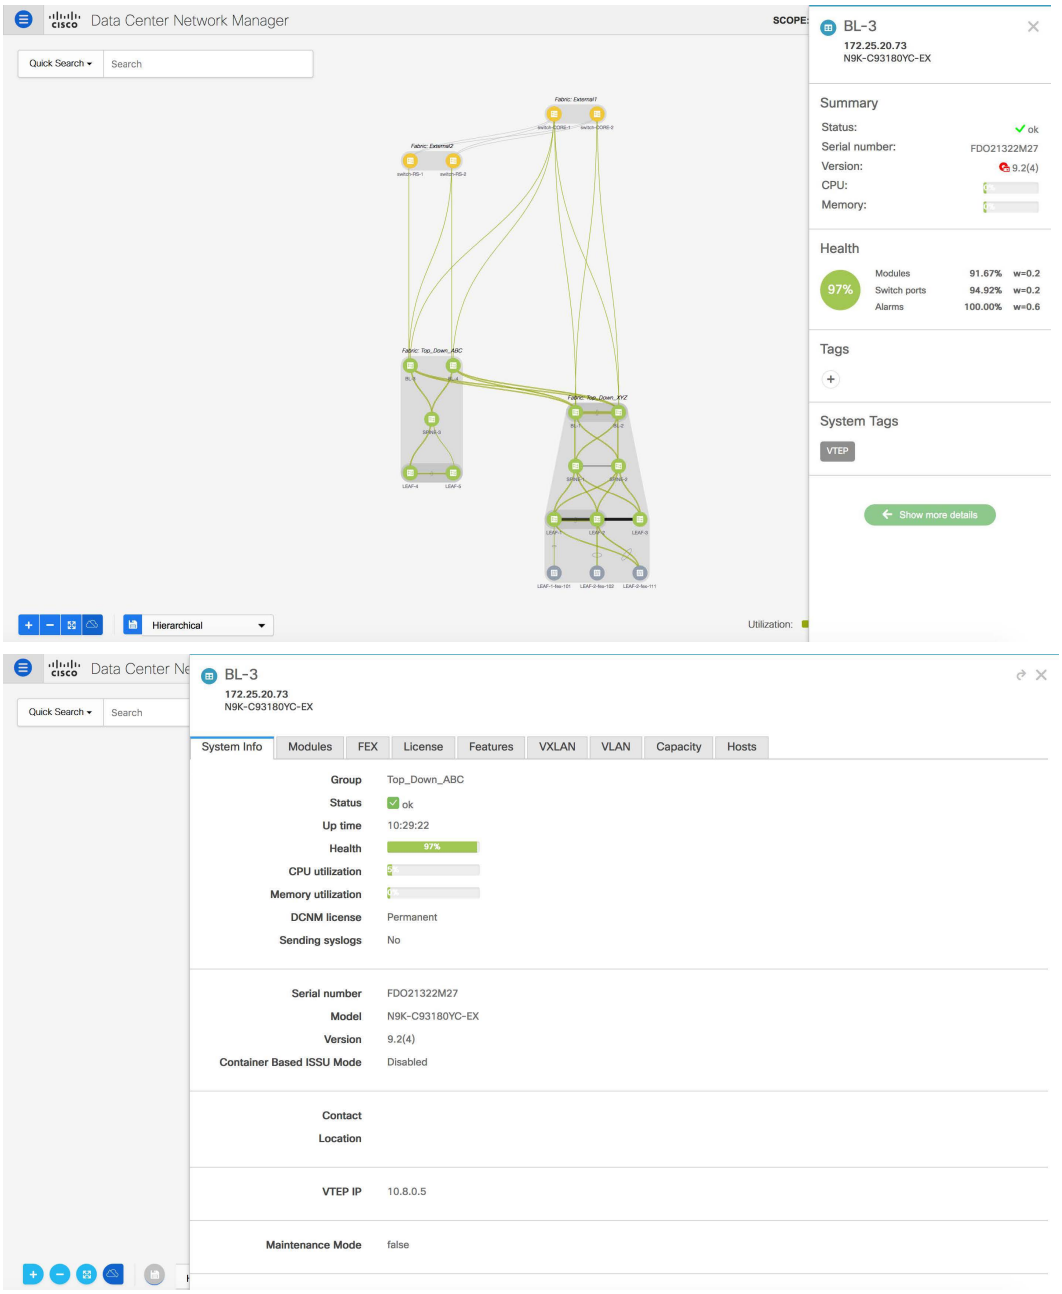

Starting from Cisco DCNM Release 11.4(1), the 400G tier has also been added to the **Physical Capacity** table under the **Capacity** tab. However, the **Physical Capacity** table under the **Capacity** tab will only show information about the physical ports that are present on the switch. For example, if the switch does not have a 400G physical port, the 400G tier is not displayed in the **Physical Capacity** table.

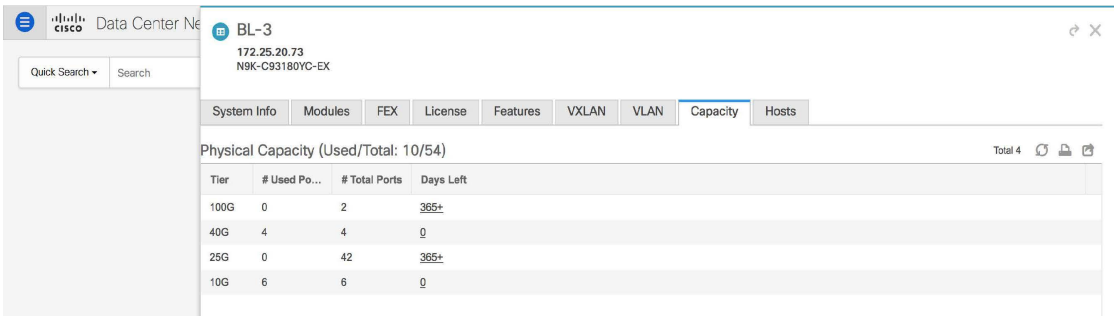

Link Slide-Out Panel

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

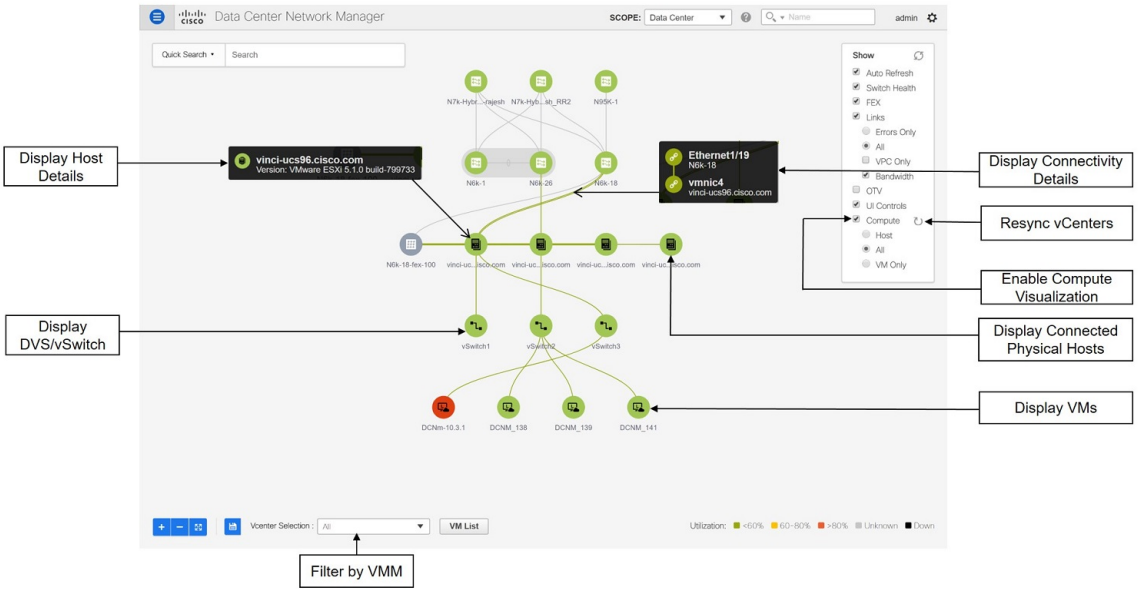

Support for Cisco UCS B-Series Blade Servers

Cisco DCNM Supports hosts running on UCS type B (chassis UCS) that are behind the Fabric interconnect. You must enable CDP of the vNIC on Cisco UCSM to use this feature.

Note By default, CDP is disabled on Cisco UCSM.

Let us consider two VMMs, VMM-A and VMM-B, for reference. After the discovery of Cisco UCS UCS B-Series Blade Servers, the Topology displays the blue colored VMM-A and VMM-B are fabric interconnect nodes. A sample topology is as shown in the figure below.

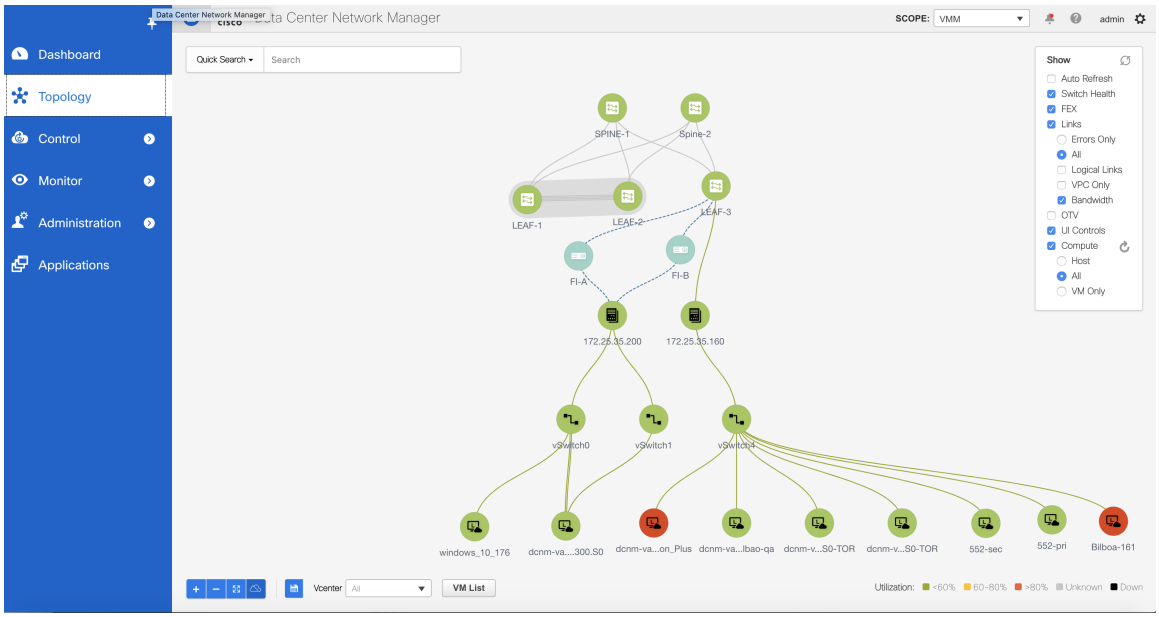

To enable CDP on UCSM, you must create a new Network Control policy using the following steps.

1. On the USCM, choose **LAN** and expand the policies.
2. Right-click on the **Network Control Policies** to create a new policy.
3. In the Name field, enter the policy name as **EnableCDP**.
4. Choose **enabled** option for CDP.

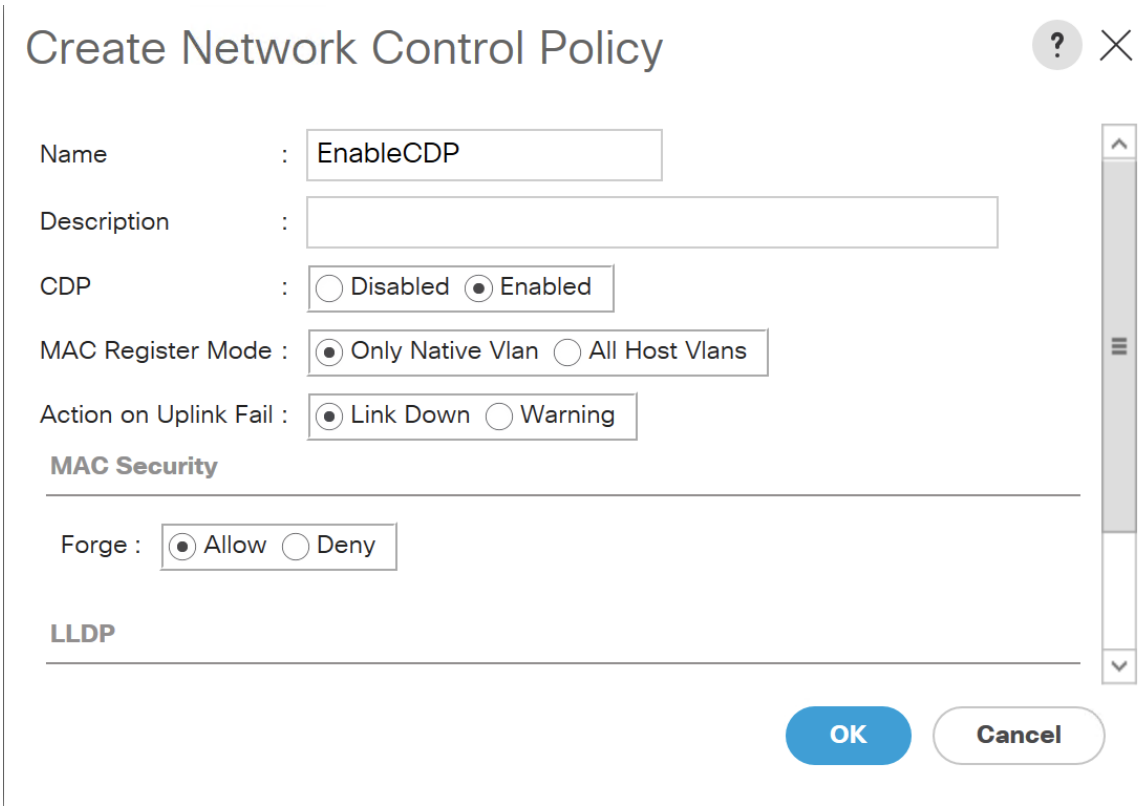

5. Click **OK** to create the policy.

To apply the new policy to the ESX NICs, perform the following steps:

- If you are using updated vNIC templates, choose each vNIC template for your ESXi vNICs, and apply the EnableCDP policy from the Network Control Policy drop-down list.
- If you are not using any vNIC templates, use the updated Service Profile Template. Apply EnableCDP policy on each of the service profile template.
- If you are using one-off Service Profiles (i.e., if each server using its own service profile), then you must go to every Service Profile and enable EnableCDP policy on every vNIC.

For more information about Cisco UCSM, refer to Cisco UCSM Network Management Guide.

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.

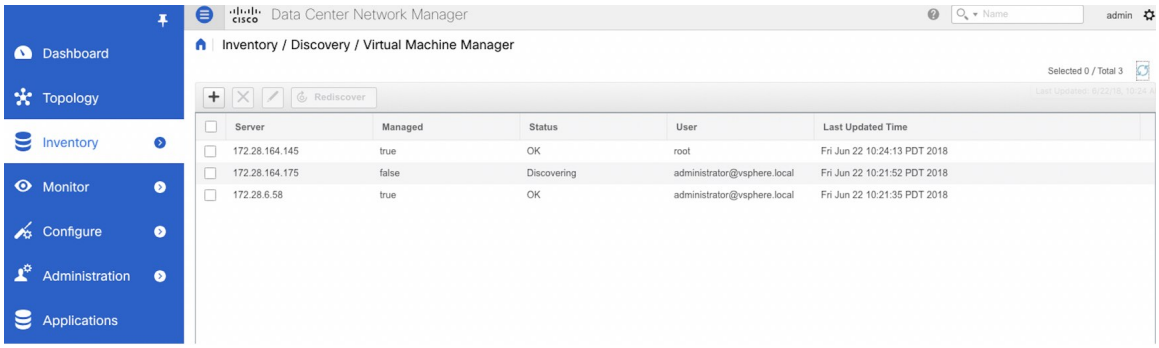

Step 2 Click the + icon to add a new VMware vSphere vCenter.

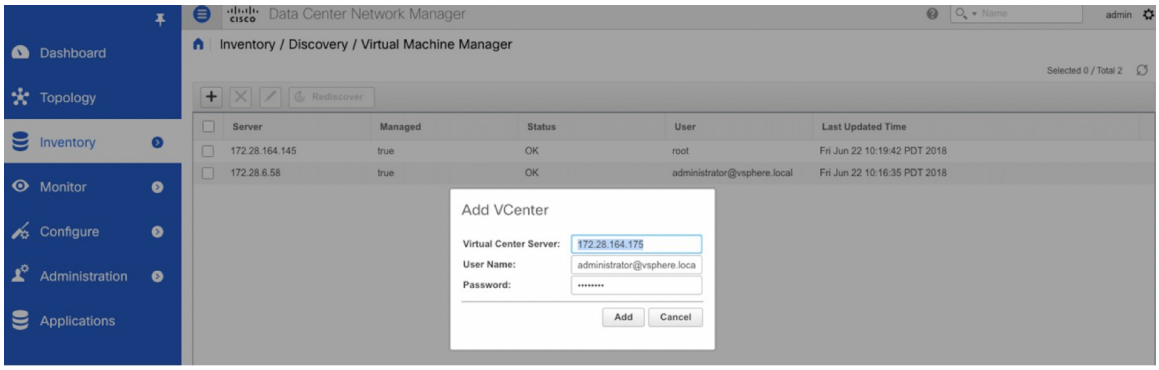

Step 3 Enter the server IP address, username, and password to the vCenter. vCenter version 5.5 or later is required.

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.

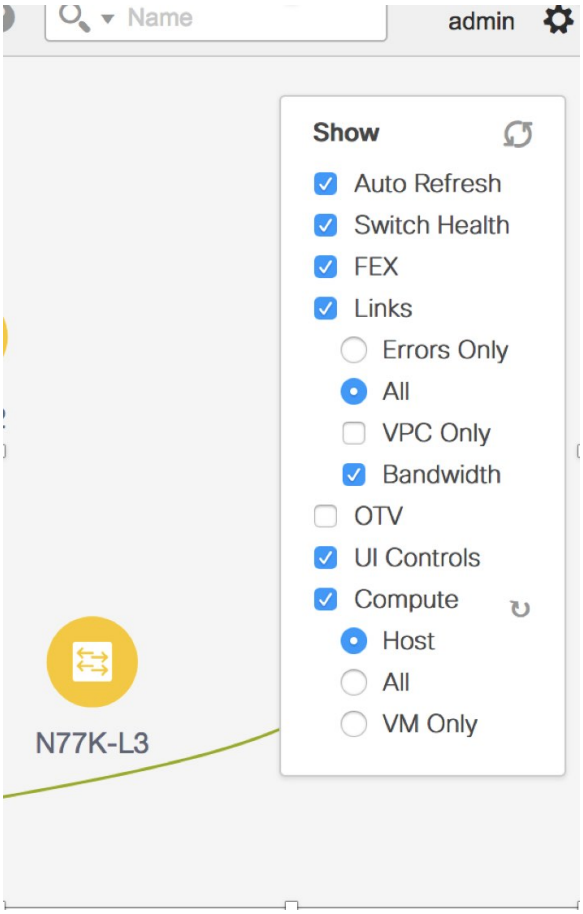

Note When you add a vCenter, you can run into a situation where an image upload is in progress, and thus the compute visualization is not complete. The **Topology** window displays the following message for more than 10 minutes:

"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**

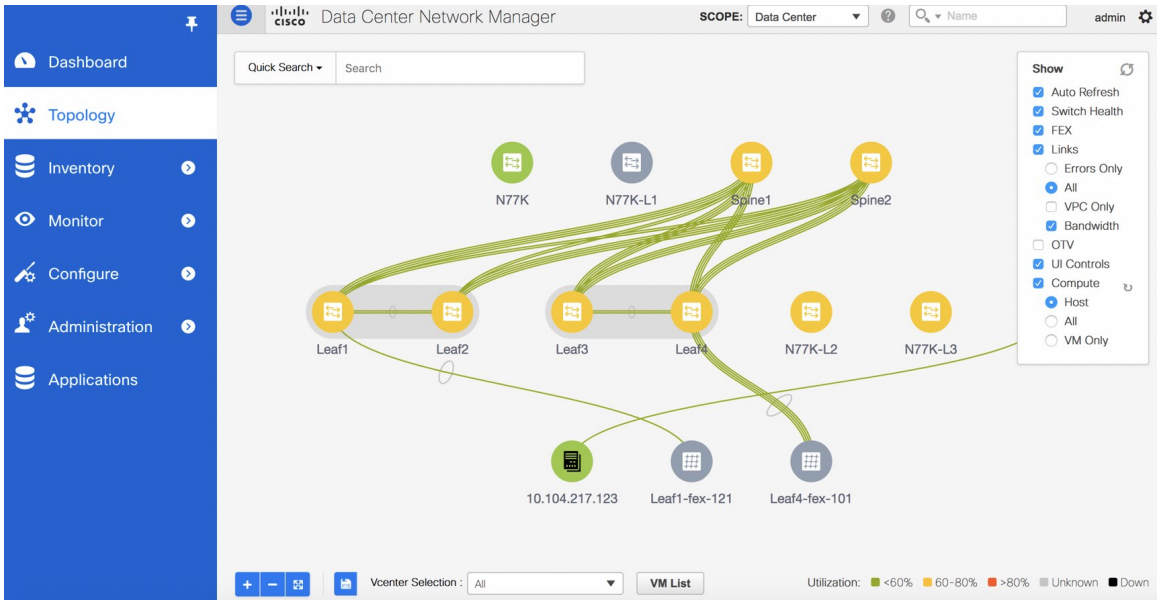

In the **All** mode, you can see double-arrows that help you to extend a node. If you double-click this node, you can see all the hidden child nodes.

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.

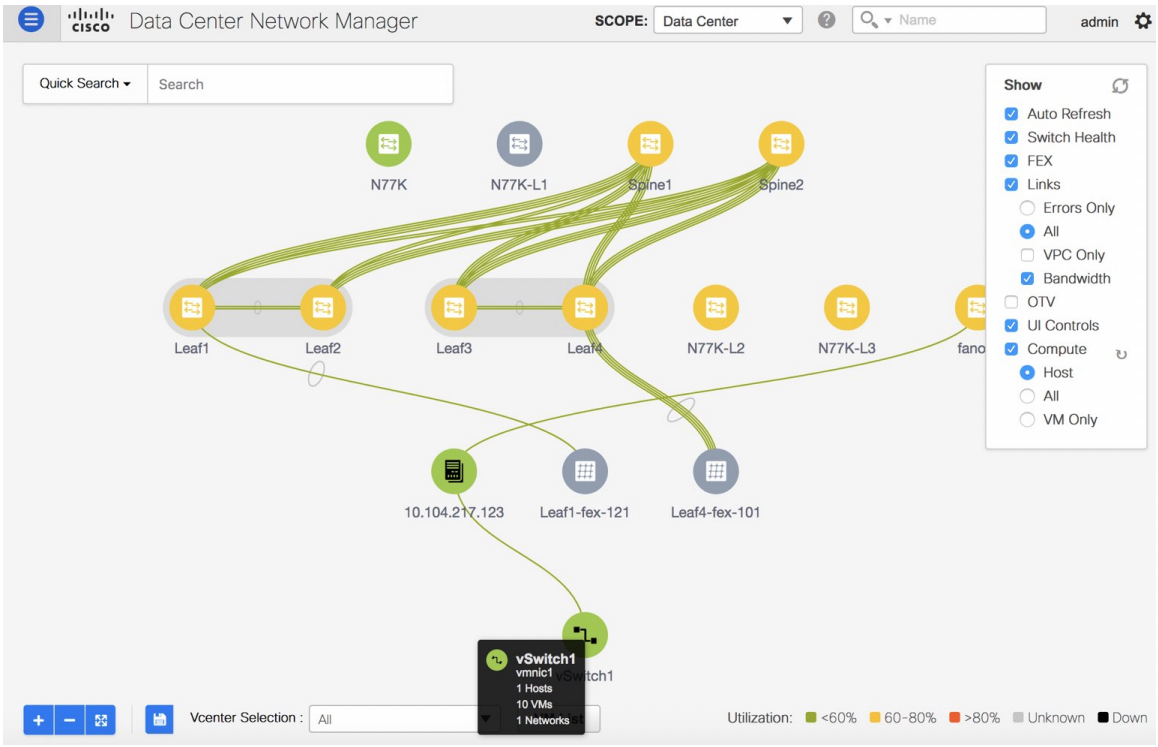

Step 4 When changing from the **Host** suboption to the **All** suboption, all the compute resources are expanded.

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**.

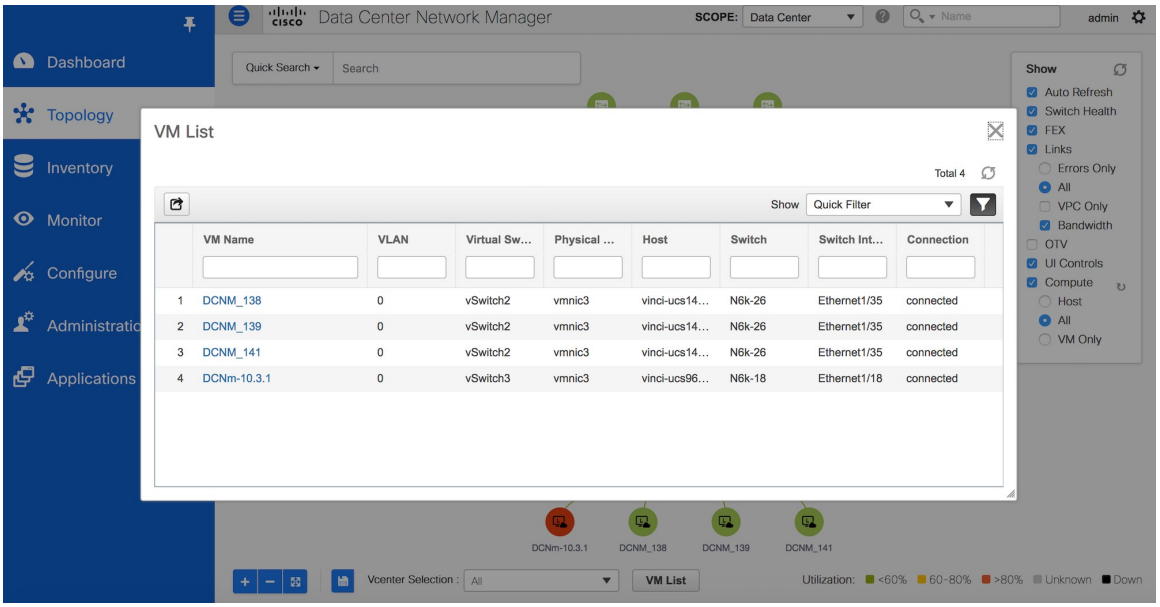

Click **Export** to export the list of virtual machines into a .csv file.

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

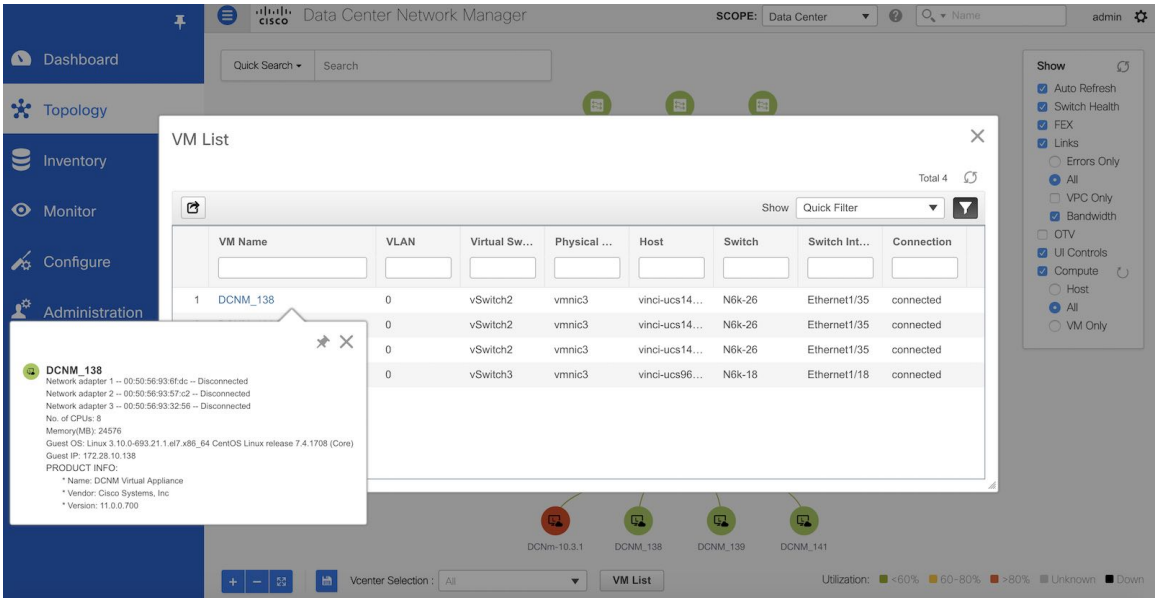

Note When you export the VM List to a .CSV file, the .CSV file may appear correct. However, when the .CSV file is imported into Microsoft Excel, it might get reformatted, for example, the VLAN column 1-1024 could be reformatted to a date 1/1/2019. Therefore ensure that columns are formatted correctly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

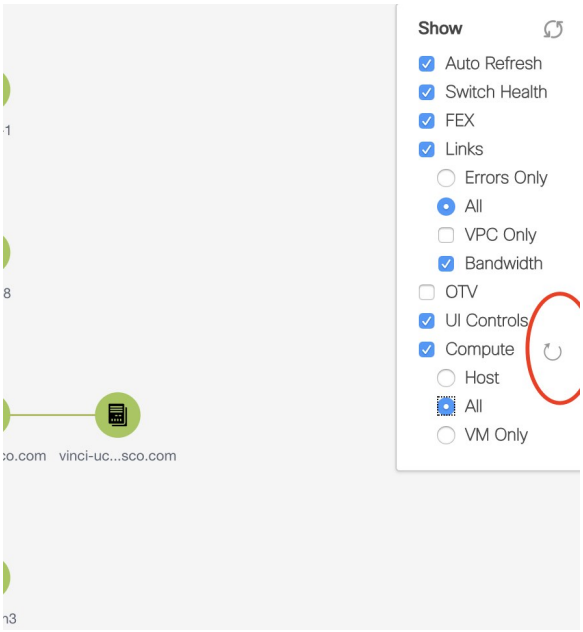

Selecting a Column in the Virtual Machine List

Procedure

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

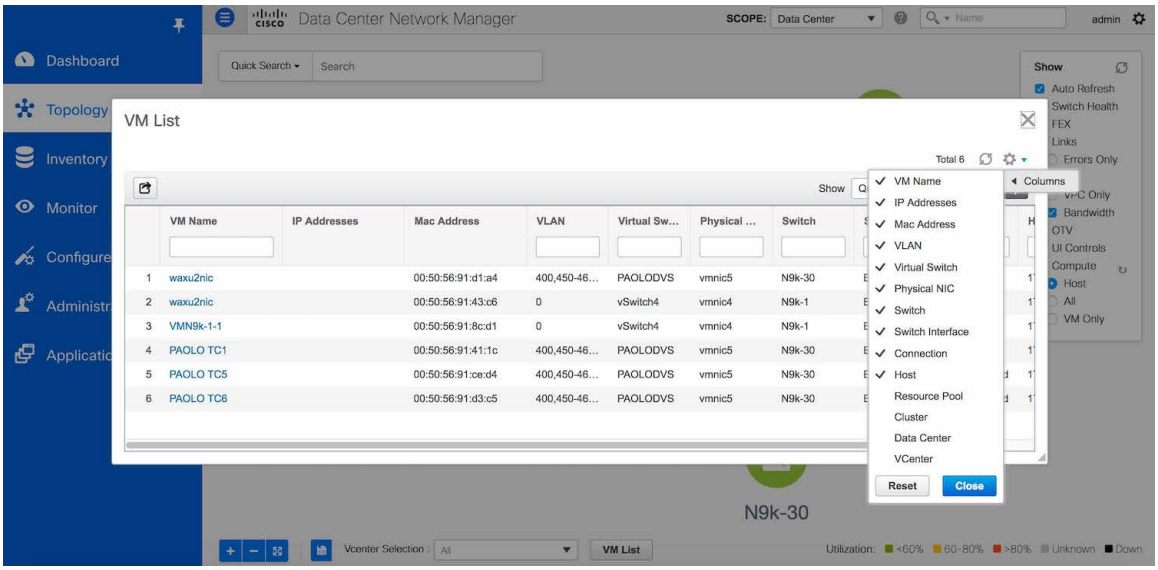

Step 2 Select the columns that you want to display in the VM list table. If you select additional columns, click **Resync vCenters** icon to refresh and view the new columns.

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.

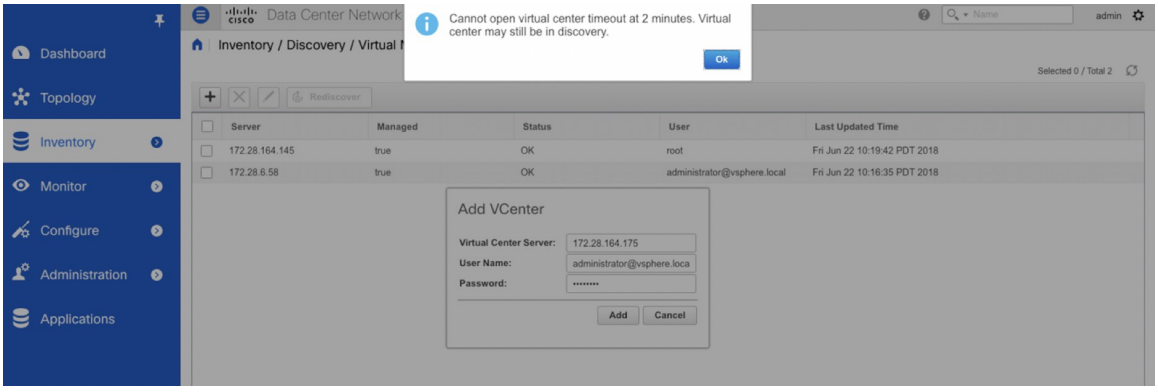

Viewing Topology in Scale Mode

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.

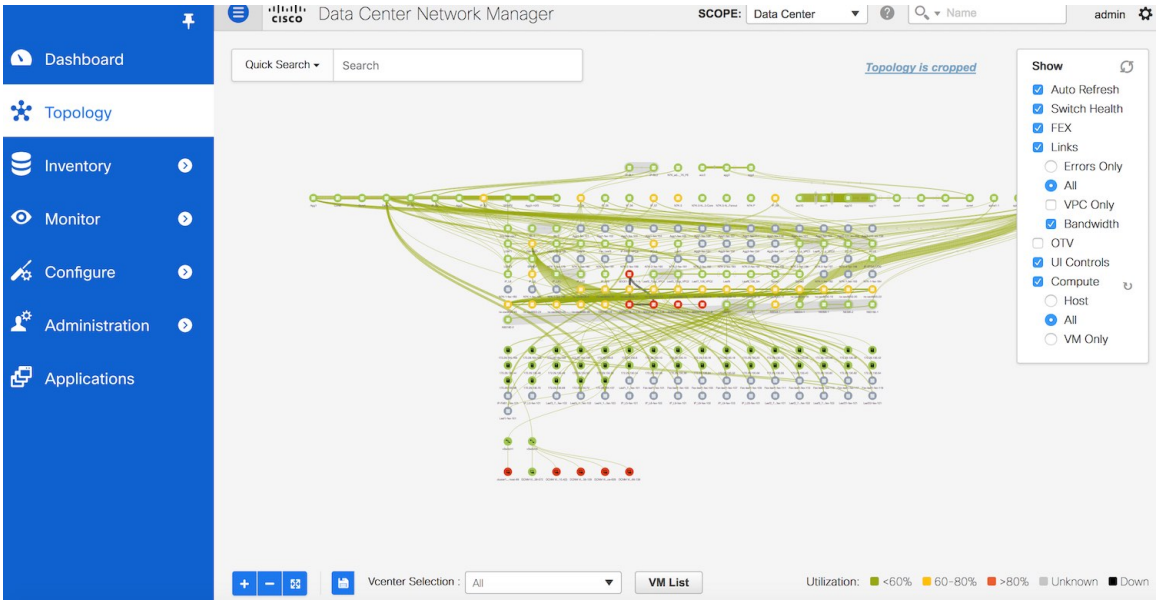

Container Orchestrator

From Release 11.3(1), Cisco DCNM allows you to configure Container Orchestrator. This feature allows you to visualize Kubernetes cluster as Container Orchestrator with the Cisco DCNM.

Ensure that you have successfully configured the VMM on Cisco DCNM before enabling Container Orchestrator Visualization feature. However, you do not need VMM for Bare-metal based Kubernetes cluster.

The Container visualization begins only after the first Kubernetes cluster is added to the container orchestrator. The information that is received from Kubernetes is appropriately organized and displayed on the main Topology window. An extra menu item labeled **Containers** appears on the Show pane.

At any point of time, click on any component in the **Topology** to view all the network paths between the selected component and fabric.

The following image details all the various features for Container Orchestrator Visualization on Cisco DCNM.

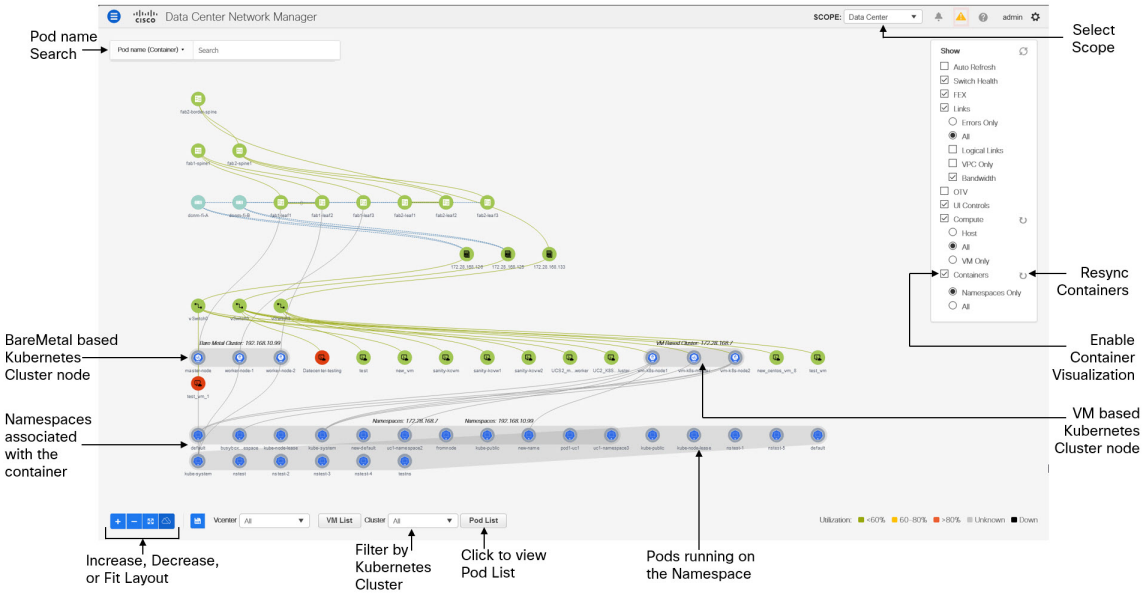

You can view the Container Orchestrator visualization based on the following important options:

- **Refresh**: Click on this icon to refresh the topology data.
- **Auto Refresh**: Select this checkbox to refresh the topology automatically.
- **Switch Health**: Select this checkbox to view the switch health status.
- **Links**: Select 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.
- **UI controls**: Select this check box to show all the various controls on the Topology window.
- **Compute**: Check the check box to enable the VCenter Compute Visualization.
  - Select **Host** to display the Compute hosts.
  - Select **All** to display all the compute nodes.
  - Select **VM Only** to display only the VMs.

  Resync to resynchronize the topology by clicking the Resync icon located next to Compute in the Show panel.
- **Containers**: Check the check box to display the containers.
  - By default, **Namespaces only** is selected to display only the Namespaces in the Kubernetes cluster

- Select **All** to view both Namespaces and Pods associated with the namespaces.

- **Resync**: You can also resynchronize the topology by clicking the **Resync** icon located next to Containers in the Show panel.

  We recommend that you wait for few minutes before resynchronizing the Containers, after Compute Resync is complete.

Using the UI Controls on Container Orchestrator Visualization

After you enable Containers on the Show panel, the cluster nodes are displayed with their Namespaces associations. For VM based Kubernetes cluster, based on the Compute selection (Hosts or VM Only or All), the topology displays the Kubernetes Clusters and associated Namespaces. For Bare-metal based Kubernetes cluster, compute selection is not required.

The icon indicates the Kubernetes nodes. The Kubernetes installation type and IP address are displayed on the Kubernetes Cluster. The icon indicates the Namespaces in the Kubernetes Cluster, and the icon indicates the Pods associated with the Namespace.

Kubernetes Clusters are of two types:

- VM based Kubernetes clusters are hosted on the VCenter.
- Kubernetes installed on Bare-metal, which is directly connected to a Switch.

On the UI Controls, from the Cluster Selection drop-down list, you can select one cluster to view the Container Visualization on that Cluster.

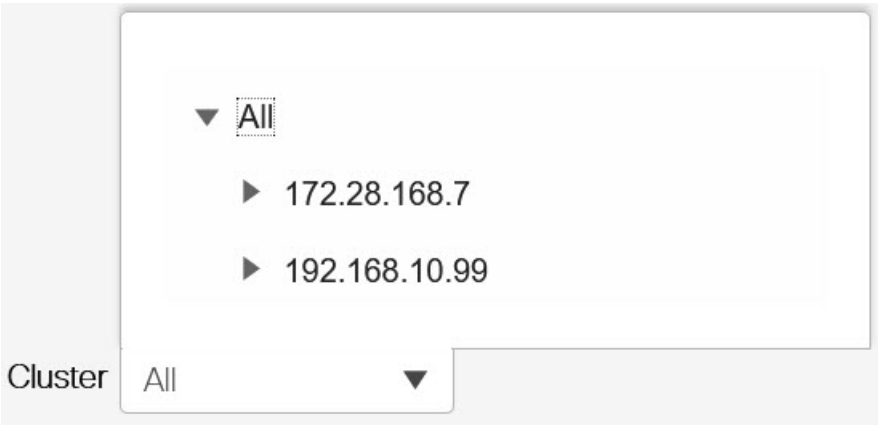

The topology now displays the Container Visualization only for the selected Cluster. Note that the other Kubernetes Clusters icon changes to the VM icon.

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

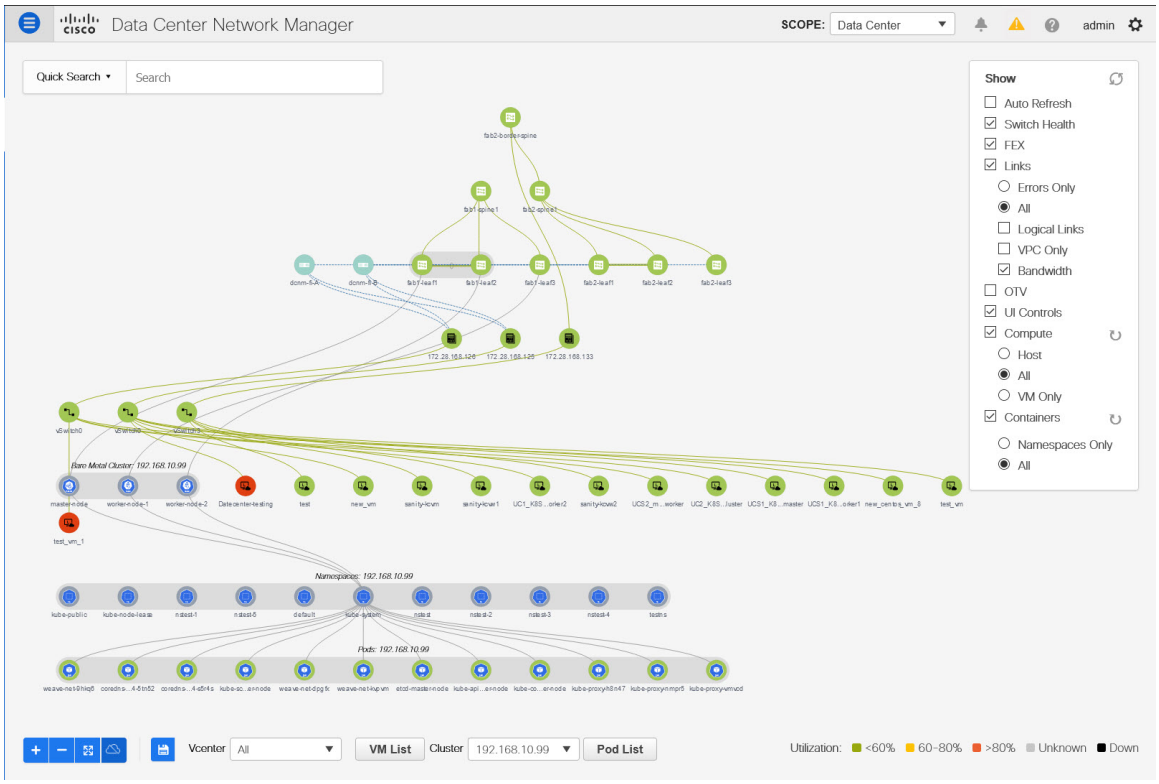

Double-click on the Nodes to view details about the node. A side panel appears, showing the Node Summary. Click **Show More Details** to view MetaData, Specifications, and Status information for the selected node.

Metadata tab consists of Kubernetes node or Pod name. Specifications tab include the desired design or configuration of the node or the Pod. Status tab indicates the running state information of the node or the pod.

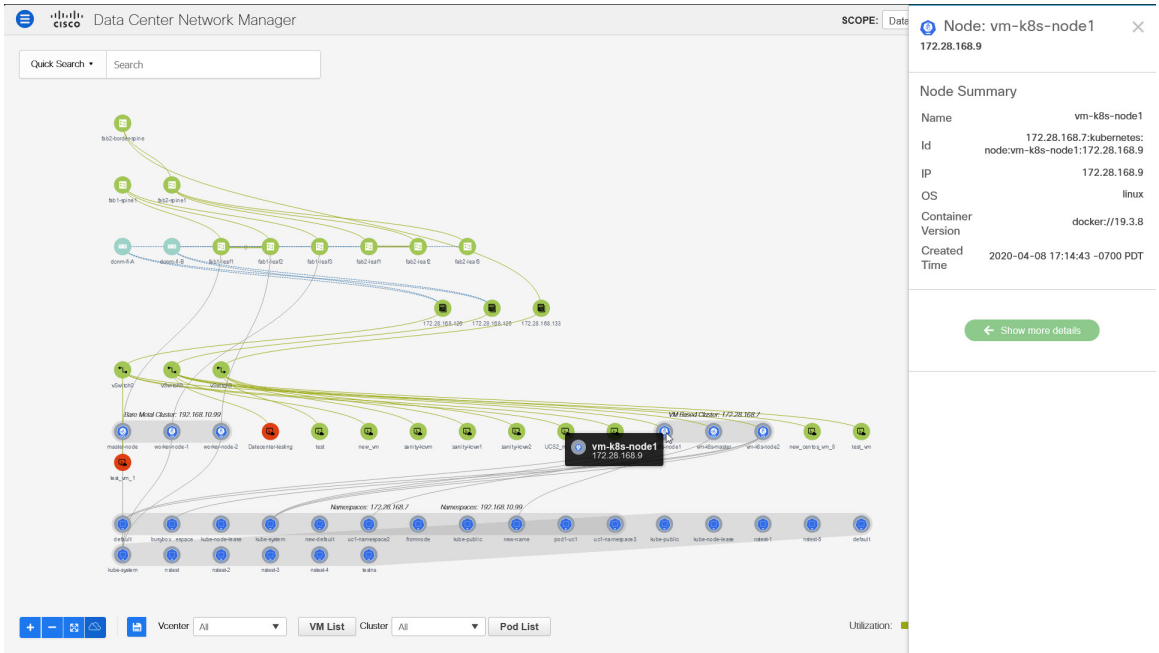

Double-click on the Namespaces to view the pods running on it. Double click on the Namespace again, to collapse the Pods associated with the Namespace.

Select **All** under the Containers on the Show panel to view all the Pods running on all the Namespaces. If there are more than 200 Pods, a new cropped view of the Topology is displayed with 5 Namespaces per Cluster, and 5 Pods per Namespace, to avoid rendering delays. An indicator appears to indicate that the Topology is cropped. You must view the Pod list to see the complete Topology details. You can also export the Pod List data for further analysis.

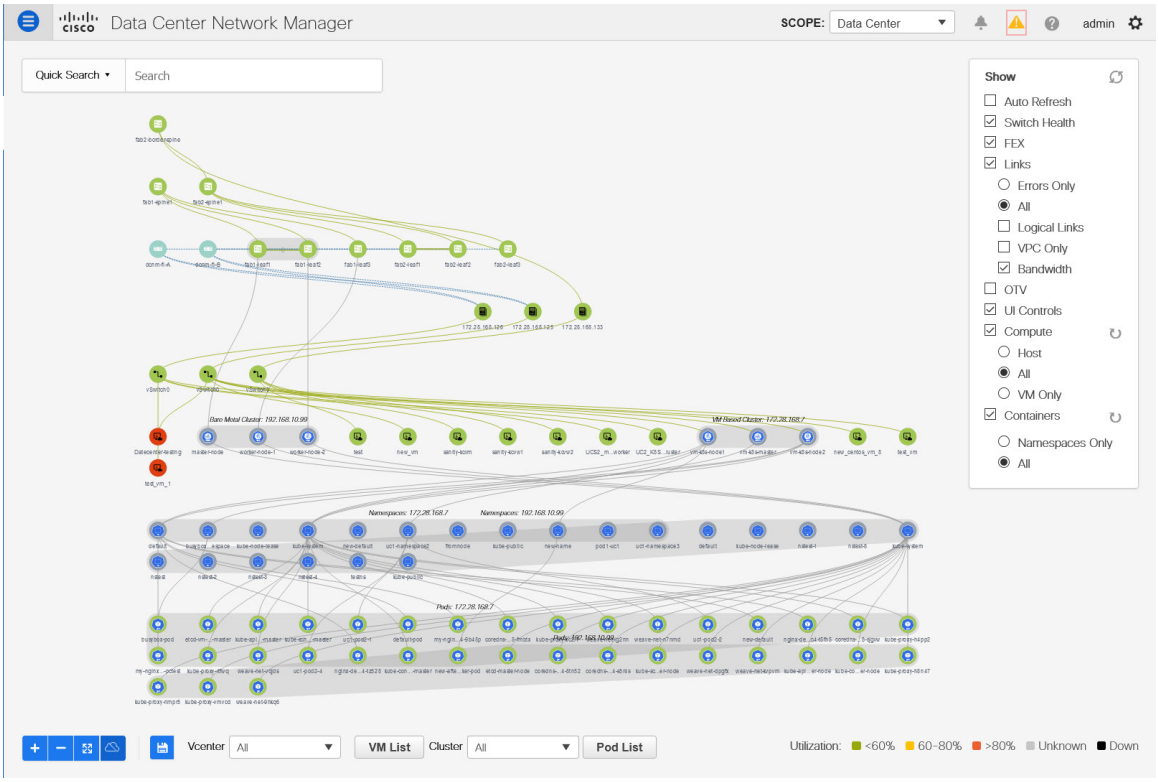

Double-click on the Pods to view details about the pod. A side panel appears, showing the Pod Summary. Click **Show More Details** to view MetaData, Specifications, and Status information for selected pod, specific to the Namespace to which it is connected.

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

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.

Pod Name (Container) Topology Search

You can locate the Pods using the Topology Search. From the Topology **Search** drop-down list, select **Pod name (Container)** Search type. In the Search field, enter the pod name. The pods and the namespace to which is connected in highlighted in the Topology.

Note The Pod search is partial unless you enter the exact name of the Pod. All the pods starting with the search string are highlighted.

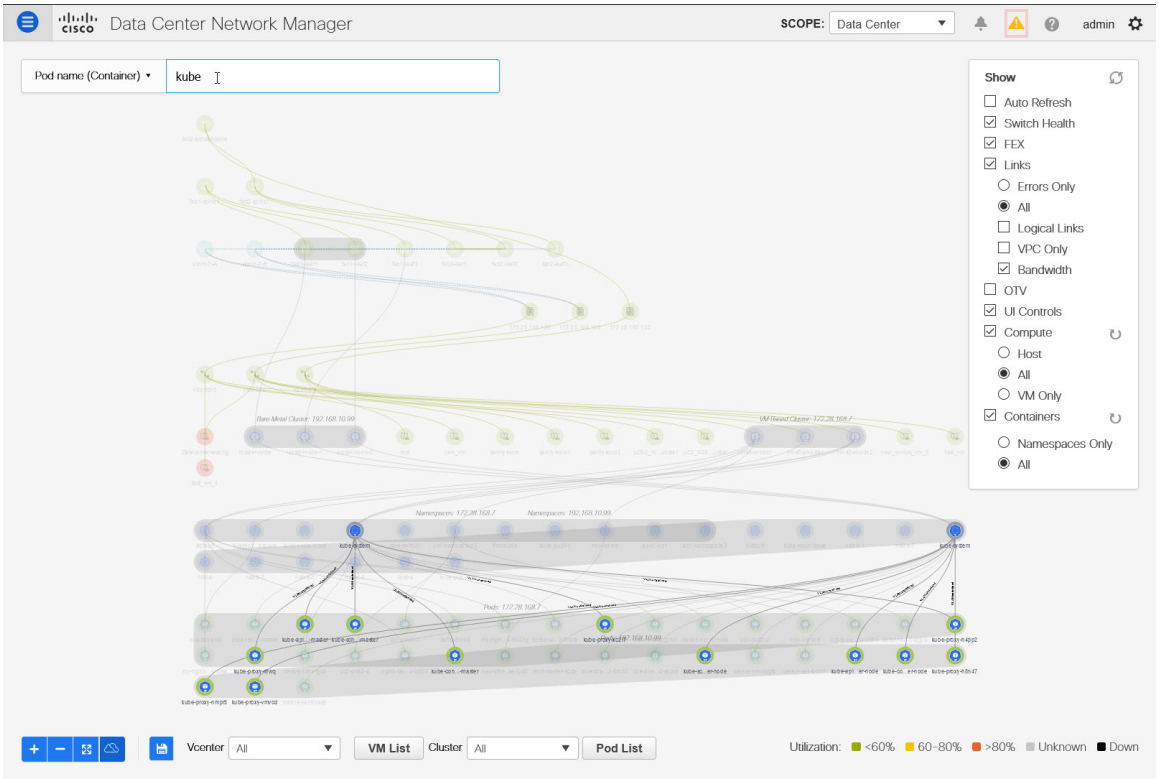

OpenStack Workload Visibility

Typically, a data center deployment comprised of various kinds of workloads such as bare-metal servers and VMs. Any kind of troubleshooting begins with the location of one or more servers on which the workload resides followed by the investigation of the compute, storage, and network resources that are servicing that workload.

From Cisco DCNM Release 11.5(1), OpenStack plugin is provided by DCNM that helps you to monitor OpenStack Clusters. You can get visibility with respect to the physical network connectivity and virtualized workloads, and debug VM networking-specific issues within the context of the data center.

Guidelines

- You can't start or stop the OpenStack application from the DCNM application catalog. The OpenStack application starts after the addition of the first OpenStack cluster. The OpenStack application is stopped when you remove the last OpenStack cluster instance from the DCNM OpenStack inventory. Any intermediate deletion of OpenStack cluster instance does not have impact on running of OpenStack plugin application.
- Based on the auto-resync feature of the OpenStack application, it retrieves information from the cluster at the configured interval.

OpenStack Topology Scale

- If there are more than 100 OpenStack VMs, only 5 VMs per host are shown, the remaining is cropped out with a message displayed. The message shows the total number of hosts and VMs.

- The OpenStack plugin can monitor up to four OpenStack clusters.
- The OpenStack plugin can monitor up to 1000 VMs across four clusters, that is, 250 VMs per cluster.

Notifications and Triggers for OpenStack

- RabbitMQ notification (oslo.messaging) bus configuration should be completed on the OpenStack cluster.

  Make the following configuration changes in the OpenStack Nova service.

  Replace the parameter values as shown. The Nova configuration file is located at the path: /etc/nova/nova.conf.

```
[notifications]
notify_on_state_change=vm_and_task_state
default_level=INFO
notification_format=both

[oslo_messaging_notifications]
driver = messagingv2
transport_url=rabbit://guest:guest@X.X.X.X:5672/
topics=notifications
retry=-1
```

  Note

  - **transport_url** is the address of the RabbitMQ endpoint hosted on the server having IP X.X.X.X at port 5672. Replace it with the appropriate server IP address.
  - **guest:guest** is the username and password to connect to the endpoint.

  Also, open port 5672 by setting the appropriate 'iptables' rule so that the monitoring application client can connect to the port and read the notification data.

- OpenStack plugin receives and handles the real-time change notifications from the OpenStack cluster and updates the topology description information. The real-time change notifications are related to the change of state of VM (for example, adding, deleting, or updating a VM) and change of state of network (for example, shutting down of a link between VM and the virtual switch).
- Powering on of a cluster node reflects in the topology view. The corresponding node is added to the cluster view. Similarly, powering down of a cluster node reflects in the topology view. The corresponding node is removed from the cluster view.
- Adding or deleting a node (controller, compute, or storage) in the OpenStack cluster is reflected automatically in DCNM in the Topology cluster view.

Using OpenStack Visualizer

Before you begin

Make sure to add an OpenStack cluster in DCNM. For more information, see OpenStack Visualizer, on page 369.

Procedure

Step 1 Navigate to **Topology**.

Step 2 In the **Show** panel, select the **OpenStack** check box to display only the OpenStack cluster nodes connecting with the fabric without displaying any of the VMs inside the cluster. This is host only view. The nodes are displayed as grouped according to the cluster. The following options are available under OpenStack:

- Select **Host** to display only OpenStack cluster hosts.
- Select **All** to display all OpenStack cluster nodes and the VMs instances hosted on the cluster nodes.
- Select **VM Only** to display only OpenStack VMs instances.

  You click the **resync** icon next to OpenStack to resync all the clusters. This icon is disabled until the resync operation is complete.

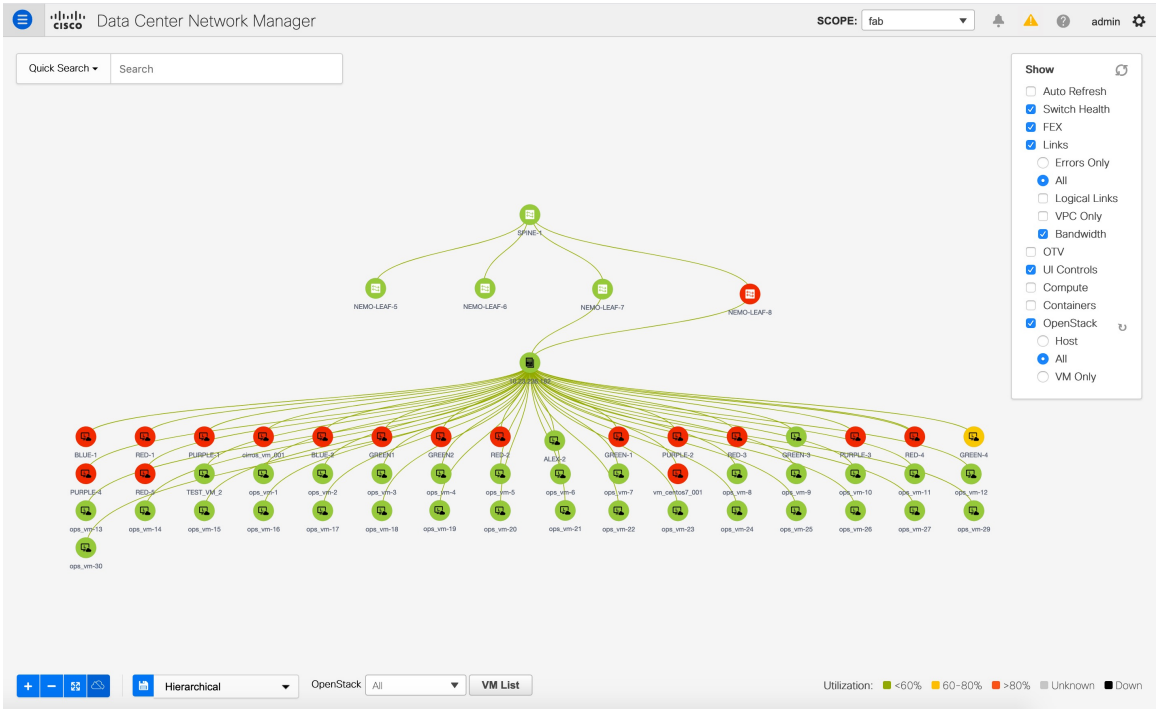

The color coding of each VM node corresponds to its state. The colors and what they indicate are described in the following list:

- Green: Indicates that the element is in good health and functioning as intended.
- Yellow: Indicates that the element is paused and suspended.
- Red: Indicates that the element is stopped, or it is shut down.

Step 3 From the **OpenStack** drop-down list, select **All** to display all OpenStack cluster nodes from all the OpenStack clusters that exist.

Select a specific OpenStack cluster IP address from this drop-down list to display all the OpenStack cluster nodes from the selected OpenStack cluster. Note that a single cluster is identified by the grey color grouping of its nodes.

Step 4 From the **SCOPE** drop-down list, select a fabric. This action shows the OpenStack cluster if it is connected to the selected fabric.

Step 5 Hover over the topology to see a tooltip information popup that shows more details of the given OpenStack resource. This action is also be applicable to the edges connecting the two resources. The following details are displayed:

- The tooltip information for an OpenStack cluster Host contains Host IP and version.
- The tooltip information for an OpenStack VM displays the VM Name, VM status: Locked/ Unlocked, VM IP, MAC, and VLAN.
- Hover over a link between a host and a VM to get information about it such as IP address of a host and VM, and the MAC address details.
- Click and hold a resource to see all the edges connecting to it.

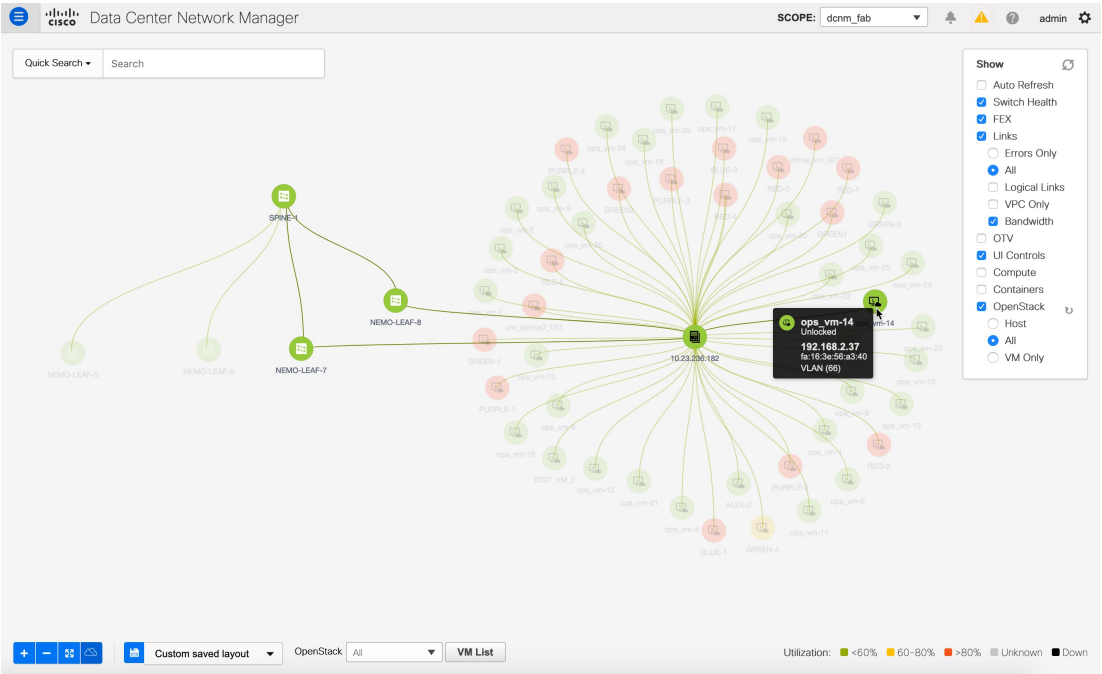

Viewing VMs in OpenStack Clusters

Procedure

Step 1 Double-click a VM to view summarized OpenStack VM data such as its name, IP, MAC, memory, segment type, locked or not, power, state, and vCPU info. Click **Show more details** to view the information such as operational state, specifications, and meta data of a VM.

Note A VM can have two interfaces connecting to a host. In this case, the VM is connecting to two different networks on the host with two different IP addresses.

Step 2 Click **VM List** to display a tabular view of the cluster topology.

The **VM List(OpenStack)** window displays the following details of VMs:

- VM Name, its IP address and MAC address
- Host name that is connected to a VM
- Switch name that is connected to a VM, switch's IP address, MAC address, and interface
- Port channel ID and VPC ID
- VLAN segment type
- Power state and status of the VM
- Allocated memory and vCPUs

You can search and filter VMs by using the search fields under each column. Click **Export** to export this data into a .CSV file.

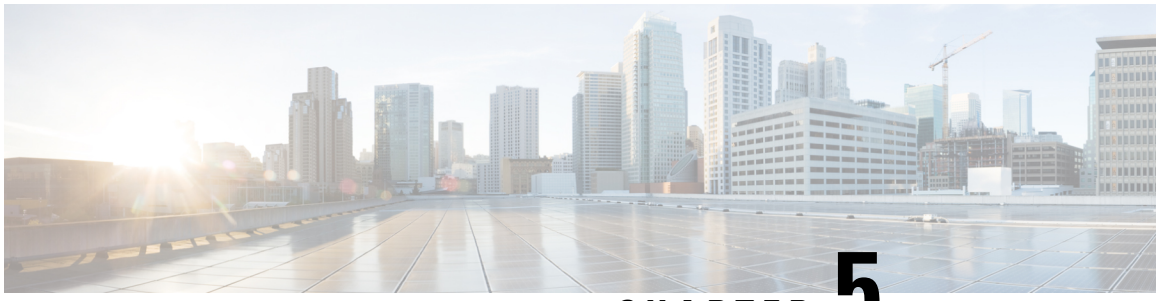

CHAPTER 5

Control

This chapter contains the following topics:

- Fabrics, on page 51
- Management, on page 359
- Template Library, on page 372
- Image Management, on page 413
- Endpoint Locator, on page 434
- Layer 4-Layer 7 Service, on page 434
- Cross Site Scripting (XSS) threat and mitigation, on page 434

Fabrics

The following terms are referred to in the document:

- Greenfield Deployments: Applicable for provisioning new VXLAN EVPN fabrics, and eBGP based Routed fabrics.
- Brownfield Deployments: Applicable for existing VXLAN EVPN fabrics:
  - Migrate CLI configured VXLAN EVPN fabrics to DCNM using the **Easy_Fabric_11_1** fabric template.
  - NFM migration to Cisco DCNM using the **Easy_Fabric_11_1** fabric template.

For information about upgrades, refer to the *Cisco DCNM Installation and Upgrade Guide for LAN Fabric Deployment*.

This section contains the following topics:

VXLAN BGP EVPN Fabrics Provisioning

DCNM 11 introduces an enhanced “Easy” fabric workflow for unified underlay and overlay provisioning of VXLAN BGP EVPN configuration on Nexus 9000 and 3000 series of switches. The configuration of the fabric is achieved via a powerful, flexible, and customizable template-based framework. Using minimal user inputs, an entire fabric can be brought up with Cisco recommended best practice configurations, in a short period of time. The set of parameters exposed in the Fabric Settings allow users to tailor the fabric to their preferred underlay provisioning options.

Border devices in a fabric typically provide external connectivity via peering with appropriate edge/core/WAN routers. These edge/core routers may either be managed or monitored by DCNM. These devices are placed in a special fabric called the External Fabric. The same DCNM controller can manage multiple VXLAN BGP EVPN fabrics while also offering easy provisioning and management of Layer-2 and Layer-3 DCI underlay and overlay configuration among these fabrics using a special construct called a Multi-Site Domain (MSD) fabric.

Note that in this document the terms switch and device are used interchangeably.

The DCNM GUI functions for creating and deploying VXLAN BGP EVPN 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).

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]).

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 (ST-FEX), Active-Active FEX (AA-FEX), loopback, subinterface, etc.
- Create breakout and unbreakout ports.
- Shut down and bring up interfaces.
- Rediscover ports and view interface configuration history.

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.

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

Provisioning of configuration on service leafs to which L4-7 service appliances may be attached. For more information, see *L4-L7 Service Basic Workflow*.

This chapter mostly covers configuration provisioning for a single VXLAN BGP EVPN fabric. EVPN Multi-Site provisioning for Layer-2/Layer-3 DCI across multiple fabrics using the MSD fabric, is documented in a separate chapter. The deployment details of how overlay Networks and VRFs can be easily provisioned from the DCNM, is covered under Creating and Deploying Networks and VRFs .

Guidelines for VXLAN BGP EVPN Fabrics Provisioning

- For any switch to be successfully imported into DCNM, the user specified for discovery/import, should have the following permissions:
  - SSH access to the switch
  - Ability to perform SNMPv3 queries
  - Ability to run the **show** commands including show run, show interfaces, etc.
- The switch discovery user need not have the ability to make any configuration changes on the switches. It is primarily used for read access.
- 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.
- LAN credentials are required to be set of any user that needs to be perform any write access to the device. LAN credentials need to be set on the DCNM, on a per user per device basis. When a user imports a device into the Easy Fabric, and LAN credentials are not set for that device, DCNM moves this device to a migration mode. Once the user sets the appropriate LAN credentials for that device, a subsequent Save & Deploy will retrigger the device import process.
- The **Save & Deploy** button triggers the intent regeneration for the entire fabric as well as a configuration compliance check for all the switches within the fabric. This button is required but not limited to the following cases:
  - A switch or a link is added, or any change in the topology
  - A change in the fabric settings that must be shared across the fabric
  - A switch is removed or deleted
  - A new vPC pairing or unpairing is done
  - A change in the role for a device

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

  When you right-click on a switch icon, you can use the **Deploy Config** option to deploy per switch configurations. This option is a local operation for a switch, that is, the expected configuration or intent for a switch is evaluated against it's current running configuration, and a config compliance check is

performed for the switch to get the **In-Sync** or **Out-of-Sync** status. If the switch is out of sync, the user is provided with a preview of all the configurations running in that particular switch that vary from the intent defined by the user for that respective switch.

- 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
  ```

  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*.

  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.

- 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.

This procedure contains descriptions for the IPv4 underlay. For information about IPv6 underlay, see IPv6 Underlay Support for Easy Fabric, on page 114.

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

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

   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 appear.

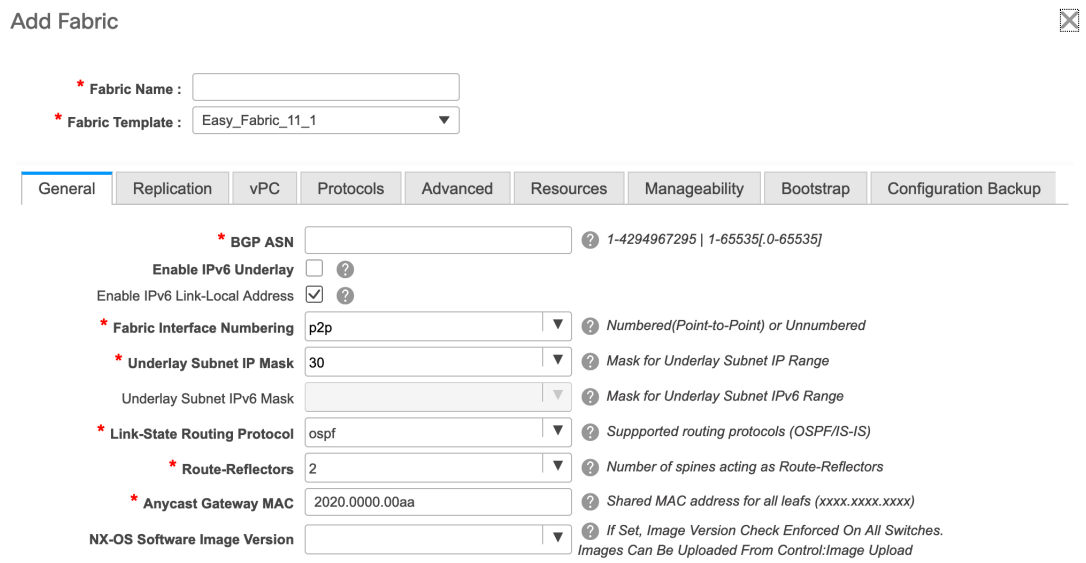

   The tabs and their fields in the screen are explained in the subsequent points. The overlay and underlay network parameters are included in these tabs.

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

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

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

Enable IPv6 Underlay: Enable the IPv6 underlay feature. For information, see IPv6 Underlay Support for Easy Fabric, on page 114.

Enable IPv6 Link-Local Address: Enables the IPv6 Link-Local address.

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.

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

Route-Reflectors (RRs) – 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 two spine devices designated as RRs.

Decreasing the count - When you reduce four route reflectors to two, 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 & Deploy** in the fabric topology window.

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, and save the Fabric Settings, the system checks that all the switches within the fabric have the selected version. If some devices do not run the image, a warning is prompted to perform an In-Service Software Upgrade (ISSU) to the specified image. The warning is also accompanied with a Resolve button. This takes the user to the image management screen with the mismatched switches auto selected for device upgrade/downgrade to the specified NX-OS image specified in Fabric Settings. Till, all devices run the specified image, the deployment process is incomplete.

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

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

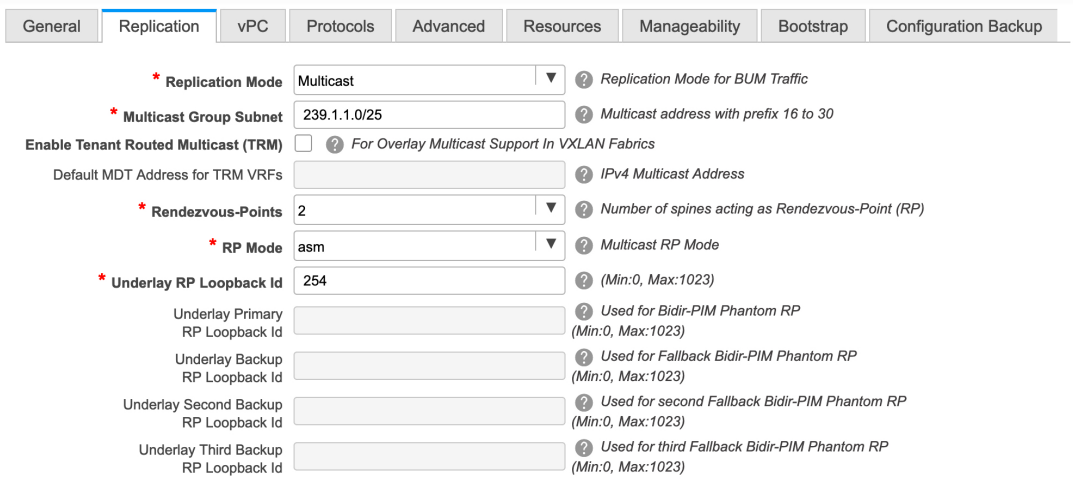

Replication Mode : The mode of replication that is used in the fabric for BUM (Broadcast, Unknown Unicast, Multicast) traffic. The choices are Ingress Replication or Multicast. When you choose Ingress replication, the multicast related fields get disabled.

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. A 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 check box to enable Tenant Routed Multicast (TRM) that allows overlay multicast traffic to be supported over EVPN/MVPN in the VXLAN BGP EVPN fabric.

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**.

For more information, see Overview of Tenant Routed Multicast, on page 196.

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 sof release 9.2(1) onwards.

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

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

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

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

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

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

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

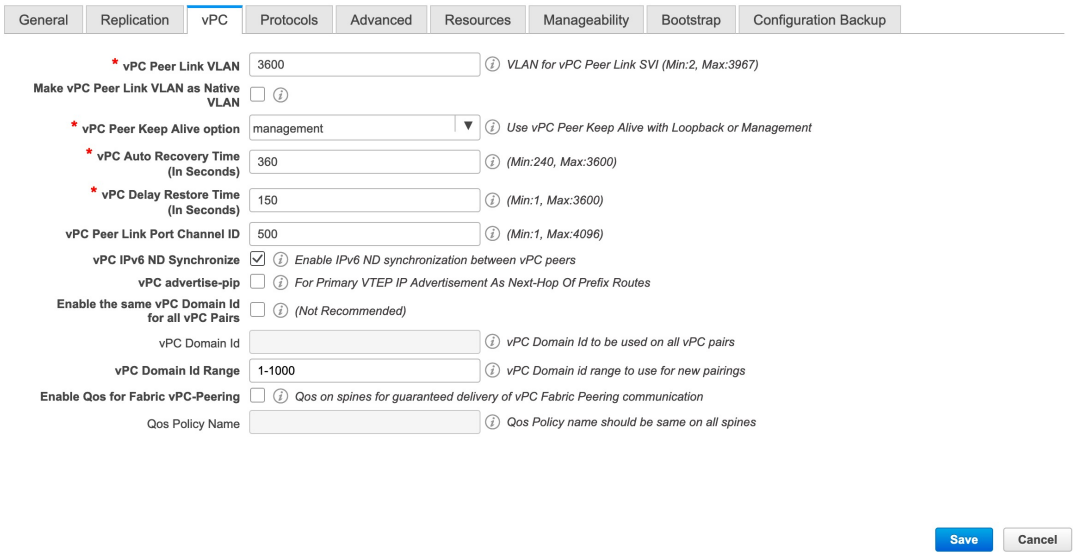

vPC Peer Link VLAN – VLAN used for the vPC peer link SVI.

Make vPC Peer Link VLAN as Native VLAN - Enables vPC peer link VLAN as Native VLAN.

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 ID - 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.

You can enable the advertise PIP feature on a specific vPC as well. For more information, see Advertising PIP on vPC, on page 240.

Enable the same vPC Domain Id for all vPC Pairs: Enable the same vPC Domain ID for all vPC pairs. When you select this field, the **vPC Domain Id** field is editable.

vPC Domain Id - Specifies the vPC domain ID to be used on all vPC pairs.

vPC Domain Id Range - Specifies the vPC Domain Id range to use for new pairings.

Enable Qos for Fabric vPC-Peering - Enable QoS on spines for guaranteed delivery of vPC Fabric Peering communication. For more information, see QoS for Fabric vPC-Peering, on page 233.

Note QoS for vPC fabric peering and queuing policies options in fabric settings are mutually exclusive.

Qos Policy Name - Specifies QoS policy name that should be same on all fabric vPC peering spines. The default name is **spine_qos_for_fabric_vpc_peering**.

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

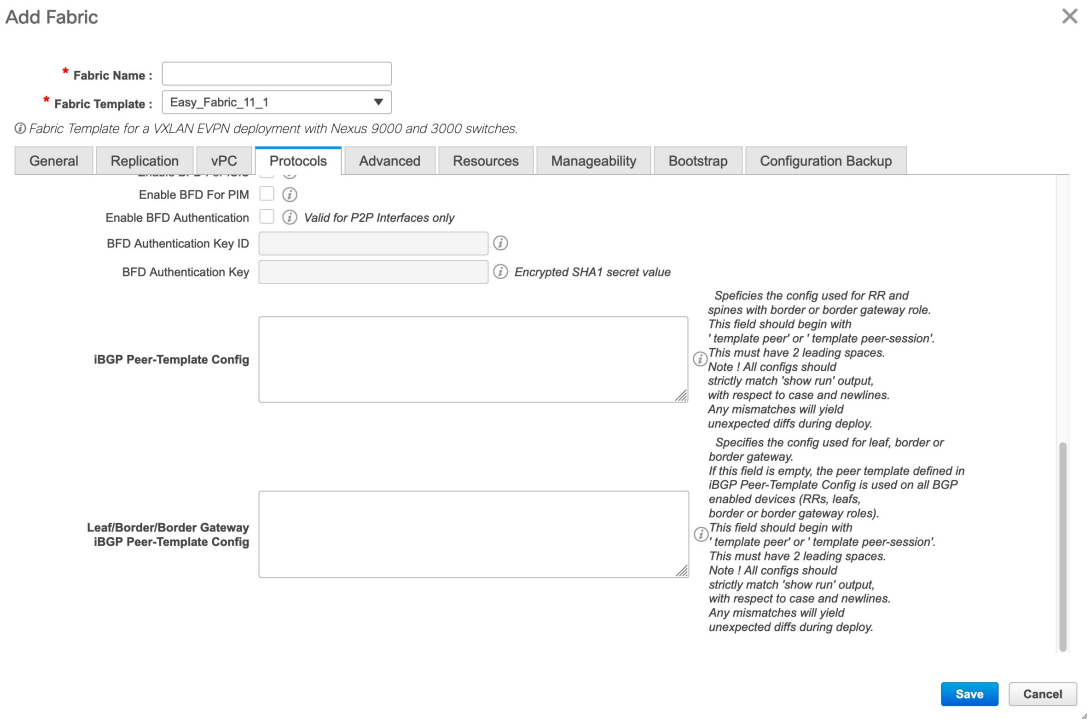

Underlay Routing Loopback Id - The loopback interface ID is populated as 0 since loopback0 is usually used for fabric underlay IGP peering purposes.

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

Underlay 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 **Underlay Routing Protoc** 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. Log in to the switch, retrieve the encrypted key and enter it in this field. Refer, *Retrieving the Authentication Key* section for details.

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, such as CiscoisisAuth.

IS-IS Authentication Key ID - The Key ID is populated.

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

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

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 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. Log in to the switch, retrieve the encrypted key and enter it in BGP Authentication Key field. Refer the Retrieving the Authentication Key section for details.

Enable PIM Hello Authentication - Enables the PIM hello authentication.

PIM Hello Authentication Key - Specifies the PIM hello authentication key.

Enable BFD: Select the check box to enable **feature bfd** on all switches in the fabric. This feature is valid only on IPv4 underlay and the scope is within a fabric.

From Cisco DCNM Release 11.3(1), BFD within a fabric is supported natively. The BFD feature is disabled by default in the Fabric Settings. If enabled, BFD is enabled for the underlay protocols with the default settings. Any custom required BFD configurations must be deployed via the per switch freeform or per interface freeform policies.

The following config is pushed after you select the **Enable BFD** check box:

```
feature bfd
```

Note After you upgrade from DCNM Release 11.2(1) with BFD enabled to DCNM Release 11.3(1), the follo configurations are pushed on all P2P fabric interfaces:

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

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 be 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 set to **unnumbered**. The BFD authentication fields will be grayed out automatically. BFD authentication valid for only for P2P interfaces.

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 Encrypted BFD Authentication Key, on page 249.

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
```

Until Cisco DCNM Release 11.3(1), iBGP peer template for iBGP definition on the leafs or border role devices and BGP RRs were same. From DCNM Release 11.4(1), the following fields can be used to specify different configurations:

- **iBGP Peer-Template Config** – Specifies the config used for RR and spines with border role.
- **Leaf/Border/Border Gateway iBGP Peer-Template Config** – Specifies the config used for leaf, border, or border gateway. If this field is empty, the peer template defined in **iBGP Peer-Template Config** is used on all BGP enabled devices (RRs, leafs, border, or border gateway roles).

In brownfield migration, if the spine and leaf use different peer template names, both **iBGP Peer-Template Config** and **Leaf/Border/Border Gateway iBGP Peer-Template Config** fields need to be set according to the switch config. If spine and leaf use the same peer template name and content (except for the “route-reflector-client” CLI), only **iBGP Peer-Template Config** field in fabric setting needs to be set. If the fabric settings on iBGP peer templates do not match the existing switch configuration, an error message is generated and the migration will not proceed.

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

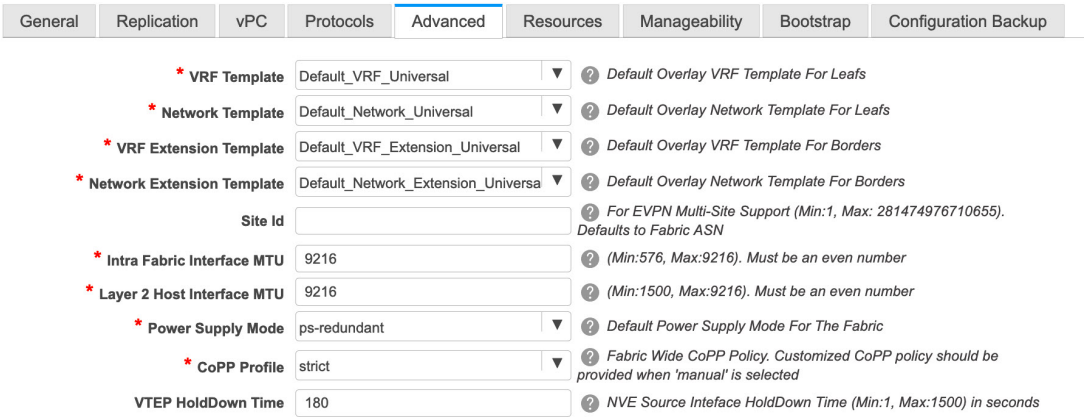

VRF Template and **VRF Extension Template**: Specifies the VRF template for creating VRFs, and the VRF extension template for enabling VRF extension to other fabrics.

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.

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

VTEP HoldDown Time - Specifies the NVE source interface hold down time.

Brownfield Overlay Network Name Format: Enter the format to be used to build the overlay network name during a brownfield import or migration. The network name should not contain any white spaces or special characters except underscore (_) and hyphen (-). The network name must not be changed once the brownfield migration has been initiated. See the *Creating Networks for the Standalone Fabric* section for the naming convention of the network name. The syntax is **[<string> | $$VLAN_ID$$] $$VNI$$ [<string>| $$VLAN_ID$$]** and the default value is **Auto_Net_VNI$$VNI$$_VLAN$$VLAN_ID$$**. When you create networks, the name is generated according to the syntax you specify. The following table describes the variables in the syntax.

Variables	Description
$$VNI$$	Specifies the network VNI ID found in the switch configuration. This is a mandatory keyword required to create unique network names.
$$VLAN_ID$$	Specifies the VLAN ID associated with the network. VLAN ID is specific to switches, hence DCNM picks the VLAN ID from one of the switches, where the network is found, randomly and use it in the name. We recommend not to use this unless the VLAN ID is consistent across the fabric for the VNI.
<string>	This variable is optional and you can enter any number of alphanumeric characters that meet the network name guidelines.

Example overlay network name: Site_VNI12345_VLAN1234

Note Ignore this field for greenfield deployments. The **Brownfield Overlay Network Name Format** applies fo the following brownfield imports:

- CLI-based overlays
- Configuration profile-based overlay where the configuration profiles were created in Cisco DCNM Release

  10.4(2).

Enable CDP for Bootstrapped Switch - Enables CDP on management (mgmt0) interface for bootstrapped switch. By default, for bootstrapped switches, CDP is disabled on the mgmt0 interface.

Enable VXLAN OAM - Enables the VXLAM OAM functionality for devices in the fabric. 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 check box to enable feature dhcp and associated configurations globally on all switches in the fabric. This is a pre-requisite for support of DHCP for overlay networks that are part of the tenant VRFs.

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

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

Enable NX-API on HTTP on Port - Specifies enabling of NX-API on HTTP. Enable this check box and the **Enable NX-API** check box to use HTTP. This check box is checked by default. If you uncheck this check box, the applications that use NX-API and supported by Cisco DCNM, such as Endpoint Locator (EPL), Layer 4-Layer 7 services (L4-L7 services), VXLAN OAM, and so on, start using the HTTPS instead of HTTP.

Note If you check the **Enable NX-API** check box and the **Enable NX-API on HTTP** check box, applicatio HTTP.

Enable Policy-Based Routing (PBR) - Select this check box to enable routing of packets based on the specified policy. Starting with Cisco NX-OS Release 7.0(3)I7(1) and later releases, this feature works on Cisco Nexus 9000 Series switches with Nexus 9000 Cloud Scale (Tahoe) ASICs. This feature is used along with the Layer 4-Layer 7 service workflow. For information on Layer 4-Layer 7 service, refer the *Layer 4-Layer 7 Service* chapter.

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.

Enable AAA IP Authorization - Enables AAA IP authorization, when IP Authorization is enabled in the remote authentication server. This is required to support DCNM in scenarios where customers have strict control of which IP addresses can have access to the switches.

Enable DCNM as Trap Host - Select this check box to enable DCNM as a SNMP trap destination. Typically, for a native HA DCNM deployment, the eth1 VIP IP address will be configured as SNMP trap destination on the switches. By default, this check box is enabled.

Greenfield Cleanup Option – Enable the switch cleanup option for switches imported into DCNM with Preserve-Config=No, without a switch reload. This option is typically recommended only for the fabric environments with Cisco Nexus 9000v Switches to improve on the switch clean up time. The recommended option for Greenfield deployment is to employ Bootstrap or switch cleanup with a reboot. In other words, this option should be unchecked.

Enable Precision Time Protocol (PTP): Enables PTP across a fabric. When you select this check box, PTP is enabled globally and on core-facing interfaces. Additionally, the **PTP Source Loopback Id** and **PTP Domain Id** fields are editable. For more information, see Precision Time Protocol for Easy Fabric, on page 100.

PTP Source Loopback Id: Specifies the loopback interface ID Loopback that is used as the Source IP Address for all PTP packets. The valid values range from 0 to 1023. The PTP loopback ID cannot be the same as RP, Phantom RP, NVE, or MPLS loopback ID. Otherwise, an error will be generated. The PTP loopback ID can be the same as BGP loopback or user-defined loopback which is created from DCNM.

If the PTP loopback ID is not found during **Save & Deploy**, the following error is generated:

Loopback interface to use for PTP source IP is not found. Create PTP loopback interface on all the devices to enable PTP feature.

PTP Domain Id: Specifies the PTP domain ID on a single network. The valid values range from 0 to 127.

Enable MPLS Handoff: Select the check box to enable the MPLS Handoff feature. For more information, see the *Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - MPLS SR and LDP Handoff* chapter.

Underlay MPLS Loopback Id: Specifies the underlay MPLS loopback ID. The default value is 101.

Enable TCAM Allocation: TCAM commands are automatically generated for VXLAN and vPC Fabric Peering when enabled.

Enable Default Queuing Policies: Check this check box to apply QoS policies on all the switches in this fabric. To remove the QoS policies that you applied on all the switches, uncheck this check box, update all the configurations to remove the references to the policies, and save and deploy. From Cisco DCNM Release 11.3(1), pre-defined QoS configurations are included that can be used for various Cisco Nexus 9000 Series Switches. When you check this check box, the appropriate QoS configurations are pushed to the switches in the fabric. The system queuing is updated when configurations are deployed to the switches. You can perform the interface marking with defined queuing policies, if required, by adding the required configuration to the per interface freeform block.

From Cisco DCNM Release 11.4(1), the DSCP mapping for QoS 5 has changed from 40 to 46 in the policy template. For DCNM 11.3(1) deployments that have been upgraded to 11.4(1), you will see the diffs that need to be deployed.

Review the actual queuing policies by opening the policy file in the template editor. From Cisco DCNM Web UI, choose **Control > Template Library**. Search for the queuing policies by the policy file name, for example, **queuing_policy_default_8q_cloudscale**. Choose the file and click the **Modify/View template** icon to edit the policy.

See the *Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide* for platform specific details.

N9K Cloud Scale Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all Cisco Nexus 9200 Series Switches and the Cisco Nexus 9000 Series Switches that ends with EX, FX, and FX2 in the fabric. The valid values are **queuing_policy_default_4q_cloudscale** and **queuing_policy_default_8q_cloudscale**. Use the **queuing_policy_default_4q_cloudscale** policy for FEXes. You can change from the **queuing_policy_default_4q_cloudscale** policy to the **queuing_policy_default_8q_cloudscale** policy only when FEXes are offline.

N9K R-Series Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all Cisco Nexus switches that ends with R in the fabric. The valid value is **queuing_policy_default_r_series**.

Other N9K Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all other switches in the fabric other than the switches mentioned in the above two options. The valid value is **queuing_policy_default_other**.

Enable MACsec - Enables MACsec for the fabric. For more information, see MACsec Support in Easy Fabric and eBGP Fabric, on page 194.

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. For more information, refer Enabling Freeform Configurations on Fabric Switches .

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

Intra-fabric Links Additional Config - Add CLIs that should be added to the intra-fabric links.

8. Click the **Resources** tab.

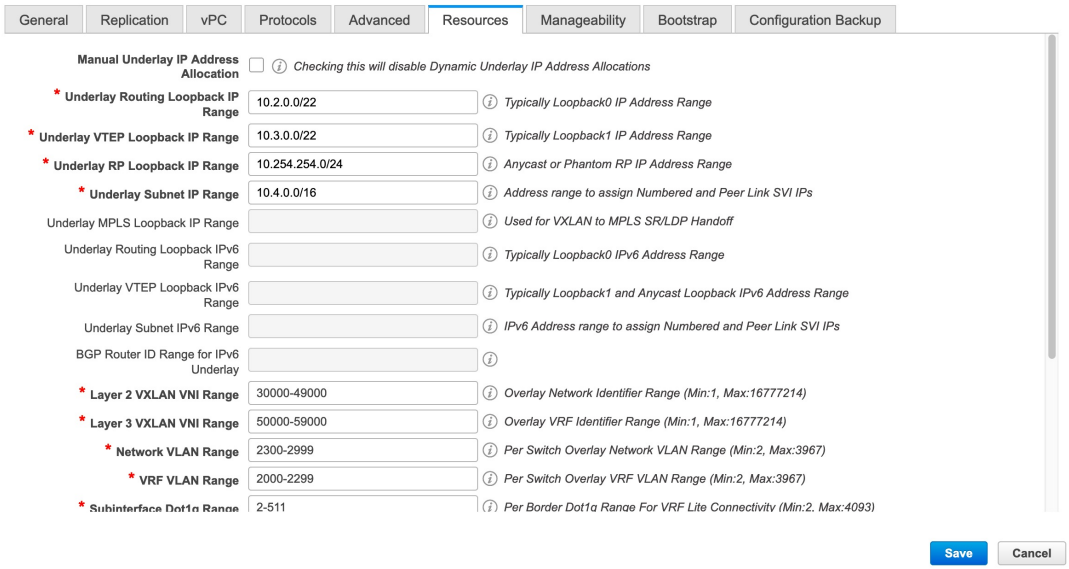

Manual Underlay IP Address Allocation – *Do not* select this check box if you are transitioning your VXLAN fabric management to DCNM.

- 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.

Underlay MPLS Loopback IP Range: Specifies the underlay MPLS loopback IP address range.

For eBGP between Border of Easy A and Easy B, Underlay routing loopback and Underlay MPLS loopback IP range must be a unique range. It should not overlap with IP ranges of the other fabrics, else VPNv4 peering will not come up.

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.

The VRF Lite Subnet IP Range field specifies resources reserved for IP address used for VRF LITE when VRF LITE IFCs are auto-created. If you select Back2BackOnly, ToExternalOnly, or Back2Back&ToExternal then VRF LITE IFCs are auto-created.

Auto Deploy Both - This check box is applicable for the symmetric VRF Lite deployment. When you select this check box, it would set the auto deploy flag to true for auto-created IFCs to turn on symmetric VRF Lite configuration.

This check box can be selected or deselected when the **VRF Lite Deployment** field is not set to Manual. In the case, a user explicitly unchecks the auto-deploy field for any auto-created IFCs, then the user input is always given the priority. This flag only affects the new auto-created IFC and it does not affect the existing IFCs.

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 upda one range of values at a time. If you want to update more than one range of values, do it in separate instance 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**.

Service Network VLAN Range - Specifies a VLAN range in the Service Network VLAN Range field. This is a per switch overlay service network VLAN range. The minimum allowed value is 2 and the maximum allowed value is 3967.

Route Map Sequence Number Range - Specifies the route map sequence number range. The minimum allowed value is 1 and the maximum allowed value is 65534.

9. Click the **Manageability** tab.

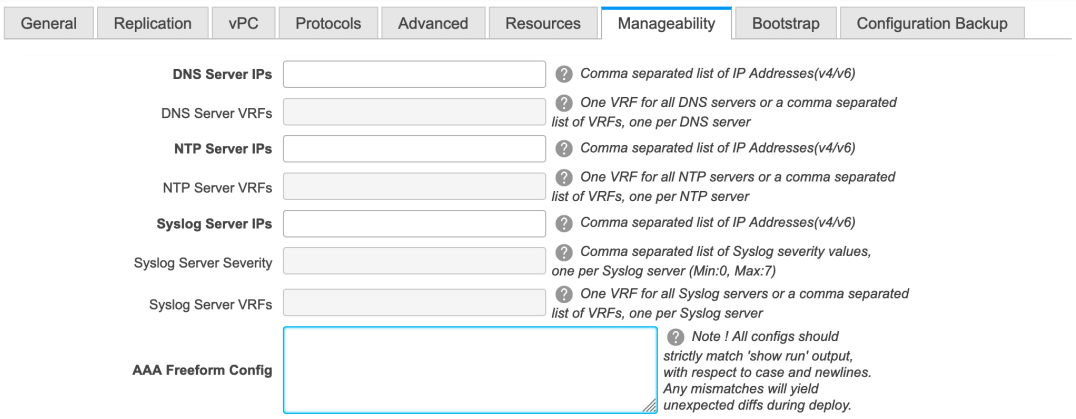

The fields in this tab are:

DNS Server IPs - Specifies the comma separated list of IP addresses (v4/v6) of the DNS servers.

DNS Server VRFs - Specifies one VRF for all DNS servers or a comma separated list of VRFs, one per DNS server.

NTP Server IPs - Specifies comma separated list of IP addresses (v4/v6) of the NTP server.

NTP Server VRFs - Specifies one VRF for all NTP servers or a comma separated list of VRFs, one per NTP server.

Syslog Server IPs – Specifies the comma separated list of IP addresses (v4/v6) IP address of the syslog servers, if used.

Syslog Server Severity – Specifies the comma separated list of syslog severity values, one per syslog server. The minimum value is 0 and the maximum value is 7. To specify a higher severity, enter a higher number.

Syslog Server VRFs – Specifies one VRF for all syslog servers or a comma separated list of VRFs, one per syslog server.

AAA Freeform Config – Specifies the AAA freeform configurations.

If AAA configurations are specified in the fabric settings, **switch_freeform** PTI with source as **UNDERLAY_AAA** and description as **AAA Configurations** will be created.

10. Click the **Bootstrap** tab.

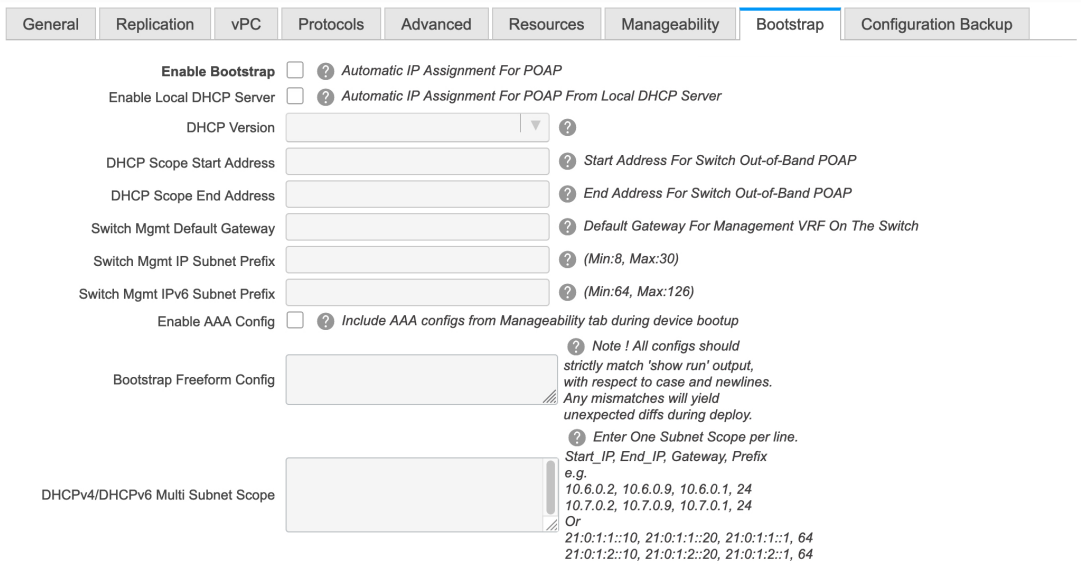

Enable Bootstrap - Select this check box to enable the bootstrap feature. Bootstrap allows easy day-0 import and bring-up of new devices into an existing fabric. Bootstrap leverages the NX-OS POAP functionality.

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 Mgmt Default Gateway** and **Switch Mgmt IP Subnet Prefix** fields.
- Local DHCP Server: Enable the **Local DHCP Server** check box 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 Version – Select DHCPv4 or DHCPv6 from this drop-down list. When you select DHCPv4, the **Switch Mgmt IPv6 Subnet Prefix** field is disabled. If you select DHCPv6, the **Switch Mgmt IP Subnet Prefix** is disabled.

Note Cisco Nexus 9000 and 3000 Series Switches support IPv6 POAP only when switches are either Layer-2 adjacent (eth1 or out-of-band subnet must be a /64) or they are L3 adjacent residing in some IPv6 /64 subn Subnet prefixes other than /64 are not supported.

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

Switch Mgmt IP 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.

Switch Mgmt IPv6 Subnet Prefix - Specifies the IPv6 prefix for the Mgmt0 interface on the switch. The prefix should be between 112 and 126. This field is editable if you enable IPv6 for DHCP.

Enable AAA Config – Select this check box to include AAA configurations from the Manageability tab as part of the device startup config post bootstrap.

Bootstrap Freeform Config - (Optional) Enter additional commands as needed. For example, if you require some additional configurations to be pushed to the device and be available post device bootstrap, they can be captured in this field, to save the desired intent. After the devices boot up, they will contain the configuration 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 347.

DHCPv4/DHCPv6 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

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

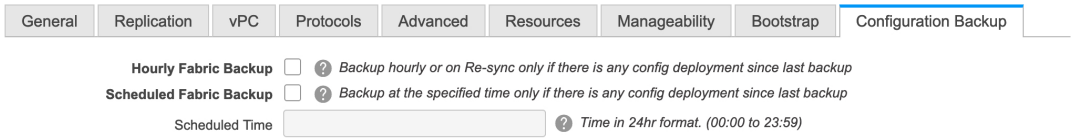

Hourly Fabric Backup: Select the check box to enable an hourly backup of fabric configurations and the intent.

The hourly backups are triggered during the first 10 minutes of the hour.

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.

The backup process is initiated after you click **Save**.

The scheduled backups are triggered exactly at the time you specify with a delay of up to two minutes. The scheduled backups are triggered regardless of the configuration deployment status.

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 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.

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.

(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.
- **Backup Now**: You can initiate a fabric backup manually by clicking **Backup Now**. Enter a name for the tag and click **OK**. Regardless of the settings you choose under the **Configuration Backup** tab in the **Fabric Settings** dialog box, you can initiate a backup using this option.

- **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 and interfaces. For more information, see Pre-provisioning a Device , on page 86 and Pre-provisioning an Ethernet Interface, on page 90.

Note When DCNM discovers a switch with the hostname containing the period character (.), it is treated as a domain-name and truncated. Only the text prior to the period character (.) is considered as a hostname. For example:

- If hostname is **leaf.it.vxlan.bgp.org1-XYZ**, DCNM shows only **leaf**
- If hostname is **leaf-itvxlan.bgp.org1-XYZ**, DCNM shows only **leafit-vxlan**

Discovering Existing Switches

1. After clicking on **Add Switches**, use the **Discover Existing Switches** tab to add one or more existing switches into the fabric. In this case, a switch with known credentials and a pre-provisioned IP address, is added to the fabric. The IP address (Seed IP), administrator username, and password (**Username** and **Password** fields) of the switch are provided as the input by a user. The **Preserve Config** knob is set to **yes** by default. This is the option that a user would select for a brownfield import of a device into the fabric. For a greenfield import where the device configuration will be cleaned up as part of the import process, the user should set the **Preserve Config** knob to **no**.

Note Easy_Fabric_eBGP does not support brownfield import of a device into the fabric.

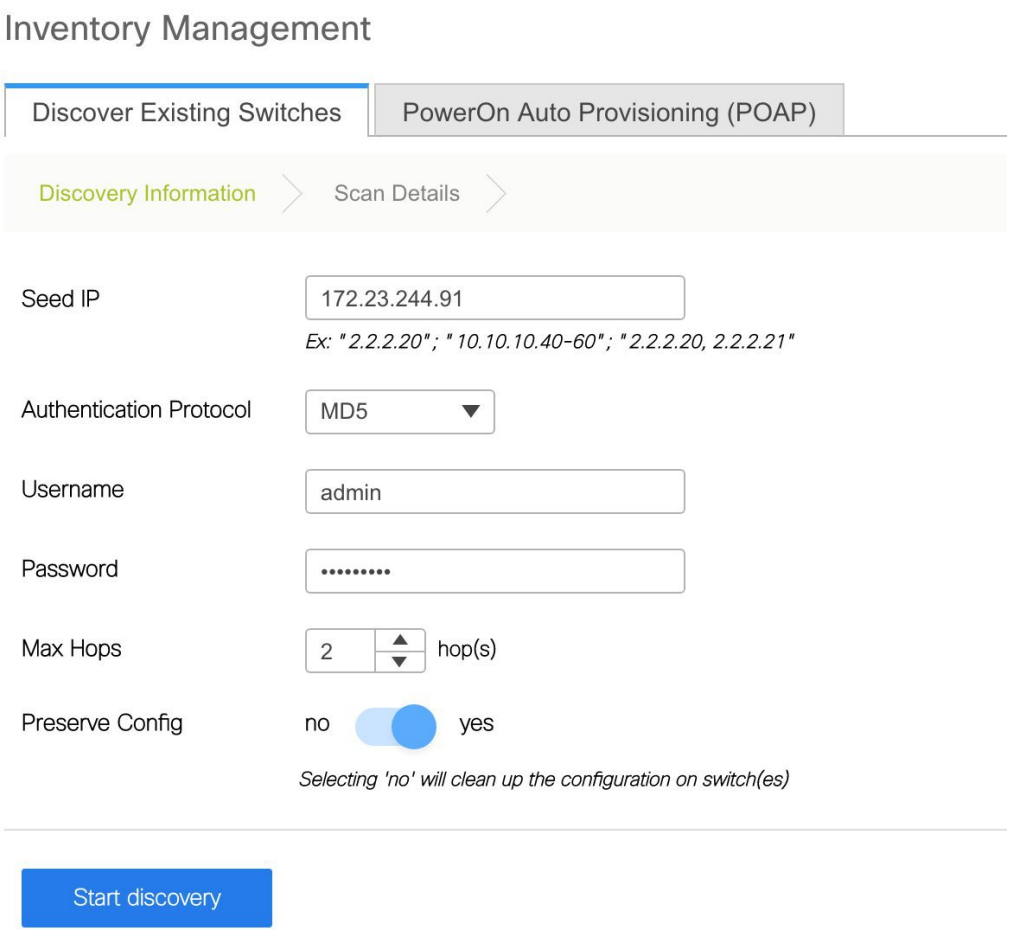

2. Click **Start discovery**. The **Scan Details** window comes up shortly. Since the **Max Hops** field was populated with 2 (by default), the switch with the specified IP address (leaf-91) and switches two hops from that switch, are populated in the **Scan Details**result.

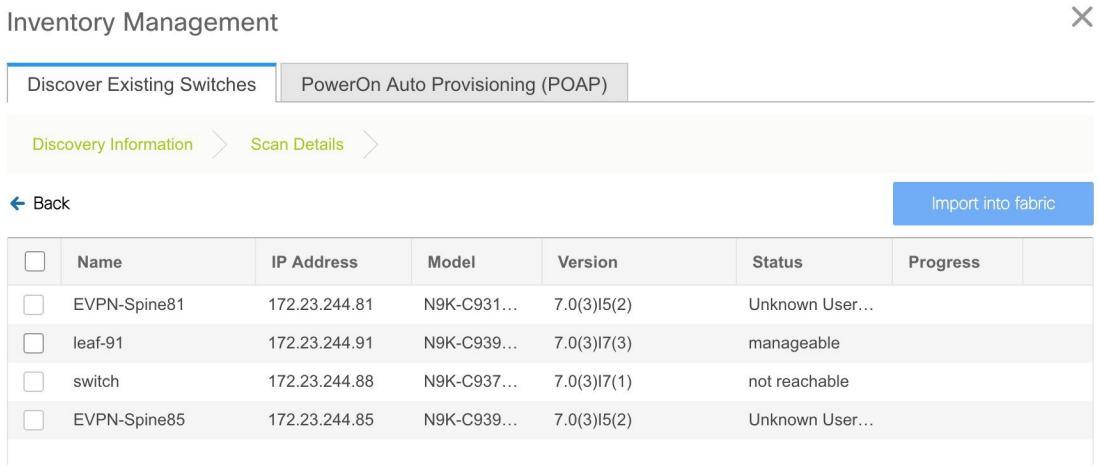

3. If the DCNM was able to perform a successful shallow discovery to a switch, the status will show up as **Manageable**. Select the check box next to the appropriate switch(es) and click **Import into fabric**.

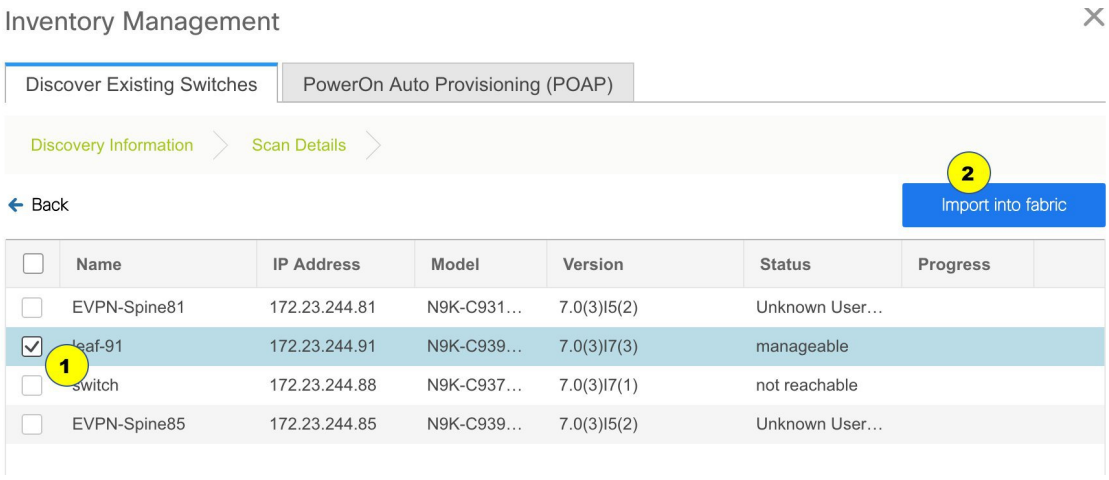

Though this example describes the discovery of one switch, multiple switches can be discovered at once.

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 c error message comes up.

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

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.

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 devices, assign an appropriate role to each device. For this purpose, tight click the device, and use the **Set role** option to set the appropriate role. Alternatively, the tabular view may be employed to assign the same role to multiple devices at one go.

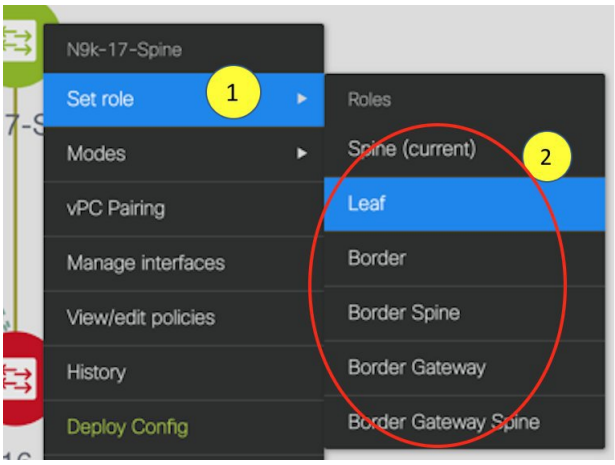

   If you choose the Hierarchical layout for display (in the Actions panel), the topology automatically gets aligned as per role assignment, with the leaf devices at the bottom, the spine devices connected on top of them, and the border devices at the top.

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

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

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

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

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

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

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

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

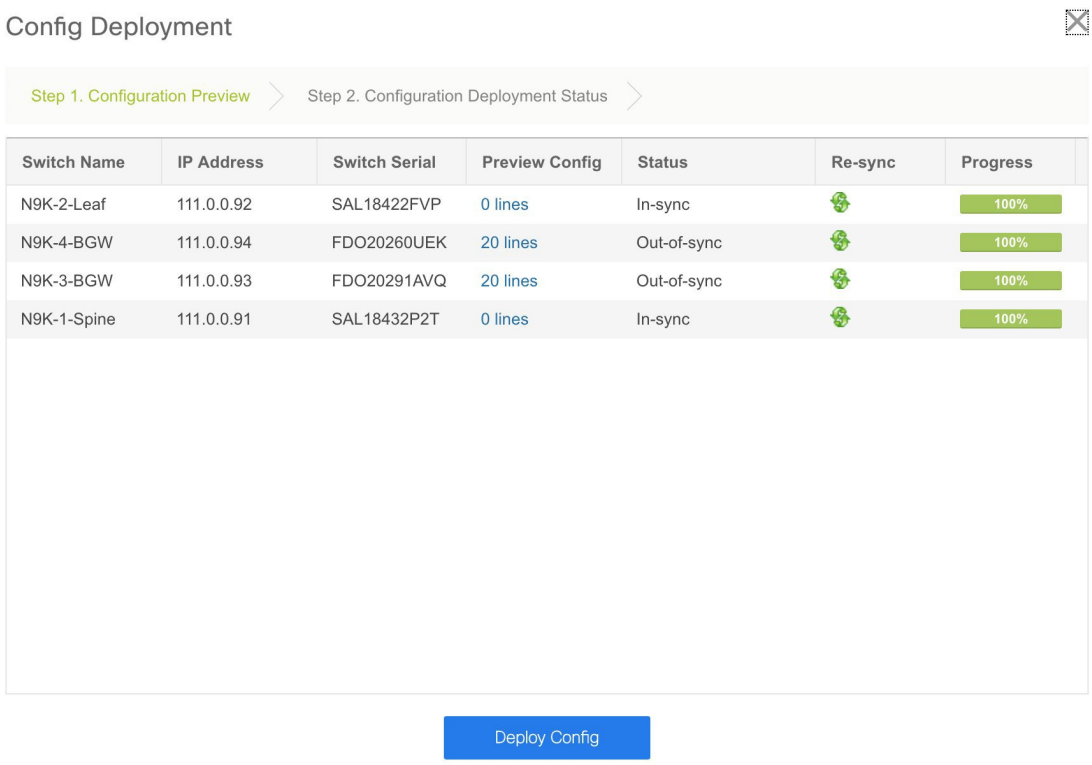

   If the status is out-of-sync, it suggests that there is inconsistency between the DCNM and configuration on the device.

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

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

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

The **PendingConfig** tab displays the pending configurations for successful deployment.

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

In DCNM 11, 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 are 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. The pending state indicates that there is a pending deployment or pending recomputation. You can click on the switch and review the pending deployments using **Preview** or **Deploy Config** options, or click **Save & Deploy** to recompute the state of the switch.

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

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.

From Cisco NX-OS Release 11.4(1), if you uncheck the **FEX** check box in the **Topology** window, FEX devices are hidden in the **Fabric Builder** topology window as well. To view FEX in **Fabric Builder**, you need to check this check box. This option is applicable for all fabrics and it is saved per session or until you log out of DCNM. If you log out and log in to DCNM, the FEX option is reset to default, that is, enabled by default. For more information, see Show Panel, on page 24.

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

Discovering New Switches

1. When a new Cisco NX-OS device is powered on, typically that device has no startup configuration or any configuration state for that matter. Consequently, it powers on with NX-OS and post initialization, goes into a POAP loop. The device starts sending out DHCP requests on all the interfaces that are up including the mgmt0 interface.
2. As long as there is IP reachability between the device and the DCNM, the DHCP request from the device, will be forwarded to the DCNM. For easy day-0 device bring-up, the bootstrap options should be enabled in the **Fabric Settings** as mentioned earlier.
3. With bootstrap enabled for the fabric, the DHCP request coming from the device will be serviced by the DCNM. The temporary IP address allocated to the device by the DCNM will be employed to learn basic information about the switch including the device model, device NX-OS version, etc.
4. In the DCNM GUI, go to a fabric (Click **Control > Fabric Builder** and click a fabric). The fabric topology is displayed.

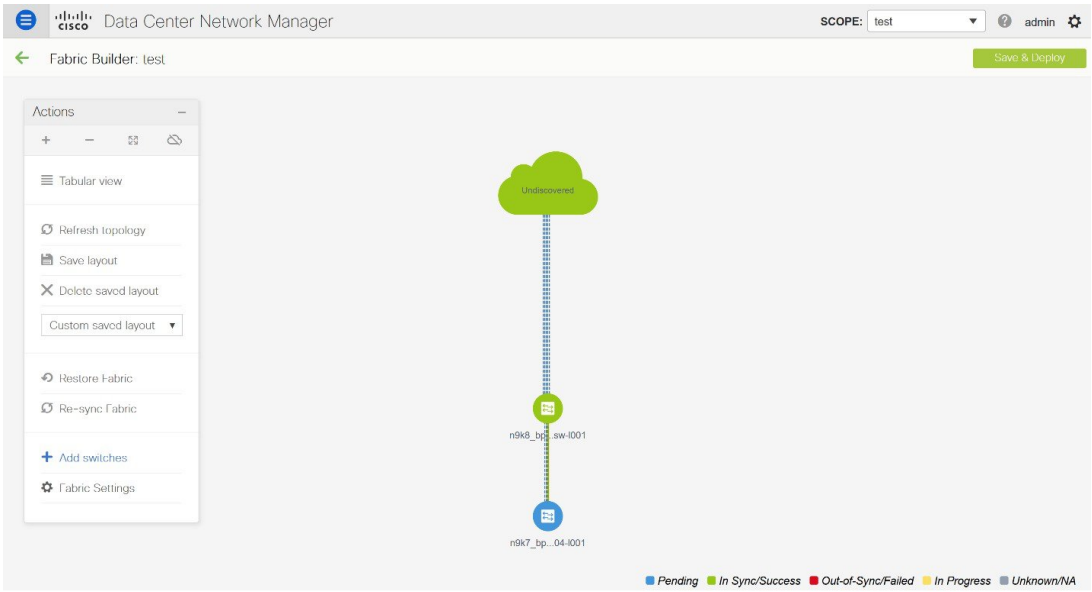

Go to the fabric topology window and click the **Add switches** option from the **Actions** panel. The Inventory Management window comes up.

5. Click the **POAP** tab.

   As mentioned earlier, DCNM retrieves the serial number, model number, and version from the device 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

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

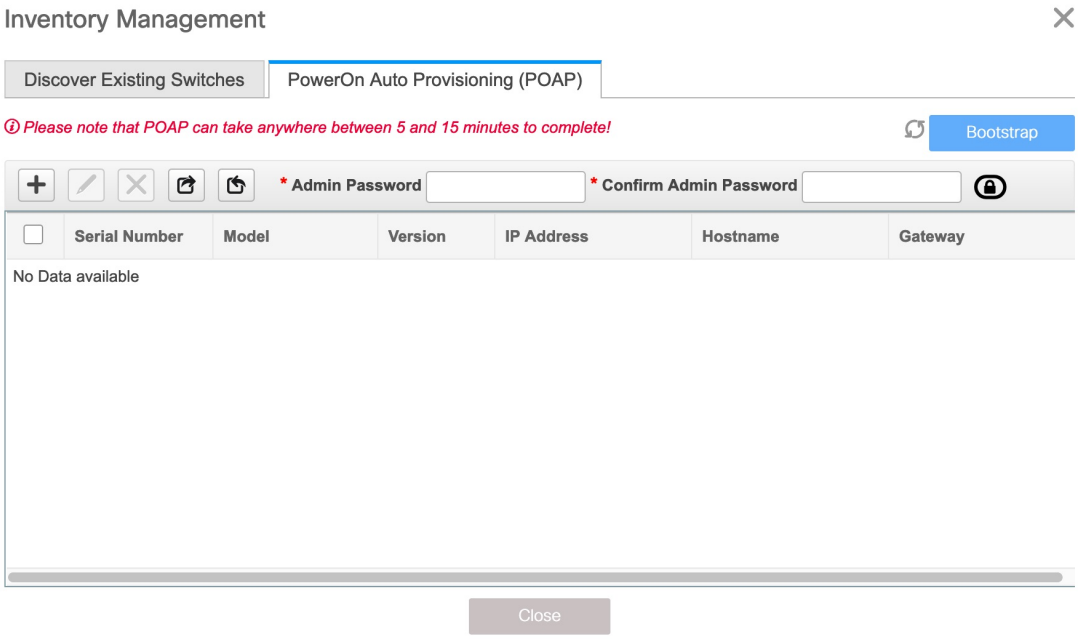

Select the checkbox next to the switch and enter the switch credentials: IP address and host name.

Based on the IP address of your device, you can either add the IPv4 or IPv6 address in the **IP Address** field.

Beginning with Release 11.2(1), you can provision devices in advance. To pre-provision devices, refer to Pre-provisioning a Device , on page 86.

6. 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 authenticatio that is, RADIUS or TACACS credentials for discovery only.

7. (Optional) Use discovery credentials for discovering switches.

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

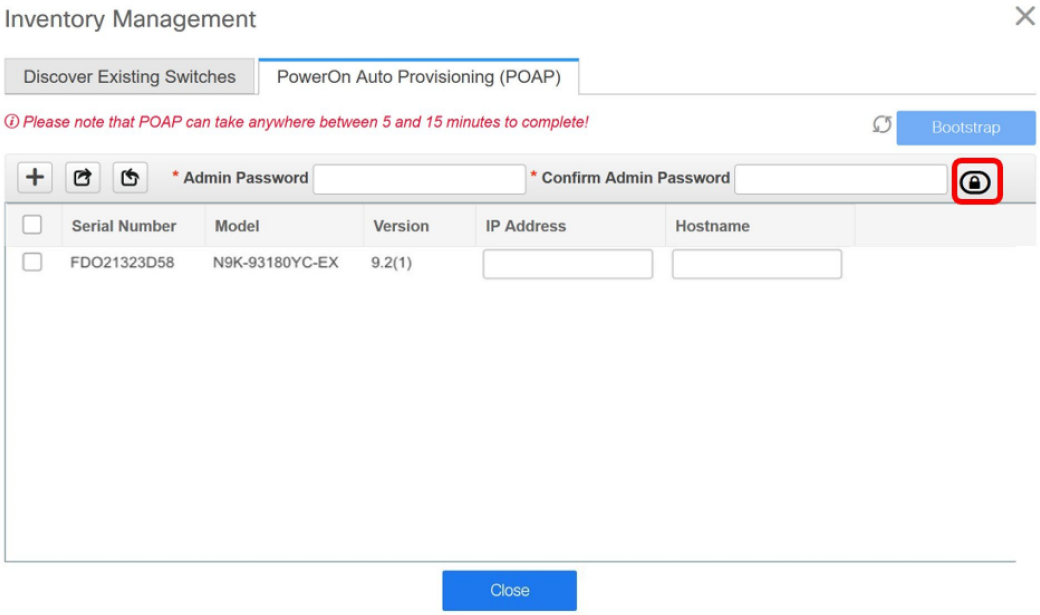

b. In the **Discovery Credentials** window, enter the discovery credentials such as discovery username and password.

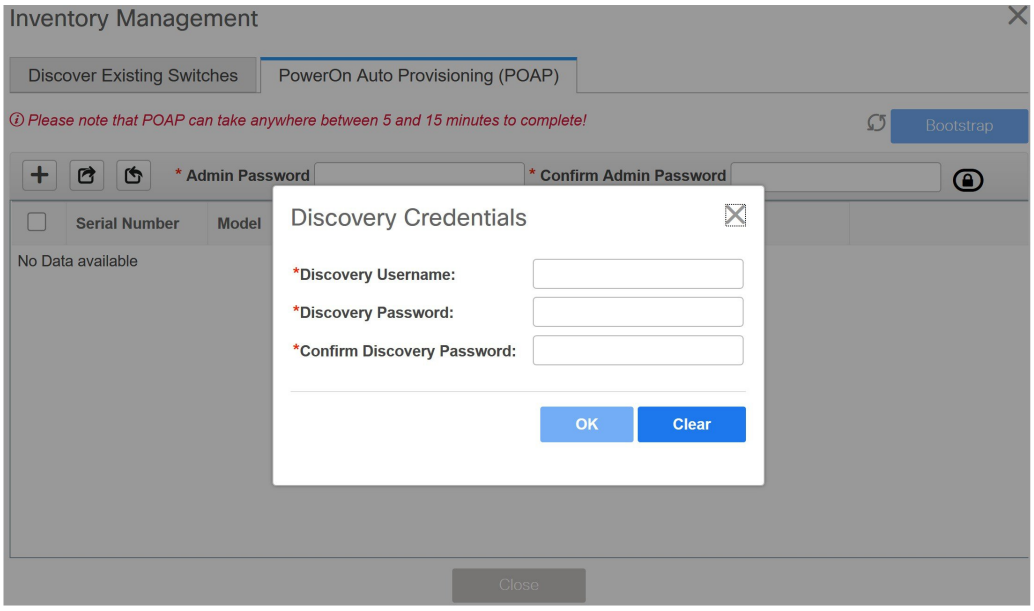

Click **OK** to save the discovery credentials.

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

8. 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.

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

10. After the added switch completes POAP, the fabric builder topology page is refreshed with the added switch thereby depicting its discovered physical connections. Set the appropriate role for the switch followed by a Save & Deploy operation at the fabric level. The Fabric Settings, switch role, the topology etc. are evaluated by the Fabric Builder and the appropriate intended configuration for the switch is generated as part of the Save operation. The pending configuration will provide a list of the configurations that need to be deployed to the new switch in order to bring it IN-SYNC with the intent.

Note For any changes on the fabric that results in the Out-of-Sync, then you must deploy the changes. The proce 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.

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

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

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

14. 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.

15. 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. 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 pairing/un-pairing 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:

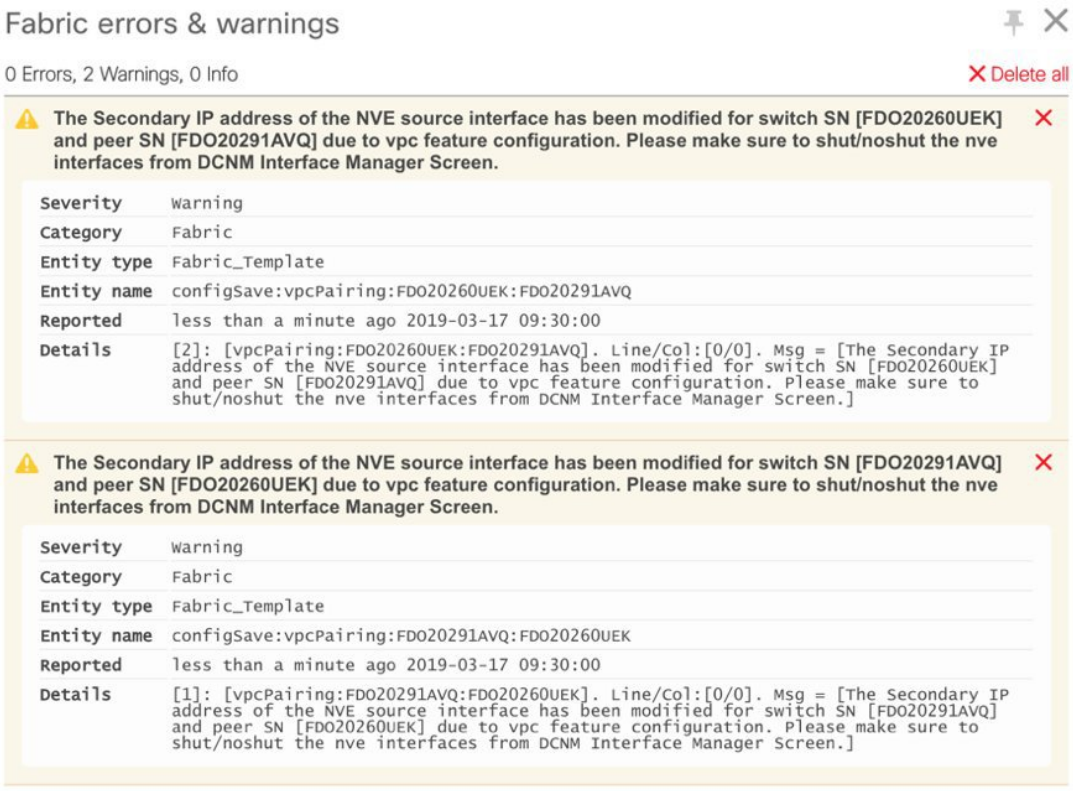

To resolve, go to the Control > Interfaces screen and deploy the Shutdown operation on the nve interface followed by a No Shutdown configuration. This is depicted in the figure below where the up arrow corresponds to a No Shutdown operation while a down arrow corresponds to a Shutdown operation.

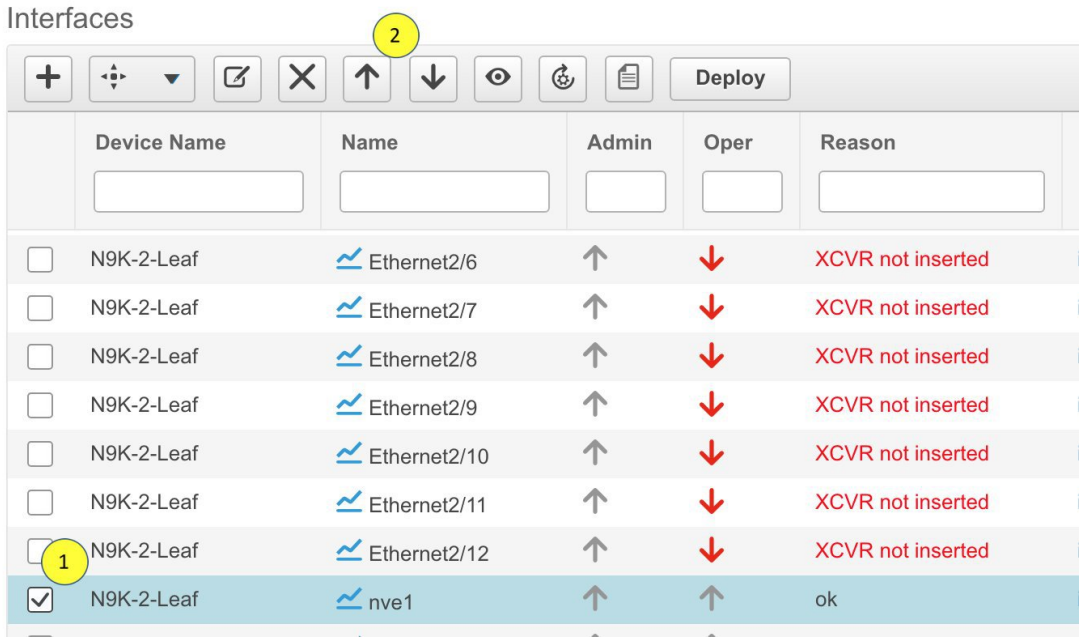

You can right-click the switch to view various options:

- **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 205.

- **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 and policy change history.

  The **Policy Change History** tab lists the history of policies along with the users who made the changes like add, update, or delete.

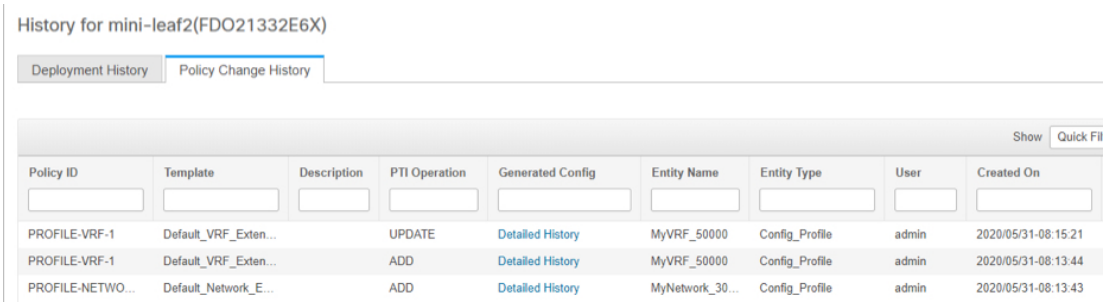

Under the **Policy Change History** tab, for a policy, click **Detailed History** under the **Generated Config** column to view the generated config before and after.

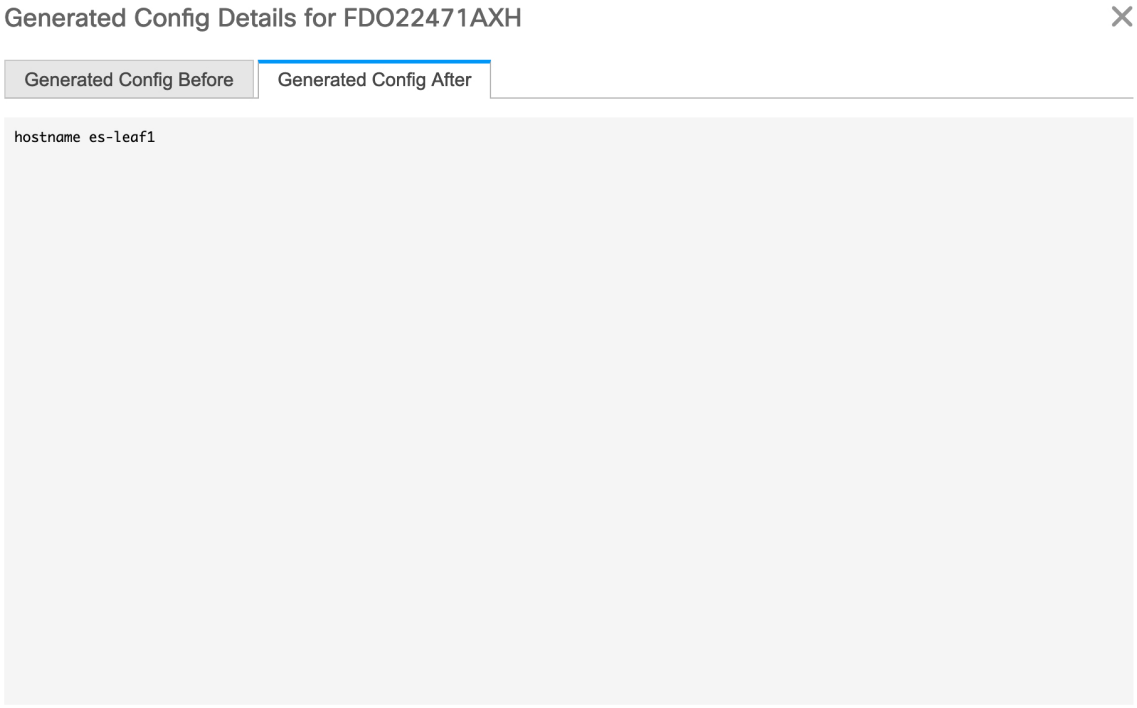

The following table provides the summary of generated config before and after for Policy Template Instances (PTIs).

PTI Operations	Generated Config Before	Generated Config After
Add	Empty	Contains the config
Update	Contains config before changes	Contains config after changes
Mark-Delete	Contains the config to be removed.	Contains the config to be removed with colour change.
Delete	Contains the config	Empty

Note When a policy or profile template is applied, an instance is created for each application of the template, which is known as Policy Template Instance or PTI.

- **Preview Config** - View the pending configuration and the side-by-side comparison of the running and expected configuration.
- **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 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 networks and deploy them on the switches. [Refer Creating and Deploying Networks and VRFs].

Pre-provisioning Support in DCNM 11

Cisco DCNM supports provisioning of device configuration in advance. This is specifically applicable for scenarios where devices have been procured, but not yet delivered or received by the Customers. The purchase order typically has information about the device serial number, device model and so on, which in turn can be used to prepare the device configuration in DCNM prior to the device connectivity to the Network. Pre-provisioning is supported for Cisco NX-OS devices in both Easy Fabric and External/Classic_LAN fabrics.

Pre-provisioning a Device

From Cisco DCNM Release 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
  - Ethernet ports
  - Port-channel
  - vPC
  - ST FEX
  - AA FEX
  - Loopback
  - Overlay network configurations
- The pre-provisioned devices do not support the following configurations in DCNM:
  - Inter-Fabric links
  - Sub-interface
  - Interface breakout configuration

- 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 must 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 Click **Control > Fabric Builder**.

The **Fabric Builder** screen is displayed.

Step 2 Click within the fabric box.

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

The **Inventory Management** screen is displayed.

Step 4 Click the **POAP** tab.

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

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

The Add a new device screen comes up.

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

c. Click **Save**.

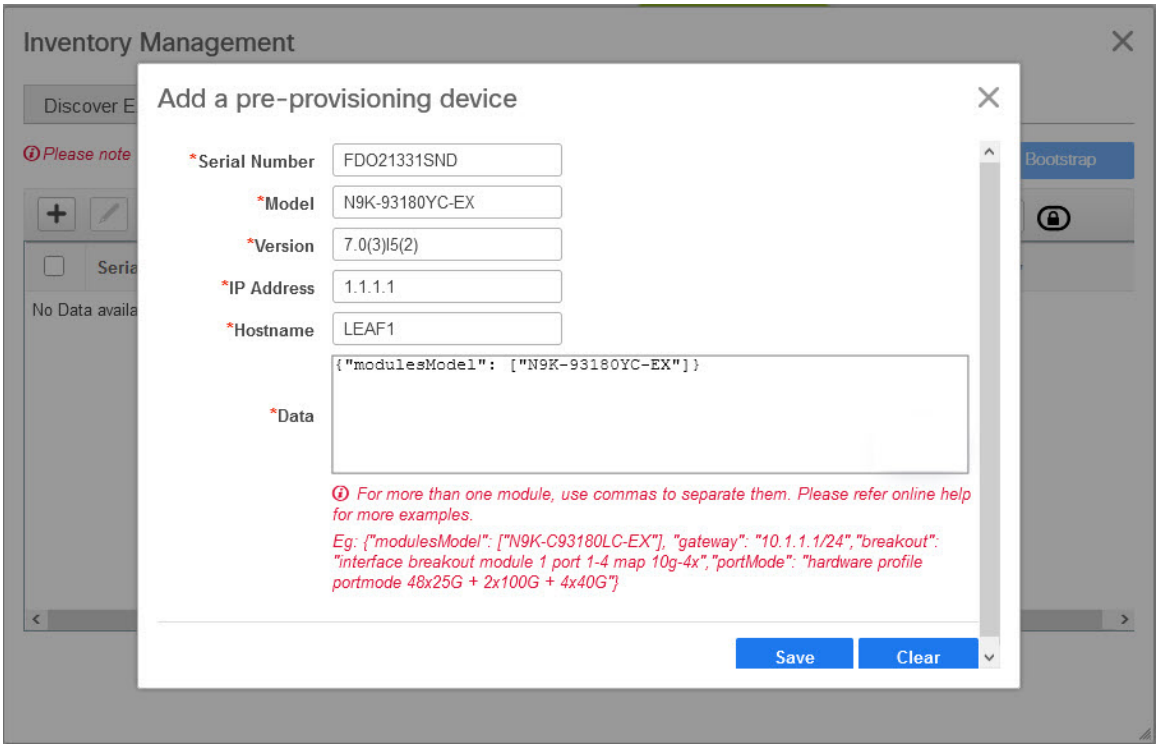

IP Address: Specify the IPv4 or IPv6 address of the new device.

Serial Number: The serial number for the new device. Serial number is found in the Cisco Build of Material Purchase and you can refer to these values while using the pre-provisioning feature.

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

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

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

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

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

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

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

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

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

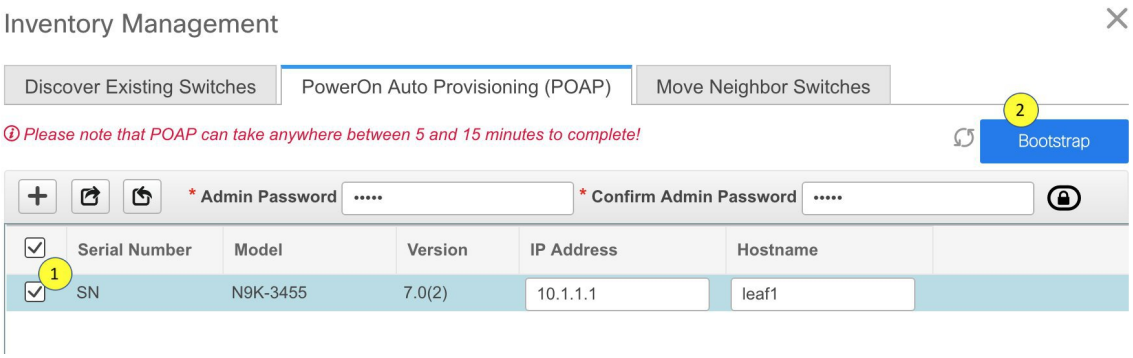

The leaf1 device appears in the fabric topology.

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

Note When a switch is in **Unreachable** discovery status, the last available information of the switch is retained in other columns.

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 253.

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 one or more switches are online to provision the host ports. This action must be performed before overlays are provisioned for the host port attachment.

Pre-provisioning an Ethernet Interface

From DCNM Release 11.4(1), you can pre-provision Ethernet interfaces in the **Interface** window. This pre-provisioning feature is supported in the Easy, External, and eBGP fabrics. You can add Ethernet interfaces to only pre-provisioned devices before they are discovered in DCNM.

Note Before attaching a network/VRF, you must pre-provision the Ethernet interface before adding it to Port-channels, vPCs, ST FEX, AA FEX, loopback, subinterface, tunnel, ethernet, and SVI configurations.

Before you begin

Make sure that you have a preprovisioned device in your fabric. For information, see Pre-provisioning a Device , on page 86.

Procedure

Step 1 Navigate to the fabric containing the pre-provisioned device from the **Fabric Builder** window.

Step 2 Right click the pre-provisioned device and select **Manage Interfaces**.

You can also navigate to the Interfaces window by selecting **Control > Fabrics > Interfaces**. From the Scope drop-down list, select the fabric containing the pre-provisioned device.

Step 3 Click **Add**.

Step 4 Enter all the required details in the **Add Interface** window.

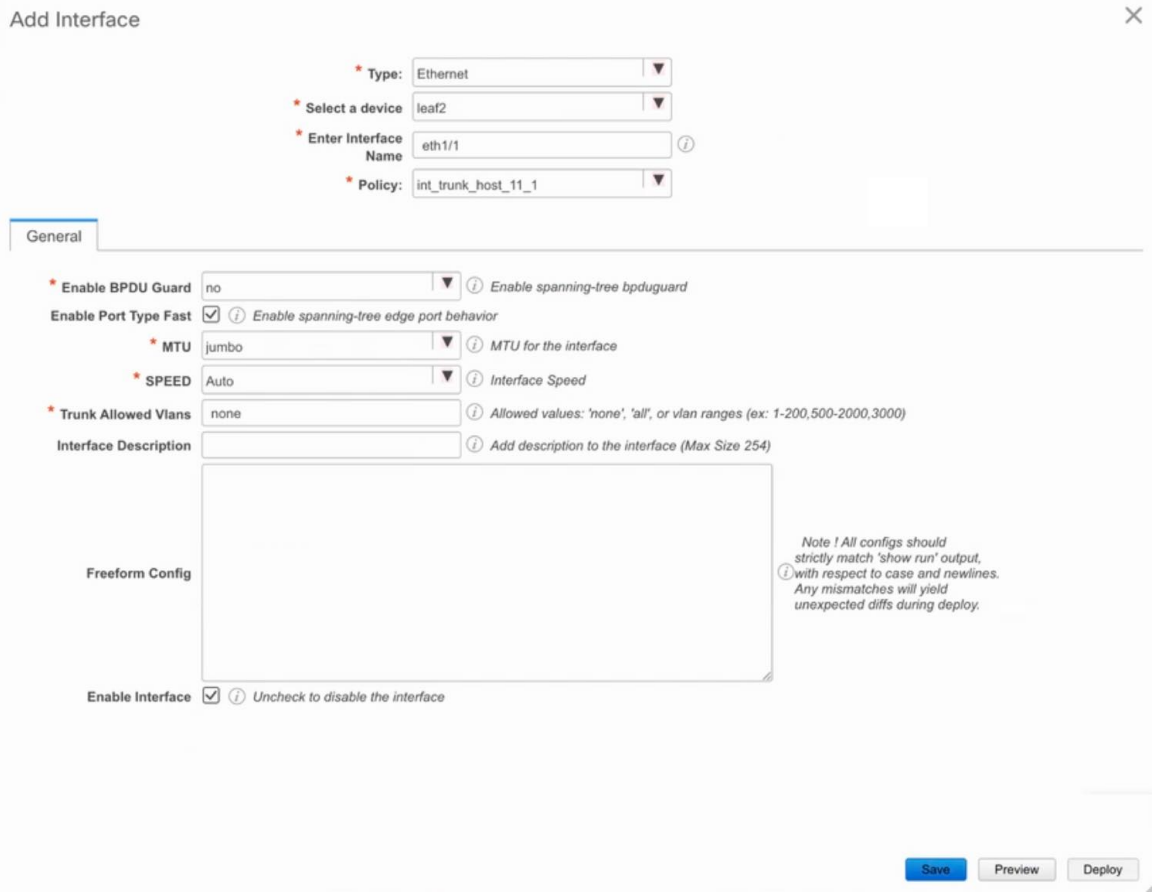

Type: Select **Ethernet** from this drop-down list.

Select a device: Select the pre-provisioned device.

Note You cannot add an Ethernet interface to an already managed device in DCNM.

Enter Interface Name: Enter a valid interface name based on the module type. For example, Ethernet1/1, eth1/1, or e1/1. The interface with same name should be available on the device after it is added.

Policy: Select a policy that should be applied on the interface.

For more information, see Adding Interfaces, on page 263.

Step 5 Click **Save**.

Step 6 Click **Preview** to check the expected configuration that will be deployed to the switch after it is added.

Note The **Deploy** button is disabled for Ethernet interfaces since the devices are pre-provisioned.

Pre-provisioning a vPC Pair

Before you begin

Ensure that you have enabled **Bootstrap** in the Fabric Settings.

Procedure

Step 1 Import both the devices into the fabric.

For instructions, see **Pre-provisioning a Device**.

The following example in the image shows two Cisco Nexus 9000 Series devices that are pre-provisioned and added to an existing Fabric. Choose **Add Switches** in the Action panel. On the Inventory Management screen, click **PowerOn Auto Provisioning (POAP)**.

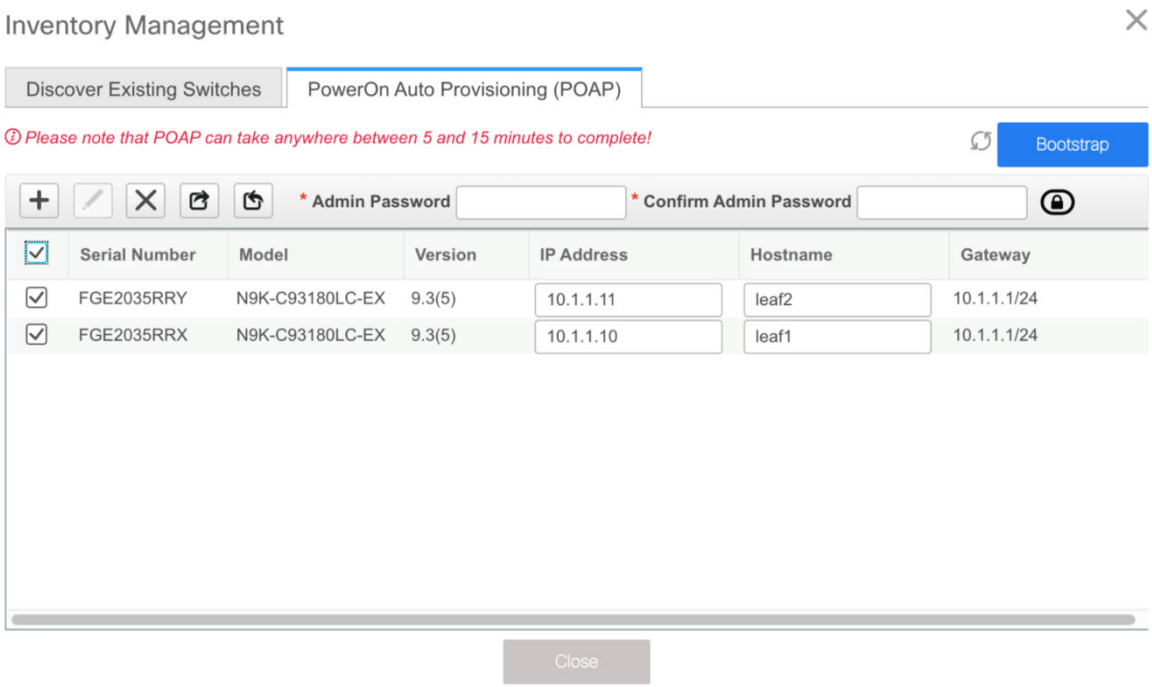

The devices will show up in the fabric as gray/undiscovered devices.

Step 2 Right click and select appropriate roles for these devices similar to other reachable devices.

Step 3 To create vPC pairing between the devices with physical peer-link or MCT, perform the following steps:

a) Provision the physical Ethernet interfaces that form the peer-link.

The vPC peer-link between leaf1-leaf2 comprises of interfaces Ethernet1/44-45 on each device. Choose **Control > Fabrics > Interfaces** to pre-provision ethernet interfaces.

For instructions, see **Preprovisioning an Ethernet Interface**.

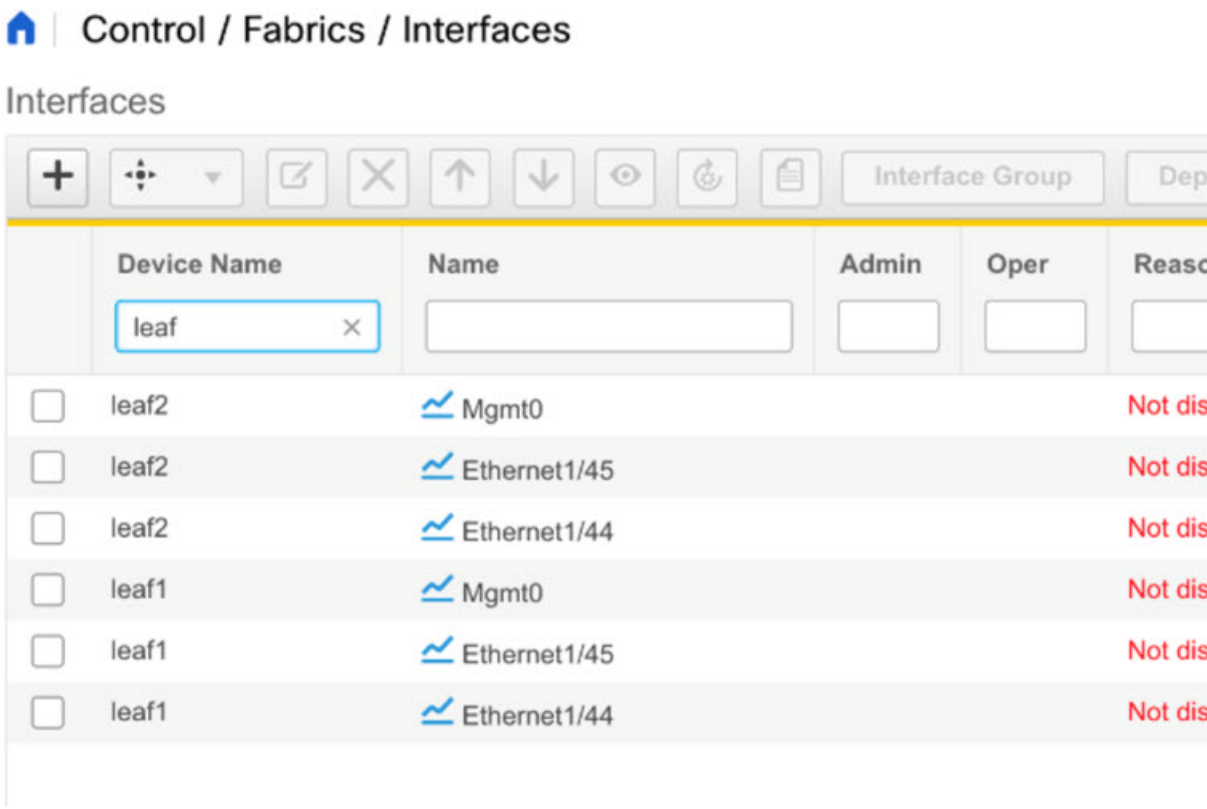

b) Create a pre-provisioned link between these interfaces.

In the Fabric Builder view, right click and select **Add** link or click **Add(+)** icon in the Links tab in the Fabric Builder Tabular view.

Create two links, one for leaf1-Ethernet1/44 to leaf2-Ethernet1/44 and another one for leaf1-Ethernet1/45 to leaf2-Ethernet1/45.

Ensure that you choose **int_pre_provision_intra_fabric_link** as link template. The Source Interface and Destination Interface field names, must match with the Ethernet interfaces pre-provisioned in the previous step.

An example of pre-provisioned link creation is as depicted in the following image.

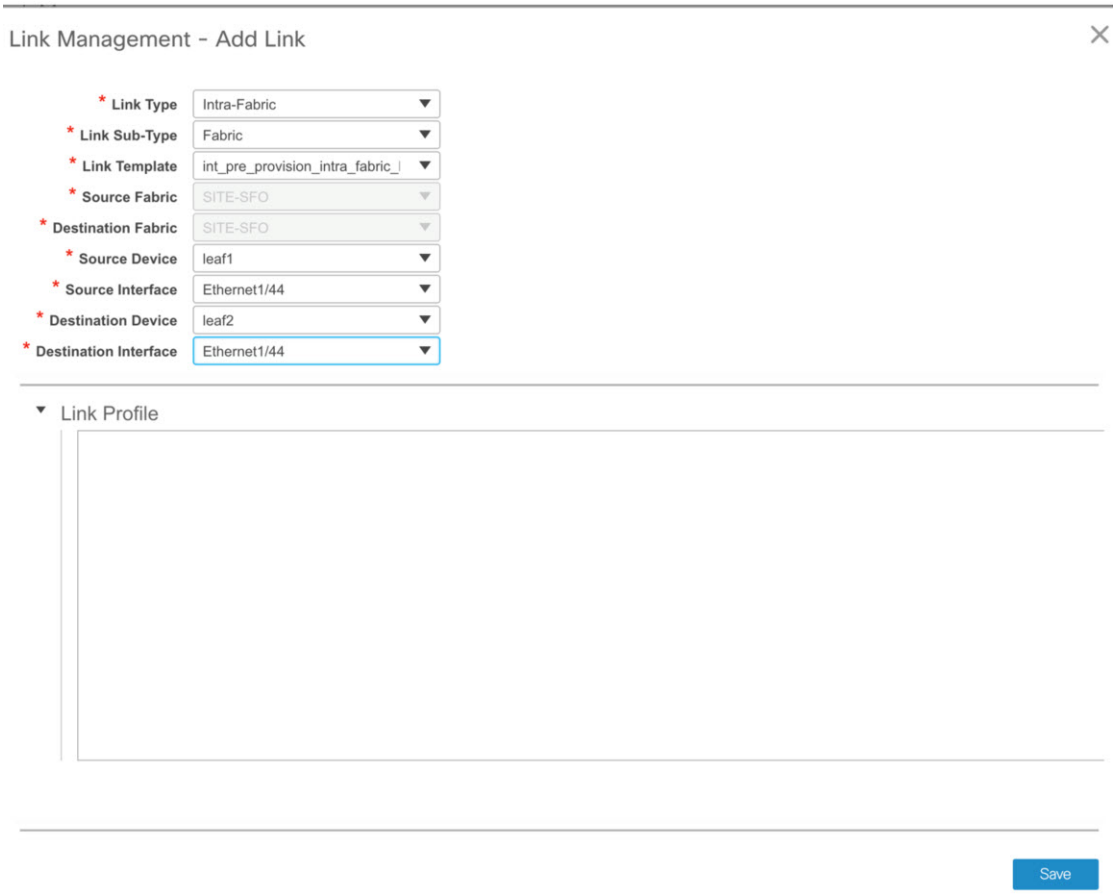

After the links are created, they are listed in the Links tab under Fabric builder as shown in the following image.

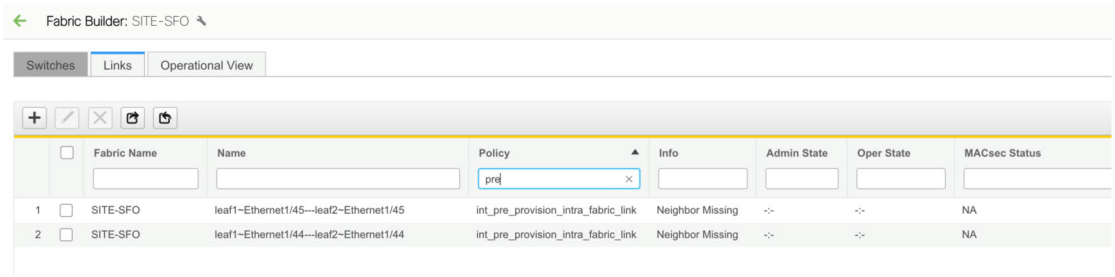

c) On Fabric topology, right click on a switch and choose vPC Pairing from the drop-down list.

Select the vPC pair and click vPC pairing for the pre-provisionsed devices.

d) Click **Save & Deploy** to generate the required intended vPC pairing configuration for the pre-provisioned devices.

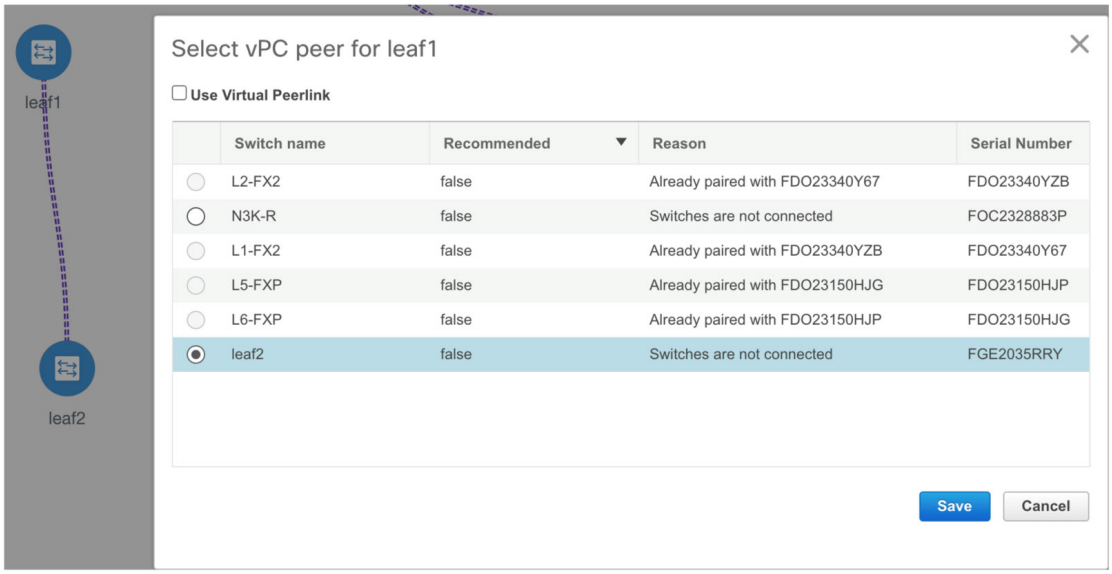

After completion, the devices will be correctly paired and the vPC pairing intent will be generated for the devices. The policies are generated as shown in the following image:

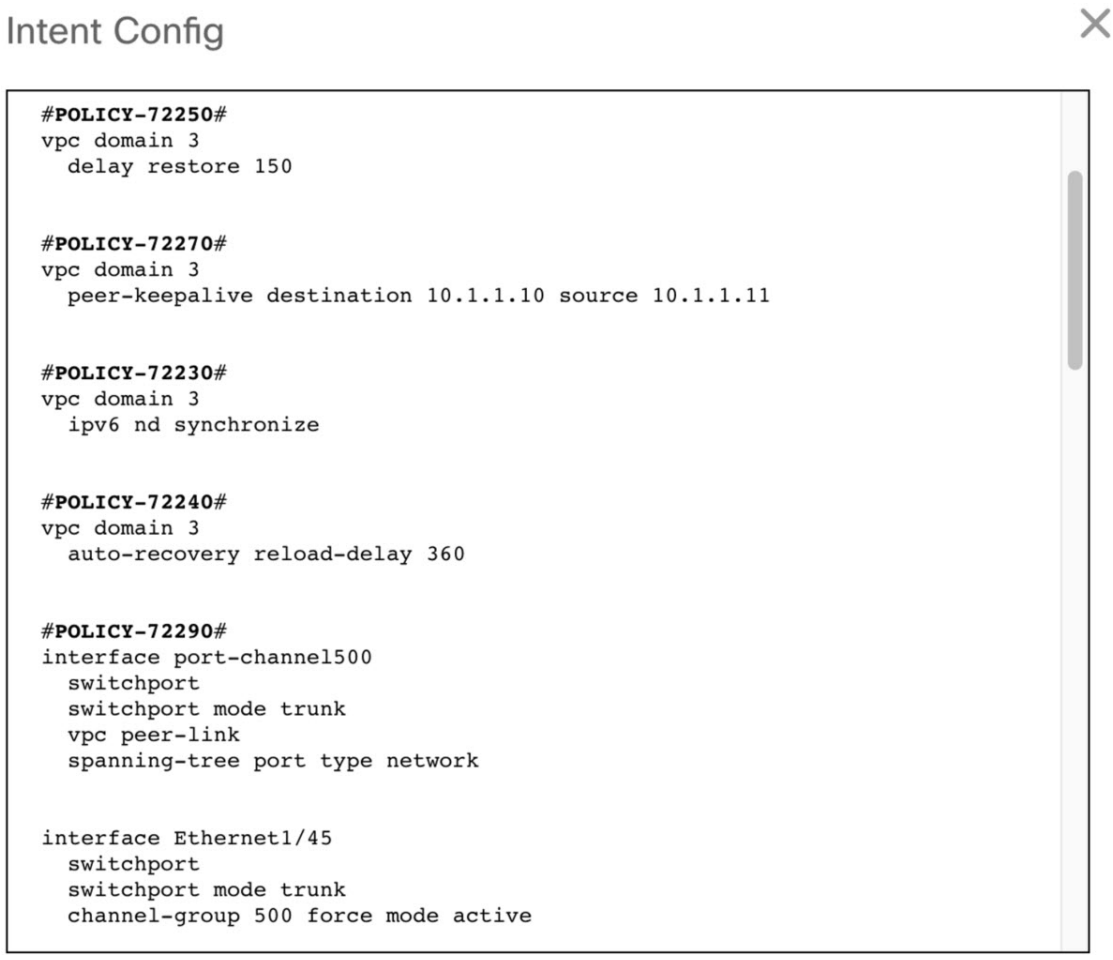

Note Because the devices are not yet operational, Config Compliance will not return any IN-SYNC or OUT-OF-SYNC status for these devices.

This is expected as CC requires the running configuration from the devices in order to compare that with the intent and calculate and report the compliance status.

Pre-provisioning a vPC Host Interface

Procedure

Step 1 Create physical ethernet interfaces on the pre-provisioned devices. Add a vPC host interface similar to a regular vPC pair or switches.

For instructions, see Pre-provisioning an Ethernet Interface, on page 90.

For example, leaf1-leaf2 represents the pre-provisioned vPC device pair, assuming that Ethernet interfaces 1/1 is already pre-provisioned on both devices leaf1 and leaf2.

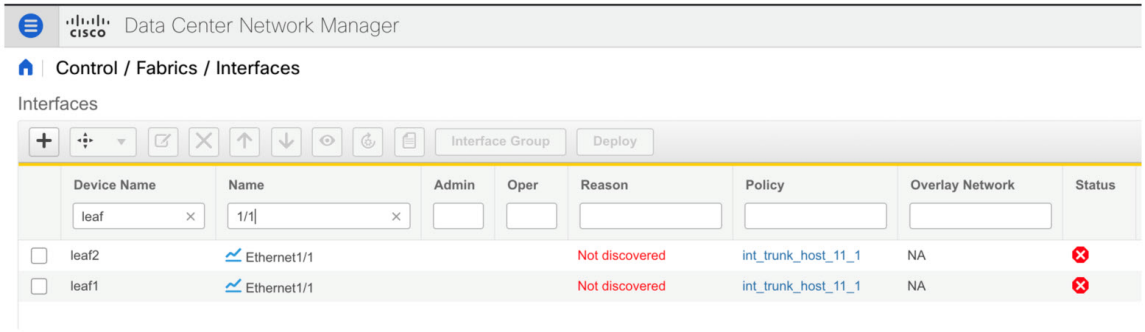

Step 2 Create a vPC host truck interface as shown in the following image.

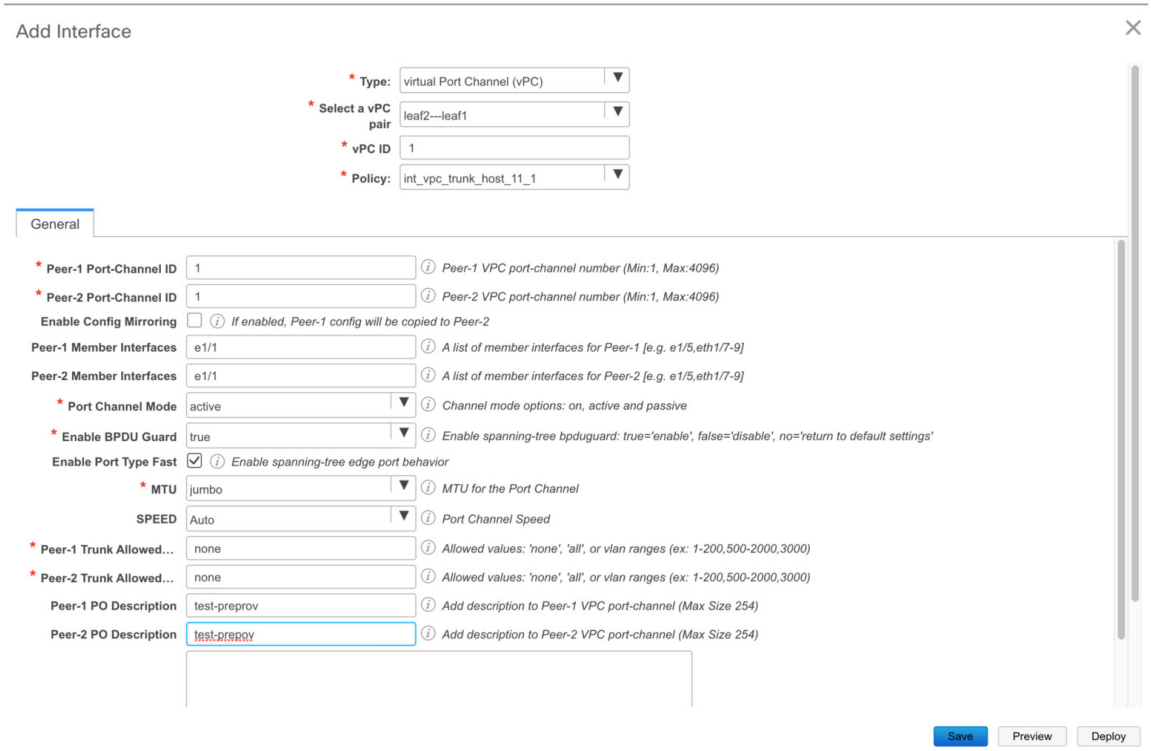

Preview and **Deploy** actions doesn't yield any result, because both require the device to be present. The vPC host interface is created and displays status as **Not discovered** as shown in the following image.

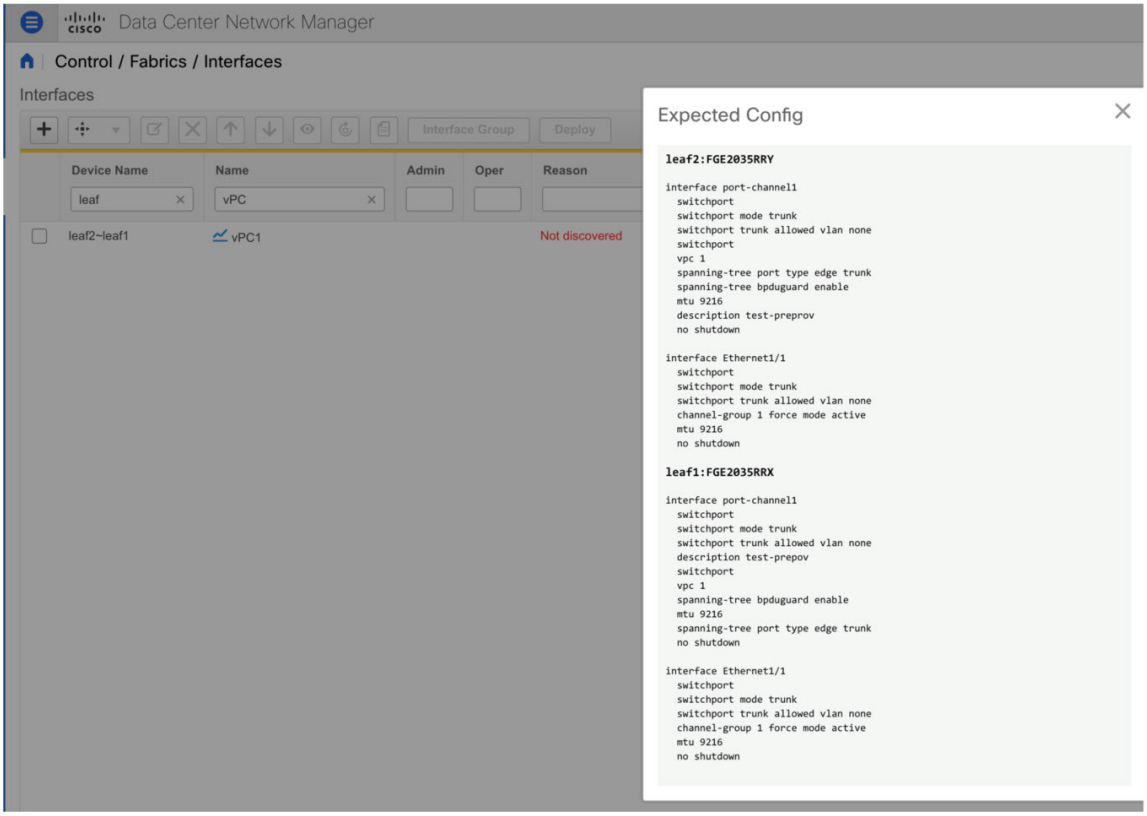

Attaching Overlays to Pre-provisioned Devices

Overlay VRFs and Networks can be attached to pre-provisioned devices similar to any other discovered device.

The following example shows where an overlay network is attached to the pre-provisioned vPC pair of leafs (leaf1-leaf2). It is also attached to the pre-provisioned vPC host interface port-channels created on leaf1-leaf2.

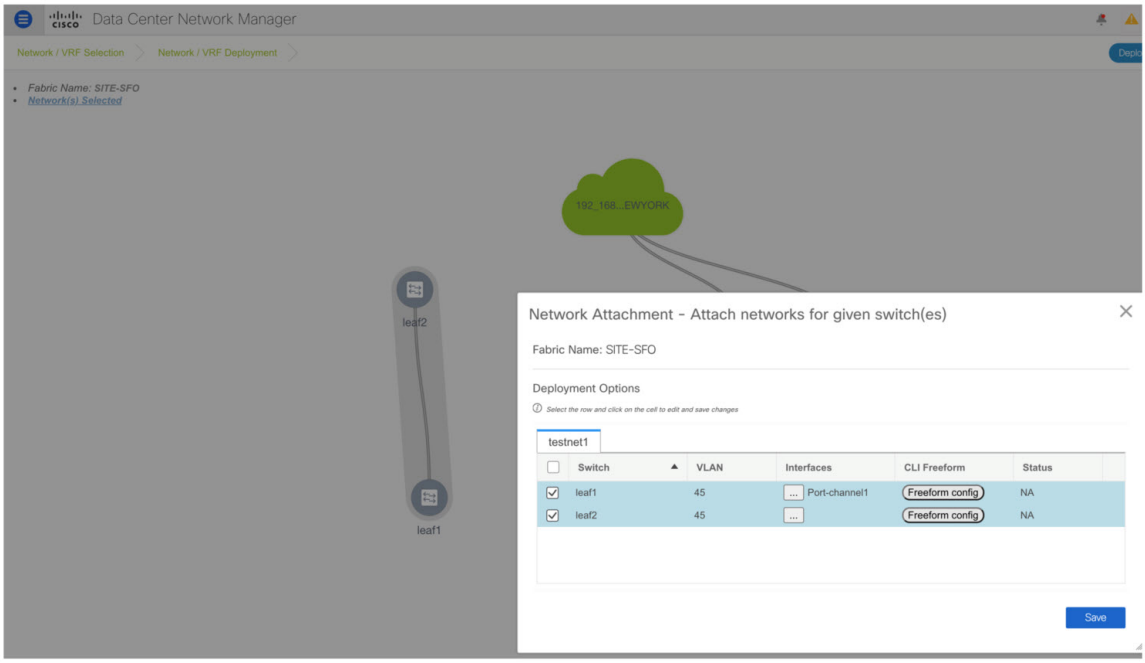

Preview and **Deploy** operations are disabled for the pre-provisioned devices, because the devices are not reachable. After the pre-provisioned device is reachable, all operations are enabled similar to other discovered devices.

On **Fabric Builder > View/Edit Policies**, you can view the entire intent generated for the pre-provisioned device, including the overlay network/VRF attachment information as shown in the following image.

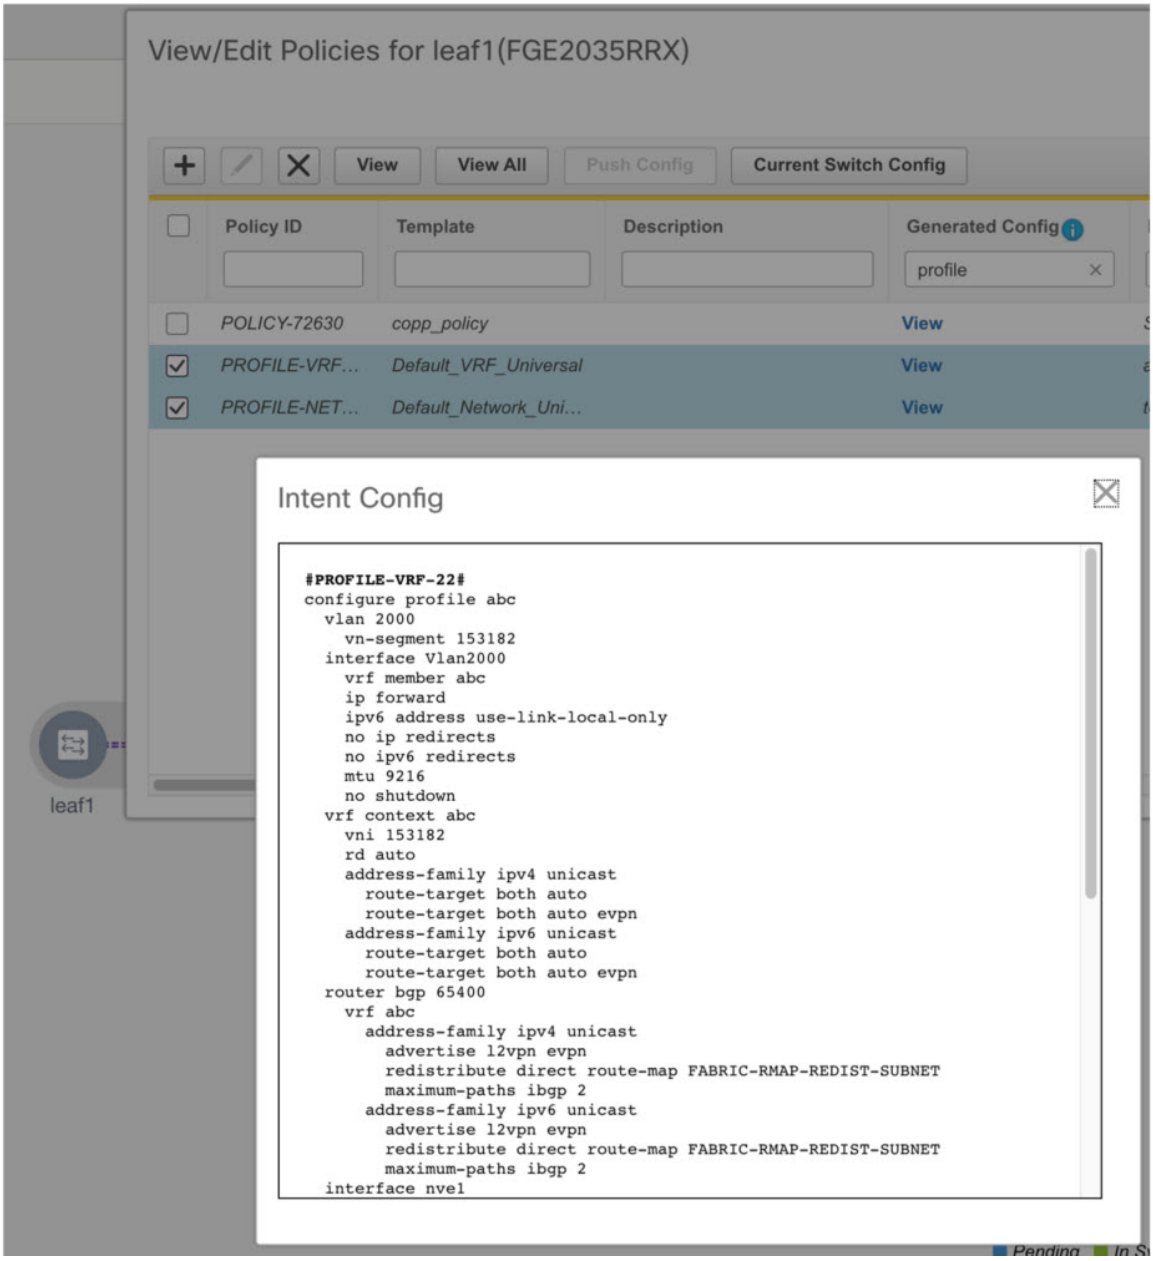

Precision Time Protocol for Easy Fabric

In the fabric settings for the **Easy_Fabric_11_1** template, select the **Enable Precision Time Protocol (PTP)** check box to enable PTP across a fabric. When you select this check box, PTP is enabled globally and on core-facing interfaces. Additionally, the **PTP Loopback Id** and **PTP Domain Id** fields are editable.

The PTP feature works only when all the devices in a fabric are cloud-scale devices. Warnings are displayed if there are non-cloud scale devices in the fabric, and PTP is not enabled. Examples of the cloud-scale devices are Cisco Nexus 93180YC-EX, Cisco Nexus 93180YC-FX, Cisco Nexus 93240YC-FX2, and Cisco Nexus 93360YC-FX2 switches.

For more information, see the *Configuring PTP* chapter in *Cisco Nexus 9000 Series NX-OS System Management Configuration Guide* and *Cisco Network Insights for Resources Application for Cisco DCNM User Guide*.

For LAN fabric deployments, specifically in a VXLAN EVPN based fabric deployments, you have to enable PTP globally, and also enable PTP on core-facing interfaces. The interfaces could be configured to the external PTP server like a VM or Linux-based machine. Therefore, the interface should be edited to have a connection with the grandmaster clock.

It is recommended that the grandmaster clock should be configured outside of Easy Fabric and it is IP reachable. The interfaces toward the grandmaster clock need to be enabled with PTP via the interface freeform config.

All core-facing interfaces are auto-enabled with the PTP configuration after you click **Save & Deploy**. This action ensures that all devices are PTP synced to the grandmaster clock. Additionally, for any interfaces that are not core-facing, such as interfaces on the border devices and leafs that are connected to hosts, firewalls, service-nodes, or other routers, the ttag related CLI must be added. The ttag is added for all traffic entering the VXLAN EVPN fabric and the ttag must be stripped when traffic is exiting this fabric.

Here is the sample PTP configuration:

```
feature ptp

ptp source 100.100.100.10 -> IP address of the loopback interface (loopback0) that is already
 created or user created loopback interface in the fabric settings

ptp domain 1 -> PTP domain ID specified in fabric settings

interface Ethernet1/59 -> Core facing interface
  ptp

interface Ethernet1/50 -> Host facing interface
  ttag
  ttag-strip
```

The following guidelines are applicable for PTP:

- The PTP feature can be enabled in a fabric when all the switches in the fabric have Cisco NX-OS Release 7.0(3)I7(1) or a higher version. Otherwise, the following error message is displayed:

  PTP feature can be enabled in the fabric, when all the switches have NX-OS Release 7.0(3)I7(1) or higher version. Please upgrade switches to NX-OS Release 7.0(3)I7(1) or higher version to enable PTP in this fabric.

- For hardware telemetry support in NIR, the PTP configuration is a prerequisite.
- If you are adding a non-cloud scale device to an existing fabric which contains PTP configuration, the following warning is displayed:

  TTAG is enabled fabric wide, when all devices are cloud scale switches so it cannot be enabled for newly added non cloud scale device(s).

- If a fabric contains both cloud scale and non-cloud scale devices, the following warning is displayed when you try to enable PTP:

  TTAG is enabled fabric wide, when all devices are cloud scale switches and is not enabled due to non cloud scale device(s).

Support for Super Spine Role in DCNM

Super Spine is a device that is used for interconnecting multiple spine-leaf PODs. Prior to the DCNM Release 11.3(1), it was possible to interconnect multiple VXLAN EVPN Easy fabrics via super spines. However, these

super spines had to be part of an external fabric. Within each Easy Fabric, an appropriate IGP is used for underlay connectivity. eBGP between the super spine layer in the external fabric and spine layer in the Easy Fabrics would be the recommended way of interconnecting multiple VXLAN EVPN Easy Fabrics. The eBGP peering can be configured via inter-fabric links or an appropriate mix of interface and eBGP configuration on the respective switches.

From DCNM Release 11.3(1), you have an extra interconnectivity option with super spines. You can have multiple spine-leaf PODs within the same Easy Fabric that are interconnected via super spines such that the same IGP domain extends across all the PODs, including the super spines. Within such a deployment, the BGP RRs and RPs (if applicable) are provisioned on the super spine layer. The spine layer becomes a pseudo interconnect between the leafs and super spines. VTEPs may be optionally hosted on the super spines if they have the border functionality.

The following Super Spine roles are supported in DCNM:

- Super Spine
- Border Super Spine
- Border Gateway Super Spine

A border super spine handles multiple functionalities including the functionalities of a super spine, RR, RP (optionally), and a border leaf. Similarly, a border gateway super spine serves a super spine, RR, RP (optional), and a border gateway. It's not recommended to overload border functionality on the super spine or RR layer. Instead, attach border leafs or border gateways to the super spine layer for external connectivity. The super spine layer serves as the interconnect with the RR or RP functionality.

The following are the characteristics of super spine switch roles in DCNM:

- Supported only for the **Easy_Fabric_11_1** template.
- Can only connect to spines and borders. The valid connections are:
  - Spines to super spines
  - Spines to border super spines and border gateway super spines
  - Super spines to border leafs and border gateway leafs
- RR or RP (if applicable) functionality is always be configured on super spines if they are present in a fabric. The maximum number of 4 RRs and RPs are supported even with Super Spines.
- Border Super Spine and Border Gateway Super Spine roles are supported for inter-fabric connections.
- vPC configurations aren't supported on super spines.
- Super spines don't support IPv6 underlay configuration.
- During the Brownfield import of switches, if a switch has the super spine role, the following error is displayed:

  Serial number: [super spine/border super spine/border gateway superspine] Role isn't supported with preserved configuration yes option.

Supported Topologies for Super Spine Switches

DCNM supports the following topologies with super spine switches.

Topology 1: Super Spine Switches in a Spine Leaf Topology

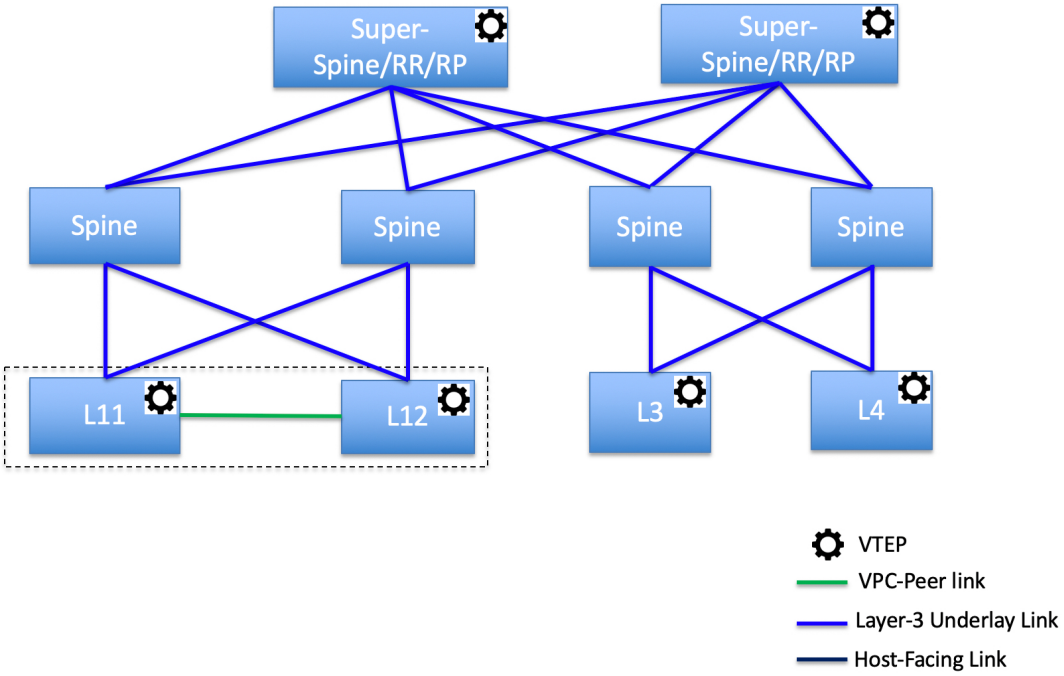

In this topology, leaf switches are connected to spines, and spines are then connected to Super Spines switches which can be super spines, border super spines, border gateway super spines.

Topology 2: Super Spine Switches Connected to Border

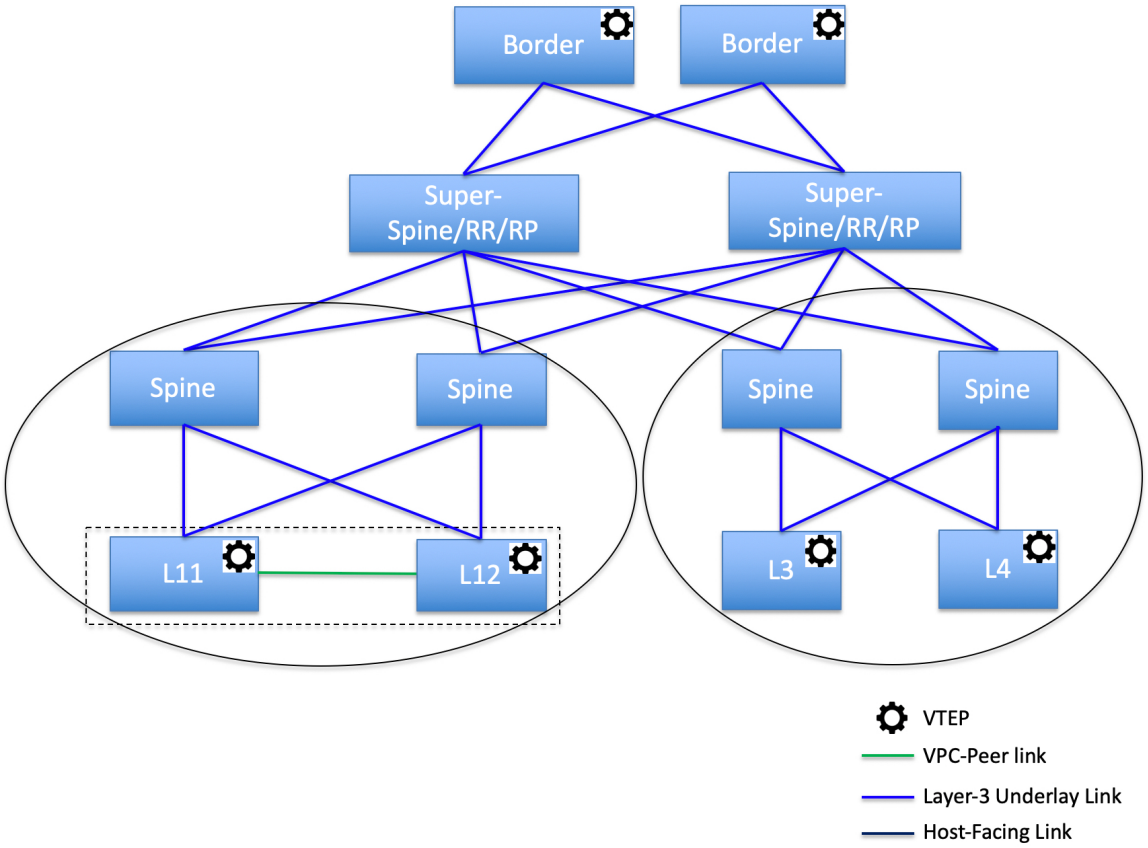

In this topology, there are four leaf switches connecting to the Spine switches, which are connected to the two Super Spine switches. These Super Spine switches are connected to the border or border gateway leaf switches.

Adding a Super Spine Switch to an Existing VXLAN BGP EVPN Fabric

Procedure

Step 1 Navigate to **Control** > **Fabric Builder**.

Step 2 From the **Fabric Builder** window, click **Add Switches** in the actions panel.

For more information, see Adding Switches to a Fabric, on page 73.

Step 3 Right-click an existing switch or the newly added switch, and use the **Set role** option to set the appropriate super spine role.

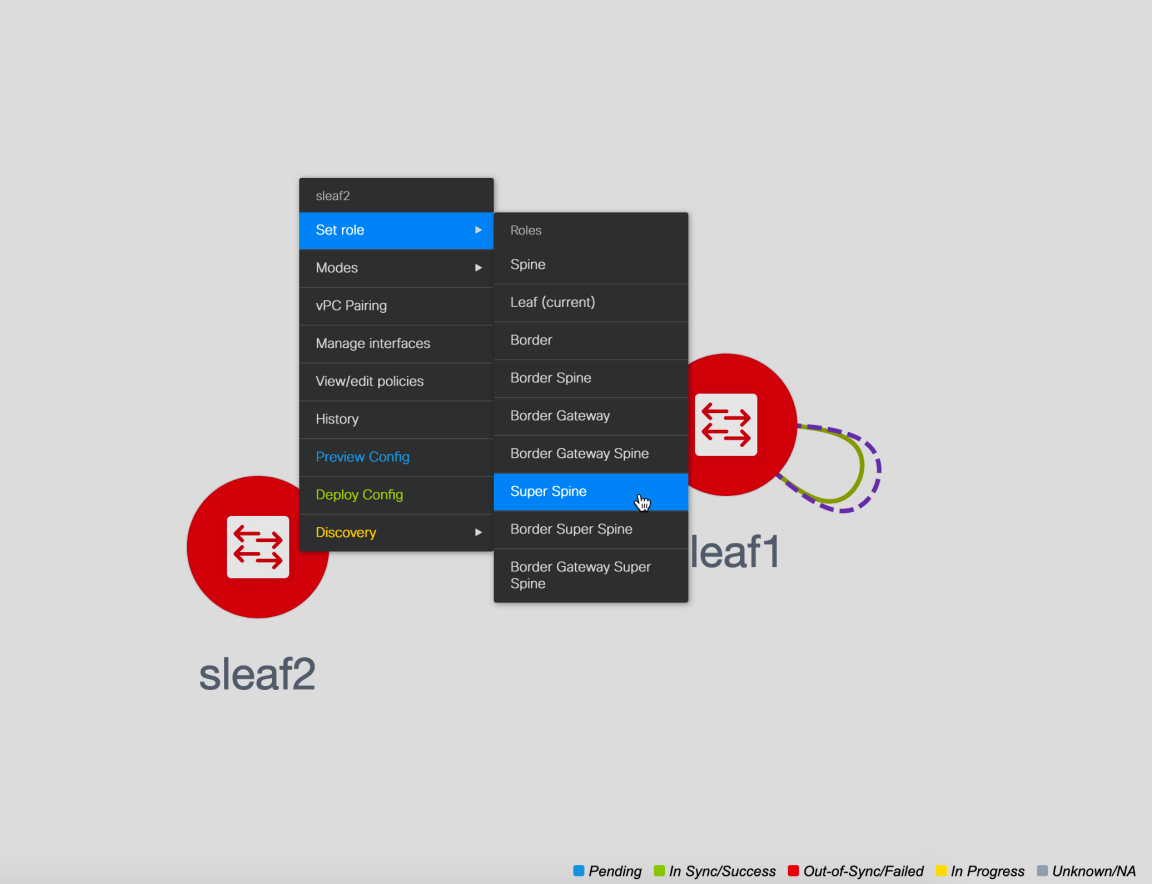

Note

- If the **Super Spine** role is present in the fabric, then the other super spine roles that you can assign for any new device are border super spine and border gateway super spine.
- If Super Spine or any of its variation role is not present in the fabric, you may assign the role to any new device provided that the same is connected to a non-border spine in the fabric. After a **Save & Deploy**, you will receive an error that can be resolved by clicking on the **Resolve** button as shown in the below steps.

Step 4 Click **Save & Deploy**.

An error is displayed saying:

Adding new switch with Super Spine role is not allowed, if save&deploy has already been performed in the fabric without any super spine role switch.

Step 5 Click the error, and click the **Resolve** button.

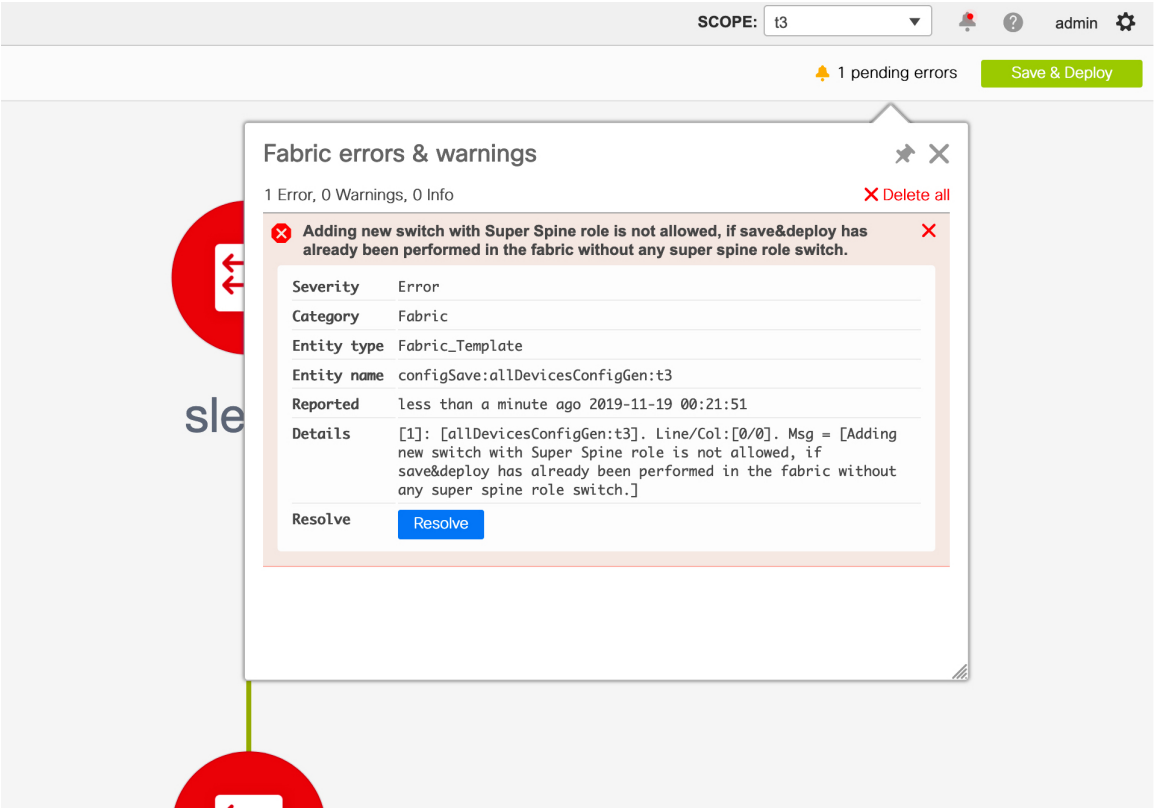

A confirmation dialog box is displayed asking whether you want to continue. If you click **Yes**, the following actions are performed by DCNM:

- Invalid connections are converted to hosts ports.
- Removes existing BGP neighborship between spines to leafs.
- Removes RRs or RPs from all the spine switches.

You should not add a device(s) with super spine, border super spine, or border gateway super spine role if the same will be connected to a border spine or border gateway spine that is already present in the fabric. This action will result in the below error after clicking **Save & Deploy**. If you want to use the existing device(s) with border spine roles, you need to remove the same and add them again with the appropriate role (spine or super spine and its variants) and valid connections.

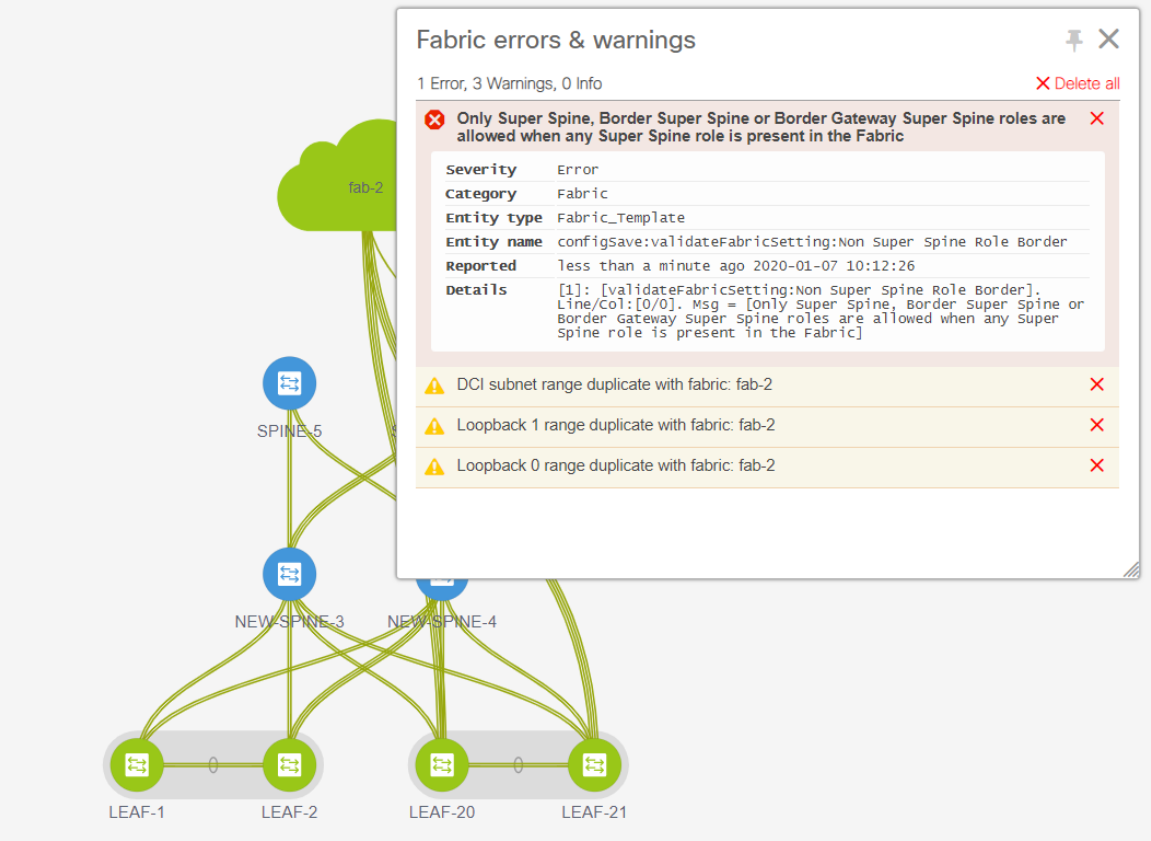

Changing the TCAM Configuration on a Device

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**.

Preselecting Switches as Route-Reflectors and Rendezvous-Points

This task shows how to preselect switches as Route-Reflectors (RRs) and Rendezvous-Points (RPs) before the first **Save & Deploy** operation.

Note This scenario is applicable when you have more than 2 spines and you want to control the preselection of RRs and RPs before the first Save & Deploy operation.

Procedure

Step 1 Import switches successfully.

Step 2 Create the **rr_state** or **rp_state** policies using **View/Edit Policies** on the spines or super spine switches, which should be preselected as RR or RP.

Note
- If there are more than 2 spines and the maximum number of RRs or RPs in the fabric settings is set to 2, then it's recommended to distribute RR and RP on different spines.
- If there are more than 4 spines and the maximum number of RRs or RPs in the fabric settings is set to 4, then it's recommended to distribute RR and RP on different spines.

Step 3 Click **Save & Deploy**, and then click **Deploy Config**.

The spines that have **rr_state** policies become RR and spines that have **rp_state** policies become RP.

Step 4 After **Save & Deploy**, if you want to replace the preselected RRs and RPs with a new set of devices, then old RR and RP devices should be removed from the fabric before performing the same steps.

Adding a vPC L3 Peer Keep-Alive Link

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

Note

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

Procedure

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

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

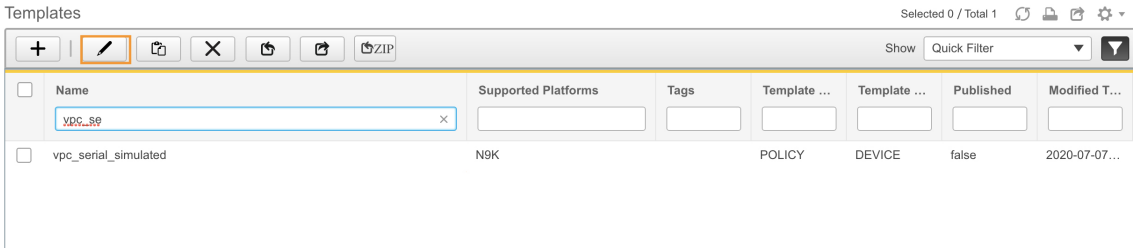

Step 3 Edit the template properties and set the **Template Sub Type** to **Device** so that this policy appears in **View/Edit Policies**.

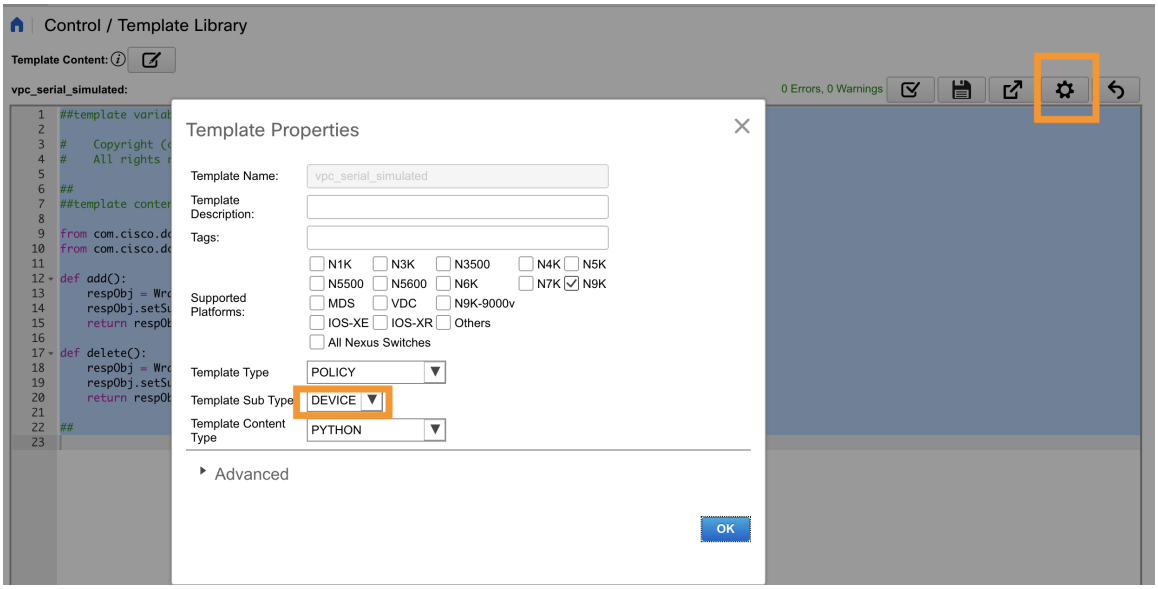

Step 4 Navigate to the **Fabric Builder** window and click on the fabric containing the vPC pair switches.

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

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

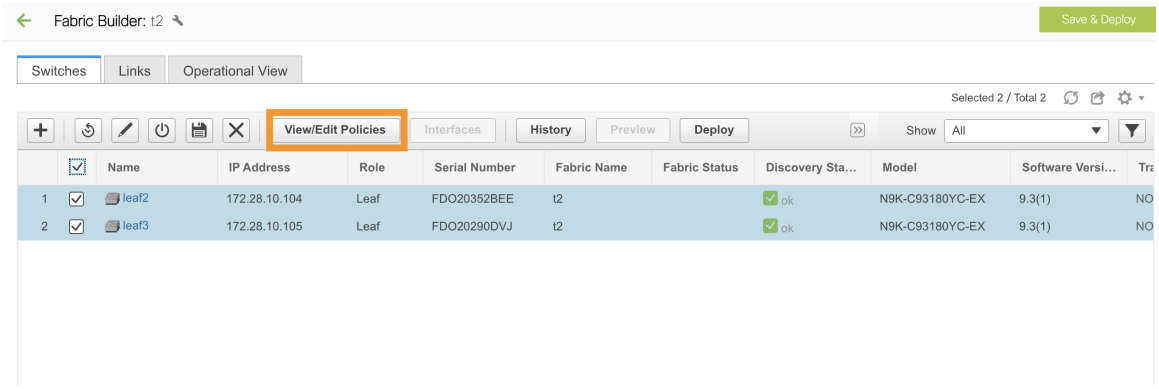

Step 6 Click + to add policies.

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

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

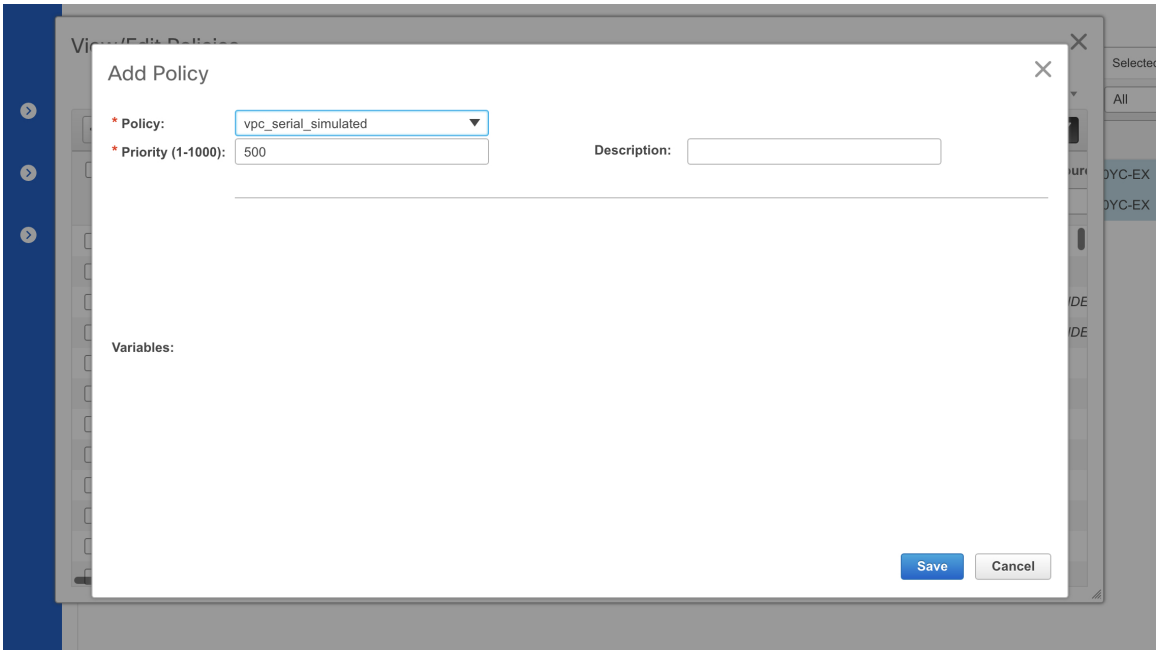

Step 8 Navigate back to **Tabular View** and click the **Links** tab.

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

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

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

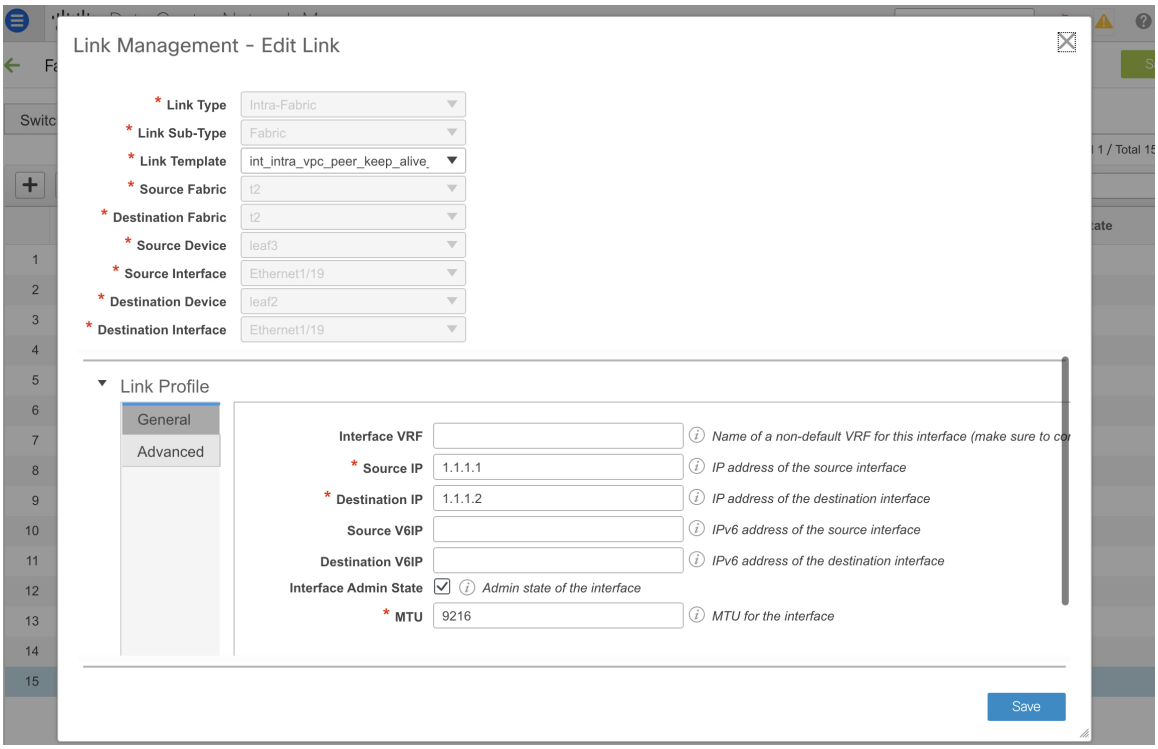

Step 11 Click **Save & Deploy**, and click **Preview Config** for one of the switches.

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

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

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

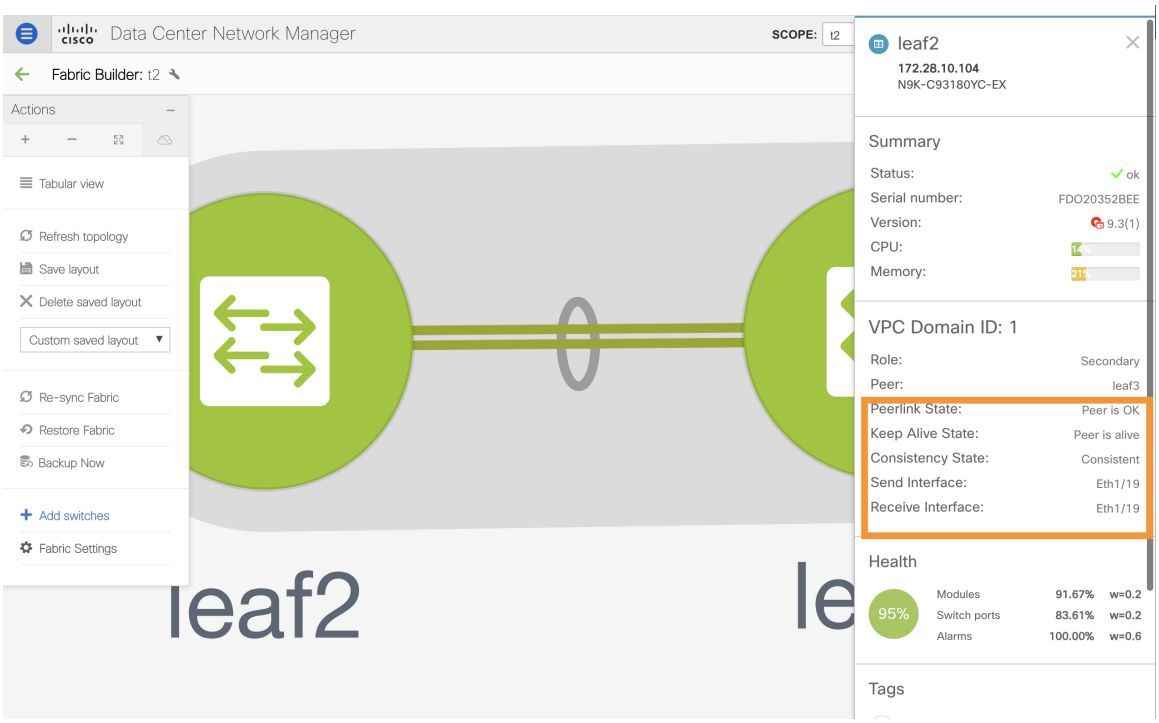

Changing the Local Authentication to AAA Authentication for Switches in a Fabric

Procedure

Step 1 Log in to DCNM and navigate to **Control > Fabric Builder**.

Step 2 Click the **Edit** icon for a fabric and add the AAA authentication commands in the **AAA Freeform Config** field under the **Manageability** tab.

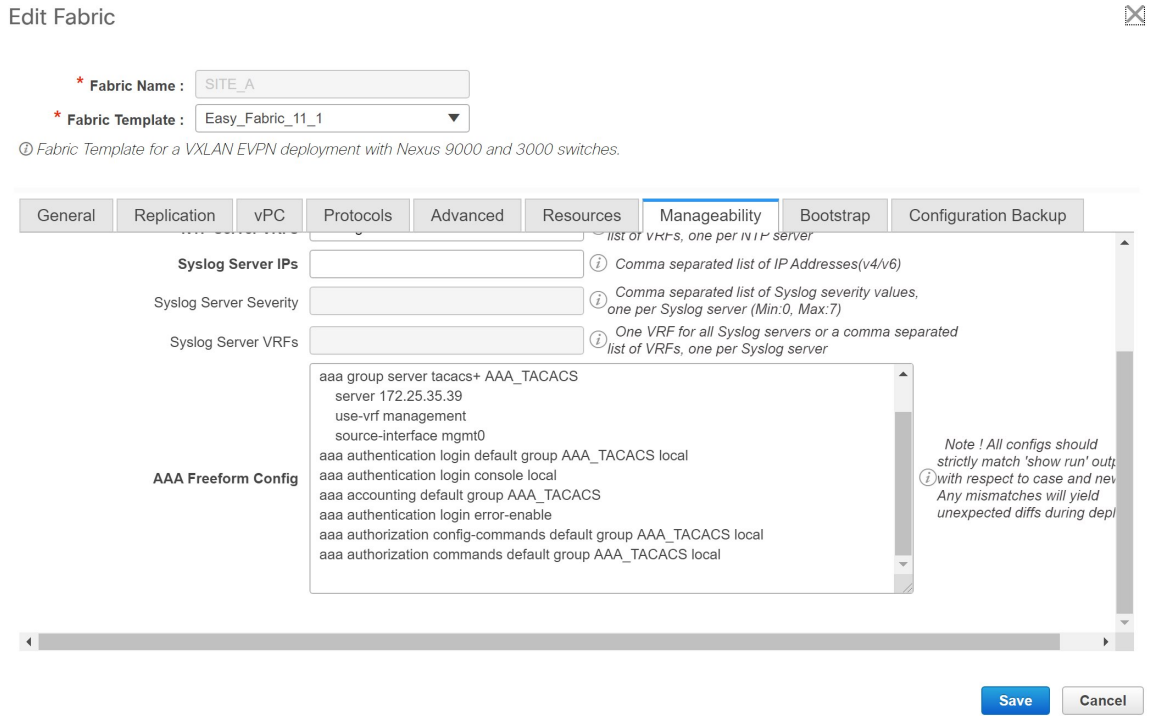

Step 3 In the **Fabric Builder** topology window, click **Add Switches**. Use the AAA credentials in this window to add switches into the DCNM.

Step 4 If you are importing switches in to the fabric via POAP, you need to have the AAA configs on the switch.

Navigate to the fabric settings and add the relevant commands in **Bootstrap Freeform Config**.

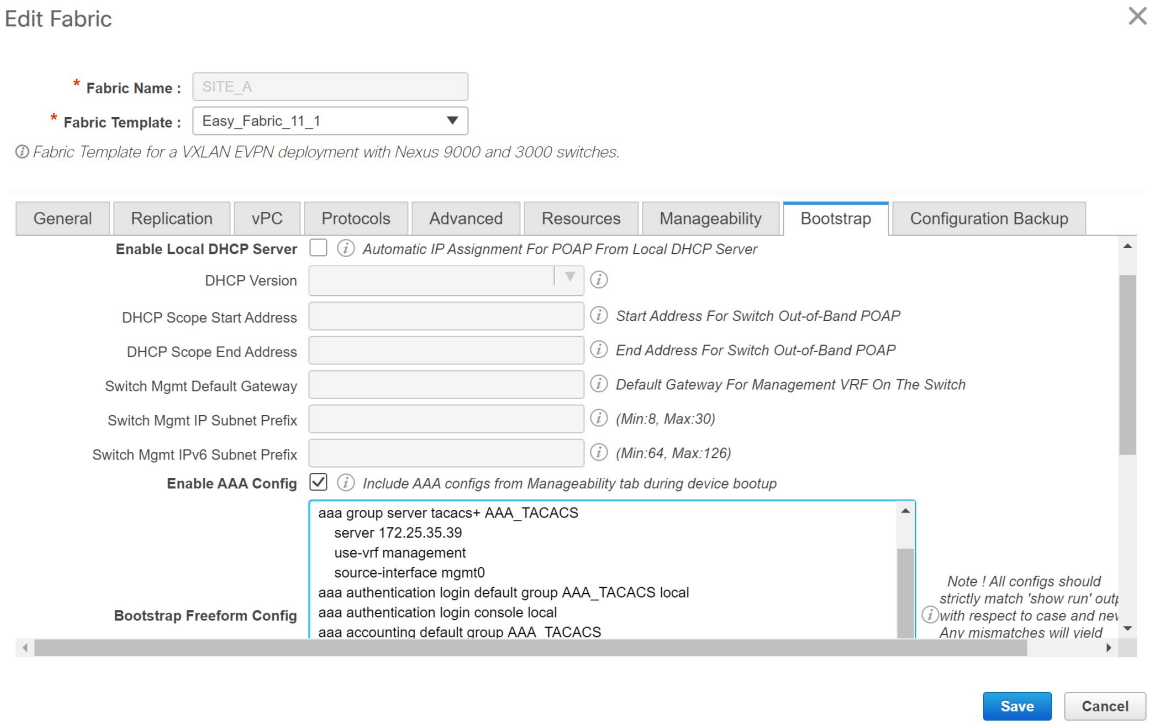

Step 5 In the **Fabric Builder** topology window, click **Add Switches**. In the **PowerON Auto Provisioning (POAP)** tab, click the **Add discovery credentials** icon and enter the discovery credentials.

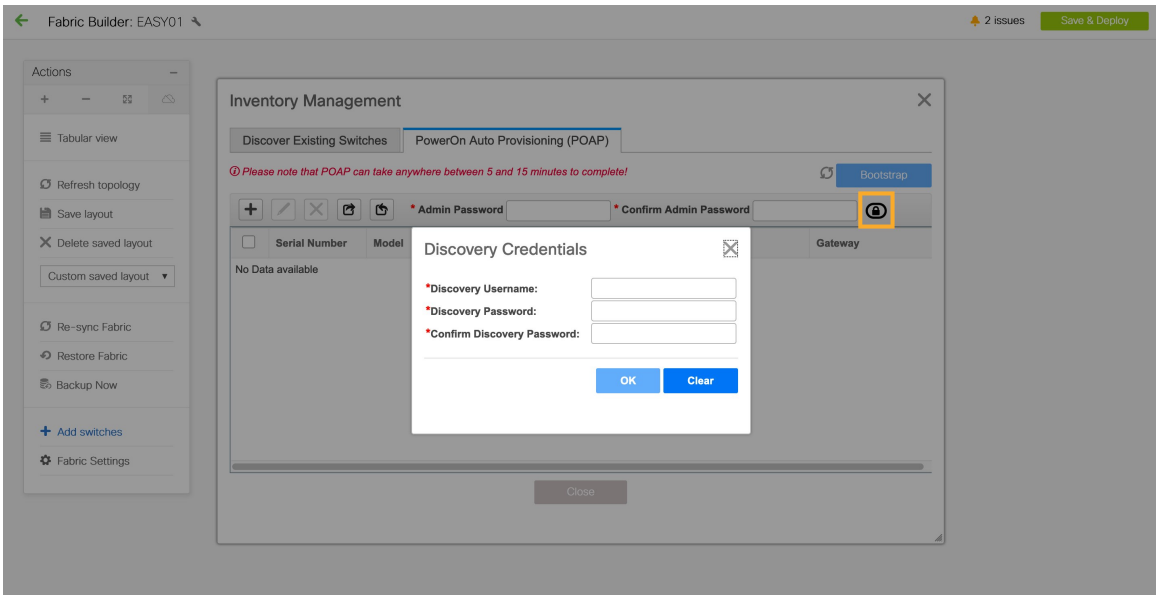

Click **Save & Deploy** after you complete adding switches.

IPv6 Underlay Support for Easy Fabric

From Cisco DCNM Release 11.3(1), you can create an Easy fabric with IPv6 only underlay. The IPv6 underlay is supported only for the **Easy_Fabric_11_1** template. For more information, see *Configuring a VXLAN Fabric with IPv6 Underlay*.

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*.

Configuring Fabrics with eBGP Underlay

You can use the **Easy_Fabric_eBGP** fabric template to create a fabric with eBGP underlay. For more information, see *Configuring a Fabric with eBGP Underlay*.

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. DCNM supports only the monitor mode for Cisco IOS-XR family devices.

- 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.
- If you are using the Cisco Nexus 7000 Series Switch with Cisco NX-OS Release 6.2(24a) on the LAN Classic or External fabrics, make sure to enable AAA IP Authorization in the fabric settings.
- You can discover the following non-Nexus devices in an external fabric:
  - IOS-XE family devices: Cisco CSR 1000v, Cisco IOS XE Gibraltar 16.10.x, Cisco ASR 1000 Series routers, and Cisco Catalyst 9000 Series Switches
  - IOS-XR family devices: ASR 9000 Series Routers, IOS XR Release 6.5.2 and Cisco NCS 5500 Series Routers, IOS XR Release 6.5.3
  - Arista 4.2 (Any model)
- Configure all the non-Nexus devices, except Cisco CSR 1000v, before adding them to the external fabric.
- From Cisco DCNM Release 11.4(1), you can configure non-Nexus devices as borders. You can create an IFC between a non-Nexus device in an external fabric and a Cisco Nexus device in an easy fabric. The interfaces supported for these devices are:
  - Routed
  - Subinterface
  - Loopback
- From Cisco DCNM, Release 11.4(1), you can configure a Cisco ASR 1000 Series routers and Cisco Catalyst 9000 Series switches as edge routers, set up a VRF-lite IFC and connect it as a border device with an easy fabric.
- Before a VDC reload, discover Admin VDC in the fabric. Otherwise, the reload operation does not occur.
- 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.
- In an external fabric, when you add the **switch_user** policy and provide the username and password, the password must be an encrypted string that is displayed in the **show run** command.

For example:

```
username admin password 5 $5$I4sapkBh$S7B7UcPH/iVTihLKH5sgldBeS3O2X1StQsvv3cmbYd1  role
 network-admin
```

In this case, the entered password should be **5$5$I4sapkBh$S7B7UcPH/iVTihLKH5sgldBeS3O2X1StQsvv3cmbYd1**.

- 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
  - LAN Classic fabric in Monitored or Managed Mode (Applicable for DCNM 11.4(1) or later)

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** tab as shown below.

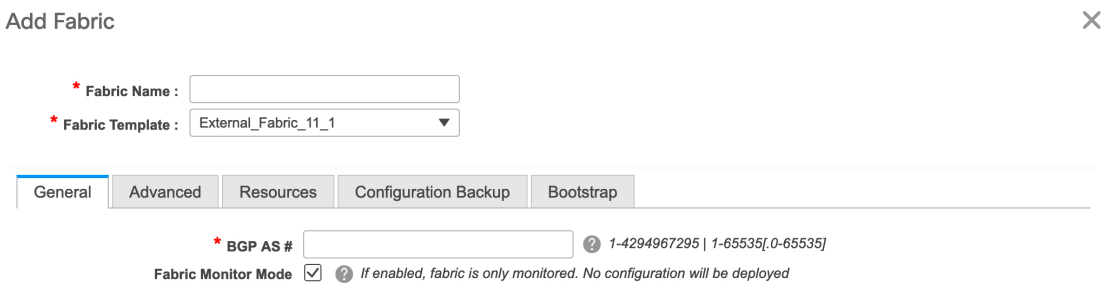

BGP AS # - Enter the BGP AS number.

Fabric Monitor Mode – Clear the check box if you want DCNM to manage the fabric. Keep the check box selected to enable a monitor only external fabric. DCNM supports only the monitor mode for Cisco IOS-XR family devices.

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.

The configurations must be pushed for non-Nexus devices before you discover them in the fabric. You cannot push configurations in the monitor mode.

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

4. Enter values in the fields under the **Advanced** tab.

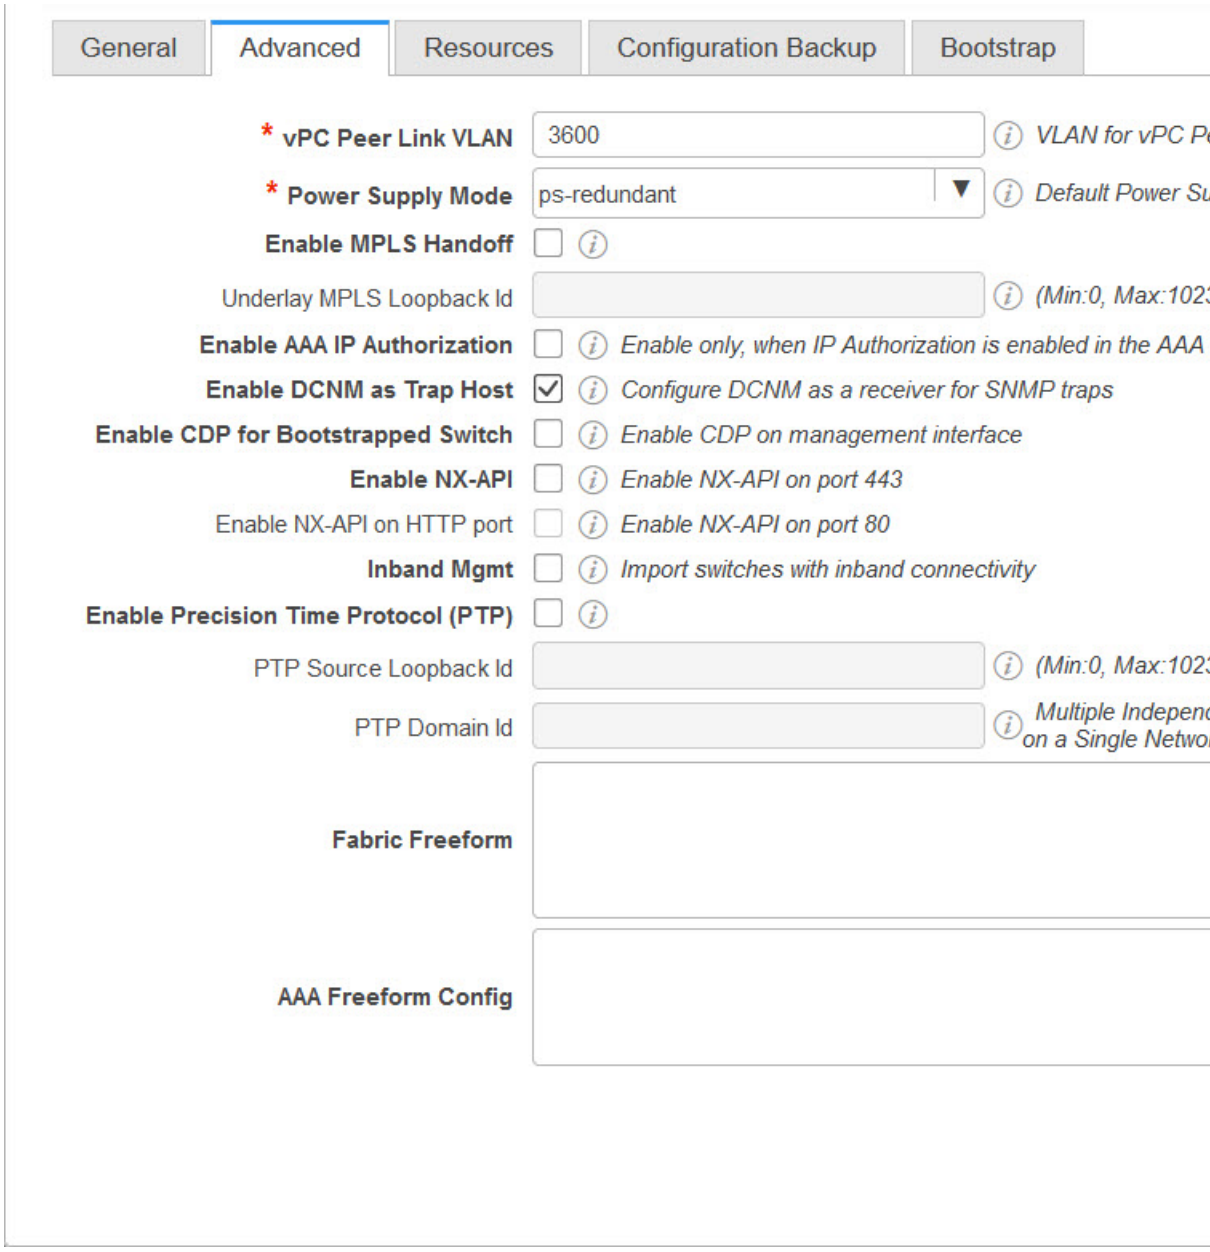

vPC Peer Link VLAN - The vPC peer link VLAN ID is autopopulated. Update the field to reflect the correct value.

Power Supply Mode - Choose the appropriate power supply mode.

Enable MPLS Handoff: Select the check box to enable the MPLS Handoff feature. For more information, see the *Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - MPLS SR and LDP Handoff* chapter.

Underlay MPLS Loopback Id: Specifies the underlay MPLS loopback ID. The default value is 101.

Enable AAA IP Authorization - Enables AAA IP authorization, when IP Authorization is enabled in the AAA Server

Enable DCNM as Trap Host - Select this check box to enable DCNM as a trap host.

Enable CDP for Bootstrapped Switch - Select the check box to enable CDP for bootstrapped switch.

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

Enable NX-API on HTTP - Specifies enabling of NX-API on HTTP. This check box is unchecked by default. Enable this check box and the **Enable NX-API** check box to use HTTP. If you uncheck this check box, the applications that use NX-API and supported by Cisco DCNM, such as Endpoint Locator (EPL), Layer 4-Layer 7 services (L4-L7 services), VXLAN OAM, and so on, start using the HTTPS instead of HTTP.

Note If you check the **Enable NX-API** check box and the **Enable NX-API on HTTP** check box, applications u HTTP.

Inband Mgmt: For External and Classic LAN Fabrics, this knob enables DCNM to import and manage of switches with inband connectivity (reachable over switch loopback, routed, or SVI interfaces) , in addition to management of switches with out-of-band connectivity (aka reachable over switch mgmt0 interface). The only requirement is that for Inband managed switches, there should be IP reachability from DCNM to the switches via the eth2 aka inband interface. For this purpose, static routes may be needed on the DCNM, that in turn can be configured via the Administration->Customization->Network Preferences option. After enabling Inband management, during discovery, provide the IPs of all the switches to be imported using Inband Management and set maximum hops to 0. DCNM has a pre-check that validates that the Inband managed switch IPs are reachable over the eth2 interface. Once the pre-check has passed, DCNM then discovers and learns about the interface on that switch that has the specified discovery IP in addition to the VRF that the interface belongs to. As part of the process of switch import/discovery, this information is captured in the baseline intent that is populated on the DCNM. For more information, see Inband Management in External Fabrics and LAN Classic Fabrics, on page 186.

Note Bootstrap or POAP is only supported for switches that are reachable over out-of-band connectivity, that is over switch mgmt0. The various POAP services on the DCNM are typically bound to the eth1 or out-of-ba interface. In scenarios, where DCNM eth0/eth1 interfaces reside in the same IP subnet, the POAP services are bound to both interfaces.

Enable Precision Time Protocol (PTP): Enables PTP across a fabric. When you select this check box, PTP is enabled globally and on core-facing interfaces. Additionally, the PTP Source Loopback Id and PTP Domain Id fields are editable. For more information, see Precision Time Protocol for External Fabrics and LAN Classic Fabrics, on page 188.

PTP Source Loopback Id: Specifies the loopback interface ID Loopback that is used as the Source IP Address for all PTP packets. The valid values range from 0 to 1023. The PTP loopback ID cannot be the same as RP, Phantom RP, NVE, or MPLS loopback ID. Otherwise, an error will be generated. The PTP

loopback ID can be the same as BGP loopback or user-defined loopback which is created from DCNM. If the PTP loopback ID is not found during Save & Deploy, the following error is generated: `Loopback interface to use for PTP source IP is not found. Please create PTP loopback interface on all the devices to enable PTP feature.`

PTP Domain Id: Specifies the PTP domain ID on a single network. The valid values range from 0 to 127.

Fabric Freeform: You can apply configurations globally across all the devices discovered in the external fabric using this freeform field. The devices in the fabric should belong to the same device-type and the fabric should not be in monitor mode. The different device types are:

- NX-OS
- IOS-XE
- IOS-XR
- Others

Depending on the device types, enter the configurations accordingly. If some of the devices in the fabric do not support these global configurations, they will go out-of-sync or fail during the deployment. Hence, ensure that the configurations you apply are supported on all the devices in the fabric or remove the devices that do not support these configurations.

5. Fill up the **Resources** tab as shown below.

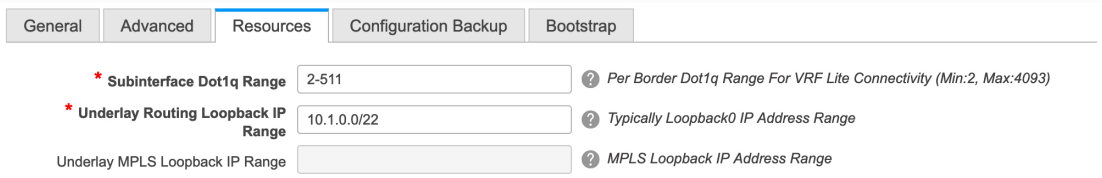

Subinterface Dot1q Range - The subinterface 802.1Q range and the underlay routing loopback IP address range are autopopulated.

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

Underlay MPLS Loopback IP Range: Specifies the underlay MPLS SR or LDP loopback IP address range.

The IP range should be unique, that is, it should not overlap with IP ranges of the other fabrics.

Enable AAA IP Authorization - Enables AAA IP authorization, when IP Authorization is enabled in the AAA Server

Enable DCNM as Trap Host - Select this check box to enable DCNM as a trap host.

6. Fill up the **Configuration Backup** tab as shown below.

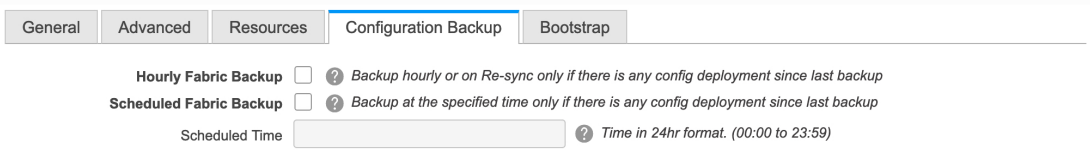

The fields on this tab are:

Hourly Fabric Backup: Select the check box to enable an hourly backup of fabric configurations and the intent.

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.

The hourly backups are triggered during the first 10 minutes of the hour.

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.

The backup process is initiated after you click **Save**.

The scheduled backups are triggered exactly at the time you specify with a delay of up to two minutes. The scheduled backups are triggered regardless of the configuration deployment status.

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:

- 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.

7. Click the **Bootstrap** tab.

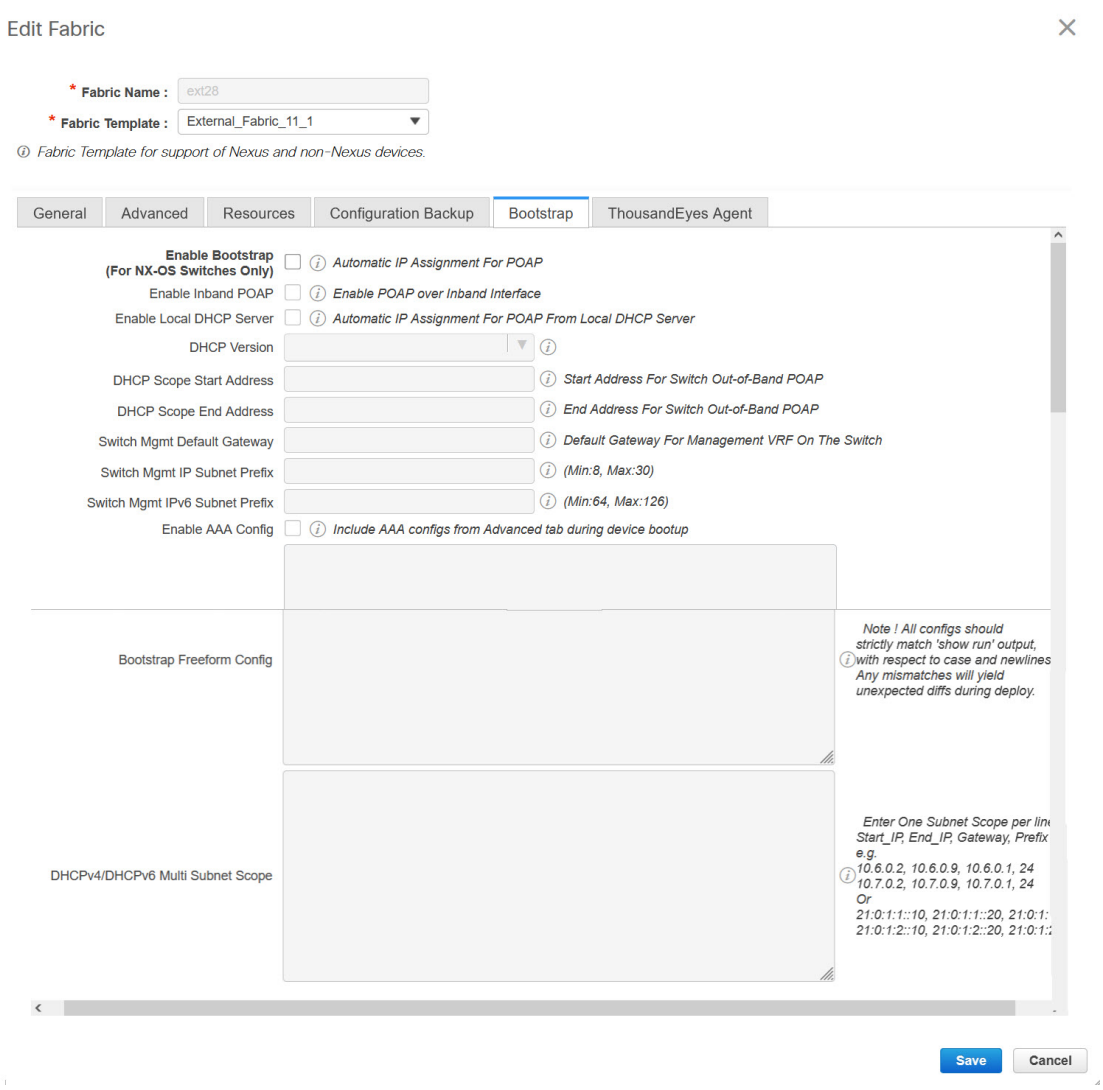

Enable Bootstrap - Select this check box to enable the bootstrap feature. After you enable bootstrap, you can enable the DHCP server for automatic IP address assignment using one of the following methods:

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

Enable Inband POAP – For External and Classic LAN Fabrics, this knob enables DCNM to import and manage of switches with inband POAP (reachable over switch loopback, routed, or SVI interfaces). As prerequisites, you must first enable Inband POAP feature using **appmgr inband-poap enable** command on the DCNM server. On DCNM Web UI, Edit Fabric to enable **Inband Mgmt** feature.

Note You must select **Inband Mgmt** checkbox in the **Advanced** tab to enable **Inband POAP feature**.

For more information, see Inband POAP Management in External Fabrics and LAN Classic Fabrics, on page 187.

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, all the remaining fields become editable.

DHCP Version – Select DHCPv4 or DHCPv6 from this drop-down list. When you select DHCPv4, the **Switch Mgmt IPv6 Subnet Prefix** field is disabled. If you select DHCPv6, the **Switch Mgmt IP Subnet Prefix** is disabled.

Note Cisco DCNM IPv6 POAP is not supported with Cisco Nexus 7000 Series Switches. Cisco Nexus 9000 an 3000 Series Switches support IPv6 POAP only when switches are either L2 adjacent (eth1 or out-of-band subnet must be a /64) or they are L3 adjacent residing in some IPv6 /64 subnet. Subnet prefixes other than /64 are not supported.

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

Switch Mgmt IP 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.

Switch Mgmt IPv6 Subnet Prefix - Specifies the IPv6 prefix for the Mgmt0 interface on the switch. The prefix should be between 112 and 126. This field is editable if you enable IPv6 for DHCP.

Enable AAA Config - Select this check box to include AAA configs from Advanced tab during device bootup.

Bootstrap Freeform Config - (Optional) Enter other commands as needed. For example, if you are using AAA or remote authentication-related configurations, 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 347.

DHCPv4/DHCPv6 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

8. 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. Choose the device type from the **Device Type** drop-down list.

   The options are **NX-OS**, **IOS XE**, **IOS XR**, and **Other**.

   - Choose **NX-OS** to discover a Cisco Nexus switch.
   - Choose **IOS XE** to discover a CSR device.
   - Choose **IOS XR** to discover an ASR device.
   - Choose **Other** to discover non-Cisco devices.

   Click the appropriate radio button. Refer the *Connecting Cisco Data Center and a Public Cloud* chapter for more information on adding Cisco CSR 1000v.

   Refer the *Adding non-Nexus Devices to External Fabrics* section for more information on adding other non-Nexus devices.

   Config compliance is disabled for all non-Nexus devices except for Cisco CSR 1000v.

4. Enter the administrator username and password of the switch.

5. 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.

6. 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.

7. Click Refresh topology to view the latest topology view.

8. *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.

Preview Config - View the pending configuration and the side-by-side comparison of the running and expected configuration.

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.

9. 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.

10. Click Deploy Config at the bottom part of the screen to initiate pending configuration onto the switch.

11. 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 followi 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 chan the SSH password to match the SNMP password on the switch. If SNMP authentication fails, discove is stopped with authentication error. If SNMP authentication passes but SSH authentication fails, DCN discovery continues, but the switch status shows a warning for the SSH error.

Move an External Fabric Under an MSD Fabric

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

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

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

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

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

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

External Fabric Depiction in an MSD Fabric Topology

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

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

External Fabric Switch Operations

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

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

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

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

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

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

External Fabric Links

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

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

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

   The Switches | Links screen comes up.

2. Choose one or more check boxes 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**.

   To delete a neighbor, select a switch and click **Delete Neighbor**.

Discovering New Switches

To discover new switches, perform the following steps:

Procedure

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

Boot the Cisco NX-OS and setup switch credentials.

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

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

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

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

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

The **Edit Fabric** screen is displayed.

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

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

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

The fabric topology screen appears.

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

The Inventory Management screen comes up.

Step 9 Click the **POAP** tab.

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

Note At the top left part of the screen, export and import options are provided to export and import the .csv file that contains the switch information. You can pre-provision a device using the import option too.

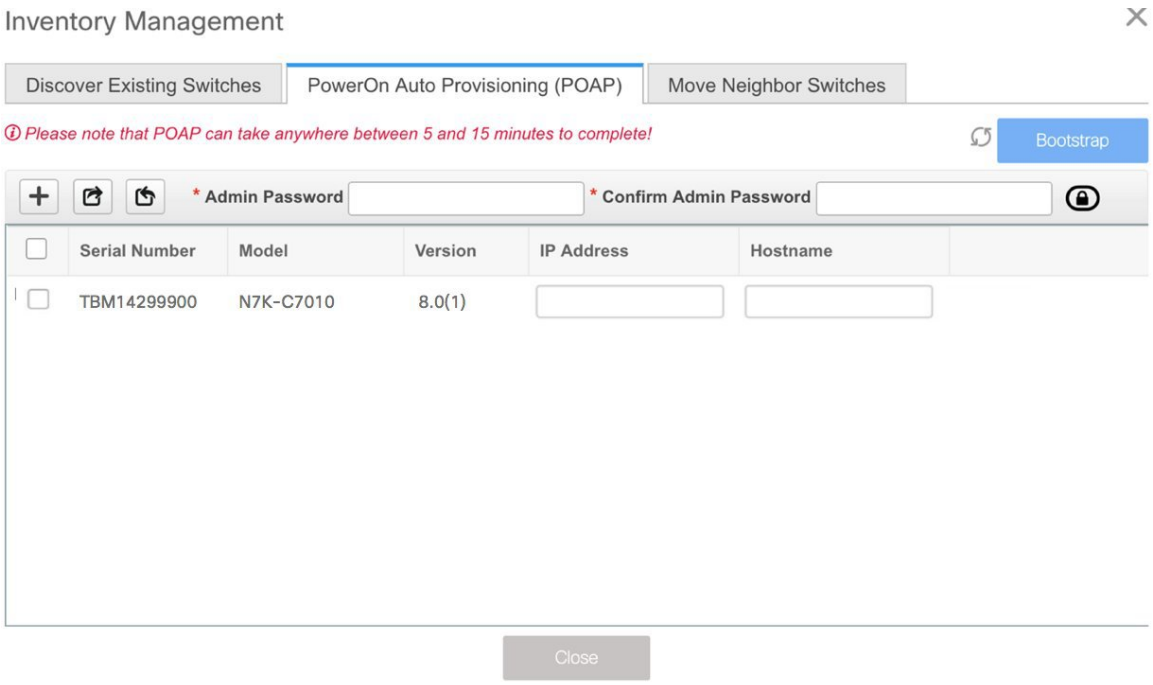

Select the checkbox next to the switch and add switch credentials: IP address and host name.

Based on the IP address of your device, you can either add the IPv4 or IPv6 address in the **IP Address** field.

Beginning with Release 11.2(1), you can provision devices in advance. To pre-provision devices, refer to Pre-provisioning a Device , on page 86.

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

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

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

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

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

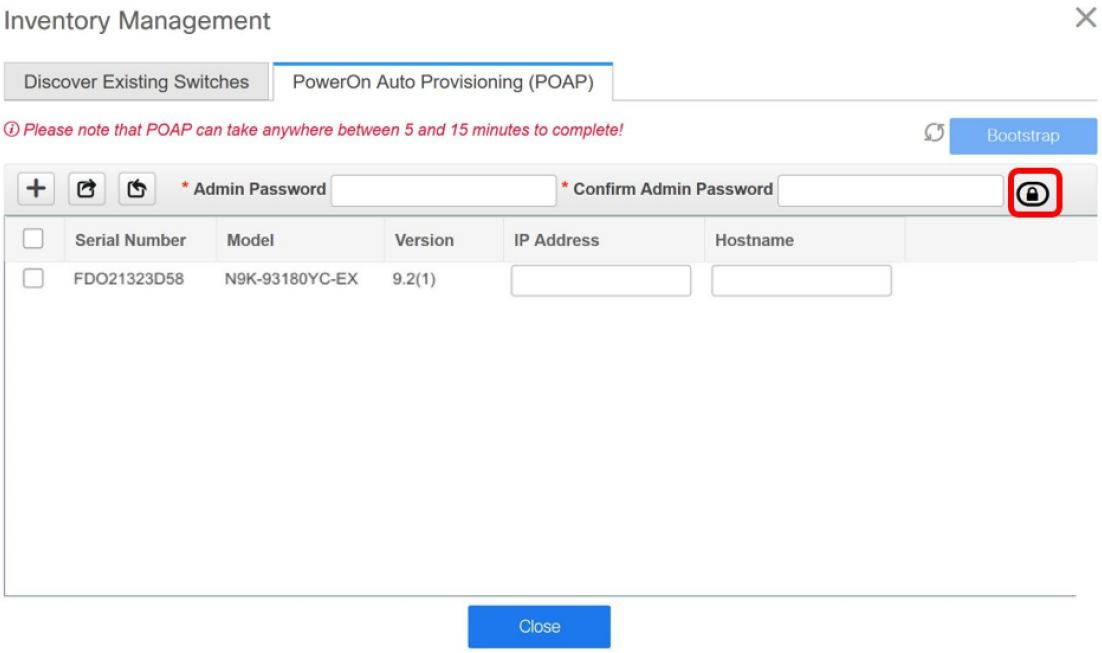

b) In the **Discovery Credentials** window, enter the discovery credentials such as discovery username and password.

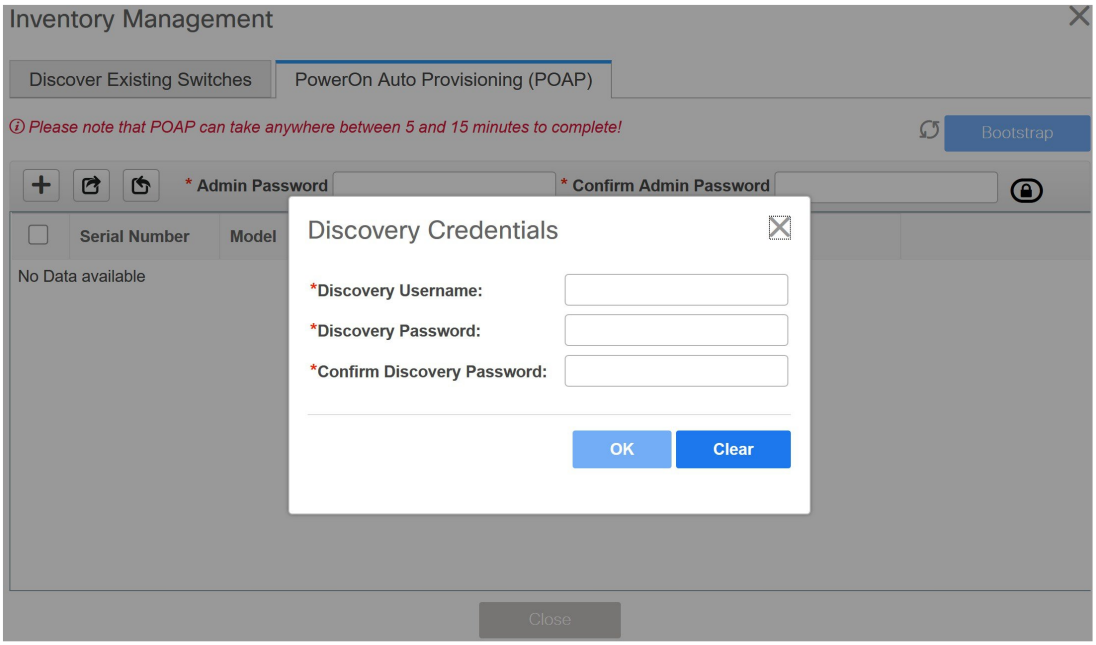

Click **OK** to save the discovery credentials.

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

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

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

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

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

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

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

Step 15 Monitor and check the switch for POAP completion.

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

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

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

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

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

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

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

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

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

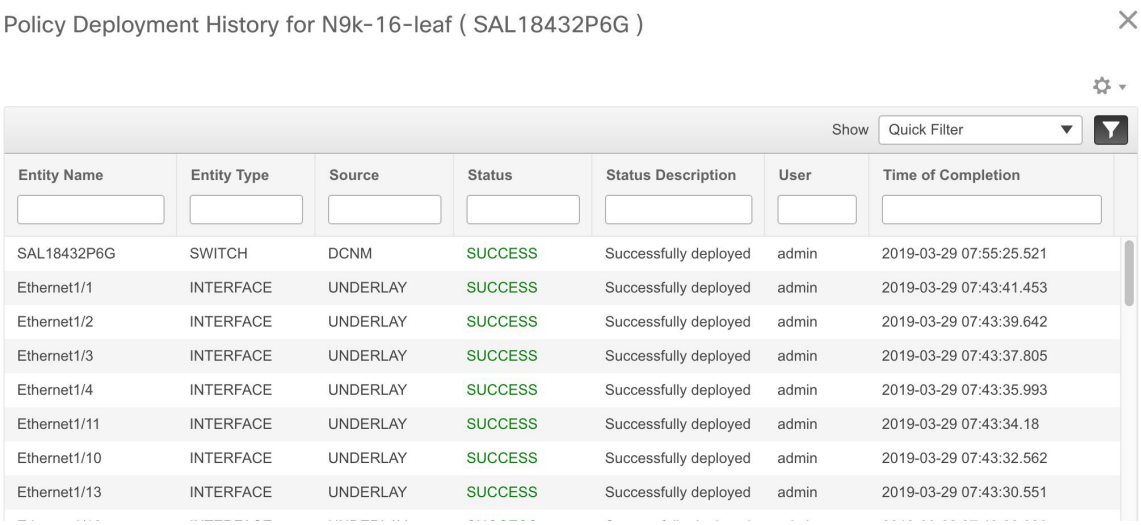

Click the **Success** link in the **Status** column for more details. An example:

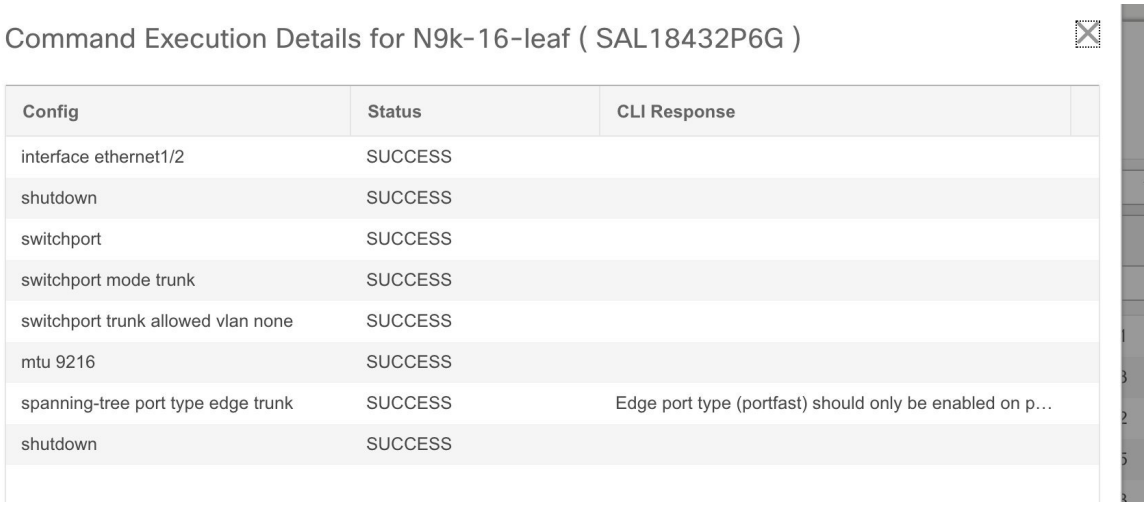

Step 21 On the DCNM UI, the discovered switches can be seen in the fabric topology.

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

- vPC pairing.
- Breakout interfaces

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

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

Adding non-Nexus Devices to External Fabrics

You can discover non-Nexus devices in an external fabric. Refer the *Cisco DCNM Compatibility Matrix* to see the non-Nexus devices supported by Cisco DCNM.

Only Cisco Nexus switches support SNMP discovery by default. Hence, configure all the non-Nexus devices before adding it to the external fabric. Configuring the non-Nexus devices includes configuring SNMP views, groups, and users. See the *Configuring non-Nexus Devices for Discovery* section for more information.

Cisco CSR 1000v is discovered using SSH. Cisco CSR 1000v does not need SNMP support because it can be installed in clouds where SNMP is blocked for security reasons. See the *Connecting Cisco Data Center and a Public Cloud* chapter to see a use case to add Cisco CSR 1000v, Cisco IOS XE Gibraltar 16.10.x to an external fabric.

However, Cisco DCNM can only access the basic device information like system name, serial number, model, version, interfaces, up time, and so on. Cisco DCNM does not discover non-Nexus devices if the hosts are part of CDP or LLDP.

The settings that are not applicable for non-Nexus devices appear blank, even if you get many options when you right-click a non-Nexus device in the fabric topology window. You cannot add or edit interfaces for ASR 9000 Series Routers and Arista switches.

From Cisco DCNM, Release 11.4(1), you can add IOS-XE devices like Cisco Catalyst 9000 Series switches and Cisco ASR 1000 Series Routers as well to external fabrics.

Configuring non-Nexus Devices for Discovery

Before discovering any non-Nexus device in Cisco DCNM, configure it on the switch console.

Configuring IOS-XE Devices for Discovery

Before you discover the Cisco IOS-XE devices in DCNM, perform the following steps:

Procedure

Step 1 Run the following SSH commands on the switch console.

```
switch (config)# hostname <hostname>
switch (config)# ip domain name <domain_name>
switch (config)# crypto key generate rsa
switch (config)# ip ssh time-out 90
switch (config)# ip ssh version 2
switch (config)# line vty 1 4
switch (config-line)# transport input ssh
switch (config)# aaa authentication login default local
switch (config)# aaa authorization exec default local none
switch (config)# username admin privilege secret <password>
switch (config)# aaa new-model
switch (config)# session-id-common
```

Step 2 Run the following command in DCNM console to perform an SNMP walk.

```
snmpbulkwalk -v3 -u admin -A <password> -l AuthNoPriv -a MD5 ,switch-mgmt-IP>
        .1.3.6.1.2.1.2.2.1.2
```

Step 3 Run the following SNMP command on the switch console.

```
snmp-server user username group-name [remote host {v1 | v2c | v3 [encrypted] [auth {md5 |
sha} auth-password]} [priv des 256 privpassword] vrf vrf-name [access access-list]
```

Configuring Arista Devices for Discovery

Enable Privilege Exec mode using the following command:

```
switch> enable
switch#

switch# show running confiruation | grep aaa        /* to view the authorization*/
aaa authorization exec default local
```

Run the following commands in the switch console to configure Arista devices:

```
switch# configure terminal
switch (config)# username dcnm privilege 15 role network-admin secret cisco123
snmp-server view view_name SNMPv2 included
snmp-server view view_name SNMPv3 included
snmp-server view view_name default included
snmp-server view view_name entity included
snmp-server view view_name if included
snmp-server view view_name iso included
snmp-server view view_name lldp included
snmp-server view view_name system included
snmp-server view sys-view default included
snmp-server view sys-view ifmib included
snmp-server view sys-view system included
snmp-server community private ro
snmp-server community public ro
snmp-server group group_name v3 auth read view_name
snmp-server user username group_name v3 auth md5 password priv aes password
```

Note SNMP password should be same as the password for username.

You can verify the configuration by running the **show run** command, and view the SNMP view output by running the **show snmp view** command.

Show Run Command

```
switch (config)# snmp-server engineID local f5717f444ca824448b00
snmp-server view view_name SNMPv2 included
snmp-server view view_name SNMPv3 included
snmp-server view view_name default included
snmp-server view view_name entity included
snmp-server view view_name if included
snmp-server view view_name iso included
snmp-server view view_name lldp included
snmp-server view view_name system included
snmp-server view sys-view default included
snmp-server view sys-view ifmib included
snmp-server view sys-view system included
snmp-server community private ro
snmp-server community public ro
snmp-server group group_name v3 auth read view_name
snmp-server user user_name group_name v3 localized f5717f444ca824448b00 auth md5
be2eca3fc858b62b2128a963a2b49373 priv aes be2eca3fc858b62b2128a963a2b49373
```

```
!
spanning-tree mode mstp
!
service unsupported-transceiver labs f5047577
!
aaa authorization exec default local
!
no aaa root
!
username admin role network-admin secret sha512
$6$5ZKs/7.k2UxrWDg0$FOkdVQsBTnOquW/9AYx36YUBSPNLFdeuPIse9XgyHSdEOYXtPyT/0sMUYYdkMffuIjgn/d9rx/Do71XSbygSn/
username cvpadmin role network-admin secret sha512
$6$fLGFj/PUcuJT436i$Sj5G5c4y9cYjI/BZswjjmZW0J4npGrGqIyG3ZFk/ULza47Kz.d31q13jXA7iHM677gwqQbFSH2/3oQEaHRq08.
username dcnm privilege 15 role network-admin secret sha512
$6$M48PNrCdg2EITEdG$iiB880nvFQQlrWoZwOMzdt5EfkuCIraNqtEMRS0TJUhNKCQnJN.VDLFsLAmP7kQBo.C3ct4/.n.2eRlcP6hij/
```

Show SNMP View Command

```
configure terminal# show snmp view
view_name SNMPv2 - included
view_name SNMPv3 - included
view_name default - included
view_name entity - included
view_name if - included
view_name iso - included
view_name lldp - included
view_name system - included
sys-view default - included
sys-view ifmib - included
sys-view system - included
leaf3-7050sx#show snmp user

User name : user_name
Security model : v3
Engine ID : f5717f444ca824448b00
Authentication : MD5
Privacy : AES-128
Group : group_name
```

Configuring Cisco IOS-XR Devices for Discovery

Run the following commands in the switch console to configure IOS-XR devices:

```
switch# configure terminal
switch (config)# snmp-server view view_name cisco included
snmp-server view view_name mib-2 included
snmp-server group group_name v3 auth read view_name write view_name
snmp-server user user_name group_name v3 auth md5 password priv des56 password SystemOwner
```

Note SNMP password should be same as password for username.

You can verify the configuration by running the show run command.

Configuration and Verification of Cisco IOS-XR Devices

```
RP/0/RSP0/CPU0:ios(config)#snmp-server view view_name cisco included
RP/0/RSP0/CPU0:ios(config)#snmp-server view view_name mib-2 included
RP/0/RSP0/CPU0:ios(config)#snmp-server group group_name v3 auth read view_name write view_name
RP/0/RSP0/CPU0:ios(config)#snmp-server user user_name group_name v3 auth md5 password priv
 des56 password SystemOwner
RP/0/RSP0/CPU0:ios(config)#commit Day MMM DD HH:MM:SS Timezone
RP/0/RSP0/CPU0:ios(config)#
RP/0/RSP0/CPU0:ios(config)#show run snmp-server Day MMM DD HH:MM:SS Timezone snmp-server
user user_name group1 v3 auth md5 encrypted 10400B0F3A4640585851 priv des56 encrypted
000A11103B0A59555B74 SystemOwner
snmp-server view view_name cisco included
snmp-server view view_name mib-2 included
snmp-server group group_name v3 auth read view_name write view_name
```

Discovering non-Nexus Devices in an External Fabric

To add non-Nexus devices to an external fabric in the fabric topology window, perform the following steps:

Before you begin

Ensure that the configurations are pushed for non-Nexus devices before adding them to an external fabric. You cannot push configurations in a fabric in the monitor mode.

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 switch. You can import more than one switch by providing the IP address range. For example: 10.10.10.40-60 The switches must be properly cabled and connected to the DCNM server and the switch status must be manageable.
Device Type	• Choose **IOS XE** from the drop-down list for adding Cisco CSR 1000v, Cisco ASR 1000 Series routers, or Cisco Catalyst 9000 Series Switches. • Choose **IOS XR** from the drop-down list for adding Cisco NCS 5500 Series Routers, IOS XR Release 6.5.3. • Choose **Other** from the drop-down list for adding non-Cisco devices, like Arista switches.

Field	Description
Username	Enter the username.
Password	Enter the password.

Note An error message appears if you try to discover a device that is already discovered.

Set the password of the device in the **LAN Credentials** window if the password is not set. To navigate to the **LAN Credentials** window from the Cisco DCNM Web UI, choose **Administration > LAN Credentials**.

Step 3 Click **Start Discovery**.

The **Scan Details** section appears with the switch details populated.

Step 4 Check the check boxes next to the switches you want to import.

Step 5 Click **Import into fabric**.

The switch discovery process is initiated. The **Progress** column displays the progress.

Discovering devices takes some time. A pop-up message appears at the bottom-right about the device discovery after the discovery progress is **100%**, or **done**. For example: **<ip-address> added for discovery.**

Step 6 Click **Close**.

The fabric topology window appears with the switches.

Step 7 (Optional) Click **Refresh topology** to view the latest topology view.

Step 8 (Optional) 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 9 (Optional) View the details of the device.

After the discovery of the device:

- The discovery status changes to **ok** in green with a check box checked next to it.
- The value of the device under the **Fabric Status** column changes to **In-Sync**.

Note When a switch is in **Unreachable** discovery status, the last available information of the switch is retained in other columns.

What to do next

Set the appropriate role. Right-click the device, choose **Set role**.

Pre-provisioning a Device

From Cisco DCNM Release 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
  - Ethernet ports
  - Port-channel
  - vPC
  - ST FEX
  - AA FEX
  - Loopback
  - Overlay network configurations
- The pre-provisioned devices do not support the following configurations in DCNM:
  - Inter-Fabric links
  - Sub-interface
  - Interface breakout configuration
- 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 must 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 Click **Control > Fabric Builder**.

The **Fabric Builder** screen is displayed.

Step 2 Click within the fabric box.

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

The **Inventory Management** screen is displayed.

Step 4 Click the **POAP** tab.

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

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

The Add a new device screen comes up.

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

c. Click **Save**.

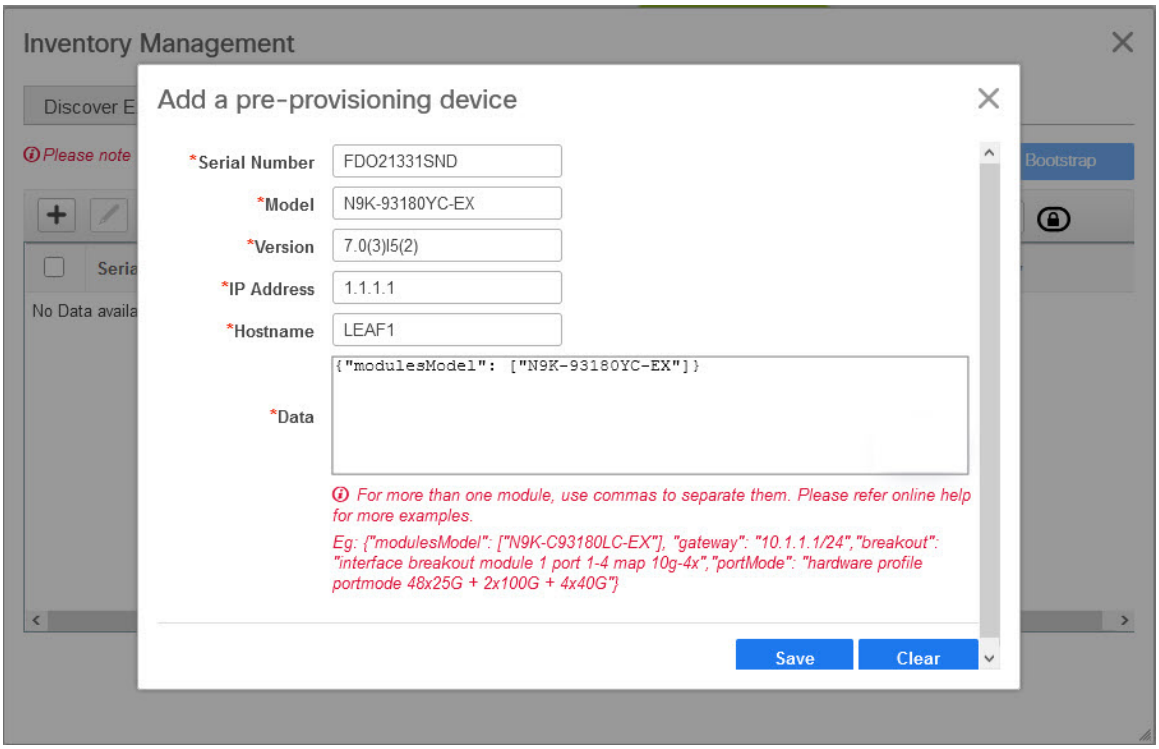

IP Address: Specify the IPv4 or IPv6 address of the new device.

Serial Number: The serial number for the new device. Serial number is found in the Cisco Build of Material Purchase and you can refer to these values while using the pre-provisioning feature.

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

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

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

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

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

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

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

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

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

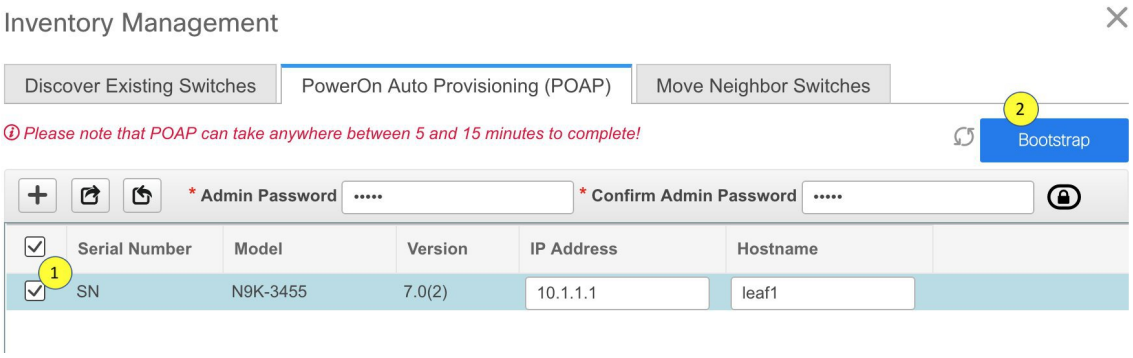

The leaf1 device appears in the fabric topology.

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

Note When a switch is in **Unreachable** discovery status, the last available information of the switch is retained in other columns.

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 253.

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 one or more switches are online to provision the host ports. This action must be performed before overlays are provisioned for the host port attachment.

Pre-provisioning an Ethernet Interface

From DCNM Release 11.4(1), you can pre-provision Ethernet interfaces in the **Interface** window. This pre-provisioning feature is supported in the Easy, External, and eBGP fabrics. You can add Ethernet interfaces to only pre-provisioned devices before they are discovered in DCNM.

Note Before attaching a network/VRF, you must pre-provision the Ethernet interface before adding it to Port-channels, vPCs, ST FEX, AA FEX, loopback, subinterface, tunnel, ethernet, and SVI configurations.

Before you begin

Make sure that you have a preprovisioned device in your fabric. For information, see Pre-provisioning a Device , on page 86.

Procedure

Step 1 Navigate to the fabric containing the pre-provisioned device from the **Fabric Builder** window.

Step 2 Right click the pre-provisioned device and select **Manage Interfaces**.

You can also navigate to the Interfaces window by selecting **Control > Fabrics > Interfaces**. From the Scope drop-down list, select the fabric containing the pre-provisioned device.

Step 3 Click **Add**.

Step 4 Enter all the required details in the **Add Interface** window.

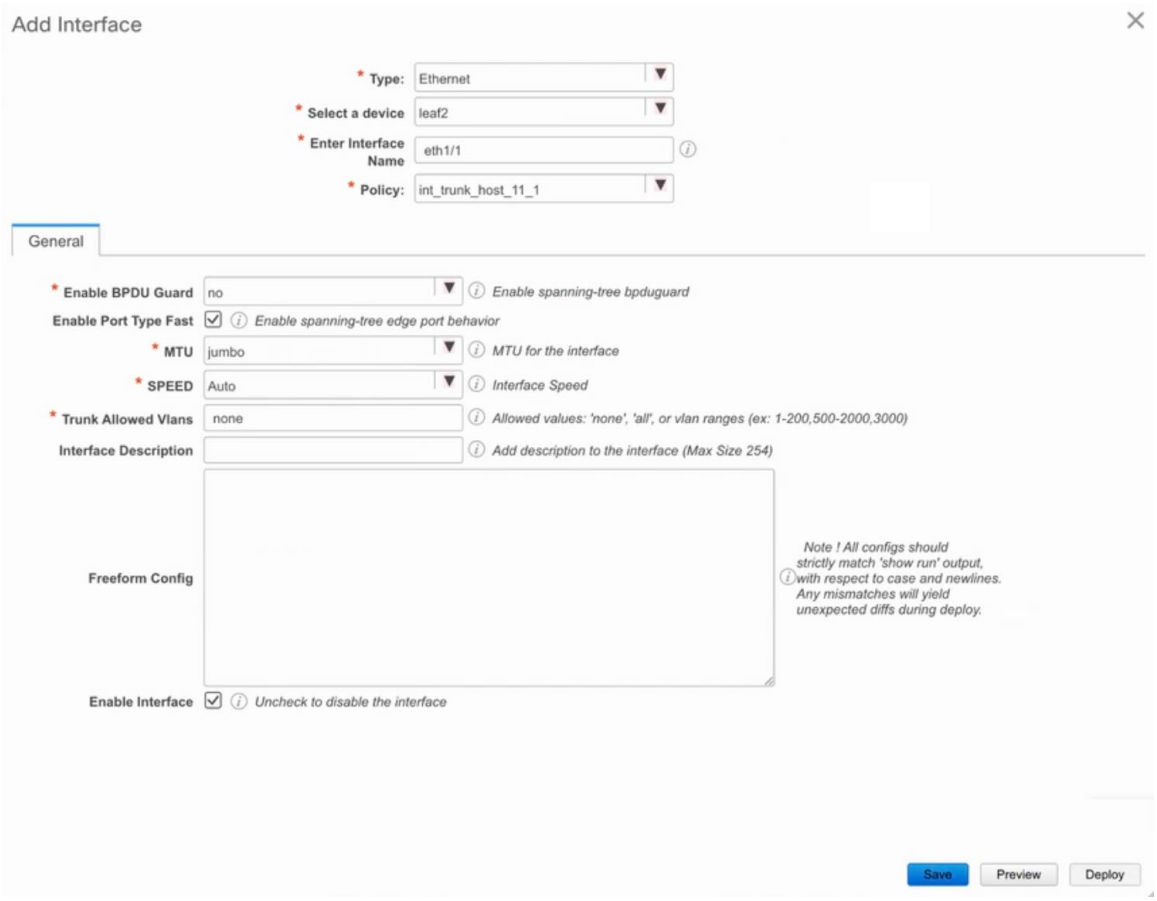

Type: Select **Ethernet** from this drop-down list.

Select a device: Select the pre-provisioned device.

Note You cannot add an Ethernet interface to an already managed device in DCNM.

Enter Interface Name: Enter a valid interface name based on the module type. For example, Ethernet1/1, eth1/1, or e1/1. The interface with same name should be available on the device after it is added.

Policy: Select a policy that should be applied on the interface.

For more information, see Adding Interfaces, on page 263.

Step 5 Click **Save**.

Step 6 Click **Preview** to check the expected configuration that will be deployed to the switch after it is added.

Note The **Deploy** button is disabled for Ethernet interfaces since the devices are pre-provisioned.

Creating a vPC Setup

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**.

Note Alternatively, you can also navigate to the **Tabular view** from the **Actions** pane. Choose a switch in the **Switches** tab and click **vPC Pairing** to create, edit, or unpair a vPC pair. However, you can use this option only when you choose a Cisco Nexus switch.

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

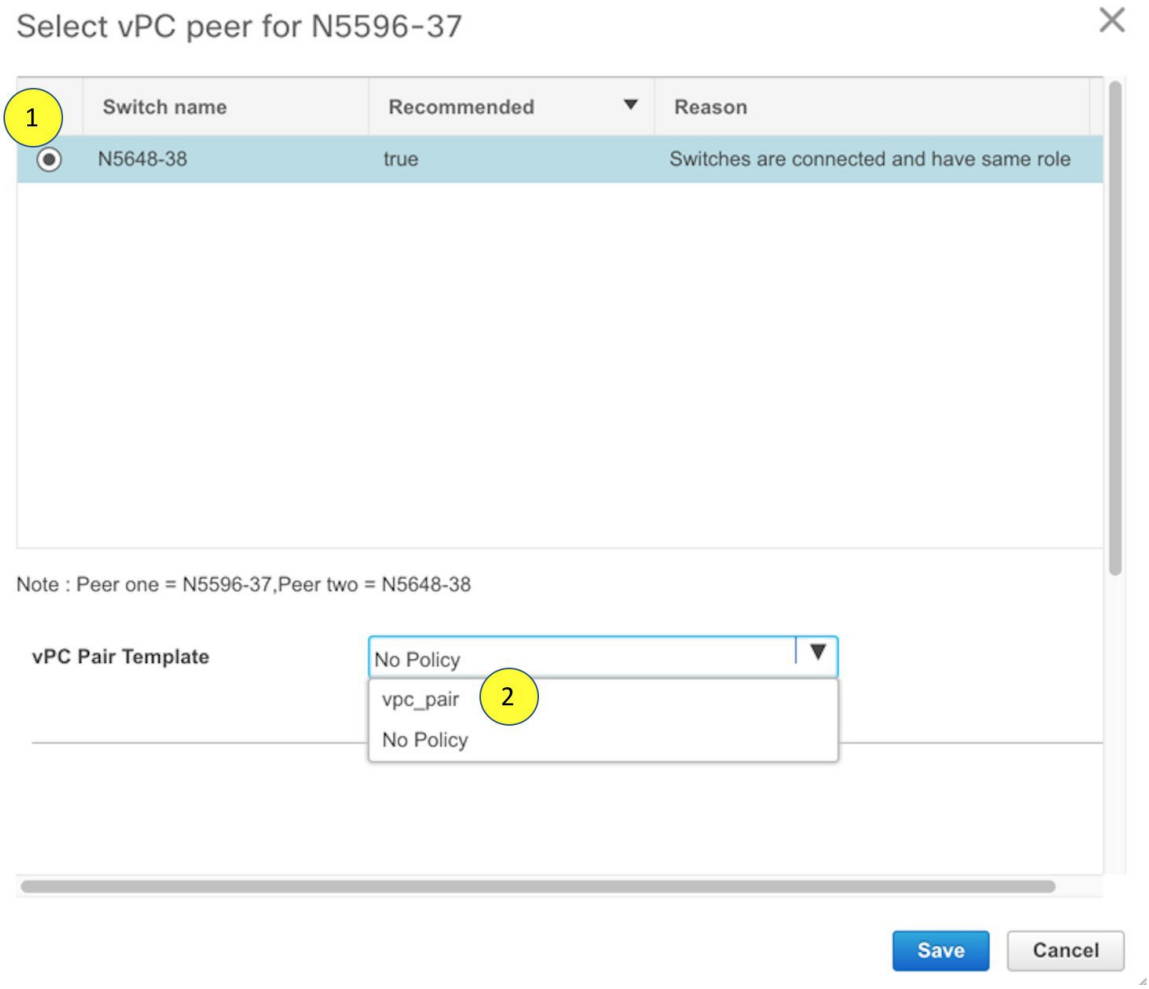

The **vPC Domain** and **vPC Peerlink** tabs appear. You must fill up the fields in the tabs to create the vPC setup. The description for each field is displayed at the extreme right.

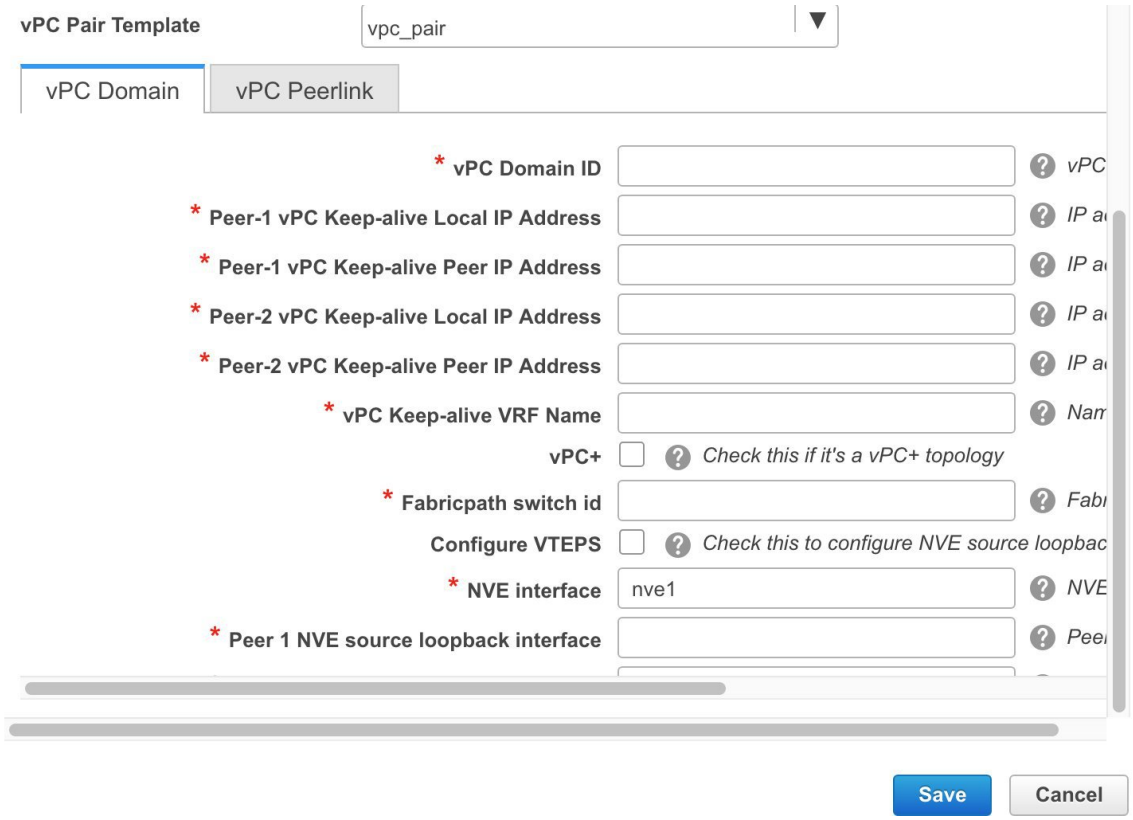

vPC Domain tab: Enter the vPC domain details.

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

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

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

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

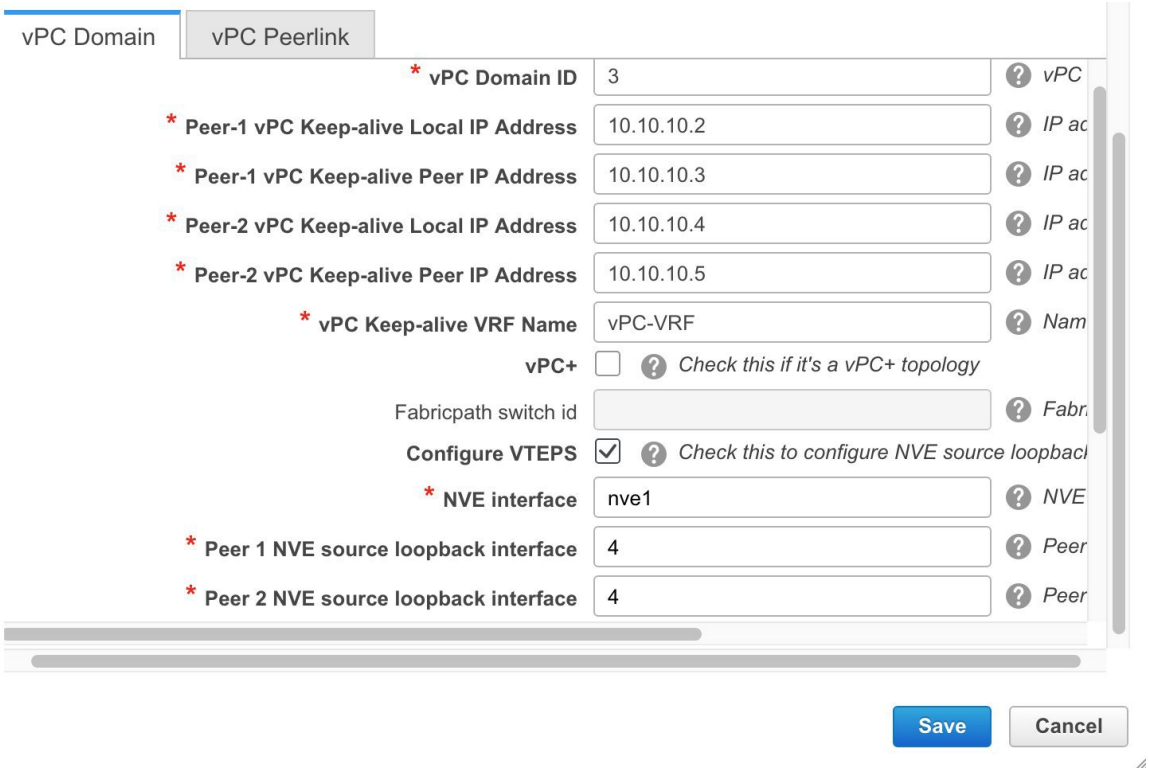

vPC Peerlink tab: Enter the vPC peer-link details.

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

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

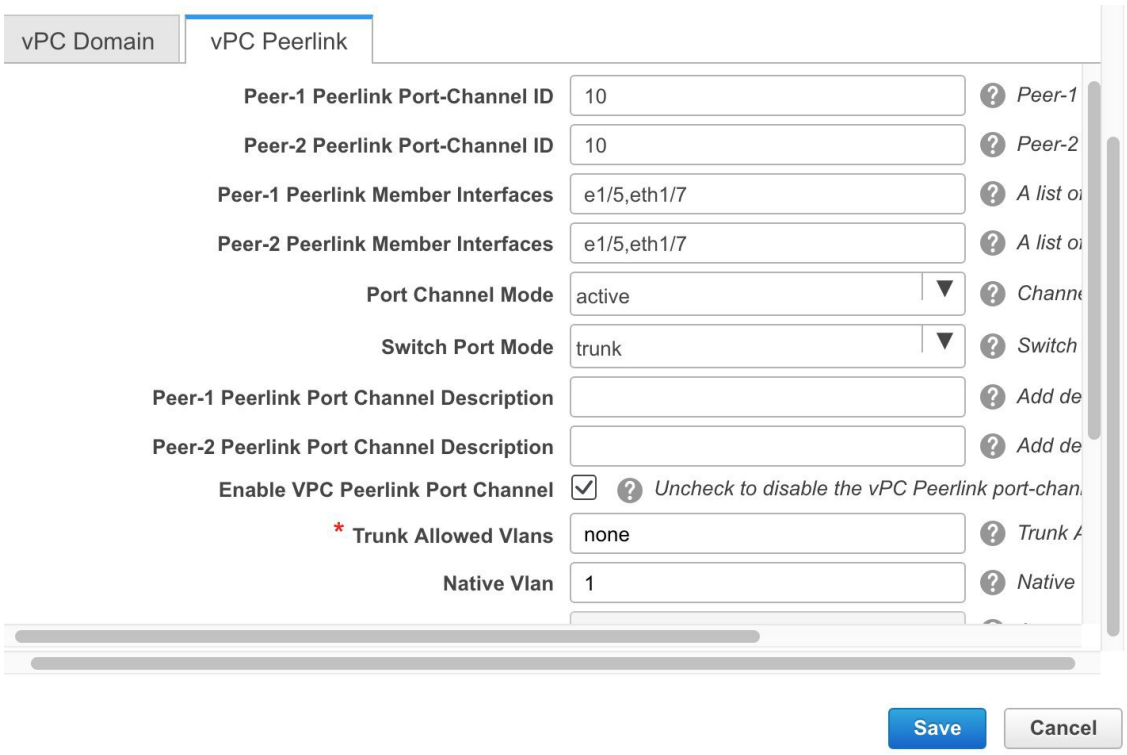

Step 3 Click **Save**.

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

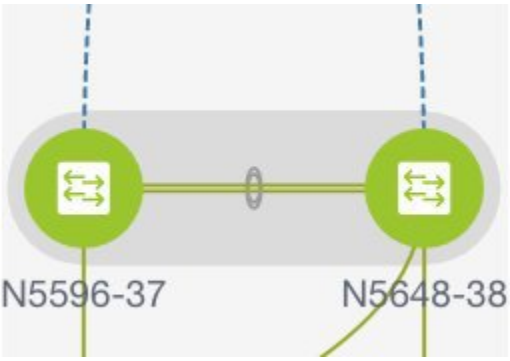

To update vPC setup details, do the following:

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

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

b. Update the field(s) as needed.

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

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

Undeploying a vPC Setup

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.
- After you unpair a BGW vPC, perform a **Save & Deploy** on the member fabric followed by a **Save & Deploy** of the MSD fabric.

A few fabric-specific terms:

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

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

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

Fabric and Switch Instance Variables

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

Fabric instance values can only be edited or updated in the fabric context from the VRFs and Networks window. The appropriate fabric should be selected in the **SCOPE** drop-down list to edit the fabric instance values. Some of the examples of fabric instance variables are BGP ASN, Multicast group per network or VRF, etc. For information about editing multicast group address, see Editing Networks in the Member Fabric, on page 168.

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

MSD and Member Fabric Process Flow

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

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

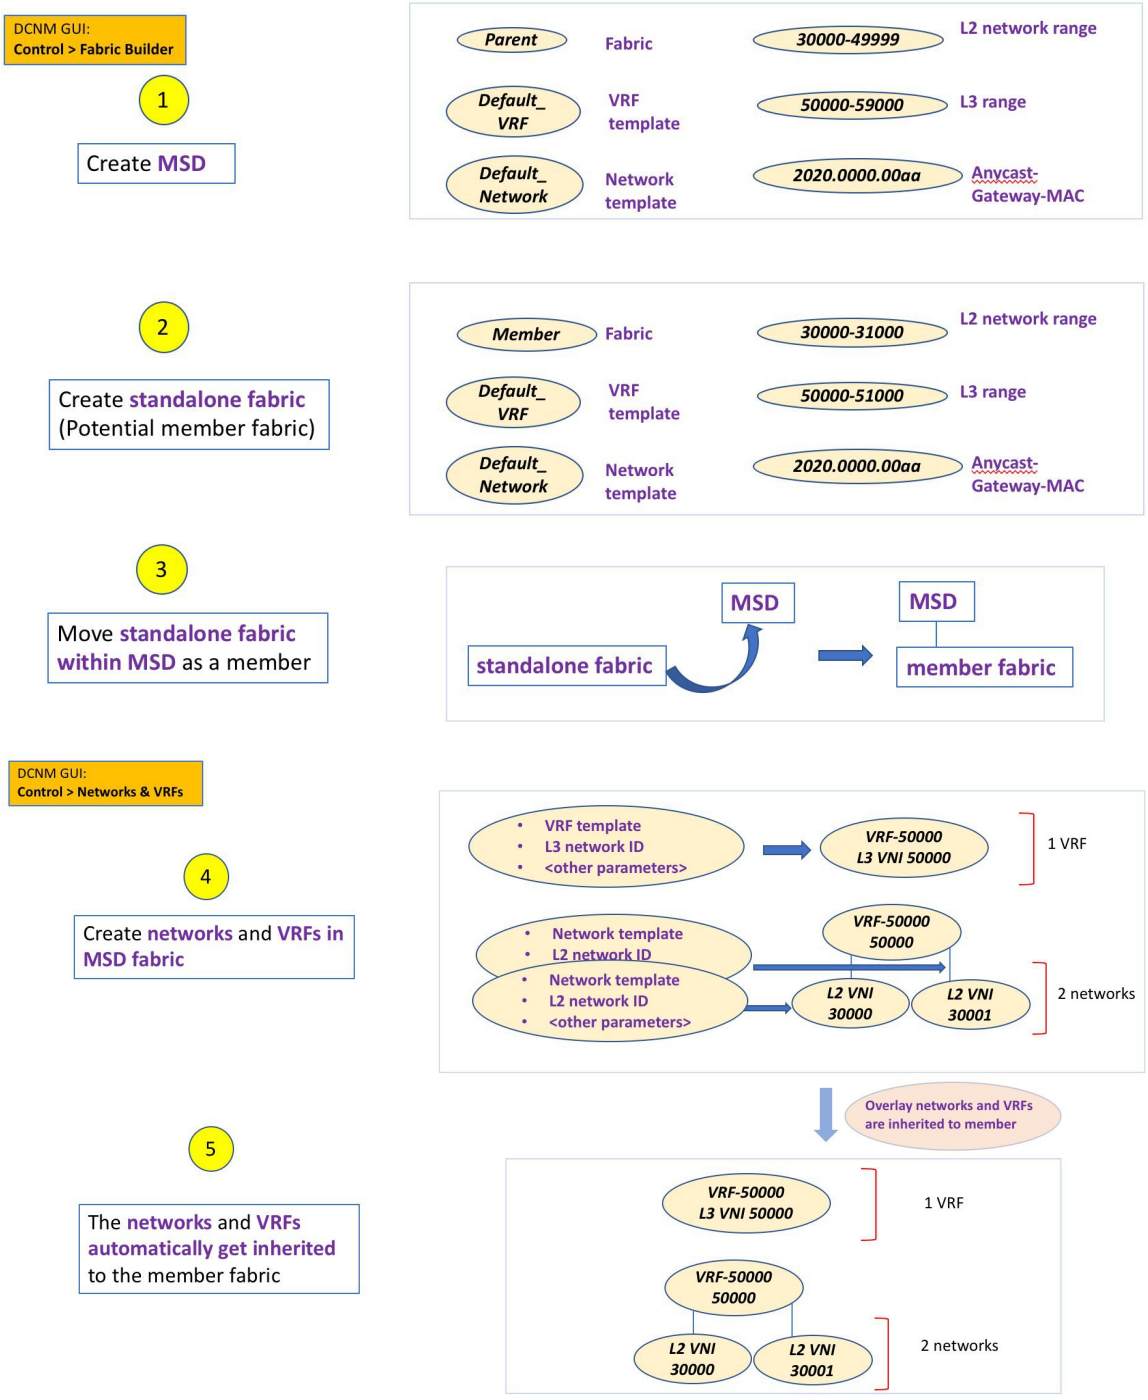

The sample flow explained the inheritance from the MSD to one member. An MSD has multiple sites (and hence, multiple member fabrics under an MSD). A sample flow from an MSD to multiple members:

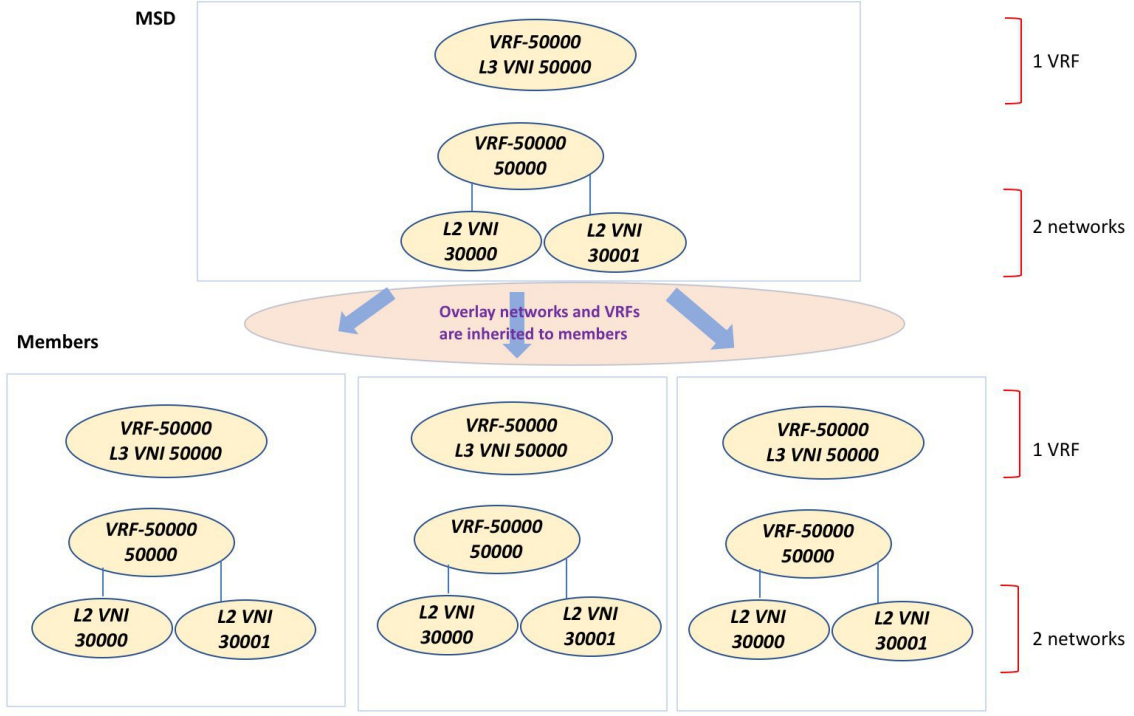

In this example, VRF-50000 (and L3 network with ID 50000), and L2 networks with IDs 30000 and 30001 are created in one go. Networks and VRFs are deployed on the member fabric switches, one after another, as depicted in the image.

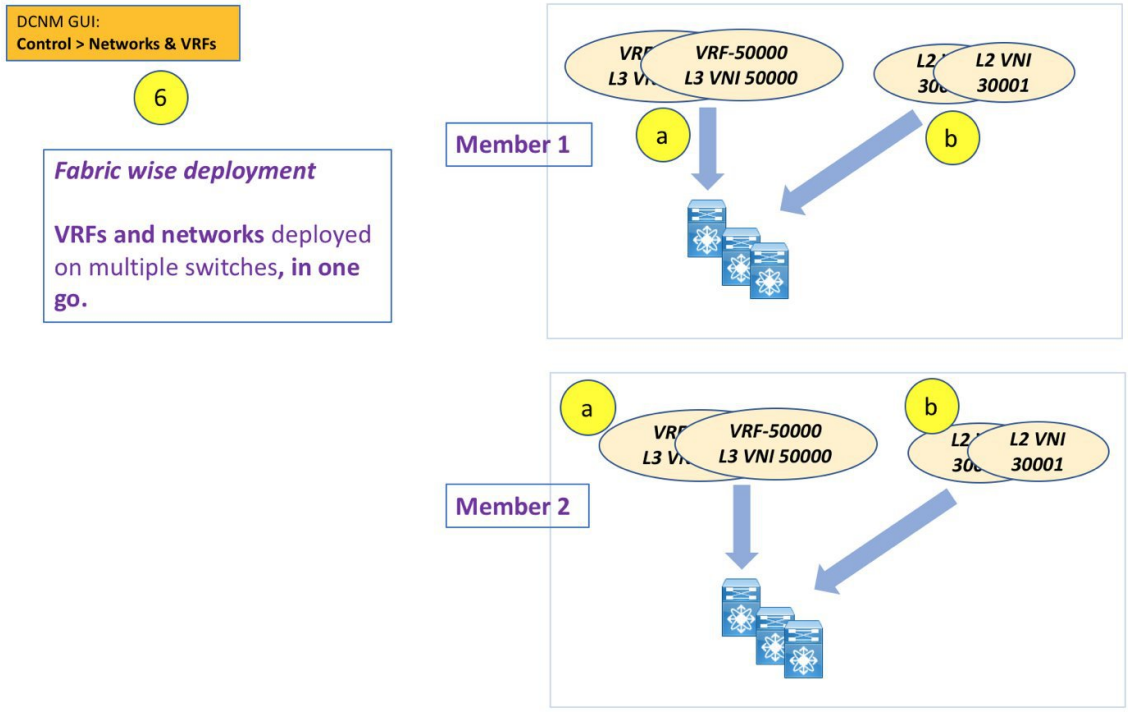

In DCNM 11.1(1), you can provision overlay networks through a single MSD deployment screen.

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

Upcoming sections in the document explain the following:

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

Creating an MSD Fabric and Associating Member Fabrics to It

The process is explained in two steps:

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

Creating an MSD Fabric

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

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

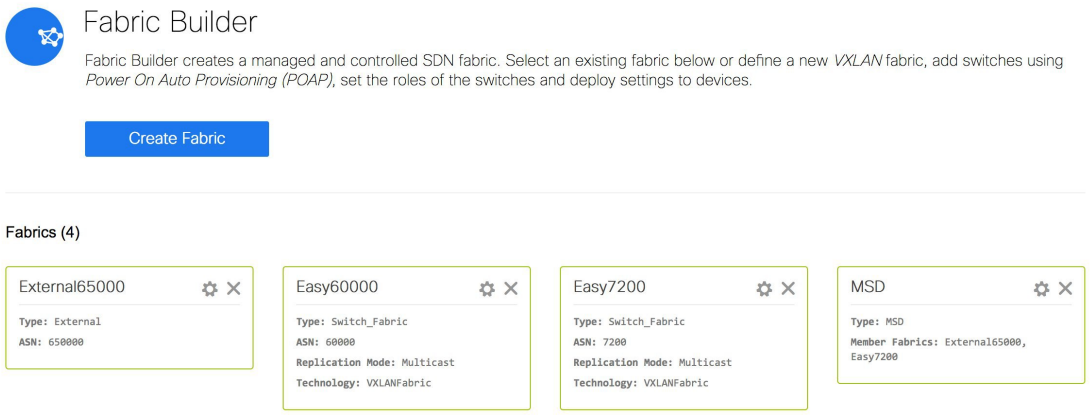

   A standalone or member fabric contains *Switch_Fabric* in the **Type** field, its AS number in the **ASN** field and mode of replication, *Multicast* or *Ingress Replication*, in the **Replication Mode** field. Since no device or network traffic is associated with an MSD fabric as it is a container,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

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

   Fabric Name - Enter the name of the fabric.

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

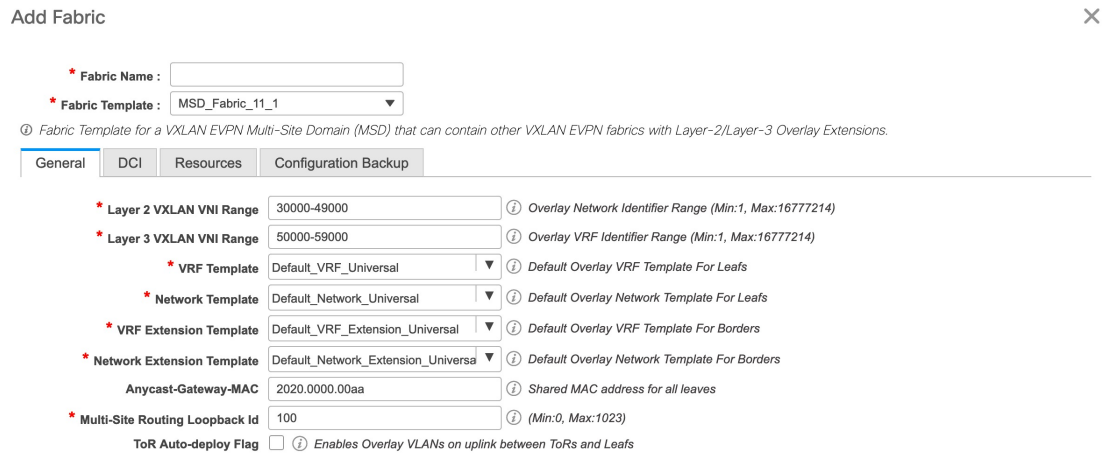

   The fields in the screen are explained:

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

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

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

VRF Template - Default VRF template.

Network Template - Default network template.

VRF Extension Template - Default VRF extension template.

Network Extension Template - Default network extension template.

Anycast-Gateway-MAC - Anycast gateway MAC address.

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

ToR Auto-deploy Flag - Select this check box to enable automatic deployment of the networks and VRFs in the Easy Fabric to the ToR switches in the External Fabric when you click **Save & Deploy** in the MSD fabric.

3. Click the **DCI** tab.

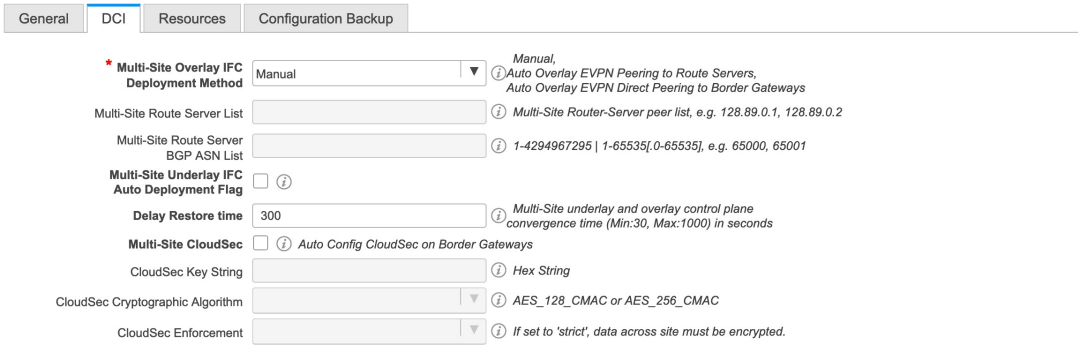

The fields are:

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.

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

Multi-Site Route Server BGP ASN List – 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 Auto Deployment Flag - Check the check box to enable auto configuration. Uncheck the check box for manual configuration.

Delay Restore Time - Specifies the Multi-Site underlay and overlay control planes convergence time. The minimum value is 30 seconds and the maximum value is 1000 seconds.

Multi-Site CloudSec – Enables CloudSec configurations on border gateways. If you enable this field, the remaining three fields for CloudSec are editable. For more information, see Support for CloudSec in Multi-Site Deployment, on page 174.

Enable Multi-Site eBGP Password - Enables eBGP password for Multi-Site underlay/overlay IFCs.

eBGP Password - Specifies the encrypted eBGP Password Hex String.

eBGP Authentication Key Encryption Type - Specifies the BGP key encryption type. It is **3** for 3DES and **7** for Cisco.

4. Click the **Resources** tab.

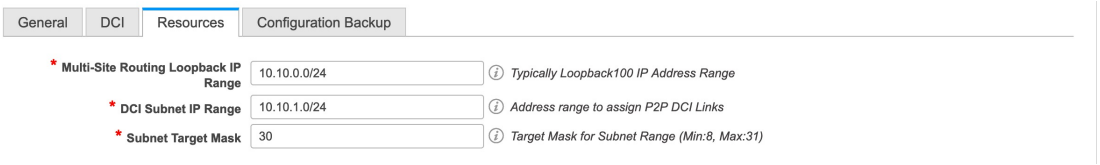

MultiSite Routing Loopback IP Range – Specify the Multi-Site loopback IP address range used for the EVPN Multi-Site function.

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

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

5. Click the **Configuration Backup** tab.

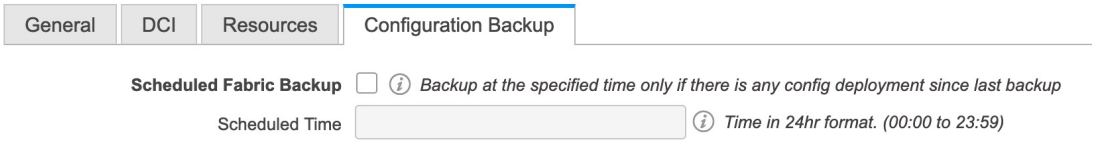

Scheduled Fabric Backup: Check the check box to enable a daily backup. This backup tracks changes in running configurations on the fabric devices that are not tracked by configuration compliance.

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.

The backup process is initiated after you click Save.

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

6. 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.

Note From Cisco DCNM Release 11.4(1), when you update the MSD fabric settings, only switches with rol relevant to MSD are updated.

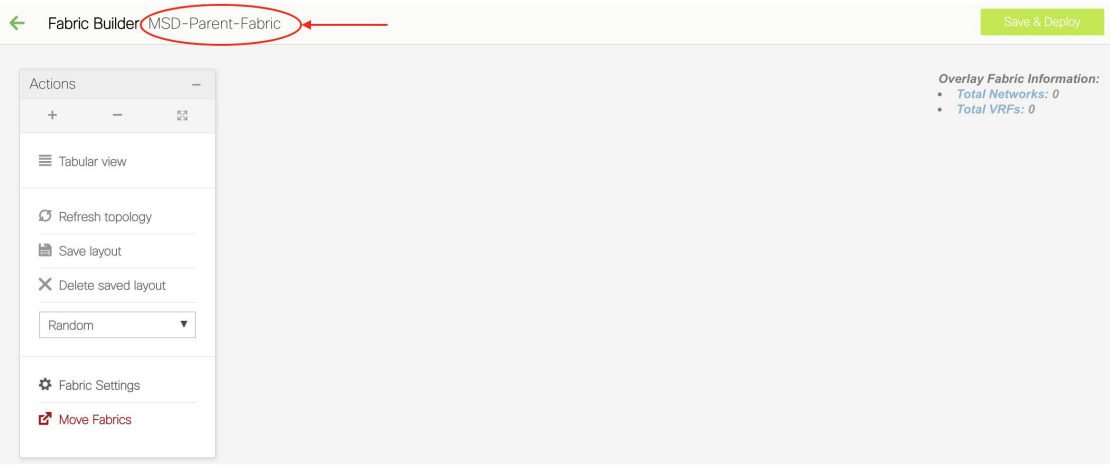

Since the MSD fabric is a container, you cannot add a switch to it. The **Add Switches** button that is available in the **Actions** panel for member and standalone fabrics is not available for the MSD fabric.

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

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

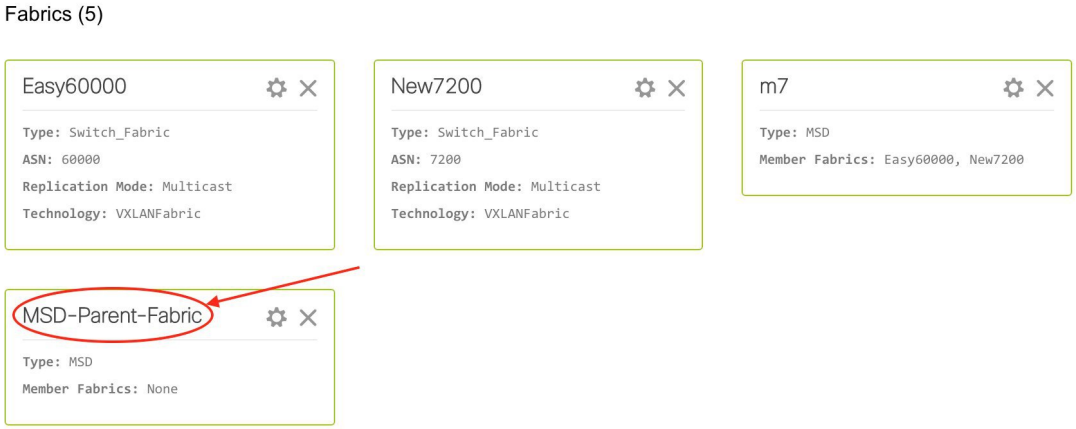

The steps for creation of an MSD fabric and moving member fabrics under it are:

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

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

Creating and Moving a New Fabric Under the MSD Fabric as a Member

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

New fabric creation is explained in the Easy Fabric creation process. In the MSD document, fabric movement is covered. However, some pointers about a standalone (potential member) fabric:

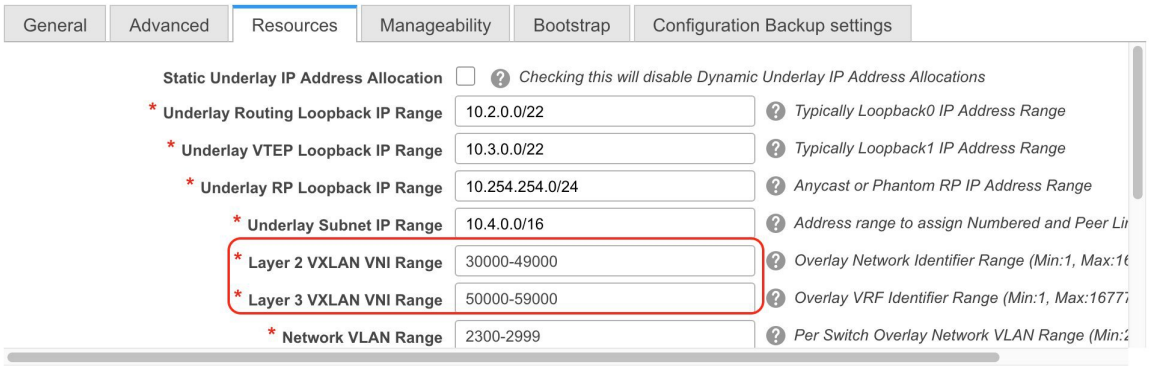

The values that are displayed in the screen are automatically generated. The VXLAN VNI ID ranges (in the L2 Segment ID Range and L3 Partition ID Range fields) allocated for new network and VRF creation are values from the MSD fabric segment ID range. If you want to update the VXLAN VNI ranges or the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

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

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

Other pointers:

- Ensure that the Anycast Gateway MAC, the Network Template and the VRF Template field values are the same as the MSD fabric. Else, member fabric movement to the MSD fail.
- The member fabric should have a Site ID configured and the Site ID must be unique among the members.
- The BGP AS number should be unique for a member fabric.
- The underlay subnet range for loopback0 should be unique.
- The underlay subnet range for loopback1 should be unique.

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

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

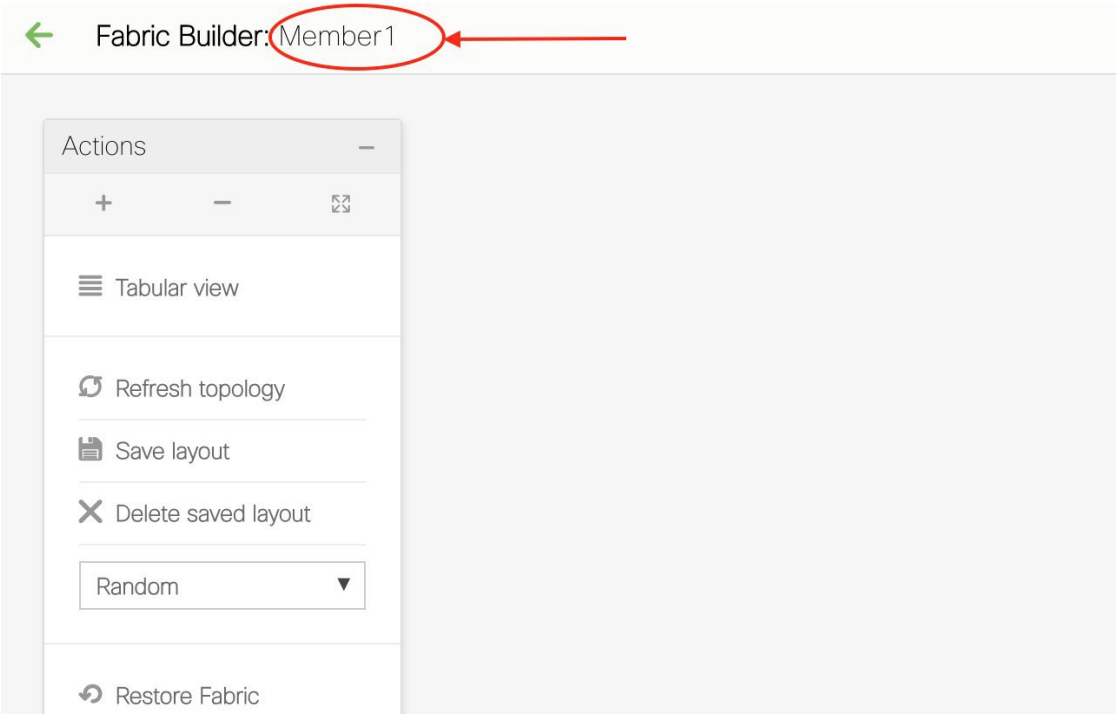

Simultaneously, the Fabric Builder page also displays the newly created fabric, Member1.

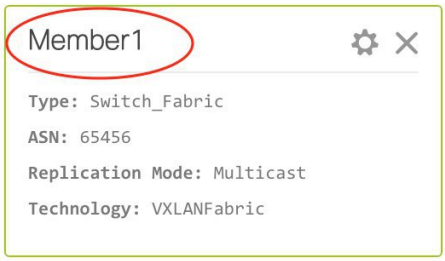

Moving the Member1 Fabric Under MSD-Parent-Fabric

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

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

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

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

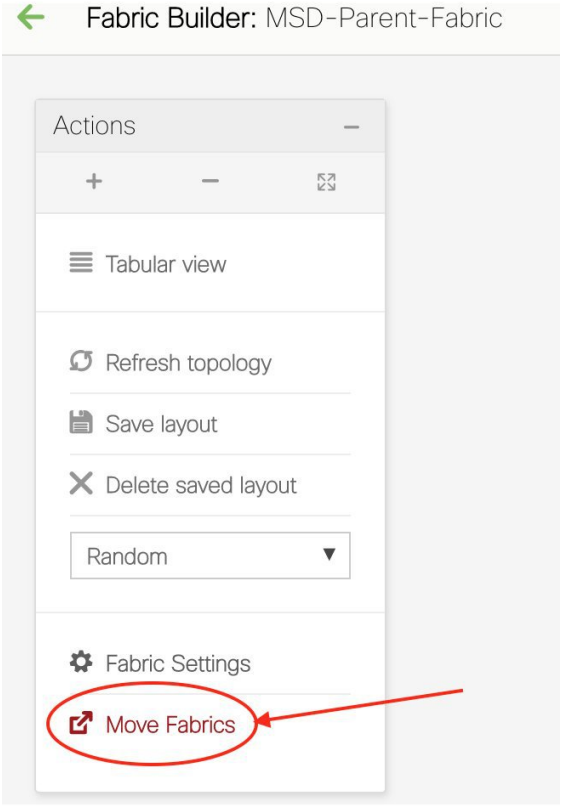

The Move Fabric screen comes up. It contains a list of fabrics.

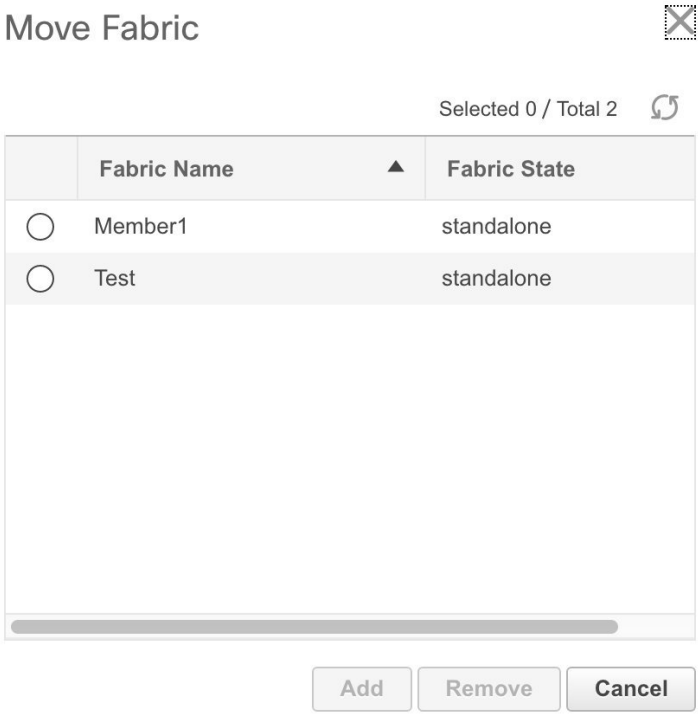

Member fabrics of other MSD container fabrics are not displayed here.

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

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

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

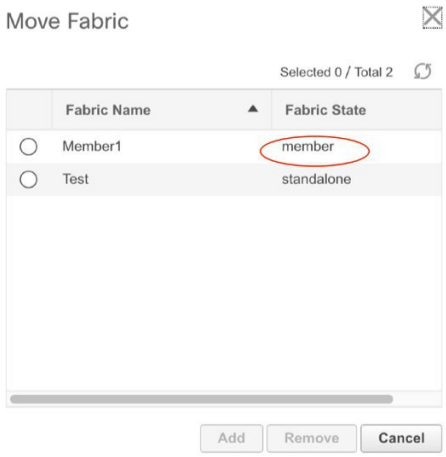

5. Close this screen.
6. Click ← above the Actions panel to go to the Fabric Builder page.

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

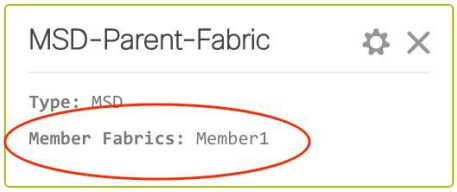

MSD Fabric Topology View Pointers

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

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

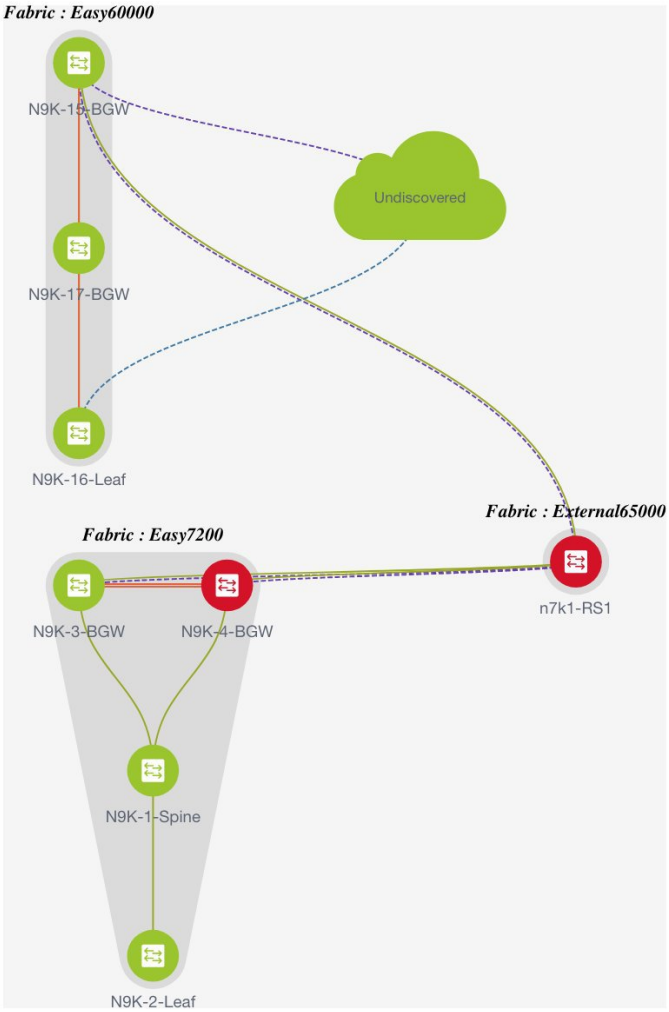

- **Member fabric topology view** - A member fabric and its switches are displayed. In addition, the connected external fabric is displayed.

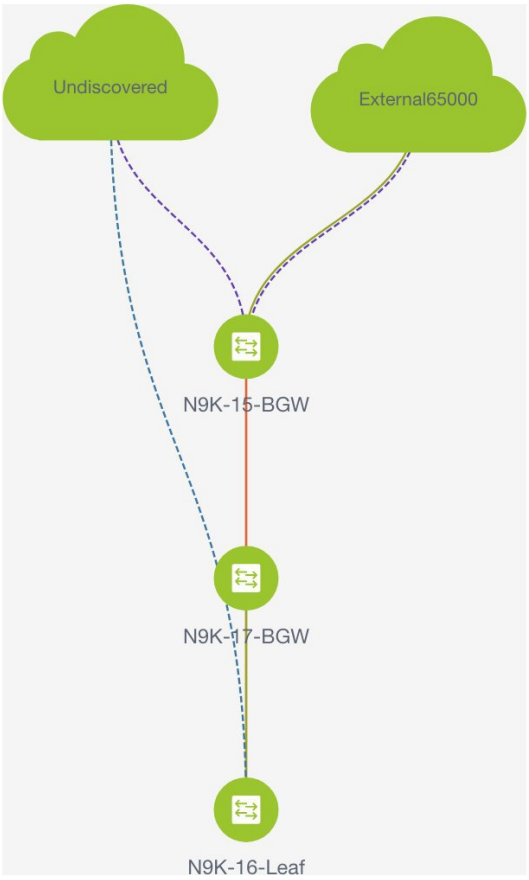

- A boundary defines a standalone VXLAN fabric, and each member fabric in an MSD fabric. A fabric's devices are confined to the fabric boundary. You can move a switch icon by dragging it. For a better user experience, in addition to switches, DNCM 11.2(1) release allows you to move an entire fabric.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

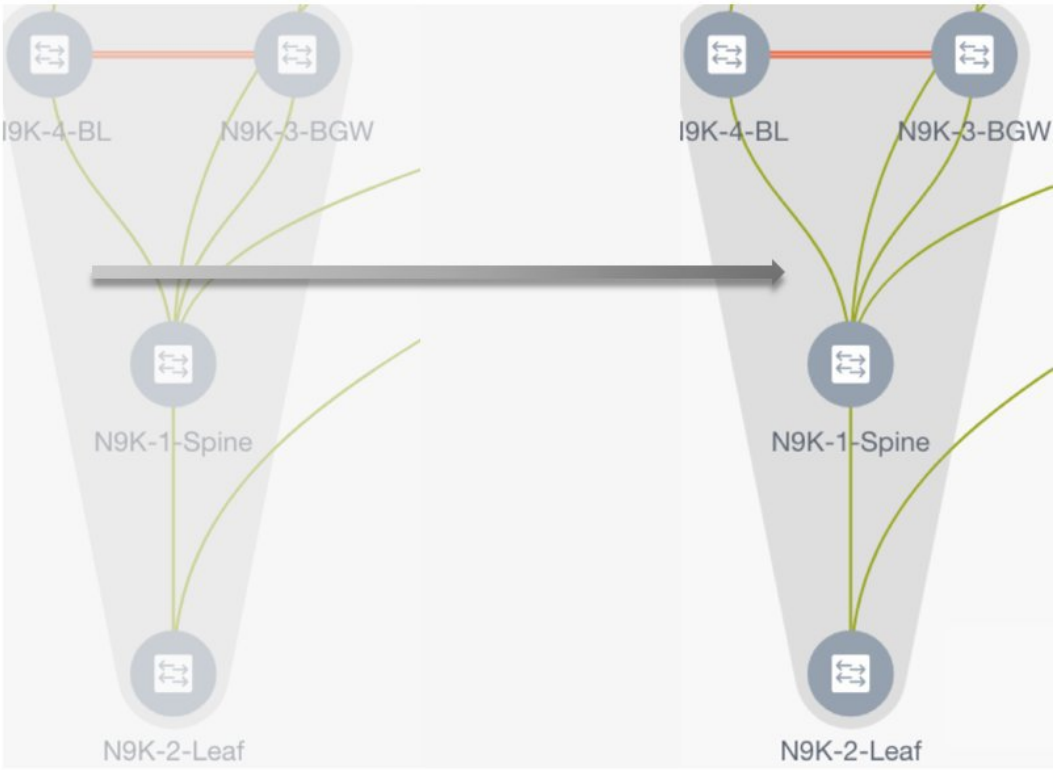

Adding and Editing Links

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

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

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

Creating and Deploying Networks and VRFs in an MSD Fabric

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

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

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

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

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

For example, if you create networks 30000 and 30001 for an MSD that contains two member fabrics, the networks are automatically created for the member fabrics and are available for deployment.

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

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

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

Creating Networks in the MSD Fabric

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

   The Networks screen comes up.

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

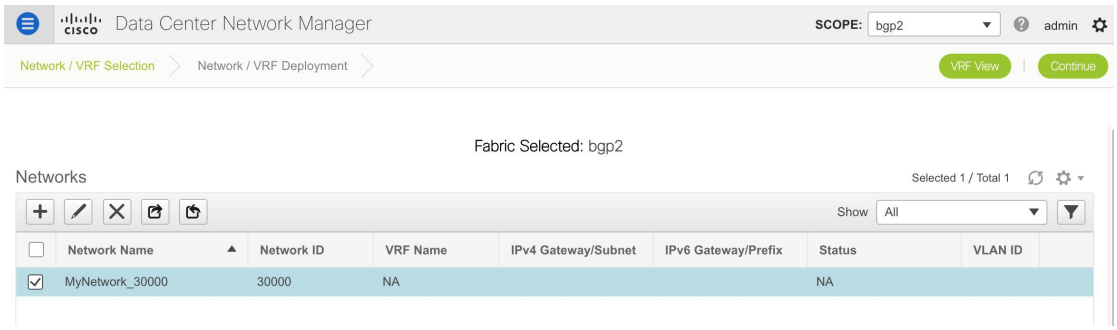

3. Select *MSD-Parent-Fabric* from the list and click **Continue** at the top right part of the screen.

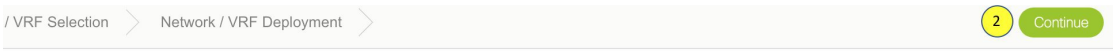

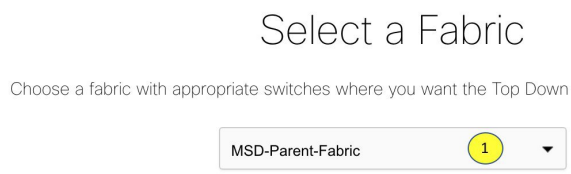

The Networks page comes up. This lists the list of networks created for the MSD fabric. Initially, this screen has no entries.

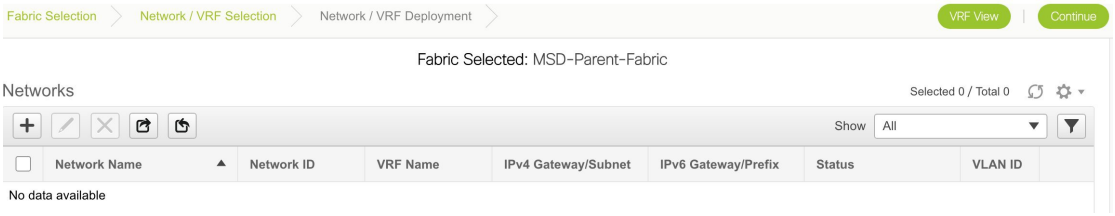

4. Click the + button at the top left part of the screen (under **Networks**) to add networks to the MSD fabric. The Create Network screen comes up. Most of the fields are autopopulated.

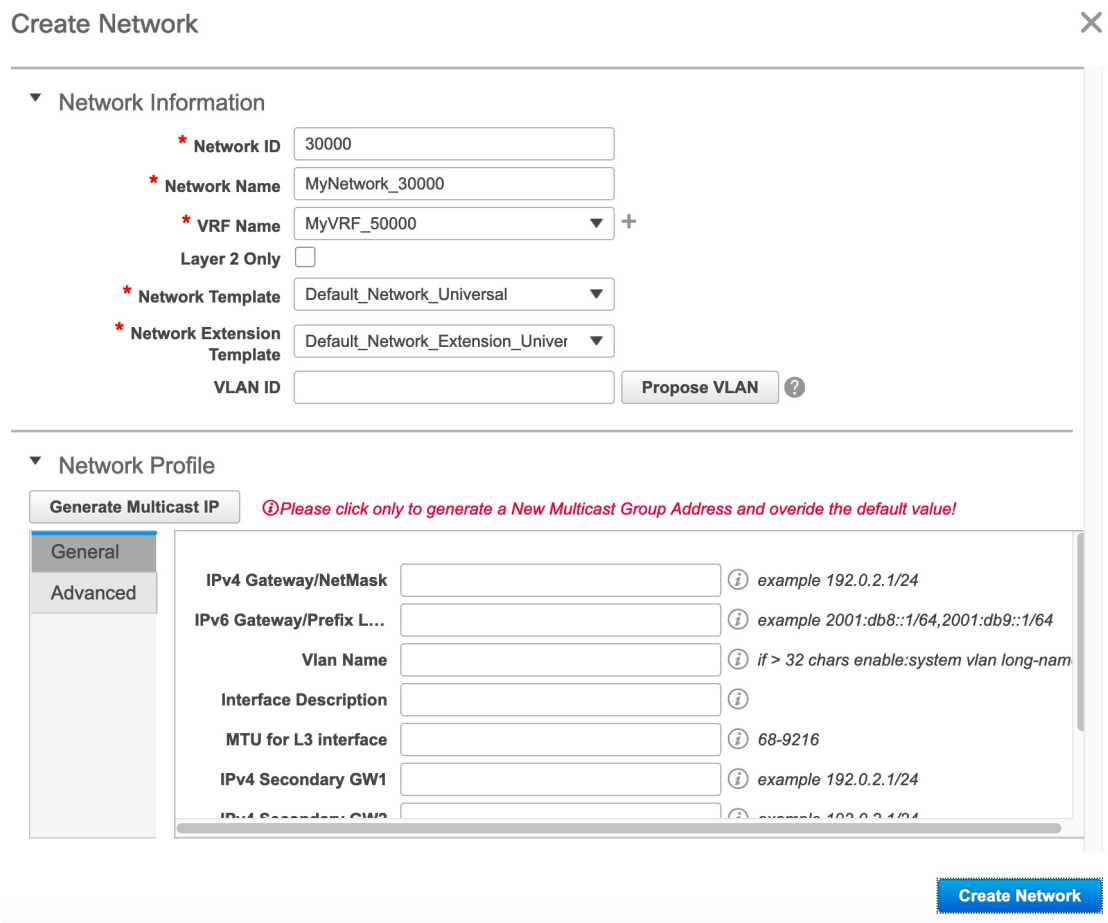

The fields in this screen are:

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

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

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

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

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

Network Template - Allows you to select a network template.

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

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

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

General tab

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

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

VLAN Name - Enter the VLAN name.

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

Interface Description - Specifies the description for the interface.

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

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

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

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

- ARP Suppression
- DHCPv4 Server 1 and DHCPv4 Server 2 - Enter the DHCP relay IP address of the first and second DHCP servers.
- DHCPv4 Server VRF - Enter the DHCP server VRF ID.
- Loopback ID for DHCP Relay interface - Enter the loopback ID of the DHCP relay interface.
- Routing Tag – The routing tag is autopopulated. This tag is associated with each gateway IP address prefix.
- TRM enable – Select the check box to enable TRM.

  For more information, see Overview of Tenant Routed Multicast, on page 196.
- L2 VNI Route-Target Both Enable - Select the check box to enable automatic importing and exporting of route targets for all L2 virtual networks.

Note From Cisco DCNM Release 11.5(1), the **Enable L3 Gateway on Border** field is not available as part of the MSD network settings. You can enable a Layer 3 gateway on the border switches at a fabric level. For more information, see Creating Networks for the Standalone Fabric, on page 277.

In the MSD fabric level, if the **Enable L3 Gateway on Border** check box is selected and you are upgrading to Cisco DCNM Release 11.5(1), then it is automatically removed from the MSD fabric level during upgrade.

- A sample of the Create Network screen:

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

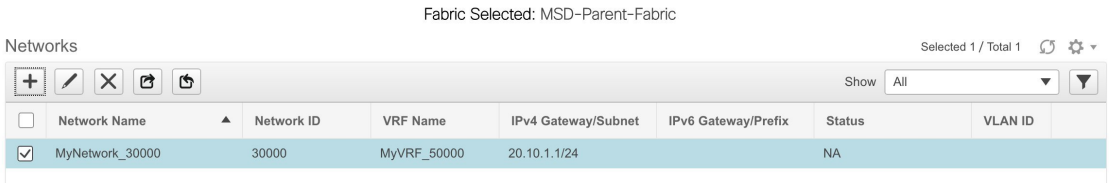

Editing Networks in the MSD Fabric

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

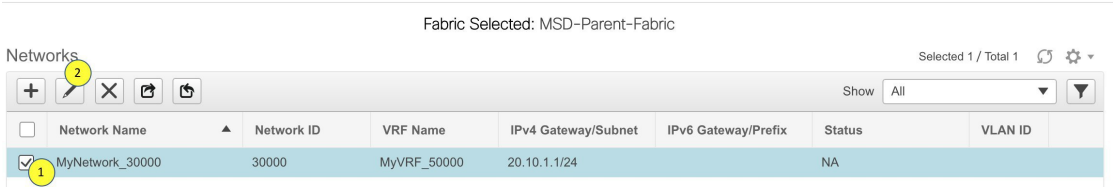

The Edit Network screen comes up.

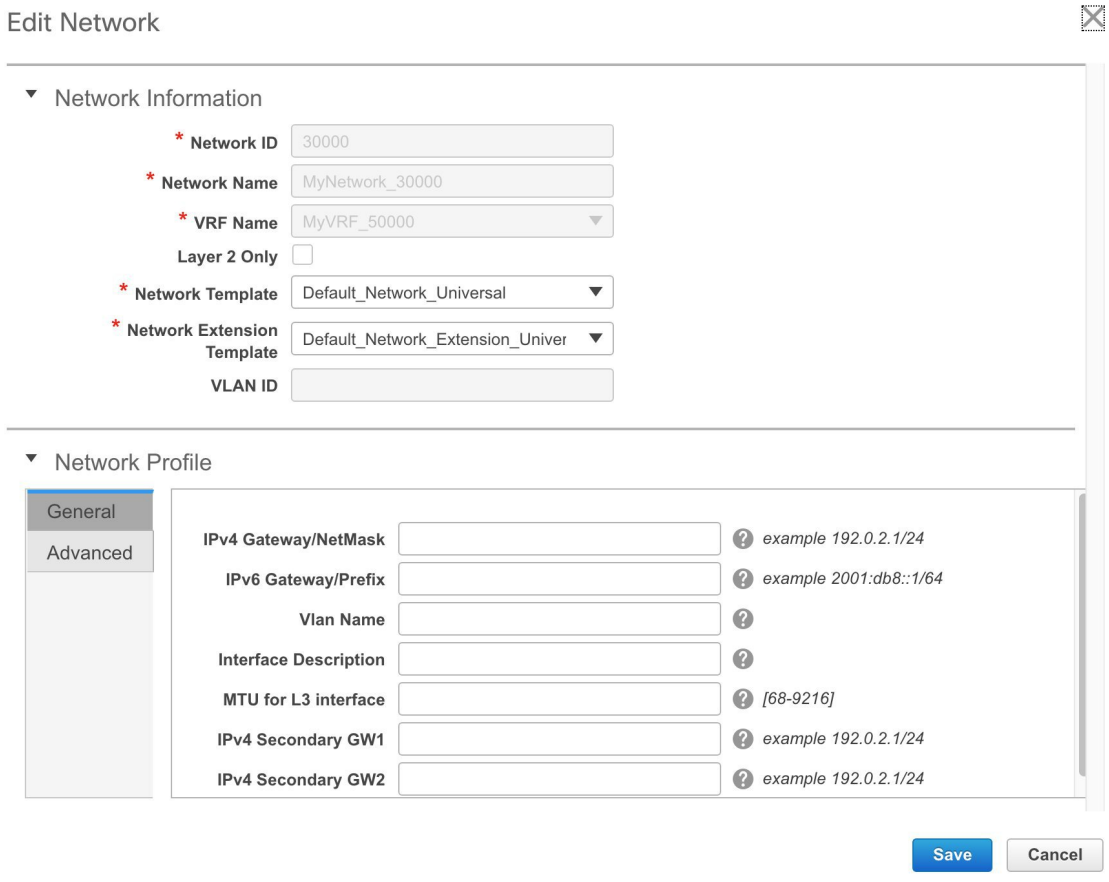

You can edit the **Network Profile** part (**General** and **Advanced** tabs) of the MSD fabric network.

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

Network Inheritance from MSD-Parent-Fabric to Member1

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

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

   The Networks screen comes up.

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

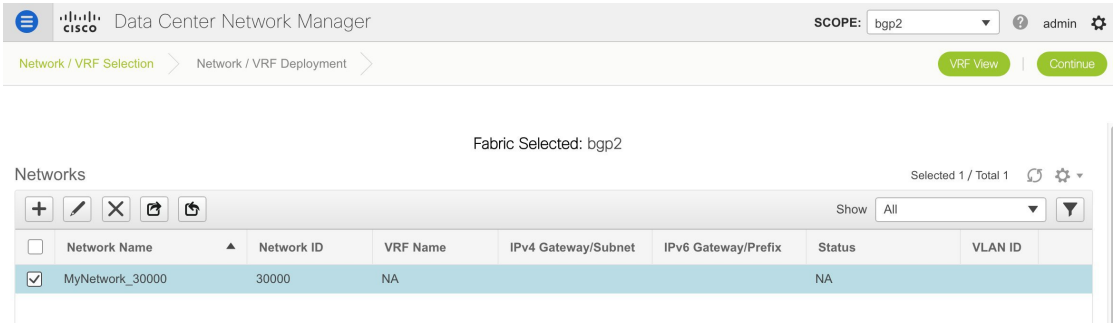

Editing Networks in the Member Fabric

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

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

1. Select the network and click the **Edit** option at the top left part of the window. The **Edit Network** window comes up.

2. Update the multicast group address in one of the following ways:

   - Under **Network Profile**, click the **Generate Multicast IP** button to generate a new multicast group address for the selected network, and click **Save**.
   - Click the **Advanced** tab in the **Network Profile** section, update the multicast group address, and click **Save**.

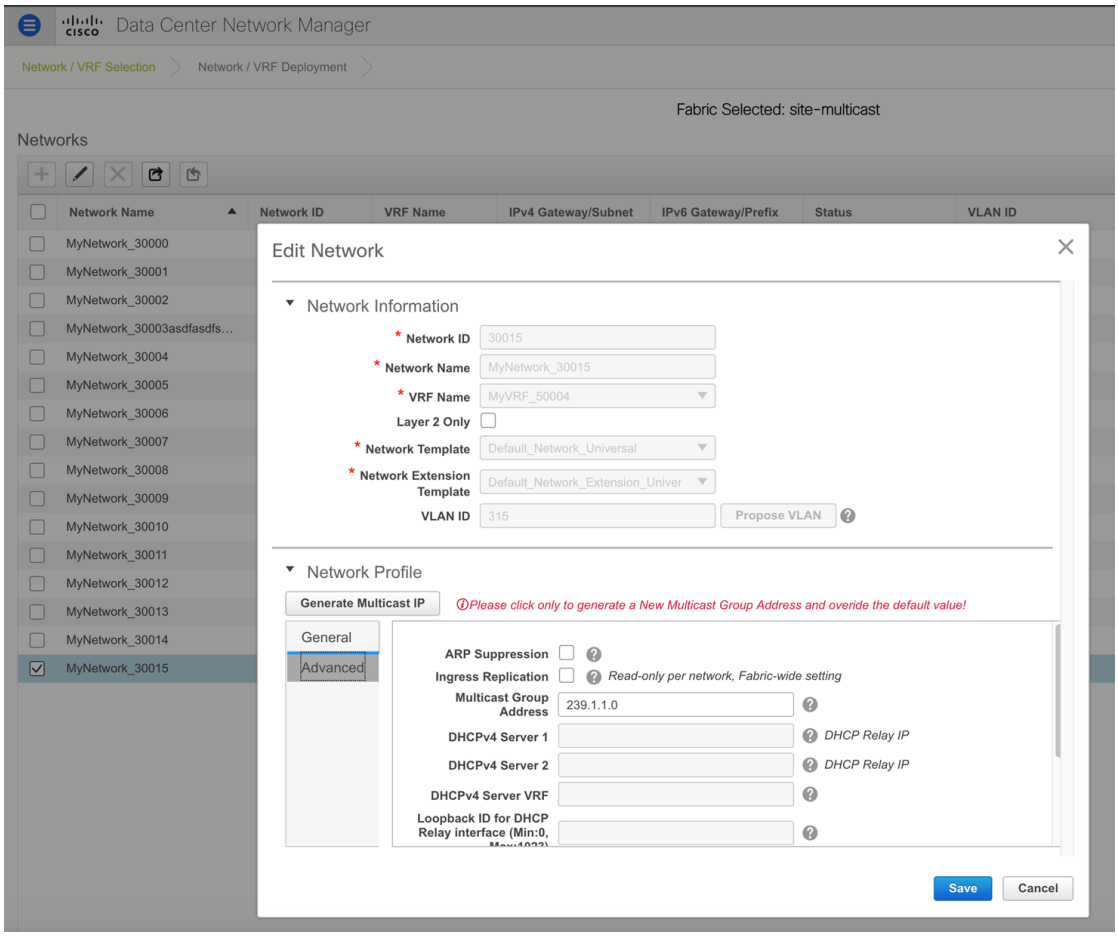

Note The **Generate Multicast IP** option is only available for member fabric networks and not MSD networks.

Deleting Networks in the MSD and Member Fabrics

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

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

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

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

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

Creating VRFs in the MSD Fabric

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

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

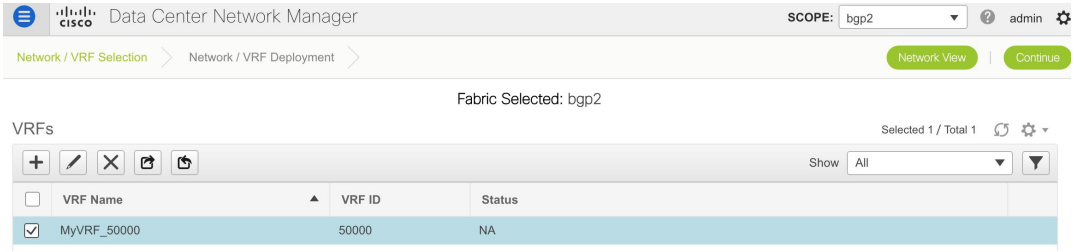

   b. Choose the MSD fabric (*MSD-Parent-Fabric*) from the drop-down box and click **Continue**. The Networks page comes up.

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

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

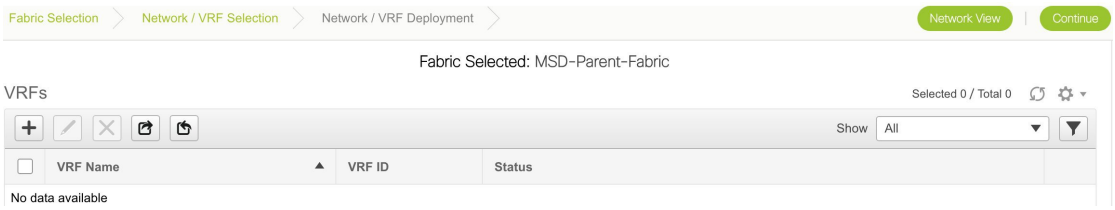

2. Click the + button at the top left part of the screen to add VRFs to the MSD fabric. The Create VRF screen comes up. Most of the fields are autopopulated.

   The fields in this screen are:

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

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

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

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

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

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

4. **Advanced** tab

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

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

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

TRM Enable – Select the check box to enable TRM.

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

For more information, see Overview of Tenant Routed Multicast, on page 196.

Is RP External – Enable this checkbox if the RP is external to the fabric. If this field is unchecked, RP is distributed in every VTEP.

RP Address – Specifies the IP address of the RP.

RP Loopback ID – Specifies the loopback ID of the RP, if **Is RP External** is not enabled.

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 **Default MDT Address for TRM VRFs** field in the fabric settings scree auto-populated in this field. You can override this field if a different multicast group address should be for this VRF.

Overlay Multicast Groups – Specifies the multicast group subnet for the specified RP. The value is the group range in "ip pim rp-address" command. If the field is empty, 224.0.0.0/24 is used as default.

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

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

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

A sample screenshot:

Advanced tab:

5. Click **Create VRF**.

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

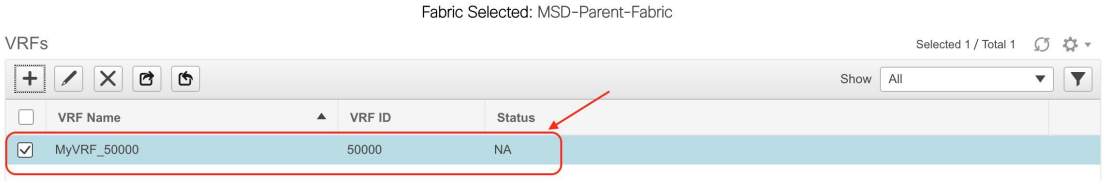

Editing VRFs in the MSD Fabric

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

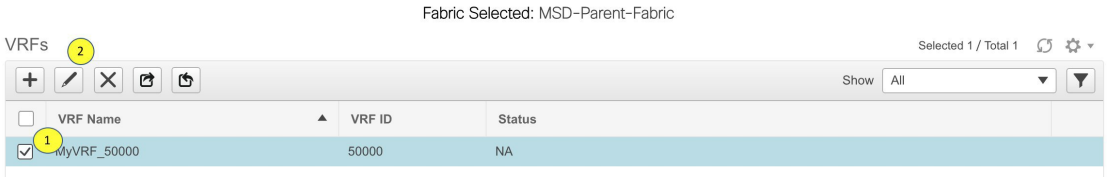

The Edit VRF screen comes up.

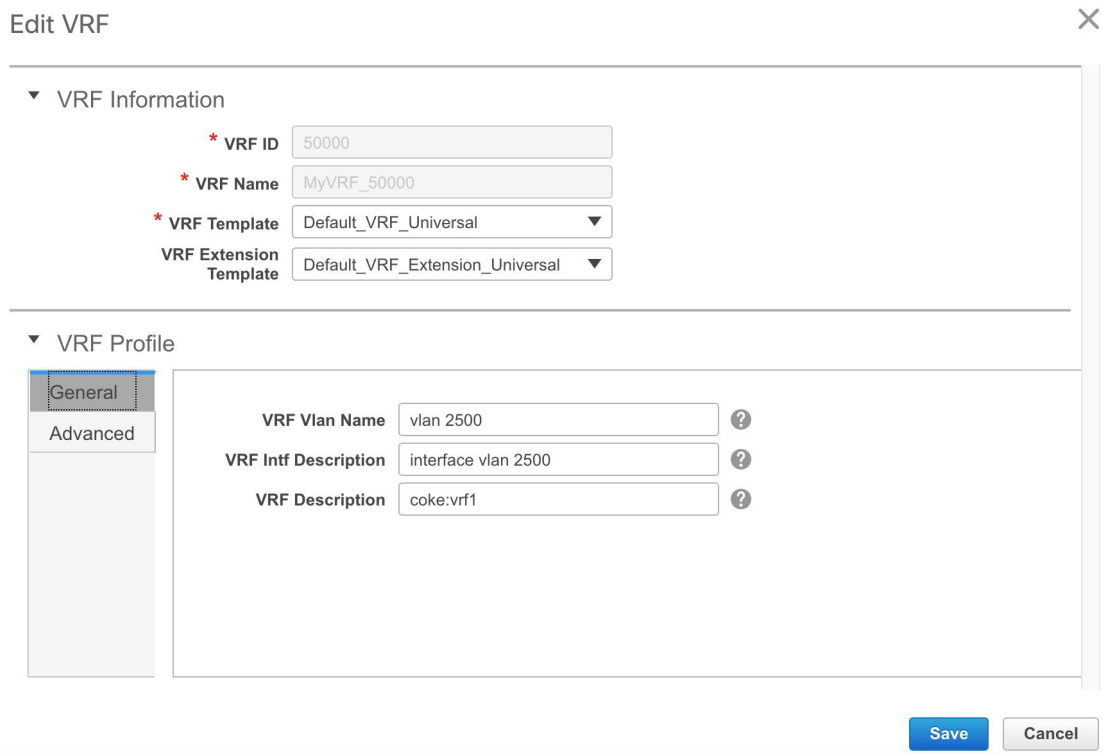

You can edit the **VRF Profile** part (**General** and **Advanced** tabs).

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

VRF Inheritance from MSD-Parent-Fabric to Member1

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

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

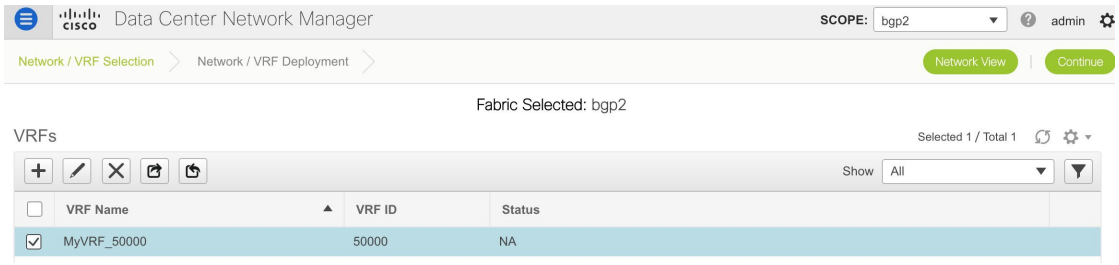

2. Click the **VRF View** button. On the VRFs page, you can see that the VRF created for the MSD is inherited to its member.

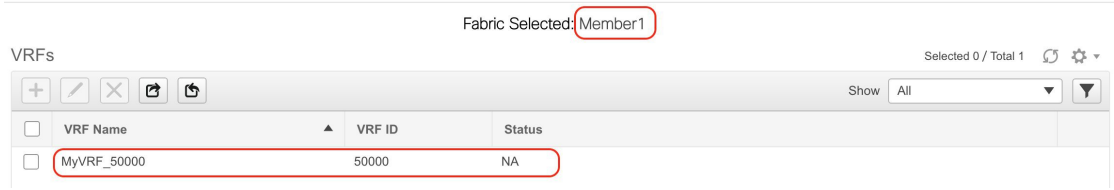

Deleting VRFs in the MSD and Member Fabrics

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

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

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

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 .

Support for CloudSec in Multi-Site Deployment

CloudSec feature allows secured data center interconnect in a multi-site deployment by supporting source-to-destination packet encryption between border gateway devices in different fabrics.

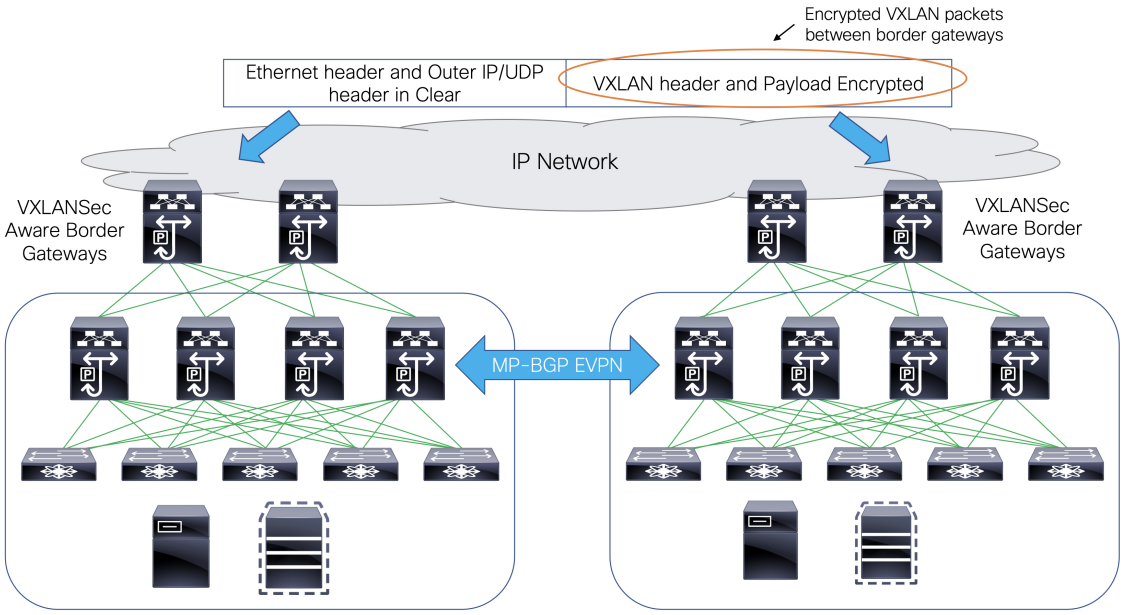

CloudSec feature is supported on Cisco Nexus 9000 Series FX2 platform with Cisco NX-OS Release 9.3(5) or later. The border gateways, border gateway spines, and border gateway superspines that are FX2 platforms, and run Cisco NX-OS Release 9.3(5) or later are referred as CloudSec capable switches.

Cisco DCNM Release 11.4(1) provides an option to enable CloudSec in an MSD fabric.

Note The CloudSec session is point to point over DCI between border gateways (BGWs) on two different sites. All communication between sites uses Multi-Site PIP instead of VIP. Enabling CloudSec requires a switch from VIP to PIP, which could cause traffic disruption for data flowing between sites. Therefore, it is recommended to enable or disable CloudSec during a maintenance window.

You can also watch the video that demonstrates how to configure the CloudSec feature. See Video: Configuring CloudSec in Cisco DCNM.

Enabling CloudSec in MSD

Navigate to **Control > Fabrics > Fabric Builder**. You can either create a new MSD fabric by clicking **Create Fabric** or edit the existing MSD fabric by clicking **Edit Fabric**.

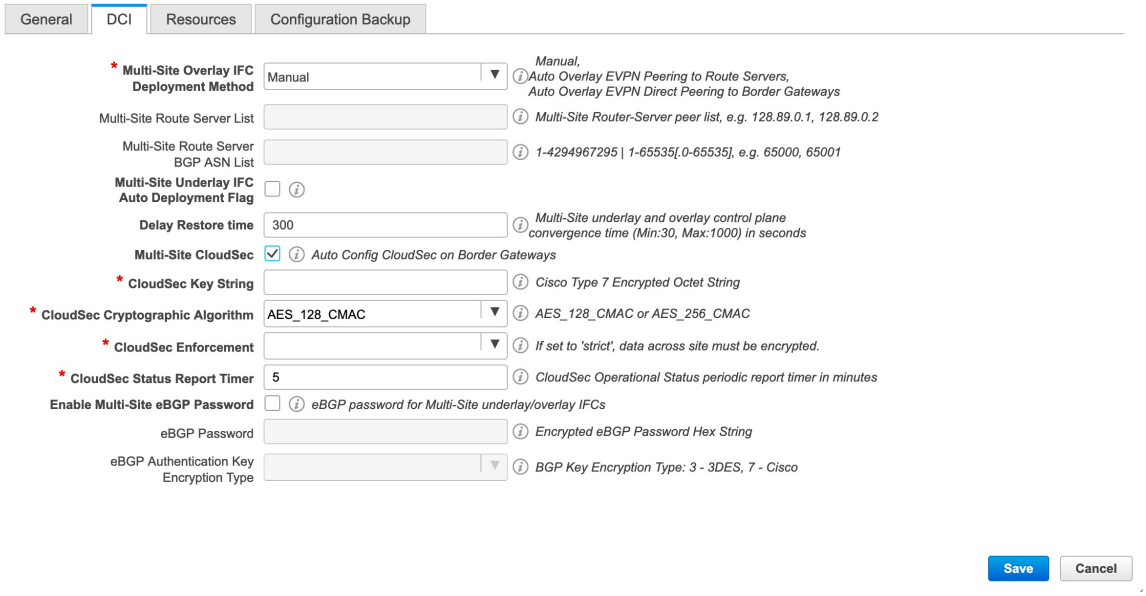

Under the **DCI** tab, you can specify the CloudSec configuration details.

Multi-Site CloudSec – Enables CloudSec configurations on border gateways. If you enable this field, the remaining CloudSec fields are editable.

Multi-Site CloudSec – Enables CloudSec configurations on border gateways. If you enable this field, the remaining three fields for CloudSec are editable.

When Cloudsec is enabled at MSD level, DCNM also enables **dci-advertise-pip** under **evpn multisite border-gateway** and **tunnel-encryption** on the uplinks for all Cloudsec capable gateways.

When you click **Save & Deploy**, you can verify theses configs in the **Preview Config** window for the border gateway switches.

Note – CloudSec isn't supported if the border gateway has vPC or TRM is enabled on it, that is, TRM enabled on multisite overlay IFC. If CloudSec is enabled in this scenario, appropriate warning or error messages are generated.

CloudSec Key String – Specifies the hex key string. Enter a 66 hexadecimal string if you choose **AES_128_CMAC** or enter a 130 hexadecimal string if you choose **AES_256_CMAC**.

CloudSec Cryptographic Algorithm – Choose **AES_128_CMAC** or **AES_256_CMAC**.

CloudSec Enforcement – Specifies whether the CloudSec enforcement should be strict or loose.

strict – Deploys the CloudSec configuration to all the border gateways in fabrics in MSD. If there are any border gateways that don't support CloudSec, then an error message is generated, and the configuration isn't pushed to any switch.

If **strict** is chosen, the **tunnel-encryption must-secure** CLI is pushed to the CloudSec enabled gateways within MSD.

loose – Deploys the CloudSec configuration to all the border gateways in fabrics in MSD. If there are any border gateways that don't support CloudSec, then a warning message is generated. In this case, the CloudSec config is only deployed to the switches that support CloudSec. If **loose** is chosen, the **tunnel-encryption must-secure** CLI is removed if it exists.

Note There should be at least two fabrics in MSD with border gateways that support CloudSec. If there's only one fabric with a CloudSec capable device, then the following error message is generated:

CloudSec needs to have at least 2 sites that can support CloudSec.

To remove this error, meet the criteria of having at least two sites that can support CloudSec or disable CloudSec.

CloudSec Status Report Timer – Specifies the CloudSec Operational Status periodic report timer in minutes. This value specifies how often the DCNM polls the CloudSec status data from the switch. The default value is 5 minutes and the range is from 5 to 60 minutes.

Using the CloudSec feature in DCNM, you can have all the gateways within the MSD to use the same keychain (and have only one key string) and policy. You can provide one key chain string for DCNM to form the key chain policy. DCNM forms the encryption-policy by taking all default values. DCNM pushes the same key chain policy, the same encryption-policy, and encryption-peer policies to each CloudSec capable gateways. On each gateway, there's one encryption-peer policy for each remote gateway that is CloudSec capable, using the same keychain and same key policy.

If you don't want to use the same key for the whole MSD fabric or want to enable CloudSec on a subset of all sites, you can use **switch_freeform** to manually push the CloudSec config to the switches.

Capture all the CloudSec config in **switch_freeform**.

For example, the below config is included in the **switch_freeform** policy:

```
feature tunnel-encryption
evpn multisite border-gateway 600
  dci-advertise-pip
tunnel-encryption must-secure-policy
tunnel-encryption policy CloudSec_Policy1
tunnel-encryption source-interface loopback20
key chain CloudSec_Key_Chain1 tunnel-encryption
  key 1000
    key-octet-string 7 075e731f1a5c4f524f43595f507f7d737062677147524054590 70b0b0701585440
cryptographic-algorithm AES_128_CMA
tunnel-encryption peer-ip 192.168.0.6
  keychain CloudSec_Key_Chain1 policy CloudSec_Policy1
```

Add **tunnel-encryption** in the Freeform Config of the uplink interface policy which will generate config like the following:

```
interface ethernet1/13
  no switchport
  ip address 192.168.1.14/24 tag 54321
  evpn multisite dci-tracking
  tunnel-encryption
  mtu 9216
  no shutdown
```

For more information, see Enabling Freeform Configurations on Fabric Switches , on page 343.

When CloudSec configuration is added to or removed from the switch, the DCI uplinks will flap, which will trigger multisite BGP session flapping. For multisite with existing cross site traffic, there will be traffic

disruption during this transition. Therefore, it is recommended to make the transition during a maintenance window.

If you're migrating an MSD fabric with the CloudSec configuration into DCNM, the Cloudsec related configuration is captured in **switch_freeform** and interface freeform config. You do not need to turn on Multi-Site Cloudsec in the MSD fabric setting. If you want to add more fabrics and establish CloudSec tunnels which share the same CloudSec policy including key as the existing one, then you can enable the CloudSec config in the MSD fabric settings. The CloudSec parameters in the MSD fabric setting need to match the existing CloudSec configuration on the switch. The CloudSec configuration is already captured in the freeform config, and enabling CloudSec in MSD will also generate config intents. Therefore, there's a double intent. For example, if you want to change the CloudSec key in the MSD settings, you need to remove the CloudSec freeform config because DCNM won't modify config in **switch_freeform**. Otherwise, the key in the MSD fabric settings is a conflict with the key in the freeform config.

Viewing CloudSec Operational State

From Cisco DCNM 11.5(1), you can use **CloudSec Operational View** to check the operational status of the CloudSec sessions if CloudSec is enabled on the MSD fabric.

Procedure

Step 1 Choose an MSD fabric.

The fabric topology window appears.

Step 2 Click **Tabular view** in the **Actions** pane.

Step 3 Choose the **CloudSec Operational View** tab.

Step 4 If CloudSec is disabled, the **CloudSec Operational View** tab isn't displayed.

The **Operational View** tab has the following fields and descriptions.

Fields	Descriptions
Fabric Name	Specifies the fabrics that have a CloudSec session.
Session	Specifies the fabrics and border gateway switches involved in the CloudSec session.
Link State	Specifies the status of the CloudSec session. It can be in one of the following states: • Up: The CloudSec session is successfully established between the switches. • Down: The CloudSec session isn't operational.
Uptime	Specifies the duration of the uptime for the CloudSec session. Specifically, it's the uptime since the last Rx and Tx sessions flapped, and the smaller value among the 2 sessions is displayed.
Oper Reason	Specifies the down reason for the CloudSec session state.

All these columns are sortable.

Note After CloudSec is enabled on a fabric, the operational status may not be available until after sessions are created, and the next status poll occurs.

Troubleshooting a CloudSec Session

If a CloudSec session is down, you can find more information about it using Programmable Report.

Procedure

Step 1 Navigate to **Applications** > **Programmable report**.

Step 2 Click the **Create Report** icon.

Step 3 Specify a report name, select the MSD fabric on which the report job should be run, and click **Next**.

Step 4 From the **Template** drop-down list, select **fabric_cloudsec_oper_status** and click **Create Job**.

The status will change to a green tick indicating Success after the report has been successfully generated.

Step 5 Click the report to view it. This report is similar to the **CloudSec Operation View** tab.

Step 6 Click **View Details** to view more information about the CloudSec session status.

Step 7 Click the operational status for a session to view the detailed info about the CloudSec session for each peer fabric and device.

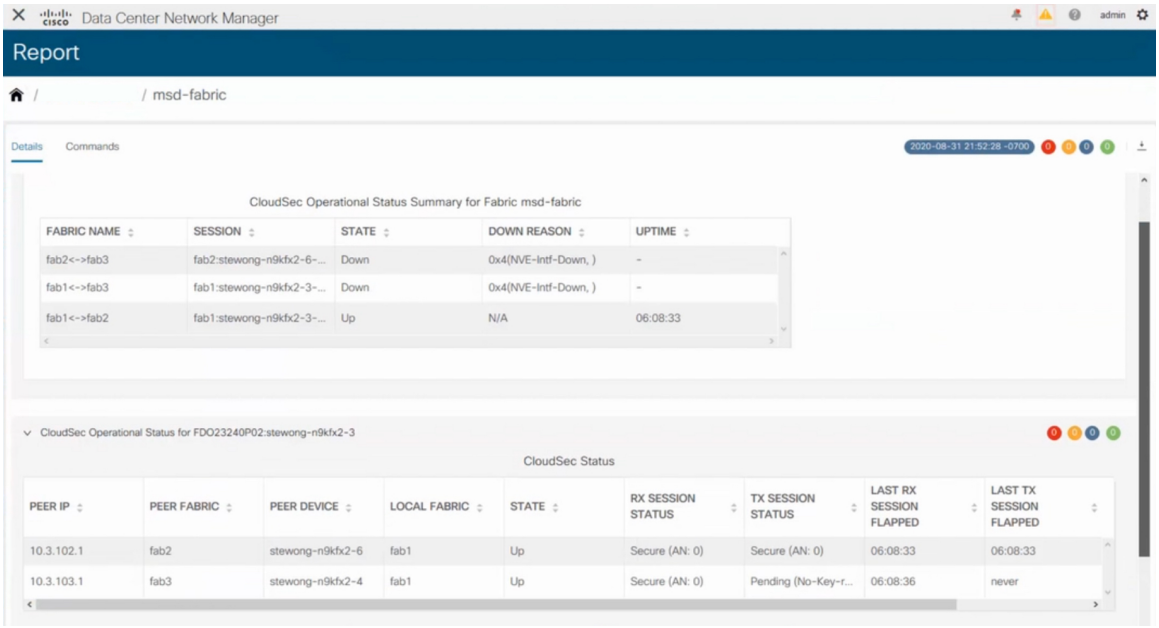

Removing a Fabric From an MSD

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 174.

Note From Cisco DCNM Release 11.4(1), after removing an individual fabric from MSD, underlay and overlay IFCs are deleted. If IFCs are extended, an error is reported to disallow the fabric remove.

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.

Managing Switches Using LAN Classic Templates

From Cisco DCNM Release 11.4(1), you can use the **LAN_Classic** and **Fabric_Group** templates to manage the switches that you used to previously manage in the DCNM Classic LAN deployment.

The **LAN_Classic** fabric template is a generic fabric template to manage Cisco Nexus switches.

Guidelines and Limitations

- Fabrics using the **LAN_Classic** fabric template can be changed to use the **External_Fabric_11_1** fabric template and then use all its associated functionalities. Note that this is the only supported fabric template conversion and it's nonreversible.
- The **LAN_Classic** fabric can be added as a member of an MSD fabric.
- Only Cisco Nexus switches are supported in the **LAN_Classic** fabric.
- The TOR Auto-Deploy functionality is supported in the **LAN_Classic** member fabric when a switch with the **ToR** role is in the fabric. For more information, see *Configuring ToR Switches and Deploying Networks*.
- If you are using the Cisco Nexus 7000 Series Switch with Cisco NX-OS Release 6.2(24a) on the LAN Classic or External fabrics, make sure to enable AAA IP Authorization in the fabric settings.
- The following features in the **LAN_Classic** template provide the same support as they do for the **External_Fabric_11_1** template:

  The following features are supported:

  - Configuration compliance
  - Backup or restore of fabric
  - Network Insights
  - Performance monitoring
  - VMM
  - Topology view
  - Kubernetes visualization
  - RBAC

  For more information, refer to the feature specific sections.

Creating a LAN Classic Fabric

Procedure

Step 1 Navigate to **Control** > **Fabrics** > **Fabric Builder**.

Step 2 Click **Create Fabric**.

Step 3 Enter the fabric name and choose **LAN_Classic** from **Fabric Template** drop - down list.

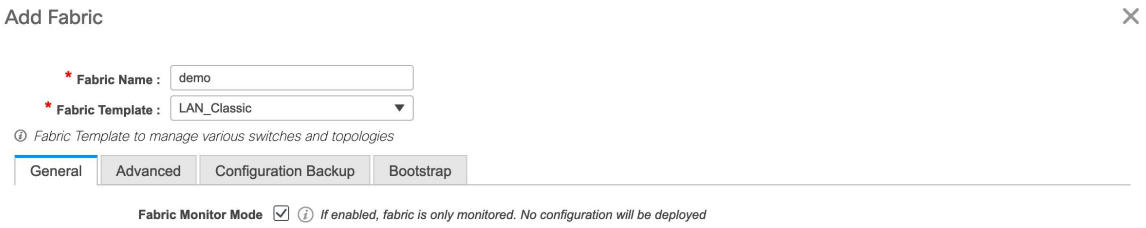

Step 4 The **General** tab is displayed by default. The field in this tab is:

Fabric Monitor Mode – Uncheck the check box if you want DCNM to manage the fabric. Keep the check box selected to enable only monitoring of the fabric. In this state, you can't deploy configurations on its switches.

The configurations must be pushed for devices before you discover them in the fabric. You can't push configurations in the monitor mode.

Step 5 Click **Advanced** tab. The fields in this tab are:

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 MPLS Handoff - Select the check box to enable the MPLS Handoff feature. For more information, see the *Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - MPLS SR and LDP Handoff* chapter.

Underlay MPLS Loopback Id: Specifies the underlay MPLS loopback ID. The default value is 101.

Enable AAA IP Authorization - Enables AAA IP authorization, when IP Authorization is enabled in the AAA Server.

Enable DCNM as Trap Host - Select this check box to enable DCNM as a trap host.

Enable CDP for Bootstrapped Switch - Enables CDP on management interface.

Enable NX-API - Specifies enabling of NX-API. This check box is unchecked by default.

Enable NX-API on HTTP port - Specifies enabling of NX-API on HTTP. This check box is unchecked by default. Enable this check box and the **Enable NX-API** check box to use HTTP. If you uncheck this check box, the applications that use NX-API and supported by Cisco DCNM, such as Layer 4-Layer 7 services (L4-L7 services), VXLAN OAM, and so on, start using the HTTPS instead of HTTP.

Note If you check the **Enable NX-API** check box and the **Enable NX-API on HTTP** check box, applications use HTTP.

Inband Mgmt: For External and Classic LAN Fabrics, this knob enables DCNM to import and manage of switches with inband connectivity (reachable over switch loopback, routed, or SVI interfaces) , in addition to management of switches with out-of-band connectivity (aka reachable over switch mgmt0 interface). The only requirement is that for Inband managed switches, there should be IP reachability from DCNM to the switches via the eth2 aka inband interface. For this purpose, static routes may be needed on the DCNM, that in turn can be configured via the Administration->Customization->Network Preferences option. After enabling Inband management, during discovery, provide the IPs of all the switches to be imported using Inband Management and set maximum hops to 0. DCNM has a pre-check that validates that the Inband managed switch IPs are reachable over the eth2 interface. Once the pre-check has passed, DCNM then discovers and learns about the interface on that switch that has the specified discovery IP in addition to the VRF that the

interface belongs to. As part of the process of switch import/discovery, this information is captured in the baseline intent that is populated on the DCNM. For more information, see Inband Management in External Fabrics and LAN Classic Fabrics, on page 186.

Note Bootstrap or POAP is only supported for switches that are reachable over out-of-band connectivity, that is, over switch mgmt0. The various POAP services on the DCNM are typically bound to the eth1 or out-of-band interface. In scenarios, where DCNM eth0/eth1 interfaces reside in the same IP subnet, the POAP services are bound to both interfaces.

Enable Precision Time Protocol (PTP): Enables PTP across a fabric. When you select this check box, PTP is enabled globally and on core-facing interfaces. Additionally, the PTP Source Loopback Id and PTP Domain Id fields are editable. For more information, see Precision Time Protocol for External Fabrics and LAN Classic Fabrics, on page 188.

Fabric Freeform - You can apply configurations globally across all the devices discovered in the external fabric using this freeform field.

AAA Freeform Config – Specifies the AAA freeform configs.

Step 6 Click the **Resources** tab. The fields in this tab are:

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.

Underlay MPLS Loopback IP Range - Specifies the underlay MPLS SR or LDP loopback IP address range.

The IP range should be unique, that is, it shouldn't overlap with IP ranges of the other fabrics.

Step 7 Click **Configuration** tab. The fields in this tab are:

Hourly Fabric Backup: Select the check box to enable an hourly backup of fabric configurations and the intent.

Scheduled Fabric Backup: Check the check box to enable a daily backup.

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.

Note Hourly or scheduled backup runs only after the next CC hourly run. Backup will run only after scheduled time is elapsed and whenever CC run happens after the elapsed time.

The backup and restore process is similar to that of an external fabric. For more information about backing up and restoring external fabrics, see Fabric Backup and Restore, on page 311.

Step 8 Click **Bootstrap** tab. The fields in this tab are:

Enable Bootstrap (For NX-OS Switches Only) - Select this check box to enable the bootstrap feature for only Cisco Nexus switches.

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 **Switch Mgmt Default Gateway** and **Switch Mgmt IP Subnet Prefix** fields.
- Local DHCP Server: Enable **Local DHCP Server** check box 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, all the remaining fields become editable.

DHCP Version – Select DHCPv4 or DHCPv6 from this drop-down list. When you select DHCPv4, the **Switch Mgmt IPv6 Subnet Prefix** field is disabled. If you select DHCPv6, the **Switch Mgmt IP Subnet Prefix** is disabled.

Note Cisco DCNM IPv6 POAP isn't supported with Cisco Nexus 7000 Series Switches. Cisco Nexus 9000 and 3000 Series Switches support IPv6 POAP only when switches are either L2 adjacent (eth1 or out-of-band subnet must be a /64) or they are L3 adjacent residing in some IPv6 /64 subnet. Subnet prefixes other than /64 aren't supported.

If you don't 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 Mgmt Default Gateway - Specifies the default gateway for the management VRF on the switch.

Switch Mgmt IP 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.

Switch Mgmt IPv6 Subnet Prefix - Specifies the IPv6 prefix for the Mgmt0 interface on the switch. The prefix should be between 112 and 126. This field is editable if you enable IPv6 for DHCP.

Enable AAA Config - Enables AAA configure. It includes AAA configs from the **Advanced** tab during device bootup.

Bootstrap Freeform Config - (Optional) Enter extra commands as needed. For example, if you're using AAA or remote authentication-related configurations, 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 ConfigErrors in Switches*.

DHCPv4/DHCPv6 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

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

Adding Switches to LAN Classic Fabric

Procedure

Step 1 Click **Add** switches. The **Inventory Management** window comes up.

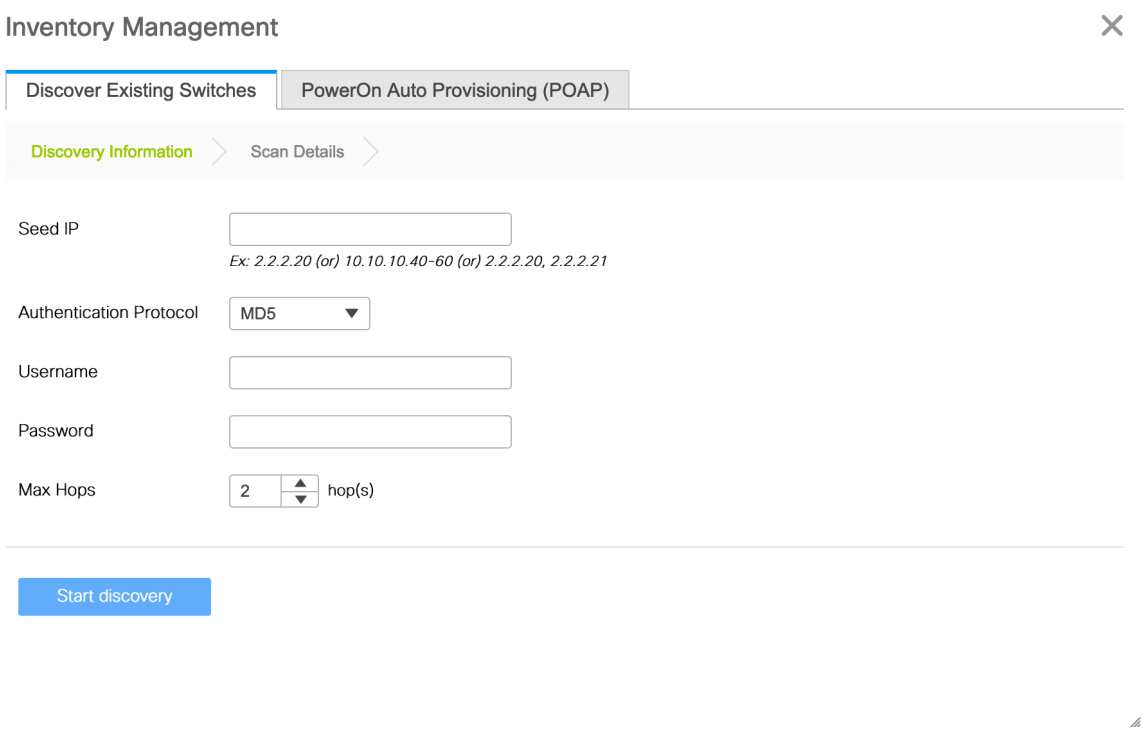

You can also add switches by clicking **Tabular View** > **Switches** > **+** .

Step 2 Enter IP address (**Seed IP**) of the switch.

Step 3 Enter the administrator username and password of the switch.

Step 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.

Step 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.

Step 6 Click **Refresh** topology to view the latest topology view.

For more information, see:

- Discovering Existing Switches, on page 73

- Discovering New Switches, on page 79

Creating a Fabric Group and Associating Member Fabrics

This procedure shows how to a create a **Fabric_Group** and add **LAN_Classic** fabrics. The **Fabric_Group** template is used for grouping **LAN_Classic** fabrics for visualization.

The following functionalities aren't supported in a **Fabric_Group**:

- Fabric backup or restore
- VXLAN overlay or IFC deployment
- Changing fabric template to and from any other fabric template
- Since **Fabric_Group** doesn't manage any configurations, clicking **Save & Deploy** reports an error.

Procedure

Step 1 Navigate to **Control** > **Fabrics** > **Fabric Builder**.

Step 2 Click **Create Fabric**.

Step 3 Enter the fabric name and choose **Fabric_Group** from the **Fabric Template** drop - down list.

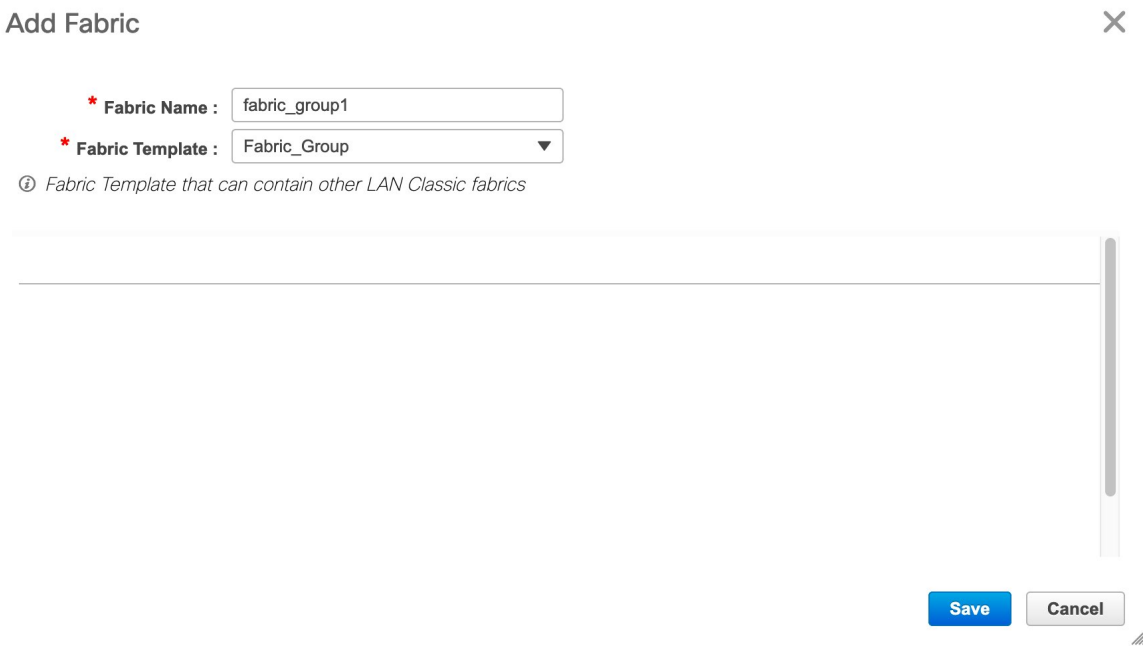

Step 4 Click **Save**.

Step 5 In the **Actions** panel, click **Move Fabrics**.

Step 6 Select a **LAN_Classic** fabric in the **Move Fabric** window.

Note You can select and add only a **LAN_Classic** fabric in a fabric group.

Step 7 Click **Add**.

Similarly, you can remove a member fabric by selecting it and clicking **Remove**.

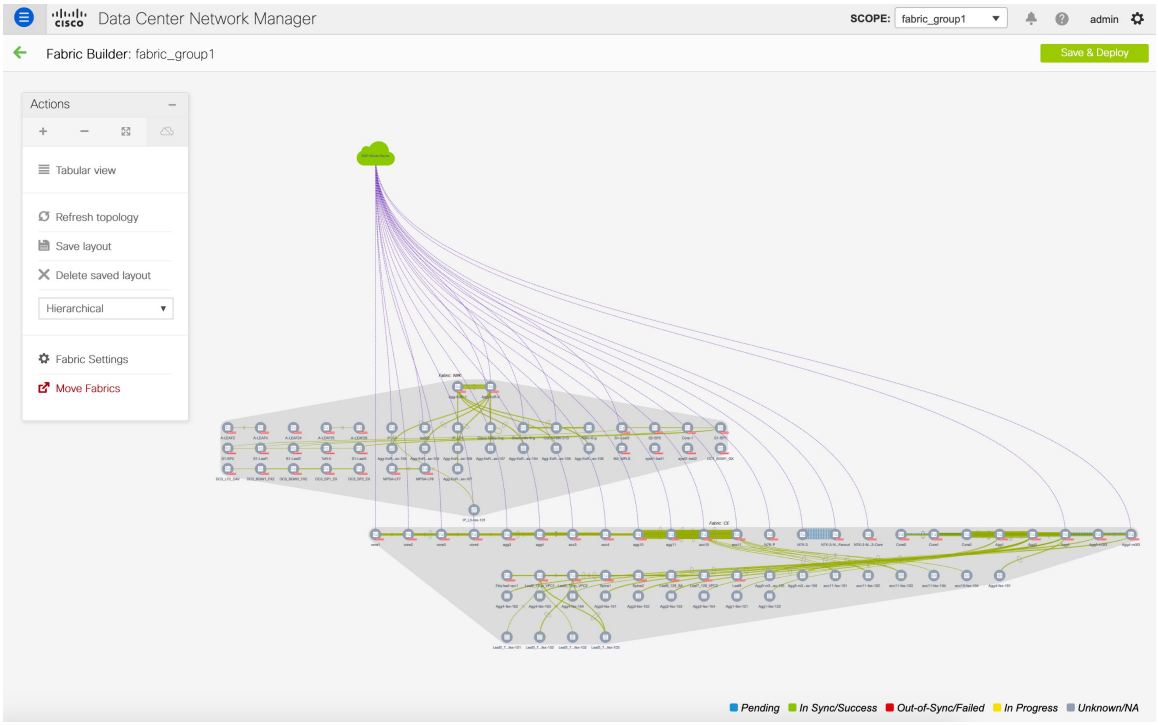

Support for Inter-Fabric Connection in LAN Classic Fabric Template

The **LAN_Classic** fabric supports VRF-Lite, Multi-Site, and MPLS IFCs with these conditions:

- The **LAN_Classic** fabric as a destination for DCI/VRF-Lite and Multi-Site IFCs is supported, but you can only manually create them by providing the required information. They won't be automatically created even when the auto deployment options are enabled in the **Easy_Fabric_11_1** and **MSD_Fabric_11_1** fabrics.
- You can't add nonexistent (meta) switches to a **LAN_Classic** fabric. A meta switch is a placeholder for a switch or device that DCNM can't discover.
- The base BGP configurations for the 'Edge Router' and 'Core Router' switch roles aren't auto generated. Configure them using the **switch_freeform** policies or other suitable means.
- If MPLS Handoff is enabled in the fabric settings, MPLS base configurations are auto generated for the 'Edge Router' and 'Core Router' switch roles.

Inband Management in External Fabrics and LAN Classic Fabrics

From Release 11.5(1), Cisco DCNM allows you to import or discover switches with inband connectivity for External and LAN Classic fabrics in Brownfield deployments only. Enable inband management, per fabric,

while configuring or editing the Fabric settings. You cannot import or discover switches with inband connectivity using POAP.

After configuration, the Fabric tries to discover switches based on the VRF of the inband management. The fabric template determines the VRF of inband switch using seed IP. If there are multiple VRFs for same seed IP, then no intent will be learnt for seed interfaces. You must create intent/configuration manually.

After configuring\editing the Fabric settings, you must Save and Deploy. You cannot change the Inband Mgmt settings after you import inband managed switches to the Fabric. If you uncheck the checkbox, the following error message is generated.

```
Inband IP <<IP Address>> cannot be used to import the switch,
please enable Inband Mgmt in fabric settings and retry.
```

After the switches are imported to the Fabric, you must manage the interfaces to create intent. Create the intent for the interfaces that you're importing the switch. Edit\update the Interface configuration. When you try to change the Interface IP, for this inband managed switch, an error message is generated:

```
Interface <<interface_name>> is used as seed or next-hop egress interface
for switch import in inband mode.
IP/Netmask Length/VRF changes are not allowed for this interface.
```

While managing the interfaces, for switches imported using inband management, you cannot change the seed IP for the switch. The following error will be generated:

```
<<switch-name>>: Mgmt0 IP Address (<ip-address>) cannot be changed,
when is it used as seed IP to discover the switch.
```

Create a policy for next-hop interfaces. Routes to DCNM from 3rd party devices may contain multiple interfaces, known as ECMP routes. Find the next-hop interface and create an intent for the switch. Interface IP and VRF changes are not allowed.

If inband management is enabled, during Image management, eth2 IP address is used to copy images on the switch, in ISSU, EPLD, RPM & SMU installations flows.

If you imported the switches using inband connectivity in the fabric, and later disable the inband Mgmt in the Fabric settings after deployment, the following error message is generated:

```
The fabric <<fabric name>> was updated with below message:
Fabric Settings cannot be changed for Inband Mgmt, when switches are already imported
using inband Ip. Please remove the existing switches imported using Inband Ip from the
fabric,
then change the Fabric Setttings.
```

However, the same fabric can contain switches imported using both inband and out-of-band connectivity.

Inband POAP Management in External Fabrics and LAN Classic Fabrics

In Release 11.5(3a), Cisco DCNM allows you to provision switches using POAP for External and LAN Classic fabrics in Brownfield deployments only. You must enable Inband POAP on the DCNM server using the **appmgr** command. Configure Inband POAP per fabric, while configuring or editing the Fabric settings. You can import or discover switches with inband connectivity using POAP.

TOR switches are basically connected to the vPC pair of leafs in Easy Fabrics. However, TOR switches have inband reachability only, and the management ports are not connected. Inband POAP allows you to bring up TOR via inband POAP pre-provisioning.

After configuring Inband POAP, devices can be pre-provisioned using the regular process. After pre-provisioning, the only the basic hostname and username configs are generated. Ensure that the seed IP address provided during bootstrap is configured appropriately on TOR switches.

The following guidelines are applicable to Inband POAP feature:

- You must enable Inband POAP on the DCNM server before configuring the feature for the Fabric.
- You cannot disable Inband Management when Inband POAP feature is enabled.
- You cannot enable Inband POAP with DHCPv6.
- This feature is supported only for External Fabric and LAN Classic Fabrics.
- Inband DHCP is also supported to devices that are layer-2 adjacent on eth2 interface.
- You can enable inband POAP with or without enabling local DHCP; however, inband POAP with external DHCP server must be supported.

To configure Inband POAP for External or Classic LAN Fabrics, you must enable Inband POAP on the DCNM server.

To enable POAP on DCNM server, perform the following steps:

1. Login to DCNM server as a **`/root`** user.
2. Enable Inband POAP feature on both Active and Standby nodes by using the **appmgr inband-poap {enable|disable|status}** command.

```
[root@dcnm-active]# appmgr inband-poap enable

[root@dcnm-standby]# appmgr inband-poap enable
```

After inband POAP is enabled on the DCNM server, launch the DCNM Web UI to edit the Fabric settings and configure Inband POAP. See Creating an External Fabric , on page 114.

Precision Time Protocol for External Fabrics and LAN Classic Fabrics

From Release 11.5(1), in the fabric settings for the **External_Fabric_11_1** or **LAN_Classic** template, select the **Enable Precision Time Protocol (PTP)** check box to enable PTP across a fabric. When you select this check box, PTP is enabled globally and on core-facing interfaces. Additionally, the **PTP Loopback Id** and **PTP Domain Id** fields are editable.

The PTP feature is supported with Cisco Nexus 9000 Series cloud-scale switches, with NX-OS version 7.0(3)I7(1) or later. Warnings are displayed if there are non-cloud scale devices in the fabric, and PTP is not enabled. Examples of the cloud-scale devices are Cisco Nexus 93180YC-EX, Cisco Nexus 93180YC-FX, Cisco Nexus 93240YC-FX2, and Cisco Nexus 93360YC-FX2 switches. For more information, refer to https://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html#~products.

Note PTP global configuration is supported with Cisco Nexus 3000 Series switches; however, PTP and ttag configurations are not supported.

For more information, see the *Configuring PTP* chapter in *Cisco Nexus 9000 Series NX-OS System Management Configuration Guide* and *Cisco Network Insights for Resources Application for Cisco DCNM User Guide*.

For External and LAN Classic fabric deployments, you have to enable PTP globally, and also enable PTP on core-facing interfaces. The interfaces could be configured to the external PTP server like a VM or Linux-based machine. Therefore, the interface should be edited to have a connection with the grandmaster clock. For PTP and TTAG configurations to be operational on External and LAN Classic Fabrics, you must sync up of Switch

Configs to DCNM using the **host_port_resync** policy. For more information, see Sync up Out-of-Band Switch Interface Configurations with DCNM, on page 190.

It is recommended that the grandmaster clock should be configured outside of Easy Fabric and it is IP reachable. The interfaces toward the grandmaster clock need to be enabled with PTP via the interface freeform config.

All core-facing interfaces are auto-enabled with the PTP configuration after you click **Save & Deploy**. This action ensures that all devices are PTP synced to the grandmaster clock. Additionally, for any interfaces that are not core-facing, such as interfaces on the border devices and leafs that are connected to hosts, firewalls, service-nodes, or other routers, the ttag related CLI must be added. The ttag is added for all traffic entering the VXLAN EVPN fabric and the ttag must be stripped when traffic is exiting this fabric.

Here is the sample PTP configuration:`feature ptp`

```
feature ptp

ptp source 100.100.100.10 -> IP address of the loopback interface (loopback0)
that is already created, or user-created loopback interface in the fabric settings

ptp domain 1 -> PTP domain ID specified in fabric settings

interface Ethernet1/59 -> Core facing interface
  ptp

interface Ethernet1/50 -> Host facing interface
  ttag
  ttag-strip
```

The following guidelines are applicable for PTP:

- The PTP feature can be enabled in a fabric when all the switches in the fabric have Cisco NX-OS Release 7.0(3)I7(1) or a higher version. Otherwise, the following error message is displayed:

  ```
  PTP feature can be enabled in the fabric, when all the switches have
  NX-OS Release 7.0(3)I7(1) or higher version. Please upgrade switches to
  NX-OS Release 7.0(3)I7(1) or higher version to enable PTP in this fabric.
  ```

- For hardware telemetry support in NIR, the PTP configuration is a prerequisite.
- If you are adding a non-cloud scale device to an existing fabric which contains PTP configuration, the following warning is displayed:

  ```
  TTAG is enabled fabric wide, when all devices are cloud-scale switches
  so it cannot be enabled for newly added non cloud-scale device(s).
  ```

- If a fabric contains both cloud-scale and non-cloud scale devices, the following warning is displayed when you try to enable PTP:

  ```
  TTAG is enabled fabric wide when all devices are cloud-scale switches
  and is not enabled due to non cloud-scale device(s).
  ```

- TTAG configuration is generated for all the devices if host configuration sync up is performed on all the devices. Ttag configuration will not be generated for any newly added devices if host configuration sync up is not performed on all newly added devices.

  If the configuration is not synced, the following warning is displayed:

  ```
  TTAG on interfaces with PTP feature can only be configured for cloud-scale devices.
  It will not be enabled on any newly added switches due to the presence of non cloud-scale
   devices.
  ```

- PTP and TTAG configurations are deployed on host interfaces.

- PTP and TTAG Configurations are supported between switches in the same fabric (intra-fabric links). PTP is created for inter-fabric links, and ttag is created for the inter-fabric link if the other fabric (Switch) is not managed by DCNM. Inter-fabric links do not support PTP or ttag configurations if both fabrics are managed by DCNM.
- TTAG configuration is configured by default after the breakout. After the links are discovered and connected post breakout, perform Save & Deploy to generate the correct configuration based on the type of port (host, intra-fabric link, or inter fabric link).

Sync up Out-of-Band Switch Interface Configurations with DCNM

From DCNM release 11.5(1), any interface level configuration made outside of DCNM (via CLI) can be synced to DCNM and then managed from DCNM. Also, the vPC pair configurations are automatically detected and paired. This applies to the External_Fabric_11_1 and LAN_Classic fabrics only. The vPC pairing is performed with the **vpc_pair** policy.

Note When DCNM is managing switches, ensure that all configuration changes are initiated from DCNM and avoid making changes directly on the switch.

When the interface config is synced up to the DCNM intent, the switch configs are considered as the reference, that is, at the end of the sync up, the DCNM intent reflects what is present on the switch. If there were any undeployed intent on DCNM for those interfaces before the resync operation, they will be lost.

Guidelines

- Supported in fabrics using the following templates: Easy_Fabric_11_1, External_Fabric_11_1, and LAN_Classic.
- Supported for Cisco Nexus switches only.
- Supported for interfaces that don't have any fabric underlay related policy associated with them prior to the resync. For example, IFC interfaces and intra fabric links aren't subjected to resync.
- Supported for interfaces that do not have any custom policy (policy template that isn't shipped with Cisco DCNM) associated with them prior to resync.
- Supported for interfaces where the intent is not exclusively owned by a Cisco DCNM feature and/or application prior to resync.
- Supported on switches that don't have Interface Groups associated with them.
- Interface mode (switchport to routed, trunk to access, and so on) changes aren't supported with overlays attached to that interface.

The sync up functionality is supported for the following interface modes and policies:

Interface Mode	Policies
trunk (standalone, po, and vPC PO)	• int_trunk_host_11_1 • int_port_channel_trunk_host_11_1 • int_vpc_trunk_host_11_1

access (standalone, po, and vPC PO)	• int_access_host_11_1 • int_port_channel_access_host_11_1 • int_vpc_access_host_11_1
dot1q-tunnel	• int_dot1q_tunnel_host_11_1 • int_port_channel_dot1q_tunnel_host_11_1 • int_vpc_ dot1q_tunnel_host_11_1
routed	int_routed_host_11_1
loopback	int_freeform
sub-interface	int_subif_11_1
FEX (ST, AA)	• int_port_channel_fex_11_1 • int_port_channel_aa_fex_11_1
breakout	interface_breakout
nve	int_freeform (only in External_Fabric_11_1/LAN_Classic)
SVI	int_freeform (only in External_Fabric_11_1/LAN_Classic)
mgmt0	int_mgmt_11_1

In an Easy fabric, the interface resync will automatically update the network overlay attachments based on the access VLAN or allowed VLANs on the interface.

After the resync operation is completed, the switch interface intent can be managed using normal DCNM procedures.

Syncing up Switch Interface Configurations to DCNM

Before you begin

- We recommend taking a fabric backup before attempting the interface resync.
- In **External_Fabric_11_1** and **LAN_Classic** fabrics, for the vPC pairing to work correctly, both the switches must be in the fabric and must be functional.
- Ensure that the switches are **In-Sync** and not in **Migration-mode** or **Maintenance-mode**.

Procedure

Step 1 In DCNM, navigate to **Control > Fabric Builder** and click a fabric.

Step 2 Ensure that switches are present in the fabric and vPC pairings are completed, and they are shown in the **Topology** view. Click **Tabular view** in the **Actions** panel.

Step 3 From **Tabular view**, select one or more switches where the interface intent resync is needed, and click **Policies**.

Note
- If a pair of switches is already paired with either **no_policy** or **vpc_pair**, select only one switch of the pair.
- If a pair of switches is not paired, then select both the switches.

Step 4 In the **Policies** window, click the **Add Policy** icon.

Step 5 In the **Add Policy** window, select **host_port_resync** from the **Policy** drop-down list. Click **Save**.

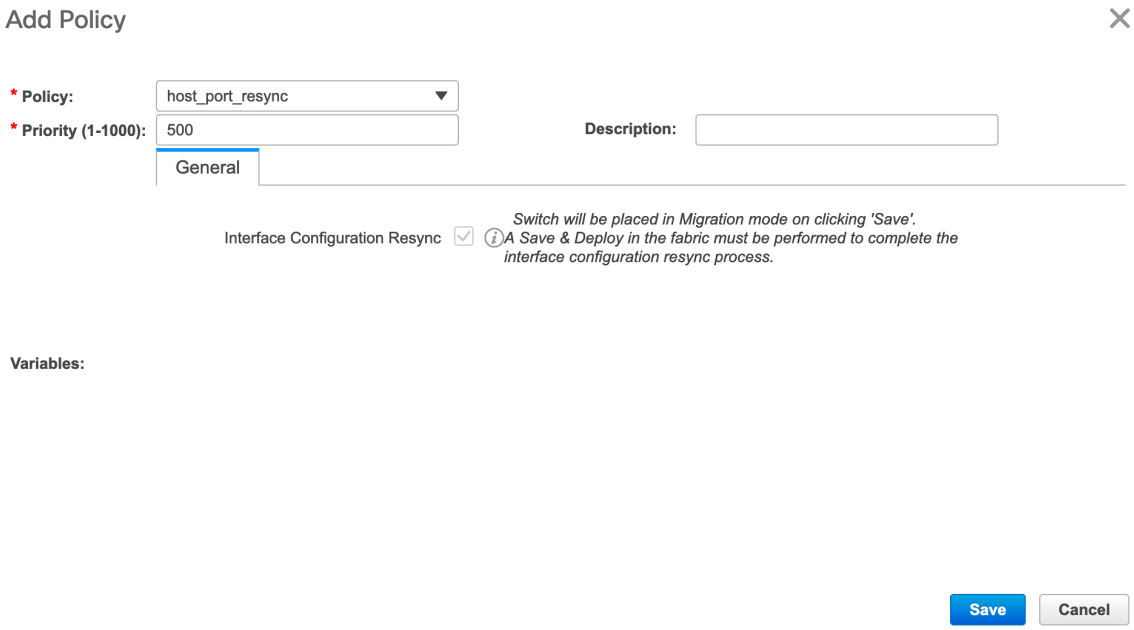

Step 6 Check the **Mode** column for the switches to ensure that they report **Migration**. For a vPC pair, both switches are in the **Migration-mode**.
- After this step, the switches in the **Topology view** are in **Migration-mode**.
- Both the switches in a vPC pair are in the migration mode even if one of the switches is placed into this mode.
- If switch(es) are unintentionally put into the resync mode, they can be moved back to the normal mode by identifying the **host_port_resync** policy instance and deleting it from the **Policies** window.

Step 7 After the configuration changes are ready to sync up to DCNM, navigate to the **Tabular view**, select the required switches, and click **Rediscover switch** to ensure that DCNM is aware of any new interfaces and other changes.

Step 8 Click **Save & Deploy** to start the resync process.

Note This process might take some time to complete based on the size of the switch configuration and the number of switches involved.

Step 9 The **Config Deployment** window is displayed if no errors are detected during the resync operation. The interface intent is updated in DCNM.

Note If the External_Fabric_11_1 or LAN_Classic fabric is in **Monitored Mode**, an error message indicating that the fabric is in the read-only mode is displayed. This error message can be ignored and doesn't mean that the resync process has failed.

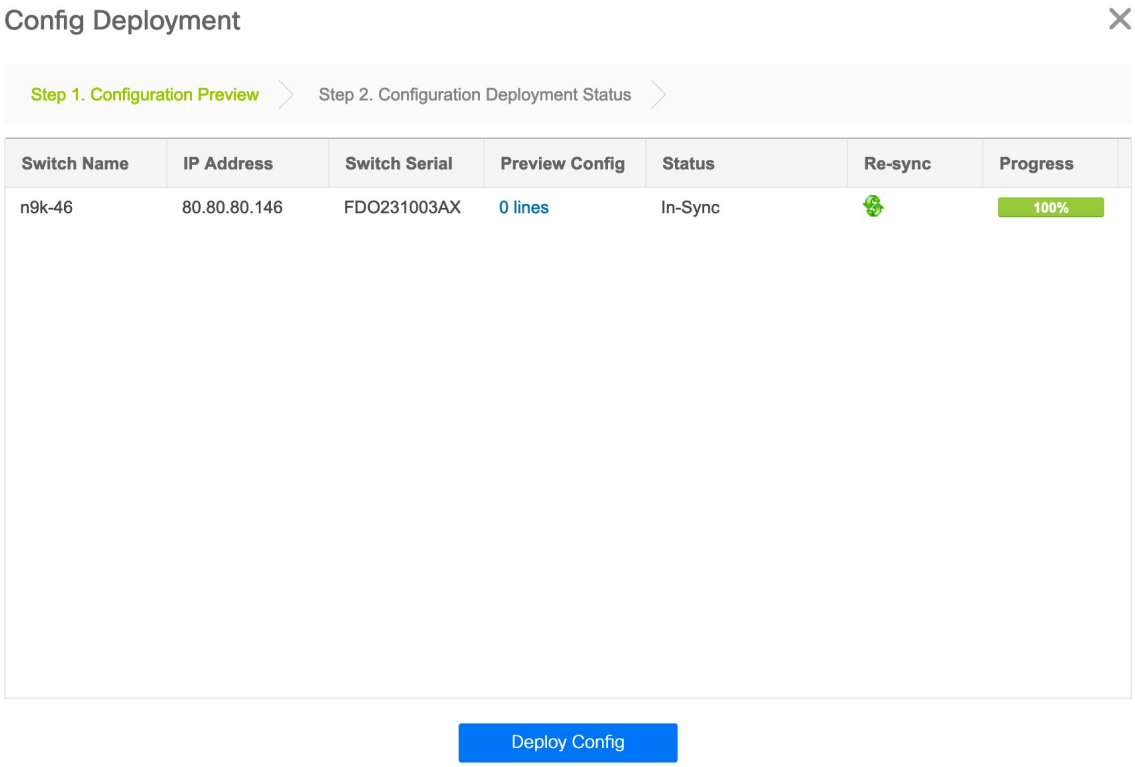

Close the **Config Deployment** window, and you can see that the switches are automatically moved out of the **Migration-mode**. Switches in a vPC pair that were not paired or paired with **no_policy** show up as paired and associated with the **vpc_pair** policy.

Note The **host_port_resync** policy that was created for the switch is automatically deleted after the resync process is completed successfully.

What to do next

The following limitations are applicable after Syncing up Switch Interface Configurations to DCNM:

- The port channel membership (once the policy exists) is not supported.
- Changing the interface mode (trunk to access etc.) that have overlays attached is not supported.
- Resync for interfaces that belong to **Interface Groups** are not supported.
- The vPC pairing in **External_Fabric_11_1** and **LAN_Classic** templates must be updated with the **vpc_pair** policy.
- Changing the interface mode that have overlays attached is not supported.

- In **Easy_Fabric** fabrics, VXLAN overlay interface attachments are performed automatically based on the allowed VLANs.

MACsec Support in Easy Fabric and eBGP Fabric

From Cisco DCNM Release 11.5(1), MACsec is supported in the Easy Fabric and eBGP Fabric on intra-fabric links. You should enable MACsec on the fabric and on each required intra-fabric link to configure MACsec. Unlike CloudSec, auto-configuration of MACsec is not supported.

MACsec is supported on switches with minimum Cisco NX-OS Releases 7.0(3)I7(8) and 9.3(5).

Note Support for MACsec is a preview feature in the Cisco DCNM Release 11.5(1).

Guidelines

- If MACsec cannot be configured on the physical interfaces of the link, an error is displayed when you click **Save**. MACsec cannot be configured on the device and link due to the following reasons:
  - The minimum NX-OS version is not met.
  - The interface is not MACsec capable.
- MACsec global parameters in the fabric settings can be changed at any time.
- MACsec and CloudSec can coexist on a BGW device.
- MACsec is not supported on Border Leaf.
- MACsec status of a link with MACsec enabled is displayed on the **Links** window.
- Brownfield migration of devices with MACsec configured is supported using switch and interface freeform configs.

  For more information about MACsec configuration, which includes supported platforms and releases, see the Configuring MACsec chapter in *Cisco Nexus 9000 Series NX-OS Security Configuration Guide*.

The following sections show how to enable and disable MACsec in DCNM:

Enabling MACsec

Procedure

Step 1 Navigate to **Control > Fabrics > Fabric Builder**.

Step 2 Click **Create Fabric** to create a new fabric or click **Edit Fabric** on an existing Easy or eBGP fabric.

Step 3 Click the **Advanced** tab and specify the MACsec details.

Enable MACsec – Select the check box to enable MACsec for the fabric.

MACsec Primary Key String – Specify a Cisco Type 7 encrypted octet string that is used for establishing the primary MACsec session. For AES_256_CMAC, the key string length must be 130 and for

AES_128_CMAC, the key string length must be 66. If these values are not specified correctly, an error is displayed when you save the fabric.

Note The default key lifetime is infinite.

MACsec Primary Cryptographic Algorithm – Choose the cryptographic algorithm used for the primary key string. It can be **AES_128_CMAC** or **AES_256_CMAC**. The default value is **AES_128_CMAC**.

You can configure a fallback key on the device to initiate a backup session if the primary session fails.

MACsec Fallback Key String - Specify a Cisco Type 7 encrypted octet string that is used for establishing a fallback MACsec session. For AES_256_CMAC, the key string length must be 130 and for AES_128_CMAC, the key string length must be 66. If these values are not specified correctly, an error is displayed when you save the fabric.

MACsec Fallback Cryptographic Algorithm - Choose the cryptographic algorithm used for the fallback key string. It can be AES_128_CMAC or AES_256_CMAC. The default value is AES_128_CMAC.

MACsec Cipher Suite – Choose one of the following MACsec cipher suites for the MACsec policy:

- GCM-AES-128
- GCM-AES-256
- GCM-AES-XPN-128
- GCM-AES-XPN-256

The default value is **GCM-AES-XPN-256**.

Note The MACsec configuration is not deployed on the switches after the fabric deployment is complete. You need to enable MACsec on intra-fabric links to deploy the MACsec configuration on the switch.

MACsec Status Report Timer - Specifies MACsec operational status periodic report timer in minutes.

Step 4 Click a fabric, click **Tabular View** in the **Actions** panel, and then click **Links**.

Step 5 Choose an intra-fabric link on which you want to enable MACsec and click **Update Link**.

Step 6 In the **Link Management – Edit Link** window, click **Advanced** in the **Link Profile** section, and select the **Enable MACsec** check box.

If MACsec is enabled on the intra fabric link but not in the fabric settings, an error is displayed when you click **Save**.

When MACsec is configured on the link, the following configurations are generated:

- Create MACsec global policies if this is the first link that enables MACsec.
- Create MACsec interface policies for the link.

Step 7 Click **Save** and then click **Save & Deploy** to deploy the MACsec configuration.

Disabling MACsec

To disable MACsec on an intra-fabric link, navigate to the **Link Management – Edit Link** window, unselect the **Enable MACsec** check box, click **Save**, and then click **Save & Deploy**. This action performs the following:

- Deletes MACsec interface policies from the link.
- If this is the last link where MACsec is enabled, MACsec global policies are also deleted from the device.

Only after disabling MACsec on links, navigate to the **Fabric Settings** and unselect the **Enable MACsec** check box under the **Advanced** tab to disable MACsec on the fabric. If there's an intra-fabric link in the fabric with MACsec enabled, an error is displayed when you click **Save & Deploy**.

Overview of Tenant Routed Multicast

Tenant Routed Multicast (TRM) enables multicast forwarding on the VXLAN fabric that uses a BGP-based EVPN control plane. TRM provides multi-tenancy aware multicast forwarding between senders and receivers within the same or different subnet local or across VTEPs.

With TRM enabled, multicast forwarding in the underlay is leveraged to replicate VXLAN encapsulated routed multicast traffic. A Default Multicast Distribution Tree (Default-MDT) is built per-VRF. This is an addition to the existing multicast groups for Layer-2 VNI Broadcast, Unknown Unicast, and Layer-2 multicast replication group. The individual multicast group addresses in the overlay are mapped to the respective underlay multicast address for replication and transport. The advantage of using a BGP-based approach allows the VXLAN BGP EVPN fabric with TRM to operate as fully distributed Overlay Rendezvous-Point (RP), with the RP presence on every edge-device (VTEP).

A multicast-enabled data center fabric is typically part of an overall multicast network. Multicast sources, receivers, and multicast rendezvous points might reside inside the data center but also might be inside the campus or externally reachable via the WAN. TRM allows a seamless integration with existing multicast networks. It can leverage multicast rendezvous points external to the fabric. Furthermore, TRM allows for tenant-aware external connectivity using Layer-3 physical interfaces or subinterfaces.

For more information, see the following:

- Guidelines and Limitations for Tenant Routed Multicast
- Guidelines and Limitations for Layer 3 Tenant Routed Multicast
- Guidelines and Limitations for Layer 2/Layer 3 Tenant Routed Multicast (Mixed Mode)

Overview of Tenant Routed Multicast with VXLAN EVPN Multi-Site

Tenant Routed Multicast with Multi-Site enables multicast forwarding across multiple VXLAN EVPN fabrics connected via Multi-Site.

The following two use cases are supported:

- Use Case 1: TRM provides Layer 2 and Layer 3 multicast services across sites for sources and receivers across different sites.
- Use Case 2: Extending TRM functionality from VXLAN fabric to sources receivers external to the fabric.

TRM Multi-Site is an extension of BGP-based TRM solution that enables multiple TRM sites with multiple VTEPs to connect to each other to provide multicast services across sites in most efficient possible way. Each TRM site is operating independently and border gateway on each site allows stitching across each site. There can be multiple Border Gateways for each site. In a given site, the BGW peers with Route Sever or BGWs of other sites to exchange EVPN and MVPN routes. On the BGW, BGP will import routes into the local VRF/L3VNI/L2VNI and then advertise those imported routes into the Fabric or WAN depending on where the routes were learnt from.

Tenant Routed Multicast with VXLAN EVPN Multi-Site Operations

The operations for TRM with VXLAN EVPN Multi-Site are as follows:

- Each Site is represented by Anycast VTEP BGWs. DF election across BGWs ensures no packet duplication.
- Traffic between Border Gateways uses ingress replication mechanism. Traffic is encapsulated with VXLAN header followed by IP header.
- Each Site will only receive one copy of the packet.
- Multicast source and receiver information across sites is propagated by BGP protocol on the Border Gateways configured with TRM.
- BGW on each site receives the multicast packet and re-encapsulate the packet before sending it to the local site.

For information about guidelines and limitations for TRM with VXLAN EVPN Multi-Site, see Configuring Tenant Routed Multicast.

Configuring TRM for Single Site Using Cisco DCNM

This section is assumes that a VXLAN EVPN fabric has already been provisioned using Cisco DCNM.

Procedure

Step 1 Enable TRM for the selected Easy Fabric. If the fabric template is **Easy_Fabric_11_1**, click the Fabric settings, navigate to the **Replication** tab, and check the **Enable Tenant Routed Multicast (TRM)** field. In addition, the default MDT multicast group field is auto-populated with a default value.

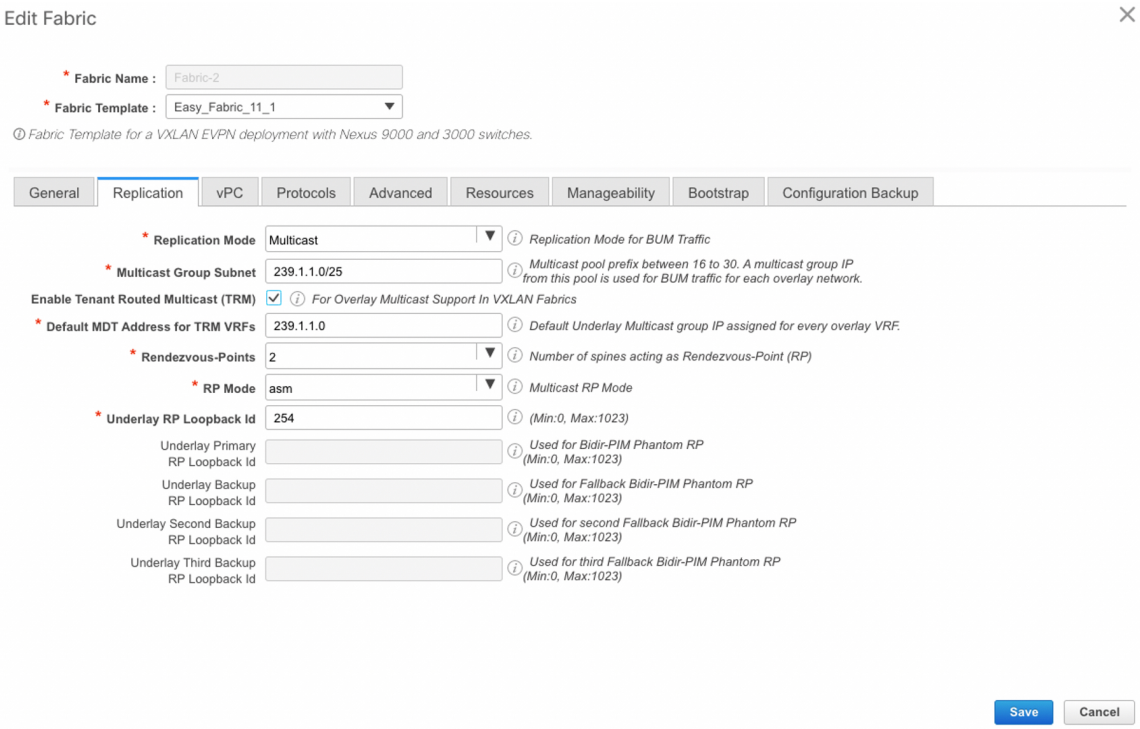

Enable Tenant Routed Multicast (TRM): Select the check box to enable Tenant Routed Multicast (TRM) that allows overlay multicast traffic to be supported over EVPN/MVPN in the VXLAN BGP EVPN fabric.

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**.

Click **Save** to save the fabric settings. At this point, all the switches turn “Blue” as it will be in the pending state. Click **Save and Deploy** to enable the following:

- Enable feature ngmvpn: Enables the Next-Generation Multicast VPN (ngMVPN) control plane for BGP peering.
- Configure ip multicast multipath s-g-hash next-hop-based: Multipath hashing algorithm for the TRM enabled VRFs.
- Configure ip igmp snooping vxlan: Enables IGMP Snooping for VXLAN VLANs.
- Configure ip multicast overlay-spt-only: Enables the MVPN Route-Type 5 on all MPVN enabled Cisco Nexus 9000 switches.
- Configure and Establish MVPN BGP AFI Peering: This is necessary for the peering between BGP RR and the Leaves.

For VXLAN EVPN fabric created using Easy_Fabric_eBGP fabric template, **Enable Tenant Routed Multicast (TRM)** field and **Default MDT Address for TRM VRFs** field can be found on the fabric settings' EVPN tab.

Step 2 Enable TRM for the VRF.

Navigate to **Control > VRFs** and edit the selected VRF. Navigate to the **Advanced Tab** and edit the following TRM settings:

TRM Enable – Select the check box to enable TRM. If you enable TRM, then the RP address and the underlay multicast address must be entered.

Is RP External – Enable this check box if the RP is external to the fabric. If this field is unchecked, RP is distributed in every VTEP.

Note If the RP is external, then select the appropriate option. If the RP is external, then RP loopback ID is greyed out.

RP Address – Specifies the IP address of the RP.

RP Loopback ID – Specifies the loopback ID of the RP, if **Is RP External** is not enabled.

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 **Default MDT Address for TRM VRFs** field in the fabric settings screen is auto-populated in this field. User can override this field if a different multicast group address should be used for this VRF.

Overlay Multicast Groups – Specifies the multicast group subnet for the specified RP. The value is the group range in "ip pim rp-address" command. If the field is empty, 224.0.0.0/24 is used as default.

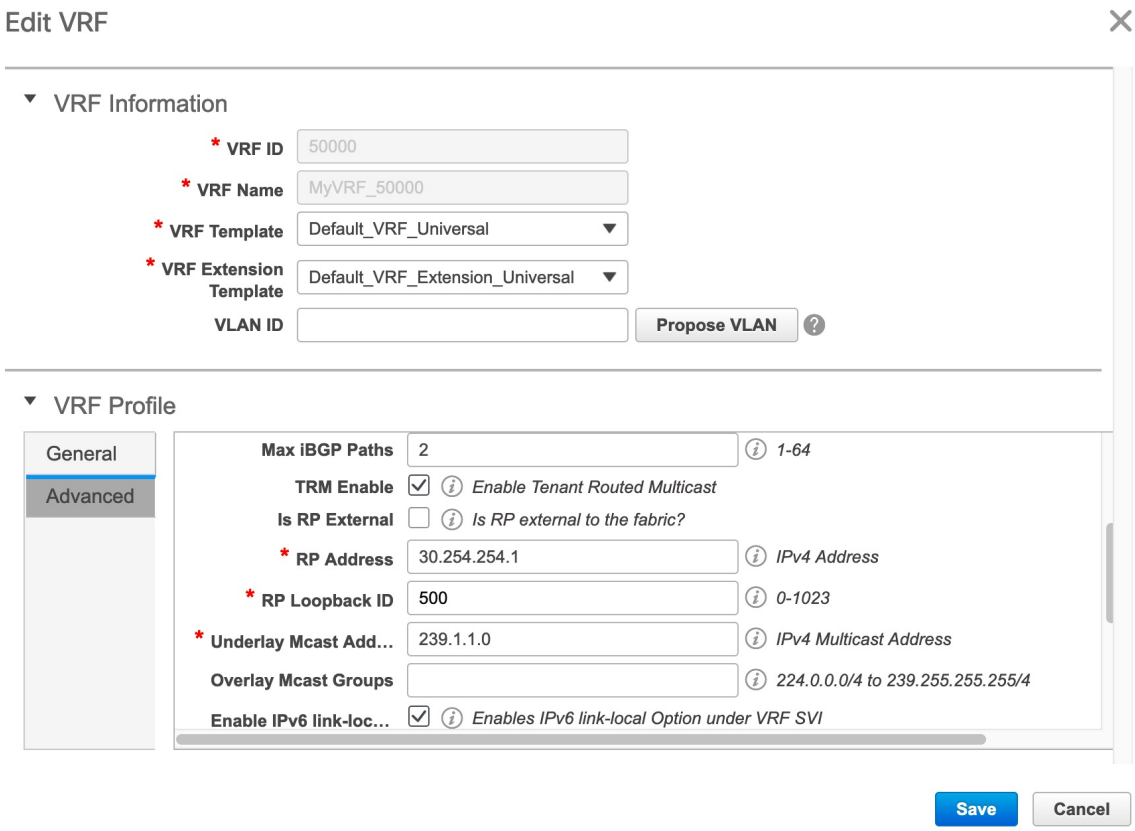

Click **Save** to save the settings. The switches go into the pending state, that is, blue color. These settings enable the following:

- Enable PIM on L3VNI SVI.
- Route-Target Import and Export for MVPN AFI.
- RP and other multicast configuration for the VRF.
- Loopback interface using the above RP address and RP loopback id for the distributed RP.

Step 3 Enable TRM for the network.

Navigate to **Control > Networks**. Edit the selected network and navigate to the **Advanced** tab. Edit the following TRM setting:

TRM enable – Select the check box to enable TRM.

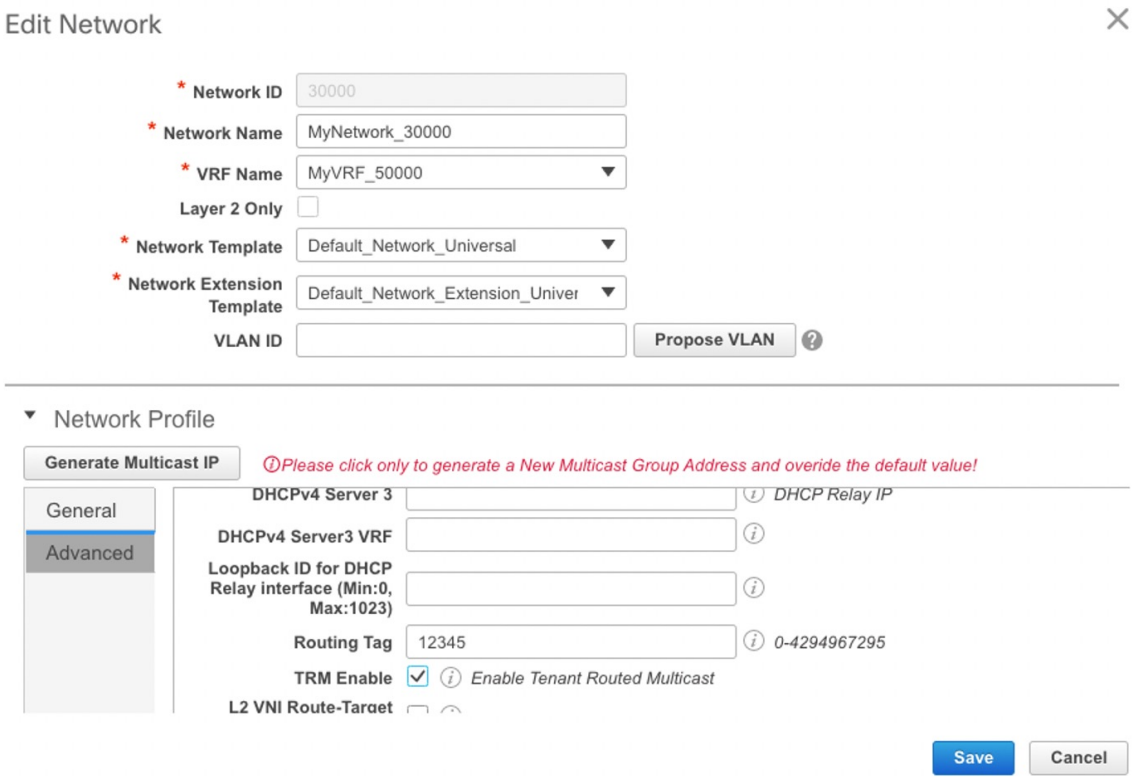

Click **Save** to save the settings. The switches go into the pending state, that is, the blue color. The TRM settings enable the following:

- Enable PIM on the L2VNI SVI.
- Create a PIM policy **none** to avoid PIM neighborship with PIM Routers within a VLAN. The **none** keyword is a configured route map to deny any ipv4 addresses to avoid establishing PIM neighborship policy using anycast IP.

Configuring TRM for Multi-Site Using Cisco DCNM

This section assumes that a Multi-Site Domain (MSD) has already been deployed by Cisco DCNM and TRM needs to be enabled.

Procedure

Step 1 Enable TRM on the BGWs.

Navigate to **Control > VRFs**. Make sure that the right DC Fabric is selected under the **Scope** and edit the VRF. Navigate to the **Advanced** tab. Edit the TRM settings. Repeat this process for every DC Fabric and its VRFs.

TRM Enable – Select the check box to enable TRM. If you enable TRM, then the RP address and the underlay multicast address must be entered.

Is RP External – Enable this check box if the RP is external to the fabric. If this field is unchecked, RP is distributed in every VTEP.

Note If the RP is external, then select the appropriate option. If the RP is external, then RP loopback ID is greyed out.

RP Address – Specifies the IP address of the RP.

RP Loopback ID – Specifies the loopback ID of the RP, if **Is RP External** is not enabled.

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 **Default MDT Address for TRM VRFs** field in the fabric settings screen is auto-populated in this field. User can override this field if a different multicast group address should be used for this VRF.

Overlay Multicast Groups – Specifies the multicast group subnet for the specified RP. The value is the group range in “ip pim rp-address” command. If the field is empty, 224.0.0.0/24 is used as default.

Enable TRM BGW MSite - Select the check box to enable TRM on Border Gateway Multi-Site.

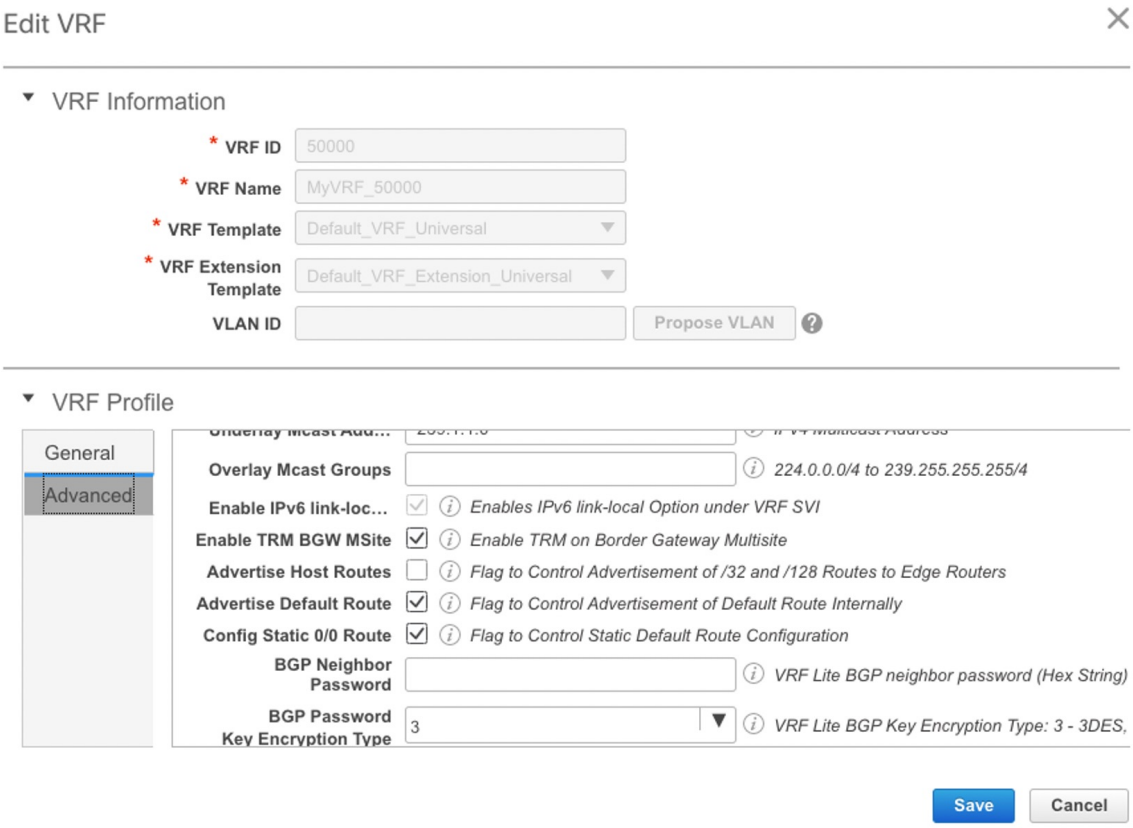

Click on **Save** to save the settings. The switches go into the pending state, that is, blue color. These settings enable the following:

- Enable feature ngmvpn: Enables the Next-Generation Multicast VPN (ngMVPN) control plane for BGP peering.
- Enables PIM on L3VNI SVI.
- Configures L3VNI Multicast Address.
- Route-Target Import and Export for MVPN AFI.
- RP and other multicast configuration for the VRF.
- Loopback interface for the distributed RP.
- Enable Multi-Site BUM ingress replication method for extending the Layer 2 VNI

Step 2 Establish MVPN AFI between the BGWs.

Navigate to **Control > Fabrics**. Select the MSD fabric. Click **Tabular view** and click **Links**. Filter it by the policy - **Overlays**.

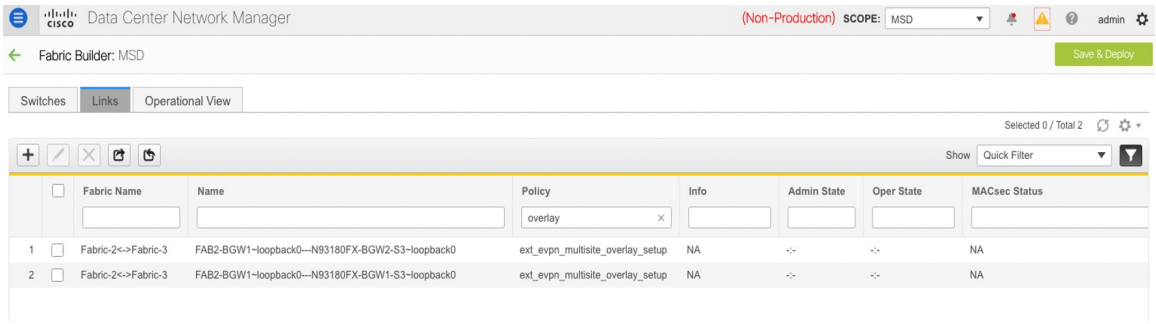

Select and edit each overlay peering to enable TRM by checking the **Enable TRM** check box.

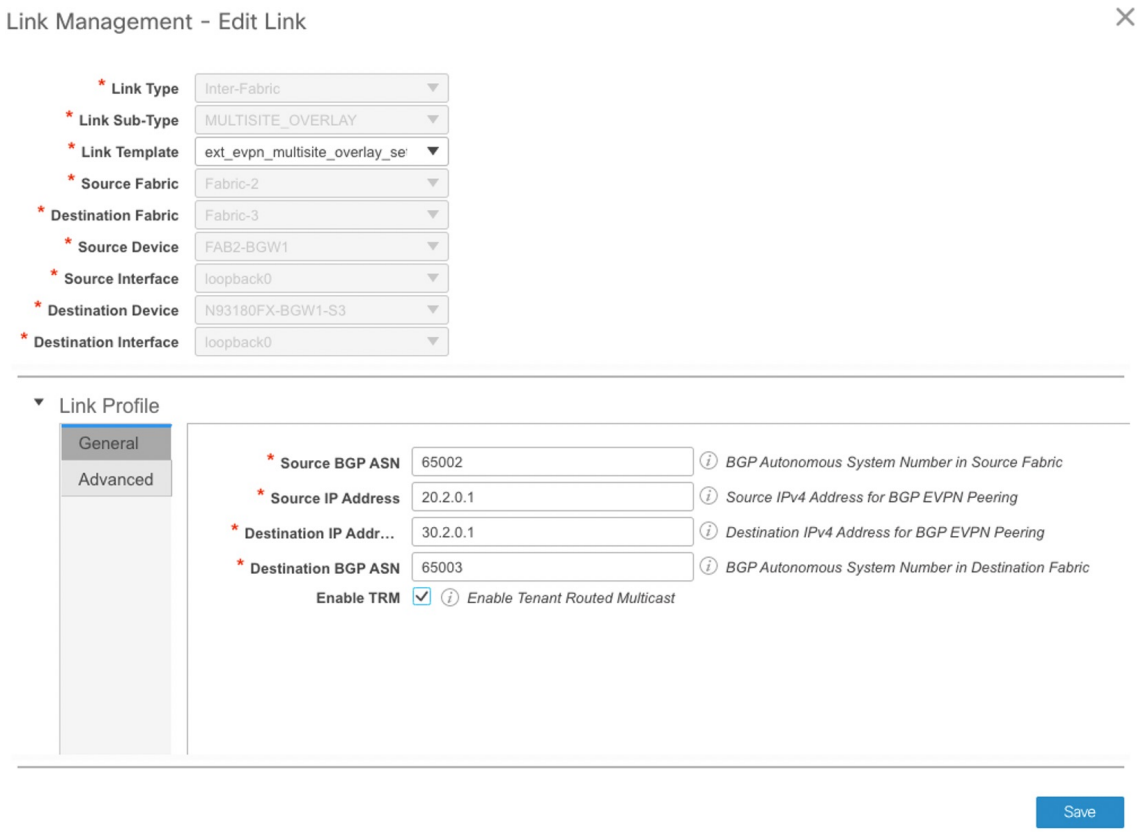

Click **Save** to save the settings. The switches go into the pending state, that is, the blue color. The TRM settings enable the MVPN peering's between the BGWs, or BGWs and Route Server.

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:

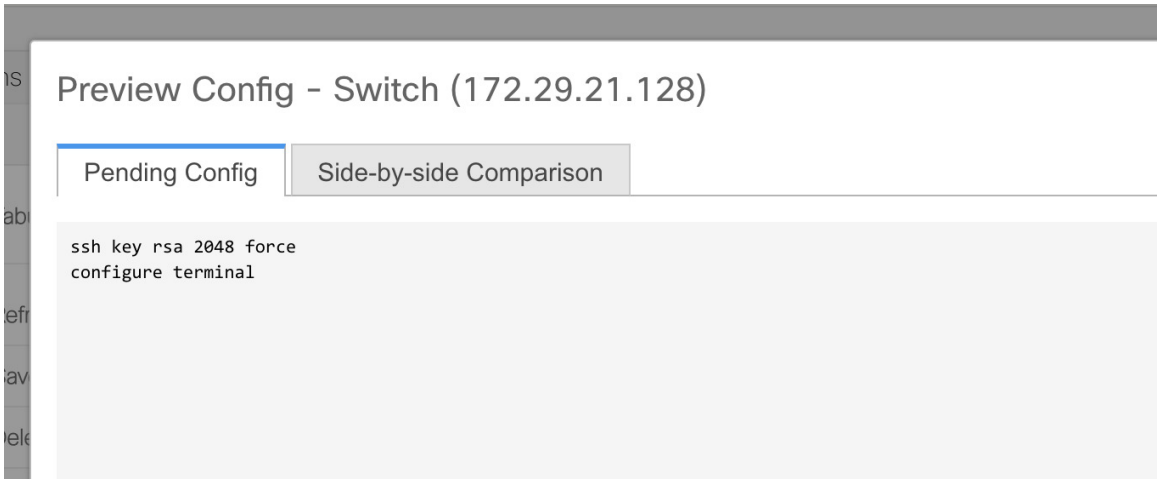

- Delete the PTI with the **ssh key rsa 2048 force** command in the intent by clicking **View/Edit Policies** in the **Tabular View** of the **Fabric Builder** window.

- 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

Alternatively, you can also navigate to the **Tabular view** from the **Actions** pane. Choose one or more devices of the same device type and click **Set Role** to set roles for devices. The device types are:

- NX-OS
- IOS XE
- IOS XR
- Other

Note Ensure that you have moved switches from maintenance mode to active mode or operational mode before setting roles.

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

You can assign one of the following roles for non-Nexus devices:

- Spine
- Leaf
- Access (This role is available only for Cisco ASR 1000 Series routers and Cisco Catalyst 9000 Series switches).
- Edge Router (Use this role for VRF-Lite).
- Core Router
- Super Spine
- Preview Config
- ToR (This role is available only for Cisco Catalyst 9000 series switches).

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

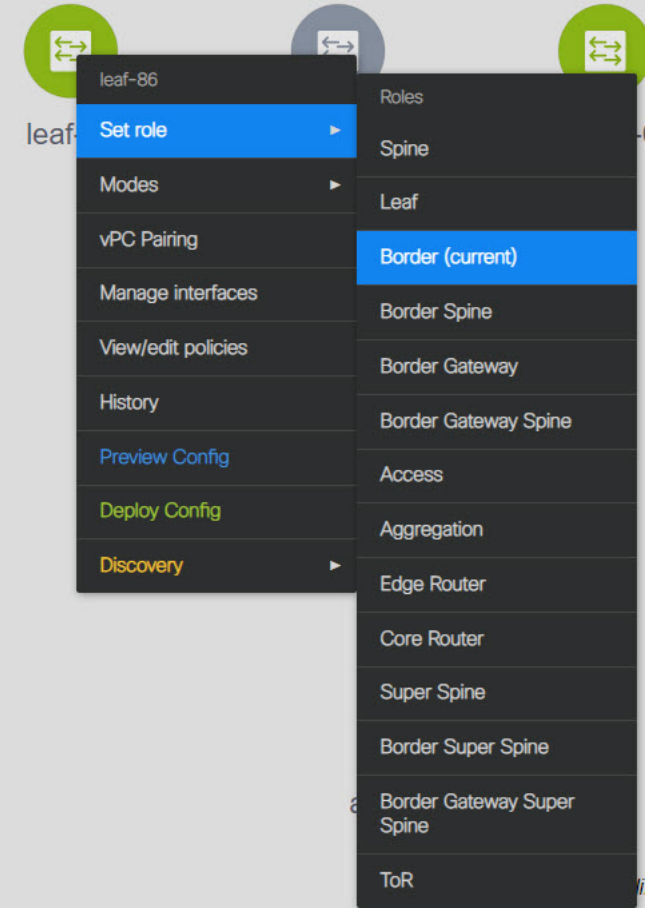

You cannot change the switch role from any Leaf role to any Spine role and from any Spine role to any Leaf role.

In case the switch role is not changed according to the allowed switch role changes mentioned above for easy fabrics, 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.
```

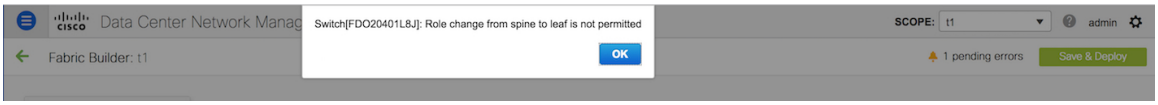

You can then change the switch role to the role that was set earlier, or set a new role, and configure the fabric.

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.

Running EXEC Mode Commands in DCNM

When you first log in, the Cisco NX-OS software places you in the EXEC mode. The commands available in the EXEC mode include the show commands that display the device status and configuration information, the clear commands, and other commands that perform actions that you do not save in the device configuration.

The following procedure shows how to run EXEC commands in DCNM:

Procedure

Step 1 From DCNM, navigate to **Control > Fabrics > Fabric Builder**.

Step 2 Click a fabric and then click **Tabular view** in the **Actions** menu.

Step 3 Select one or more switches and click the **Play** button (Execute Commands).

Step 4 From the **Template** drop-down list, select **exec_freeform**.

Step 5 Enter the commands in the **Freeform CLI** field.

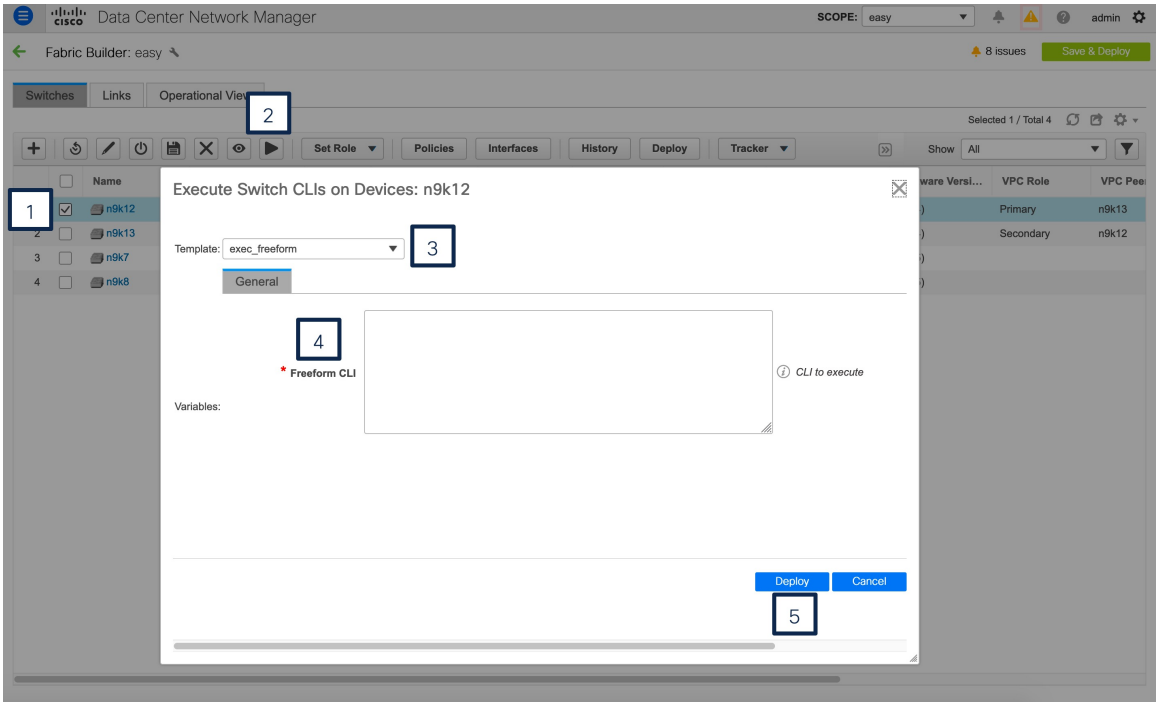

Step 6 Click **Deploy** to run the EXEC commands.

Step 7 In the **CLI Execution Status** window, you can check the status of the deployment. Click **Detailed Status** under the **Command** column to view details.

Step 8 In the **Command Execution Details** window, click the info under the **CLI Response** column to view the output or response.

Fabric Multi Switch Operations

Click **Tabular view** from the **Actions** pane in the fabric topology window. The tabular view has the following tabs:

- Tabular View - Switches
- Tabular View - Links
- Tabular View - Operational View

Tabular View - Switches

You can manage switch operations in this tab. Each row represents a switch in the fabric, and displays switch details, including its serial number.

Some of the actions that you can perform from this tab are also available when you right-click a switch in the fabric topology window. However, the **Switches** tab enables you to provision configurations on multiple switches, like deploying policies, simultaneously.

The **Switches** tab has following information of every switch you discover in the fabric:

- Name: Specifies the switch name.
- IP Address: Specifies the IP address of the switch.
- Role: Specifies the role of the switch.
- Serial Number: Specifies the serial number of the switch.
- Fabric Name: Specifies the name of the fabric, where the switch is discovered.
- Fabric Status: Specifies the status of the fabric, where the switch is discovered.
- Discover Status: Specifies the discovery status of the switch.
- Model: Specifies the switch model.
- Software Version: Specifies the software version of the switch.
- Last Updated: Specifies when the switch was last updated.
- Mode: Specifies the current mode of the switch.
- VPC Role: Specifies the vPC role of the switch.
- VPC Peer: Specifies the vPC peer of the switch.

The **Switches** tab has the following icons and buttons:

- Add switches: Click this icon to discover existing or new switches to the fabric. The **Inventory Management** dialog box appears.

This option is also available in the fabric topology window. Click **Add switches** in the **Actions** pane.

Refer the following sections for more information:

- Adding Switches to a Fabric: Provides information on adding switches to easy fabrics.
- Discovering New Switches: Provide information on adding Cisco Nexus switches to external fabrics.
- Adding non-Nexus Devices to External Fabrics: Provide information on adding non-Nexus switches to external fabrics.

- Rediscover switch: Initiate the switch discovery process by DCNM afresh.
- Update discovery credentials: Update device credentials such as authentication protocol, username and password.
- Saving config and Reload: Save the configurations and reload the switch.

Note This option is grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.

- Copy running to startup config: From Cisco DCNM, Release 11.4(1), you can perform an on-demand copy running-configuration to startup-configuration operation for one or more switches.

Note This option will be grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.

- Remove switches: Remove the switch from the fabric.

Note This option will be grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.

- Preview: You can preview the pending configurations and the side-by-side comparison of running configurations and expected configurations.
- 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 **Policies** window. 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.

- Interfaces: Deploy configurations on the switch interfaces.
- History: View the deployment history and the policy change history using this button. Choose one or more switches and click **History**.

The **Policy Change History** tab lists the history of policies along with the users who made the changes like add, update, or delete.

Under the **Policy Change History** tab, for a policy, click **Detailed History** under the **Generated Config** column to view the generated config before and after.

The following table provides the summary of generated config before and after for Policy Template Instances (PTIs).

PTI Operations	Generated Config Before	Generated Config After
Add	Empty	Contains the config
Update	Contains config before changes	Contains config after changes
Mark-Delete	Contains the config to be removed	Contains the config to be removed with colour change
Delete	Contains the config	Empty

Note When a policy or profile template is applied, an instance is created for each application of the template. This instance is known as Policy Template Instance or PTI.

- Deploy: Deploy switch configurations. From Cisco DCNM Release 11.3(1), you can deploy configurations for multiple devices using the **Deploy** button.

Note

- This option grays out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.
- In an MSD fabric, you can deploy configurations only on the Border Gateway, Border Gateway Spine, Border Gateway Super-Spine, or External Fabric switches.

- Set Role: Choose one or more devices of the same device type and click **Set Role** to set roles for devices. The device types are:
  - NX-OS
  - IOS XE
  - IOS XR
  - Other

  Ensure that you have moved switches from maintenance mode to active mode or operational mode before setting roles. See the Switch Operations section for more information on setting roles.
- vPC Pairing: Choose a switch and click **vPC Pairing** to create, edit, or unpair a vPC pair. However, you can use this option only when you choose a Cisco Nexus switch. Refer the following sections for more information:

- Creating a vPC Setup: Provides information on how to create a vPC pair in external fabrics.
- vPC Fabric Peering: Provides information on how to create a vPC pair in easy fabrics.

Tabular View - Links

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

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

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

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

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

Creating Intra-Fabric Links

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

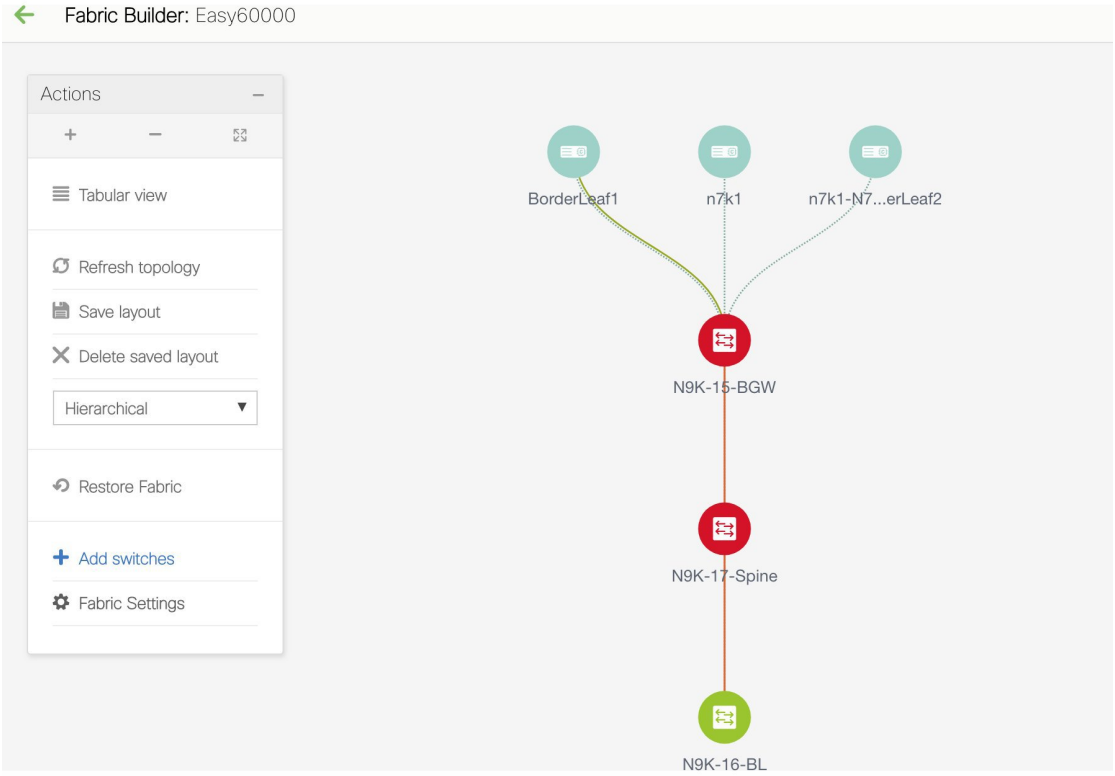

A screen with the tabs Switches and Links appears. They list the fabric switches and links in a table.

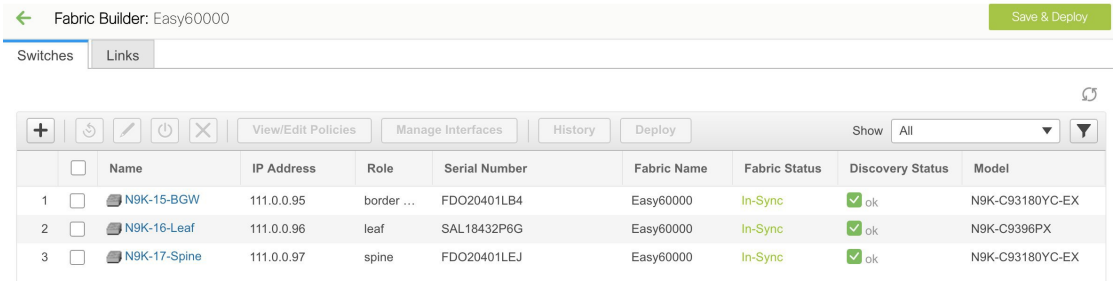

4. Click the Links tab. You can see a list of links.

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

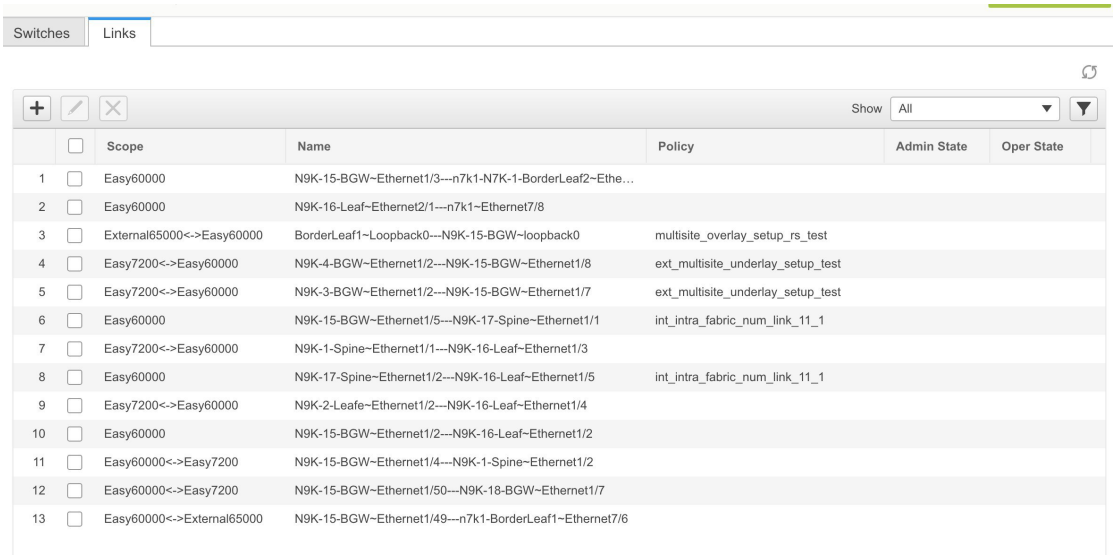

5. Click the Add (+) button at the top left part of the screen to add a link.

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

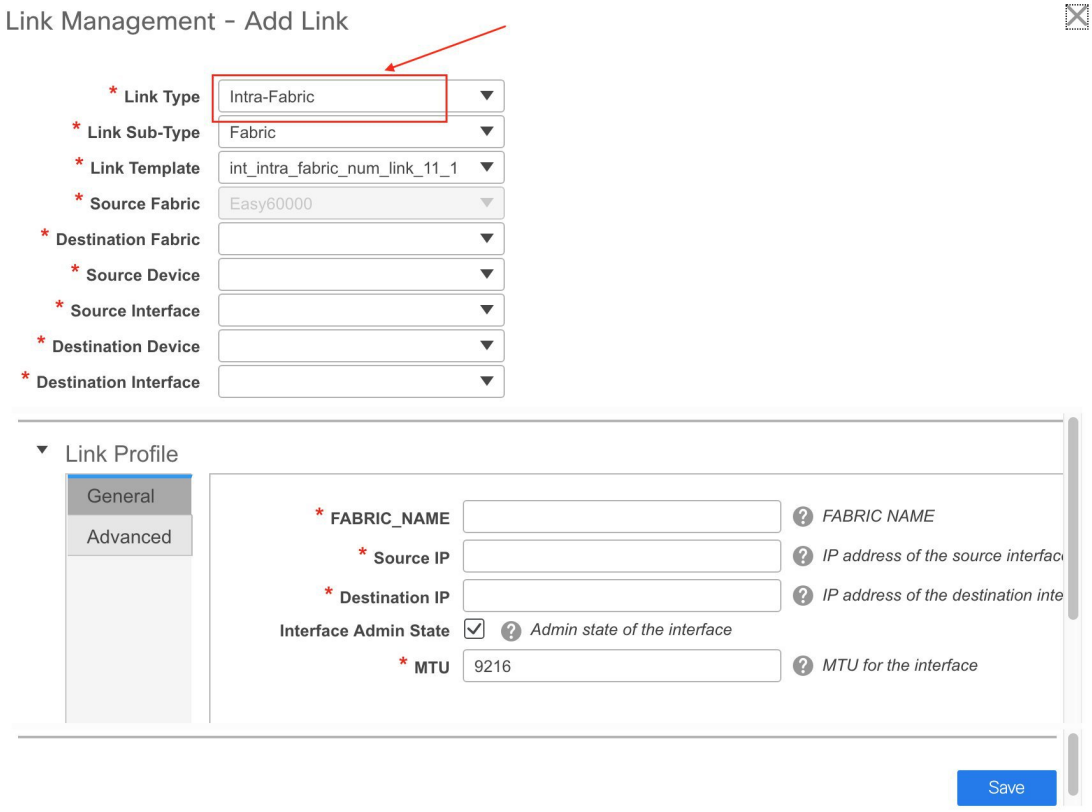

   The fields are:

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

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

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

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

Correspondingly, the Link Profile section fields is updated.

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

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

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

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

Note Select the pre-provisioned device as the destination device if you are creating a link between an existing d 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_** template.

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

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

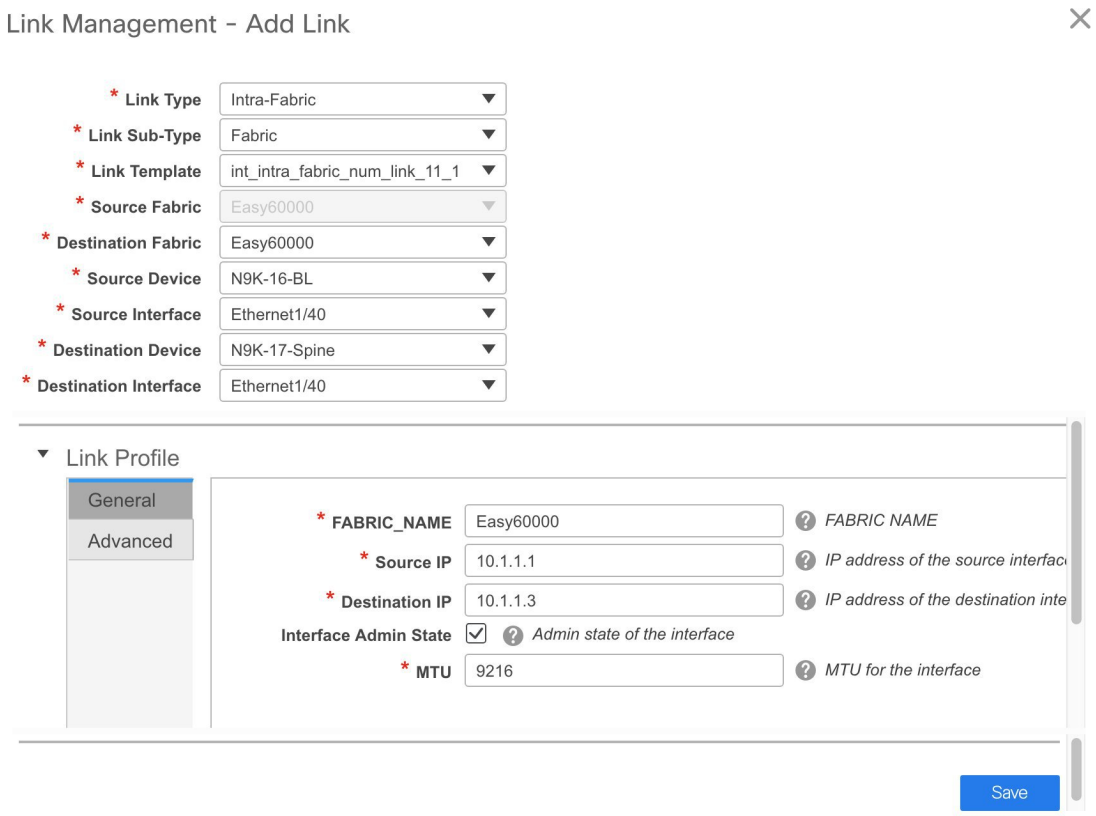

Advanced tab.

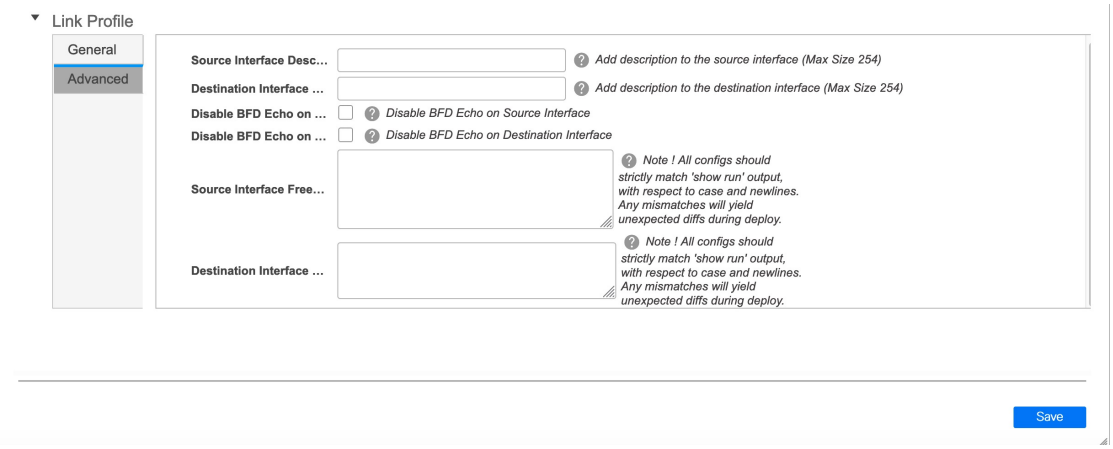

Source Interface Description and Destination Interface Description – Describe the links for later use. For example, if the link is between a leaf switch and a route reflector device, you can enter the information in these fields (Link from leaf switch to RR 1 and Link from RR 1 to leaf switch). This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

Disable BFD Echo on Source Interface and **Disable BFD Echo on Destination Interface** – Select the check box to disable BFD echo packets on source and destination interface.

Note that the BFD echo fields are applicable only when you have enabled BFD in the fabric settings.

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

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

   The new link appears in the Links tab.

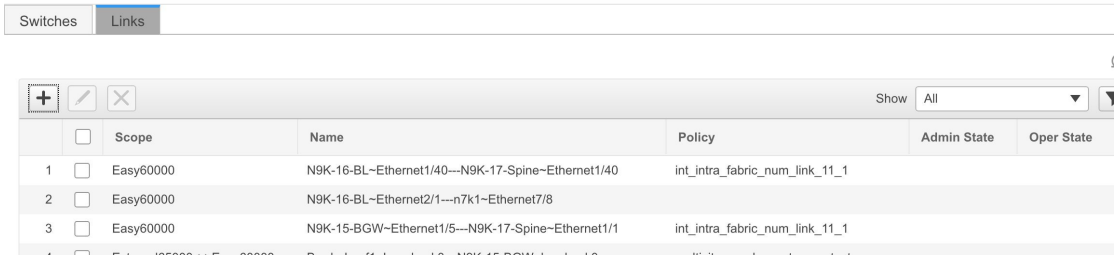

7. Click **Save & Deploy** to deploy the link configurations on the switches.

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

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

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

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

Creating Inter-Fabric Links

1. Click the Links tab in the Switches | Links page. The list of previously created links is 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).

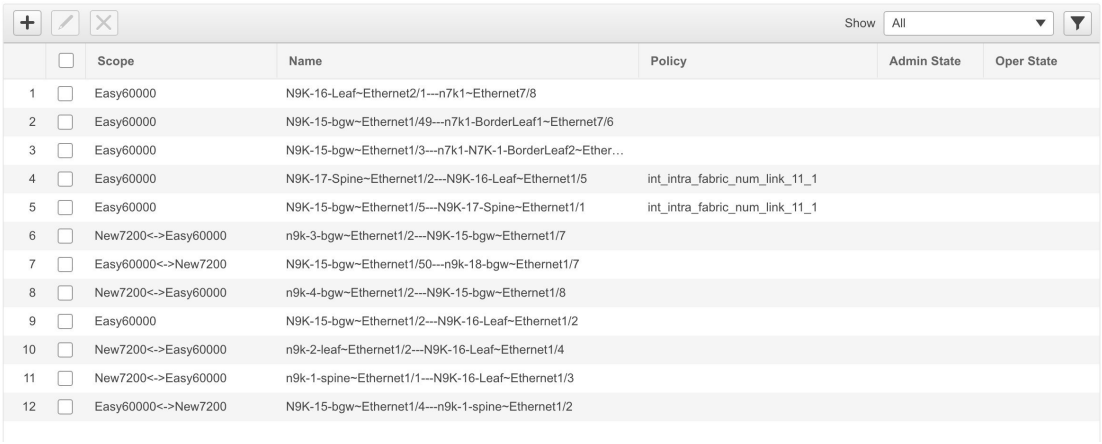

2. Click the Add (+) button at the top left part of the screen to add a link. The Add Link screen comes up. By default, the Intra-Fabric option is chosen as the link type.

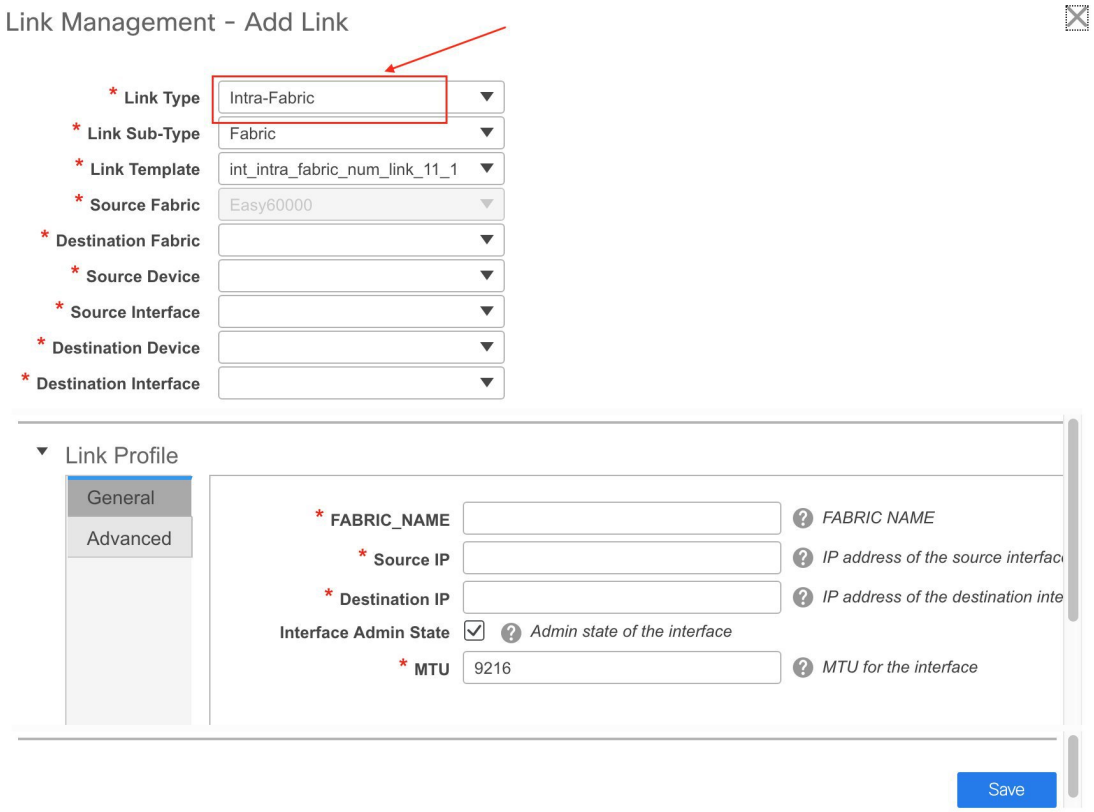

3. From the Link Type drop-down box, choose Inter-Fabric since you are creating an IFC. The screen changes correspondingly.

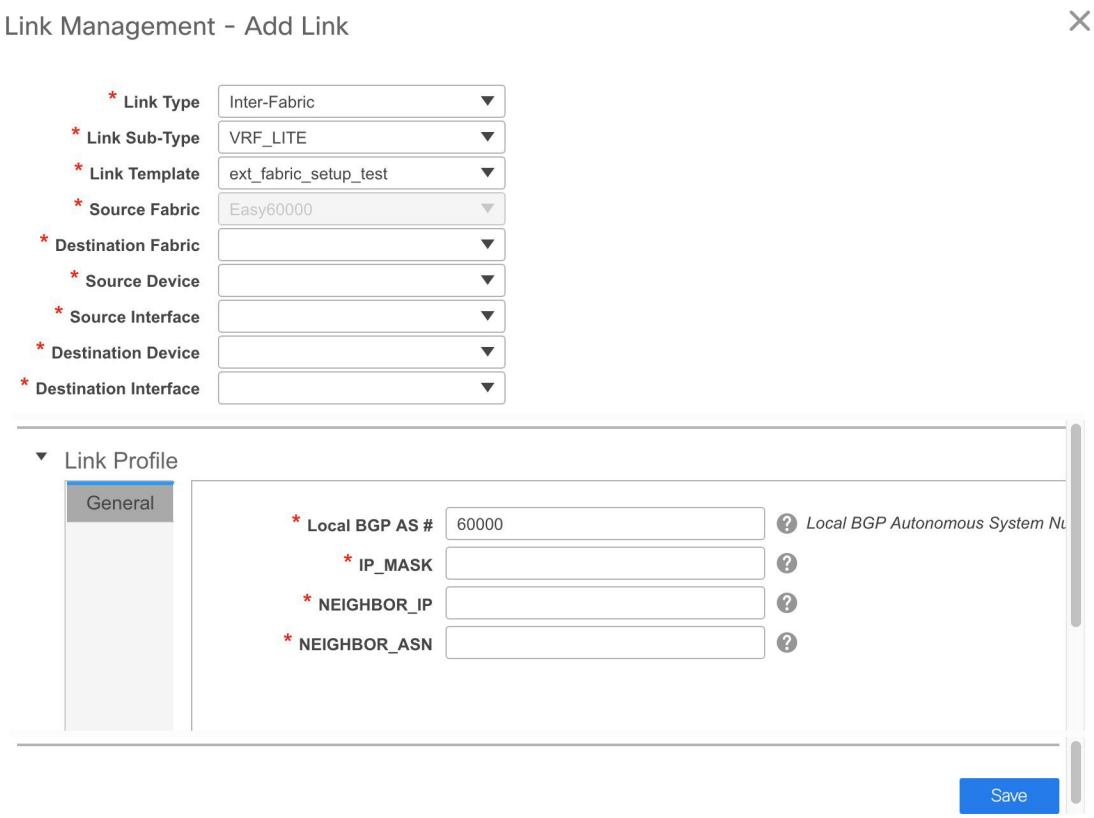

The fields for inter-fabric link creation are explained:

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

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

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

For information about VXLAN MPLS interconnection, see the *Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - MPLS SR and LDP Handoff* chapter.

For information about routed fabric interconnection, see the *Creating Inter-Fabric Links Between a Routed Fabric and an External Fabric* section in the *Configuring a Fabric with eBGP Underlay* chapter.

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 chapte more details.

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

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

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

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

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

General tab in the Link Profile section.

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

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

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

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

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

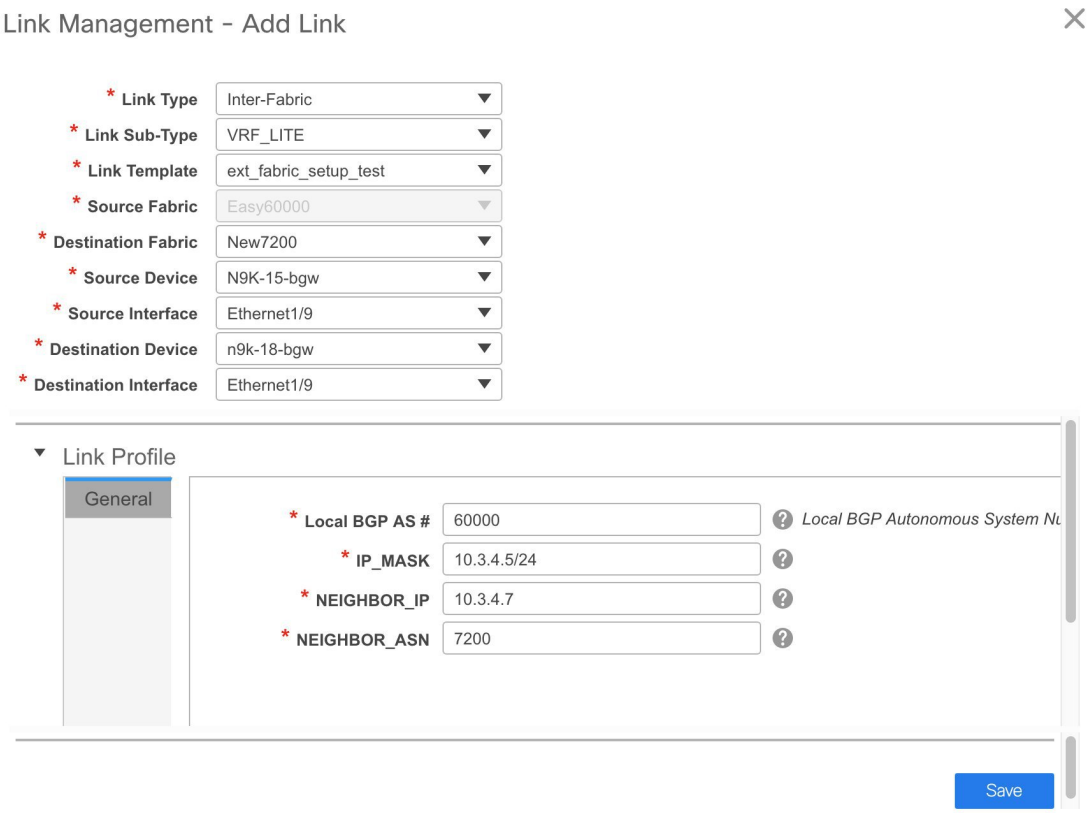

4. Click Save at the bottom right part of the screen.

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

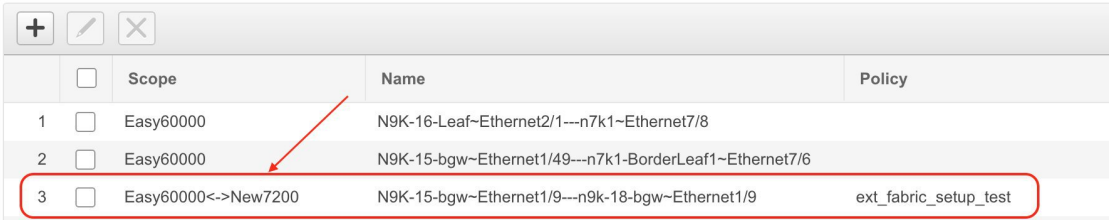

5. Click on Save & Deploy to deploy the link configurations on the switches.

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

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

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

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

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

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

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

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

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

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

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

Exporting Links

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

The fabric topology window appears.

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

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

3. Click the **Links** tab.

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

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

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

Symmetric Automatic VRF Lite

- Check the **Auto Deploy Flag** check box in the **Link Management** dialog box. Checking this check box enables VRF lite deployment on both ends of the link for managed devices.
- When you extend the VRF lite in a back-to-back scenario, the VRF should already be present in the peer fabric and the VRF name should be the same. An error message appears if the VRF is not present in the peer fabric and if you try to extend the VRF lite.
- When you extend the VRF lite between an easy fabric and an external fabric, the VRF name can be the same as that of the source fabric, default, or another VRF name. However, the child PTIs for the subinterface and the VRF creation or peering on the external fabric has the source. Hence, you cannot edit or delete the policies from the **View/Edit policies** window.
- If you perform a DCNM upgrade and notice that the policies are not attached to the IFC, edit the policies and VRF to attach them again.
- Besides the IPv6 address, enter the IP mask, IPv4 address, and the neighbor IP address as well to deploy VRF from topdown using symmetric VRF lite.
- Deploy configurations in both the fabrics.

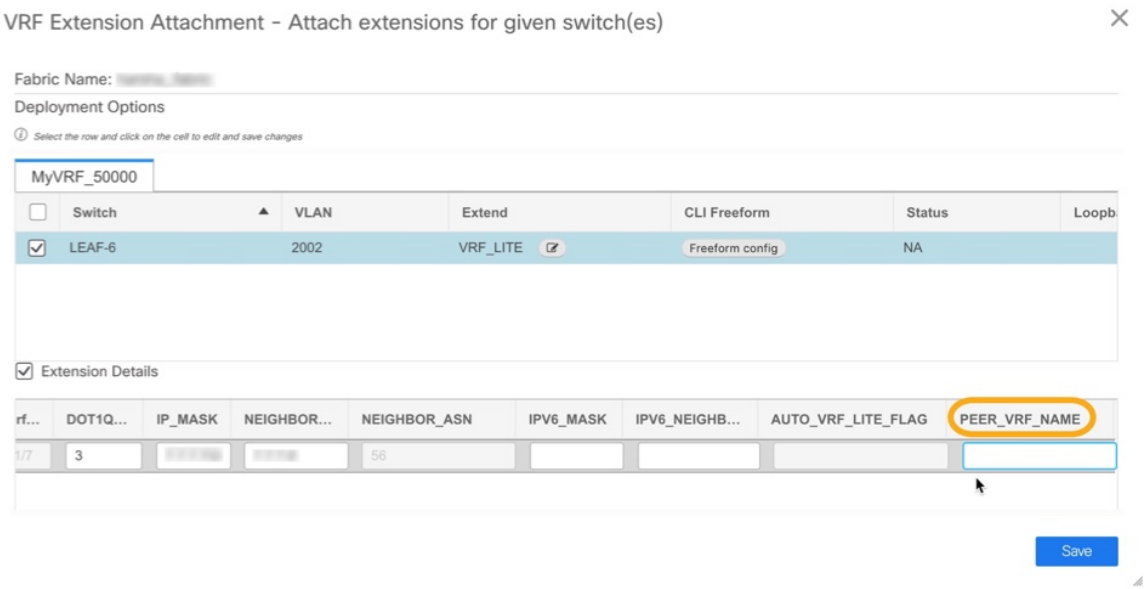

- You can edit or delete IFCs in the **Link** tab in the VXLAN fabric. The extra consideration for auto configured IFCs is that, in order to prevent the regeneration of IFC on next save and deploy, change the mode back to manual mode, or save the configuration only on the relevant devices.
- In a back-to-back 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 is automatically deleted from the external fabric.
- Deploy the configurations in both the fabrics.

See the *Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - VRF Lite* chapter for a use case on VRF Lite.

Layer 3 Port Channels

From Cisco DCNM Release 11.3(1), Layer 3 port channels are supported in external links and interfaces. In the **Interfaces** window, you can select a port channel and a corresponding Layer 3 port channel interface template. This template allows you to configure various options related to Layer 3 port channels including an ability to specify all Layer 3 interface-related configurations. Layer 3 port channels are supported only in easy fabrics and external fabrics.

External connectivity using VRF_LITE will also be supported using Layer 3 port-channels. For physical routed interfaces and LAYER 3 port channel interfaces, you can set the MTU.

You can also watch the video that demonstrates how to extend symmetric VRF Lite using Layer 3 port channels in Cisco DCNM. See the Extending Symmetric VRF Lite Using Layer 3 Port Channels video.

Configuring Layer 3 Port Channel on Interfaces

To configure a Layer 3 port channel on an interface from the Cisco DCNM Web UI, perform the following steps:

Procedure

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

The **Interfaces** window appears.

Step 2 Click **Add Interface**.

The **Add Interface** dialog box appears.

Step 3 Choose the **Port Channel** type and a device.

The port-channel ID is autopopulated.

Step 4 Choose the **int_l3_port_channel** policy.

The fields under the **General** area changes accordingly.

Step 5 Enter the values in the fields and click **Save**.

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, the **Resource could not be allocated** error appears.

Step 6 (Optional) Click the **Preview** option to preview the configurations to be deployed.

Step 7 Click **Deploy** to deploy the specified logical interface.

The newly added interface appears in the screen. You can break out and unbreakout an interface by using the breakout option at the top left.

Configuring Layer 3 Port Channel on Interfaces for IOS XE Devices

To configure a Layer 3 port channel on an interface for IOS XE devices, perform the following steps:

Procedure

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

The **Interfaces** window appears.

Step 2 Click **Add Interface**.

The **Add Interface** dialog box appears.

Step 3 Choose the **Port Channel** type and a device.

The port-channel ID is autopopulated.

Step 4 Choose the **ios_xe_int_l3_port_channel** policy.

The fields under the **General** area changes accordingly.

Step 5 Enter the values in the fields and click **Save**.

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, the **Resource could not be allocated** error appears.

Note The port-channel ID range for Cisco Catalyst 9000 Series switches is from 1 to 128 and for Cisco ASR 1000 Series routers the range is from 1 to 64.

Step 6 (Optional) Click the **Preview** option to preview the configurations to be deployed.

Step 7 Click **Deploy** to deploy the specified logical interface.

The newly added interface appears in the screen.

Deploying Policies on Physical Interfaces for non-Nexus Devices

More policies are added to support non-Nexus devices from Cisco DCNM, Release 11.4(1). When you import any non-Nexus device into an external fabric, few physical interfaces are created by default based on the number of ports. The policy is created only for the management port. For the Cisco Catalyst 9000 Series switches, the management port is GigabitEthernet0/0, and for the Cisco ASR 1000 Series routers, the management port is GigabitEthernet0.

The following table lists the policies added for different non-Nexus devices:

Devices	Policies
Cisco CSR 1000V Series Router	GigabitEthernet
Cisco IOS-XE Devices	• GigabitEthernet_mgmt • ios_xe_int_access_host • ios_xe_int_freeform • ios_xe_int_routed_host • ios_xe_int_trunk_host **Note** Use the GigabitEthernet_mgmt policy only for the management port, which is GigabitEthernet0/0.

To deploy policies on physical interfaces in the **Interfaces** window of Cisco DCNM Web UI, perform the following steps:

Before you begin

Import and discover non-Nexus devices into the external fabric. Ensure that the fabric isn't in monitor mode.

Procedure

Step 1 Check the check box of the interface on which you want to deploy the policy.

Step 2 Click the **Edit Configuration** icon.

Step 3 Choose a policy from the **Policy** drop-down list.

The valid options are:

- **GigabitEthernet**
- **GigabitEthernet_mgmt**
- **ios_xe_int_access_host**
- **ios_xe_int_freeform**
- **ios_xe_int_routed_host**
- **ios_xe_int_trunk_host**

Note
- Based on the option you choose, the fields under the **General** area vary.
- If you choose the **ios_xe_int_routed_host** policy, ensure you have configured the VRF manually, which is out-of-band, or using the **ios_xe_switch_freeform** policy in the **View/Edit Policies** window.
- DCNM doesn't support NVE or BDI interfaces. However, if you have already created them manually or out-of-band, use the **ios_xe_int_freeform** policy to define their configurations.

Step 4 Enter values for all the mandatory fields.

Note Choose the speed based on your device.

Step 5 Click **Save**.

Step 6 Click **Preview** to preview the pending configurations.

Step 7 Click **Deploy** to deploy the policy on the interface.

Configuring Layer 3 Port Channel on Subinterfaces

To configure a Layer 3 port channel on an interface from the Cisco DCNM Web UI, perform the following steps:

Procedure

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

The **Interfaces** window appears.

Step 2 Choose a Layer 3 port channel interface.

Step 3 Click **Add Interface**.

The **Add Interface** dialog box appears.

Step 4 Choose the **Subinterface** type.

The subinterface ID and policy are autopopulated, and the fields under the **General** area changes accordingly.

Step 5 Enter the values in the fields and click **Save**.

Only saved configurations are pushed to the device.

Step 6 (Optional) Click the **Preview** option to preview the configurations to be deployed.

Step 7 Click **Deploy** to deploy the specified logical interface.

A confirmation window appears, and the newly added subinterface appears in the list.

Configuring Layer 3 Port Channel for Inter-fabric Connectivity

To configure a Layer 3 port channel link from the **Fabric Builder** window, perform the following steps:

Before you begin

Ensure Layer 3 port channels are created on interfaces.

Procedure

Step 1 Choose an easy fabric or an external fabric, for which you want to extend the VRF-Lite.

The fabric topology window appears.

Step 2 Click **Tabular view** in the **Actions** pane.

All the components of this fabric are listed with their status and other details accordingly in different tabs.

Step 3 Choose the **Links** tab.

Step 4 Click the **Add Link** icon.

The **Add Link** dialog box appears.

Step 5 Choose **Inter-Fabric** link type.

Step 6 Choose **VRF_LITE** link sub-type.

Step 7 Choose the link template from the **Link Template** drop-down list.

Valid values are **ext_fabric_setup_11_1** and **service_link_trunk**.

Step 8 Enter the details for all other fields accordingly.

Step 9 Enter the details for fields in the **Link Profile** area wherever necessary.

You can set the MTU. The **Ext_VRF_Lite_Jython** auto-deploy template is used for VRF-Lite configuration on the device in the fabric.

Step 10 Click **Save**.

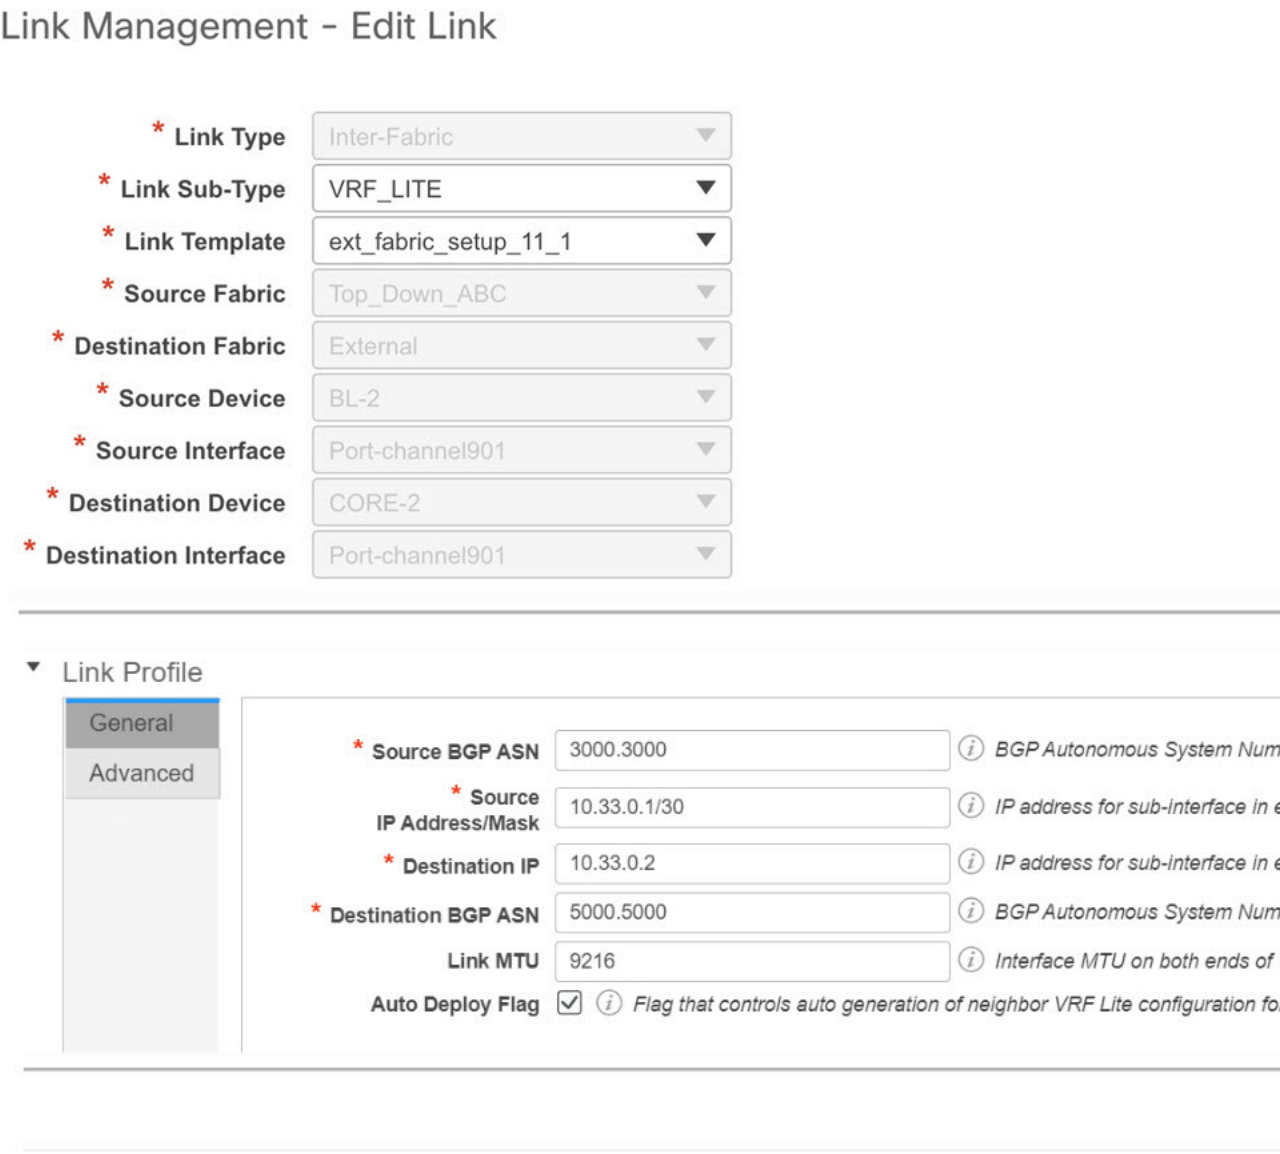

What to do next

After creating a VRF Lite IFC with the Layer 3 port-channel, using the top-down flow, when a VRF is extended using VRF Lite, a sub-interface is created on the Layer 3 port-channel. You can edit the Layer 3 port channel links even after VRFs are extended. However, Layer 3 port channels are not supported for intra-fabric links.

Tabular View - Operational View

From Cisco DCNM 11.3(1), the operational support for a fabric is provided. This feature provides the following information:

- Operational status of a fabric
- Alarm and event notifications

You can view the operational status information in the **Operational View** tab. You can view the alarm and event notifications by clicking the **Alerts and Notifications** icon, next to the **Help** icon, in the top pane of Cisco DCNM.

Viewing the Operational Status

To view the operational status of a fabric from the **Fabric Builder** window, perform the following steps:

Procedure

Step 1 Choose a fabric.

The fabric topology window appears.

Step 2 Click **Tabular view** in the **Actions** pane.

Step 3 Choose the **Operational View** tab.

The Operational View tab has the following fields and descriptions.

Fields	Descriptions
Fabric Name	Specifies the fabrics that have links.
Name	Specifies the link name.
Is Present	Specifies if the link is present or not. Valid values are **true** and **false**.
Link State	Specifies the status of the logical link. A logical link can be in one of the following states. • **Established**: When a link is in the **Established** state the peers send update messages to exchange information about each route advertised to the BGP peer. A notification is sent if there is an error and the state changes to **Idle**. Only a link using the BGP routing protocol can be in the **Established** state. • **Idle**: A link using BGP protocol will be in Idle state when there is an error between peers. • **UP**: A link using ISIS protocol will be in the **UP** state, when the link is successfully established between peers. • **FULL**: A link using the OSPF protocol will be in the **FULL** state when the link is successfully established between peers. • **peer-alive**: Specifies the link as a peer keepalive link that monitors the vitality of a vPC peer switch.

Fields	Descriptions
Link Type	Specifies the type of logical link. The link can be of the following type: • **BGP** • **ISIS** • **OSPF** • **VPC_KEEPALIVE**
Uptime	Specifies the duration of the uptime for the link type.

All these columns are sortable.

Viewing Logical Links

The logical links appear in the **Topology** window. To view the logical links from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Topology**.

The **Topology** window appears.

Step 2 Check the **Logical Links** check box in the Show pane.

The logical links between devices appear in blue color.

Note The color of the link will change based on its state.

Step 3 (Optional) Hover over the link to know the link type.

Viewing Alerts and Event Notifications

Alert and event notifications includes health score, topology node display, alarm view, alarm policies, and notification services. An event is any action that impacts network, devices or Cisco DCNM. An alert is a notification that is triggered as part of an event to make it visible.

Support for ToR Switches

From Cisco DCNM 11.3(1), support for the Top-of-Rack (ToR) switches is added in DCNM. You can add the Layer 2 ToR switches in an external fabric, and they can be connected to the Leaf switches in the Easy Fabric. For more information, see *Configuring ToR Switches and Deploying Networks*.

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). However, both greenfield as well as brownfield deployments support vPC fabric peering in Cisco DCNM, Release 11.3(1). This feature is applicable for **Easy_Fabric_11_1** and **Easy_Fabric_eBGP** fabric templates.

Note The **Easy_Fabric_eBGP** fabric does not support brownfield import.

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, and FX2 support vPC fabric peering.
- From Cisco DCNM, Release 11.4(1), Cisco Nexus N9K-C93180YC-FX3S and N9K-C93108TC-FX3P platform switches support vPC fabric peering.
- Cisco Nexus 9300-EX, and 9300-FX/FXP/FX2/FX3 platform switches support vPC Fabric Peering. Cisco Nexus 9200 and 9500 platform switches do not 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.
- Brownfield deployments and greenfield deployments support vPC fabric peering in Cisco DCNM, Release 11.3(1).
- 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.

QoS for Fabric vPC-Peering

From Cisco DCNM Release 11.4(1), in the **Easy_Fabric_11_1** fabric settings, you can enable QoS on spines for guaranteed delivery of vPC Fabric Peering communication. Additionally, you can specify the QoS policy name.

Note the following guidelines for a greenfield deployment:

- If QoS is enabled and the fabric is newly created:
  - If spines or super spines neighbor is a virtual vPC, make sure neighbor is not honored from invalid links, for example, super spine to leaf or borders to spine when super spine is present.
  - Based on the Cisco Nexus 9000 Series Switch model, create the recommended global QoS config using the **switch_freeform** policy template.
  - Enable QoS on fabric links from spine to the correct neighbor.
- If the QoS policy name is edited, make sure policy name change is honored everywhere, that is, global and links.
- If QoS is disabled, delete all configuration related to QoS fabric vPC peering.
- If there is no change, then honor the existing PTI.

For more information about a greenfield deployment, see the *Creating a New VXLAN BGP EVPN Fabric* section.

Note the following guidelines for a brownfield deployment:

Brownfield Scenario 1:

- If QoS is enabled and the policy name is specified:

Note You need to enable only when the policy name for the global QoS and neighbor link service policy is same for all the fabric vPC peering connected spines.

  - Capture the QoS config from switch based on the policy name and filter it from unaccounted configuration based on the policy name and put the configuration in the **switch_freeform** with PTI description.
  - Create service policy configuration for the fabric interfaces as well.
  - Greenfield config should make sure to honor the brownfield config.
- If the QoS policy name is edited, delete the existing policies and brownfield extra configuration as well, and follow the greenfield flow with the recommended config.
- If QoS is disabled, delete all the configuration related to QoS fabric vPC peering.

Note No cross check for possible or error mismatch user configuration, and user might see the diff.

Brownfield Scenario 2:

- If QoS is enabled and the policy name is not specified, QoS configuration is part of the unaccounted switch freeform config.
- If QoS is enabled from fabric settings after **Save & Deploy** for brownfield, QoS configuration overlaps and you will see the diff if fabric vPC peering config is already present.

For more information about a brownfield deployment, see the *Creating a New VXLAN BGP EVPN Fabric* section..

Fields and Description

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

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

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

Creating a Virtual Peer Link

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

Procedure

Step 1 Choose **Control > Fabrics**.

The **Fabric Builder** window appears.

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

The fabric topology window appears.

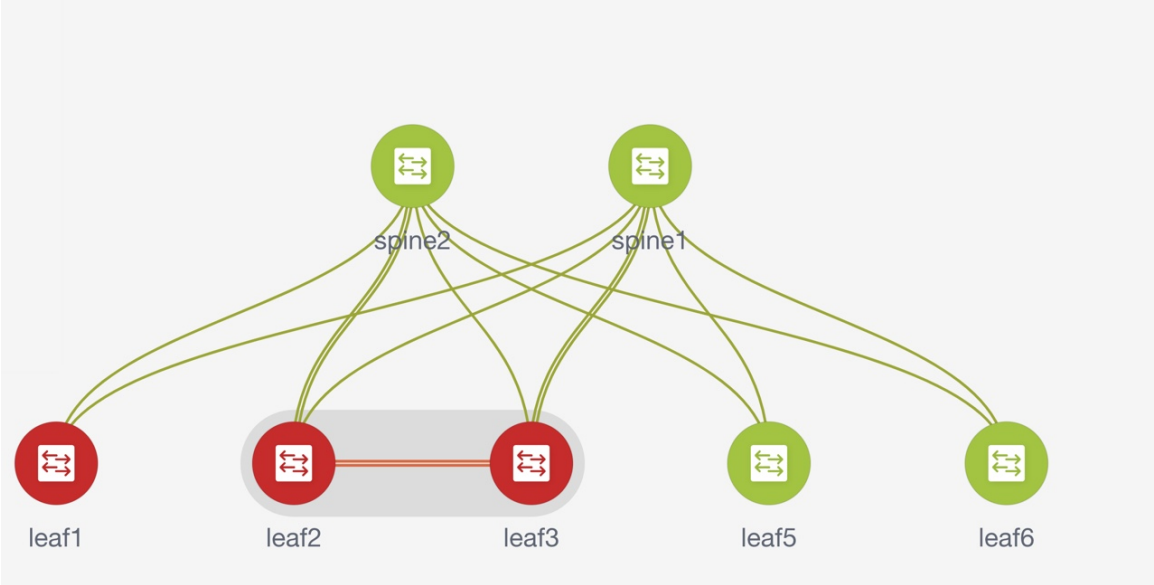

Step 3 Right-click a switch and choose **vPC Pairing** from the drop-down list.

The window to choose the peer appears.

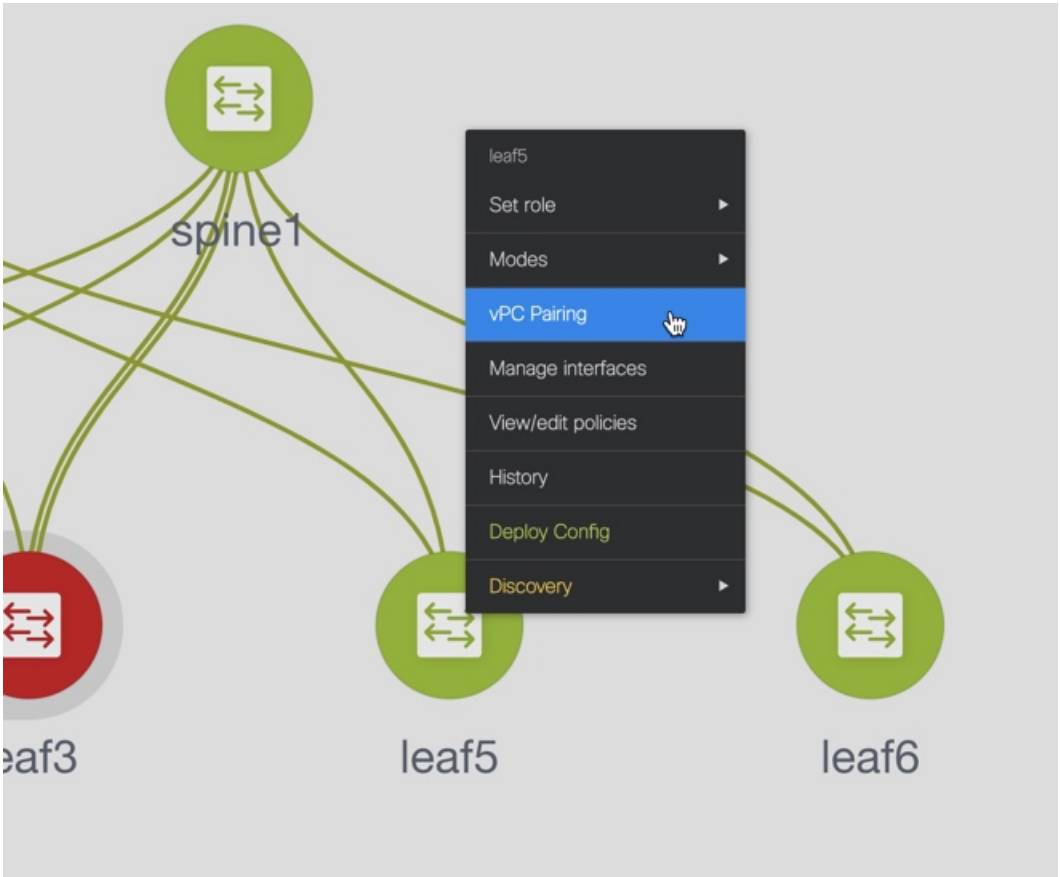

Note Alternatively, you can also navigate to the **Tabular view** from the **Actions** pane. Choose a switch in the **Switches** tab and click **vPC Pairing** to create, edit, or unpair a vPC pair. However, you can use this option only when you choose a Cisco Nexus switch.

You will get the following error when you choose a switch with the border gateway leaf role.

```
<switch-name> has a Network/VRF attached. Please detach the Network/VRF before vPC
Pairing/Unpairing
```

Step 4 Check the **Use Virtual Peerlink** check box.

Step 5 Choose a peer switch and check the **Recommended** column to see if pairing is possible.

If the value is **true**, pairing is possible. You can pair switches even if the recommendation is **false**. However, you will get a warning or error during **Save & Deploy**.

Step 6 Click **Save**.

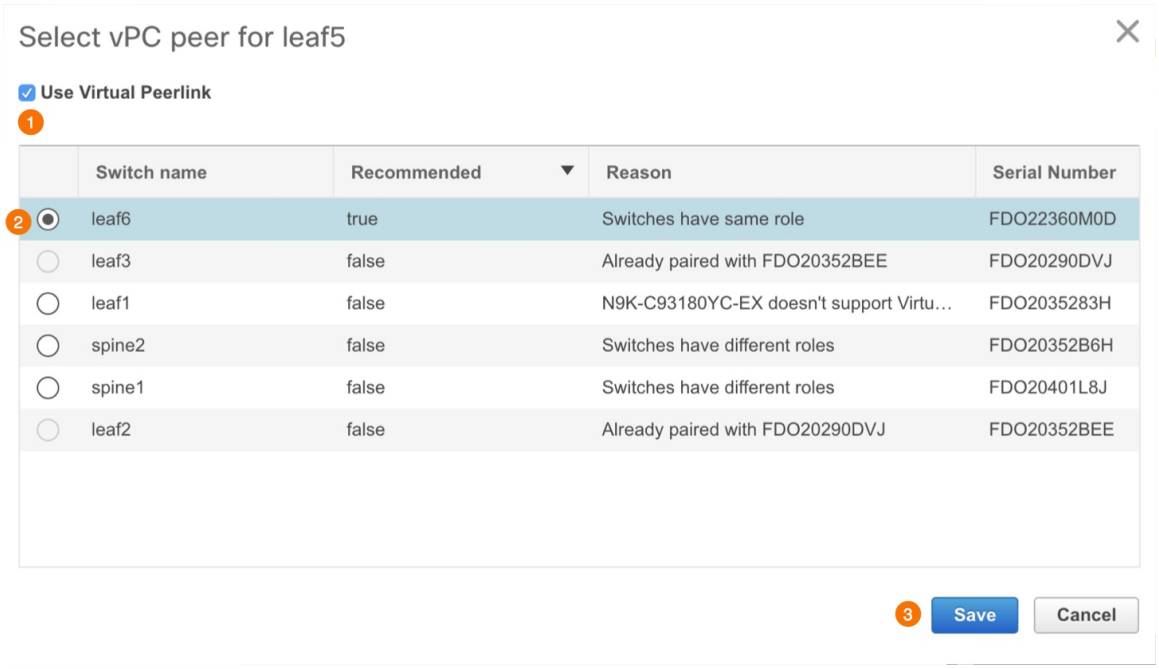

Step 7 In the **Fabric Topology** window, click **Save & Deploy**.

The **Config Deployment** window appears.

Step 8 Click the field against the switch in the **Preview Config** column.

The **Config Preview** window appears for the switch.

Step 9 View the vPC link details in the pending configuration and the side-by-side configuration.

Step 10 Close the window.

Step 11 Click the pending errors icon next to the **Save & Deploy** icon to view errors and warnings, if any.

If you see any warnings that are related to TCAM, click the **Resolve** icon. A confirmation dialog box about reloading switches appears. Click **OK**. You can also reload the switches from **Tabular view** in the fabric topology window.

The switches that are connected through vPC fabric peering, are enclosed in a gray cloud.

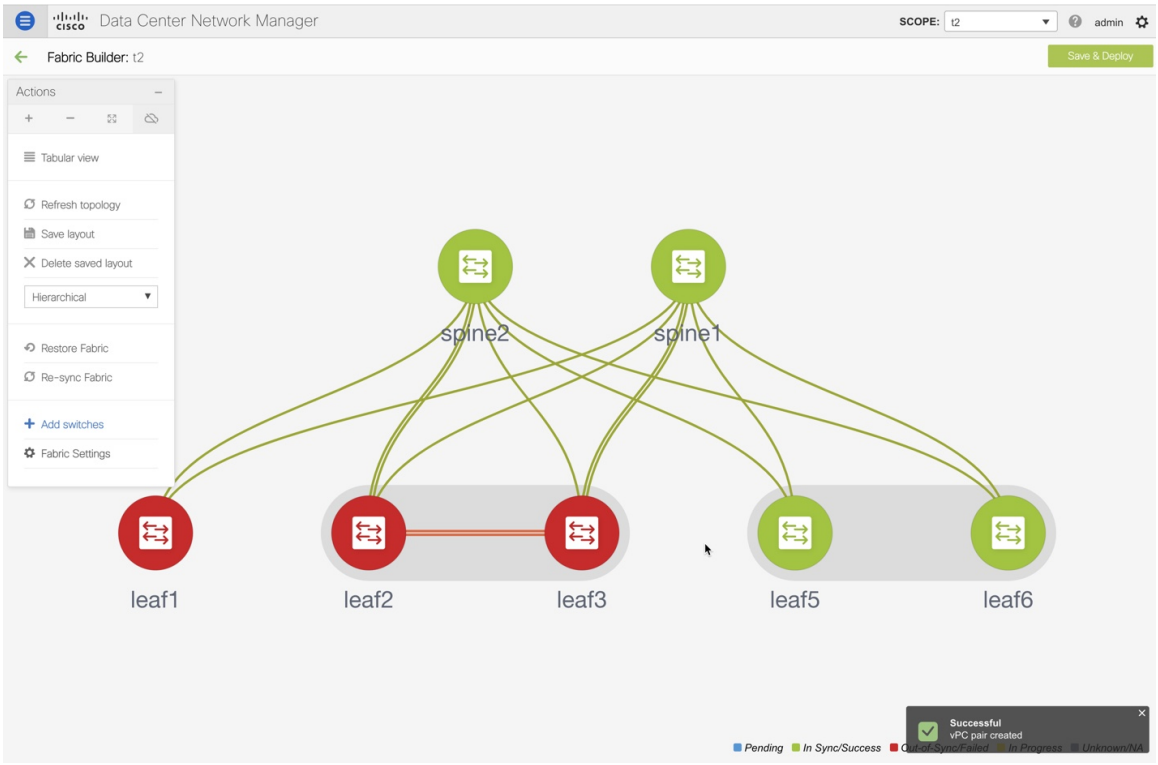

Converting a Physical Peer Link to a Virtual Peer Link

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

Before you begin

- Plan the conversion from physical peer link to virtual peer link during the maintenance window of switches.
- Ensure the switches support vPC fabric peering. Only the following switches support vPC fabric peering:
  - Cisco Nexus N9K-C9332C Switch, Cisco Nexus N9K-C9364C Switch, and Cisco Nexus N9K-C9348GC-FXP Switch
  - Cisco Nexus 9000 Series Switches that ends with FX, FX2, and FX2-Z

Procedure

Step 1 Choose **Control > Fabrics**.

The **Fabric Builder** window appears.

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

Step 3 Right-click the switch that is connected using the physical peer link and choose **vPC Pairing** from the drop-down list.

The window to choose the peer appears.

Note Alternatively, you can also navigate to the **Tabular view** from the **Actions** pane. Choose a switch in the **Switches** tab and click **vPC Pairing** to create, edit, or unpair a vPC pair. However, you can use this option only when you choose a Cisco Nexus switch.

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.

Note Alternatively, you can also navigate to the **Tabular view** from the **Actions** pane. Choose a switch in the **Switches** tab and click **vPC Pairing** to create, edit, or unpair a vPC pair. However, you can use this option only when you choose a Cisco Nexus switch.

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.

Advertising PIP on vPC

In the fabric settings, you can check the **vPC advertise-pip** check box to enable the Advertise PIP feature on all vPCs in a fabric. From Cisco DCNM Release 11.4(1), you can use the **vpc_advertise_pip_jython** policy to enable Advertise PIP feature on specific vPCs in a fabric.

Note the following guidelines:

- If advertise-pip is not globally enabled or vPC peer is not using fabric peering, only then the vpc_advertise_pip_jython policy can be created on specific peers.
- Enabling **vpc advertise-pip** doesn't affect the current behavior.
- Disabling advertise pip for a fabric doesn't affect this policy.
- Unpairing of switches deletes this policy.
- You can manually delete this policy from the peer switch where it was created.

Procedure

Step 1 From the **Fabric Builder** window, click a fabric, and then right-click on a switch with vPC and select **View/Edit Policies**.

Step 2 Click **Add** and select the **vpc_advertise_pip_jython** policy template and enter the mandatory parameters data.

Note You can add this policy on one vPC peer, and it will create respective commands for vpc advertise on both peers.

Step 3 Click **Save**, and then deploy this policy.

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.

Cisco DCNM Release 11.5(3a) allows you to establish eBGP underlay connectivity between two fabrics which can be managed by the same or different DCNM appliances. A new policy template **ebgp_underlay_dci** establishes ebgp connectivity between device in one fabric (for example - border gateway in Easy Fabric) with another device in a different fabric (for example - core router in External Fabric) to exchange underlay routes.

While configuring the policy using this template, specify the neighbor IP address and the BGP ASN to establish the BGP session. You can also override Local ASN for Device in External Fabric.

You must create this policy for each interface. You must create a policy using the same **ebgp_underlay_dci** template on the other fabric.

After configuring the policy on both the interfaces on the links, you can review the intent configuration and deploy to the switches. The generated intent is pushed to establish eBGP session to exchange routes.

Use-case: In a Multi-Site fabric, you can add fabrics as members. For multiple fabrics within the Multi-Site to exchange routes with eachother, you must configure both multi-site overlay IFC and multi-site underlay IFC. You can create an inter-fabric link to establish connectivity. If the fabrics are not in the same DCNM instance, **ebgp_underlay-dci** policy template allows you to establish multi-site underlay connectivities.

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**.

Note **View/Edit Policies** is not enabled for an MSD fabric.

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.

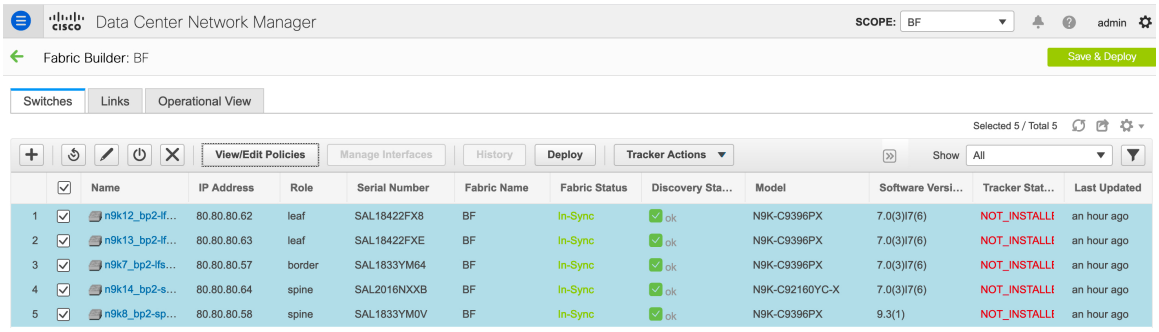

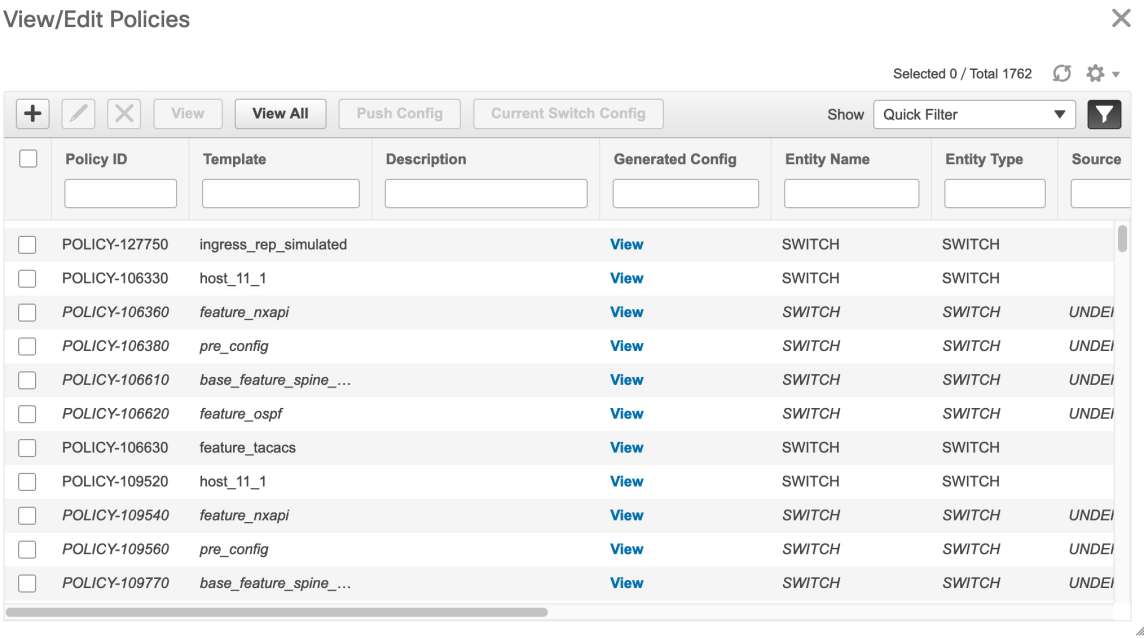

Note You can view the generated config for a device by hovering over the **View** button under the **Generated Config** column. Additionally, you can enter a config in the search field under this column to filter policies.

Step 4 Select a policy and click the **View** button to view its configs.

Note Python policies are used to place logic and control CLI policies. From DCNM Release 11.3(1), multiple CLI child policies are aggregated for each Python policy.

Step 5 In the **View/Edit Policies** window, click **View All** to view all the configurations pushed to the switches using policies.

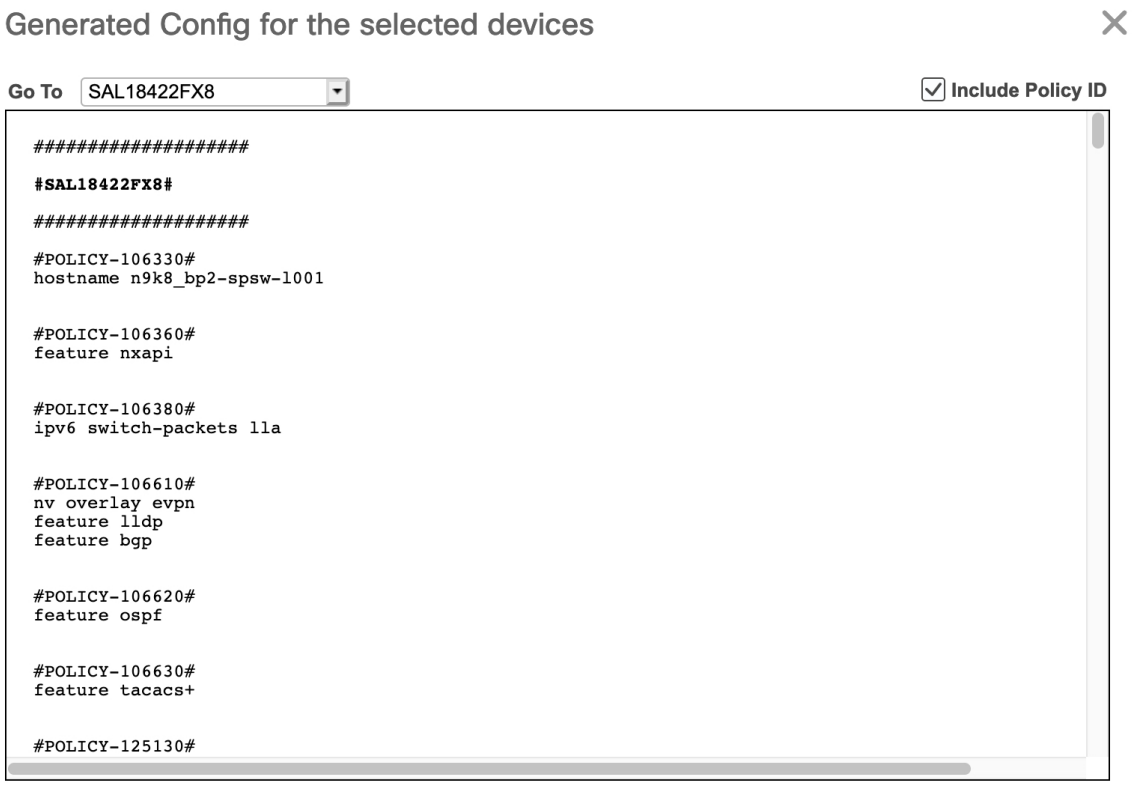

Go To: Select a device from this drop-down list to navigate to its starting config.

This option is applicable only when you view policies for multiple devices.

Include Policy ID: Select this check box to view policy IDs for all the policies. By default, this check box is selected.

Adding a Policy

Procedure

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

Step 2 Select any available fabric, and then click Tabular view.

Step 3 Select a single or multiple switches in the **Switches** tab, and click the **View/Edit Policies** button.

Step 4 Click the **Add** icon.

Step 5 Select a policy template and enter the mandatory parameters data and click **Save**. PTI is added per each device based on n-number of devices selection.

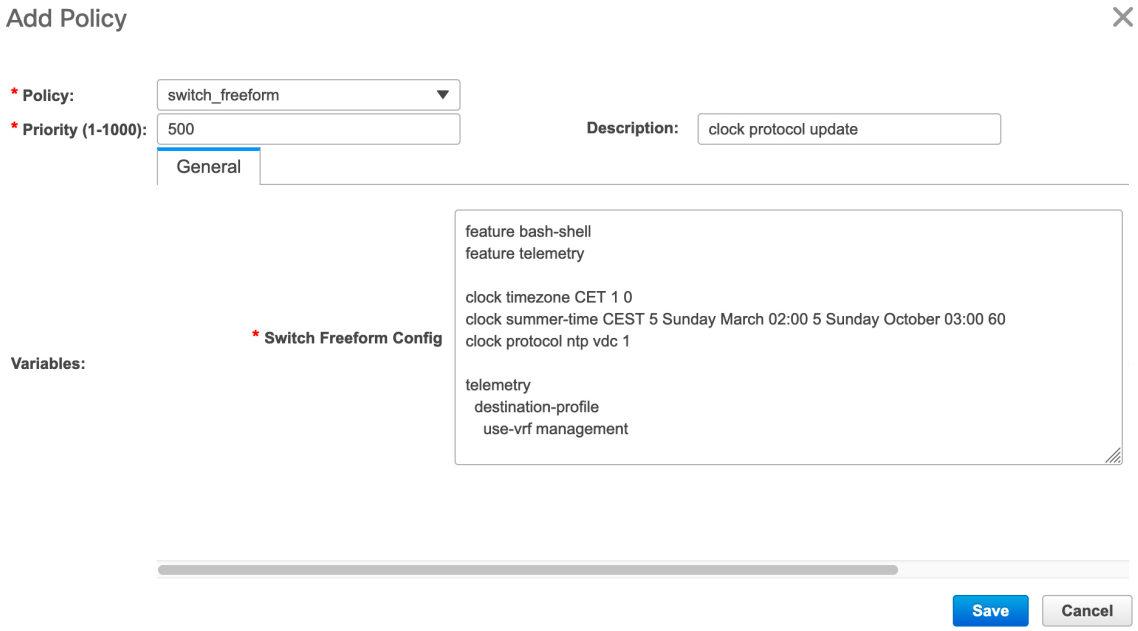

Policy: Select a policy from this drop-down list.

Priority: Specify a priority for the policy. The applicable values are from 1 to 1000. The default value is 500. The lower number in the **Priority** field means that there is a higher priority for the generated configuration and POAP startup-configuration. For example, features are 50, route-maps are 100, and vpc-domain is 200.

Description: (Optional) Specify a description for the policy. This field is used to differentiate multiple freeform policies. The **Description** column is added in the **View/Edit Policies** window, which you can use to filter or find policies based on description.

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.

- If the external fabric is in the monitor mode, the **Push Config** option is disabled.
- This option will be greyed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.

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.

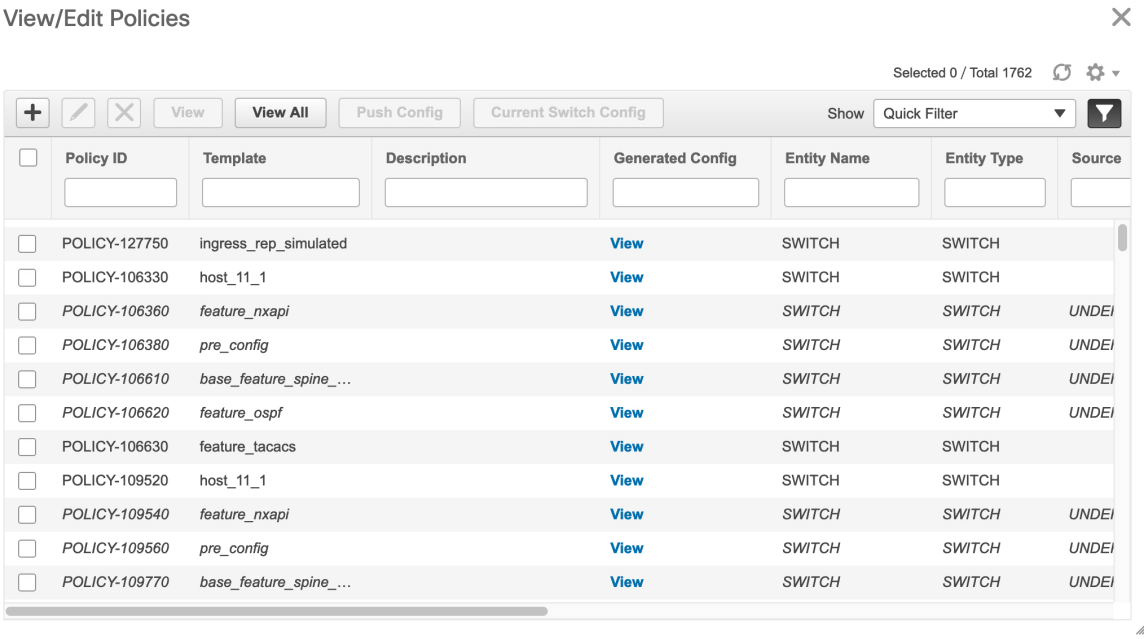

Note The policies in the italics font cannot be edited. The value under the **Editable** and **Mark Deleted** columns for these policies is **false**.

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

- This option will be greyed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.
- A warning appears if you push config for a Python policy.
- A warning appears if you edit, delete, or push config a mark-deleted policy. A mark-deleted policy is set to **true** under the **Mark Deleted** column. The switch freeform child policies of **Mark Deleted** policies appears in the **View/Edit Policies** dialog box. You can edit only **Python** switch_freeform policies. You cannot edit **Template_CLI** switch_freeform_config policies.

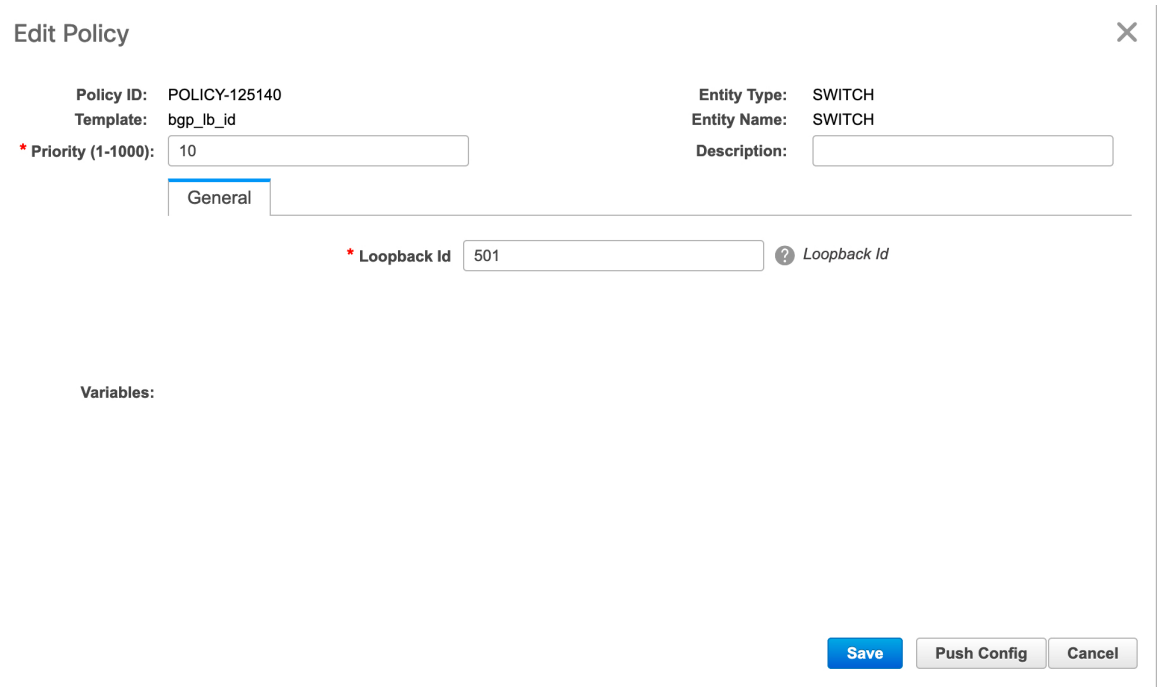

Current Switch Configuration

Procedure

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

Step 2 Select any available fabric, and then click **Tabular** view.

Step 3 Select multiple switches in the switches tab, and click **View/Edit Policies**.

Step 4 Click **Current Switch Config**.

The current switch configuration appears in the **Running Config** dialog box.

Note The running configuration will not appear for the Cisco CSR 1000v when you click **Current Switch Config** if the user role cannot access the enable prompt by default.

Retrieving the Authentication Key

Retrieving the 3DES Encrypted OSPF Authentication Key

1. SSH into the switch.
2. On an unused switch interface, enable the following:

```
config terminal
    feature ospf
    interface Ethernet1/1
       no switchport
       ip ospf message-digest-key 127 md5 ospfAuth
```

In the example, **ospfAuth** is the unencrypted password.

Note This Step 2 is needed when you want to configure a new key.

3. Enter the **show run interface Ethernet1/1** command to retrieve the password.

```
Switch # show run interface Ethernet1/1
   interface Ethernet1/1
     no switchport
     ip ospf message-digest key 127 md5 3 sd8478f4fsw4f4w34sd8478fsdfw
     no shutdown
```

The sequence of characters after **md5 3** is the encrypted password.

4. Update the encrypted password into the **OSPF Authentication Key** field.

Retrieving the Encrypted IS-IS Authentication Key

To get the key, you must have access to the switch.

1. SSH into the switch.
2. Create a temporary keychain.

```
config terminal
   key chain isis
   key 127
   key-string isisAuth
```

In the example, **isisAuth** is the plaintext password. This will get converted to a Cisco type 7 password after the CLI is accepted.

3. Enter the **show run | section "key chain"** command to retrieve the password.

```
key chain isis
   key 127
       key-string 7 071b245f5a
```

The sequence of characters after key-string 7 is the encrypted password. Save it.

4. Update the encrypted password into the ISIS Authentication Key field.
5. Remove any unwanted configuration made in Step 2.

Retrieving the 3DES Encrypted BGP Authentication Key

1. SSH into the switch and enable BGP configuration for a non-existent neighbor.

Note Non-existent neighbor configuration is a temporary BGP neighbor configuration for retrieving the passwor

```
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.

Retrieving the Encrypted BFD Authentication Key

1. SSH into the switch.
2. On an unused switch interface, enable the following:

```
switch# config terminal
switch(config)# int e1/1
switch(config-if)# bfd authentication keyed-SHA1 key-id 100 key cisco123
```

In the example, **cisco123** is the unencrypted password and the key ID is **100**.

Note This Step 2 is needed when you want to configure a new key.

3. Enter the **show running-config interface** command to retrieve the key.

```
switch# show running-config interface Ethernet1/1

interface Ethernet1/1
description connected-to- switch-Ethernet1/1
no switchport
mtu 9216
bfd authentication Keyed-SHA1 key-id 100 hex-key 636973636F313233
no ip redirects
ip address 10.4.0.6/30
no ipv6 redirects
ip ospf network point-to-point
ip router ospf 100 area 0.0.0.0
no shutdown
```

The BFD key ID is **100** and the encrypted key is **636973636F313233**.

4. Update the key ID and key in the **BFD Authentication Key ID** and **BFD Authentication Key** fields.

Custom Maintenance Mode Profile Policy

When you place a switch in maintenance mode using DCNM, only a fixed set of BGP and OSPF isolate CLIs are configured in the maintenance mode profile. Starting from Cisco DCNM Release 11.3(1), you can create a **custom_maintenance_mode_profile** PTI with customized configurations for maintenance mode and normal mode profile, deploy the PTI to the switch, and then move the switch to maintenance mode.

Creating and Deploying a Custom Maintenance Mode Profile Policy

Procedure

Step 1 Select **Control>Fabric Builder**, click **Tabular View**, and select a switch in the **Name** column or select **Control>Fabric Builder** and right-click the switch.

Step 2 Click **View/Edit Policies** and click on **+** to add a new policy. The **Add Policy** window comes up.

Step 3 Select **custom_maintenance_mode_profile** from the **Policy** dropdown list.

Step 4 Fill in the **Maintenance mode profile contents** with the desired configuration CLIs.

Example:

```
configure maintenance profile maintenance-mode
ip pim isolate
```

Fill in the **Normal mode profile contents** with the desired configuration CLIs.

Example:

```
configure maintenance profile normal-mode
no ip pim isolate
```

Step 5 Click **Save**.

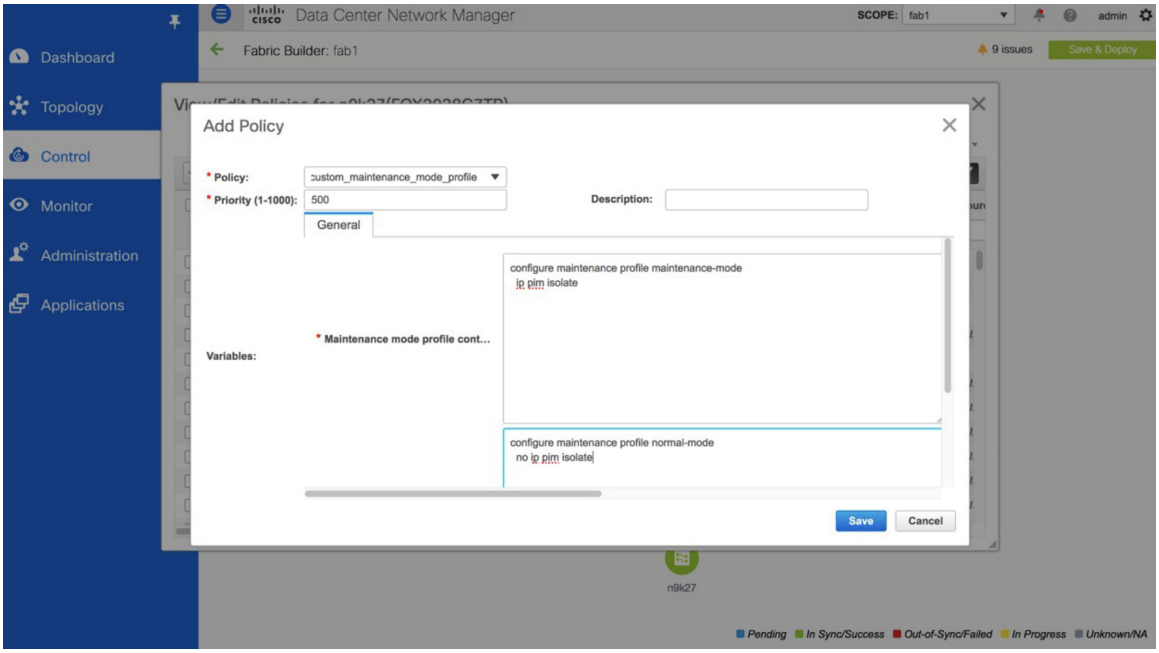

Step 6 Right-click the switch in the **Fabric Builder** window and select **Deploy Config**. Verify the configuration in the **Pending Config** window and then deploy the configuration to the switch.

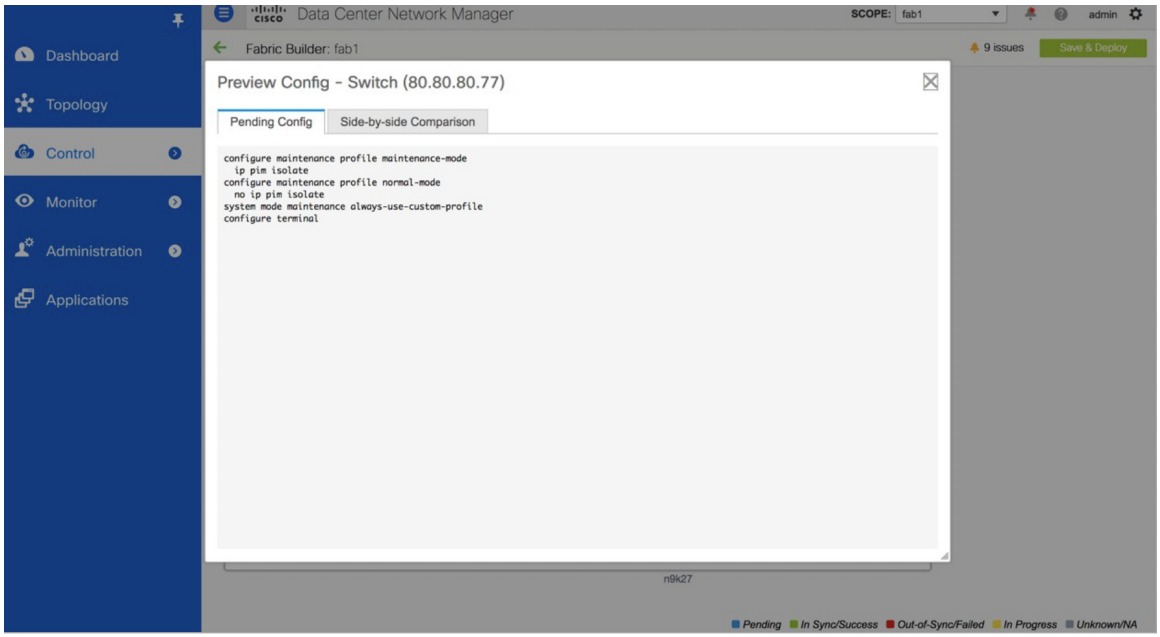

Step 7 Then, right-click the switch and select **Modes>Maintenance Mode** to move the switch to maintenance mode.

Deleting a Custom Maintenance Mode Profile Policy

Procedure

Step 1 The switch has to be moved to active/operational or normal mode before deleting the custom maintenance mode profile policy. To do this, right-click the switch in the **Fabric Builder** window and select **Modes>Active/Operational Mode**.

Step 2 After the switch has been moved to active/operational or normal mode, click **Tabular View** in the **Fabric Builder** window, and select the switch in the **Name** column or right-click the switch in the **Fabric Builder** window.

Step 3 Click **View/Edit Policies**, and select the **custom_maintenance_mode_profile** policy that has to be deleted.

Step 4 Click **X** to delete the policy.

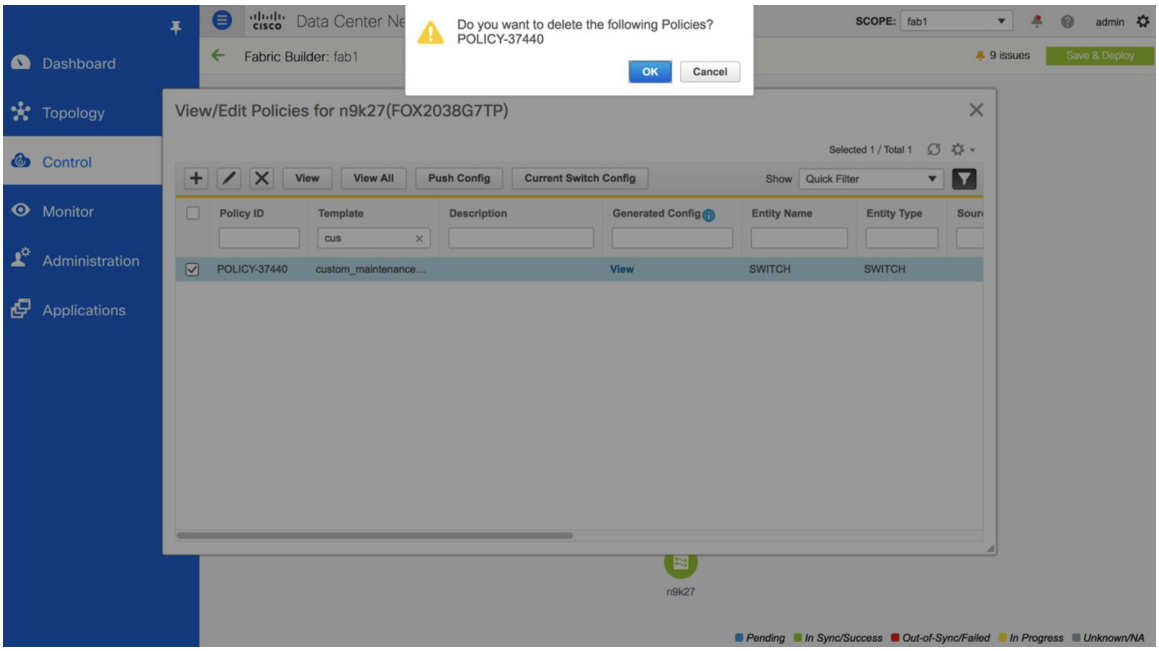

Step 5 Right-click the switch in the **Fabric Builder** window and select **Deploy Config**. Verify the configuration in the **Pending Config** window and deploy the configuration to the switch.

```
no system mode maintenance always-use-custom-profile
no configure maintenance profile normal-mode
no configure maintenance profile maintenance-mode
configure terminal
```

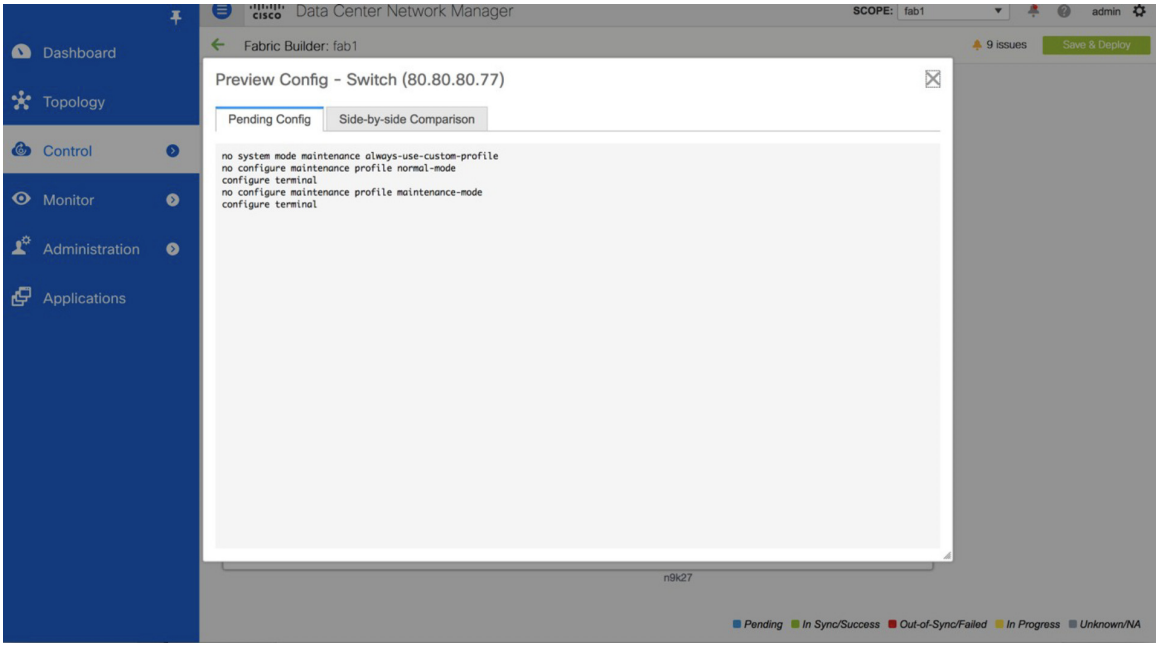

Return Material Authorization (RMA)

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

Prerequisites

- Ensure that the fabric is up and running with minimal disruption while replacing the switch.
- To use the POAP RMA flow, configure the fabric for bootstrap (POAP).
- Perform save and deploy more than once, if needed, to copy the FEX configurations for the RMA of switches that have FEX deployed.

Guidelines and Limitations

- To replace the switch, remove the old switch from the fabric and discover the new switch in the fabric. For example, if you want to replace a Cisco Nexus 9300-EX switch with a Cisco Nexus 9300-FX switch, remove the 9300-EX switch from the fabric followed by discovering the 9300-FX switch in the same fabric.
- When GIR is enabled before upgrading Cisco Nexus 7000 Series switches, DCNM pushes the **system mode maintenance** command to the switches when the DCNM RMA procedure is initiated. This command applies the configuration that is present in the default maintenance mode profile to the switches. For more information on performing Graceful Insertion and Removal (GIR) on the Cisco Nexus 7000 Series switches, refer Configuring GIR.

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**.

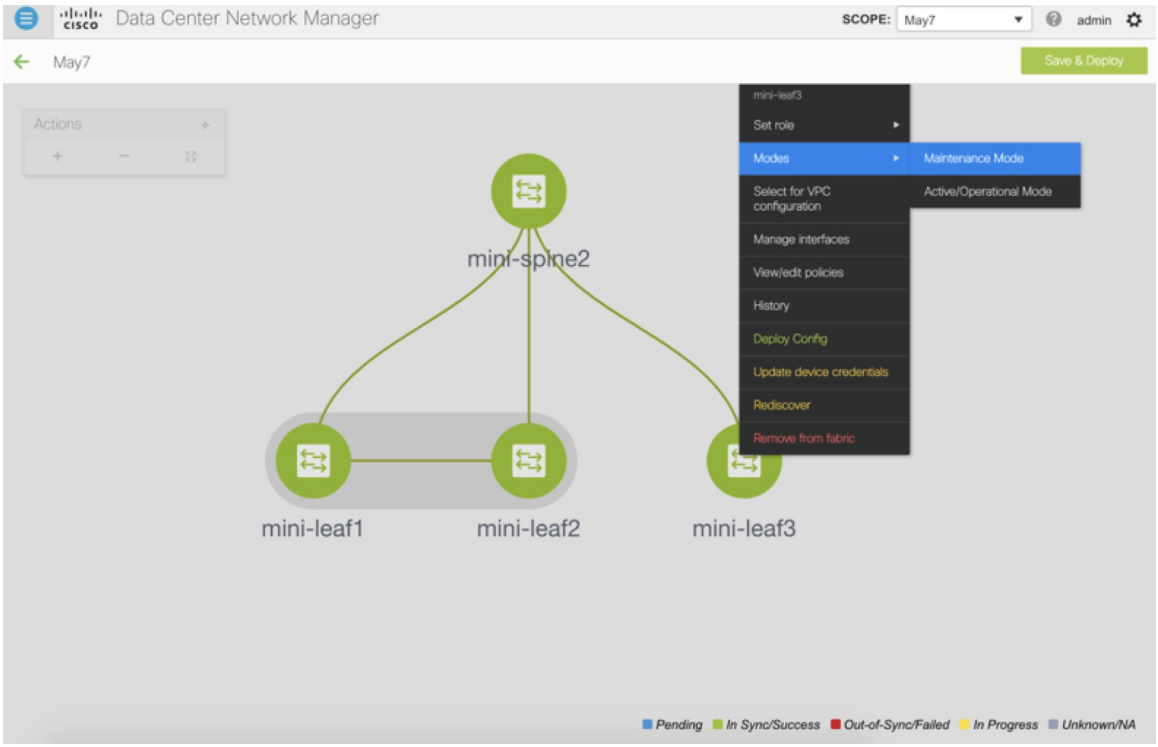

Step 4 Physically replace the device in the network. Physical connections should be made in the same place on the replacement switch as they existed on the original switch.

Step 5 Provision RMA flow and select the replacement device.

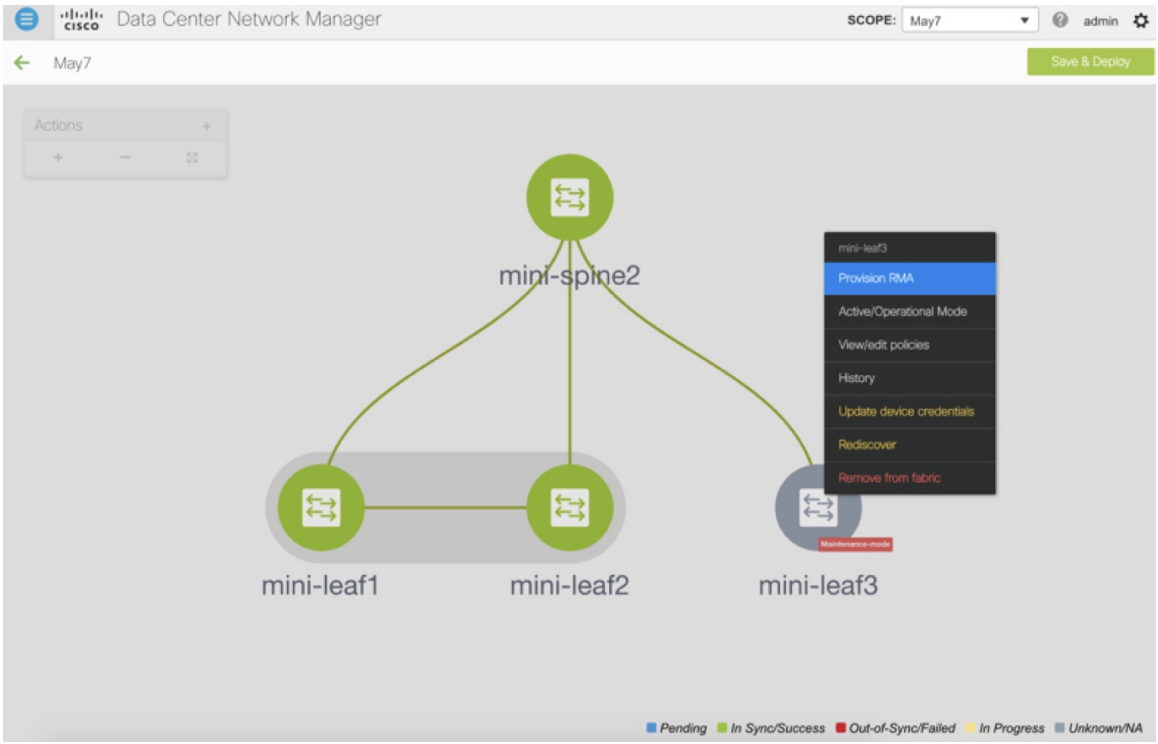

Step 6 The Provision RMA UI will show the replacement device 5-10 minutes after it is powered on.

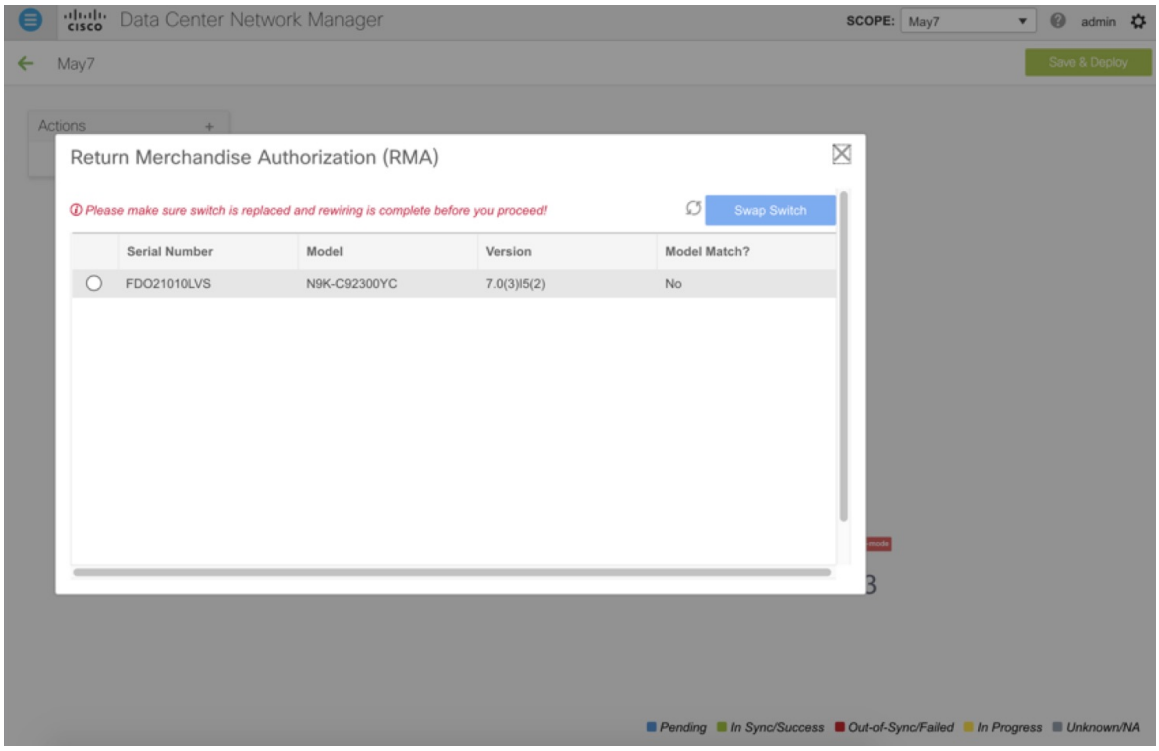

Step 7 Select the correct replacement device and click **Swap Switch**. This begins POAP with the full "expected" configuration for that device. Total POAP time is generally around 10-15 minutes.

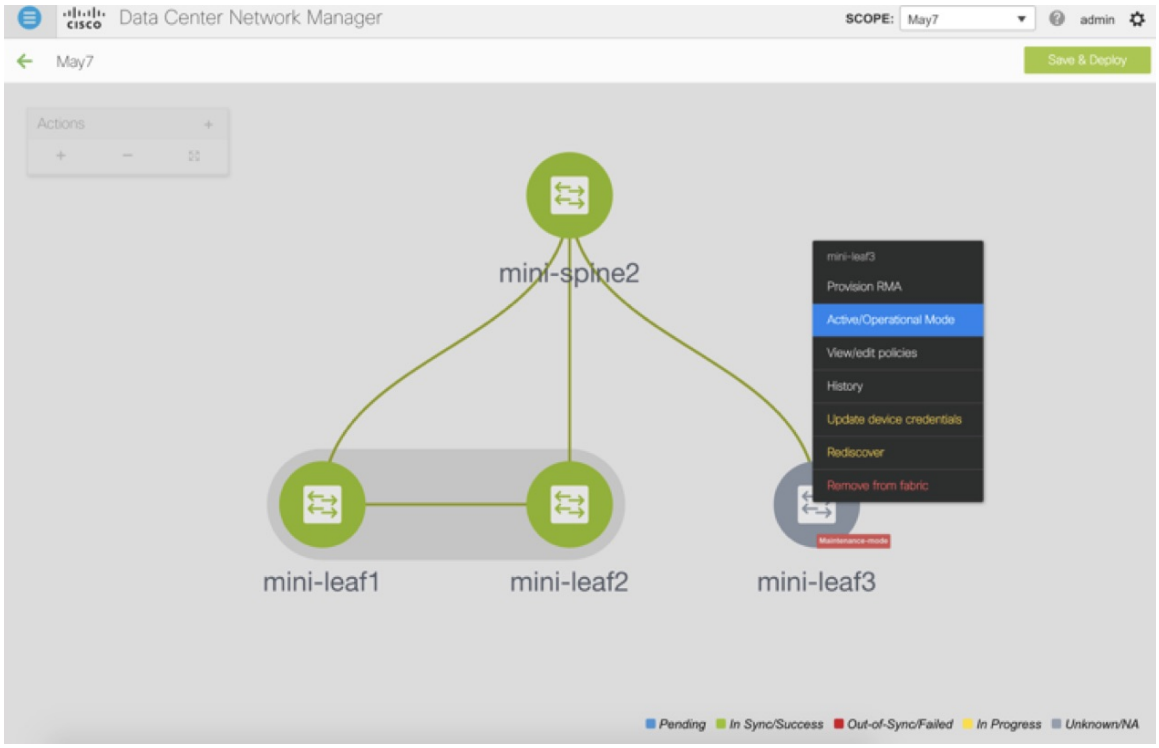

Manual RMA Flow

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).

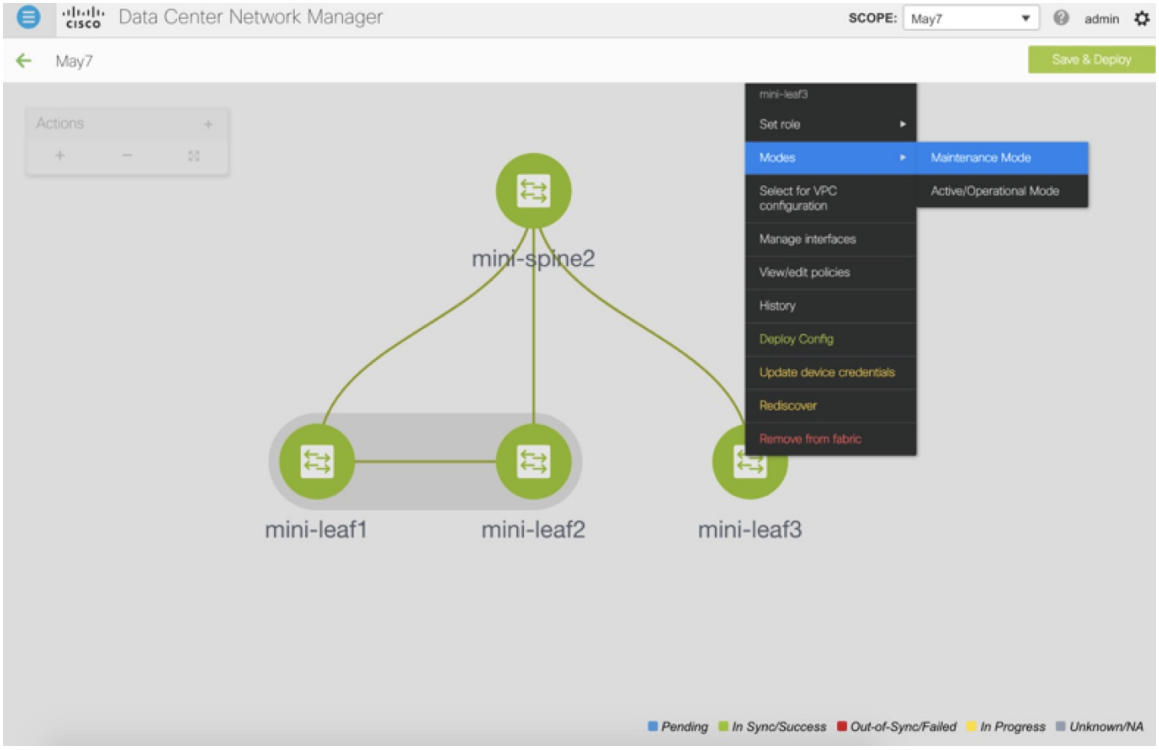

Step 2 Physically replace the device in the network.

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**.

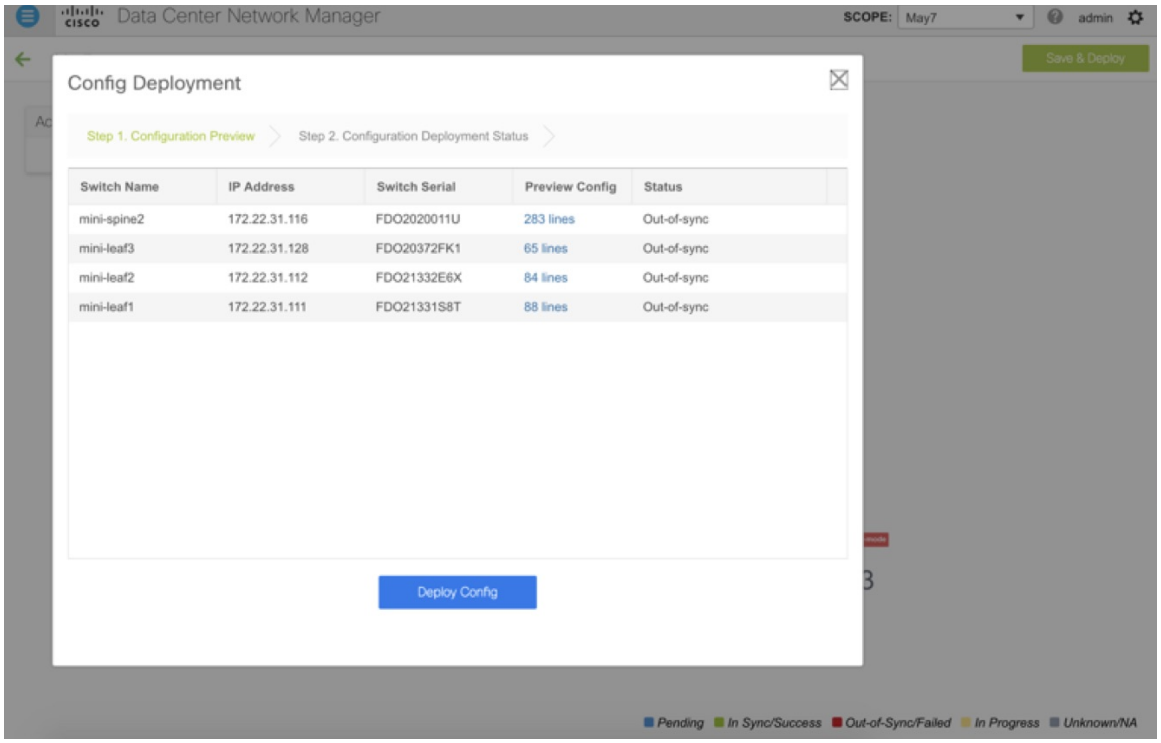

Step 6 Depending on the configuration, if breakout ports or FEX ports are in use, you have to deploy again to completely restore the configuration.

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

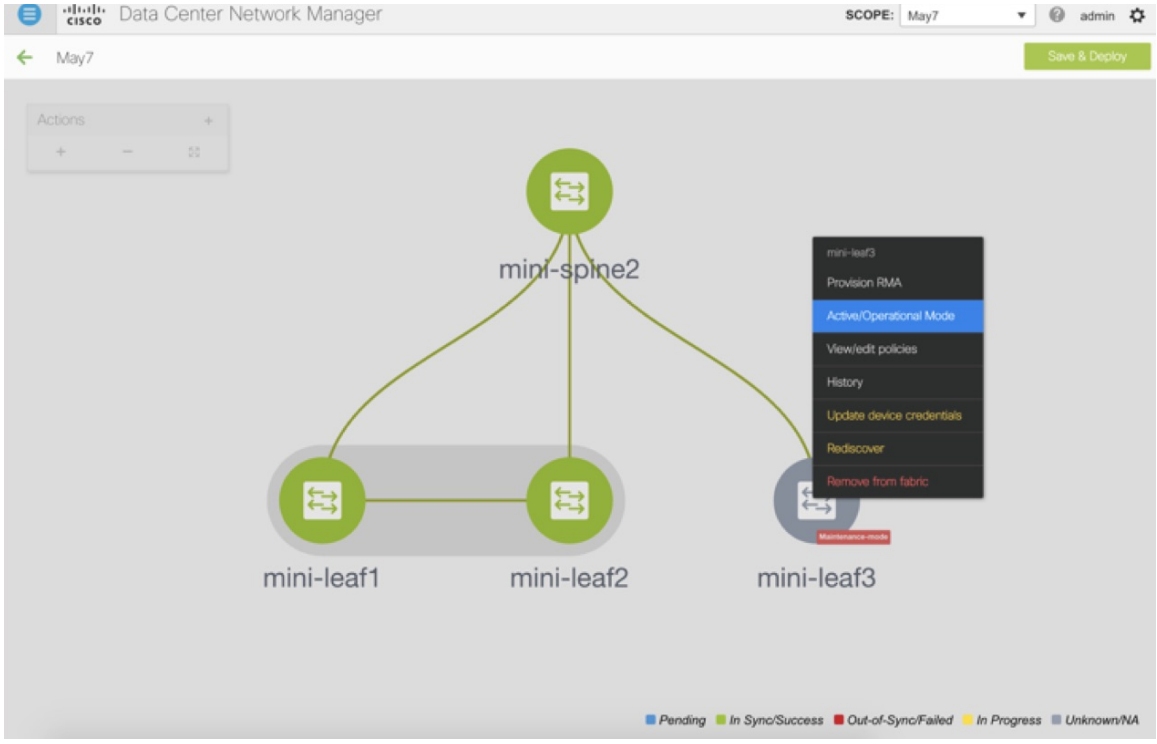

Custom Maintenance Mode Profile Policy

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 266.
- The **flowcontrol** or **priority-flow-control** config is not supported for HIF ports or PO with HIF ports as members.

- 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 aren't 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
- If an interface is created out-of-band, you need to perform fabric resync or wait for Config Compliance poling before this interface can be deleted. Otherwise, Config Compliance does not generate the correct diff.

However, you cannot add or edit interfaces for ASR 9000 Series Routers and Arista switches.

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

Note

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

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

The following table describes the different user roles and the operations these roles support in the **Interfaces** window from Cisco DCNM Release 11.4(1).

Operations	User Roles		
	network-admin	network-operator	network-stager
Add	Save, Preview, Deploy	Blocked	Save, Preview
Breakout	Supported	Blocked	Blocked
Unbreakout	Supported	Blocked	Blocked
Edit	Save, Preview, Deploy	Preview	Save, Preview
Delete	Save, Preview, Deploy	Blocked	Save, Preview
Shutdown	Save, Preview, Deploy	Blocked	Save, Preview
No Shutdown	Save, Preview, Deploy	Blocked	Save, Preview
Show	Supported	Supported	Supported
Rediscover	Supported	Supported	Supported
Deploy	Preview, Deploy	Blocked	Blocked

The following table describes the new user role access-admin operations support in the host facing port of **Interfaces** window from Cisco DCNM Release 11.5(1).

Operations	User Roles
	access-admin
Add	Save, Preview, Deploy
Breakout	Blocked
Unbreakout	Blocked
Edit	Save, Preview, Deploy
Delete	Save, Preview, Deploy
Shutdown	Save, Preview, Deploy
No Shutdown	Save, Preview, Deploy
Show	Supported
Rediscover	Supported
Deploy	Preview, Deploy

From Cisco DCNM Release 11.4(1), you can disable deployments, or freeze, a fabric in DCNM as a network administrator. However, you cannot perform all actions when you freeze the fabric or if the fabric is in monitor mode.

The following table describes the actions you can perform when you freeze a fabric and when you enable the monitor mode for a fabric.

Operations	DCNM Mode	
	Freeze Mode	Monitor Mode
Add	Save, Preview	Blocked
Breakout	Blocked	Blocked
Unbreakout	Blocked	Blocked
Edit	Save, Preview	Blocked
Delete	Save, Preview	Blocked
Shutdown	Save, Preview	Blocked
No Shutdown	Save, Preview	Blocked
Show	Supported	Supported
Rediscover	Supported	Supported
Deploy	Blocked	Blocked

The buttons for the associated operations are grayed out accordingly.

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 Ethernet, and Switch Virtual Interface (SVI). The respective interface ID field is displayed when you select an interface type.

- 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.
- You can preprovision Ethernet interfaces in the Interface window. This preprovisioning feature is supported in Easy, eBGP, and External fabrics. For more information, see Pre-provisioning an Ethernet Interface, on page 90.

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, select the policy to apply 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.

Note The policies are filtered based on the interface type you choose in the **Type** drop-down list and the device you choose in the **Select a Device** drop-down list.

Step 7 Enter values in the required fields under the **General** tab.

The fields vary according to the interface type you choose.

Note From Cisco DCNM, Release 11.5(1) you can mirror the configurations of Peer-1 on Peer-2 while creating a vPC. When you check the **Enable Config Mirroring** check box, the Peer-2 fields will be grayed out. The configurations that you enter in the Peer-1 fields will be copied to Peer-2 fields.

From Release 11.5(1), **int_vlan** template allows you to create SVI interfaces on Nexus switches in External Fabric and LAN Classic fabric. DCNM release 11.5(3a) allows you to deploy SVI interfaces using **int_vlan** template for Easy Fabrics also. Additionally, up to 3 DHCP relay IPs can be provided for this interface.

If **int_vlan** is created in Easy_Fabric, you can advertise the SVI subnet to underlay IGP (OSPF or ISIS). Select the **Advertise Subnet into Underlay IGP** checkbox.

Note **adv_static_prefix_bgp** template will create a static route for the specified IPv4 prefix and advertise it to eBGP peers.

Step 8 Click **Save** to save the configurations.

Note To apply QoS policies on the interface, create the interface freeform with references accordingly.

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

Step 10 Click **Deploy** to deploy the specified logical interface.

The newly added interface appears in the screen.

Breakout

Click the drop-down arrow next to the **Breakout** icon to display a list of the available breakout options. The available options are **10g-4x**, **25g-4x**, **50g-2x**, **50g-4x**, **100g-2x**, **100g-4x**, **200g-2x**, and **Unbreakout**. Choose the required option.

Editing Interfaces

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

Note The **Edit Interface** allows you to change the policy and add or remove an interface from a port channel or vPC.

Procedure

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

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

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

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

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

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

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

During overlay network deployment on switches, the network can be associated with trunk interfaces. The trunk interface to network association is reflected in the **Interfaces** screen. You can update such interfaces.

For interface policies that are not created from the **Control** > **Interfaces** screen, you can edit some configurations but not change the policy itself. The policy and fields that cannot be edited are grayed out.

The following are some examples of policies that cannot be edited:

- Loopback interface policies - The int_fabric_loopback_11_1 policy is used to create a loopback interface. You can edit the loopback IP address and description but not the int_fabric_loopback_11_1 policy instance.
- Fabric underlay network interface policies (int_fabric_num_11_1, for example) and fabric overlay network interface (NVE) policies.
- Policies associated with port channels and member ports of port channels, including the port channels and member ports associated with a vPC.

- SVIs created during network and VRF creation. The associated VLANs appear in the interfaces list.

Editing Interfaces Associated with Links

There are two types of links, namely intra-fabric links and inter-fabric links. As the name implies, intra-fabric links are set up between devices within the same Easy fabric and are typically used for spine-leaf connectivity. Inter-fabric links are set up between the Easy fabric, and typically other external or Easy fabrics. They are used for external WAN and/or DCI connectivity. A policy is associated with each link that effectively states the configuration that is applied to both ends of the link. In other words, the link policy becomes the parent of the individual child interface policies that are associated with the two interfaces that form the link. In this scenario, you must edit the link policy to edit the interface policy fields such as description, IP address, and any per interface freeform config. The following procedure shows how to edit the interfaces associated with links:

Procedure

Step 1 Choose **Control > Fabric Builder**, and select the fabric containing the link.

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

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

Step 3 Click the **Links** tab.

Step 4 Select the link that you want to edit and click the **Update Link** icon.

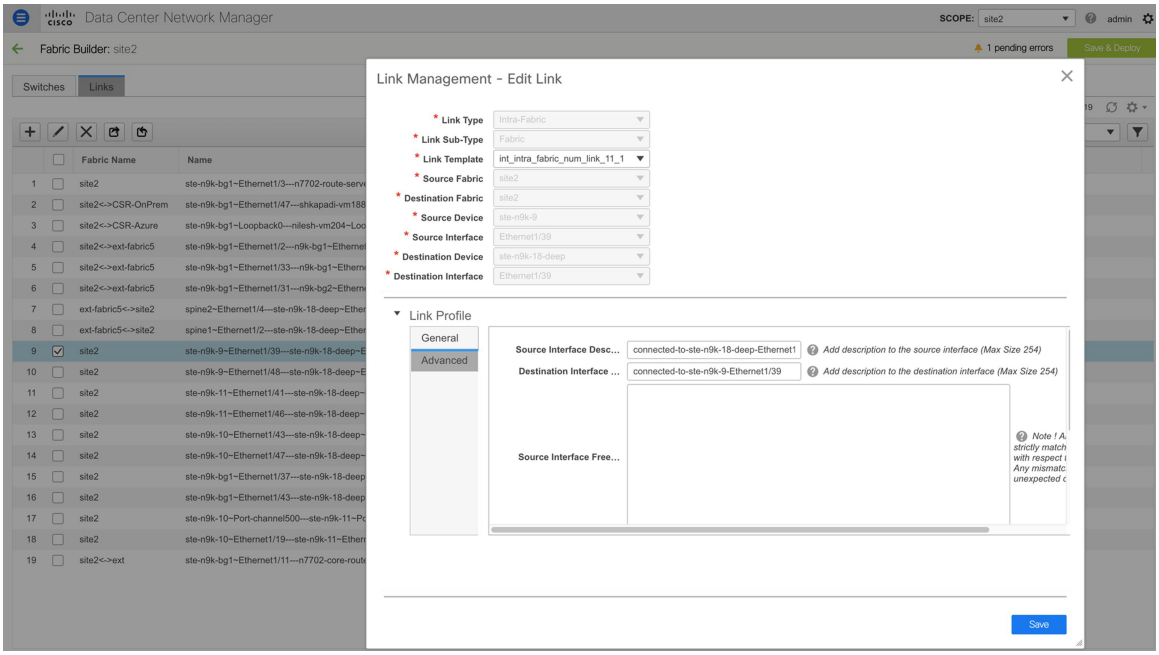

Update the link based on your requirements and click **Save**.

Deleting Interfaces

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

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

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

Procedure

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

Step 2 Select the interfaces.

Step 3 Click **Delete**.

You cannot delete logical interfaces created in the fabric underlay.

Step 4 Click **Save**.

Step 5 (Optional) Click **Preview** to view all the changes before deleting the interface.

The deletion will be highlighted in red colour with strikethrough under the **Expected Config** tab.

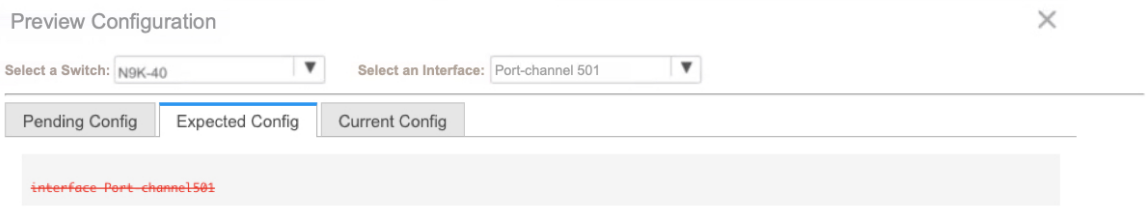

Step 6 Click **Deploy** to delete the interface.

Shutting Down and Bringing Up Interfaces

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

Procedure

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

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

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

A confirmation window appears where you can save, preview, and deploy the changes. Click **Save** to preview of deploy the changes.

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

A confirmation window appears where you can save, preview, and deploy the changes. Click **Save** to preview or deploy the changes.

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.

After you deploy the interface configuration, the interface status information is updated. However, the overall switch-level state may be in the pending state, which is in blue. The overall switch-level state goes to the pending state whenever there is a change in intent from any module, such as interface, link, policy template update, top-down, or so on. In the pending state, a switch may have pending configurations or switch-level recomputation. The switch-level recomputation occurs when:

- You preview or deploy for the switch
- During a save and deploy
- During hourly sync

Preview or deploy the switches to review their state and to understand the root cause of their pending state. Save and deploy for a fabric-wide recomputation.

Click **Preview** to preview the configurations before you click **Deploy**.

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 the Cisco Nexus 9000, 3000, and 7000 Series Switches in the external fabric.

When you add an interface to an external fabric device, the Resource Manager is not in sync with the device. So, ensure that the value populated in the ID field (Port-channel ID, vPC ID, Loopback ID, etc) is not previously configured on the switch.

If you want to configure a portchannel in the external fabric, you should add and deploy the **feature_lacp** policy on the switches where the portchannel will be configured.

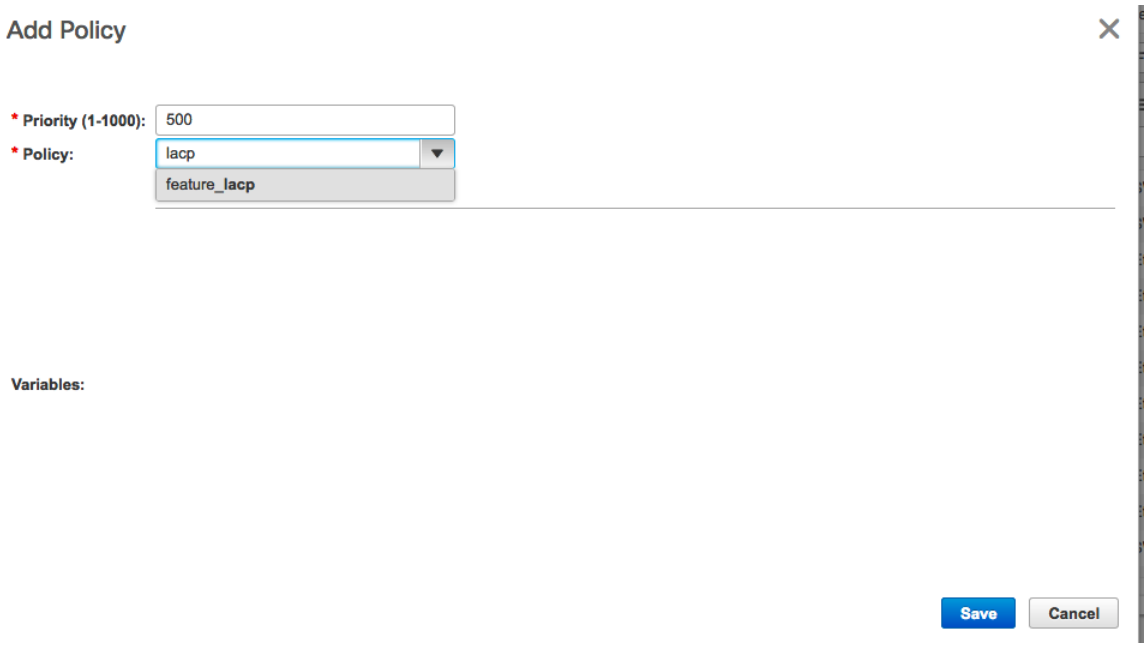

When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on its switches. If you click **Save & Deploy** in the fabric topology screen, it displays an error message. However, the following settings (available when you right-click the switch icon) are allowed:

vPC pairing - You can designate a vPC switch pair, but it is only for reference.

View/edit policy - You can add a policy but you cannot deploy it on the switch.

Manage interfaces – You can only create intent for adding interfaces. If you try to deploy, edit, or delete interfaces, it results in an error message.

Interface Groups

From Cisco DCNM Release 11.5(1), you can create an interface group that allows grouping of host-facing interfaces at a fabric level. Specifically, you can create an interface group for physical Ethernet interfaces, L2 port-channels, and vPCs. You can attach or unattach multiple overlay networks to the interfaces in an interface group.

Guidelines

- Interface groups are only supported for the fabrics with the **Easy_Fabric_11_1** template.
- An interface group is specific to a fabric. For example, consider two fabrics: Fab1 and Fab 2. The interface group IG1 in Fab1 isn't applicable to Fab 2.
- An interface group can only have interfaces of a certain type. For example, you need three separate interface groups if you want to group three types of interfaces such as IG1 for physical Ethernet trunk interfaces, IG2 for L2 trunk port-channels, and IG3 for vPC host trunk ports.
- An interface group can also be created using preprovisioned interfaces.
- Interface groups are limited to switches with the leaf role. They aren't supported for other roles such as Border, BGW, and other related variants.
- For L2 port-channels and vPCs that are part of an interface group, they can't be deleted until they are de-associated from the interface group even if there are no networks associated with the interface group. Similarly, a trunk port that has no overlay networks but is part of an IG can't be converted to an access port. In other words, you can't change policies for interfaces that are part of an interface group. However, you can edit certain fields for policies.
- For L4-L7 services configuration on leaf switches, trunk ports that are used for services attachment can't be part of interface groups.
- When you perform a per fabric backup of an easy fabric, if there are interface groups created in that fabric, all the associated interface group state is backed up.
- If an easy fabric contains an interface group, then this fabric can't be imported into the MSO. Similarly, if an easy fabric has been added to the MSO, you can't create interface groups for interfaces that belong to switches in the easy fabric.
- The **Interface Group** button is enabled only for Admin and Stager users. For all other users, this button is disabled.
- The **Interface Group** button is disabled in the following circumstances:
    - Select **Data center** from the **SCOPE** drop-down list.
    - Select a fabric without any switches.
    - Select any other interface apart from vPC, Port-channel, and Ethernet.
    - If the interface has a policy attached from another source, for example:
        - If the interface is member of a port-channel or vPC.
        - If the port-channel is member of vPC.
        - If the interface has a policy from underlay or links.

Note If you select different types of interfaces, the **Interface Group** button is enabled. However, when you try to create or save different types of interfaces to an interface group, an error is displayed.

Creating an Interface Group

Procedure

Step 1 From DCNM, navigate to **Control > Fabrics > Interfaces**.

Step 2 From the **SCOPE** drop-down list, select a fabric.

Step 3 Select the interfaces that have to be grouped and click **Interface Group**.

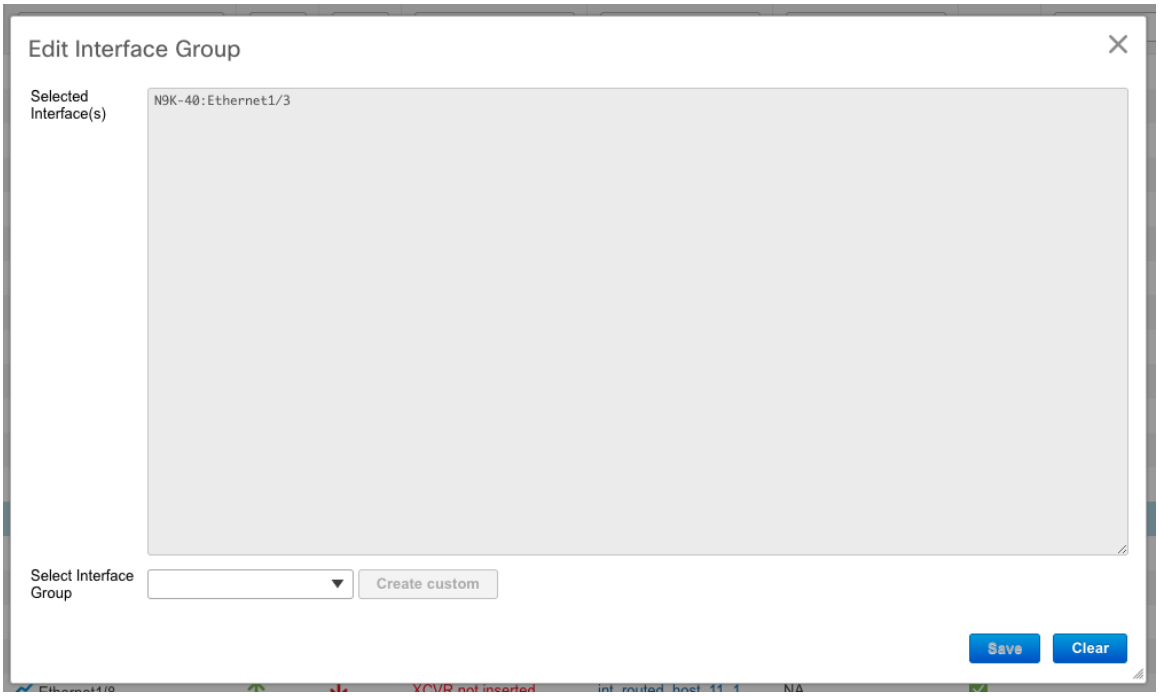

Step 4 4. In the **Edit Interface Group** window, create a custom interface group by entering an interface group name in the **Select Interface Group** field and click **Create custom**. An interface group name can have a maximum length of 64 characters.

If you have already created an interface group, select it from the **Select Interface Group** drop-down list. Also, if an interface is already part of an interface group, you can move it to a different interface group by selecting the new group from the **Select Interface Group** drop-down list.

Note An interface can belong to only a single interface group.

You can create interface groups from either the **Interfaces** window or the **Networks** window. For more information, see Attaching Networks to an Interface Group, on page 273.

Step 5 Click **Save**.

In the **Interfaces** window, you can see the interface group name under the **Interface Group** column.

Removing Interfaces from an Interface Group

Procedure

Step 1 From DCNM, navigate to **Control > Fabrics > Interfaces**.

Step 2 From the **SCOPE** drop-down list, select a fabric.

Step 3 Select the interfaces to disassociate from an interface group and click **Interface Group**.

Step 4 In the **Edit Interface Group** window, make sure that nothing is selected in the **Select Interface Group** drop-down list, and click **Clear**.

A dialog box pops up asking whether you want to clear all the associated interfaces. Click **Yes** to proceed. Note that if there are any networks attached to these interfaces, they are detached as well when you click **Clear**.

Attaching Networks to an Interface Group

Procedure

Step 1 From DCNM, navigate to **Control > Fabrics > Networks**.

Step 2 From the **SCOPE** drop-down list, select a fabric.

Step 3 In the **Networks** window, select the networks that you need to attach to an interface group and click **Interface Group**.

Note
- An overlay network can belong to multiple interface groups.
- You can select only the networks with a VLAN ID. Otherwise, an appropriate error message is displayed.

Step 4 In the **Interface Groups** window, you can perform the following:

- Select an existing interface group from the **Select Interface Group** drop-down list and click **Save**.

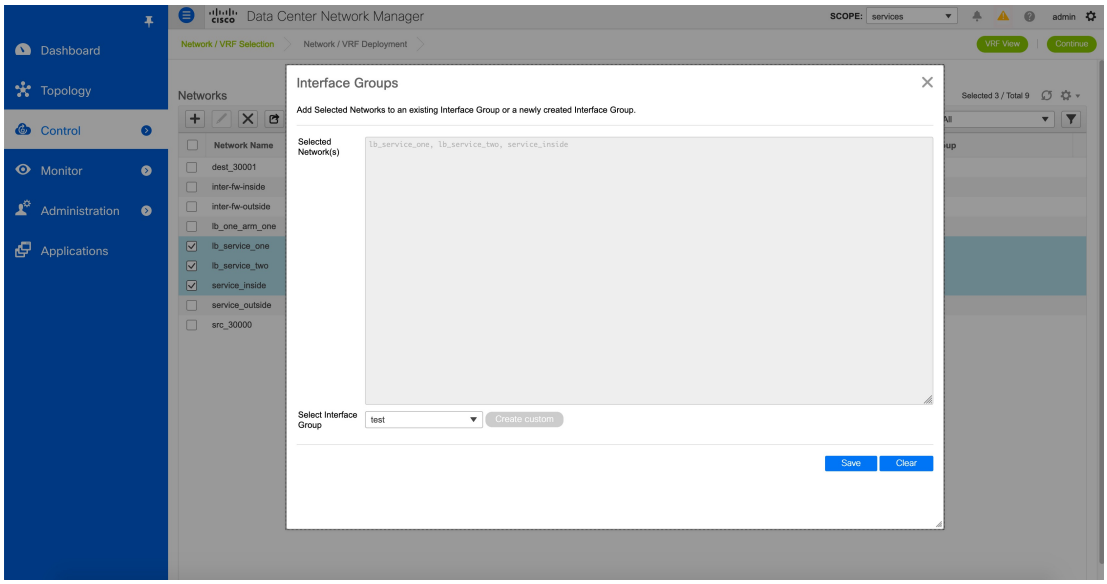

For example, you select three networks and the interface group **test**, and click the **Save** button, the following operations are performed in the background:

a. DCNM retrieves interfaces that are part of the interface group **test**.

b. DCNM determines that three networks are added to the interface group **test**. Therefore, it autoattaches these networks to all the interfaces that are part of the interface group **test**.

c. For each interface, DCNM pushes the "**switchport trunk allowed vlan add xxxx**" command three times for each selected network.

Note DCNM ensures that there's no duplicate configuration intent.

If you click the **Clear** button, DCNM pushes "**switchport trunk allowed vlan remove xxx**' config intent.

- Create a custom interface group by entering an interface group name in the **Select Interface Group** field and click **Create custom**. Click **Save**.

  If you choose this option, make sure to add interfaces to this Interface Group in the **Interfaces** window. As a result, DCNM performs the following operations:

  a. Removes all existing overlay networks that don't belong to the interface group from these interfaces.

  b. Adds new overlay networks to these interfaces that are part of the interface group but not yet attached to these interfaces.

  For more information about associating interfaces to interface groups, see Creating an Interface Group, on page 272.

Step 5 Click **Continue** and click **Save & Deploy** to deploy the selected networks on the switches.

Unattaching a Network from an Interface Group

This procedure shows how to unattach a network from an interface group in the Networks window. Also, you can unattach networks when you remove an interface from an interface group in the **Interfaces** window. For more information, see *Removing Interfaces from an Interface Group*.

Procedure

Step 1 1. From DCNM, navigate to **Control > Fabrics > Networks**.

Step 2 From the **SCOPE** drop-down list, select a fabric.

Step 3 In the **Networks** window, select the networks that you need to unattach to an interface group and click **Interface Group**.

Step 4 In the **Interface Groups** window, select the interface group from the **Select Interface Group** drop-down list and click **Clear** to unattach a network.

Step 5 (Optional) Navigate to **Control > Fabrics > Interfaces**.

Under the **Overlay Network** column, you can see the unattached network in the red color for the corresponding interface. Click the network to view the expected config that is struck through.

Step 6 Navigate to the **Fabric Builder** or **Networks** window, and click **Save & Deploy**.

Deleting an Interface Group

An interface group is automatically deleted when it's not in use. DCNM performs an implicit delete of an interface group if there are no interfaces and no networks mapped to the interface group. This check is performed whenever you click the **Clear** button in the **Edit Interface Group** window. There may be exception scenarios where you need to clean up the interface groups explicitly.

For example, you create an interface group **storageIG** and add an interface to it. Later, you want to change the interface mapping to another group. Therefore, you select the interface and click **Interface Group** to open the **Edit Interface Group** window. Select the other interface group named **diskIG**. Now, the **storageIG** interface group doesn't have any associated member interfaces or networks. In this case, perform the following steps:

Procedure

Step 1 Select an interface that doesn't belong to an interface group.

Step 2 Click **Interface Group** to open the **Edit Interface Group** window.

Step 3 Select the **StorageIG** interface group from the **Select Interface Group** drop-down list.

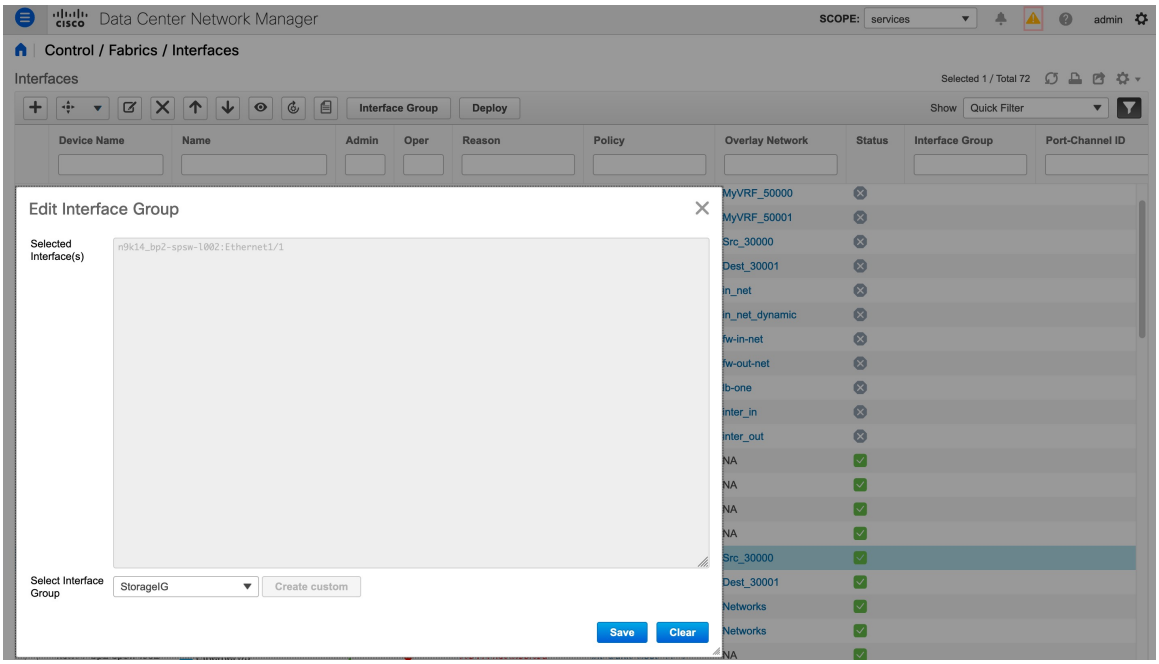

Step 4 Click **Clear**.

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.

For information about creating interface groups and attaching networks, see Interface Groups, on page 270.

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.
- From a fabric topology window: Right-click anywhere in the fabric topology window. Choose **Overlay View > VRF View** or **Overlay View > Network View**, accordingly. This option is applicable only for switch fabrics, easy fabrics, and MSD fabrics.

You can toggle between the network view and VRF view in both the windows by clicking the **VRF View** or **Network View** button. When you are in the networks or VRFs window, ensure you choose the appropriate fabric from the **Scope** drop-down list before you create any networks or VRFs.

Viewing Networks and VRFs for a Fabric

- Click **Control** > **Networks** from the main menu.

  The **Networks** screen comes up. The SCOPE drop down box (at the top right part of the screen) lists all fabrics managed by the DCNM instance, in alphabetical order. You can choose the correct fabric from SCOPE. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

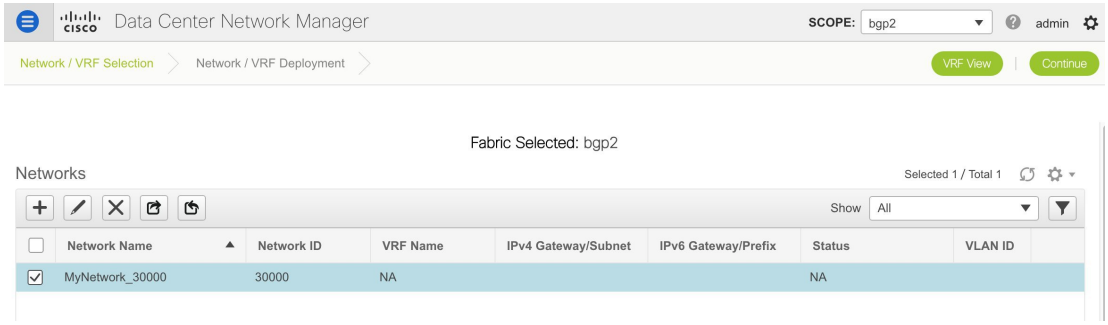

- Click **Control** > **VRFs** from the main menu.

  The VRFs screen comes up. The SCOPE drop down box (at the top right part of the screen) lists all fabrics managed by the DCNM instance, in alphabetical order. You can choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

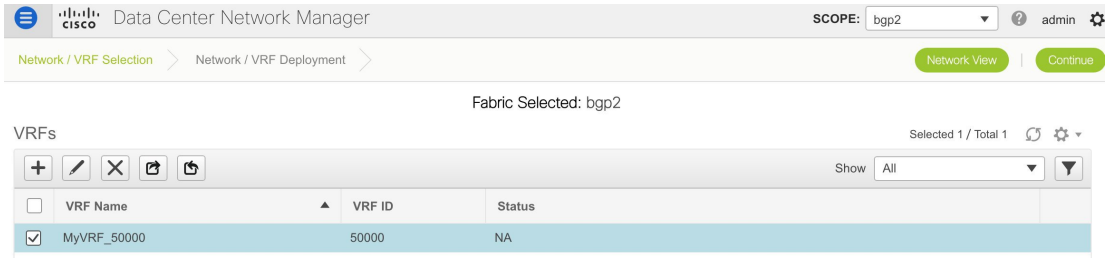

Note The **Networks** or **VRFs** windows are applicable only for the Easy or MSD fabrics.

Creating Networks for the Standalone Fabric

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

   The Networks screen comes up.

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

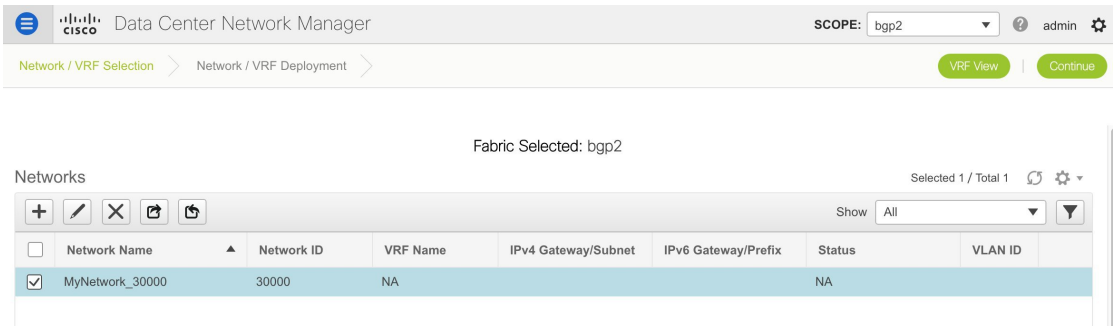

3. Click the + button at the top left part of the screen (under **Networks**) to add networks to the fabric. The Create Network screen comes up. Most of the fields are autopopulated.

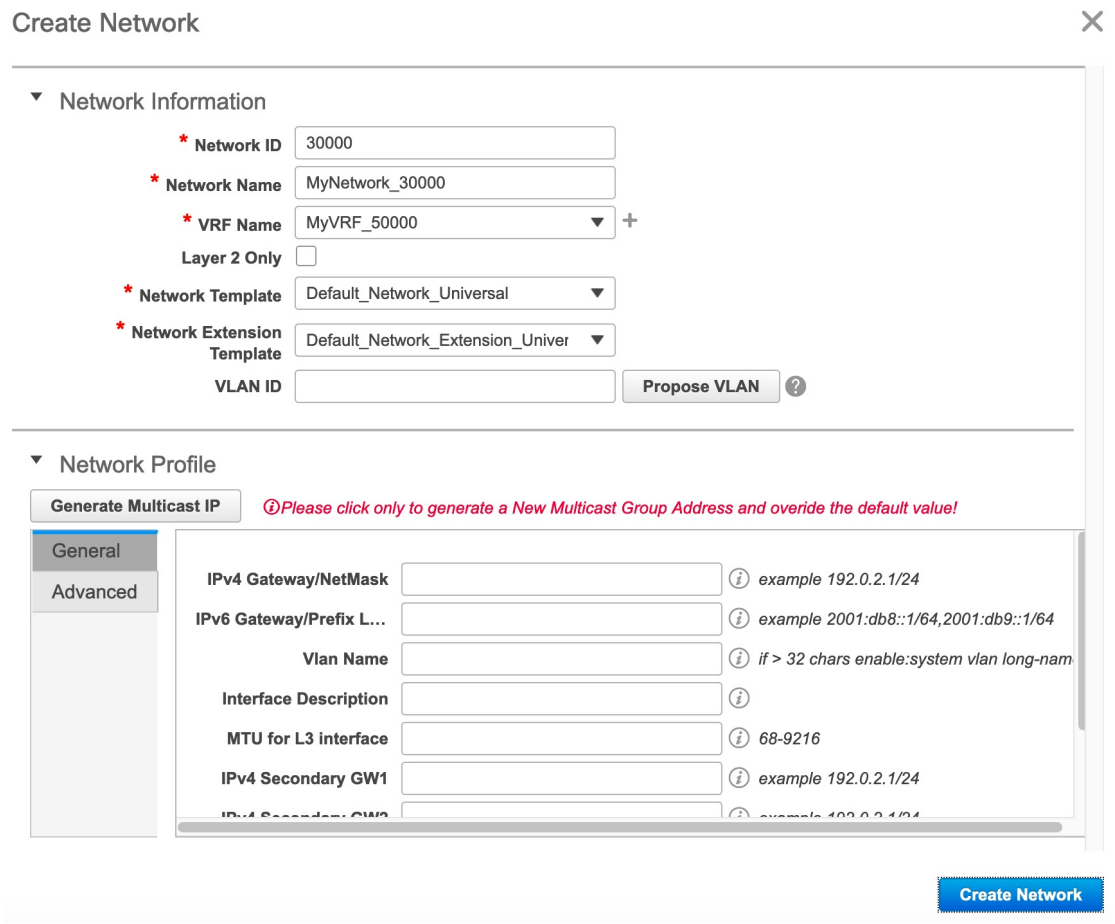

The fields in this screen are:

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

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

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

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

Network Template: A universal template is autopopulated. This is only applicable for leaf switches.

Network Extension Template: A universal extension template is autopopulated. This allows you to extend this network to another fabric. The methods are VRF Lite, Multi Site, and so on. The template is applicable for border leaf switches and BGWs.

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

The VLAN ID default range is 2 to 3967. From DCNM Release 11.5(2), you can use a VLAN range greater than default value 3967. The reserved VLAN range must be set to a different range. In switch command enter "**system vlan <vlan> reserve**". Save the configuration to startup configuration and reload the switch for the new reserved VLAN range to reflect.

From Cisco DCNM Web UI, choose **Administration > DCNM Server > Server Properties**, enter the value as 4094 for **RM.TOP_DOWN_NETWORK_VLAN.MAX** and **RM.TOP_DOWN_VRF_VLAN.MAX**, click **Apply Changes** and then restart DCNM. Once the DCNM is up, you can create VRF and network using the VLAN value greater than 3967.

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 t 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 configu 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 fiel

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.

Loopback ID for DHCP Relay interface (Min:0, Max:1023) - Specifies the loopback ID for DHCP relay interface.

Routing Tag – The routing tag is autopopulated. This tag is associated with each gateway IP address prefix.

TRM enable – Select the check box to enable TRM.

For more information, see Overview of Tenant Routed Multicast, on page 196.

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 check box to enable a Layer 3 gateway on the border switches.

A sample of the Create Network screen is given below.

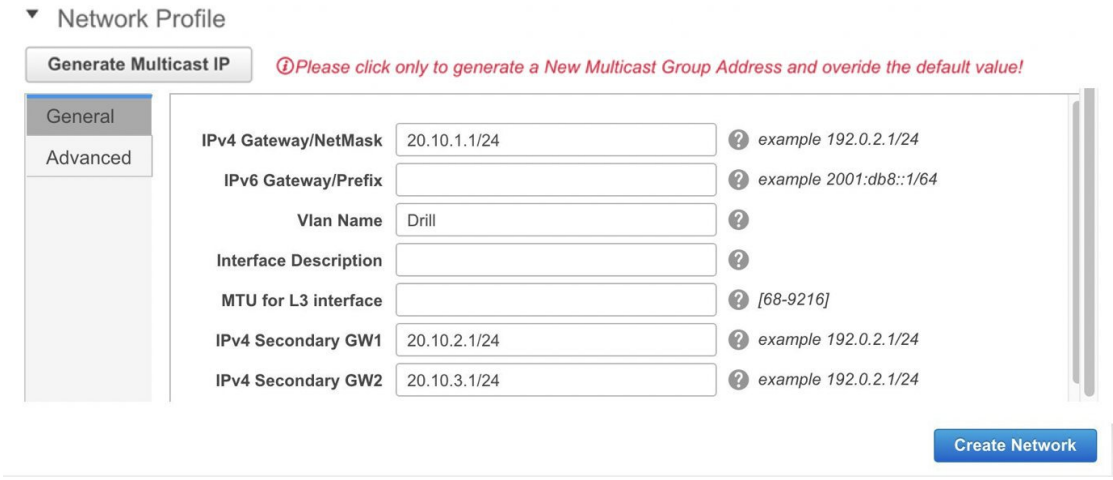

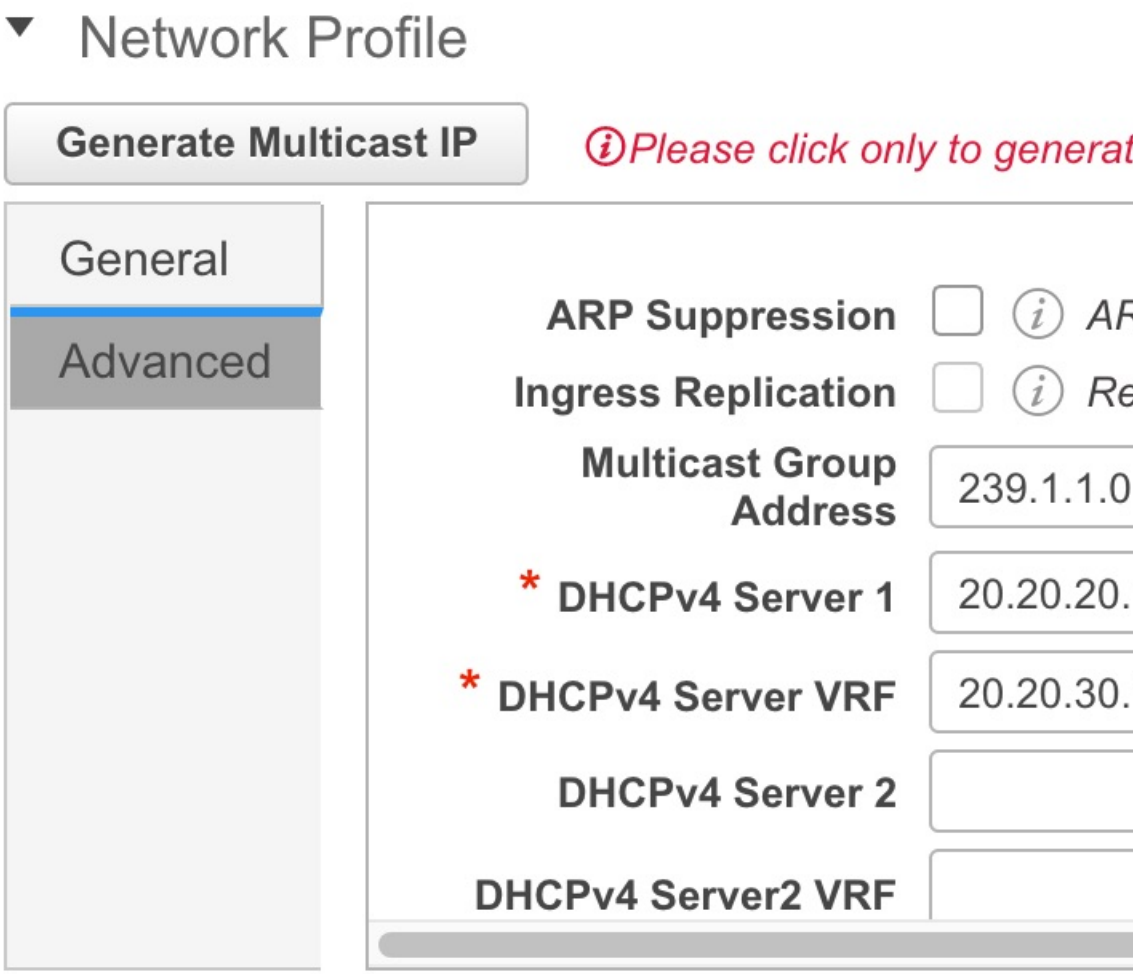

4. Click **Create Network**. A message appears at the bottom right part of the screen indicating that the network is created.

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

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

Export and Import Network Information

You can export network information to a .CSV file. The exported file contains information pertaining to each network, including the fabric it belongs to, the associated VRF, the network templates used to create the network, and all other configuration details that you saved during network creation.

In the Networks screen, click the Export icon to export network information as a .CSV file.

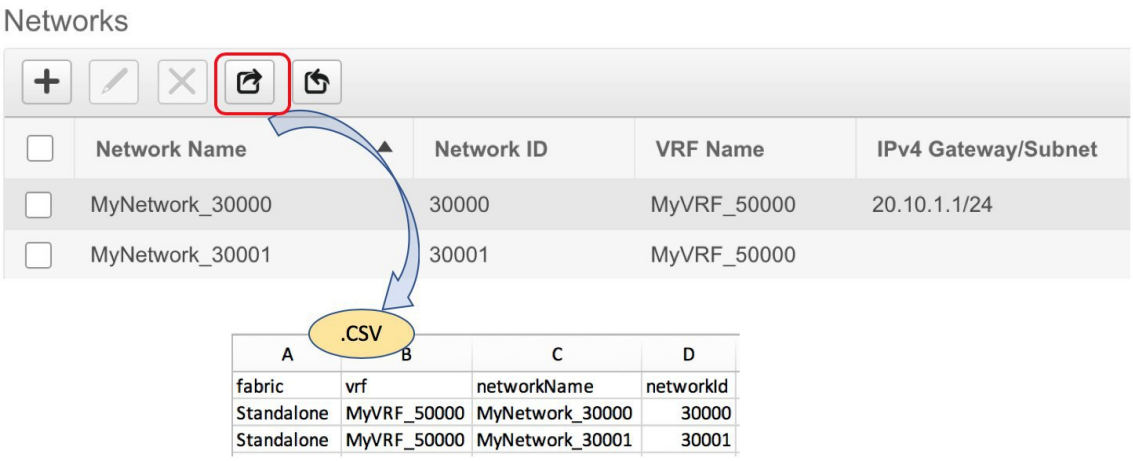

You can use the exported .CSV file for reference or use it as a template for creating new networks. To import networks, do the following:

1. Update new records in the .CSV file. Ensure that the networkTemplateConfig field contains the JSON Object. A message at the bottom right part of the screen displays errors and success messages. This screenshot depicts two new networks being imported.

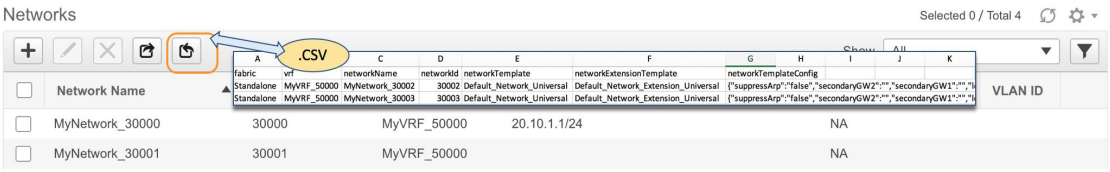

2. In the Networks screen, click the Import icon and import the .CSV file into DCNM.

You can see that the imported networks are displayed in the Networks screen.

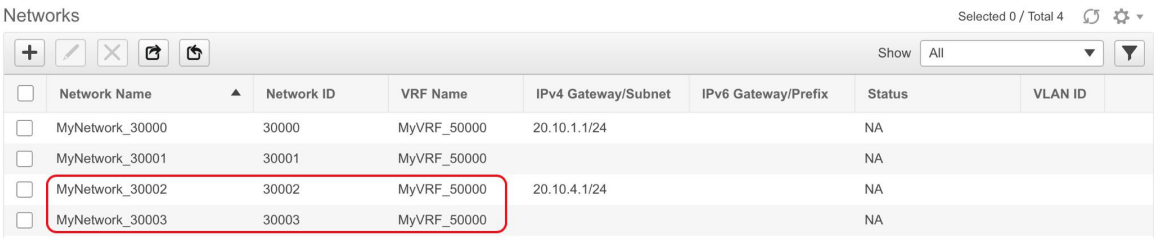

Editing Networks for the Standalone Fabric

To edit networks for standalone fabrics from Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Click **Control** > **Networks**.

The **Networks** window appears.

Step 2 Choose a fabric from the **SCOPE** drop-down list.

The **Networks** window refreshes and lists the networks in the fabric.

Step 3 Choose a network.

Step 4 Click the **Edit** icon.

The **Edit Network** window appears.

Step 5 Update the fields in the **General** and **Advanced** tabs of the **Network Profile** area as needed.

Note You can edit the network name. The edited network name appears in the **Network Name** column in the **Networks** window. The original name, which you used while creating a network, appears in the **Display Name** column. To view the original network name from the **Display Name** column in the **Networks** window, click **Settings**. Expand the **Columns** drop-down list, and choose the **Display Name** option. Click **Close**. You can also view the original network name in the network topology view.

Step 6 Click **Save** at the bottom right part of the window to save the updates.

Creating VRFs for the Standalone Fabric

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

   The VRFs screen comes up.

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

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

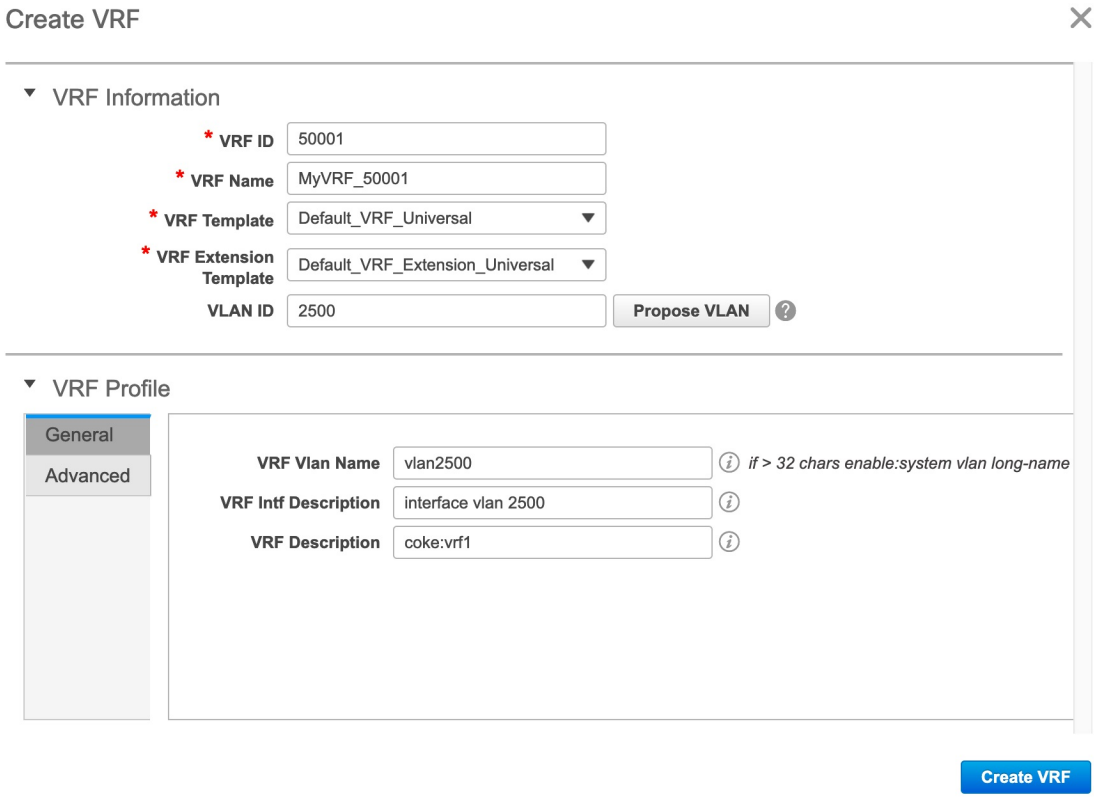

The fields in this screen are:

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

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

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

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

Fill the fields in the **VRF Profile** section.

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

The VLAN ID default range is 2 to 3967. From DCNM Release 11.5(2), you can use a VLAN range greater than default value 3967. The reserved VLAN range must be set to a different range. In switch command enter "**system vlan <vlan> reserve**". Save the configuration to startup configuration and reload the switch for the new reserved VLAN range to reflect.

From Cisco DCNM Web UI, choose **Administration > DCNM Server > Server Properties**, enter the value as 4094 for **RM.TOP_DOWN_NETWORK_VLAN.MAX** and **RM.TOP_DOWN_VRF_VLAN.MAX**, click **Apply Changes** and then restart DCNM. Once the DCNM is up, you can create VRF and network using the VLAN value greater than 3967.

Advanced tab – The fields in the tab are autopopulated.

VRF Intf MTU - Specifies VRF interface MTU.

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 check box to enable TRM.

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

For more information, see Overview of Tenant Routed Multicast, on page 196.

Is RP External – Enable this checkbox if the RP is external to the fabric. If this field is unchecked, RP is distributed in every VTEP.

RP Address – Specifies the IP address of the RP.

RP Loopback ID – Specifies the loopback ID of the RP, if **Is RP External** is not enabled.

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 **Default MDT Address for TRM VRFs** field in the fabric settings scree auto-populated in this field. You can override this field if a different multicast group address should be for this VRF.

Overlay Multicast Groups – Specifies the multicast group subnet for the specified RP. The value is the group range in “ip pim rp-address” command. If the field is empty, 224.0.0.0/24 is used as default.

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.

Enable TRM BGW MSite - Select the check box to enable TRM on Border Gateway Multisite.

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.

Config Static 0/0 Route - Select the check box to enable static default route configuration.

BGP Neighbor Password - Specifies the VRF Lite BGP neighbor password.

BGP Password Key Encryption Type - Select the encryption type from this drop-down list.

Sample screenshots of the Create VRF screen:

Advanced tab:

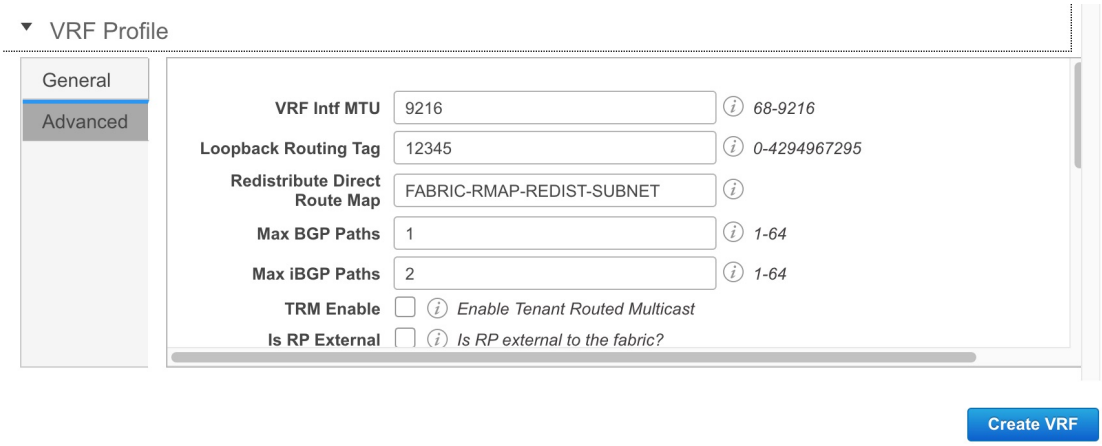

From Release 11.5(3a), DCNM allows you to configure Route Target (RT) relevant parameters for VRF. You can manually update RT parameters, and also choose to disable auto-generating RT parameters.

Route Target tab:

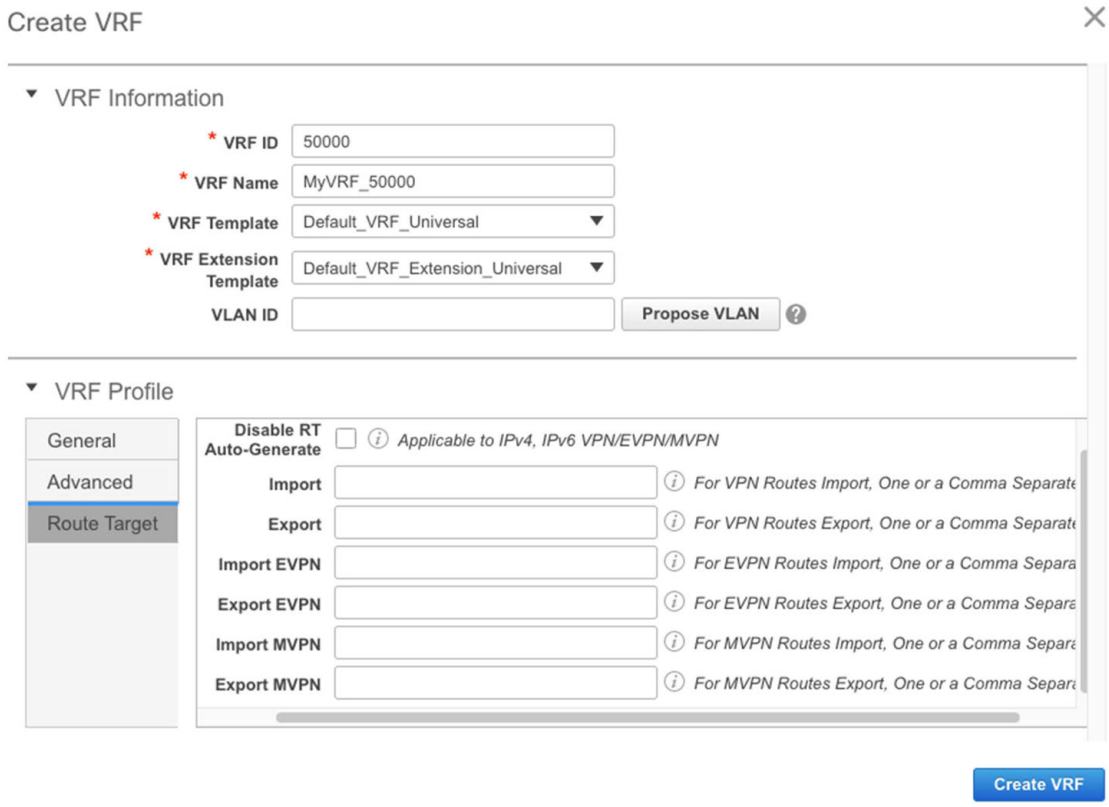

The fields in this screen are:

Disable RT Auto-Generate – Disables auto RT. If you do not select the check box and allow RT to autogenerate, auto-RT will have import and export route target based on the BGP ASN and VNI values.

Import – Route-Targets for VPN Routes Import

Export – Route-Targets for VPN Routes Export

Import EVPN – Route-Targets for EVPN Routes import

Export EVPN – Route-Targets for EVPN Routes export

If the Easy fabric is a member of the Multi-Site domain, all the above 4 parameters are controlled by the Multi-Site domain. Only the MVPN parameters are configurable.

Import MVPN – Route-Targets for MVPN routes import. This field is available at Easy Fabric, for both standalone fabrics, and a member fabric of the Multi-Site domain. You can configure this parameter only when TRM is enabled.

Export MVPN – Route-Targets for MVPN routes export. This field is available at Easy Fabric, for both standalone fabrics, and a member fabric of the Multi-Site domain. This parameter can be configured only when TRM is enabled.

4. Click **Create VRF**.

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

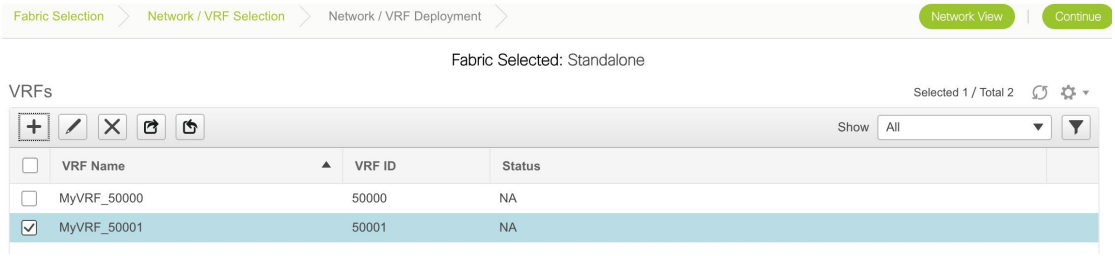

Export and Import VRF Information

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.

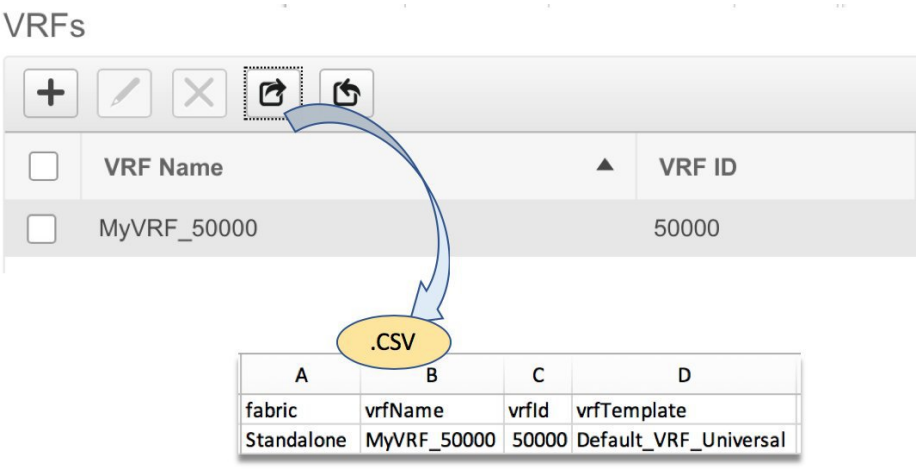

You can use the exported .CSV file for reference or use it as a template for creating new VRFs. To import VRFs, do the following:

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.

Note When you create a VRF using the **Import** option on the VRF window or using the DCNM APIs, you mig see an error saying: Instance name is not specified.

This error is because of a tagging issue. To remove this error, edit the VRF in DCNM Web UI and then deplo

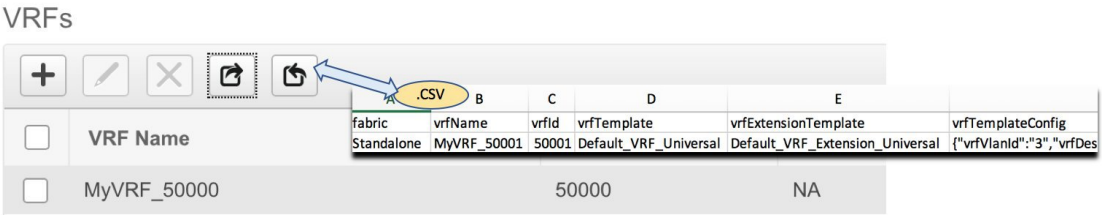

You can see that the imported VRF is displayed in the VRFs screen.

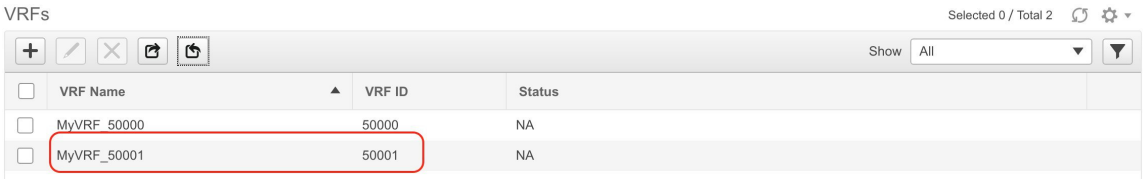

Editing VRFs for the Standalone Fabric

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

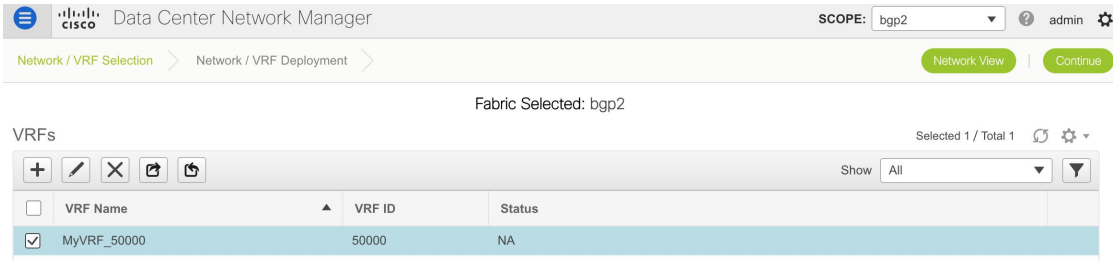

2. From the **Select a Fabric** drop-down list, select the fabric *Standalone*, and click **Continue** on the top right part of the screen. The Networks page is displayed.
3. Click the **VRF View** at the top right part of the screen. The VRFs page appears.

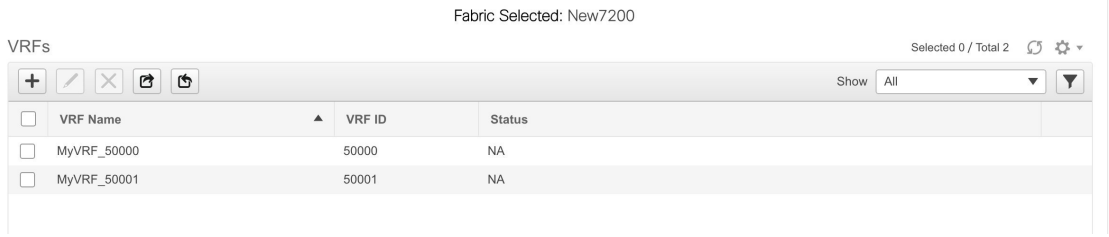

4. Select the **VRF** and click the **Edit** option at the top left part of the screen. The **Edit VRF** screen comes up.
5. Update the fields in the **General** and **Advanced** tabs of the **VRF Profile** section as needed.
6. Click **Save** at the bottom right part of the screen to save the updates.

Deploying Networks for the Standalone and MSD Fabrics

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

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

   The Networks screen comes up.

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

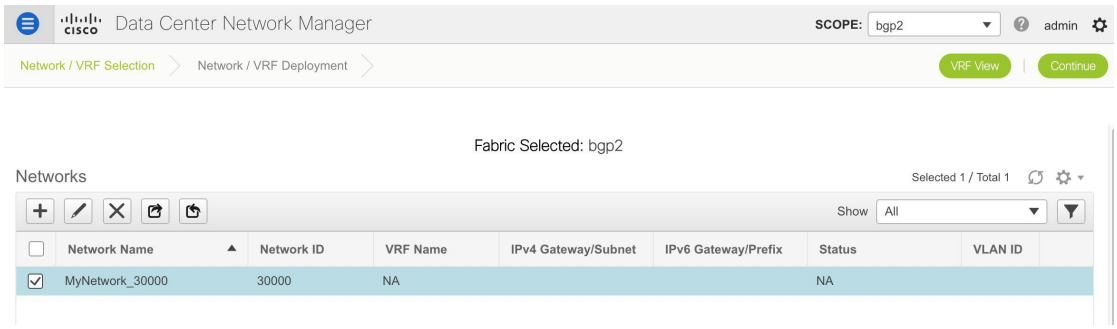

3. Select networks that you want to deploy. In this case, select the check boxes next to both the networks and click **Continue** at the top right part of the screen.

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

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

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

   At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly. Blue for *Pending* state, yellow for *In Progress* when the provisioning is in progress, green when successfully deployed, and so on. From DCNM 11.3(1), the pending state indicates that there is a pending deployment or pending recomputation. You

can click on the switch and review the pending deployments using **Preview** or **Deploy Config** options, or click **Save & Deploy** to recompute the state of the switch.

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.

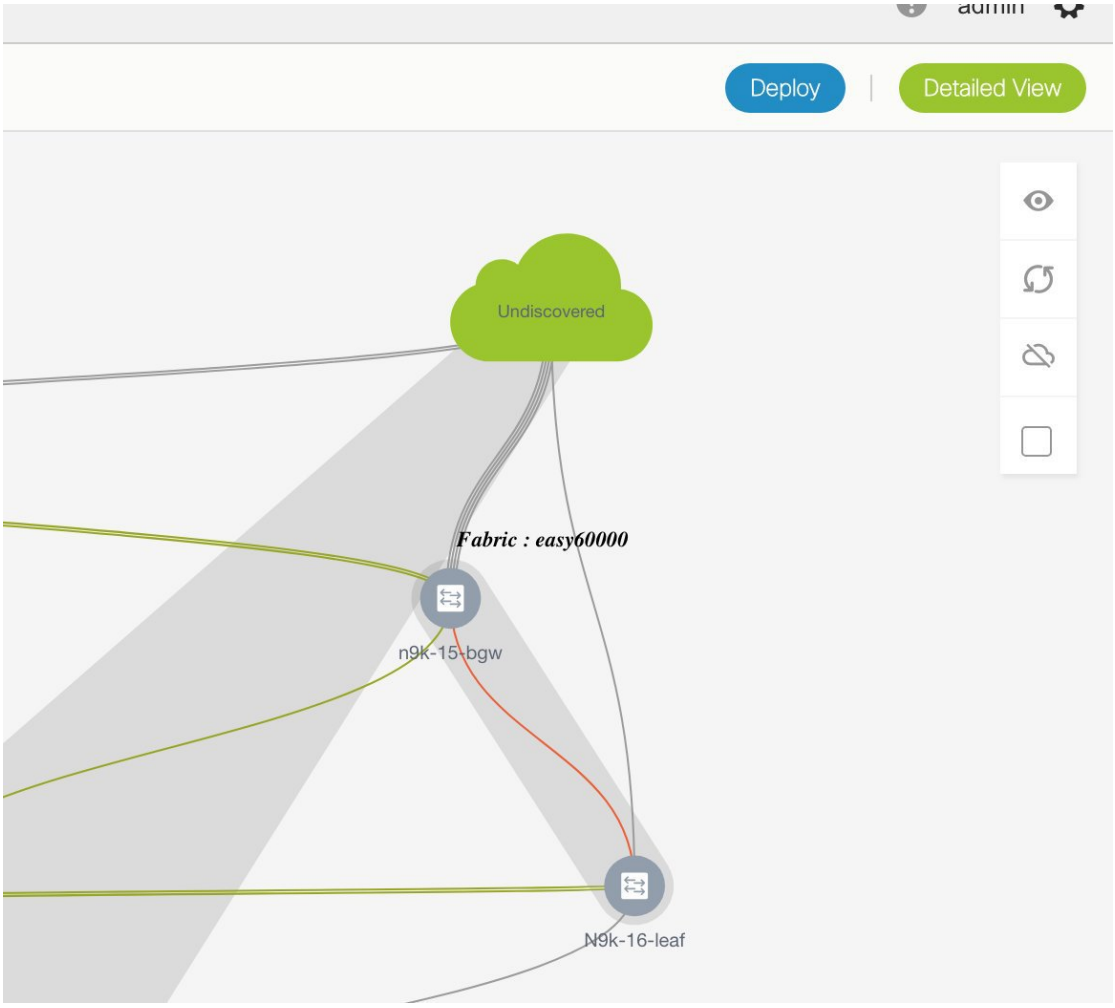

You can move the topology around the screen by clicking the left mouse button on the screen and moving it in the direction you desire. You can enlarge or shrink the switch icons proportionately by moving the cursor roller. You can also use corresponding alternatives on the touchpad.

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.

From Cisco DCNM Release 11.5(1), the **Interfaces** window doesn't list interfaces that are part of an interface group. Specifically, the trunk ports, access ports, and dot1q tunnel ports.

If you try to perform a network attachment to a switch and an interface is part of an interface group, an appropriate error is displayed.

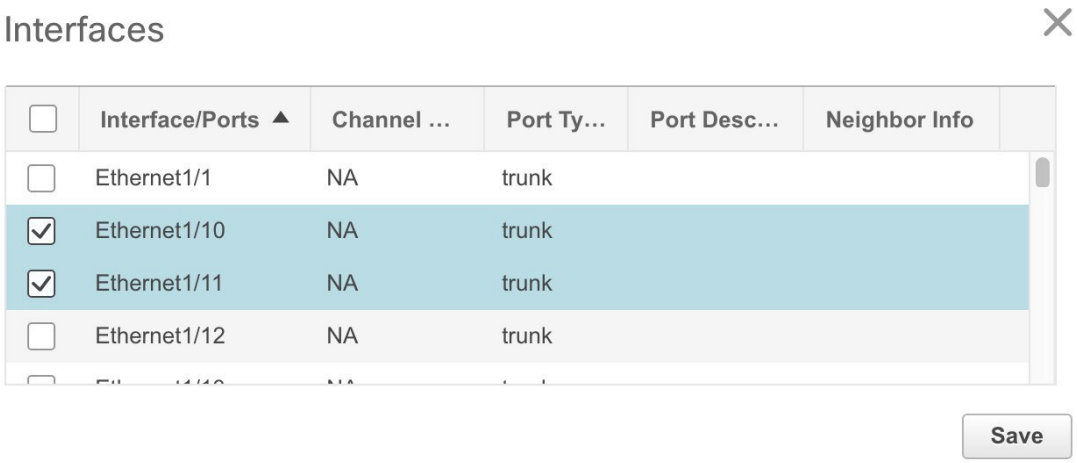

5. Double-click a switch to deploy the networks on it. For deployment of networks on multiple switches, click Multi-Select from the panel at the top right part of the screen (the topology freezes to a static state), and drag the cursor across the switches.

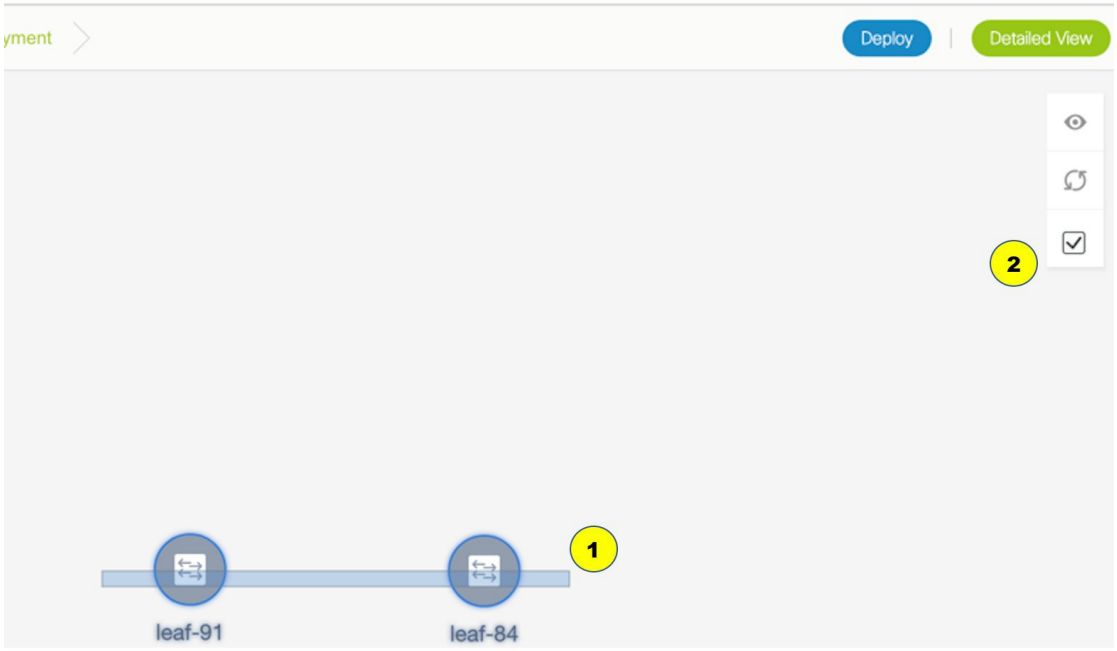

Immediately the Network Attachment dialog box appears.

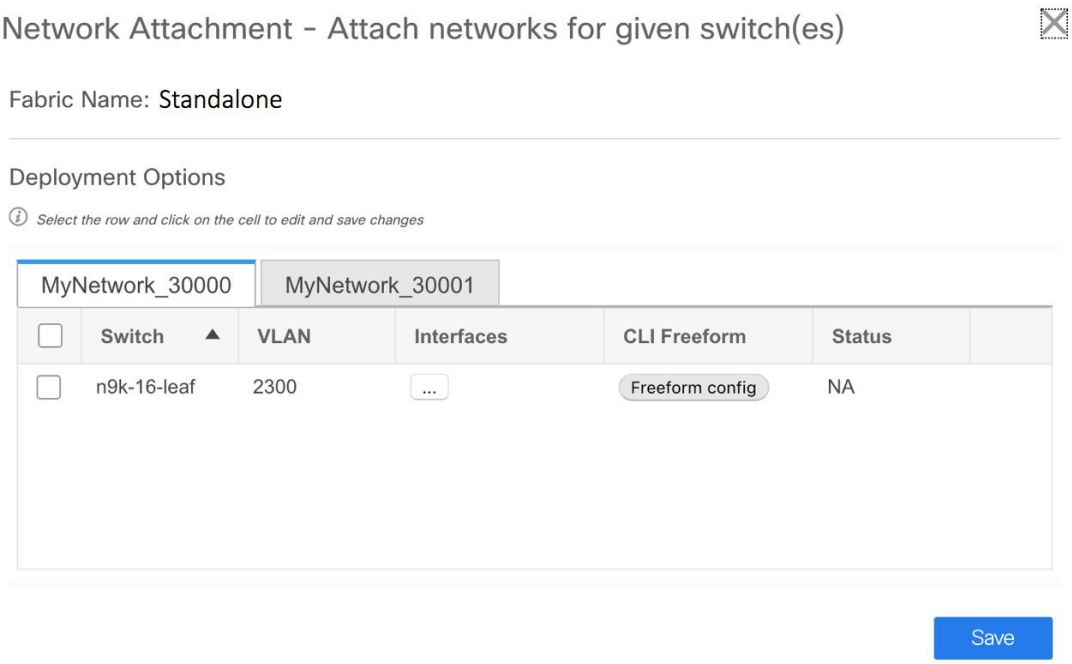

A tab represents each network (the first network is displayed by default) that is being deployed. In each network tab, the switches are displayed. Each row represents a switch.

Click the check box next to the **Switch** column to select all switches. The network is ready to be provisioned on the switches.

VLAN - Update the VLAN ID if needed.

When you update a VLAN ID and complete the network deployment process, the old VLAN is not automatically removed. To complete the process, you should go to the fabric topology screen (click **Control > Fabric Builder** and click within the corresponding fabric box to go to the screen) and use the Save and Deploy option.

When updating the VLAN ID for a given network, the original VLAN ID is not automatically removed from the attached trunk interface. In order to remove the old or original VLAN ID, you must perform **Save and Deploy** + **Config Deploy** operation from within the fabric in Fabric Builder. For this, go to the fabric topology screen (click **Control > Fabric Builder** and click within the corresponding fabric box to go to the screen) and execute the **Save and Deploy** operation. Verify that config compliance is removing the expected config, then execute **Deploy Config** operation to remove the configs.

Interfaces – Click … in the column to add interfaces associated with the selected network.

VLAN to trunk port mapping – The selected trunk ports include the VLAN as an allowed VLAN on the port.

VLAN to vPC domain mapping - If you want to associate the VLAN to port channels of a vPC domain, add the port channels from the list of interfaces. The vPC port channels include the VLAN as an allowed VLAN.

Freeform configurations – Click Freeform config to enable additional configurations on the switch. After the configurations are saved, the Freeform config button gets highlighted.

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

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

Note Addition and removal of interfaces are displayed in the **Interfaces** column of the Switches Deploy scr
Though the interface-related updates (like addition or removal of trunk ports) are provisioned on the swi
the correct configurations will not reflect in the preview screen. When you add or remove a trunk or ac
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. From DCNM 11.3(1), the pending state indicates that there is a pending deployment or pending recomputation. You can click on the switch and review the pending deployments using **Preview** or **Deploy Config** options, or click **Save & Deploy** to recompute the state of the switch.

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

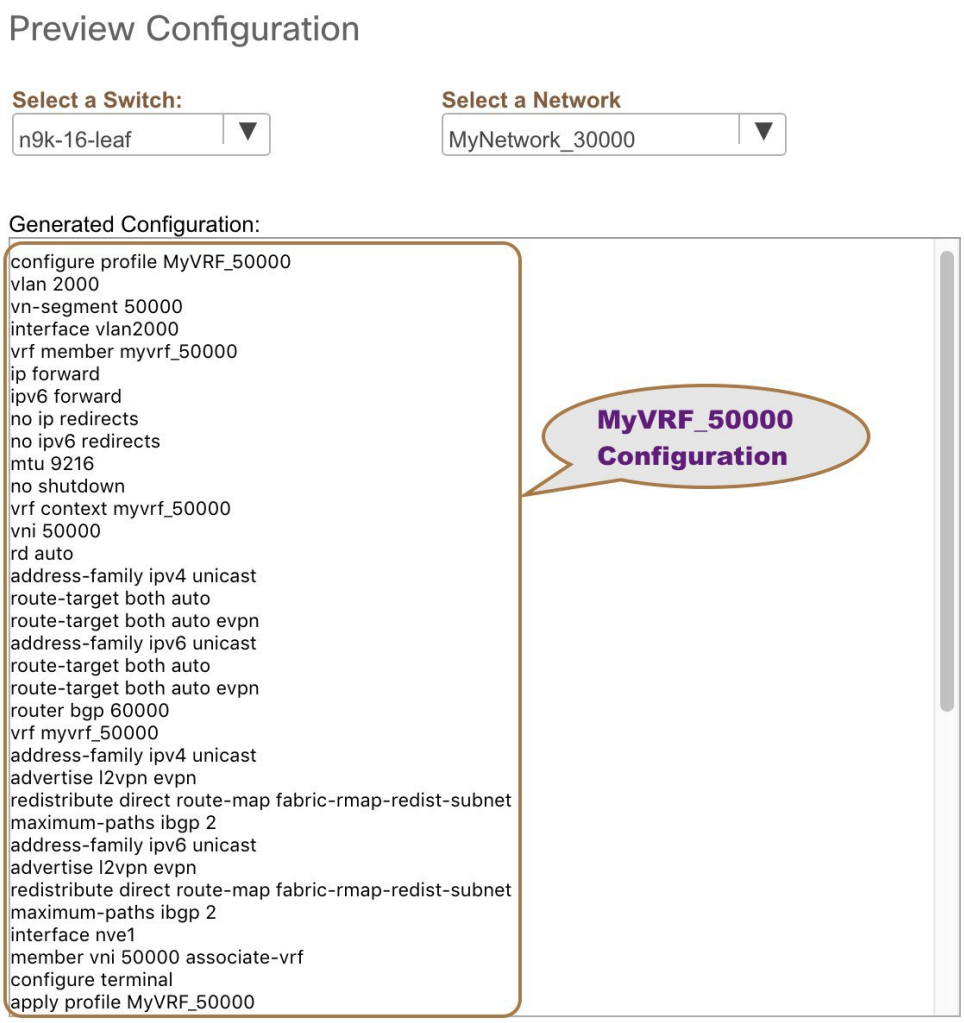
Preview Configuration
Select a Switch:
n9k-16-leaf
Select a Network
MyNetwork_30000
Generated Configuration:
configure profile MyVRF_50000
vlan 2000
vn-segment 50000
interface vlan2000
vrf member myvrf_50000
ip forward
ipv6 forward
no ip redirects
no ipv6 redirects
mtu 9216
no shutdown
vrf context myvrf_50000
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 l2vpn evpn
redistribute direct route-map fabric-rmap-redist-subnet
maximum-paths ibgp 2
address-family ipv6 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redist-subnet
maximum-paths ibgp 2
interface nve1
member vni 50000 associate-vrf
configure terminal
apply profile MyVRF_50000
MyVRF_50000
Configuration

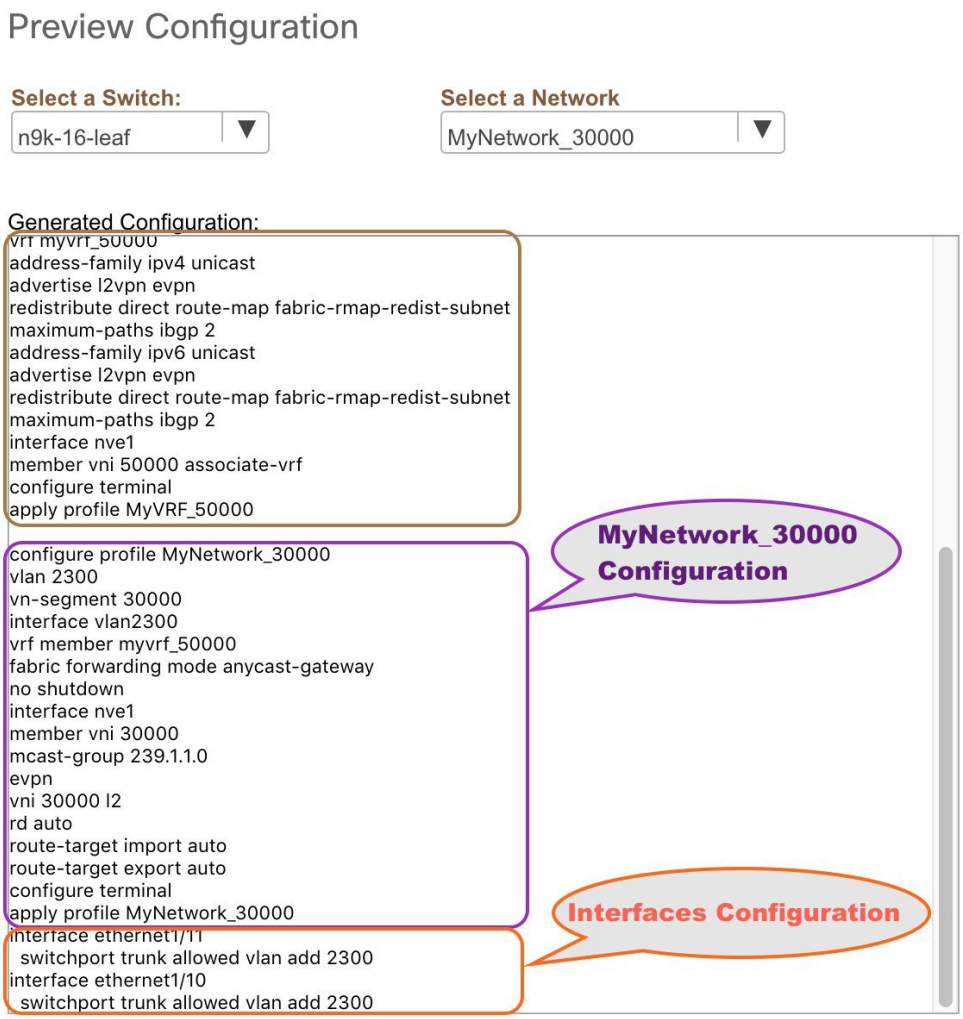

On the preview screen, you can select from the **Select a switch** and **Select a network** drop-down boxes at the top of the screen to view other network configurations.

After checking the configurations, close the screen. The Topology screen appears again.

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

Note In case you click **Deploy** and there is no configuration diff that has to be deployed, a pop-up window c up stating **No switches PENDING for deployment**.

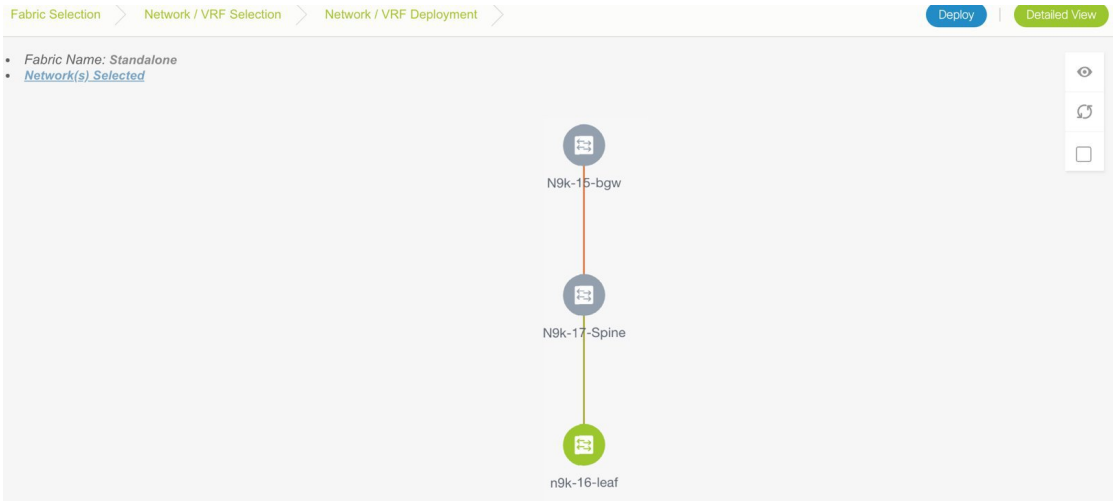

Note The status of the switch is determined by the aggregated status of the selected networks or VRFs in the following hierarchy: **Pending**, **In Progress**, **Out-of-Sync/Failed**, **In Sync/Success**, and **Unknown/NA**. F example, if any one of the networks or VRFs is in the **Out-of-Sync/Failed** status and others are not in the **Pending** or **In Progress** status, then the switch status is **Out-of-Sync/Failed**. The default status is **Unknown/NA**, when the status is not known.

Go to the Networks page to view the individual status for all networks.

Network Deployment for an MSD Fabric

Consider a scenario wherein you are deploying the same networks on different member fabric border devices. You can choose one fabric, deploy networks on its border devices, and then choose the second fabric and deploy networks.

Alternatively, you can choose the MSD fabric, and deploy the networks from a single topology view of all member fabric border devices.

This is a topology view of an MSD fabric wherein the two member fabrics topologies and their connections are depicted. You can deploy networks on the BGWs of the fabrics at once.

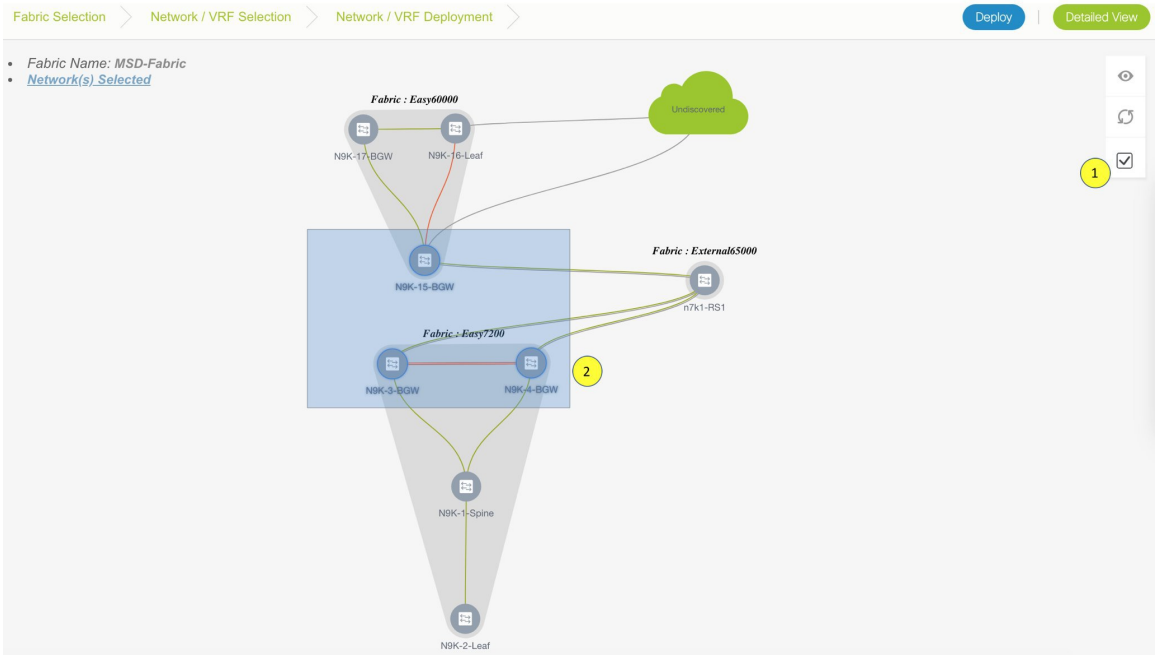

Detailed View

You can also use the Detailed View option to deploy networks and VRFs. Click **Detailed View** at the top right part of the screen. The Detailed View window appears. This lists the networks in a tabular view.

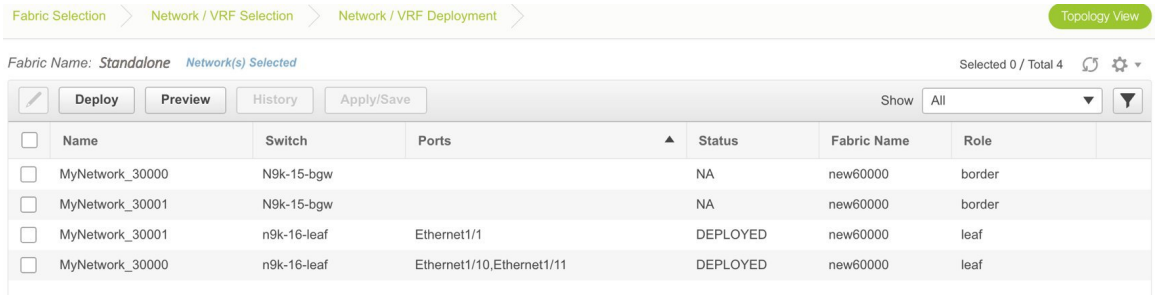

The options:

Edit - Select a network and click the Edit icon at the top left part of the screen.

Note If you select one network/switch entry and click on Edit, the Network Attach dialog box appears. To maintain consistency across the Topology View and Detailed View screens, the Network Attach screen displays all networks, and not just the selected network/switch.

Preview – Click Preview to preview configurations before deployment. You can only preview pending configurations, and not uninitiated or deployed configurations.

Deploy – Click Deploy to provision networks onto the switches.

History – Select a row and click History to view the configuration instances and status. Network and VRF-wise configurations are displayed. Click in the Status column of any instance for more details.

The fields in the table contain the configuration instance in each row, the associated switch and fabric names, the switch role, trunk ports (if any), and the deployment status.

Quick Attach – Choose a network and click **Quick Attach**. A confirmation window appears. Click **OK**. The network will be attached to the selected switch.

Quick Detach– Choose a network and click **Quick Detach**. A confirmation window appears. Click **OK**. The network will be detached from the selected switch.

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

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

Deploying VRFs for the Standalone and MSD Fabrics

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

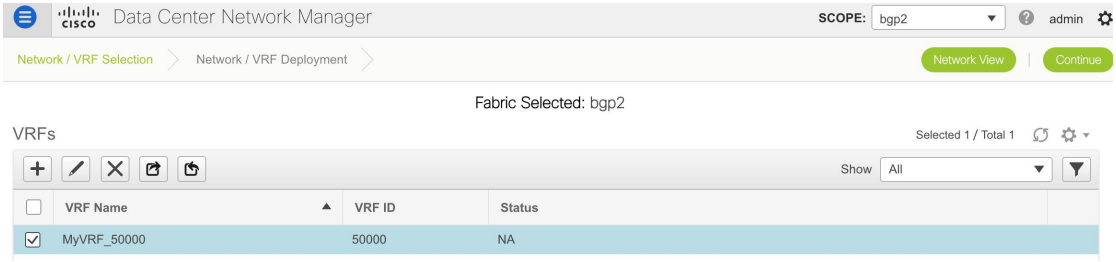

2. Select check boxes next to the VRFs that you want to deploy and click Continue at the top right part of the screen.

   The VRF Deployment screen appears. On this page, you can see the topology of the Standalone fabric. The following example shows you how to deploy the VRFs MyVRF_50000 and MyVRF_50001 on the leaf switch. You can deploy VRFs simultaneously on multiple switches but of the same role (Leaf, Border Gateway, and so on).

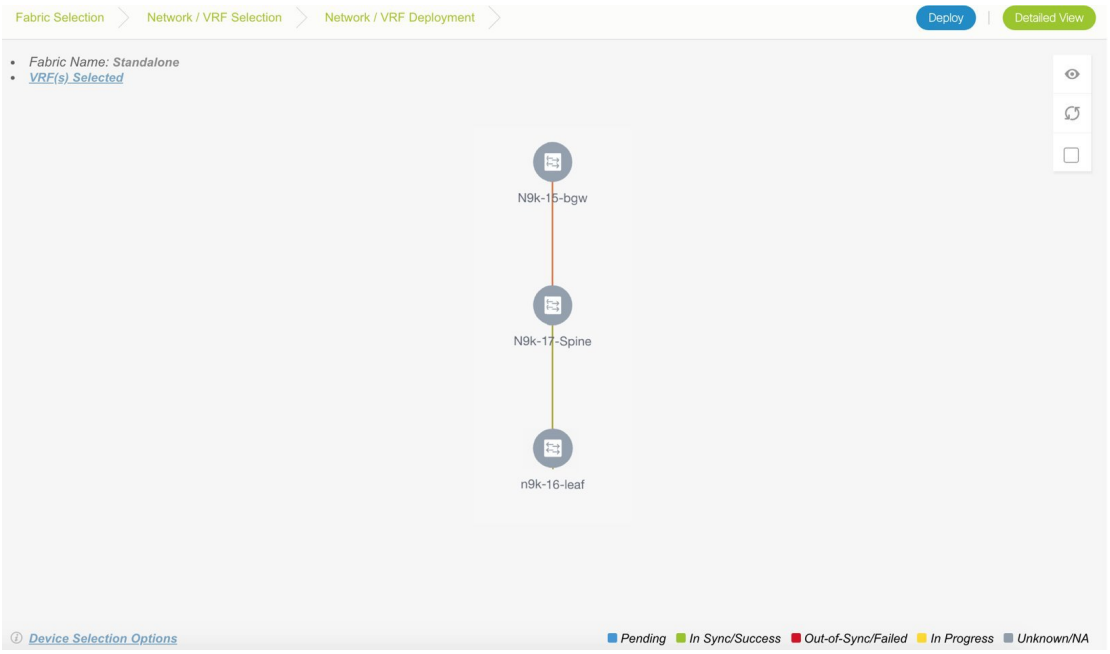

At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly. Blue for *Pending* state, yellow for *In Progress* state when the provisioning is in progress, red for failure state, green when successfully deployed, and so on. From DCNM 11.3(1), the pending state indicates that there is a pending deployment or pending recomputation. You can click on the switch and review the pending deployments using **Preview** or **Deploy Config** options, or click **Save & Deploy** to recompute the state of the switch.

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

Undiscovered cloud display – To display (or not display) an **Undiscovered** cloud in this screen, click the cloud icon in the vertical panel, at the top-right part of the screen. When you click the icon, the **Undiscovered** cloud and its links to the fabric topology are not displayed. Click the icon again to display the **Undiscovered** cloud.

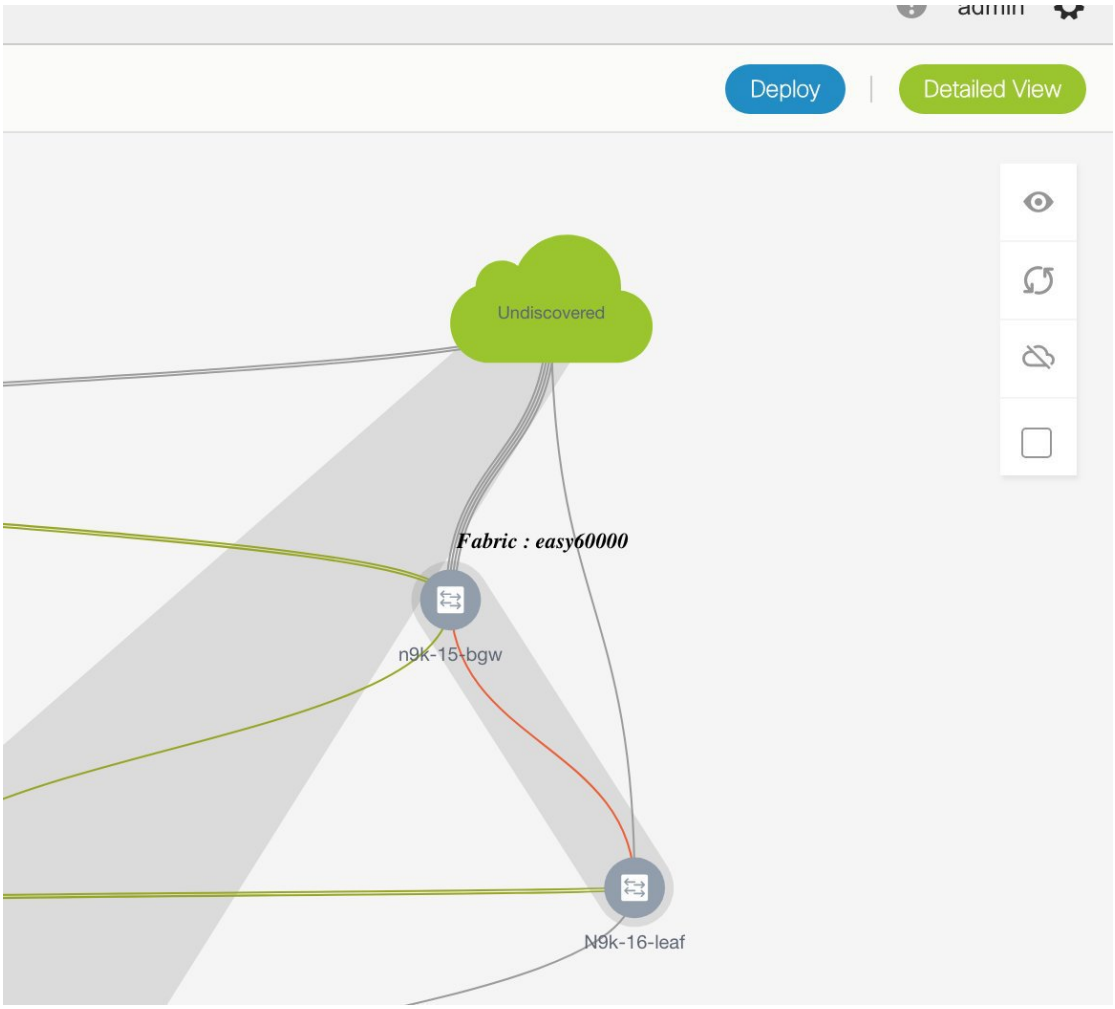

You can move the topology around the screen by clicking the left mouse button on the screen and moving it in the direction you desire. You can enlarge or shrink the switch icons proportionately by moving the cursor roller. You can also use corresponding alternatives on the touchpad.

3. Double-click a switch to deploy VRFs on it. The VRF Attachment screen comes up.

Note For deployment of VRFs on multiple switches, click the Multi-Select option from the panel at the top right part of the screen (This freezes the topology to a static state), and drag the cursor across the switches.

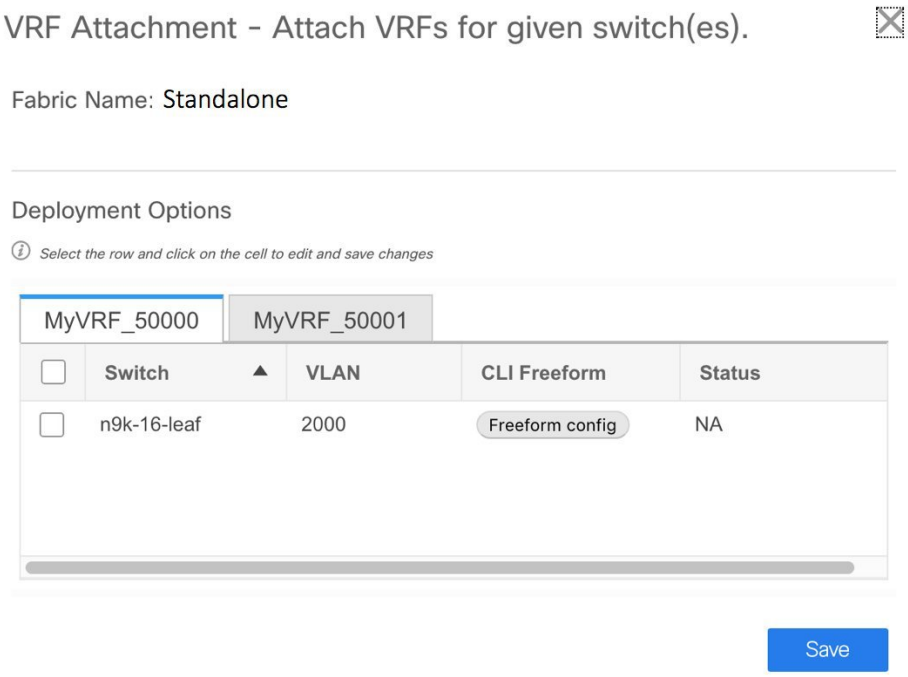

A tab represents each VRF that is being deployed (the first selected VRF is displayed by default). In each VRF tab, the selected switches are displayed. Each row represents a switch.

VLAN ID - Click within the VLAN column to update the VRF VLAN ID, if needed.

Freeform configurations – Click Freeform config to enable additional configurations on the switch. After you save freeform configurations, the Freeform config button gets highlighted.

Click the checkbox next to the Switch column to select all switches. VRF MyVRF_50000 is ready to be provisioned on the switch

4. Select the other VRF tab and make the same selections.
5. Click **Save** (at the bottom right part of your screen) to save VRF configurations.

   The topology screen comes up again. Click the *Refresh* button in the vertical panel at the top right part of the screen. The blue color on the switch icons indicates that the deployment is pending. From DCNM 11.3(1), the pending state indicates that there is a pending deployment or pending recomputation. You can click on the switch and review the pending deployments using **Preview** or **Deploy Config** options, or click **Save & Deploy** to recompute the state of the switch.

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

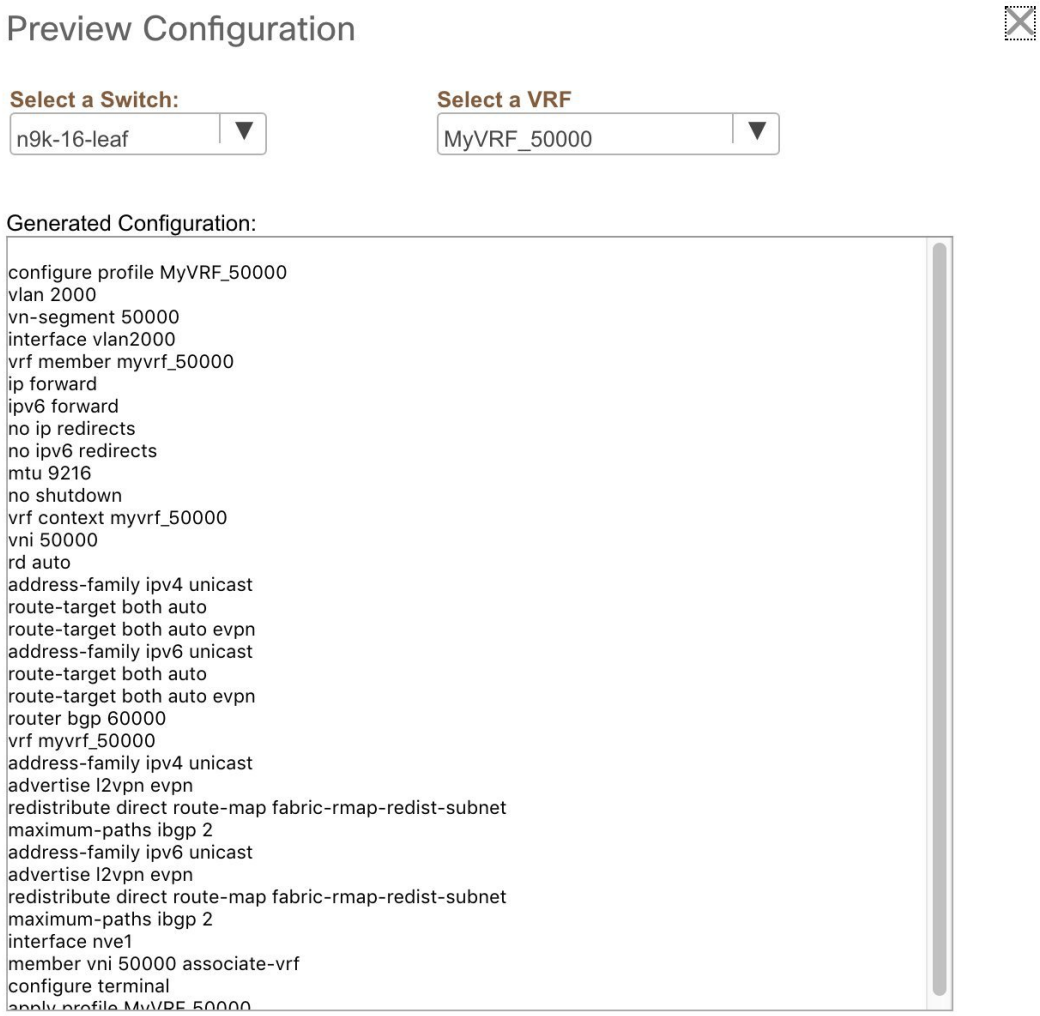

After checking the configurations, close the screen. The *Topology View* screen appears.

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

Note In case you click **Deploy** and there is no configuration diff that has to be deployed, a pop-up window com up stating **No switches PENDING for deployment**.

VRFs Deployment for an MSD Fabric

Consider a scenario wherein you are deploying the same VRFs on different member fabric border devices. You can choose one fabric, deploy VRFs on its border devices, and then choose the second fabric and deploy the VRFs.

Alternatively, you can choose the MSD fabric, and deploy the VRFs from a single topology view of all member fabric border devices at once.

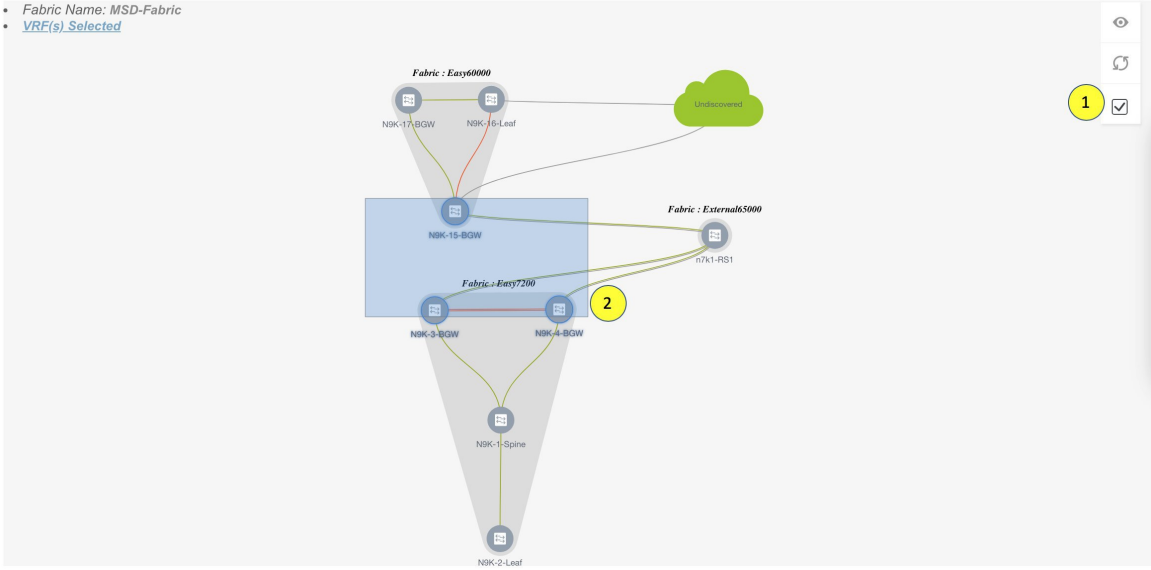

Detailed View

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

Click **Detailed View** at the top right part of the screen. The Detailed View screen comes up. This lists the VRFs in a tabular view.

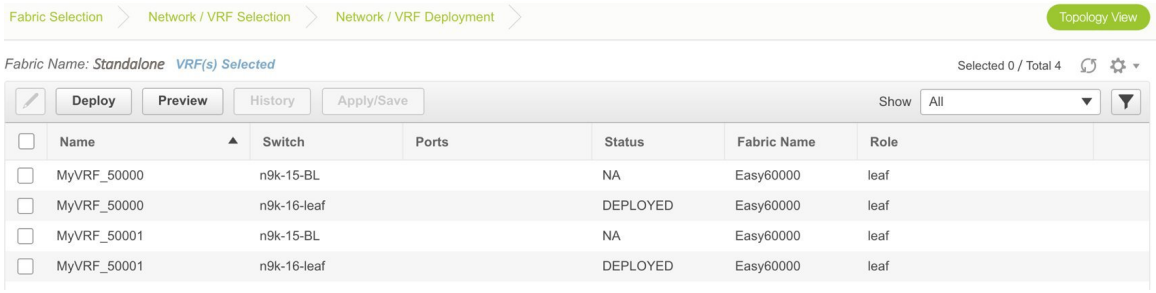

The options:

Edit - Select a VRF and click the Edit icon at the top left part of the screen.

Note If you select one VRF/switch entry, the VRF Attach screen comes up. To maintain consistency across the Topology View and Detailed View screens, the VRF Attach screen displays all VRFs, and not just the selected VRF/switch entry.

Preview – Click Preview to preview configurations before deployment. You can only preview pending configurations, and not uninitiated or deployed configurations.

Deploy – Click Deploy to provision VRFs onto the switches.

History – Select a row and click History to view the configuration instances and status. Network and VRF-wise configurations are displayed. Click in the Status column of any instance for more details.

The fields in the table contain the configuration instance in each row, the associated switch and fabric names, the switch role, and the deployment status.

Quick Attach: Choose a VRF and click **Quick Attach**. A confirmation window appears. Click **OK**. The VRF will be attached to the selected switch.

Quick Detach: Choose a VRF and click **Quick Detach**. A confirmation window appears. Click **OK**. The VRF will be detached from the selected switch.

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

Undeploying Networks for the Standalone Fabric

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

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

   The Networks screen comes up.

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

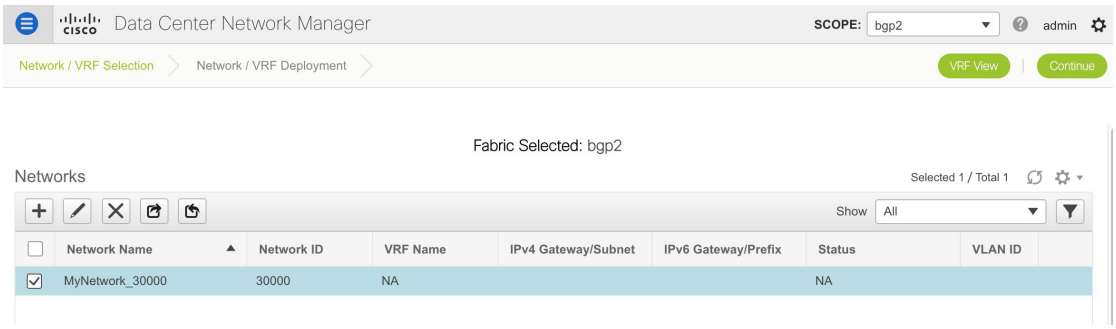

3. Select the networks that you want to undeploy and click Continue. The topology view comes up.
4. Select the Multi-Select button (if you are undeploying the networks from multiple switches), and drag the cursor across switches with the same role. The Network Attachment screen comes up.

   (For a single switch, double-click the switch and the Network Attachment screen comes up).

   (For a single switch, double-click the switch and the Switches Deploy screen comes up).

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

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

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

Undeploying VRFs for the Standalone Fabric

You can undeploy VRFs from the deployment screen. The DCNM screen flow for undeployment is similar to the deployment process flow.

1. Choose **Control > Fabrics > VRFs**.
2. Choose the correct fabric from **SCOPE**. When you select a fabric, the **VRFs** screen refreshes and lists networks of the selected fabric.
3. Select the VRFs that you want to undeploy and click **Continue**. The *Topology View* page comes up.
4. Select the Multi-Select option (if you are undeploying the VRFs from multiple switches), and drag the cursor across switches with the same role. The VRF Attachment screen comes up.

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

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

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

Deleting Networks and VRFs

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

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

Configuring Multiple VLAN IDs to a Single VNI

The following procedure shows how to tag multiple VLAN IDs to a single VNI in DCNM.

Procedure

Step 1 Navigate to **Control > Networks**.

Step 2 Select the fabric from the **Scope** drop-down list and then select the network. Click **Continue**.

Step 3 Check the **Multi-Select** check box and drag the cursor over the switches that needs to be updated with VLAN IDs.

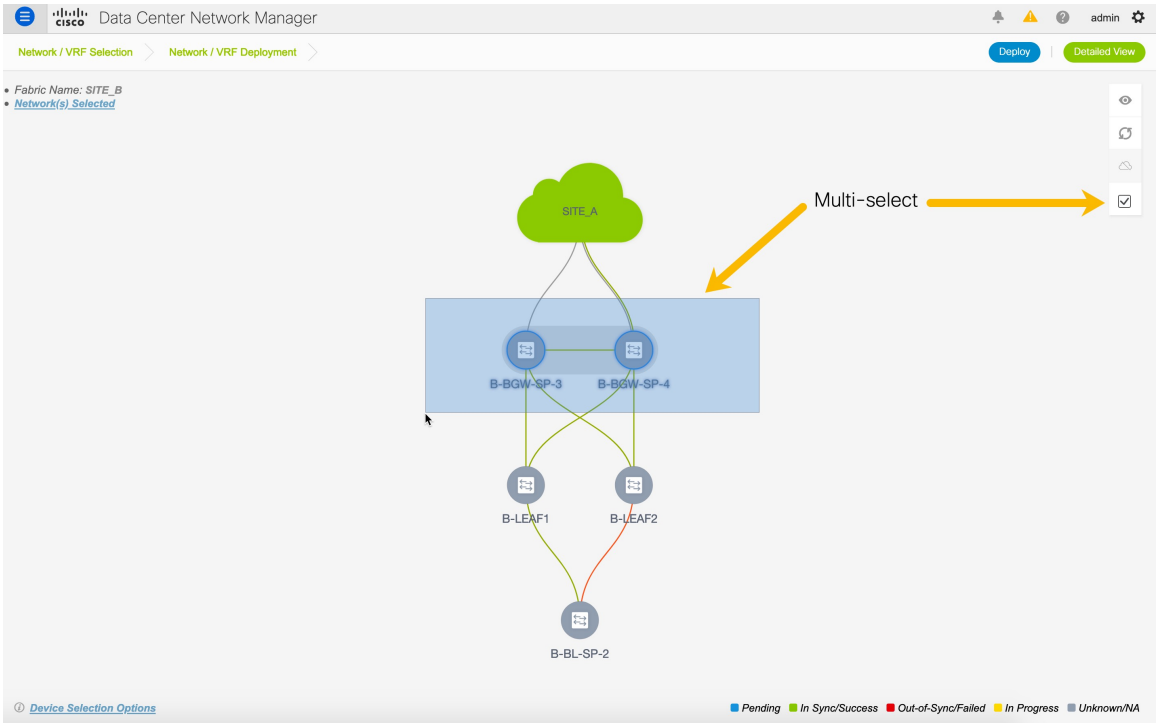

Step 4 In the **Network Attachment** window, edit the VLAN ID for the switches and click **Save**.

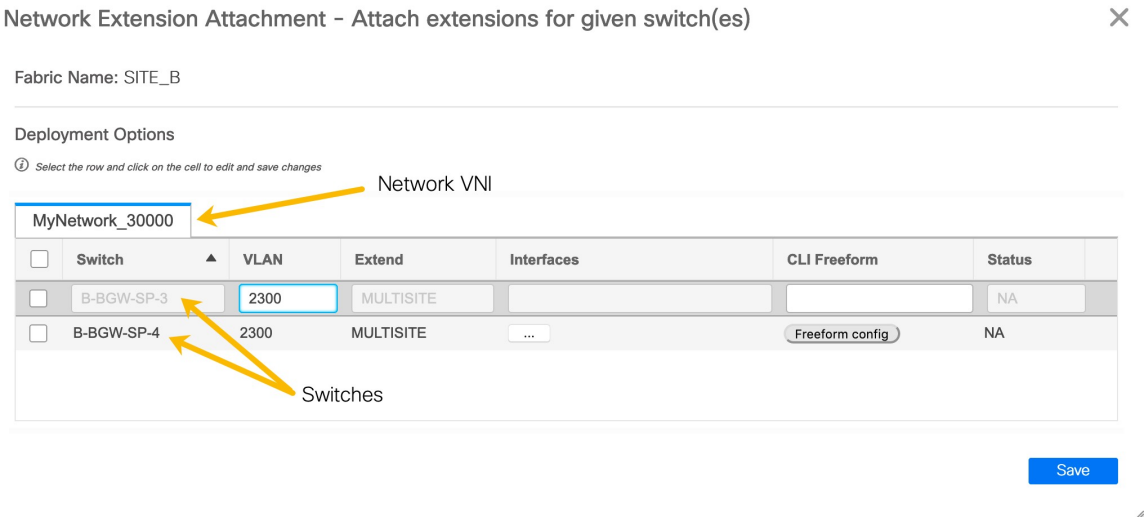

Step 5 Click **Deploy** to deploy the configuration.

Enhanced Role-based Access Control in Cisco DCNM

From Cisco DCNM Release 11.4(1), you can see the following role-based access control (RBAC) changes:

- Read-only access to the Cisco DCNM Web UI and APIs for the **network-operator** user role
- A new user role called **network-stager**.
- Freeze deployment for a particular fabric or all fabrics in DCNM as a user with the **network-admin** role.

From Cisco DCNM Release 11.5(1), you can see new user roles, **device-upg-admin**, and **access-admin** are added.

Note Actions that cannot be performed by a selected user role is grayed out.

You can also watch the video that demonstrates some of the operations performed by a network stager and how to freeze a fabric in Cisco DCNM. See the Enhanced Role-based Access Control (RBAC) video.

Device-upg-admin Role

A user with the **device-upg-admin** role can perform operations only in **Image Management** window.

See the Image Management, on page 413 section for more information.

Access-admin Role

A user with the **access-admin** role can perform operations only in **Interface Manager** window for all fabrics.

An access-admin can perform the following actions:

- Add, edit, delete and deploy layer 2 port channels, and vPC.

- Edit host vPC, and ethernet interfaces.
- Save, preview, and deploy from management interfaces.
- Edit interfaces for LAN classic, and external fabrics if it isn't associated with policy.

  Apart from nve, management, tunnel, subinterface, SVI, interface grouping, and loopback interfaces

However, a user with the access-admin role can't perform the following actions:

- Cannot edit layer 3 port channels, ST FEX, AA FEX, loopback interfaces, nve interfaces, and subinterfaces.
- Cannot edit member interfaces and port channels of Layer 3, ST FEX, AA FEX.
- Cannot edit interfaces with policy associated from underlay and link.
- Cannot edit peer link port channel.
- Cannot edit management interface.
- Cannot edit tunnel.

Note The icons and buttons are grayed out for this role when the fabric or DCNM is in deployment-freeze mode.

Network-Operator Role

A user with the **network-operator** role has access to the following menu in the DCNM Web UI:

- Dashboard
- Topology
- Monitor
- Applications

From Cisco DCNM, Release 11.4(1), a user with this role has read-only access to the **Control** menu as well.

A network operator can view fabric builder, fabric settings, preview configurations, policies, and templates. However, a network operator cannot perform the following actions:

- Cannot change expected configurations of any switch within any fabric.
- Cannot deploy any configurations to switches.
- Cannot access the administration options like licensing, creating more users, and so on.

Network-Stager Role

A user with the network-stager role can make configuration changes on DCNM. A user with the network-admin role can deploy these changes later. A network stager can perform the following actions:

- Edit interface configurations.
- View or edit policies.
- Create interfaces.

- Change fabric settings.
- Edit or create templates.

However, a network stager cannot perform the following actions:

- Cannot make any configuration deployments to switches.
- Cannot perform deployment-related actions from the DCNM Web UI or the REST APIs.
- Cannot access the administration options like licensing, creating more users, and so on.
- Cannot move switches in and out of maintenance mode.
- Cannot move fabrics in and out of deployment-freeze mode.
- Cannot install patches.
- Cannot upgrade switches.
- Cannot create or delete fabrics.
- Cannot import or delete switches.

The difference between a network operator and a network stager is as a network stager, you can only define intent for existing fabrics, but cannot deploy those configurations.

Only a network admin can deploy the changes and edits that are staged by a user with the **network-stager** role.

Viewing Policy Change History

Different users can simultaneously change expected configuration of switches in the DCNM. You can view the history of these staged changes in the **Policy Change History** tab. The deployment history captures the changes that are pushed or deployed from DCNM to switches.

Note Only deployment history is supported for non-Nexus devices.

To view the changes by different users, perform the following steps:

Procedure

Step 1 Log into Cisco DCNM with the **network-admin**, **network-stager**, or **network-operator** user role.

Step 2 Navigate to the fabric topology window.

Step 3 Right-click the switch for which you intent change history.

Step 4 Choose **History**.

Step 5 Click the **Policy Change History** tab.

Step 6 Search for the interface to which you made changes in the **Generated Config** column.

Step 7 The **PTI Operation** column will have the value **UPDATE** for the changes made by different users.

Step 8 Scroll horizontally to the **User** column. You can see the user names populated with the timestamp.

For each configurable entity, the detailed history under the **Generated Config** column, provides delta of configuration changes made by every user.

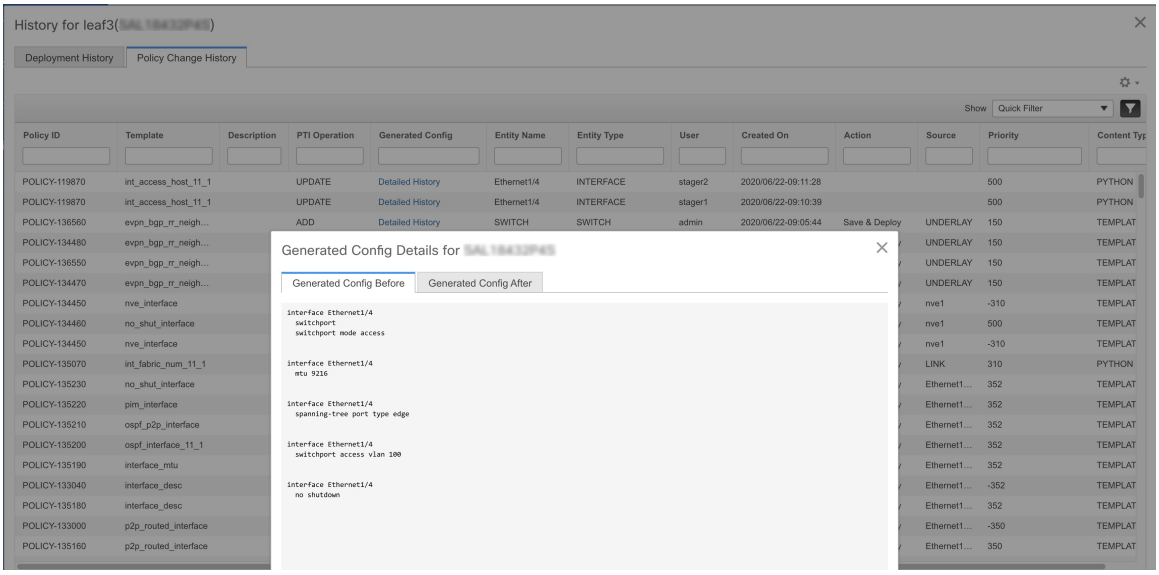

Freezing Fabrics in Cisco DCNM

As a network admin, you can disable, or freeze, deployments for LAN classic fabrics, easy fabrics, and external fabrics. Deployment freeze disables configuration or write access from the DCNM to the switches. When you freeze a fabric, switches cannot be reloaded, moved in and out of maintenance mode, and you cannot add or delete switches within the fabric. This feature provides complete control for a network admin to disable inadvertent changes to the physical network from the DCNM, unless a maintenance window is scheduled.

Freezing a Fabric

To disable deployment for a fabric from Cisco DCNM Web UI, perform the following:

Procedure

Step 1 Navigate to the **Fabric Builder** window or the fabric topology window.

Step 2 Click the spanner (🔧) icon.

The spanner icon is next to the fabric name in the fabric topology window. A confirmation window appears asking you if you want to disable all deployments for the fabric.

Step 3 Click **Yes**.

Note When you hover over the spanner icon before freezing the fabric, the tooltip will read **Deployment Enabled**. When you hover over the spanner icon after freezing the fabric, the tooltip will read **Deployment Disabled**.

After you disable deployments or freeze a fabric, you can only save, edit, or preview changes but not deploy them. All deploy related actions from the DCNM to this fabric will be grayed out.

To enable all deployments for the fabric, click the same spanner ( ) icon and unfreeze the fabric.

Freezing All Fabrics

In addition to the per-fabric deployment freeze knob, the network admin can freeze deployments for all fabrics within the DCNM at the same time.

To freeze all fabrics in your DCNM setup from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Administration > DCNM Server > Server Properties**.

Step 2 Search for the **DEPLOYMENT_FREEZE** field.

Step 3 Set the value as **true**.

The default value is **false**.

Note When you freeze DCNM, you cannot deploy any changes to switches. However, users with appropriate roles, like the network-admin role or the network-stager role, with appropriate access, can make changes in the DCNM for deployment at a later stage.

Actions that cannot be performed when you freeze a fabric or DCNM are grayed out.

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. However, if such external fabrics are member fabrics of an MSD fabric, the backup is taken at MSD-fabric level.

Note From Cisco, DCNM Release 11.4(1), you can take a backup of external fabrics in monitor-only mode, but can't restore them. You can restore this backup when the external fabric isn't in monitor-only mode.

- 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.

From Cisco DCNM Release 11.4(1), the backup captures the intent related to IFC as well. 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`

The backed-up config files can be found in the corresponding directory with the fabric name. Each backup of a fabric is treated as a different version, regardless if it is backed up manually or automatically. You can find all versions of the backup in the corresponding fabric directories. Hence, the backed up intent configuration file, running configuration file and PTIs can be found at location: `/usr/local/cisco/dcm/dcnm/data/archive/<fabric_name>/Version_x`, where `x` is the version number. The valid value is between 1 and the limit you set in the **archived.versions.limit** field. The default value is 50, which means only 50 backups are archived, and the oldest backups are removed. The minimum value is 10. If you specify a value lesser than 10, it will be overwritten to 10. 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.

You can also watch the video that demonstrates how to back up and restore an MSD fabric in Cisco DCNM. See the MSD Fabric Backup and Restore video.

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.

Note MSD fabrics don't support 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**.

DCNM initiates the backup process after you click **Save**.

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. You cannot initiate standalone backups for a member fabric on an MSD fabric.

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.

The manual backup is indicated in midnight blue. When you hover over the backup, the name has the tag you mentioned in *Step 4* confirming that it's a manual backup.

Golden Backup

You can now mark the backups that you don't want to delete even after you reach the archiving limit. These backups are the golden backups. You can't delete golden backups of fabrics. However, Cisco DCNM archives only up to 10 golden backups. You can mark a backup as golden backup while restoring the fabric. To mark a backup as golden 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 Click **Restore Fabric** from the **Actions** menu.

The **Restore Fabric** window appears.

Step 3 Choose the time period from where you want to choose the backup.

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 choose 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 Choose the backup you want to mark as golden by clicking the backup.

You can choose the automatic or manual backup. These backups are color-coded. Automatic backups are indicated in blue color. Manual backups are indicated in midnight blue color. Golden backups are indicated in orange color. The automatic backups have only the versions in their names. Whereas the manual backups have tag names, which you gave when you initiated a manual backup, along with the version in the backup name. Hover over a backup to see the name. The automatic backup is initiated from the **Configuration Backup** tab in the **Fabric Settings** dialog box. The manual backup is initiated by clicking **Backup Now** from the **Actions** pane in the fabric topology window.

Step 5 Check the **Mark backup as golden backup** check box to mark the backup as a golden backup.

A confirmation dialog box appears.

Step 6 Click **Yes**.

Step 7 Continue with rest of the fabric restore procedure as mentioned in the *Restoring Fabrics* section or exit the window.

Validating Backups

DCNM validates all the backups when you initiate a fabric restore process. The validation includes the following checks:

- The DCNM release from which you want to restore: You can restore backups only from Cisco DCNM, Release 11.3(1), and Cisco DCNM, Release 11.4(1). Hence, if you upgrade from Cisco DCNM, Release 11.3(1) to Cisco DCNM, Release 11.4(1), you can restore a backup, which you archived before upgrading.
- Member fabrics composition: DCNM checks the name or ID of the member fabrics of an MSD fabric. If you change them after you back up, the restore won't proceed.
- Template validation: DCNM checks if templates from the backup match the templates in the current version. If you delete or rename any templates, you can't proceed with the restoring.
- Device composition of a fabric: If there are any changes to the inventory of switches after you back them up, you can't restore.

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 Choose the backup you want to restore.

You can choose the automatic or manual backup. These backups are color-coded. Automatic backups are indicated in blue color. Manual backups are indicated in midnight blue color. Golden backups are indicated in orange color. The automatic backups have only the versions in their names. Whereas the manual backups have tag names, which you gave when you initiated a manual backup, along with the version in the backup name. Hover over a backup to see the name. The automatic backup is initiated from the **Configuration Backup** tab in the **Fabric Settings** dialog box. The manual backup is initiated by clicking **Backup Now** from the **Actions** pane in the fabric topology window.

Note If the fabric was a member of an MSD fabric and a backup was taken at the MSD-fabric level, that backup doesn't appear here. Only the standalone backups of the fabric taken before it was part of an MSD fabric appear here.

Step 5 Check the **Mark backup as golden backup** check box to mark the backup as a golden backup.

Step 6 Click **Next** to see the selected backup information of the devices in sync.

The switch name, switch serial number, IP address, and the delta 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 7 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 8 Go back to **View Backup Summary** window.

Step 9 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 10 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 11 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 12 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 Choose the backup you want to restore.

You can choose the automatic or manual backup. These backups are color-coded. Automatic backups are indicated in blue color. Manual backups are indicated in midnight blue color. Golden backups are indicated in orange color. The automatic backups have only the versions in their names. Whereas the manual backups have tag names, which you gave when you initiated a manual backup, along with the version in the backup name. Hover over a backup to see the name. The automatic backup is initiated from the **Configuration Backup** tab in the **Fabric Settings** dialog box. The manual backup is initiated by clicking **Backup Now** from the **Actions** pane in the fabric topology window.

Note If the fabric was a member of an MSD fabric and if any backup was taken for the MSD fabric, that backup does not appear here. Only the standalone backups of the fabric taken before it was part of an MSD fabric appear here.

Step 5 (Optional) Check the **Mark backup as golden backup** check box to mark the backup as a golden backup.

Step 6 Click **Next** 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 7 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 8 Go back to the **View Backup Summary** window.

Step 9 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 10 Click **Close** after the restore process is complete.

Restoring MSD Fabrics

When you restore an MSD fabric, the overlay information related to the MSD fabric is restored before restoring information related to the child fabrics. If there's any change in the inventory of the MSD fabric, the backup is considered to be invalid and the restore is blocked. You can't initiate a restore process for a member fabric. You get an error stating that the fabric is currently a member fabric of an MSD fabric. Move the member fabrics out of the MSD fabric to restore the previous standalone backups. Restoring an MSD fabric involves restoring fabric intent, underlay or interface intent, overlay intent, and intent regeneration.

You can also watch the video that demonstrates how to backup and restore an MSD fabric in Cisco DCNM. See the MSD Fabric Backup and Restore video.

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**.

Step 2 Choose an MSD fabric.

Step 3 Click **Restore Fabric** from the **Actions** menu.

The **Restore Fabric** wizard appears and you will be in the **Select Backup** step.

Note This option is not available for member fabrics, of an MSD fabric, from its corresponding fabric topology window.

Step 4 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 5 Choose the backup you want to restore.

You can choose the automatic or manual backup. These backups are color-coded. Automatic backups are indicated in blue color. Manual backups are indicated in midnight blue color. Golden backups are indicated in orange color. The automatic backups have only the versions in their names. Whereas the manual backups have tag names, which you gave when you initiated a manual backup, along with the version in the backup name. Hover over a backup to see the name. The automatic backup is initiated from the **Configuration Backup** tab in the **Fabric Settings** dialog box. The manual backup is initiated by clicking **Backup Now** from the **Actions** pane in the fabric topology window.

Note The standalone backups that you took for a member fabric before importing it to the MSD fabric do not appear here. Only the MSD backups appear here.

Step 6 Click the backup you want to restore.

The **Backup Summary** area appears. It includes the following information:

- Backup taken on: Timestamp when the backup was taken.
- DCNM version: DCNM version when the backup was taken.
- Backup version: Version of the backup, including the tag name if it's a manual backup.
- Total number of Fabrics: Specifies the total number of member fabrics imported in the MSD fabric.
- Total number of Easy Fabrics: Specifies the number of member fabrics that are easy fabrics.
- Total number of External Fabrics: Specifies the number of member fabrics that are external fabrics.
- Total number of devices: Specifies the total number of switches in all member fabrics.
- Number of devices in out of sync status: Specifies the number of devices that aren't in sync.
- Number of devices in unknown status: Specifies the number of devices whose status isn't known.
- Member fabrics: Specifies the names of the member fabrics Mark back as golden backup check box: (Optional) Check the **Mark backup as golden backup** check box to mark the backup as a golden backup.

Note If there are any devices with Out-of-Sync or Unknown status, the restore process will be blocked.

Step 7 Click **Next** to move to the **Restore Preview** step.

The **Easy Fabric** tab has information about the switch name, fabric name, switch serial, IP address, and the delta configuration of the member easy fabric. The **External Fabric** tab contains information about the switch name, fabric name, switch serial, IP address, switch status, configuration, and if the restore is supported for the member external fabric.

Note The backup isn't valid if devices are added or removed from the fabric. You can restore only the valid backups.

Step 8 Click **Restore Intent** to proceed to the **Restore Status** step in restoring.

The restore status and description appears for the member fabrics. Click the member fabric radio button to view the fabric-level progress of that fabric. The progress is automatically updated every 5 seconds.

Step 9 Click **Next** after the status is successful.

The **Configuration Preview** window appears. You can view the details of the switch name, IP address, switch serial number, preview configuration, status, and the progress in this window.

Note
- You can click **Next** only if the status is **Completed**.
- You can't go back to the previous step because the fabric configurations change.
- If the restoring failed, the fabric will be rolled back to the previous configuration.

Step 10 Click **Deploy** to deploy the restored configuration.

The **Configuration Deployment Status** window appears. You can view the following details:

- Switch name
- IP address
- Status

- Status description
- Progress

Step 11 Click **Close** after the restoring process is complete.

Restoring a Switch

From Cisco DCNM, Release 11.5(1), you can restore a Cisco Nexus switch in external fabrics and LAN classic fabrics from the Cisco DCNM Web UI. The information you restore at switch-level is extracted from the fabric-level backups. The switch-level restoring doesn't restore fabric-level intents and other configurations applied using the fabric settings. Only switch-level intents are restored. Therefore, after you restore a switch, it might go out-of-sync because the fabric-level intents aren't restored. Perform a fabric-level restore to restore the intents as well. You can restore only one switch at a time. You can't restore a switch if the fabric where it's discovered is part of an MSD fabric.

To restore a switch in Cisco DCNM, perform the following steps from the Cisco DCNM Web UI:

Procedure

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

Step 2 Choose an external fabric or LAN classic fabric.

Step 3 Right-click a Cisco Nexus switch for which you want to restore the configurations.

Step 4 Choose the **Restore Config** option.

Alternatively, you can click **Tabular view** in the **Actions** pane and navigate to the **Switches** tab. Choose a Cisco Nexus switch by checking the check box and click **Restore**.

For non-Nexus switches, the **Restore Config** option doesn't appear and the **Restore** button grays out.

This option does not appear when you log in with the **network-operator** role or when the fabric is in monitor mode or freeze mode.

The **Restore Switch** wizard appears and you are in the **Select Backup** step.

Step 5 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.

Step 6 Choose the backup you want to restore.

You can choose the automatic, manual, or golden backup. These backups are color-coded. Automatic backups are indicated in blue color. Manual backups are indicated in midnight blue color. Golden backups are indicated in orange color. The automatic backups have only the versions in their names. Whereas the manual backups have tag names, which you gave when you initiated a manual backup, along with the version in the backup name. Hover over a backup to see the name. The automatic backup is initiated from the **Configuration Backup** tab in the **Fabric Settings** dialog box. The manual backup is initiated by clicking **Backup Now** from the **Actions** pane in the fabric topology window.

Step 7 Click the backup you want to restore.

The **Backup Summary** area appears. It includes the following information:

- Backup taken on: Timestamp when the backup was taken.
- DCNM version: DCNM version when the backup was taken.
- Backup version: Version of the backup, including the tag name if it's a manual backup.
- Total number of devices: Specifies the total number of switches in the fabric when the backup was taken.
- Number of devices in sync status: Specifies the number of devices that are in sync.
- Number of devices in out of sync status: Specifies the number of devices that aren't in sync.
- Number of devices in unknown status: Specifies the number of devices whose status isn't known.
- Mark backup as golden backup check box: (Optional) Check the **Mark backup as golden backup** check box to mark the backup as a golden backup. If you mark a backup as golden backup, the fabric-level backup is also marked as a golden backup.

Note Most of this information is at the fabric level, and may or may not directly impact the proceedings of the switch-level restore.

Step 8 Click **Next** to move to the **Restore Preview** step.

You can view information about the switch name, switch serial, IP address, status, restore supported, delta configuration and the VRF details.

Step 9 (Optional) 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, which is the expected configuration.

Step 10 Click **Restore** to proceed to the **Restore Status** step in restoring.

The restore status and description appears for the switch.

Step 11 Click **Close** after the restoring process is complete.

Note
- You can't go back to the previous step because the fabric configurations change.
- If the restoring failed, the switch rolls back to the previous configuration.

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 11.5(1) Upgrade for VXLAN BGP EVPN, External, and MSD Fabrics

Note the following guidelines after you upgrade to the DCNM Release 11.5(1):

- As part of the upgrade from an earlier DCNM release, the fabric and associated templates are carried over to the DCNM Release 11.5(1).
- From DCNM 11.3(1), some of the policy templates from earlier DCNM releases are deprecated and are actively updated in each newer DCNM release. These policy templates are automatically removed after the upgrade if they are not found to be in use. This removal does not affect any operations and helps in reducing the number of policies displayed in the DCNM template library.
- Navigate to each fabric from the **Fabric Builder** window, 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 324.

Caution Some configuration changes can be expected as part of this step. Therefore, perform it only during a scheduled maintenance window.

- Post DCNM upgrade from Release 11.2(1), you could see the following diff if the fabric has a border device (border, border spine, border gateway, etc):

```
route-map extcon-rmap-filter-v6 deny 20
  no match ip address prefix-list host-route-v6
route-map extcon-rmap-filter-v6 deny 20
  match ipv6 address prefix-list host-route-v6
```

  The above config is expected and it is meant to correct the route-map definition. Deployment of this diff will correct the switch configuration. If the fabric was created as Greenfield before upgrade, no additional action is needed. If the fabric was created as Brownfield before upgrade with the wrong route-map configuration on the device, this config will be captured in a **switch_freeform** policy. Post upgrade, you should edit the freeform policy to remove the CLI **match ip address prefix-list host-route-v6** before the deployment.
- After a multi-level upgrade from Cisco DCNM 10.4(2) or 11.0(1), you can change the VRF templates to **Default_VRF_Universal** or **Default_VRF_Extension_Universal** to enable **ipv6 address use-link-local-only**.

Changing ISIS Configuration from Level 1 to Level 2

This procedure shows how to change ISIS configuration on switches from Level 1 to Level 2 in a VXLAN fabric deployment.

1. Choose **Control** > **Fabrics** > **Fabric Builder**.
2. Click a fabric in the **Fabric Builder** window.
3. Click **Tabular view** under **Actions** menu.
4. Search for all the **base_isis** policies in the **Template** search field.
5. Select all the **base_isis** policies and click the **Delete** icon to delete policies
6. Click **Save & Deploy**.

After all the **base_isis** policies are deleted, DCNM considers the migrated brownfield fabric as a greenfield fabric and creates the **base_isis_level2** policies on the switches.

Configuration Compliance in DCNM

The entire intent or expected configuration defined for a given switch is stored in DCNM. When you want to push this configuration down to one or more switches, the configuration compliance (CC) module is triggered. CC takes the current intent, the current running configuration, and then comes up with the set of configurations that are required to go from the current running configuration to the current expected config so that everything will be In-Sync.

When performing a software or firmware upgrade on the switches, the current running configuration on the switches is not changed. Post upgrade, if CC finds that the current running configuration does not have the current expected configuration or intent, it reports an Out-of-Sync status. There is no auto deployment of any configurations. You can preview the diffs that will get deployed to get one or more devices back In-Sync.

With CC, the sync is always from the DCNM to the switches. There is no reverse sync. So, if you make a change out-of-band on the switches that conflicts with the defined intent in DCNM, CC captures this diff, and indicates that the device is Out-of-Sync. The pending diffs will undo the configs done out-of-band to bring back the device In-Sync. Note that such conflicts due to out-of-band changes are captured by the periodic CC run that occurs every 60 mins by default, or when you click the RESYNC option either on a per fabric or per switch basis. Note that you can also capture the out-of-band changes for the entire switch by using the CC REST API. For more information, see *Cisco DCNM REST API Guide, Release 11.2(1)*.

From Cisco DCNM Release 11.2(1), to improve ease of use and readability of deployed configurations, CC in DCNM has been enhanced with the following:

- All displayed configurations in DCNM are easily readable and understandable.
- Repeated configuration snippets are not displayed.
- Pending configurations precisely show only the diff configuration.
- Side-by-side diffs has greater readability, integrated search or copy, and diff summary functions.

The CC engine computes diff by comparing the intent with the running configuration on the switch ensuring that any configuration that is defined in the intent exists on the switch. For any component or configuration snippet that is defined in the intent, the CC engine ensures that the same component or configuration snippet exists on the switch by generating appropriate commands, if required, to match the switch configuration with the intent configuration.

Top-level configuration commands on the switch that do not have any associated DCNM intent are not checked for compliance by Configuration Compliance (CC). However, CC performs compliance checks, and attempts removal, of the following commands even if there is no DCNM intent:

- **configure profile**
- **apply profile**
- **interface vlan**
- **interface loopback**
- **interface Portchannel**
- Sub-interfaces, for example, **interface Ethernet** *X/Y.Z*
- **fex**
- **vlan** <*vlan-ids*>

CC performs compliance checks, and attempts removal, of these commands only when *Easy_Fabric_11_1* and *Easy_Fabric_eBGP* fabric templates are used. On *External_Fabric* templates, top-level configuration commands on the switch, including the commands mentioned above, that do not have any associated DCNM intent are not checked for compliance by CC.

We recommend using the DCNM freeform configuration template to create additional intent and deploy these commands to the switches to avoid unexpected behavior

Now, consider a scenario in which the configuration that exists on the switch has no relationship with the configuration defined in the intent. Examples of such configurations are a new feature that has not been captured in the intent but is present on the switch or some other configuration aspect that has not been captured in the intent. Configuration compliance does not consider these configuration mismatches as a diff. In such cases, Strict Configuration Compliance ensures that every configuration line that is defined in the intent is the only configuration that exists on the switch. However, configuration such as boot string, rommon configuration, and other default configurations are ignored during strict CC checks. For such cases, the internal configuration compliance engine ensures that these config changes are not called out as diffs. These diffs are also not displayed in the **Pending Config** window. But, the Side-by-side diff utility compares the diff in the two text files and does not leverage the internal logic used in the diff computation. As a result, the diff in default configurations are highlighted in red in the **Side-by-side Comparison** window.

Starting from Cisco DCNM Release 11.4(1), such diffs are not highlighted in the **Side-by-side Comparison** window. The auto-generated default configuration that is highlighted in the **Running config** window is not visible in the **Expected config** window.

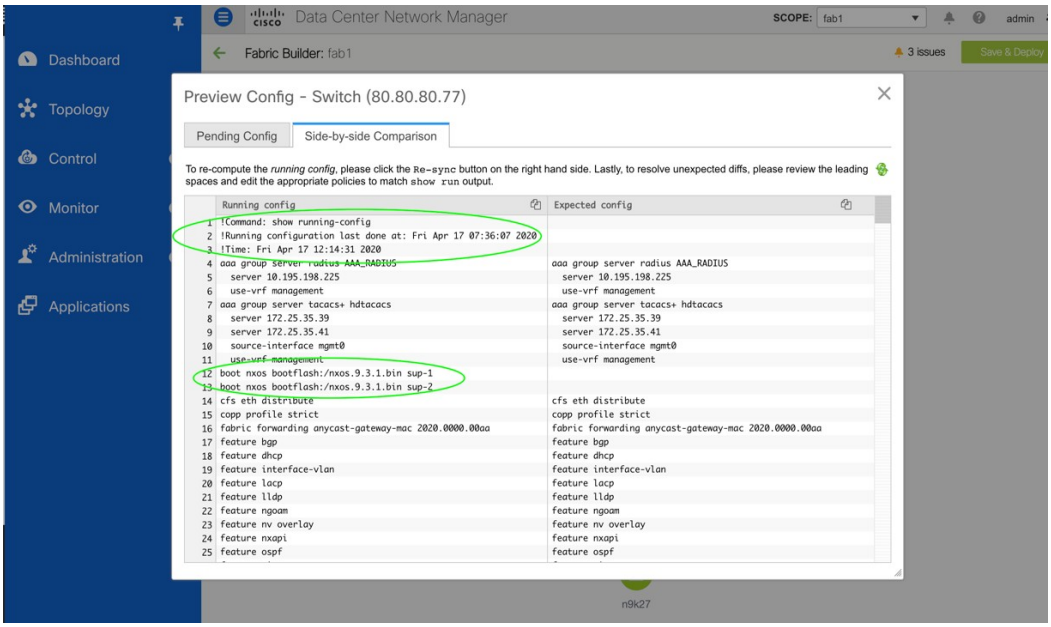

Any configurations that are shown in the **Pending Config** window are highlighted in red in the **Side-by-side Comparison** window if the configurations are seen in the **Running config** window but not in the **Expected config** window. Also, any configurations that are shown in the **Pending Config** window are highlighted in green in the **Side-by-side Comparison** window if the configurations are seen in the **Expected config** window but not in the **Running config** window. If there are no configurations displayed in the **Pending Config** window, no configurations are shown in red in the **Side-by-side Comparison** window.

All freeform configurations have to strictly match the **show running configuration** output on the switch and any deviations from the configuration will show up as a diff during **Save & Deploy**. You need to adhere to the leading space indentations.

You can typically enter configuration snippets in DCNM using the following methods:

- User-defined profile and templates
- Switch, interface, overlay, and vPC freeform configurations
- Network and VRF per switch freeform configurations
- Fabric settings for Leaf, Spine, or iBGP configurations

Caution The configuration format should be identical to the **show running configuration** of the corresponding switch. Otherwise, any missing or incorrect leading spaces in the configuration can cause unexpected deployment errors and unpredictable pending configurations. If any unexpected diffs or deployment errors are displayed, check the user-provided or custom configuration snippets for incorrect values.

If DCNM displays the "Out-of-Sync" status due to unexpected pending configurations, and this configuration is either unable to be deployed or stays consistent even after a deployment, perform the following steps to recover:

1. Check the lines of config highlighted under the **Pending Config** tab in the **Config Preview** window.

2. Check the same lines in the corresponding **Side-by-side Comparison** tab. This tab shows whether the diff exists in "intent", or "show run", or in both with different leading spaces. Leading spaces are highlighted in the **Side-by-side Comparison** tab.
3. If the pending configurations or switch with an out-of-sync status is due to any identifiable configuration with mismatched leading spaces in "intent" and "running configuration", this indicates that the intent has incorrect spacing and needs to be edited.
4. To edit incorrect spacing on any custom or user-defined policies, navigate to the switch and edit the corresponding policy:
   a. If the source of the policy is **UNDERLAY**, you will need to edit this from the Fabric settings screen and save the updated configuration.
   b. If the source is blank, it can be edited from the **View/Edit policies** window for that switch.
   c. If the source of the policy is **OVERLAY**, but it is derived from a switch freeform configuration. In this case, navigate to the appropriate **OVERLAY** switch freeform configuration and update it.
   d. If the source of the policy is **OVERLAY** or a custom template, perform the following steps:
      1. Navigate to **Administration > DCNM Server > Server Properties**, set the **template.in_use.check** property to **false**. This allows the profiles or templates to be editable.
      2. Edit the specific profile or template from the **Control > Template Library** edit window, and save the updated profile template with the right spacing.
      3. Click **Save & Deploy** to recompute the diffs for the impacted switches.
      4. After the configurations are updated, set the **template.in_use.check** property to **true**, as it slows down the performance of the DCNM system, specifically for **Save & Deploy** operations.

To confirm that the diffs have been resolved, click **Save & Deploy** after updating the policy to validate the changes.

Note DCNM checks only leading spaces, as it implies hierarchy of the command, especially in case of multi-command sequences. DCNM does not check any trailing spaces in command sequences.

Example 1: Configuration Compliance in Switch Freeform Policy

Let us consider an example with an incorrect spacing in the Switch Freeform Config field.

The switch freeform policy is created as shown:

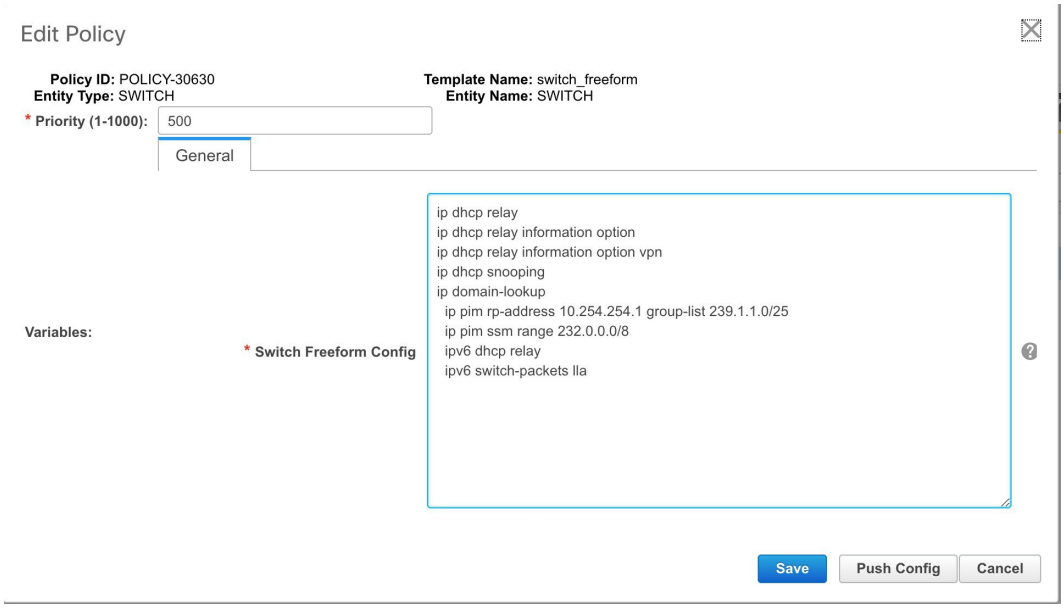

After deploying this policy successfully to the switch, DCNM persistently reports the following diffs:

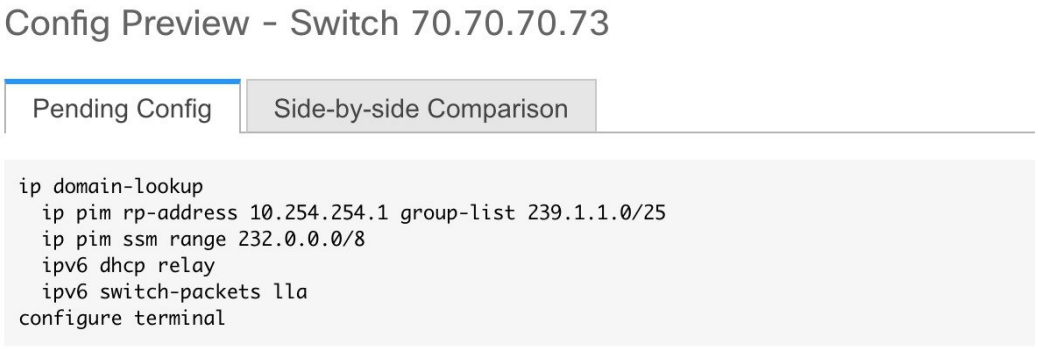

After clicking the **Side-by-side Comparison** tab, you can see the cause of the diff. As seen below, the **ip pim rp-address** line has 2 leading spaces, while the running configuration has 0 leading spaces.

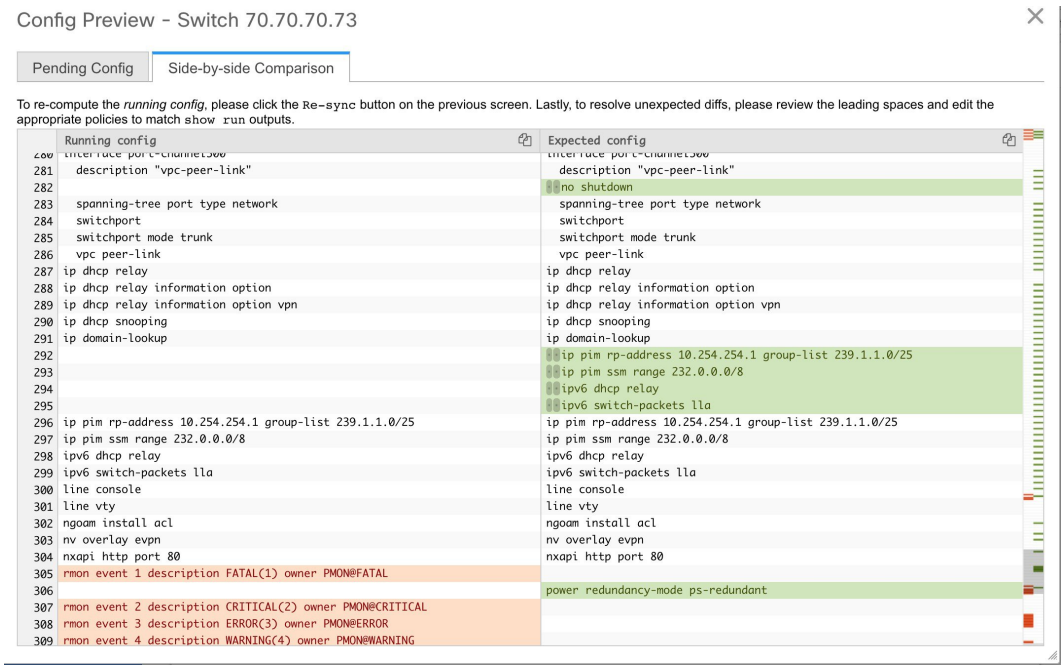

To resolve this diff, edit the corresponding Switch Freeform policy so that the spacing is correct.

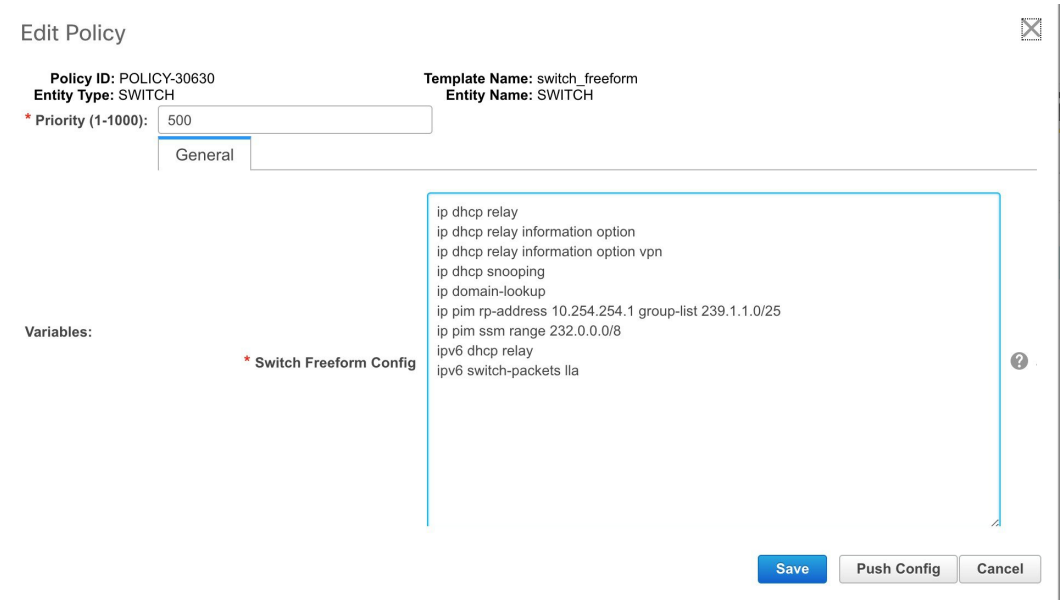

After you save, you can use the **Push Config** or **Save & Deploy** option to re-compute diffs.

As shown below, the diffs are now resolved. The **Side-by-side Comparison** tab confirms that the leading spaces are updated.

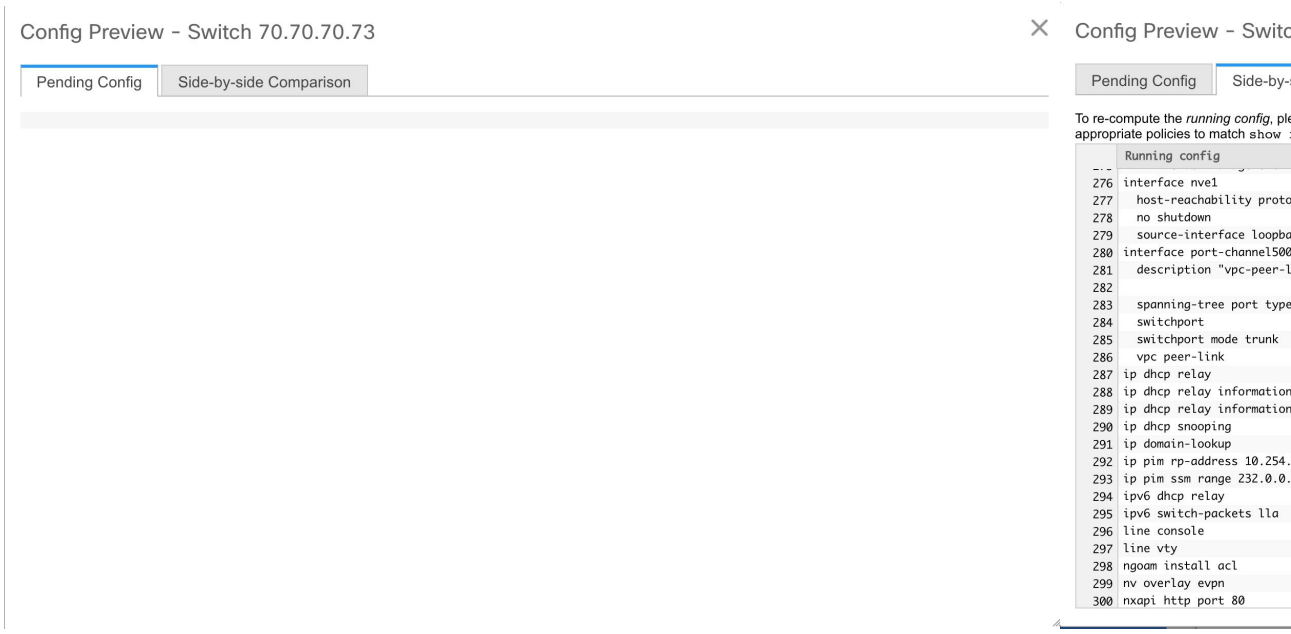

Example 2: Resolving a Leading Space Error in Overlay Configurations

Let us consider an example with a leading space error that is displayed in the **Pending Config** tab.

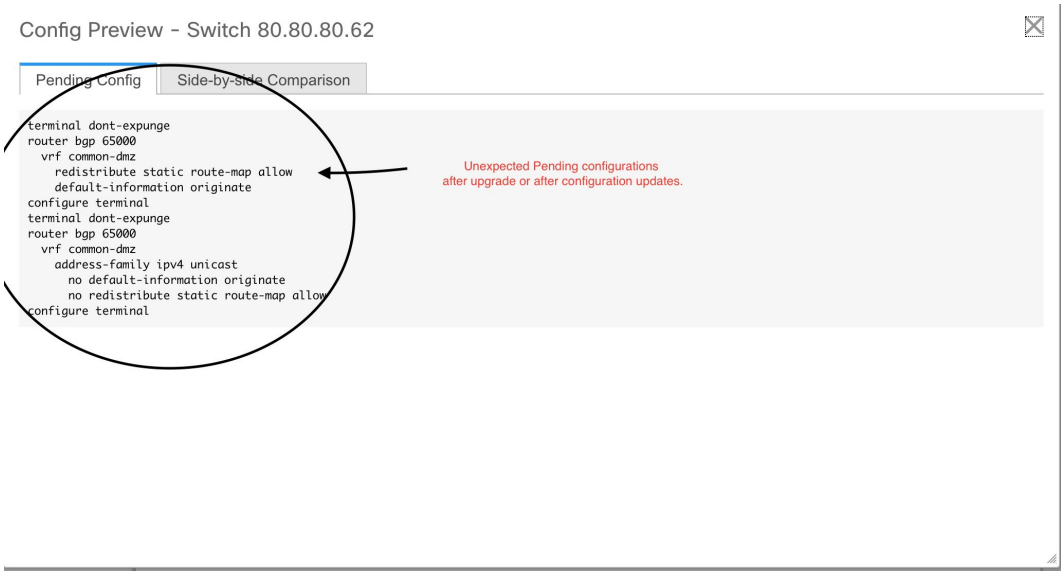

In the **Side-by-side Comparison** tab, search for diffs line by line to understand context of the deployed configuration.

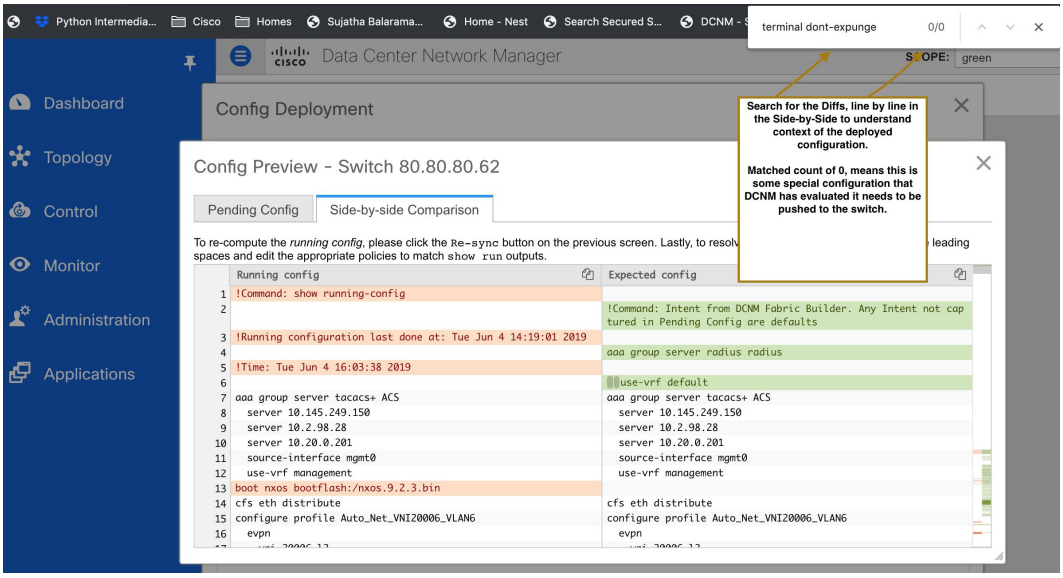

A matched count of 0 means that it is a special configuration that DCNM has evaluated to push it to the switch.

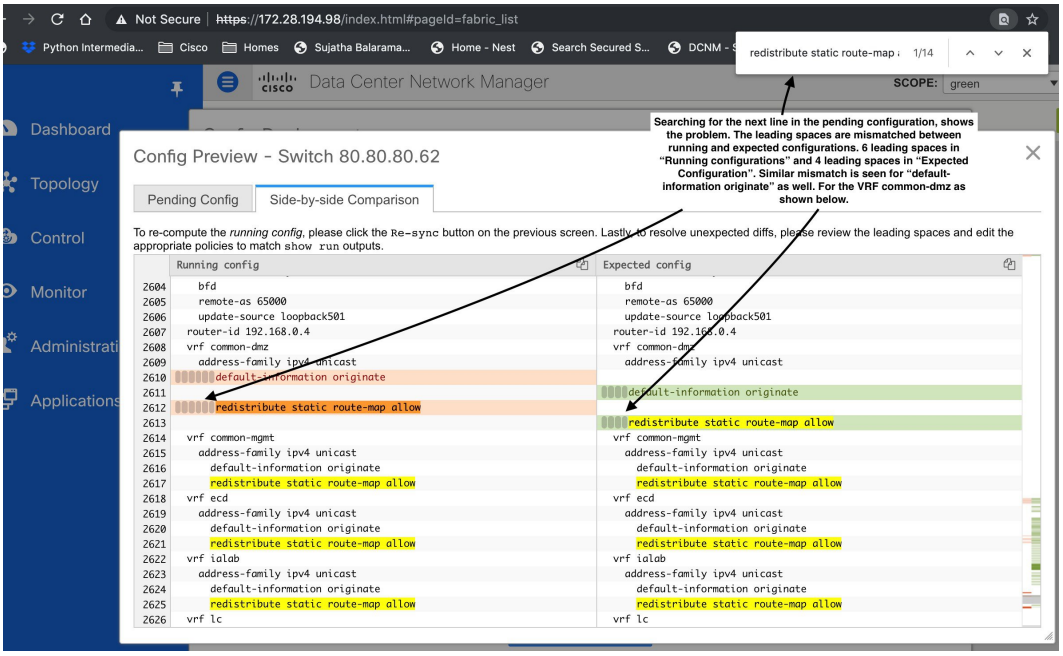

You can see that the leading spaces are mismatched between running and expected configurations.

Navigate to the respective freeform configs and correct the leading spaces, and save the updated configuration.

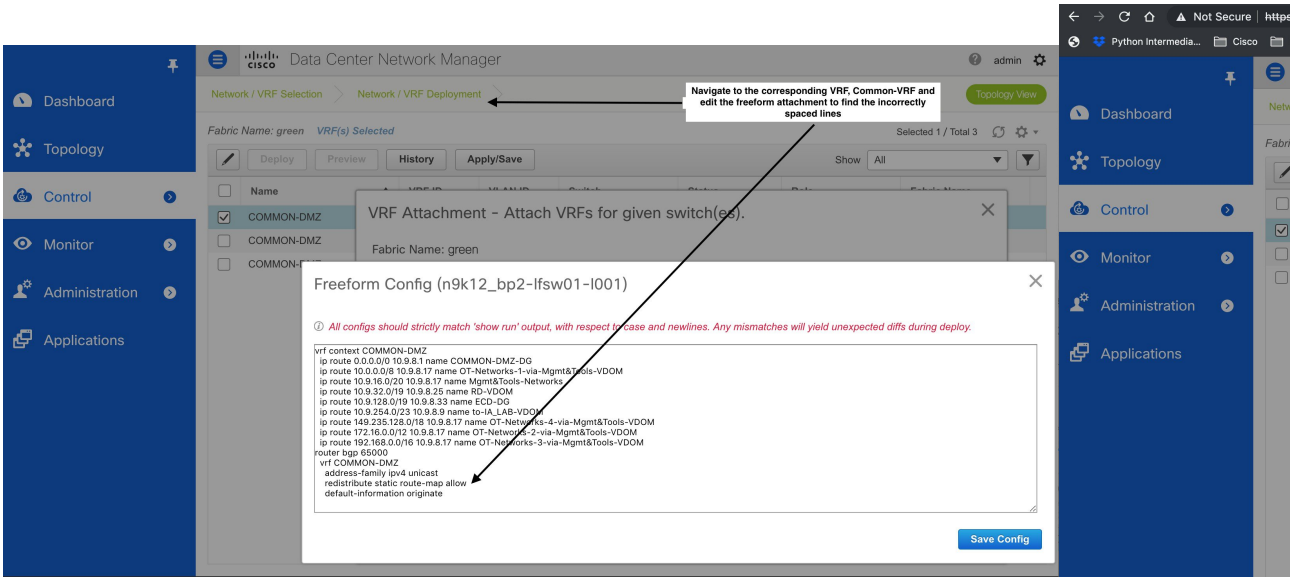

Navigate to the **Fabric Builder** window for the fabric and click **Save & Deploy**.

In the **Config Deployment** window, you can see that all the devices are in-sync.

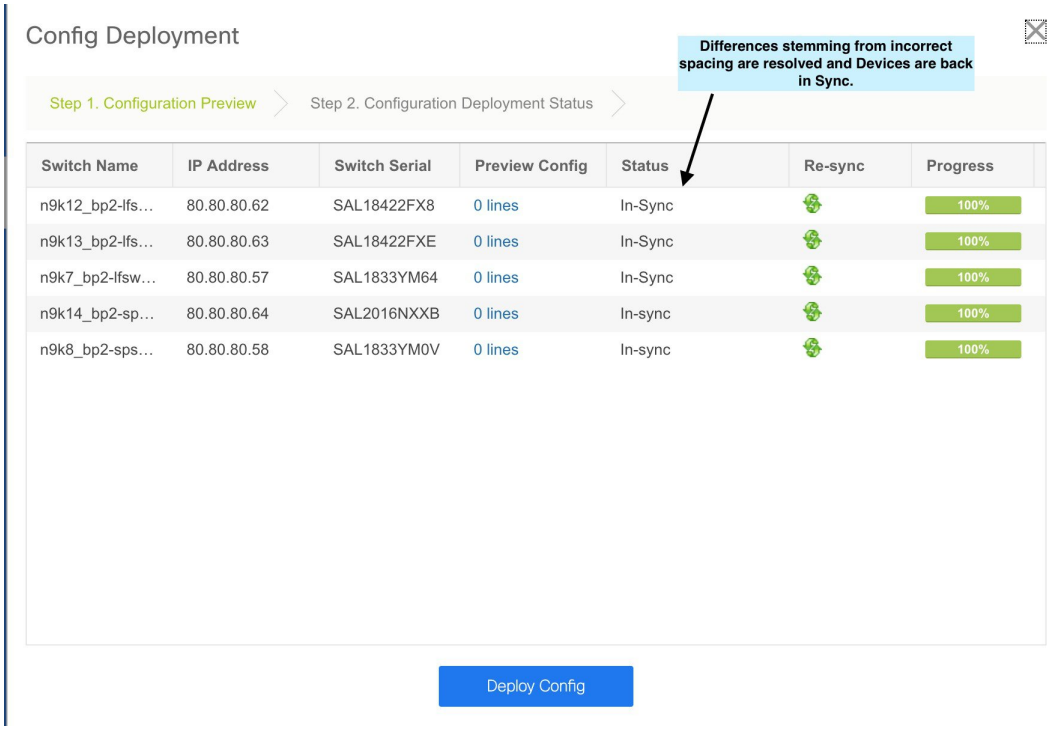

Configuration Compliance in External Fabrics

With external fabrics, any Nexus switch can be imported into the fabric, and there is no restriction on the type of deployment. It can be LAN Classic, VXLAN, FabricPath, vPC, HSRP, etc. When switches are imported into an external fabric, the configuration on the switches is retained so that it is non-disruptive. Only basic policies such as the switch username and mgmt0 interface are created after a switch import.

In the external fabric, for any intent that is defined in the DCNM, configuration compliance (CC) ensures that this intent is present on the corresponding switch. If this intent is not present on the switch, CC reports an Out-of-Sync status. Additionally, there will be a Pending Config generated to push this intent to the switch to change the status to In-Sync. Any additional configuration that is on the switch but not in intent defined in DCNM, will be ignored by CC, as long as there is no conflict with anything in the intent.

When there is user-defined intent added on DCNM and the switch has additional configuration under the same top-level command, as mentioned earlier, CC will only ensure that the intent defined in DCNM is present on the switch. When this user defined intent on DCNM is deleted as a whole with the intention of removing it from the switch and the corresponding configuration exists on the switch, CC will report an Out-of-Sync status for the switch and will generate **Pending Config** to remove the config from the switch. This **Pending Config** includes the removal of the top-level command. This action leads to removal of the other out-of-band configurations made on the switch under this top-level command as well. If you choose to override this behavior, the recommendation is that, you create a freeform policy and add the relevant top-level command to the freeform policy.

Let us see this behavior with an example.

1. A **switch_freeform** policy defined by the user in DCNM and deployed to the switch.

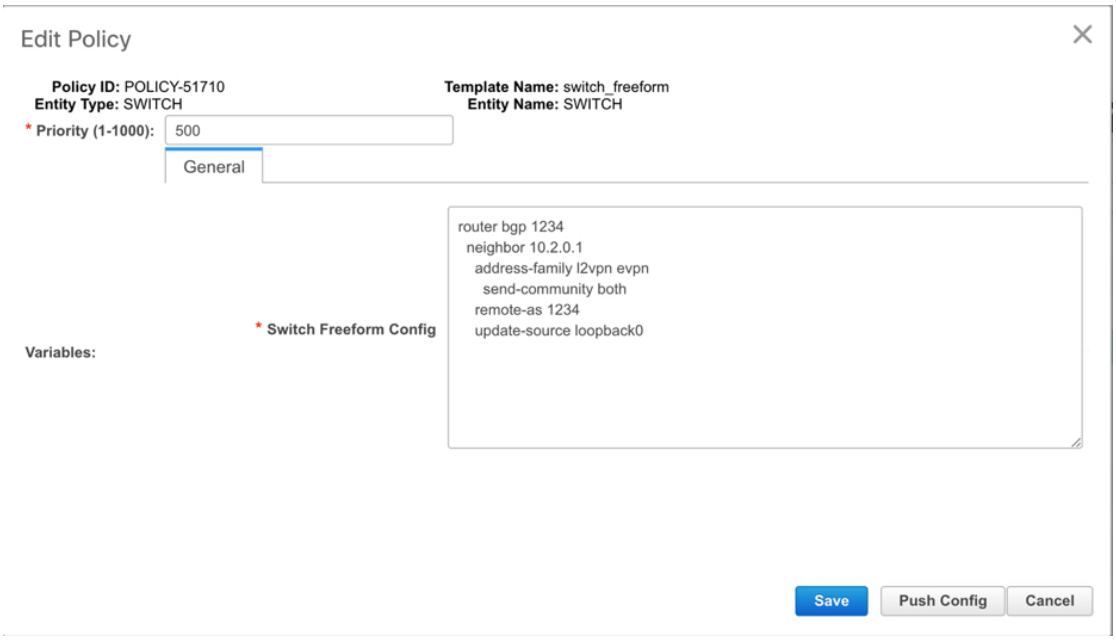

2. Additional configuration exists under **router bgp** in **Running config** that does not exist in user-defined DCNM intent **Expected config**. Note that there is no **Pending Config** to remove the additional config that exists on the switch without a user defined intent on DCNM.

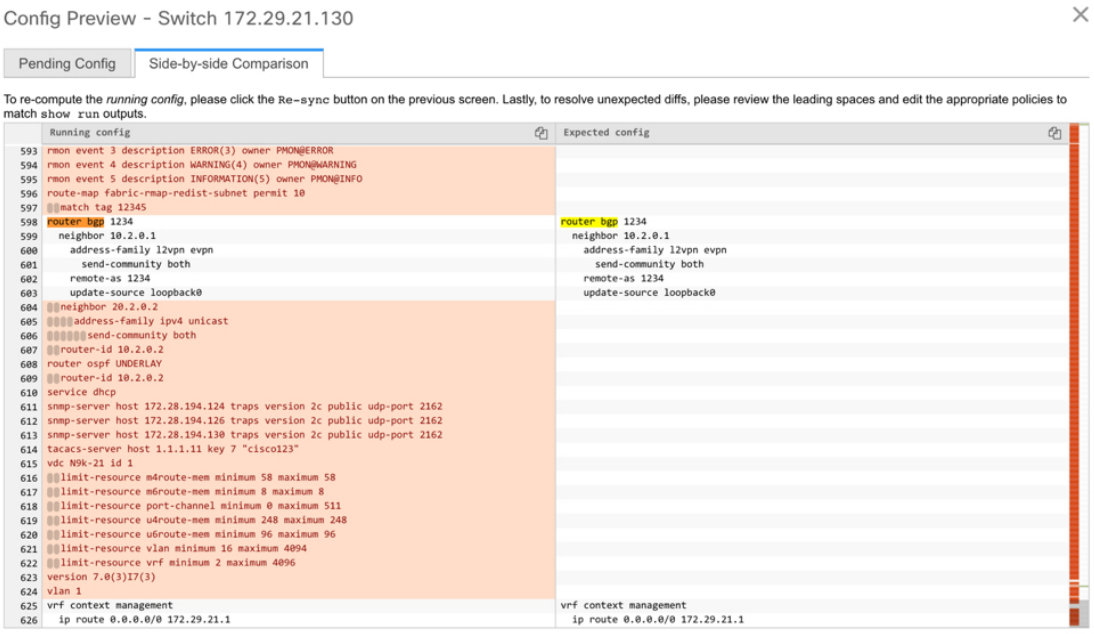

3. The **Pending Config** and the **Side-by-side Comparison** when the intent that was pushed earlier via DCNM is deleted from DCNM by deleting the switch_freeform policy that was created in the Step 1.

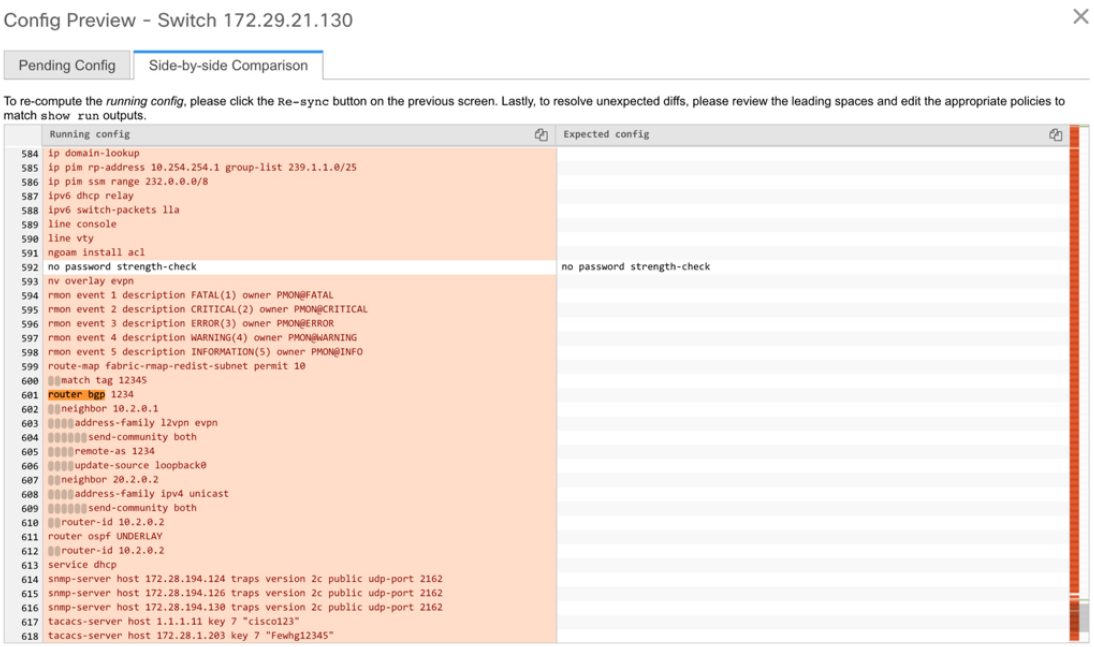

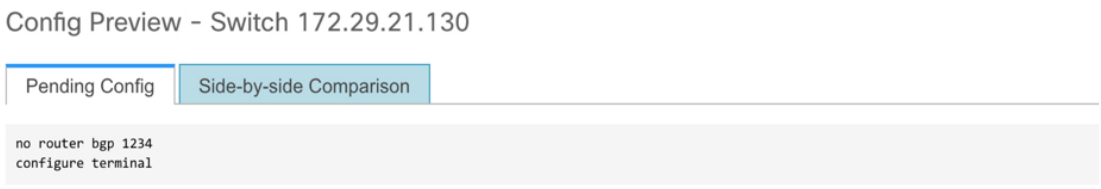

4. A **switch_freeform** policy with the top-level **router bgp** command needs to be created. This enables CC to generate the configuration needed to remove only the desired sub-config which was pushed from DCNM earlier.

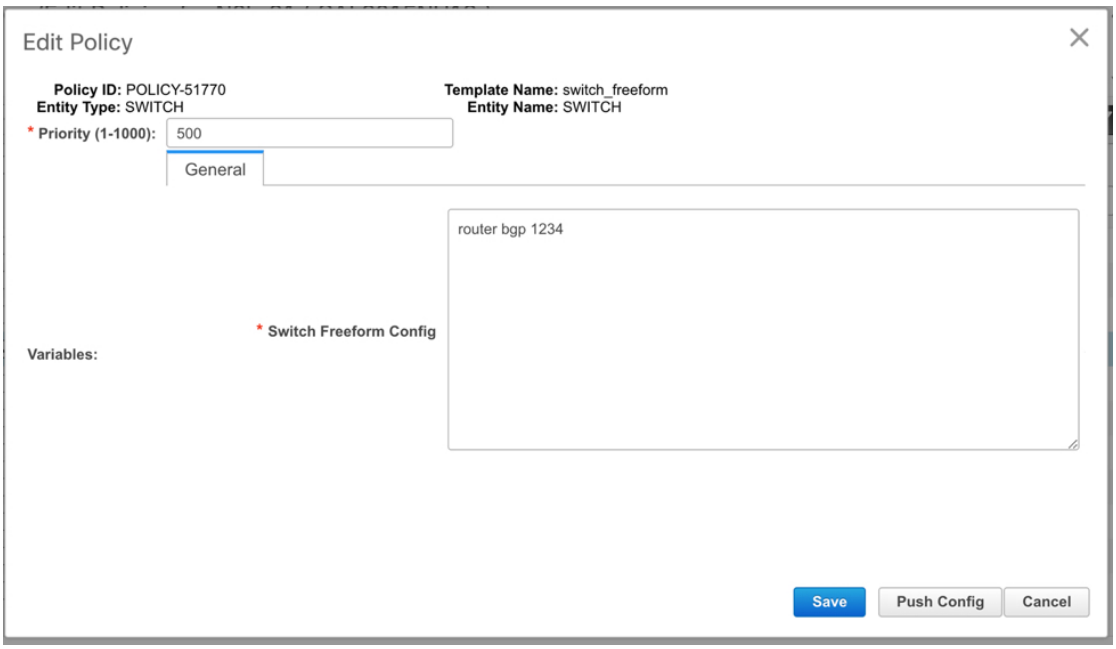

5. The removed configuration is only the subset of the configuration that was pushed earlier from DCNM.

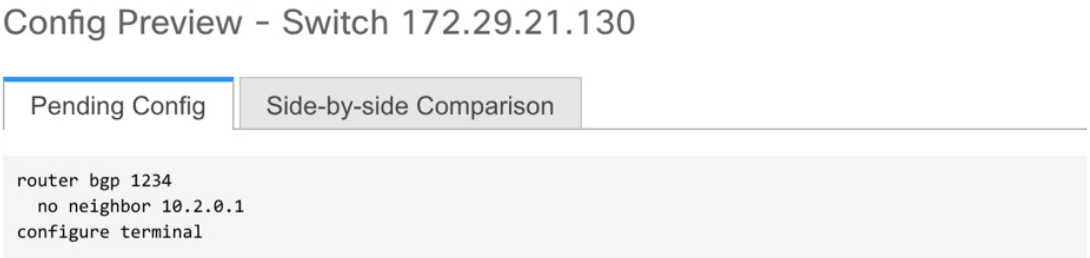

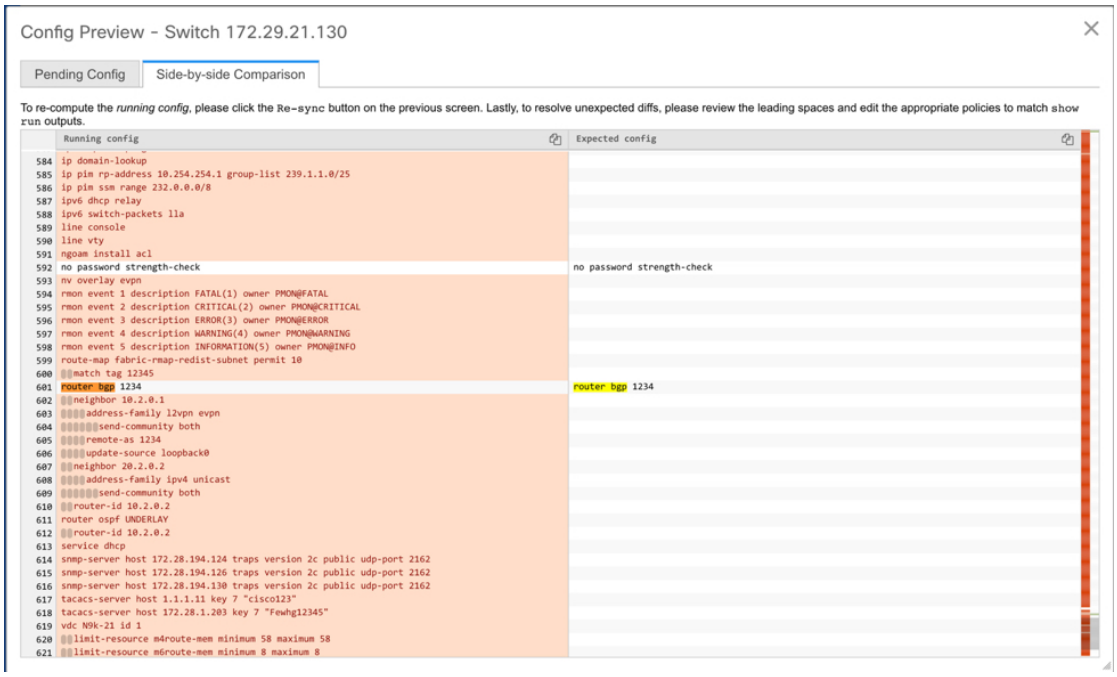

For interfaces on the switch in the external fabric, DCNM either manages the entire interface or does not manage it at all. CC checks interfaces in the following ways:

- For any interface, if there is a policy defined and associated with it, then this interface is considered as managed. All configurations associated with this interface must be defined in the associated interface policy. This is applicable for both logical and physical interfaces. Otherwise, CC removes any out-of-band updates made to the interface to change the status to **In-Sync**.
- Interfaces created out-of-band (applies for logical interfaces such as port-channels, sub interfaces, SVIs, loopbacks, etc.), will be discovered by DCNM as part of the regular discovery process. However, since there is no intent for these interfaces, CC will not report an **Out-of-Sync** status for these interfaces.
- For any interface, there can always be a monitor policy associated with it in DCNM. In this case, CC will ignore the interface's configuration when it reports the **In-Sync** or **Out-of-Sync** config compliance status.

Special Configuration CLIs Ignored for Configuration Compliance

The following configuration CLIs are ignored during configuration compliance checks:

- Any CLI having 'username' along with 'password'
- Any CLI that starts with 'snmp-server user'

Any CLIs that match the above will not show up in pending diffs and clicking **Save & Deploy** in the **Fabric Builder** window will not push such configurations to the switch. These CLIs will not show up in the **Side-by-side Comparison** window also.

To deploy such configuration CLIs, perform the following procedure:

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

Resolving Diffs for Case Insensitive Commands

By default, all diffs generated in DCNM while comparing intent, also known as Expected Configuration, and Running Configuration, are case sensitive. However, the switch has many commands that are case insensitive, and therefore it may not be appropriate to flag these commands as differences. These outlier cases are captured in the **compliance_case_insensitive_clis.txt** text file.

There could be additional commands not included in the existing **compliance_case_insensitive_clis.txt** file that should be treated as case insensitive. If the pending configuration is due to the differences of cases between the Expected Configuration in DCNM and the Running Configuration, you can configure DCNM to ignore these case differences as follows:

1. Modify the following file on the DCNM file system:

   `/usr/local/cisco/dcm/dcnm/model-config/compliance_case_insensitive_clis.txt`

   The sample entries in **compliance_case_insensitive_clis.txt** file are displayed as:

   ```
   [root@dcnm98 model-config]# pwd
   /usr/local/cisco/dcm/dcnm/model-config
   [root@dcnm98 model-config]# cat compliance_case_insensitive_clis.txt
   "^(no |)interface\s+Port(.)"
   "^(no |)interface\s+Loo(.)"
   "^(no |)interface\s+Eth(.)"
   "^update-source\s+Loo(.)"
   "^vrf\s+"
   "^hardware profile portmode\s+"
   "^(.*)route-map\s+(.)"
   "^(.*)neighbor-policy(.)"
   "(no |)encapsulation\s+(.)"
   "(.*)alert-group\s+(.)"
   "^streetaddress\s+(.)"
   "^transport email\s+(.)"
   "(no |)action\s+(.)"
   "(no|)\s+\d*\s+remark.*"
   [root@dcnm98 model-config]#
   ```

   If newer patterns are detected during deployment, and they are triggering pending configurations, you can add these patterns to this file. The patterns need to be valid regex patterns.

   This enables DCNM to treat the documented configuration patterns as case insensitive while performing comparisons.

2. Click **Save & Deploy** for fabrics to see the updated comparison outputs.

Resolving Config Compliance After Importing Switches

After importing switches in Cisco DCNM, configuration compliance for a switch can fail because of an extra space in the management interface (mgmt0) description field.

For example, before importing the switch:

```
interface mgmt0
  description SRC=SDS-LB-LF111-mgmt0, DST=SDS-LB-SW001-Fa0/5
```

After importing the switch and creating a config profile:

```
interface mgmt0
  description SRC=SDS-LB-LF111-mgmt0,DST=SDS-LB-SW001-Fa0/5
```

In this example, the space after the comma (,) is removed.

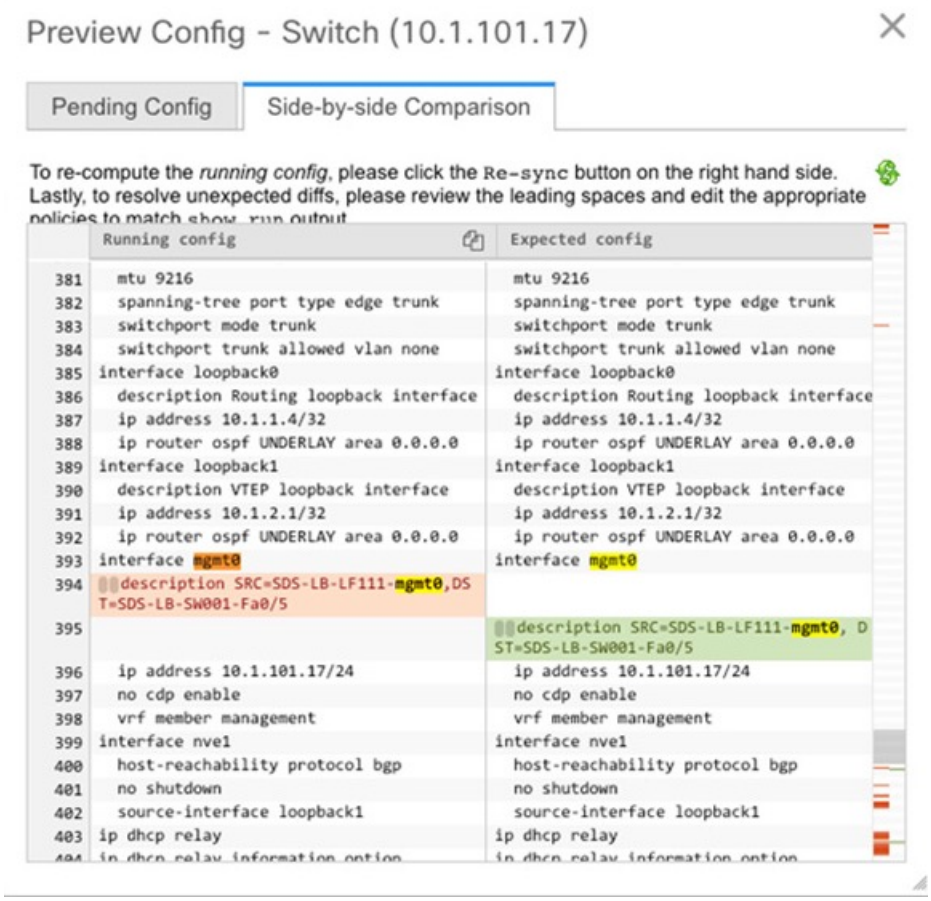

Navigate to Interface Manager and click the **Edit** icon after selecting the mgmt0 interface. Remove the extra space in the description.

Strict Configuration Compliance

From Cisco DCNM Release 11.3(1), strict configuration compliance checks for diff between the switch configuration and the associated intent and generates **no** commands for the configurations that are present on the switch but are not present in the associated intent. When you click **Save and Deploy**, switch configurations that are not present on the associated intent are removed. You can enable this feature by selecting the **Enable Strict Config Compliance** checkbox under the **Advanced** tab in the **Add Fabric** or **Edit Fabric** window. By default, this feature is disabled.

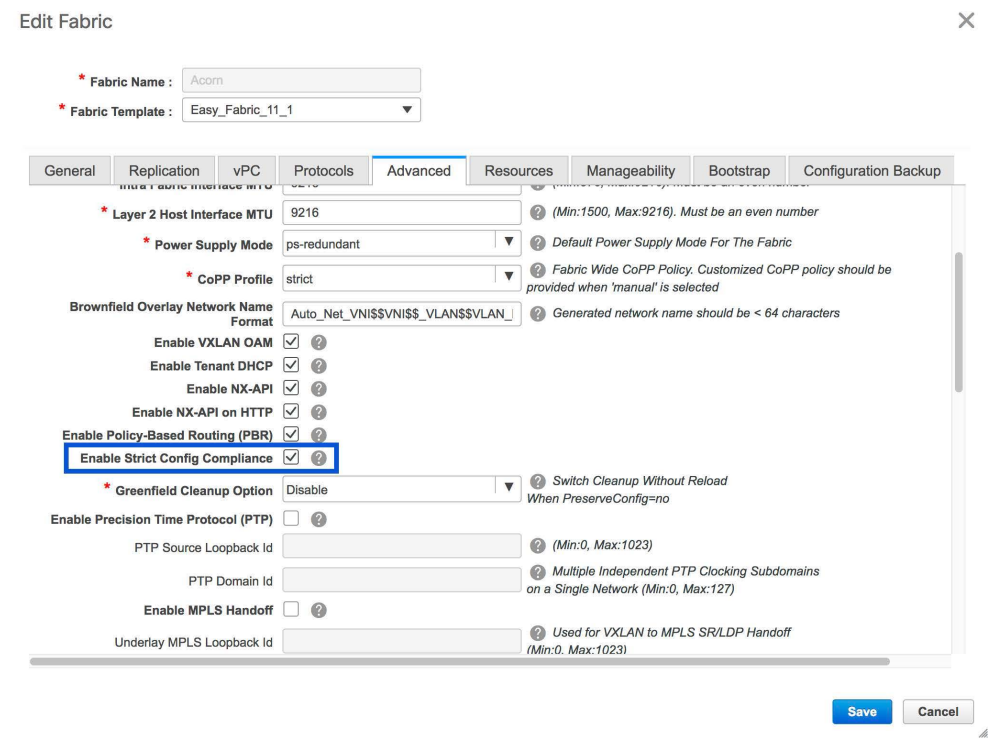

The strict config compliance feature is supported on the Easy Fabric templates - Easy_Fabric_11_1 and Easy_Fabric_eBGP. To avoid generating diff for commands that are auto-generated by the switch, such as vdc, rmon, and so on, a file that has a list of default commands is used by CC to ensure that diffs are not generated for these commands. This file is located at **/usr/local/cisco/dcm/dcnm/model-config/strict_cc_exclude_clis.txt**.

Note

- In case any diffs are generated after strict configuration compliance is enabled, the switch icon turns blue in color in the **Fabric Builder** window.

Example: Strict Configuration Compliance

Let us consider an example in which the **feature telnet** command is configured on a switch but is not present in the intent. In such a scenario, the status of the switch is displayed as **Out-of-sync** after a CC check is done.

Now, click **Preview Config** of the out-of-sync switch. As the strict config compliance feature is enabled, the **no** form of the **feature telnet** command appears under **Pending Config** in the **Preview Config** window.

Click the **Side-by-side Comparison** tab to display the differences between the running configuration and the expected configuration. Starting from Cisco DCNM Release 11.3(1), the **Re-sync** button is also displayed at the top right corner under the Side-by-side Comparison tab in the Preview Config window. 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.

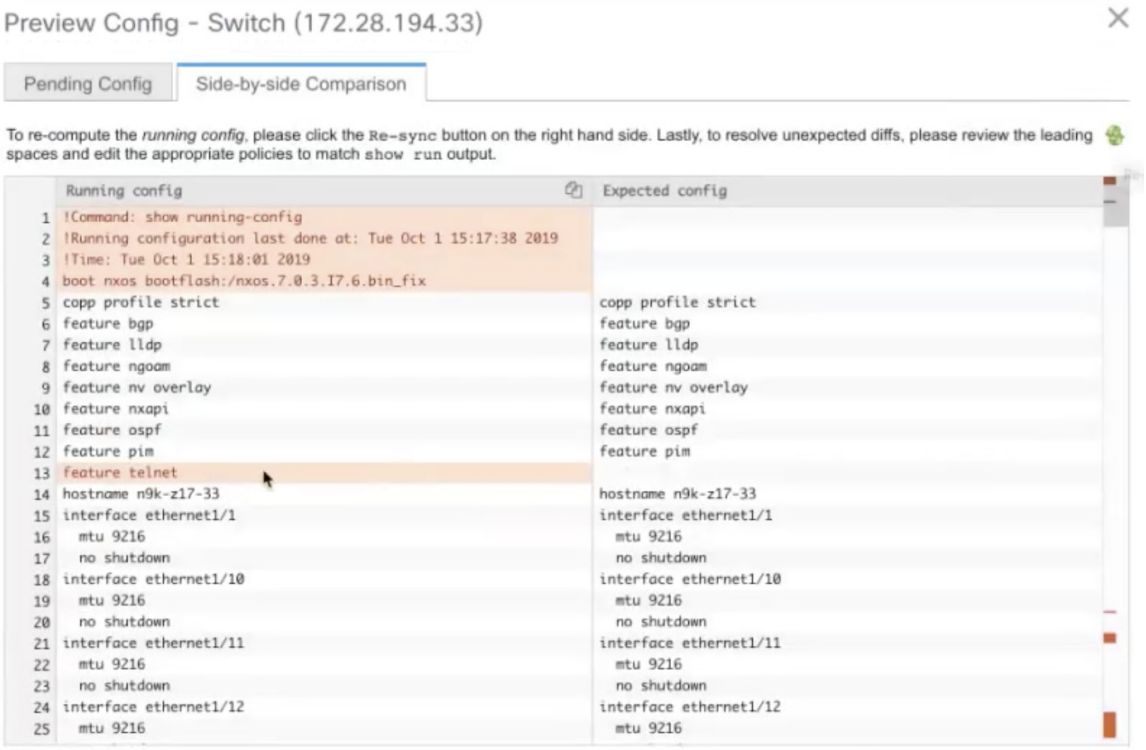

The re-sync operation does a full CC run for the switch and recollects “show run” and “show run all” commands from the switch. When you initiate the re-sync process, a progress message is displayed on the screen. During the re-sync, the running configuration is taken from the switch. The Out-of- Sync/In-Sync status for the switch is recalculated based on the intent defined in DCNM.

Now, close the **Preview Config** window and click **Save and Deploy**. The Strict configuration compliance feature then ensures that the running config on the switch does not deviate from the intent by pushing the **no** form of the **feature telnet** command to the switch. The diff between the configurations is highlighted. The diff other than the **feature telnet** command are default switch and boot configurations and are ignored by the strict CC check.

In Cisco DCNM Release 11.2(1) and earlier releases, you had to right-click on a switch in the Fabric builder window and select **Deploy Config** to display the **Config Deployment** window. You then had to click **Preview Config** for a specific switch to bring up the **Preview Config** window that displays the pending configuration for that switch. This leads to a scenario in which the user may think that the preview config is inadvertently being deployed on the switch. Starting from Cisco DCNM Release 11.3(1), you can right-click on a switch in the **Fabric Builder** window and select **Preview Config** to display the **Preview Config** window. This window displays the pending configuration that has to be pushed to the switch to achieve configuration compliance with the intent.

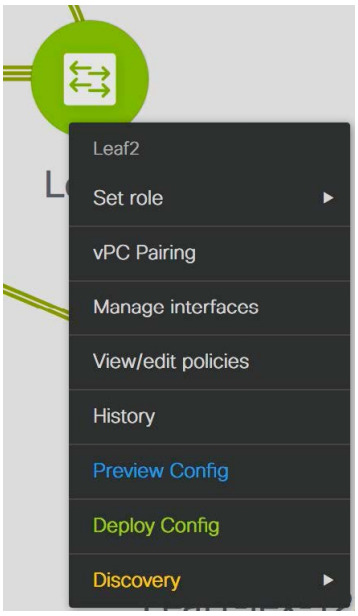

Custom freeform configurations can be added in DCNM to make the intended configuration on DCNM and the switch configurations identical. The switches are then in In-Sync status. For more information on how to add custom freeform configurations on DCNM, refer Enabling Freeform Configurations on Fabric Switches .

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 th Nexus switches. For more information, see Resolving Freeform Config Errors in Switches, on page 347.
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 347.

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
```

IP Prefix List/Route-map configuration

IP prefix list and route-map configurations are typically configured on border devices. These configurations are global because they can be defined once on a switch and then applied to multiple VRFs as needed. The intent for this configuration can be captured and saved in a switch_freeform policy. As mentioned earlier, note that the config saved in the policy should match the **show run** output. This is especially relevant for prefix lists where the NX-OS switch may generate sequence numbers automatically when configured on the CLI. An example snippet is shown below:

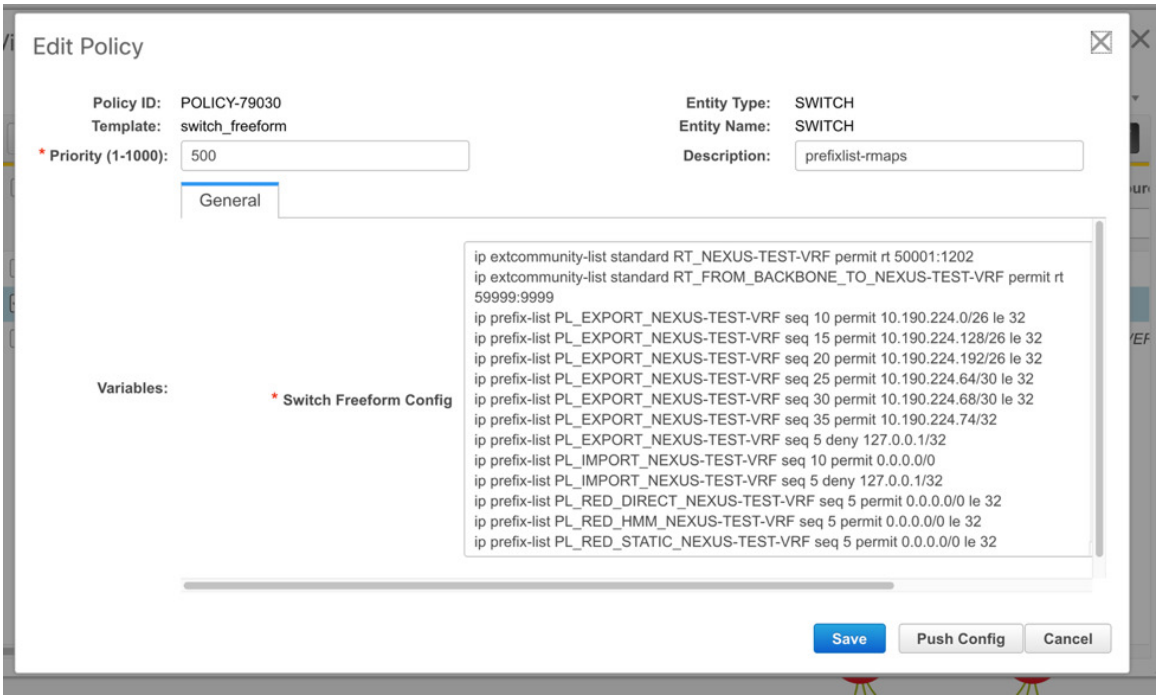

ACL configuration

ACL configurations are typically configured on specific switches and not fabric-wide (leaf/spine switches). When you configure ACLs as freeform CLIs on a switch, you should include sequence numbers. Else, there will be a mismatch between the intended and running configuration. A configuration sample with sequence numbers:

```
ip access-list ACL_VTY
  10 deny tcp 172.29.171.67/32 172.29.171.36/32
  20 permit ip any any
ip access-list vlan65-acl
  10 permit ip 69.1.1.201/32 65.1.1.11/32
  20 deny ip any any

interface Vlan65
  ip access-group vlan65-acl in
line vty
  access-class ACL_VTY in
```

If you have configured ACLs without sequence numbers in a **switch_freeform** policy, update the policy with sequence numbers *as shown in the running configuration of the switch*.

After the policy is updated and saved, right click the device and select the per switch **Deploy Config** option to deploy the configuration. Alternatively, use the **Save and Deploy** option in the fabric topology screen

(within Fabric Builder) so that the fabric triggers Configuration Compliance and resolves the configuration mismatch.

Resolving Freeform Config Errors in Switches

Copy-paste the running-config to the freeform config with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. Otherwise, configuration compliance in DCNM marks switches as out-of-sync.

Let us see an example of the freeform config of a switch.

```
feature bash-shell
feature telemetry

clock timezone CET 1 0
# Daylight saving time is observed in Metropolitan France from the last Sunday in March
(02:00 CET) to the last Sunday in October (03:00 CEST)
clock summer-time CEST 5 Sunday March 02:00 5 Sunday October 03:00 60
clock protocol ntp

telemetry
  destination-profile
    use-vrf management
```

The highlighted line about the daylight saving time is a comment that is not displayed in the **show running config** command output. Therefore, configuration compliance marks the switch as out-of-sync because the intent does not match the running configuration.

Let us check the running config in the switch for the clock protocol.

```
spine1# show run all | grep "clock protocol"
clock protocol ntp vdc 1
```

You can see that **vdc 1** is missing from the freeform config.

In this example, let us copy-paste the running config to the freeform config.

Here is the updated freeform config:

```
feature bash-shell
feature telemetry

clock timezone CET 1 0
clock summer-time CEST 5 Sunday March 02:00 5 Sunday October 03:00 60
clock protocol ntp vdc 1

telemetry
  destination-profile
    use-vrf management
```

After you copy-paste the running config and deploy, the switch will be in-sync. When you click **Save & Deploy**, the **Side-by-side Comparison** tab in the **Config Preview** window provides you information about the difference between the defined intent and the running config.

Deploying Freeform CLIs on a Specific Switch on a Per VRF/Network basis

1. Click **Control > VRFs**. After choosing the appropriate fabric scope, the listing of the currently defined VRFs for the fabric shows up.
2. Create a new VRF by clicking the + button or select an existing VRF and click the **Continue** button on the top right.
3. The topology view for the fabric shows up. Switches to which the VRF is already deployed are highlighted in green. Other switches will be in gray color.
4. Select an individual switch. The VRF attachment form shows up listing the switch that is selected. In case of a vPC pair, both switches belonging to the pair will show up.
5. Under the CLI Freeform column, select the button labelled **Freeform config**. This option allows a user to specify additional configuration that should be deployed to the switch along with the VRF profile configuration.
6. Add or update the CLIs in the **Free Form 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.

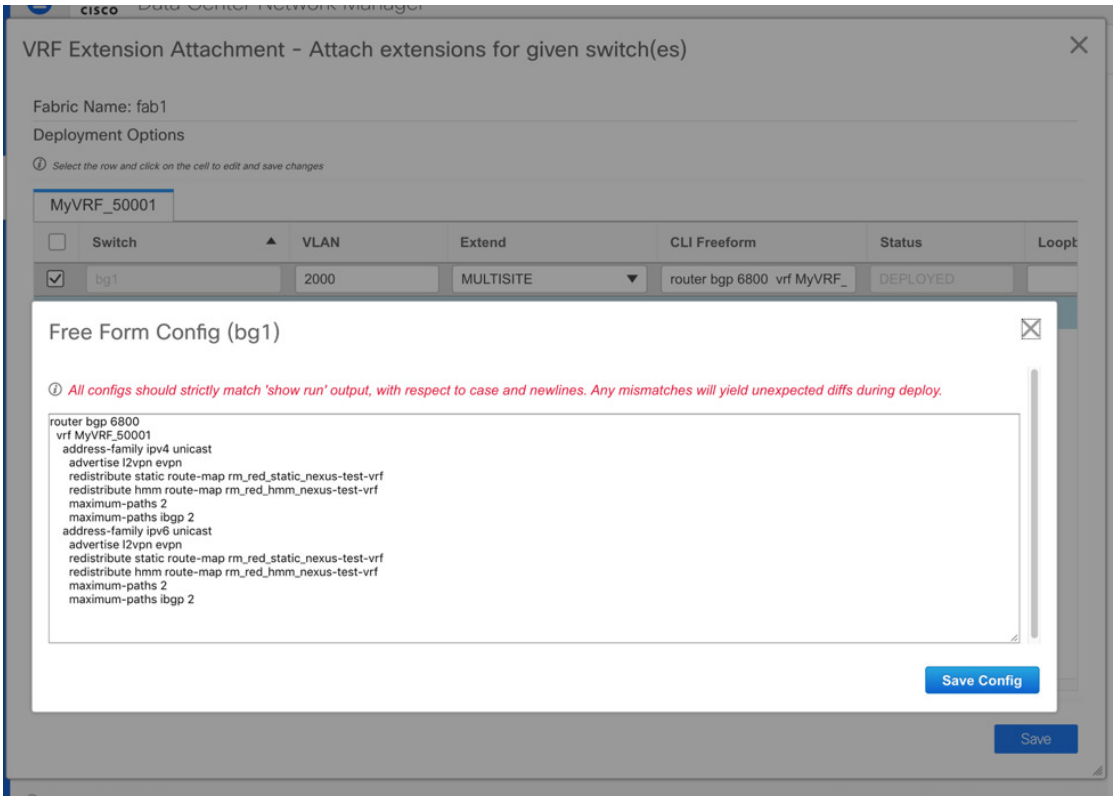

7. Click **Save Config**.

Note The **Freeform config** button will be gray when there is no per VRF per switch config specified. The butto will be blue when some config has been saved by the user.

After the policy is saved, Click **Save** on the VRF Attachment pop-up to save the intent to deploy the VRF to that switch. Ensure that the checkbox on the left next to the switch is checked.

8. Now, optionally, click **Preview** to look at the configuration that will be pushed to the switch.

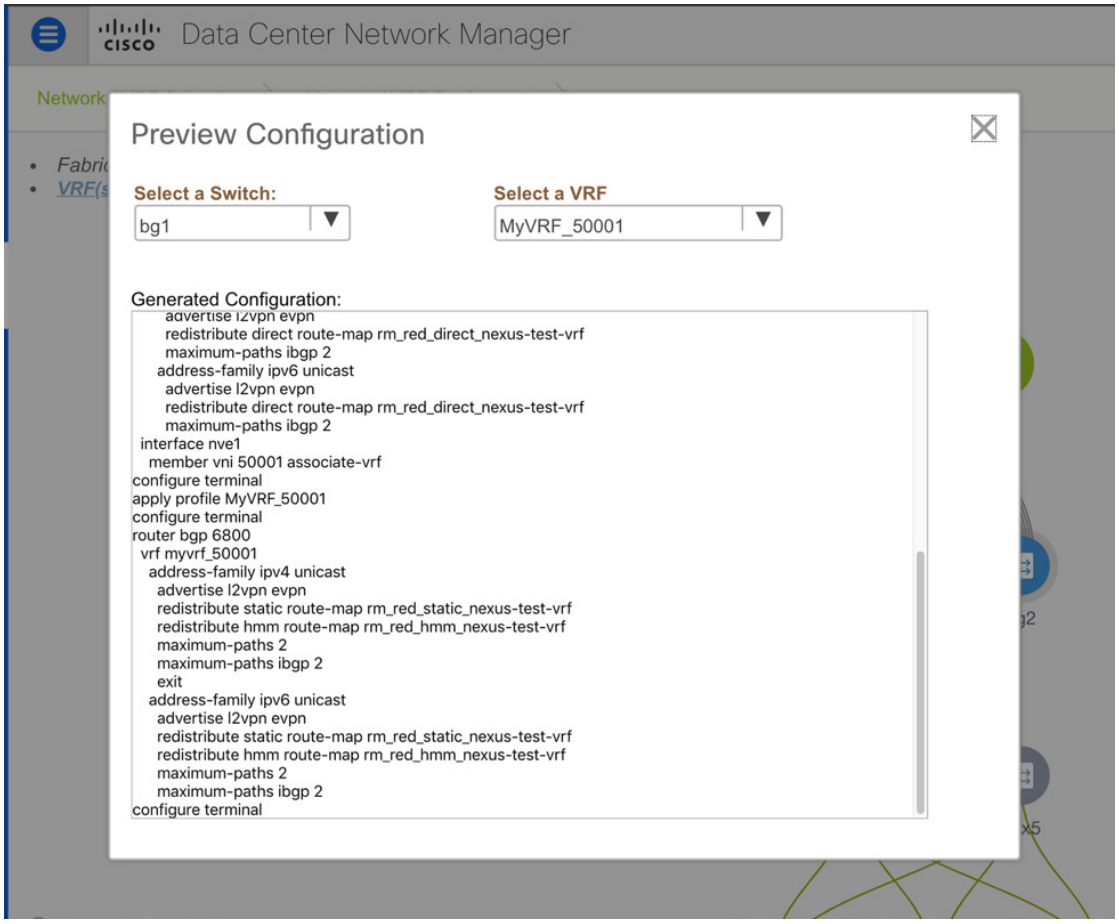

9. Click **Deploy** to push the configuration to the switch.

The same procedure can be used to define a per Network per Switch configuration.

VMM Workload Automation

VMM workload automation is about the automation of network configuration in Cisco's Nexus switches for workloads spawned in a VMware environment. Note that this is a preview feature in the Cisco DCNM Release 11.4(1).

You can also watch the video that demonstrates this automation. See Video: VMM Workload Automation in Cisco DCNM.

Overview of Network Objects in vCenter

VMM workload automation involves the mapping of network objects in vCenter to the network objects in DCNM. The following network objects in vCenter are considered:

- Virtual Switch (VS): A regular VS runs in an ESXi host that performs software-based switching. A VS can have multiple port-groups (PG), where each PG has the network port configuration properties

connecting to the network, like a VLAN. Each VS can have multiple uplink ports connecting to the leaf switches. Workloads spawned in the ESXi host can attach to the PG created in this VS.

- Distributed Virtual Switch (DVS): A DVS is a virtual switch that spans across multiple ESXi hosts. Similar to a regular VS, a DVS has multiple port-groups referred to as Distributed Port Groups (DPG). A DPG has the network port configuration properties connecting to the network, like a VLAN. Each DVS can have multiple uplink ports, which can be connected to the leaf switches. Workloads spawned in any of the hosts that are members of the DVS, can attach to the DPG. Through this document, and in the config file, a DPG is also referred as Distributed Virtual Port Group (DV-PG).

Let us consider the following topology in DCNM:

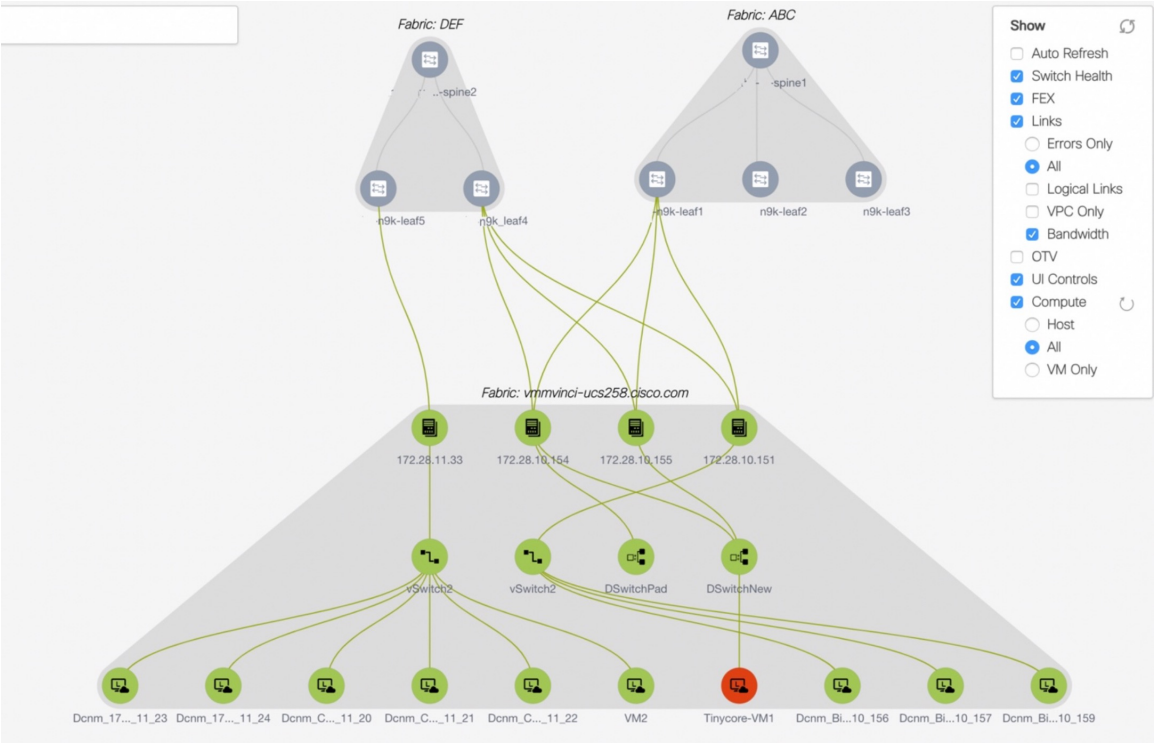

- There are four hosts in the setup with IP addresses:
  - 172.28.11.33
  - 172.28.10.154
  - 172.28.10.151
  - 172.28.10.155

- A DVS named **DSwitchNew** spawns across hosts 172.28.10.154 and 172.289.10.155. This DVS has a DPG named **DPGNew** that isn't shown in the figure. This DVS connects to switches n9k-leaf1 and n9k-leaf4 through uplink ports <vmnic3, vmnic1> and switch interfaces <e1/25, e1/7> respectively (not shown in the figure). A VLAN value of 571 is associated with the DPG **DPGNew**.

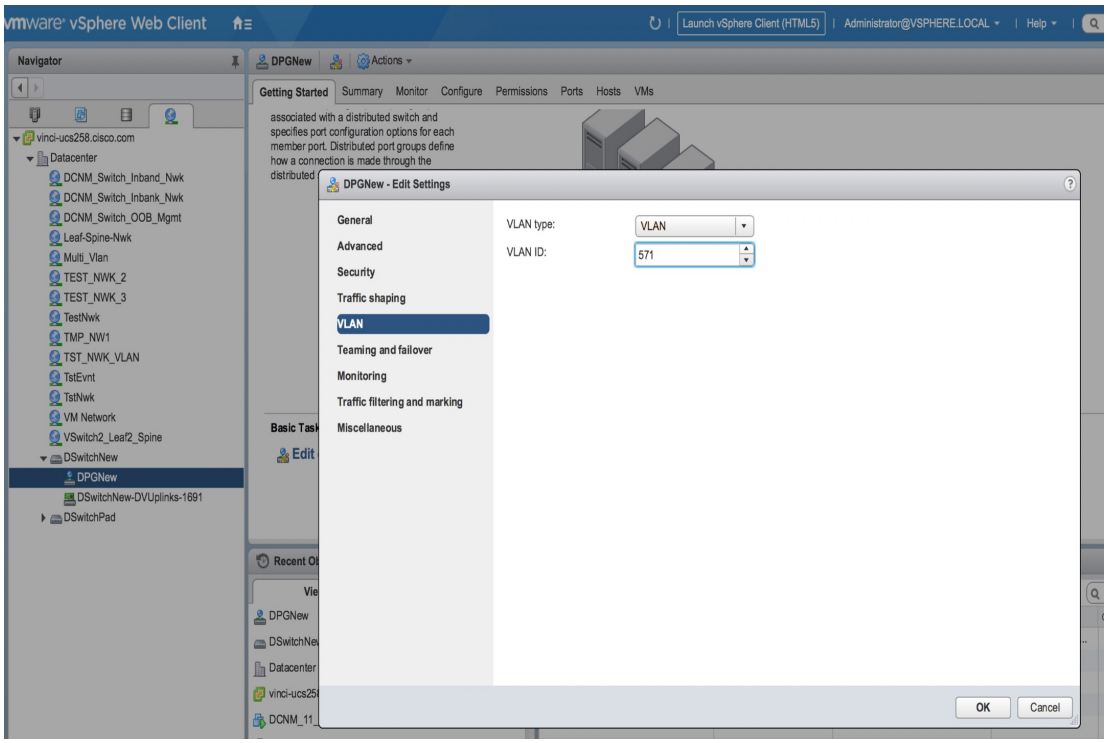

- There's also a regular vSwitch named vSwitch2 in host 172.28.11.33. This VS has a PG named DCNM_Inband. This VS connects to the leaf switch n9k-leaf5 through uplink port vmnic4 and switch interface e1/23. A VLAN value of 540 is associated with the PG **DCNM_Inband**.

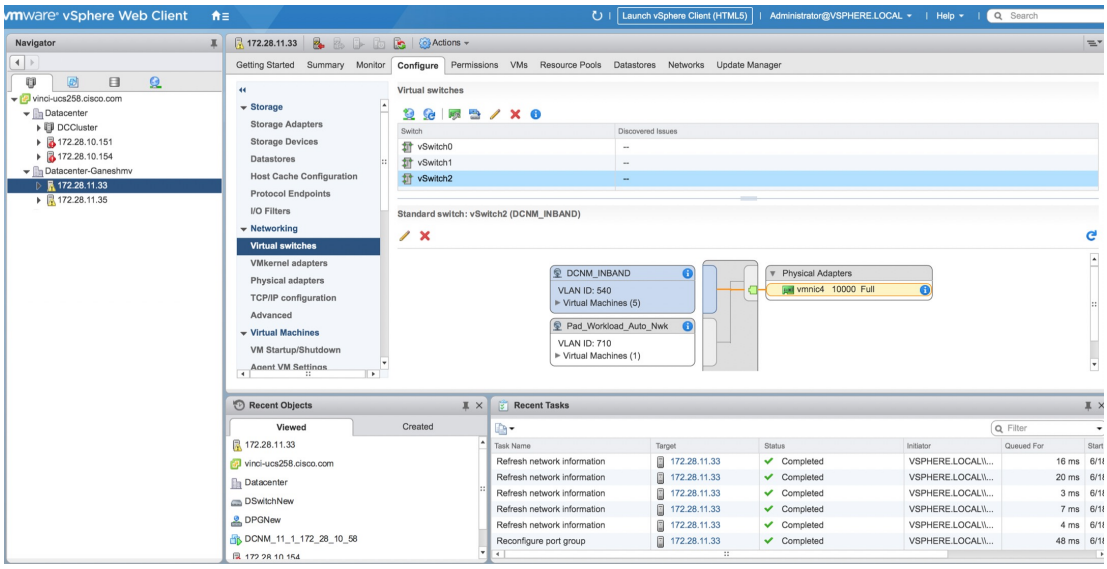

How VMM Workload Automation Works

When workloads are spawned, they require provisioning in the network or fabric. The workloads spawned in the vCenter are either associated with a DPG or a PG. This DPG or PG in the VMWare should be mapped to a corresponding DCNM network. As an example in the earlier topology, when workloads are spawned in the

host 172.28.11.33 with PG **Inband**, then network provisioning that includes overlay and underlay configuration needs to happen in the leaf switch n9k-leaf5 with the relevant interface e1/23.

For the network provisioning to happen, each network object in a vCenter (a DPG or a PG) should be mapped to a network object in DCNM. A network object in DCNM has the following characteristics:

- VRF Name
- VLAN ID
- IPv4/IPv6 subnet and gateway information
- Secondary IPv4/v6 and gateway information
- BGP-EVPN configuration

A static mapping should be defined in a config file that maps the network object in a vCenter to a network object in DCNM. For information, see Configuration Files for VMM Workload Automation, on page 352.

After the config file is populated, the workload automation module can be started. The module scans all the vCenters specified in the config file (conf.yml) and collects the following information for each vCenter:

- List of DVS and the DPG configured in all data centers.
- List of PGs configured in every host in all data centers.
- For every DPG or PG specified in the conf file, it finds the configured VLANs and the directly connected neighbor switches along with its interface information.
- For every <DVS, DPG> or <Host, PG> specified in the config file, it gets the associated network mapping in DCNM.

The module merges all the information and calls the DCNM APIs to provision or amend the networks in all the switches that are discovered as neighbors in one of the earlier steps.

Provisioning of a network or fabric uses DCNM top-down provisioning and it consists of the following steps:

1. Attaching the network configuration to the relevant interfaces of one or more switches that are discovered as neighbors. This attachment is done by the workload automation module.
2. After the configuration is attached, you can review the exact CLIs that are pushed to the switches.
3. After review, you can deploy the configuration to the switches. This deployment can be either done by the script based on the configuration file setting (default is **False**) or you can do it through DCNM. After this step, the configuration appears in the switches.

For more information, see https://pypi.org/project/vmm-workload-auto/.

Configuration Files for VMM Workload Automation

The following configuration files are used for VMM workload automation:

- Global YML File (conf.yml): This file has global configuration and access or credential information of the DCNM and the vCenters. Also, the location of the CSV file for each DCNM is specified in this file. For more information, see Configuration Files for VMM Workload Automation, on page 352.
- CSV file (sample.csv): This file has the mapping of <DVS, DVS-PG> or <Host, PG> in vCenter to the Network name in DCNM. There's a separate CSV file for each DCNM. For more information, see CSV File for Mapping Networks in vCenter and DCNM, on page 354.

Configuration File for Mapping DCNM and vCenter

The configuration file (conf.yml) specifies the DCNM IP address, username, and password. For each DCNM, the list of vCenter information like the IP address, username, and password is also specified. Multiple DCNMs can be specified in this conf.yml file. For every DCNM instance, there's an associated CSV file. The multi-DCNM case is only applicable when the script isn't run in a DCNM, but run in a server that has connectivity to all the DCNMs and vCenters specified in the config file.

In the config file, the hierarchy of information that should be specified is as follows:

```
Global config parameters
DCNM1
      DCNM1 config parameters including location of the CSV file
 vCenter1
  vCenter1 config parameters
  …
 vCenter2
              vCenter2 config parameters
  ….
DCNM2
       ….
  …
```

The location of this configuration file depends on the installation method of the VMM workload automation script. For more information, see Installing VMM Workload Automation Script. This file contains example entries. Modify it based on your environment.

The config file has the following entries:

LogFile: Specifies the name of the log file including the absolute path that will be used by the workload automation module for logging the errors and debug information. Make sure that the directory has write permission for creating the log file. For example, /tmp/workloadauto.log.

ListenPort: Specifies the port that the workload automation module uses to listen for the REST APIs, for example, 9590. Make sure that this port isn't used by any other application. You can check the same by running the **sudo netstat -tulpn** command.

AutoDeploy: Specifies whether the script should automatically deploy the configuration in the switches after attaching the networks. By default, it's set to **False** so that you can review the config and deploy it in the DCNM.

NwkMgr: Specifies the top-level section that contains the DCNM information. For multiple DCNM instances, repeat the fields with the appropriate values. For an example, see conf_multiple_dcnm.yml file that handles multiple DCNMs.

Ip: Specifies the IP address of DCNM, for example, 172.28.10.156.

User: Specifies the username used to log in to DCNM, for example, admin.

Password: Specifies the password of DCNM.

CsvFile: Specifies the absolute path of the location of the CSV file for this DCNM, for example, /etc/vmm_workload_auto/sample.csv.

ServerCntrlr: Specifies the information for the server controller, that is, vCenter/vSphere. For multiple vCenters that fall under this DCNM, this section repeats. For an example, refer the conf_multiple_vcenter.yml file, which contains multiple vCenters under a DCNM.

Ip: Specifies the IP address of the vCenter.

Type: Specifies the type of the server controller. The default is vCenter.

User: Refers to the username used to log in to the vCenter. For example, administrator@vsphere.local.

Password: Refers to the password for the vCenter.

The following example shows the contents of a conf.yml file:

```
LogFile: /tmp/workloadauto.log
ListenPort: 9590
AutoDeploy: false
NwkMgr:
- Ip: 172.28.10.151
  User: admin
  Password: C1sco_123
  CsvFile: /etc/sample.csv
  ServerCntrlr:
    - Ip: 172.28.10.194
      Type: vCenter
      User: administrator@vsphere.local
      Password: Cisc0!23
```

CSV File for Mapping Networks in vCenter and DCNM

The CSV file contains the mapping of the network object in vCenter to the network created in DCNM. This file has the following entries in CSV format, that is, comma-separated entries. The reason for having a CSV file is to specify the mapping between a PG (or DPG) of vSphere to the network name of DCNM. It's a 1-1 mapping. However, since a PG or a DPG can't be identified on its own (not unique), you need an extra DVS name or Hostname to map it.

The CSV file contains the following fields:

vCenter - Specifies the IP address of vCenter

Dvs - Specifies the name of the DVS.

Dvs_pg - Specifies the DVS PG (DPG) in the DVS

Host - Specifies the Host/Server (IP address)

Host_pg - Specifies the port-group in the host.

Fabric - Specifies the fabric in DCNM.

Network - Specifies the name of the network already created in DCNM.

The network object is identified by a unique pair of either <DVS, DVS_PG> or <Host, Host_PG>.

Consider the following example:

vCenter Params					**DCNM Params**	
vCenter	**DVS**	**DVPortGroup/ Network**	**ESXi Host**	**Port Group/ Network**	**Fabric Name**	**Network Name**
172.28.12.123	DVS1	DPG1			Fab1	Network 10
172.28.12.123	DVS1	DPG1			Fab2	Network 30
172.28.12.123			172.28.12.11	PG10	Fab1	Network 20
172.28.12.123			172.28.12.12	PG20	Fab1	Network 20

This table has the mapping for vCenter 172.28.12.123. It has four entries:

- The first entry specifies that for DPG1 in DVS1, the network in DCNM is ‘Network10’ in fabric ‘Fab1’. There can be cases where in the hosts of the DVS can connect to switches in multiple fabrics. The network name in each fabric can be different, so you need the fabric name as well. The example in the table shows one such case in the second entry.
- The second entry specifies the same <DVS1, DPG1> pair being mapped to Network 30 in the fabric ‘Fab2’.
- The third entry specifies that for PG10 in the host 172.28.12.11, the network in DCNM is ‘Network20’ in fabric ‘Fab1’.
- The fourth entry specifies that for PG20 in host 172.28.12.11, the network in DCNM is ‘Network20’ in fabric ‘Fab1’.

As seen in the earlier table, the network object is identified by a unique pair of either <DVS, DVS_PG> or <Host, Host PG>. If there’s a value specified for DVS, DVS_PG, then the values for <Host, Host_PG> are blank. In other words, <DVS, DVS_PG> and <Host, Host_PG) are mutually exclusive.

When the earlier table is specified in a CSV format, it appears as below in the CSV file:

```
172.28.12.123,DVS1,DPG1,,,Fab1,Network10
172.28.12.123,DVS1,DPG1,,,Fab2,Network30
172.28.12.123,,,172.28.12.11,PG10,Fab1,Network20
172.28.12.123,,,172.28.12.12,PG20,Fab1,Network20
```

Let’s consider more examples:

- **172.28.10.184,DSwitchPad,DSPad-PG2,,,DEF,MyNetwork_30000**

  This line in the CSV file specifies the IP address of vCenter as 172.28.10.184 and the <DVS, DVS PG> values are DSwitchPad, DSPad-PG2 respectively. Since the values for DVS, DVS-PG is specified, the values for Host, Host-PG are blank as seen in this example. The Fabric name is DEF and the network in DCNM is MyNetwork_30000.

- **172.28.10.184,,,172.28.11.33,Pad_Workload_Auto_Nwk,DEF,MyNetwork_60000**

  In this example, the values for <DVS, DVS-PG> is left blank and the values for <Host, Host_PG> is specified as 172.28.11.33 and Pad_Workload_Auto_Nwk respectively. The fabric in DCNM is DEF and the network name in DCNM is MyNetwork_60000.

An example CSV file is as follows:

```
vCenter,Dvs,Dvs_pg,Host,Host_pg,Fabric,Network
172.28.10.184,DSwitchNew,DPGNew,,,DEF,MyNetwork_30000
172.28.10.184,DSwitchNew,DPGNew,,,ABC,MyNetwork_30000
172.28.10.184,,,172.28.11.33,Pad_Workload_Auto_Nwk,DEF,MyNetwork_60000
```

Installing and Starting the VMM Workload Automation Module

You can install the VMM workload automation module by using the PIP install or the install script.

Using PIP Install

Before you begin

This installation method is for users who are familiar with **pip install** and know how to set up the proxy or handle cases when there's a conflict in the python packages.

Procedure

Step 1 Decide whether you want to run this module in a virtual environment or on a physical server. If you decide to run this on a server, ensure that the you have the write permission for doing pip install.

Step 2 Setup the http_proxy, https_proxy, and no_proxy appropriately.

For example:

```
export http_proxy=http://proxy.esl.cisco.com:80
export https_proxy=https://proxy.esl.cisco.com:80
export no_proxy=127.0.0.1,172.28.10.0/24
```

In this example, 172.28.10.0 specified in the no_proxy is the management subnet of DCNM.

Step 3 Download and install the module from https://pypi.org/.

```
pip3 install vmm-workload-auto
```

Similarly, you can uninstall the module using the command: **pip3 uninstall vmm-workload-auto**.

Step 4 By default, the installation will happen in the following directories unless the you override by giving options in the **pip** command.

The package is installed under:
`/usr/local/lib/python3.7/site-packages/vmm_workload_auto-0.1.1.dist-info`

The config files are installed under:
`/usr/local/lib/python3.7/site-packages/etc/vmm_workload_auto`

The source code is placed under: `/usr/local/lib/python3.7/site-packages/workload_auto`

Step 5 Edit the config files in:
`/usr/local/lib/python3.7/site-packages/etc/vmm_workload_auto`.

For more information, see Configuration File for Mapping DCNM and vCenter.

Make sure that the path of the CSV file specified in the **conf.yml** file is correct.

Step 6 Start the VMM Workload automation module.

The entry point for the python module is: `/usr/local/bin/vmm_workload_auto`.

You can either run it as:

```
/usr/local/bin/vmm_workload_auto
```

Or

```
vmm_workload_auto
```

If `/usr/local/bin/` is already in **$PATH**.

Provide the config file as a command-line option.

```
/usr/local/bin/vmm_workload_auto
--config=/usr/local/lib/python3.7/site-packages/etc/vmm_workload_auto/conf.yml
```

Using the Install Script

Using the install script is an alternate method for users who don't want to use **pip install**. The install script performs the installation and starts the python module.

Procedure

Step 1 Navigate to https://pypi.org/project/vmm-workload-auto/ and download the latest .tar.gz file.

Step 2 Untar it. For example:

```
tar -xvf vmm_workload_auto-0.1.0.tar.gz
```

Step 3 Modify the config/conf.yml and config/sample.csv according to your environment.

Step 4 Run the setup script as "source setup.sh".

Step 5 The install script initially prompts the user to edit the conf.yaml and .csv files. The script will then prompt the user for proxy and other details. After everything is complete, the script installs the python packages and starts the module automatically.

Step 6 For information about installing the script, see the Installation section in the README file at https://pypi.org/project/vmm-workload-auto/.

Post Installation

After running the workload automation module, navigate to the DCNM Networks window and check whether the network attachments are completed. Review the configuration and deploy it, if **AutoDeploy** is set to **false** in the configuration file (conf.yml).

Additional Functionalities Using REST APIs

The workload automation module also provides the following REST APIs:

Note The REST APIs execute in a separate window after the VMM Workload automation module is running. Make sure that the automation module is running before running the REST APIs.

- Refresh - When the CSV file is changed, a refresh operation needs to be performed. This operation rereads the file and applies any new configuration if needed. The refresh API is as follows:

  ```
  curl -XPOST http://127.0.0.1:{port}/workload_auto/refresh
  ```

- Resync - When there's any change in the DVS-PG, PG, VLAN, or neighbor switches, then a resync operation is needed. If there are any changes found, the configuration is reapplied accordingly. The resync API is as follows:

  ```
  curl -XPOST http://127.0.0.1:{port}/workload_auto/resync
  ```

- Clean - In order to clean up the network provisioning that was previously done using the module, a cleanup operation is needed. The clean API is as follows:

```
curl -XPOST http://127.0.0.1:{port}/workload_auto/clean
```

Events in vCenter

In the DCNM Release 11.4(1), real-time event processing isn't done using the module. The various relevant events and its significance to this module are as follows:

Refresh

The refresh API enables the module to read the CSV file again and apply the network configuration to one or more the relevant switches. The refresh operation needs to be performed for the following events:

- Add PG: Create an entry in the CSV file that specifies the associated network in DCNM for this PG. After the entry is added, call the refresh REST API.
- Add DPG: Create an entry in the CSV file that specifies the associated network in DCNM for this DPG. After the entry is added, call the refresh REST API.

Resync

The resync API enables the module to discover the network objects and its associated properties again. The result of this resync operation is applying the network configuration to the new or changed switches or interfaces. Perform the resync operation for the following events:

- Add host to a DVS.
- Modify VLAN in a DPG or PG.
- Change in topology: When any of the following information is changed, issue the Resync REST API to rediscover the topology and applying the REST API.
    - Neighbor switch change: This can happen if the attached leaf switch is replaced with a new switch or rewired to a different switch.
    - Interface change: This can happen due to rewiring to a different interface in the switch.
    - Host pNIC change.
    - Add an extra connection: This can happen when:
        - A regular interface in the host is made a port-channel by connecting an extra interface from the host to the switch.
        - An extra interface in the host connecting to a different switch forming a vPC pair.

No Action Required

You need not perform any action for the following events:

- Add stand-alone host.
- Add vSwitch.

- Add DVS.
- Delete DVS.

Mapping Change

The different scenarios for a mapping change in a CSV are as follows:

- If a new mapping is added, run the refresh API after adding the mapping in the CSV file.
- If the mapping between the vCenter network to the DCNM network needs to change, then run the clean REST API, modify the mapping the CSV file, and run the refresh REST API.
- If the existing mapping needs to be deleted, then run the clean REST API, delete the mapping in the CSV file and run the refresh API.

Other Events

The other events and the operation that aren't part of a category are as follows:

- Host removed from DVS: When a host is removed from the DVS, the network configuration in the associated leaf switch and connected interface needs to be removed. This needs to be done for all the DPGs of this DVS. Navigate to DCNM and unattach the appropriate networks.
- DPG or PG Delete: For all the network mappings specified in the spec file that are associated with this DPG or PG, remove the network configuration in the relevant switches and interfaces. Navigate to DCNM and unattach the appropriate networks.
- Port Down or Switch Down: If the port or switch is permanently going to be offline, the configuration needs to be removed out of band. If the switch isn't reachable from the host, but it's still managed by DCNM, navigate to DCNM and unattach the appropriate networks.

Management

The Management menu includes the following submenus:

Resources

Cisco DCNM allows you to manage the resources. The following table describes the fields that appear on this page.

Field	Description
Scope Type	Specifies the scope level at which the resources are managed. The scope types can be **Fabric**, **Device**, **DeviceInterface**, **DevicePair**, **Fabric**, and **Link**.
Scope	Specifies the resource usage scope. Valid values are the switch serial numbers or fabric names. Resources with serial numbers are unique, and can be used on the serial number of the switch only.

Field	Description
Allocated Resource	Specifies if the resources are managed with device, device interface, or fabric. Valid values are ID type, subnet, or IP addresses.
Allocated To	Specifies the entity name for which the resource is allocated.
Resource Type	Specifies the resource type. The valid values are **TOP_DOWN_VRF_LAN**, **TOP_DOWN_NETWORK_VLAN**, **LOOPBACK_ID**, **VPC_ID**, and so on.
Is Allocated?	Specifies if the resource is allocated or not. The value is set to **True** if the resource is permanently allocated to the given entity. The value is set to **False** if the resource is reserved for an entity and not permanently allocated.
Allocated On	Specifies the date and time of the resource allocation.

Allocating a Resource

To allocate a resource 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 **Edit Fabric** icon in the fabric where you want to allocate the resource.

The **Edit Fabric** dialog box appears.

Note Alternatively, you can navigate to the **Edit Fabric** dialog box from the fabric topology window. Click **Fabric Settings** in the **Actions** pane.

Step 3 Choose the **Resources** tab.

Step 4 Uncheck the **Manual Underlay IP Address Allocation** check box.

If you check this check box, provide the IP addresses manually to all resources using the **Resource Allocation** window.

Step 5 Click **Save**.

Step 6 Choose **Control > Management > Resources**.

The **Resource Allocation** window appears. This window lists all the resources under the selected scope.

Step 7 Click the **Allocate Resource** icon.

The **Allocate Resource** dialog box appears.

Step 8 Choose the pool type, pool name, and scope type from the drop-down lists accordingly.

The options for pool type are **ID**, **IP**, and **SUBNET**. Based on the pool type you choose, the values in the **Pool Name** drop-down list changes.

Step 9 Choose the serial number in the **Serial Number** drop-down list.

This field appears for all scope types except for the fabric scope type.

Step 10 Enter the entity name in the **Entity Name** field.

The embedded help gives example names for different scope types.

Step 11 Enter the ID, IP address, or the subnet in the **Resource** field based on what pool type you chose in *Step 3*.

Step 12 Click **Save** to allocate the resource.

Examples to Allocate Resources

Example 1: Assigning an IP to loopback 0 and loopback 1

```
#loopback 0 and 1
    L0_1: #BL-3
        pool_type: IP
        pool_name: LOOPBACK0_IP_POOL
        scope_type: Device Interface
        serial_number: BL-3(FDO2045073G)
        entity_name: FDO2045073G~loopback0
        resource : 10.7.0.1


    # L1_1: #BL-3
    #     pool_type: IP
    #     pool_name: LOOPBACK1_IP_POOL
    #     scope_type: Device Interface
    #     serial_number: BL-3(FDO2045073G)
    #     entity_name: FDO2045073G~loopback1
    #     resource : 10.8.0.3
```

Example 2: Assigning a Subnet

```
#Link subnet
    Link0_1:
        pool_type: SUBNET
        pool_name: SUBNET
        scope_type: Link
        serial_number: F3-LEAF(FDO21440AS4)
        entity_name: FDO21440AS4~Ethernet1/1~FDO21510YPL~Ethernet1/3
        resource : 10.9.0.0/30
```

Example 3: Assigning an IP to an Interface

```
    #Interface IP
    INT1_1: #BL-3
        pool_type: IP
        pool_name: 10.9.0.8/30
        scope_type: Device Interface
        serial_number: BL-3(FDO2045073G)
        entity_name: FDO2045073G~Ethernet1/17
        resource : 10.9.0.9
```

Example 4: Assigning an Anycast IP

```
    #ANY CAST IP
    ANYCAST_IP:
        pool_type: IP
```

```
    pool_name: ANYCAST_RP_IP_POOL
    scope_type: Fabric
    entity_name: ANYCAST_RP
    resource : 10.253.253.1
```

Example 5: Assigning a Loopback ID

```
#LOOPBACK ID
LID0_1: #BL-3
    pool_type: ID
    pool_name: LOOPBACK_ID
    scope_type: Device
    serial_number: BL-3(FDO2045073G)
    entity_name: loopback0
    resource : 0
```

Releasing a Resource

To release a resource from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Management > Resources**.

The **Resource Allocation** window appears. This window lists all the resources under the selected scope.

Step 2 Choose a resource that you want to delete.

Note You can delete multiple resources at the same time by choosing multiple resources.

Step 3 Click the **Release Resource(s)** icon.

A confirmation dialog box appears.

Step 4 Click **Yes** to release the resource.

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.

Container Orchestrator

On Cisco DCNM Web UI, choose **Control > Management > Container Orchestrator**. You can add, delete, edit, and rediscover container types.

You can also watch the video that demonstrates how to use Container Visualization with Cisco DCNM. See Video: Using Container Visualization in Cisco DCNM.

The following table describes the fields and description on Container Orchestrator window.

Field	Description
Container Type	Displays the type of orchestrator.
Cluster IP	Displays the IP address of the Kubernetes cluster
Cluster Name	Specifies the name of the cluster
Managed	Specifies that the cluster is managed.
Status	Displays the status of the cluster. • **Cert expired** implies that the certificate is expired. You must add the correct certificate again. • **Not reachable** implies that DCNM can't reach the Kubernetes cluster. • **Ok** implies that the cluster is functioning correctly. • **Discovering** implies that the cluster is being discovered. • **Blank** implies that the cluster isn't managed. **Note** Note: If the status is empty, it implies that the cluster isn't managed.
User	Specifies the role of the Kubernetes cluster
Last Updated Time	Displays the time elapsed since the last change.

The following table describes the action you can perform on the Container Orchestrator window.

Field	Description
Add	Click **Add** icon to add a new cluster to the container orchestration. You can add up to four containers.
Delete	Select the Kubernetes cluster and click **Delete** icon to delete.

Field	Description
Edit	Select the Kubernetes cluster and click on the **Edit** icon to edit the cluster details.
Rediscover	Select the Kubernetes clusters and click **Rediscover** to refresh the cluster.

You can perform the following actions on the Container Orchestrator:

Adding Container Orchestrator

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

Before you begin

To add VM-based Kubernetes cluster, ensure that you have successfully configured the VMM on Cisco DCNM before enabling Container Orchestrator Visualization feature. You must add the vCenter, to the VMM, which hosts the VMs on which the VM-based Kubernetes cluster is running.

Ensure that the hostname is unique across all the clusters nodes.

You don't need VMM for Bare-metal-based cluster. For Bare-metal-based cluster, perform the following:

- Edit the server properties on **Web UI > Administration > DCNM Server > Server Properties** to enable LLDP on DCNM. In the **cdp.discover-lldp** field, enter **true** to enable LLDP.
- Ensure that the LLDP feature is enabled on all LEAF switches in the Fabric.
- On the Kubernetes cluster, ensure that LLDP and SNMP services are enabled on all Bare-metal nodes.
- If the Cisco UCS is using an Intel Nic, LLDP neighborship fails to establish due to FW-LLDP.

  Workaround – For selected devices based on the Intel® Ethernet Controller (for example, 800 and 700 Series), disable the LLDP agent that runs in the firmware. Use the following command to disable LLDP:

  echo 'lldp stop' > /sys/kernel/debug/i40e/<bus.dev.fn>/command

  To find the *bus.dev.fn* for a given interface, run the following command and select the ID associated with the interface. The ID is highlighted in the below sample output.

```
[ucs1-lnx1]# dmesg | grep enp6s0
[ 12.609679] IPv6: ADDRCONF(NETDEV_UP): enp6s0: link is not ready
[ 12.612287] enic 0000:06:00.0 enp6s0: Link UP
[ 12.612646] IPv6: ADDRCONF(NETDEV_UP): enp6s0: link is not ready
[ 12.612665] IPv6: ADDRCONF(NETDEV_CHANGE): enp6s0: link becomes ready
[ucs1-lnx1]#
```

Note LLDP feature is enabled on those fabric switches, to which the bare-metal cluster nodes are connected. They can also be connected to the border gateway switches.

If the Fabric, to which the Kubernetes cluster is connected to, is discovered after the Cluster was discovered, you must rediscover the cluster to display the topology correctly.

If the Bare-metal-based Kubernetes cluster is discovered after configuring LLDP, you must rediscover the Baremetal cluster to display the topology correctly.

Procedure

Step 1 Choose **Control > Management > Container Orchestrator**.

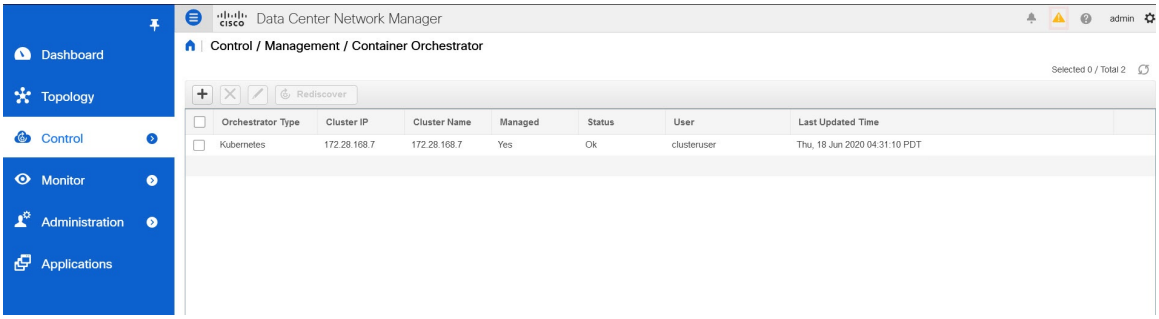

Step 2 Click **Add**.

The Add Container Orchestrator appears.

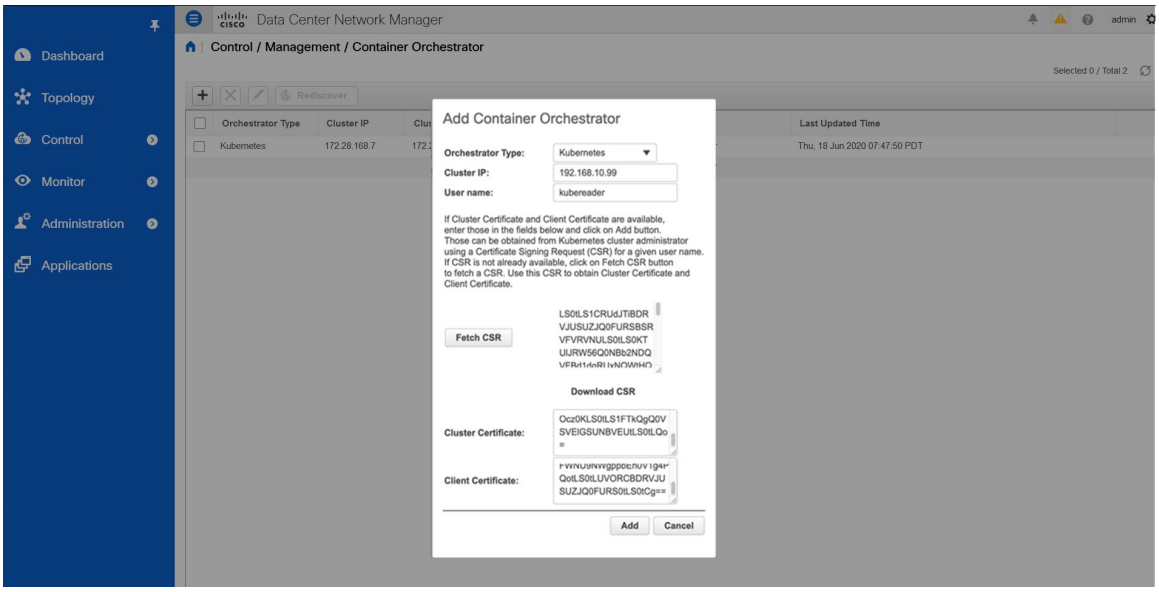

Step 3 From the **Orchestrator** drop-down list, choose **Kubernetes**.

Step 4 In the **Cluster IP** field, enter the IP address of the Master node of the Kubernetes cluster.

Step 5 In the **User Name** field, enter the username of the API Client to connect to Kubernetes.

Step 6 Click **Fetch CSR** to obtain a Certificate Signing Request (CSR) from the Kubernetes Visualizer application.

Note This option is disabled until you enter a valid Cluster IP address and username.

Use the **Fetch CSR** only if you haven't obtained the SSL certificate. If you already have a valid certificate, you need not fetch the CSR.

Click **Download CSR**. The certificate details are saved in the **<username>.csr** in your directory. Paste the contents of the CSR to a file **kubereader.csr**, where, *kubereader* is the username of the API Client to connect to Kubernetes.

The CSR file name must adhere to naming convention <<*username*>>.*csr*.

Note As the certificates are generated on the Kubernetes cluster, you need Kubernetes admin privileges to generate certificates.

The script to generate the certificate **genk8clientcert.sh** is located on the DCNM server at `./root/packaged-files/scripts/genk8sclientcert.sh` location.

Step 7 Login to the Kubernetes cluster controller node.

Note You need admin privileges to generate the certificates.

Step 8 Copy the `genk8clientcert.sh` and `kubereader.csr` from the DCNM server location to the Kubernetes Cluster controller node.

Note Perform a "vnc cut and paste" operation to ensure that all the characters are copied correctly.

Step 9 Generate the CSR for the user name, by using the **genk8sclientcert.sh** script.

```
(k8s-root)# ./genk8sclientcert.sh kubereader 10.x.x.x
```

where,

- `kubereader` is the username of the API Client to connect to Kubernetes. (as defined in Step Step 5, on page 366).
- `10.x.x.x` is the IP address of the DCNM server.

The following message is displayed, after the certificates are generated successfully:

```
------------------------------------------------------------
The K8s CA certificate is copied into k8s_cluster_ca.crt file.
    This to be copied into "Cluster CA" field.
The client certificate is copied into kubereader_10.x.x.x.crt file.
    This to be copied into "Client Certificate" field.
------------------------------------------------------------
```

There are two new certificates generated in the same location:

- k8s_cluster_ca.crt
- *username_dcnm-IP*.crt

  For example: `kubereader_10.x.x.x.crt` (where, `kubereader` is the username, and `10.x.x.x` is the DCNM IP address)

Step 10 Use the **cat** command to extract the certificate from these 2 files.

```
dcnm(root)# cat kubereader_10.x.x.x.crt
dcnm(root)# cat k8s_cluster_ca.crt
```

Provide these two certificates to the user, who is adding the Kubernetes cluster on Cisco DCNM.

Step 11 Copy the content in the `kubereader_10.x.x.x.crt` to **Client Certificate** field.

Note Perform a "vnc cut and paste" operation to ensure that all the characters are copied correctly.

Step 12 Copy the content in the `k8s_cluster_ca.crt` to the **Cluster Certificate** field.

Note Perform a "vnc cut and paste" operation to ensure that all the characters are copied correctly.

Step 13 Click **Add** to add the container orchestrator.

Click **Cancel** to discard adding container orchestrator.

Deleting Container Orchestrator

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

Procedure

Step 1 Choose **Control** > **Management** > **Container Orchestrator**.

Step 2 Select the **Container Orchestrator** that you want to delete.

You can select more than one Cluster at a time.

Click **Delete**.

Note All the data will be deleted if you delete the Cluster. The Cluster will be removed from the Topology view also.

Step 3 Click **Yes** on the confirmation message to delete the Container Orchestrator.

Click **No** to discard.

Editing Container Orchestrator

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

Procedure

Step 1 Choose **Control** > **Management** > **Container Orchestrator**.

Step 2 Select the **Container Orchestrator** that you want to modify. Click **Edit**.

The Edit Container Orchestrator window appears.

Step 3 Modify the values appropriately.

You can update the Cluster and the Client certificates. You can also update the Managed status of the Kubernetes cluster. If you choose to update the Managed status, certificates are not required.

Step 4 Click **Apply** to save the changes.

Click **Cancel** to discard.

Rediscover Kubernetes Cluster

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

Procedure

Step 1 Choose **Control** > **Management** > **Container Orchestrator**.

Step 2 Select the **Container Orchestrator** that you want to rediscover.

You can select more than one Cluster at a time.

Click **Rediscover**.

This action may take some time to refresh the container information.

OpenStack Visualizer

On Cisco DCNM Web UI, choose **Control** > **Management** > **OpenStack Visualizer**. You can add, delete, edit, and rediscover OpenStack Clusters. Note that this is a preview feature.

For information about how to view OpenStack clusters in **Topology**, see OpenStack Workload Visibility, on page 46.

The following table describes the fields and description on the **OpenStack Visualizer** window.

Field	Description
Cluster Type	Specifies the type of cluster.
Cluster IP	Specifies the Controller IP address of the cluster.
Managed	Specifies whether the cluster is managed or unmanaged.
Status	Specifies the status of the cluster.
Username	Specifies the username for the cluster.
Project Name	Specifies the project name.
Region	Specifies the region.
User Domain	Specifies the user domain.
Project Domain	Specifies the project domain.
Last Updated Time	Specifies the last updated time.

The following table describes the action you can perform on the **OpenStack Visualizer** window.

Field	Description
Add	Click **Add** icon to add a new OpenStack cluster to the container orchestration.
Delete	Select the OpenStack cluster and click **Delete** icon to delete.
Edit	Select the OpenStack cluster and click on the **Edit** icon to edit the cluster details.
Rediscover	Select the OpenStack clusters and click **Rediscover** to refresh the cluster.

Adding OpenStack Cluster

This task show how to add an OpenStack cluster.

Before you begin

- Navigate to **Administration > DCNM Server > Server Properties**. Make sure to the set the **cdp.discover-lldp** property to **True** and click **Apply Changes**.

  On the OpenStack cluster, ensure that the LLDP service is enabled on all the bare-metal nodes. LLDP feature is enabled on those fabric switches, to which the bare-metal cluster nodes are connected. They can also be connected to the border gateway switches.

- You can change the resync timer by using the **openstackviz.resync.timer** property. The default value is 60 minutes. Note that you can't set this value below 60 minutes. The resync function restarts the OpenStack plugin and rediscovers all the OpenStack clusters.

- For selected devices based on the Intel® Ethernet Controller (for example, 800 and 700 Series), disable the Link Layer Discovery Protocol (LLDP) agent that runs in the firmware. Use the following command to achieve the same:

  ```
  # echo 'lldp stop' > /sys/kernel/debug/i40e/bus.dev.fn/command
  ```

  To find *bus.dev.fn* for a given interface, run the following command and select the ID associated with the interface. The ID is highlighted in the below output.

  ```
  # dmesg | grep eth0
  [ 8.269557] enic 0000:6a:00.0 eno5: renamed from eth0
  [ 8.436639] i40e 0000:18:00.0 eth0: NIC Link is Up, 40 Gbps Full Duplex, Flow Control:
   None
  [ 10.968240] i40e 0000:18:00.0 ens1f0: renamed from eth0
  [ 11.498491] ixgbe 0000:01:00.1 eno2: renamed from eth0
  ```

Procedure

Step 1 Navigate to **Control > Management > OpenStack Visualizer**.

Step 2 Click the **Add** icon to add an OpenStack Cluster.

- You should at least have read permissions to fetch the cluster information (for example, VMs and Host information).
- In DCNM Release 11.5(1), you can add a cluster only with a single project and region.

Step 3 In the **Add OpenStack Cluster** window, specify the following details:

- **Orchestrator Type**: Specifies the type of orchestrator. By default, OpenStack is selected from this drop-down list.
- **Server IP**: Specifies the Controller IP address of the OpenStack cluster.
- **Port**: Specifies the port number.
- **Version**: Specifies the version.
- **Username** and **Password**: Specifies the username and password of the OpenStack cluster.

- **Project**: Specifies the project name.
- **Region**: Specifies the region. The default region is **RegionOne**.
- **User Domain**: Specifies the user domain. The default user domain is **default**.
- **Project Domain**: Specifies the project domain. The default project domain is **default**.
- **AMQP Endpoint**: Specifies a colon (:) separated multi-valued field containing the address details of an AMQP endpoint. The value should be specified in the format: **username:password:port**. The fields specify the following information:
  - **username**: Specifies the username of the AMQP endpoint.
  - **password**: Specifies the password of the AMQP endpoint.
  - **port**: Specifies the port number of the AMQP endpoint.

  The default value for this field is **guest:guest:5672**.

Step 4 Click **Add**.

After discovery, the status changes from **Discovering** to **Ok**. The information that is received from the OpenStack Cluster is appropriately organized and displayed on the main **Topology** window. An extra menu item labeled **OpenStack** appears on the **Show** pane.

Editing OpenStack Cluster

Procedure

Step 1 Navigate to **Control > Management > OpenStack Visualizer**.

Step 2 Select the OpenStack cluster that you want to modify. Click **Edit**.

In the **Edit OpenStack Cluster** window, you can edit the following fields:

- **Username** and **Password**: Specifies the username and password of the OpenStack cluster.
- **Managed**: You can select **unmanaged** to unmanage an OpenStack cluster.
- **AMQP Endpoint**: Specifies a colon (:) separated multi-valued field containing the address details of an AMQP endpoint. The value should be specified in the format: **username:password:port**. The fields specify the following information:
  - **username**: Specifies the username of the AMQP endpoint.
  - **password**: Specifies the password of the AMQP endpoint.
  - **port**: Specifies the port number of the AMQP endpoint.

  The default value for this field is **guest:guest:5672**.

Step 3 Click **Apply** to save the changes.

Click **Cancel** to discard.

Deleting OpenStack Cluster

Procedure

Step 1 Navigate to **Control > Management > OpenStack Visualizer**.

Step 2 Select the OpenStack cluster that you want to delete. Click **Delete**.

Upon deletion of a cluster from the inventory view, OpenStack plugin stops fetching and receiving the change notifications from the cluster, shuts down the connection with the removed cluster, and releases all software resources.

Step 3 Click **Yes** on the confirmation message to delete the OpenStack cluster.

Click **No** to discard.

Rediscovering OpenStack Cluster

Procedure

Step 1 Navigate to **Control > Management > OpenStack Visualizer**.

Step 2 Select a specific cluster or all the clusters that you want to rediscover. Click **Rediscover**.

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 3: 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.

Field	Description
Save Template As	Allows you to save the selected template in a different name. You can edit the template as required.
Delete Template	Allows you to delete a template
Import Template	Allows you to import a template from your local directory, one at a time.
Export template	Allows you to export the template configuration to a local directory location.
Import Template Zip File	Allows you to import `.zip` file, that contains more than one template that is bundled in a `.zip` format All the templates in the ZIP file are extracted and listed in the table as individual templates.

Note Notifications appear next to **Import Template Zip File** if there are issues while loading templates after restarting the server. Click the notifications to see the errors in the **Issues in loading Template** window. Templates with errors are not listed in the **Templates** window. To import these templates, correct the errors, and import them.

From Cisco DCNM Release 11.4(1), you can only view templates with the **network-operator** role. You cannot modify or save templates with this role. However, you can create or modify templates with the **network-stager** role.

Table 4: Template Properties

Field	Description
Template Name	Displays the name of the configured template.
Template Description	Displays the description that is provided while configuring templates.
Tags	Displays the tag that is assigned for the template and aids to filter templates based on the tags.
Supported Platforms	Displays the supported Cisco Nexus platforms compatible with the template. Check the check box of platforms that are supported with the template. **Note** You can select multiple platforms.
Template Type	Displays the type of the template.
Template Sub Type	Specifies the sub type that is associated with the template.
Template Content Type	Specifies if it is Jython or Template CLI.

Table 5: Advanced Template Properties

Field	Description
Implements	Displays the abstract template to be implemented.
Dependencies	Specifies the specific feature of a switch.
Published	Specifies if the template is published or not.
Imports	Specifies the base template for importing.

In addition, from the menu bar, choose **Control > Template Library > Templates** and you can also:

- Click **Show Filter** to filter the templates that is based on the headers.
- Click **Print** to print the list of templates.
- Click **Export to Excel** to export the list of template to a Microsoft Excel spreadsheet.

This section contains the following:

Template Structure

The configuration template content mainly consists of four parts. Click the **Help** icon next to the **Template Content** for information about editing the content of the template.

This section contains the following:

Template Format

This section describes the basic information of the template. The possible fields are as detailed in the table below.

Property Name	Description	Valid Values	Optional?
name	The name of the template	Text	No
description	Brief description about the template	Text	Yes
userDefined	Indicates whether the user created the template. Value is ‘true’ if user created.	“true” or “false”	Yes
supportedPlatforms	List of device platforms supports this configuration template. Specify ‘All’ to support all platforms.	N1K, N3K, N3500, N4K, N5K, N5500, N5600, N6K, N7K, N9K, MDS, VDC, N9K-9000v, IOS-XE, IOS-XR, Others, All Nexus Switches list separated by comma.	No

Property Name	Description	Valid Values	Optional?
templateType	Specifies the type of Template used.	• CLI • POAP **Note** POAP option is not applicable for Cisco DCNM LAN Fabric deployment. • POLICY • SHOW • PROFILE • FABRIC • ABSTRACT • REPORT	Yes

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

Property Name	Description	Valid Values	Optional?

		• CLI
		  • N/A
		• POAP
		  • N/A
		  • VXLAN
		  • FABRICPATH
		  • VLAN
		  • PMN
		Note POAP option is not applicable for Cisco DCNM LAN Fabric deployment.
		• POLICY
		  • VLAN
		  • INTERFACE_VLAN
		  • INTERFACE_VPC
		  • INTERFACE_ETHERNET
		  • INTERFACE_BD
		  • INTERFACE_PORT_CHANNEL
		  • INTERFACE_FC
		  • INTERFACE_MGMT
		  • INTERFACE_LOOPBACK
		  • INTERFACE_NVE
		  • INTERFACE_VFC
		  • INTERFACE_ANYCAST_PORT_CHANNEL
		  • DEVICE
		  • FEX
		  • INTRA_FABRIC_LINK
		  • INTER_FABRIC_LINK	

Property Name	Description	Valid Values	Optional?
		• INTERFACE • SHOW • VLAN • INTERFACE_VLAN • INTERFACE_VPC • INTERFACE_ETHERNET • INTERFACE_BD • INTERFACE_PORTCHANNEL • INTERFACE_FC • INTERFACE_MGMT • INTERFACE_LOOPBACK • INTERFACE_NVE • INTERFACE_VFC • INTERFACE_SAN_PORTCHANNEL • DEVICE • FEX • INTRA_FABRIC_LINK • INTER_FABRIC_LINK • INTERFACE • PROFILE • VXLAN • FABRIC • NA	

Property Name	Description	Valid Values	Optional?
		• ABSTRACT • VLAN • INTERFACE_VLAN • INTERFACE_VPC • INTERFACE_ETHERNET • INTERFACE_BD • INTERFACE_PORT_CHANNEL • INTERFACE_FC • INTERFACE_MGMT • INTERFACE_LOOPBACK • INTERFACE_NVE • INTERFACE_VFC • INTERFACE_SAN_PORT_CHANNEL • DEVICE • FEX • INTRA_FABRIC_LINK • INTER_FABRIC_LINK • INTERFACE • REPORT • UPGRADE • GENERIC	

Property Name	Description	Valid Values	Optional?
contentType		• CLI • TEMPLATE_CLI • POAP • TEMPLATE_CLI **Note** POAP option is not applicable for Cisco DCNM LAN Fabric deployment. • POLICY • TEMPLATE_CLI • PYTHON • SHOW • TEMPLATE_CLI • PROFILE • TEMPLATE_CLI • PYTHON • FABRIC • PYTHON • ABSTRACT • TEMPLATE_CLI • PYTHON • REPORT • PYTHON	Yes
implements	Used to implement the abstract template.	Text	Yes
dependencies	Used to select the specific feature of a switch.	Text	Yes

Property Name	Description	Valid Values	Optional?
published	Used to Mark the template as read only and avoids changes to it.	“true” or “false”	Yes

Template Variables

This section contains declared variables, the data type, default values, and valid values conditions for the parameters that are used in the template. These declared variables are used for value substitution in the template content section during the dynamic command generation process. Also these variables are used in decision making and in iteration blocks in the template content section. Variables have predefined data types. You can also add a description about the variable. The following table describes the syntax and usage for the available datatypes.

Variable Type	Valid Value	Iterative?
boolean	true|false	No
enum	`Example: running-config, startup-config`	No
float	Floating number format	No
floatRange	`Example: 10.1,50.01`	Yes
Integer	Any number	No
integerRange	Contiguous numbers separated by “-“ Discrete numbers separated by “,” `Example: 1-10,15,18,20`	Yes
interface	Format: <if type><slot>[/<sub slot>]/<port> `Example: eth1/1, fa10/1/2 etc.`	No
interfaceRange	`Example: eth10/1/20-25, eth11/1-5`	Yes
ipAddress	IPv4 OR IPv6 address	No

Variable Type	Valid Value	Iterative?
ipAddressList	You can have a list of IPv4, IPv6, or a combination of both types of addresses. Example 1: 172.22.31.97, 172.22.31.99, 172.22.31.105, 172.22.31.109 Example 2: 2001:0db8:85a3:0000:0000:8a2e:0370:7334, 2001:0db8:85a3:0000:0000:8a2e:0370:7335, 2001:0db8:85a3:1230:0000:8a2f:0370:7334 Example 3: 172.22.31.97, 172.22.31.99, 2001:0db8:85a3:0000:0000:8a2e:0370:7334, 172.22.31.254	Yes
ipAddressWithoutPrefix	Example: 192.168.1.1 or Example: 1:2:3:4:5:6:7:8	No
ipV4Address	IPv4 address	No
ipV4AddressWithSubnet	Example: 192.168.1.1/24	No
ipV6Address	IPv6 address	No
ipV6AddressWithPrefix	Example: 1:2:3:4:5:6:7:8 22	No
ipV6AddressWithSubnet	IPv6 Address with Subnet	No
ISISNetAddress	Example: 49.0001.00a0.c96b.c490.00	No
long	Example: 100	No
macAddress	14 or 17 character length MAC address format	No

string	Free text, for example, used for the description of a variable
	
	Example:
	string scheduledTime
	{
	
	regularExpr=^([01]\d|2[0-3]):([0-5]\d)$;
	}
		No

Variable Type	Valid Value	Iterative?
string[]	Example: {a,b,c,str1,str2}	Yes

struct	Set of parameters that are bundled under a single variable.
	
	struct <structure name declaration > {
	<parameter type> <parameter 1>;
	<parameter type> <parameter 2>;
	…..
	} [<structure_inst1>] [, <structure_inst2>] [, <structure_array_inst3 []>];
	
	struct interface_detail {
	string inf_name;
	string inf_description;
	ipAddress inf_host;
	enum duplex {
	validValues = auto, full, half;
	};
	}myInterface, myInterfaceArray[];
		No
		Note If the struct variable is declared as an array, the variable is iterative.

wwn (Available only in Cisco DCNM Web Client)	Example: 20:01:00:08:02:11:05:03	No

Variable Meta Property

Each variable that is defined in the template variable section has a set of meta properties. The meta properties are mainly the validation rules that are defined for the variable.

The following table describes the various meta properties applicable for the available variable types.

Variable Type	Description	Variable Meta Property											
		default Value	valid Values	decimal Length	min	max	min Slot	max Slot	min Port	max Port	min Length	max Length	regular Expr
boolean	A boolean value. Example: true	Yes											
enum			Yes										

Variable Type	Description	Variable Meta Property											
		default Value	valid Values	decimal Length	min	max	min Slot	max Slot	min Port	max Port	min Length	max Length	regular Expr
float	signed real number Example: 75.56, -8.5	Yes	Yes	Yes	Yes	Yes							
floatRange	range of signed real numbers Example: 50.5 - 54.75	Yes	Yes	Yes	Yes	Yes							
integer	signed number Example: 50, -75	Yes	Yes		Yes	Yes							
integerRange	Range of signed numbers Example: 50-65	Yes	Yes		Yes	Yes							
interface	specifies interface/port Example: Ethernet 5/10	Yes	Yes				Yes	Yes	Yes	Yes			
interfaceRange		Yes	Yes				Yes	Yes	Yes	Yes			
ipAddress	IP address in IPv4 or IPv6 format	Yes											

Variable Type	Description	Variable Meta Property											
		default Value	valid Values	decimal Length	min	max	min Slot	max Slot	min Port	max Port	min Length	max Length	regular Expr
ipAddrList	You can have a list of IPv4, IPv6, or a combination of both types of addresses. Example 1: 172.22.31.97, 172.22.31.99, 172.22.31.105, 172.22.31.109 Example 2: [redacted] Example 3: 172.22.31.97, 172.22.31.99, [redacted], 172.22.31.254 **Note** Separate the addresses in the list using commas and not hyphens.	Yes											

Variable Type	Description	Variable Meta Property											
		default Value	valid Values	decimal Length	min	max	min Slot	max Slot	min Port	max Port	min Length	max Length	regular Expr
ipAddressWithPrefix	IPv4 or IPv6 Address (does not require prefix/subnet)												
ipV4Address	IPv4 address	Yes											
ipV4AddressWithSubnet	IPv4 Address with Subnet	Yes											
ipV6Address	IPv6 address	Yes											
ipV6AddressWithPrefix	IPv6 Address with prefix	Yes											
ipV6AddressWithSubnet	IPv6 Address with Subnet	Yes											
ISISNetAddress	Example: 49.0000.0000.0000.0000.00												
long	Example: 100	Yes			Yes	Yes							
macAddress	MAC address												

Variable Type	Description	Variable Meta Property											
		default Value	valid Values	decimal Length	min	max	min Slot	max Slot	min Port	max Port	min Length	max Length	regular Expr
string	literal string Example for string Regular expression: string shellcmd { [illegible] }	Yes									Yes	Yes	Yes
string[]	string literals that are separated by a comma (,) Example: {string1, string2}	Yes											

Variable Type	Description	Variable Meta Property											
		default Value	valid Values	decimal Length	min	max	min Slot	max Slot	min Port	max Port	min Length	max Length	regular Expr
struct	Set of parameters that are bundled under a single variable. `struct <structure name declaration > { <parameter type> <parameter 1>; <parameter type> <parameter 2>; ….. } [structarray1] [, structarray2] [, structarray3 []>];`												
wwn	WWN address												

Example: Meta Property Usage

```
##template variables

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

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

struct interface_a{
```

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

##
```

Variable Annotation

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

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

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

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

Annotation Key	Valid Values	Description
IsFEXID	“true” or “false”	
IsGateway	“true” or “false”	Validates if the IP address is a gateway
IsInternal	“true” or “false”	Makes the fields internal and does not display them on the window **Note** Use this annotation only for the ipAddress variable.
IsManagementIP	“true” or “false” **Note** This annotation must be marked only for variable “ipAddress”.	
IsMandatory	“true” or “false”	Validates if a value should be passed to the field mandatorily
IsMTU	“true” or “false”	
IsMultiCastGroupAddress	“true” or “false”	
IsMultiLineString	“true” or “false”	Converts a string field to multiline string text area
IsMultiplicity	“true” or “false”	
IsPassword	“true” or “false”	
IsPositive	“true” or “false”	Checks if the value is positive
IsReplicationMode	“true” or “false”	
IsShow	“true” or “false”	Displays or hides a field on the window
IsSiteId	“true” or “false”	
IsSourceDevice	“true” or “false”	
IsSourceFabric	“true” or “false”	
IsSourceInterface	“true” or “false”	

Annotation Key	Valid Values	Description
IsSourceSwitchName	“true” or “false”	
IsSwitchName	“true” or “false”	
IsRMID	“true” or “false”	
IsVPCDomainID	“true” or “false”	
IsVPCID	“true” or “false”	
IsVPCPeerLinkPort	“true” or “false”	
IsVPCPeerLinkPortChannel	“true” or “false”	
IsVPCPortChannel	“true” or “false”	
Password	Text	Validates the password field
PeerOneFEXID	“true” or “false”	
PeerTwoFEXID	“true” or “false”	
PeerOnePCID	“true” or “false”	
PeerTwoPCID	“true” or “false”	
PrimaryAssociation		
ReadOnly	“true” or “false”	Makes the field read-only
ReadOnlyOnEdit	“true” or “false”	
SecondaryAssociation	Text	
Section		
UsePool	“true” or “false”	
UseDNSReverseLookup		
Username	Text	Displays the username field on the window
Warning	Text	Provides text to override the Description annotation

Example: AutoPopulate Annotation

```
##template variables
string BGP_AS;
 @(AutoPopulate="BGP_AS")
```

```
  string SITE_ID;
##
```

Example: DisplayName Annotation

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

Example: IsMandatory Annotation

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

Example: IsMultiLineString Annotation

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

IsShow Annotation

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

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

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

Example: Warning Annotation

```
##template variables
@(Warning="This is a warning msg")
  string SITE_ID;
##
```

Templates Content

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

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

Template content is governed by the usage of variables.

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

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

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

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

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

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

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

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

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

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

```
Syntax: if(<operand 1>  <logical operator>  <operand 2>){
command1 ..
command2..
..
}
else  if (<operand 3> <logical operator> <operand 4> )
{
Command3 ..
Command4..
..
}
else
{
```

```
Command5 ..
Command6..
..
}
Example: if-else if-else statement
if($$USER_NAME$$ == 'admin'){
Interface2/10
no shut
}
else {
Interface2/10
shut
}
```

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

```
Syntax:
foreach <loop index variable> in $$<loop variable>$$ {
@<loop index variable> ..
}
Example: foreach Statement
foreach ports in $$MY_INF_RANGE$${
interface @ports
no shut
}
```

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

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

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

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

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

Template Content Editor

The template content editor has the following features:

- Syntax highlighting: The editor highlights the syntax, like different types of statements, keywords, and so on, for Python scripting.
- Autocompletion: The editor suggests the template datatypes, annotations, or metaproperties when you start typing.
- Go to line: You can navigate to the exact line in the template content editor instead of scrolling. Press **Command-L** in Mac or **Ctrl-L** in Windows, and enter the line number to which you want to navigate to in the pop-up window.

  If you enter a value greater than the number of lines in the editor, you will be navigated to the last line in the editor window.

- Template search and replace: Press **Command-F** in Mac or **Ctrl-F** in Windows, enter the search term in the **Search for** field, and select the type of search in the search window. You can perform the following searches in the editor:
  - **RegExp Search**: You can perform the regular expression search in the editor.
  - **CaseSensitive Search**: You can perform a case-sensitive search in the editor.
  - **Whole Word Search**: You can perform a whole word search to find the exact words in the editor. For example, a regular search for the word "play" returns results where it is part of words like "display," but the whole word search returns results only when there is an exact match for the word "play".
  - **Search In Selection**: You can perform a search in the selected content. Select the content to which you want to limit the search and enter the search term.

  Choose the + icon in the search window to use the replace option. Enter the replacing word in the **Replace with** field. You can replace the selected word once by selecting **Replace**. To replace all the occurrences of the selected word, select **All**.
- Code folding: You can expand or group code blocks in the editor by clicking the arrow next to their line numbers.
- Other features: The editor automatically indents the code, the closing braces, and highlights the matching parenthesis.

Template Editor Settings

You can edit the following features of a template editor by clicking **Template Editor Settings**.

- **Theme**: Select the required theme for the editor from the drop-down list.
- **KeyBinding**: Select the editor mode from the **KeyBinding** drop-down list to customize the editor. **Vim** and **Ace** modes are supported. The default is **Ace**.
- **Font Size**: Select the required font size for the editor.

Advanced Features

The following are the advanced features available to configure templates.

- Assignment Operation

  Config template supports assignment of variable values inside the template content section. The values are validated for the declared data type of the variable. If there is a mismatch, the value is not assigned.

  Assignment operation can be used under the following guidelines:

  - The operator on the left must be any of the template parameters or a for loop parameter.
  - The operator on the right values can be any of the values from template parameters, for loop parameters, literal string values surrounded by quotes or simple string values.

  If a statement does not follow these guidelines, or if it does not suit this format, it will not be considered as assignment operation. It is substituted during command generation like other normal lines.

  ```
  Example: Template with assignment operation
  ```

```
##template properties
name =vlan creation;
userDefined= true;
supportedPlatforms = All;
templateType = CLI;
published = false;
##
##template variables
integerRange vlan_range;
@(internal=true)
integer vlanName;
##
##template content
foreach vlanID in $$vlan_range$${
vlan @vlanID
$$vlanName$$=@vlanID
name myvlan$$vlanName$$
}
##
```

- Evaluate methods

  Config template uses the Java runtime provided Java script environment to perform arithmetic operations (such as ADD, SUBTRACT, and so on), string manipulations, and so on.

  Locate the JavaScript file in the template repository path. This file contains primary set of arithmetic, string functions. You can also add custom JavaScript methods.

  These methods can be called from config template content section in below format:

```
Example1:
$$somevar$$ = evalscript(add, "100", $$anothervar$$)
```

  Also the *evalscript* can be called inside if conditions as below:

```
if($$range$$ > evalscript(sum, $$vlan_id$$,  -10)){
do something...
}
```

  You can call a method that is located at the backend of the Java script file.

- Dynamic decision

  Config template provides a special internal variable “LAST_CMD_RESPONSE”. This variable stores the last command response from the device during the execution of the command. This can be used in the config template content to make dynamic decisions to deliver the commands that are based on the device condition.

Note The if block must be followed by an else block in a new line, which can be empty.

  An example use case to create a VLAN, if it is does not exist on the device.

```
Example: Create VLAN
##template content
show vlan id $$vlan_id$$
if($$LAST_CMD_RESPONSE$$ contains "not found"){
vlan $$vlan_id$$
}
else{
```

```
}
##
```

This special implicit variable can be used only in the “IF” blocks.

- Template referencing

  You can have a base template with all the variables defined. This base template can be imported to multiple templates. The base template content is substituted in the appropriate place of the extending template. The imported template parameters and the contents can be accessed inside the extending template.

```
Example: Template Referencing
Base template:
##template properties
 name =a vlan base;
 userDefined= true;
 supportedPlatforms = All;
 templateType = CLI;
 published = false;
 timestamp = 2015-07-14 16:07:52;
 imports = ;
##
##template variables
 integer vlan_id;
##
##template content
 vlan $$vlan_id$$
##

Derived Template:
##template properties
 name =a vlan extended;
 userDefined= true;
 supportedPlatforms = All;
 templateType = CLI;
 published = false;
 timestamp = 2015-07-14 16:07:52;
 imports = a vlan base,template2;
##
##template variables
 interface vlanInterface;
##
##template content
 <substitute a vlan base>
 interface $$vlanInterface$$
 <substitute a vlan base>
##
```

When you launch the extended template, the parameter inputs for the base template are also obtained. In addition, the substituted content is used for complete CLI command generation.

Report Template

Starting from Cisco DCNM 11.3(1) Release, a new template type, REPORT, has been added. This template has two subtypes, UPGRADE and GENERIC. The template type is python.

UPGRADE

The UPGRADE template is used for pre-ISSU and post-ISSU scenarios. These templates are listed in the ISSU wizard.

Refer to the default upgrade template packaged in DCNM for more information on pre-ISSU and post-ISSU handling. The default upgrade template is issu_vpc_check.

GENERIC

The GENERIC template is used for any generic reporting scenarios, such as, collecting information about resources, switch inventory, SFPs, and NVE VNI counters. You can also use this template to generate troubleshooting reports.

Resources Report

This report displays information about resource usage for a specific fabric.

The **Summary** section shows all resource pools with the current usage percentages. Use the horizontal scroll bar at the bottom of the window to display more columns.

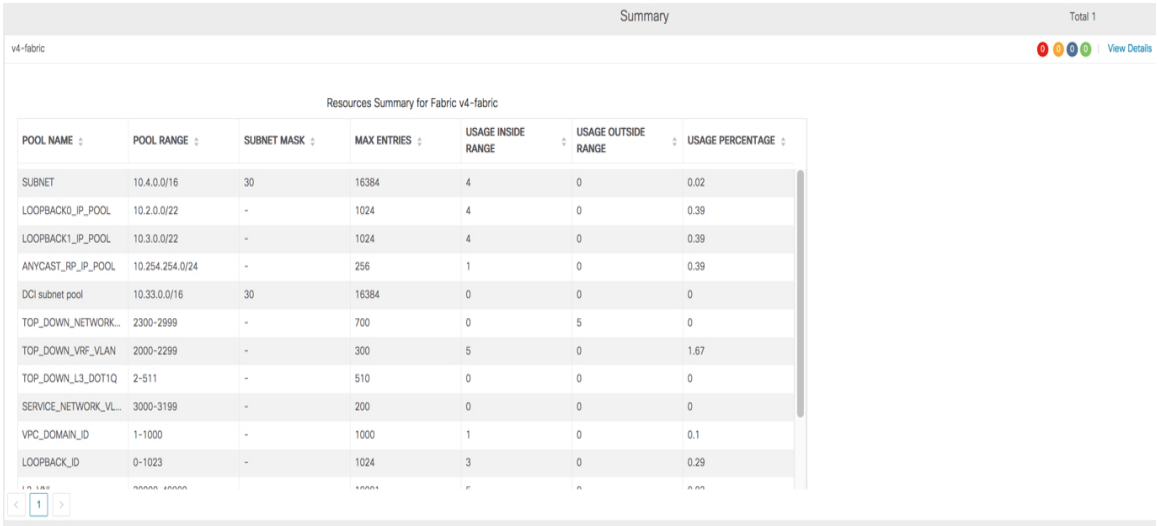

POOL NAME	POOL RANGE	SUBNET MASK	MAX ENTRIES	USAGE INSIDE RANGE	USAGE OUTSIDE RANGE	USAGE PERCENTAGE
SUBNET	10.4.0.0/16	30	16384	4	0	0.02
LOOPBACK0_IP_POOL	10.2.0.0/22	-	1024	4	0	0.39
LOOPBACK1_IP_POOL	10.3.0.0/22	-	1024	4	0	0.39
ANYCAST_RP_IP_POOL	10.254.254.0/24	-	256	1	0	0.39
DCI subnet pool	10.33.0.0/16	30	16384	0	0	0
TOP_DOWN_NETWORK...	2300-2999	-	700	0	5	0
TOP_DOWN_VRF_VLAN	2000-2299	-	300	5	0	1.67
TOP_DOWN_L3_DOT1Q	2-511	-	510	0	0	0
SERVICE_NETWORK_VL...	3000-3199	-	200	0	0	0
VPC_DOMAIN_ID	1-1000	-	1000	1	0	0.1
LOOPBACK_ID	0-1023	-	1024	3	0	0.29

POOL NAME: Specifies the name of the pool.

POOL RANGE: Specifies the IP address range of the pool.

SUBNET MASK: Specifies the subnet mask.

MAX ENTRIES: Specifies the maximum number of entries that can be allocated from the pool.

USAGE INSIDE RANGE: Specifies the current number of entries allocated inside the pool range.

USAGE OUTSIDE RANGE: Specifies the current number of entries set outside the pool range.

USAGE PERCENTAGE: This is calculated by using the formula: (Usage Inside Range/Max Entries) *100.

Click **View Details** to display a view of resources allocated or set in each resource pool. For example, the detailed section for a SUBNET has information about the resources that have been allocated within the subnet.

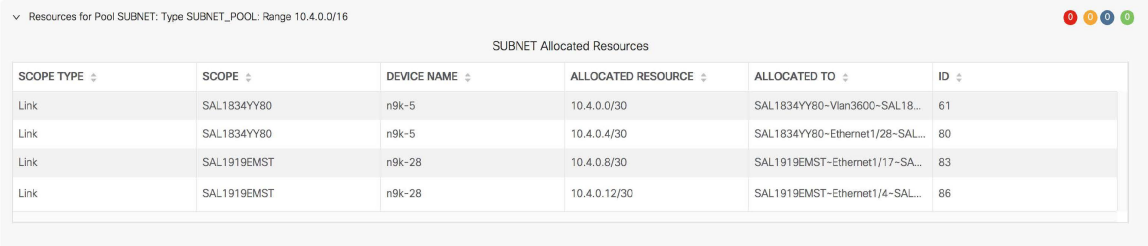

Switch Inventory Report

This report provides a summary about the switch inventory.

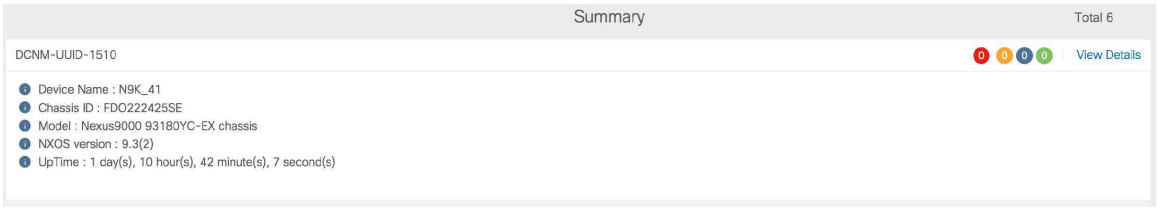

Click **View Details** to display more information about the modules and licenses.

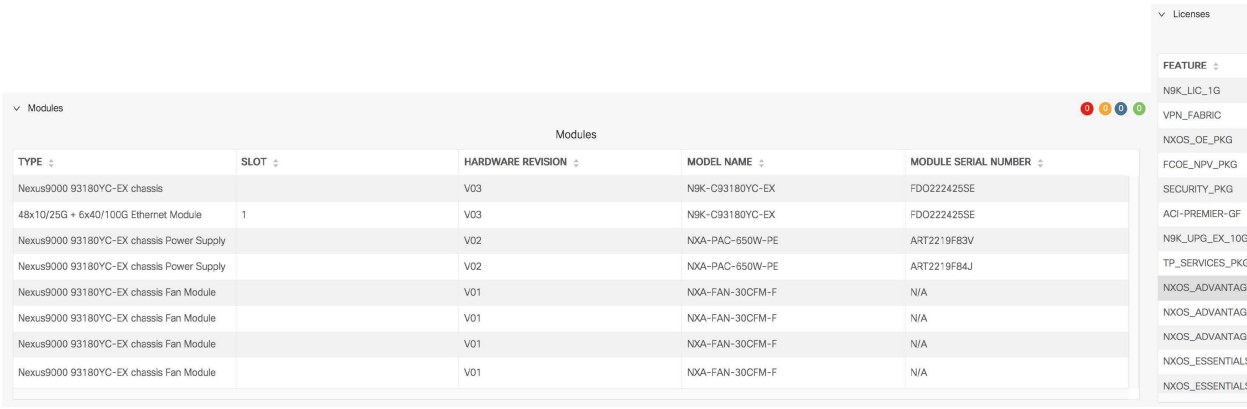

SFP Report

This report provides information about utilization of SFPs at a fabric and device level.

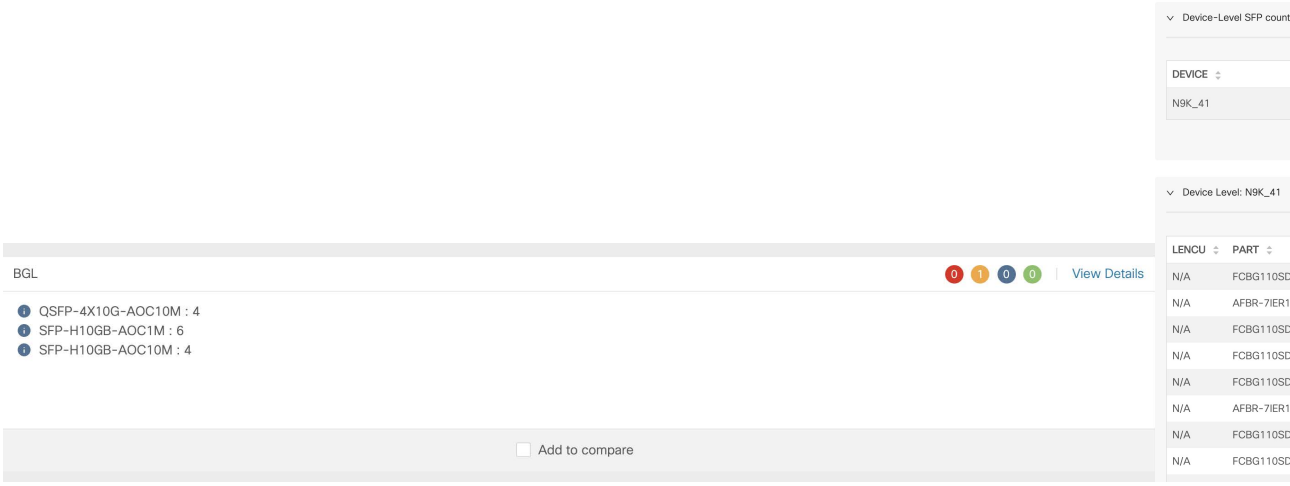

Note The switch inventory and SFP reports are supported only on Cisco Nexus devices.

Troubleshooting Reports

These reports are generated to help in troubleshooting scenarios. Currently, the **NVE VNI Counters** report is the only pre-defined troubleshooting report. Generating **NVE VNI Counters** reports involves performing periodic checks to identify the VNIs that are among the top hits based on network traffic. In a large-scale setup, we recommend limiting the report generation frequency to a minimum of 60 minutes.

NVE VNI Counters Report

This report collects the **show nve vni counters** command output for each VNI in the fabric.

After comparing the oldest report and the newest report, the **Summary** section shows the top-10 hit VNIs. The top hit VNIs are displayed in these categories:

- L2 or L3 VNIs for unicast traffic
- L2 or L3 VNIs for multicast traffic
- L2 only VNIs for unicast traffic
- L2 only VNIs for multicast traffic
- L3 only VNIs for unicast traffic
- L3 only VNIs for multicast traffic

The oldest report refers to the first report that is saved in the current reporting task. If you want to select a specific report as the first report against which the current report has to be compared, delete all reports that are older than the one selected so that the selected report becomes the first and oldest report.

For example, three reports were run yesterday at 8:00 a.m, 4:00 p.m. and 11:00 p.m. If you want to use the report at 11:00 p.m. as the first and oldest report for today's reporting, delete the two reports that were run yesterday at 8:00 a.m. and 4:00 p.m.

For a periodic report, the oldest report is the first report that is run at the start time of a period. For daily and weekly reports, the current report is compared against the previously generated report.

The **Summary** section displays a column-wise report with information about the total transmitted bytes and the VNIs. Use the horizontal scroll bar at the bottom of the window to display more columns.

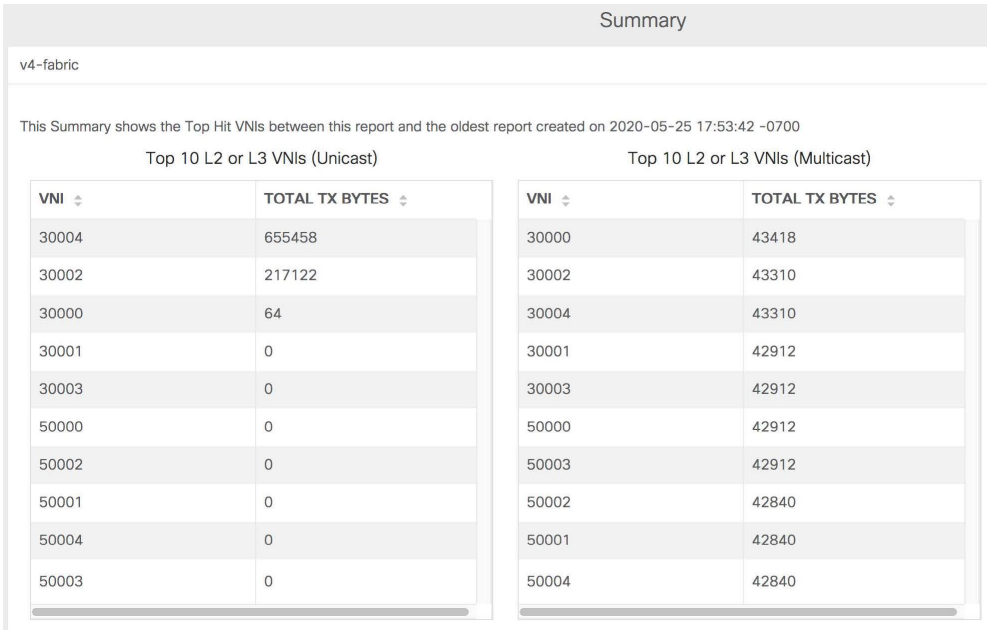
Summary

v4-fabric

This Summary shows the Top Hit VNIs between this report and the oldest report created on 2020-05-25 17:53:42 -0700

Top 10 L2 or L3 VNIs (Unicast)

VNI	TOTAL TX BYTES
30004	655458
30002	217122
30000	64
30001	0
30003	0
50000	0
50002	0
50001	0
50004	0
50003	0

Top 10 L2 or L3 VNIs (Multicast)

VNI	TOTAL TX BYTES
30000	43418
30002	43310
30004	43310
30001	42912
30003	42912
50000	42912
50003	42912
50002	42840
50001	42840
50004	42840

VNI
30004
30002
30000
30001
30003

Note The **Summary** section in the NVE VNI Counters report displays negative numbers in the TOTAL TX BYTES column if a report is generated after a switch reload or after clearing the counters on the switch. The numbers are displayed correctly in the subsequent reports. As a workaround, we recommend deleting all old reports or creating a new job before reloading switches or clearing counters.

Click **View Details** to display more information. This section shows NVE VNIs and counters on a per-switch basis.

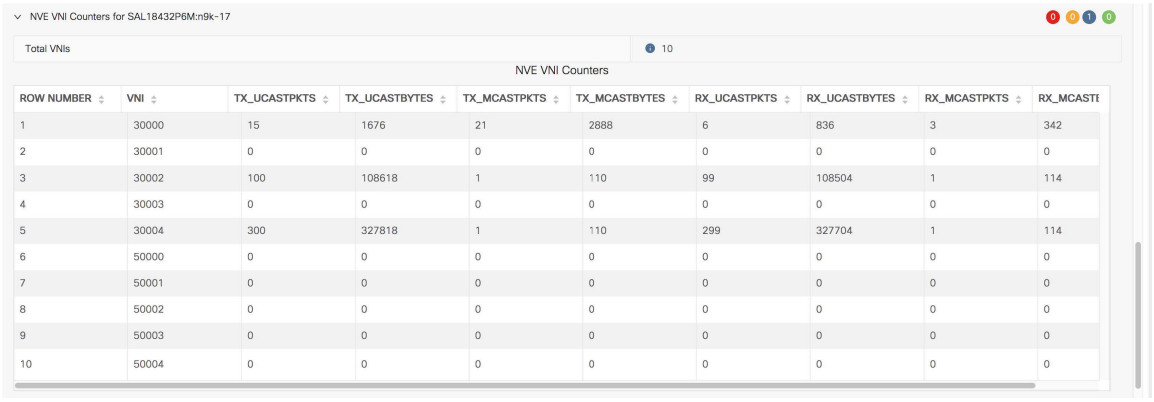
NVE VNI Counters for SAL18432P6M:n9k-17

Total VNIs 10

NVE VNI Counters

ROW NUMBER	VNI	TX_UCASTPKTS	TX_UCASTBYTES	TX_MCASTPKTS	TX_MCASTBYTES	RX_UCASTPKTS	RX_UCASTBYTES	RX_MCASTPKTS	RX_MCASTI
1	30000	15	1676	21	2888	6	836	3	342
2	30001	0	0	0	0	0	0	0	0
3	30002	100	108618	1	110	99	108504	1	114
4	30003	0	0	0	0	0	0	0	0
5	30004	300	327818	1	110	299	327704	1	114
6	50000	0	0	0	0	0	0	0	0
7	50001	0	0	0	0	0	0	0	0
8	50002	0	0	0	0	0	0	0	0
9	50003	0	0	0	0	0	0	0	0
10	50004	0	0	0	0	0	0	0	0

For more information on how the reports are displayed, refer Programmable Report, on page 614.

Report Template Functions

generateReport method

The generateReport method is invoked while generating a report and contains the report implementation logic. This method accepts any context object. and returns a WrappersResp object. For more information on WrappersResp refer link.

Validation method

The validation method is optional. If the template defines this method, the Programmable Report application calls this method to perform pre-validation checks while creating the job. This method is called only when the job is created and invoked only once irrespective of the number of devices or fabrics selected. If the validation passes, this method returns a WrappersResp object with a SuccessRetCode. If the validation fails, this method returns a FailureRetCode along with an error list.

Examples of a successful validation and a failed validation are as follows:

*Successful validation

```
def validate(context):
    respObj = WrappersResp.getRespObj()

 ## Validation logic here

    respObj.setSuccessRetCode()
    return respObj
```

*Failed validation

```
def validate(context):
    respObj = WrappersResp.getRespObj()

 ## Validation logic here

    respObj.setFailureRetCode()
    respObj.addErrorReport(template_name,error)
    return respObj
```

We can also perform validation based on the content of the context parameter.

Context parameter

The Context parameter consists of the following attributes:

- User name - Name of the user who created the job
- User role - Role of the user who created the job
- Job ID
- Recurrence - NOW, ONCE, DAILY, WEEKLY, MONTHLY, ONDEMAND, PERIODIC
- Period - If the recurrence is periodic, then the period will display the selected frequency.

For more information on job context APIs, refer the *Job Context Information* section.

Report Python Library

A REPORT has the following components:

- Summary
  - Key and Values
  - Messages- Inferences
- Details/Sections
  - Key and Values
  - JSON document – Cards
  - Array of JSON Documents – Tables
- Command Log

A python library is provided to generate the **report** JSON model. To use these APIs, the following import statement has to be added to the template:

```
from reportlib.preport import Report
```

Report APIs

Create Report

To create a 'Report' object, use this API –

```
report = Report ("Report title")
```

Add Summary

Each report can have a summary. This is a python dictionary. To add a summary, use this API –

```
summary = report.add_summary ()
```

Adding Content to the Summary

To add content to the summary, use the following APIs-

Key and Values -

```
summary ['NXOS Version'] = '8.4(1)'
```

Messages and Inferences -

```
summary.add_message ("Simple message")
```

Note In Cisco DCNM Release 11.4(1), adding a JSON object as a value in summary is not supported. The following example is not supported-

```
summary["info"] = {"key":"value","key-2":"value-2"}
```

Adding tables in Summary

To add table to the summary, use this API –

```
table = summary.add_table(title,_id)
```

title: Title of the table.

_id: Unique identifier for the table.

Adding rows to the table

To add rows to the table, use this API –

```
table.append(value, _id)
```

value: is a JSON object. Nested JSON is not supported.

_id: is the unique identifier for the row.

Adding a Section

A section is a logical grouping of report content. Sections are created and configured by you to display the required information. To add a section, use this API –

```
section = report.add_section ("Section title",_id)
```

_id: Unique identifier for the section.

Adding Content to a Section Key and Values

To add a simple key and value pair to a section, use this API –

```
section['key'] = 'value'
```

JSON Document – Cards

A JSON document can be added in the same way as a simple key and value pair is added.

Note Nested JSON is not supported in Cisco DCNM Release 11.4(1).

An example of a JSON document displayed as a card widget is as shown below:

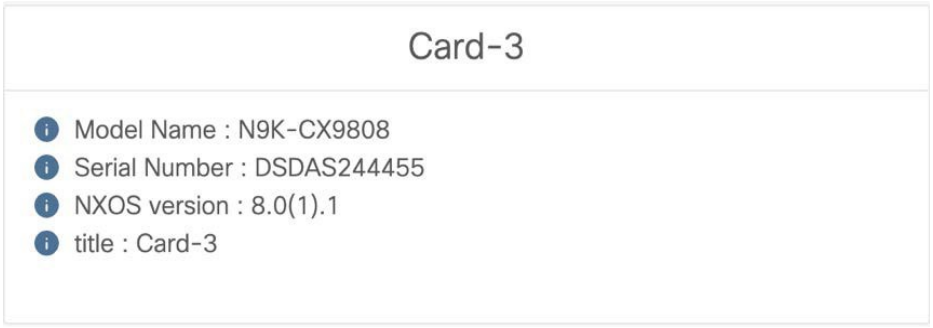

Array of JSON Documents – Tables

To create a table and add rows to this table, use this API –

```
section.append(key, dictionary, _id)
```

_id: Unique identifier that identifies a row in a table. A duplicate *_id* results in a unique id violation error.

Example-

```
section.append('Switch Details', {'name': 'N5K'},'DSDAS244455')
```

The creation of tables using this API has the following limitations:

- All JSON documents should have the same set of keys/columns. Any difference in the number of columns or column names may result in the table not being rendered on the Web UI.

- Nested JSON is not supported.

Formatters

A Formatter enables additional formatting for values that are displayed on the user interface. For example, you can mark values as ERROR, SUCCESS, WARNING, and INFO. These values are color-coded and displayed on the Web UI. Errors are displayed in red, Warnings in yellow, Info in blue and Success in green.

To configure formatting, use this API-

```
Formatter.add_marker(value,marker)
```

value: Value to add a marker

marker: Marker.ERROR, Marker.SUCCESS, Marker.WARNING, and Marker.INFO

Example-

```
Formatter.add_marker ("NXOS version",Marker.INFO)
```

Charts

You can add charts to both the summary and section.

To add a chart to Summary, use this API-

```
report = Report("title")
summary = report.add_summary()
summary.add_chart(ChartType, _id)
```

ChartType: ChartTypes.COLUMN_CHART, ChartTypes.PIE_CHART, and ChartTypes.LINE_CHART

_id: Unique ID for the chart

To add a chart to a section, use this API-

```
report = Report("title")
section = report.add_section ("section_title",_id)
section.add_chart(ChartType, _id)
```

ChartType: ChartTypes.COLUMN_CHART, ChartTypes.PIE_CHART, and ChartTypes.LINE_CHART

_id: Unique ID for the chart

Note Ensure that the classes are imported in the import section.

Pie chart

To display information in a pie chart, use this API-

To set title and subtitle:

```
pie_chart.set_title("Chart title")
pie_chart.set_subtitle("Sub title")
```

To add value:

```
pie_chart.add_value("key",value)
```

key: String key

value: Numeric value

Column chart

To display information in a column chart, use this API-

To set title and subtitle title:

```
column_chart.set_title("Chart title")
column_chart.set_subtitle("Sub title")
```

To set X-Axis and Y-Axis title

```
column_chart.set_xAxis_title("X-Axis title")
column_chart.set_yAxis_title("y-Axis title")
```

To add value:

```
bar_chart.add_value("key",value,category)
```

key: String key

value: Numeric value

category: The column chart divides the data into a logical group that is called a category . A given key should have a value in each category. For example, Device count is a key and Fabric Names are categories. A chart should have a Device count for each fabric.

Line Chart

To display information in a line chart, use this API-

To set title and subtitle title:

```
line_chart.set_title("Chart title")
line_chart.set_subtitle("Sub title")
```

To set X-Axis and Y-Axis title

```
line_chart.set_xAxis_title("X-Axis title")
line_chart.set_yAxis_title("y-Axis title")
```

To add value:

```
line_chart.add_value("key",value,category)
```

key: String key

value: Numeric value

category: The line chart divides the data into logical group called category. A given key should have a value in each category. For example, 'Device count' is a key and 'Fabric Names' are categories. A chart should have a Device count for each fabric or category.

Running CLIs on a Device

To configure running of CLIs on a device, use this API-

```
from reportlib.preport import show
cli_responses = show(serial_number,*commands)
```

serial_number: Serial number of the device on which the commands have to be run. In case of a VDC instance, the serial number is **serial_number:vdc_name**.

**commands*: Commands that are run on the device. These are variable arguments.

Examples-

- Executing a command on single switch:

```
cli_responses = show("FOX1816G0S9",'show version | xml', 'show inventory | xml', 'show
license usage | xml')
```

- Executing a command on multiple switches:

```
cli_responses = show( ["FOX1816G0S9","SSI15470HJ5"],'show version | xml', 'show inventory
 | xml', 'show     license usage | xml')
```

Show commands and store responses

To configure the show commands and store responses, use this API-

```
from reportlib.preport import show_and_store
cli_responses = show_and_store(report,serial_number,*commands)
```

report: Report object created earlier.

serial_number: serial number of the device to run commands. In case of VDC, serial number should be *serial_number:vdc_name*. You can add a list of serial numbers to run the same set of commands on multiple devices.

commands: Commands to run on the device. These are variable arguments. You can specify multiple commands.

Examples-

- Executing a command on single switch:

```
cli_responses = show_and_store(report, "FOX1816G0S9", 'show version | xml', 'show
inventory | xml', 'show license usage | xml')
```

- Executing a command on multiple switches:

```
cli_responses = show_and_store(report, ["FOX1816G0S9","SSI15470HJ5"], 'show version |
xml', 'show inventory | xml', 'show license usage | xml')
```

Note This API stores the response from the device in elasticsearch along with the report. We recommend being cautious while using this API, as storing all responses may reduce available storage space.

Return value

The API mentioned above returns a list of responses. Each response is a dictionary with the following structure-

```
{
'status': 'success' | 'failed,
'response':<response from device>,
'command':<cli command>,
'serial_number': <device serial number>
}
```

In case of multiple switches, the response is a list of responses with separate entries for each switch.

Example-

```
[
{
'status': 'success',
'response':<response from device>,
'command':'show version',
'serial_number': 'FOX1816G0S9'
},
{
'status': 'success',
'response':<response from device>,
```

```
'command':'show version',
'serial_number': 'SSI15470HJ5'
}
]
```

Job context information

To display the recurrence while scheduling the job from the application, use this API-

```
get_recurrence(context)
```

Return values can be NOW,ONCE,DAILY,WEEKLY,MONTHLY,ONDEMAND, and PERIODIC.

If a job is scheduled as Periodic and information about a specific period has to be retrieved, use this API-

```
period = get_period(context)
```

period.get_period(): Returns the period.

period. get_time_unit(): Returns the time unit (HOURS, MINUTES).

Analysis with Historical Reports

Retrieve previously generated reports

Use the *get_previous_reports()* method to get reports that have been generated in the past. This can be used to perform analysis based on current data and historical data. This API will return a list of reports in descending order of the time at which the reports were created.

```
List of reports = get_previous_reports(context,entity,count)
```

context: The object received as input from the generateReport (context) method

entity: serial_number or fabric name

count: Number of reports to fetch

Get oldest report

To retrieve the oldest report, use this API-

```
oldest_report = get_oldest_report(context,entity)
```

context: The object received as input from the generateReport(context) method

entity: serial_number or fabric name

Both the APIs listed above return a Report object with the following APIs to retrieve information-

- Get summary : *report.get_summary()*
- Get section : *report.get_section(_id)* where *_id* is the unique identifier for the section as mentioned in *Adding a Section*.

XML Utilities

The XML utilities are based on xml.etree.elementtree (https://docs.python.org/2/library/xml.etree.elementtree.html).

getxmltree

To return the XML tree under the specified tag, use this API-

```
from reportlib.preport import getxmltree
xml_element_tree = getxmltree(xml_string,tag)
```

xml_string: XML response from device.

tag: XML tag. The complete XML under this tag will be returned as the ElementTree.

xml_element_tree: The API that returns the xml.etree.ElementTree object

getxmlrows

To get an array of rows if the CLI response contains rows, use this API-

```
from reportlib.preport import getxmlrows
rows = getxmlrows(xml_tree,tag_xpath)
```

xml_tree: xml.etree.ElementTree object.

tag_xpath: xpath of the XML record. Please refer https://docs.python.org/2/library/xml.etree.elementtree.html#xpath-support.

rows: An array of rows

getnodevalue

To read the XML node value, use this API-

```
from reportlib.preport import getnodevalue
value = getnodevalue(xml_tree,node_xpath)
```

xml_tree: The **xml.etree.ElementTree** object

node_xpath: xpath of the XML record. Please refer https://docs.python.org/2/library/xml.etree.elementtree.html#xpath-support.

Check for existence of node

This API returns True or False depending on whether the given tag is present or not in the XML tree.

```
from reportlib.preport import
has_tag has_tag(xml_tree,tag)
```

xml_tree: The **xml.etree.ElementTree** object

WrappersResp

Each report has to return an object of the type **WrappersResp**. This can be initiated by using the API given below. Import this from com.cisco.dcbu.vinci.rest.services.jython import WrappersResp.

```
respObj = WrappersResp.getRespObj()
```

The return code in WrapperResp indicates whether the report ran successfully or not.

- If all commands are run and the required information is extracted, then the report returns a success API - *respObj.setSuccessRetCode()*
- In case of any exception such as a command failure, then the report returns a failure API - *respObj.setFailureRetCode()*

Setting a failure code indicates that there is an issue with report execution and the report is not generated.

To return a report with errors, use Formatter to mark the error and set the *WrapperResp* to *Success*. Refer *Formatters* for more information.

For any errors that may come up, you can use this API to specify the reason for the error-

```
respObj.addErrorReport(template_name,error_message)
```

The *report* object created by you should be set to the value of *WrappersResp* as shown below:

```
respObj.setValue(report)
```

Logger

Logger enables logging of messages from the report template. Information that is logged using the logger is logged to this location- "/usr/local/cisco/dcm/fm/logs/preport_jython.log".

```
Logger.info("message")
Logger.debug("message")
Logger.error("message")
Logger.trace("message")
Logger.warn("message")
```

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 411.

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 and options:

Table 6: Image Management Menu

Submenu	Options	Actions
Image Upload	Smart Image Management	Image Upload
		Deleting an Image

Submenu	Options	Actions	
Install & Upgrade	Upgrade History Window Name: **Software Upgrade Tasks**	View	
		Delete	
		New Installation	New ISSU Installation
			EPLD Installation
		Finish Installation	
	Switch Level History	View Device Upgrade Tasks	
		Refresh Switch Level History Table	
Package [SMU/RPM]	Packages	Installing Packages and Patches	
		Uninstalling Packages and Patches	
		Activating Packages and Patches	
		Deactivate	
Image Management Policies	Image Management Policies	Adding an Image Management Policy	
		Deleting an Image Management policy	

Ensure that your user role is **network-admin** or **device-upg-admin** and you didn't freeze the DCNM to perform the following operations:

- Upload or delete images.
- Install, delete, or finish installation of an image.
- Install or uninstall packages and patches.
- Activate or deactivate packages and patches.
- Add or delete image management policies (applicable only for network-admin user role).
- View management policies.

You can view any of the image installations or device upgrade tasks if your user role is **network-admin**, **network-stager**, **network-operator**, or **device-upg-admin**. You can also view them if your DCNM is in freeze mode.

Smart Image Management

This feature allows you to upload or delete images that are used during POAP and switch upgrade. You can also upload or delete RPMs and SMUs that are used for installing in the **Packages** window. To view the **Smart Image Management** window from the Cisco DCNM Web UI homepage, choose **Control > Image Management > Image Upload**.

You can view the following details in the **Smart Image Management** window.

Field	Description
Platform	Specifies the name of the platform. Images, RPMs, or SMUs are categorized as follows: • N9K/N3k • N6K • N7K • N77K • N5K • Other • Third Party The images are the same for N9K and N3K platforms. The platform will be **Other** if the uploaded images are not mapped to any of the existing platforms. The platform will be **Third Party** for RPMs.
Image Name	Specifies the filename of the image, RPM, or SMU that you uploaded.
Image Type	Specifies the file type of the image, EPLD, RPM, or SMU.
Image Subtype	Specifies the file type of the image, EPLD, RPM, or SMU. The file type EPLDs are **epld**. The file types of images are **nxos**, **system** or **kickstart**. The file type for RPMs is **feature** and for SMUs the file type is **patch**.
NXOS Version	Specifies the NXOS image version for only Cisco switches.
Image Version	Specifies the image version for all devices, including the non-Cisco devices as well.
Size (Bytes)	Specifies the size of the image, RPM, or SMU files in bytes.
Checksum	Specifies the checksum of the image. The checksum checks if there's any corruption in the file of the image, RPM, or SMU. You can validate the authenticity by verifying if the checksum value is same for the file you downloaded from the Cisco website and the file you upload in the **Image Upload** window.

You can sort all columns.

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. RPMs and SMUs are used in the **Packages** window. All the images, RPMs, and SMUs are used in the **Image Management Policies** window.

Your user role should be **network-admin**, or **device-upg-admin** to upload an image. You can't perform this operation with the **network-stager** user role.

Procedure

Step 1 Choose **Control > Image Management > Image Upload**.

The **Smart Image Management** 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**.

You can upload a ZIP file as well. Cisco DCNM processes and validate the image file and categorize it under the existing platforms accordingly. If it doesn't fall under **N9K/N3K**, **N6K**, **N7K**, **N77K**, or **N5K** platforms, the image file is categorized under **Third Party** or **Other** platform. The **Third Party** platform is applicable only for RPMs.

Step 5 Click **OK**.

The EPLD images, RPMs, and SMUs are uploaded to the repository in the following path: **/var/lib/dcnm/upload/<platform_name>**

All NX-OS, kickstart and system images are uploaded to the repository in the following paths: **/var/lib/dcnm/images** and **/var/lib/dcnm/upload/<platform_name>**

The upload takes some time depending on the file size and network bandwidth.

Note You can upload images for all Cisco Nexus Series Switches.

You can upload EPLD images only for Cisco Nexus 9000 Series Switches.

If your network speed is slow, increase the wait time of Cisco DCNM to 1 hour so that the image upload is complete. To increase the wait time from Cisco DCNM Web UI, perform the following steps:

a. Choose **Administrator > DCNM Server > Server Properties**.

b. Search for the **csrf.refresh.time** property, and set the value as **60**.

Note The value is in minutes.

c. Click **Apply Changes**.

d. Restart the DCNM server.

Deleting an Image

To delete an image from the repository from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Image Management > Image Upload**.

The **Smart Image Management** 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.

Install & Upgrade

The **Install & Upgrade** menu includes the following submenus:

Upgrade History

This feature enables you to upgrade the Cisco Nexus Platform Switches using In-Service Software Upgrade (ISSU). This upgrade procedure may be disruptive or non-disruptive based on the device configuration. You can select the Kickstart, System, or NX-OS images from image repository or the file system on the device. To select the images from the repository, the same needs to be uploaded from **Control > Image Management > Image upload** tab.

The following table describes the fields that appear on **Control > Image Management > Upgrade History**.

Field	Description
Task Id	Specifies the serial number of the task. The latest task will be listed in the top. **Note** If Failover is triggered in Native HA, the Task Id sequence number is incremented by 32.
Task Type	Specifies the type of task. • Compatibility • Upgrade
Owner	Based on the Role-Based Authentication Control (RBAC), specifies the owner who initiated this task.
Devices	Displays all the devices that were selected for this task.

Field	Description
Job Status	Specifies the status of the job. • Planned • In Progress • Completed • Completed with Exceptions **Note** If the job fails on a single or multiple devices, the status field shows COMPLETED WITH EXCEPTION indicating a failure.
Created Time	Specifies the time when the task was created.
Scheduled At	Specifies the time when the task is specified to be executed. You can also choose to schedule a task to be executed at a later time.
Completed Time	Specifies the time when the task was completed.
Comment	Shows any comments that the Owner has added while performing the task.

Note After a fresh Cisco DCNM installation, this page will have no entries.

You can perform the following:

View

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

Procedure

Step 1 Choose **Control > Image Management > Install & Upgrade > Upgrade History**, check the task ID check box.

Select only one task at a time.

Step 2 Click **View**.

The **Installation Task Details** window appears.

Step 3 Click **Settings**. Expand the **Columns** menu and choose the details you want to view.

You can view the following information in this window:

- Location of the kickstart and system images
- Compatibility check status

- Installation status
- Pre-ISSU report status and post-ISSU report status
- Descriptions
- Report summary
- Version check results
- Logs

The columns changes according to the task you choose to view. You can see the switch name, IP address, platform details, image name, and installation status for an EPLD task. The report status includes the report summary as well. The report summary includes hyperlinks to detailed pre-ISSU reports and post-ISSU reports. Clicking these hyperlinks takes you to a new tab or window to view the reports. The report summary will also include the commands that you defined in the report templates.

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.
- It takes some time for the upgraded EPLD information to appear. A job is scheduled to fetch updates from the switch to DCNM every five minutes until the switch is reachable.

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

You can install ISSU and EPLD images in Cisco DCNM.

New ISSU Installation

To upgrade the devices that are discovered from the Cisco DCNM, perform the following steps:

Before you begin

Add report templates in the **Template Library** window if you want pre-ISSU and post- ISSU reports. Refer to the default upgrade template packaged in DCNM for more information on pre-ISSU and post-ISSU handling. The default upgrade template is **issu_vpc_check**.

Procedure

Step 1 Choose **Control > Image Management > Install & Upgrade > Upgrade History**.

Step 2 Choose **New Installation > ISSU** 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.

Note Switches that are part of fabrics in the freeze mode or monitor mode aren't listed here. An error message appears if you choose a fabric, which is in freeze mode or monitor mode, from the **Device Scope** drop-down menu.

Step 3 Select the check box to the left of the switch name.

You can select more than one switch.

Step 4 Click **Next**.

The **Pre-Post ISSU Reports** window appears.

Note Pre-Post ISSU Reports are not supported in SAN and Media Controller installations.

Step 5 (Optional) Check the **Skip Pre-Post ISSU Reports** check box to skip the pre-post ISSU reports on switches and go to *Step 8*.

By default, this check box isn't checked.

Step 6 Choose a report template from the **Select Report Template** drop-down list.

Only the templates of **REPORT** template type with **UPGRADE** sub-type that are listed in the **Control > Template Library** window appear in the **Select Report Template** drop-down list.

Step 7 Fill in the required fields in the **General** tab based on the template you chose in *Step 6*.

Step 8 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 9 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 10 Click **Select Image** in the **System Image** column.

The **Software Image Browser** dialog box appears.

Step 11 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) Choose a VRF from the **Select Vrf** drop-down list.

Note This field does not appear for Cisco MDS switches.

This VRF is selected for other selected devices by default. The default value is **management**.

d) Click **OK**.

This image is selected for all other selected devices of same platform type.

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 12 The **Vrf** column indicates the name of the virtual routing and forwarding (VRF).

Step 13 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 14 **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 15 Drag and drop the switches to reorder the upgrade task sequence.

Step 16 (Optional) Uncheck **Skip Version Compatibility** check box if you want to check the compatibility of Cisco NX-OS software version on your device with the upgraded images that you chose.

Step 17 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 18 Click **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:

- **Disruptive**
- **Bios force**
- **Allow non-disruptive**
- **Force non-disruptive**

Disruptive is the default value for Cisco Nexus 9000 Series switches. The upgrade option is **Not Applicable** for other switches.

When you choose **Allow non-disruptive** under **Upgrade Option 1** and if the switch does not support non-disruptive upgrade, then it will go through a disruptive upgrade.

When you choose **Force non-disruptive** under **Upgrade Option 1**, the **Skip Version Compatibility** check box will be unchecked because compatibility check is mandatory for non-disruptive upgrade. If the switches you choose do not support non-disruptive upgrade, a warning message appears asking you to review the switch selection. Use the check boxes to choose or remove switches.

The drop-down list for **Upgrade Option 2** has the following options when you choose **Allow non-disruptive** or **Force non-disruptive** under **Upgrade Option 1**:

- **NA**
- **bios-force**

When you choose **Disruptive** or **Bios-force** under **Upgrade Option 1**, **Upgrade Option 2** is disabled.

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 **Allow 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 19 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 **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 **Compatibility Verification** column displays **Skipped**.

The **Pre-ISSUReport Status** column specifies if the pre-ISSU reports were generated. You can view the compatibility log and the report summary in the **Compatibility Status** column. Click the hyperlink in the report summary to see a detailed report of the pre-ISSU check.

Note The status might take some time to reflect in the Web UI depending in the internet bandwidth.

You can review the switch selection and check or uncheck the switches for upgrading accordingly.

Step 20 Click **Finish Installation Later** to perform the upgrade later.

Step 21 Click **Next**.

Step 22 Check the check box to save the running configuration to the startup configuration before upgrading the device.

Step 23 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 24 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.

Note This option is disabled if you put the device in maintenance mode.

b. Select **Concurrent** to upgrade all the devices at the same time.

Step 25 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.

EPLD Installation

Cisco DCNM supports two types of EPLD image installations or upgrade on Cisco Nexus 9000 Series Switches:

- Upgrade all modules from an EPLD image.
- Upgrade only specific modules from an EPLD image.

To select the images from the repository, upload them from **Control > Image Management > Image Upload**.

To install or upgrade EPLD images in Cisco DCNM, perform the following steps:

Procedure

Step 1 Choose **Control > Image Management > Install & Upgrade > Upgrade History.**

Step 2 Choose **New Installation > EPLD**.

The Cisco Nexus 9000 Series Switches are displayed in the **Select Switches** window.

Note Switches that are part of fabrics in the freeze mode or monitor mode are not listed here. An error message appears if you choose a fabric, which is in freeze mode or monitor mode, from the **Device Scope** drop-down menu.

Step 3 Check the check box to the left of the switch name.

You can choose more than one device.

Step 4 Click **Next**.

The **Specify EPLD Images** window appears. This tab displays the switches, that you selected in the previous screen, and allows you to choose the EPLD images for upgrade.

Step 5 Click **Select Image** in the **EPLD image** column.

The **EPLD Image Browser** dialog box appears.

Step 6 Choose the EPLD image file from the file server or switch file system.

If you choose **File Server**:

a) From the **Select Image** list, choose the appropriate image.

Note
- Only files with IMG 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.
- Only image files present in the **Image Upload** window can be selected. You cannot select images present in any other paths.

b) Click **OK** to choose the EPLD 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 IMG 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 EPLD image or **Cancel** to revert to the **Specify Software Images** dialog box.

Step 7 Choose the VRF from the Select Vrf drop-down list.

The valid values are management, default, and keepalive.

Step 8 (Optional) Check the **Use this Vrf for all other selected devices** check box to use the VRF for all other devices you chose.

Step 9 (Optional) Check the **Use this image for all other selected devices of same platform type** check box to use this image for all other devices you chose.

Step 10 The **Vrf** column indicates the name of the virtual routing and forwarding (VRF).

Step 11 Specify the available space for the **Primary Supervisor** and **Secondary Supervisor** modules of the switch in the **Available Space** column.

The **Available Space** column shows the available memory in MB on the switch (for less than 1 MB, it is shown and marked as KB).

Bootflash browser shows the filename, size, and last modified date for all the files and directories on the switch bootflash. You can delete files by selecting them and clicking **Delete** to increase the available space on the switch.

Step 12 Check if the total size of selected images is greater than available space on a switch in the **Selected Files Size** column.

The **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.

Note The EPLD upgrade fails if the version that is supposed to be returned is not returned.

Step 13 Drag and drop the switches to reorder the upgrade task sequence.

Step 14 Click the hyperlink in the **Module Options** column to choose the module for corresponding switch to upgrade EPLD modules.

The **Module Options** dialog box appears. The default value is **All**, which installs or upgrade all EPLD modules for the switch you chose.

Step 15 Choose the modules.

Step 16 Click **OK**.

Step 17 Choose the FPGA region by clicking the hyperlinks under the **FPGA Region** column.

The valid options are **Primary** and **Golden**.

If you choose the golden upgrade, ensure the BIOS is updated and all the prerequisites are met. See the *Cisco Nexus 9000 Series FPGA/EPLD Upgrade Release Notes* for more information.

Step 18 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** window. You can identify the EPLD upgrade tasks by the task type.

What to do next

After you complete the upgrade on the switch, ensure that you wait for 20 minutes to allow the switch to reboot, and stabilize the SNMP agent. Cisco DCNM discovers polling cycles in order to display the new version of the switch in the **Switch Level History** window in Cisco DCNM Web UI.

You can view the EPLD golden upgrade notifications in the **Events** window. From the homepage of the Cisco DCNM Web UI, choose **Monitor > Switch > Events**.

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 Review the switch selection and check or uncheck the switches for upgrading accordingly.

Step 4 Click **Next**.

Step 5 Check the check box to save the running configuration to the startup configuration before upgrading the device.

Step 6 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 7 You can schedule the upgrade process to occur immediately or later.

a. Select **Deploy Now** to upgrade the device immediately.

b. Select **Choose time to Deploy** and specify the time in DD/MM/YYYY HH:MM:SS format to perform the upgrade later.

Step 8 You can choose the execution mode that is based on the devices and the line cards that you have chosen to upgrade.

a. Select **Sequential** to upgrade the devices in the order in which they were chosen.

Note This option is disabled if you put the device in maintenance mode.

b. Select **Concurrent** to upgrade the devices at the same time.

Step 9 Click **Finish** to complete the upgrade process.

Switch Level History

You can view the history of the upgrade process at a switch level. You can view the current version of the switch and other details.

The following table describes the fields that appear on **Control > Image Management > Install & Upgrade > Switch Level History**.

Field	Description
Switch Name	Specifies the name of the switch
IP Address	Specifies the IP Address of the switch
Platform	Specifies the Cisco Nexus switch platform
Current Version	Specifies the current version on the switch software

Click the radio button next to a switch name to select the switch and view its upgrade history. Click **View** to view the upgrade task history for the selected switch.

The following table describes the fields that appear on **Control > Image Management > Install & Upgrade > Switch Level History > View Device Upgrade Tasks**:

Field	Description
Owner	Specifies the owner who initiated the upgrade.
Job Status	Specifies the status of the job. • Planned • In Progress • Completed
KickStart Image	Specifies the kickStart image that is used to upgrade the Switch.
System Image	Specifies the system image that is used to upgrade the switch.
Completed Time	Specifies the date and time at which the upgrade was successfully completed.
Status Description	Specifies the installation log information of the job.

Packages

Image Management also helps you to install or uninstall the required packages and patches. All RPM packages and SMU patches installed on switches appear in the **Package [SMU/RPM]** window. You can now perform the following actions on packages or patches:

- Install
- Uninstall
- Activate
- Deactivate

You need admin privileges to perform this operation. The following table describes the fields that appear on **Control > Image Management > Package [SMU/RPM]**.

Field	Description
Switch Name	Specifies the name of the switch for which the file is installed.
Serial Number	Specifies the serial number of the switch.
IP Address	Specifies the IP address of the device.
Release	Specifies the release version of the switch OS.
Name	Specifies the name of the file.
Version	Specifies the version the file.
Type	Specifies if the file is a base package, non-base package, or a patch.
Status	Specifies if the package or patch is activated or not. Valid values are **Active** and **Inactive**.

You can perform the following tasks from the **Packages** window:

Installing Packages and Patches

To install a package or a patch from Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Image Management > Package [SMU/RPM]** and click the **Install** icon.

The **Select Devices** window appears.

Note Switches that are part of fabrics in the freeze mode or monitor mode are not listed here. An error message appears if you choose a fabric, which is in freeze mode or monitor mode, from the **Device Scope** drop-down menu.

If the switches are in migration mode, the check boxes will be disabled.

Step 2 Check the check box on the left of the switch name.

You can select more than one switch.

Step 3 Click **Next**.

Step 4 Click **Select Packages** in the **Packages/Patches** column.

The **Packages/Patches Browser** dialog box appears.

Step 5 Choose the file from **File Server** or **Switch File System**.

If you choose **File Server**:

a) From the **Select Image** list, choose the appropriate package or patch that must be installed on the device.

The packages or patches that are uploaded for a particular platform will be listed in this file selector. You can select more than one file to be installed, but select only one patch or package if installation needs reload of the switch.

Check the check box to use the same package for all other selected devices of the same platform.

This package or patch image is selected for other selected devices by default.

b) Click **OK** to choose the patch image.
c) Choose the VRF from the drop-down list.

You can use this VRF for all other selected devices.

This VRF is selected for other selected devices by default.

If you choose **Switch File System**:

a) From the **Select Image** list, choose the appropriate file image that is located on the flash memory of the device.

You can select more than one file to be installed on the device, but select only one patch or package if installation needs reload of device. Only files with RPM or SMU 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**.

Step 6 Click **Finish**.

You can view the list of packages that are installed on the switch in the **Packages** window.

Note When you install a package, it is activated as well.

Uninstalling Packages and Patches

The uninstallation process deactivates the selected package or patch followed by its removal. Only non-base RPM packages and SMU patches can be removed. When you uninstall a base RPM package, it only gets deactivated. Base RPM packages cannot be removed. Select only one patch or package if uninstallation needs reload of device.

To uninstall a package or patch on your devices from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Image Management > Package [SMU/RPM]**.

Step 2 Choose a package or patch and click the **Uninstall** icon.

A confirmation window appears

Step 3 Click **OK**.

You can uninstall more than one package or patch at a time, but all the selected packages or patches should have the same status.

Activating Packages and Patches

You can activate the inactive packages or patches. To activate a package or a patch from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Image Management > Package [SMU/RPM]**.

Step 2 Choose an inactive package or patch, and click the **Activate** icon.

A confirmation dialog box appears.

Step 3 Click **OK**.

The **Installation Task Details** dialog box appears. You can click the hyperlink under the **Status** column to view the installation status details.

Deactivate

You can deactivate the active packages or patches. To deactivate a package or patch from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Image Management > Package [SMU/RPM]**.

Step 2 Choose one or more active packages or patches, and click the **Deactivate** icon.

A confirmation dialog box appears.

Step 3 Click **OK**.

Image Management Policies

The image management policies will have the information of intent of NX-OS images along with RPMs or SMUs. The policies can belong to a specific platform or to an umbrella of different types of platforms. An umbrella type policy can have policies for one or more platforms. Regardless of a switch's platform, you can associate an umbrella image management policy with a group of switches. You can choose only one platform

policy per platform under an umbrella type policy. Based on the policy applied on a switch, Cisco DCNM checks if the required NXOS and RPMs or SMUs are present on the switch. If there is any mismatch between the policy and images on the switch, a fabric warning is generated.

The following table has the fields and descriptions of the **Policies** window.

Field	Description
Policy Name	Specifies the policy name.
Policy Type	Specifies if the policy type is **PLATFORM** or **UMBRELLA**.
Release	Specifies the platform release for platform policies. The field is empty for umbrella policies.
Policy / Package Name	Specifies the patch or package name. The package names are displayed for platform policies and the associated platform policies for umbrella policies.
Platform	Specifies the platform for platform policies.
Policy Description	Specifies the user-defined policy description.

You can perform the following tasks from the **Policies** window:

Adding an Image Management Policy

To add an image management policy from the Cisco DCNM Web UI, perform the following steps:

Before you begin

Upload the images in the **Image Upload** window before creating an image management policy. See the Image Upload, on page 415 section for more information about uploading images.

Procedure

Step 1 Choose **Control > Image Management > Image Management Policies**.

The **Policies** window appears.

Step 2 Click the **Add** icon.

The **Create Image Management Policy** dialog box appears.

Step 3 Choose the policy type.

Valid values are **Platform** and **Umbrella**.

Step 4 a) If you chose the **Platform** policy type, the following fields appear in the **Create Image Management Policy** dialog box.

Fields	Actions
Policy Name	Enter the policy name.
Platform	Choose a platform from the Platform drop-down list. The options will be populated based on the images you upload in the **Image Upload** window. The options for the **Release** drop-down list will be autopopulated based on the platform you choose.

Fields	Actions
Release	Choose the NX-OS version from the **Release** drop-down list. The options for **Package Name** will be autopopulated based on the release you choose.
Package Name	(Optional) Choose the packages.
Policy Description	(Optional) Enter a policy description.

b) If you chose **Umbrella** policy type, the following fields appear in the **Create Image Management Policy** dialog box.

Fields	Actions
Policy Name	Enter the policy name.
Platform Policies	Choose the platform policies under this umbrella policy. Choose only one policy per platform.
Policy Description	(Optional) Enter a policy description.

Step 5 Click **OK**.

A confirmation window appears.

What to do next

Attach the policy to a device. See Attaching an Image Management Policy to Devices, on page 432 section for more information.

Attaching an Image Management Policy to Devices

To attach an image management policy from the Cisco DCNM Web UI, perform the following steps:

Before you begin

Create an image management policy for the switch platforms to which you want to attach the policies in the **Image Management Policies** window. See the Adding an Image Management Policy, on page 431 section for more information.

Procedure

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

The **Fabric Builder** window appears.

Step 2 Choose a fabric.

The fabric topology window appears.

Step 3 Click **Tabular view** in the **Actions** pane.

Step 4 Choose the switches to which you want to attach image management policies in the **Switches** tab.

Step 5 Click the **Image Management Policies** icon.

The **Attach Policy to Device** dialog box appears. The IP address, switch name, serial number, and the policy name of the switches you selected appears in this dialog box.

Step 6 Choose the switches to which the image management policies should be attached.

Step 7 Click the **Add** icon.

You will get a warning if no policies are created for the selected platforms.

Step 8 Choose a policy from the **Select Policy** drop-down list.

All the platform policies and umbrella policies, listed in the **Image Management Policies** window, compatible with the selected switches appear in the drop-down list. Ensure the policy you choose has the information related to the platform of the selected switch. Do not attach policies for non-default VDC.

Step 9 Click **OK**.

The policy name is updated for the switches in the **Attach Policy to Device** dialog box.

Step 10 (Optional) Navigate to the fabric topology window.

Step 11 (Optional) Click **Re-sync Fabric** in the **Actions** pane.

Alternatively, you can wait for the scheduled CC check and verify if the intended NX-OS images, RPMs, or SMUs are installed on the switches.

Step 12 (Optional) Check for any pending errors and resolve them by clicking **Resolve**.

To remove a policy from a switch, follow the above procedure till *Step 6* and click the **Delete** icon in *Step 7*.

Deleting an Image Management policy

To delete an image management policy from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Image Management > Image Management Policies**.

The **Policies** window appears.

Step 2 Click the **Delete** icon.

A confirmation dialog box appears.

Note
- You cannot delete a platform policy that is used in an umbrella policy. Delete the umbrella policy before deleting such platform policies.
- You cannot delete a policy that is in use. Before deleting detach the policy from devices.

Step 3 Click **OK**.

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 is dependent on the scope selected by you from the **SCOPE** drop-down list.

- Endpoint Locator , on page 559
- Monitoring Endpoint Locator, on page 583

Layer 4-Layer 7 Service

Cisco DCNM Release 11.3(1) introduces the ability to insert Layer 4-Layer 7 (L4-L7) service devices in a data center fabric, and also enables selectively redirecting traffic to these service devices. You can add a service node, create route peering between the service node and the service leaf switch, and then selectively redirect traffic to these service nodes.

On the Cisco Web UI, choose **Control > Services**. For information regarding configuring Service Nodes, refer Layer 4-Layer 7 Service, on page 861.

Cross Site Scripting (XSS) threat and mitigation

Cross-Site Scripting (XSS) attacks are a type of injection. Malicious scripts are injected into otherwise benign and trusted websites. XSS attacks occur when an attacker uses a web application to send malicious code. The malicious code is in the form of a browser script to a different end user.

An attacker can use XSS to send a malicious script to an unsuspecting user. The browser can’t realize that the script shouldn’t be trusted and executes the script. Because the browser thinks the script came from a trusted source, the malicious script can access any cookies, session tokens, or other sensitive information retained by the browser and used with that site.

XSS attacks occur when accessibility to DCNM is established. It must have been authorization to access the system and inject a malicious string to DCNM as a data, which can be read back by unsuspecting users on their browser. Therefore, the malicious code is executed. OWASP XSS Cheatsheet provides complete list of special characters which may cause XSS.

Cross Site Scripting (XSS) threat and Handling of special Characters in Policy Fields

Various policy fields traditionally have used values which include strings having special characters.

Example

```
Port mode = "40G+10G"
Shared secret = <A password having many special characters>
Description = "NYC & SFO, >100G"
```

Note Some of the fields, like 'description' may not have special characters. Other fields such as, 'Port mode' and 'Shared secret' need special characters, as they are tied to NXOS CLI command format or required for interworking of systems.

Handling on DCNM 11.5(1)

DCNM Release 11.5(1) sanitizes the policy-related field contents for special characters based on OWASP guidelines, therefore avoiding the Cross Site Scripting (XSS) attacks. The policy template variables values are scanned for a special set of XSS characters and reported as errors. Because some of the special characters are needed by policy, as per NXOS requirements, DCNM Release 11.5(2) provides a mechanism to allow special characters.

The following image shows a typical error message:

Handling on DCNM 11.5(2)

Cisco DCNM Release 11.5(2) provides a Server Property **ef.sanitize.state** to control the sanitization behavior. The following keywords describe the functionality.

- **Strict**—Sanitizes the content for XSS threat characters as per OWASP guidelines.

  This implies that there are no exceptions. All special characters such as **@ & + \ + % = < >** causes XSS failure.

- **Default**—Sanitizes the content for a reduced set of characters.

  The allowed characters are **@ % & \ + ' = .**

  However, this sanitizes if the allowed characters prefixed with **$** or **< >**.

  Example: **$@** or **<>@** isn't allowed; however, **@** is allowed.

- **Loose**—Disables the sanitization completely.

To update the Server Property on the Cisco DCNM Web UI, choose **Administration** > **Server Properties**.

Default value for this server property is **Default**.

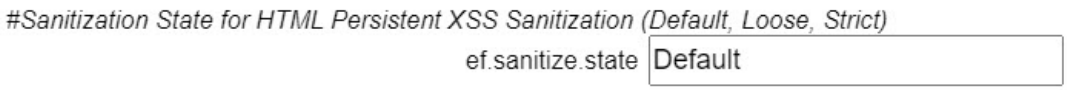

Strict mode provides the efficient defense against XSS, as it prevents storing XSS vulnerable data on Cisco DCNM. However, for practical reasons where traditional templates are used, and/or NXOS CLI command mandating use of special characters, use one of the following mechanisms:

- Set the property value to **Loose** using the following procedure to allow special characters; however, this increases XSS threat. In this case, ensure to consider the following precautions:
  - You can access DCNM using a secure machine, such as, within Data Center VPN. This ensures that the malicious user doesn't reach DCNM easily.
  - Users with role secure**admin** the password avidly, as these operations require admin privilege.
- Create custom policy templates that contain the XSS unsafe content directly in the **Template Content** and then deploy these policies to switches.

Example

Below CLI added to GUI **switch_freeform** policy errors out upon saving the policy owing to XSS threat mitigation enforcements.

```
ip as-path access-list ORIGIN-ACL seq 10 permit "^$"
```

Perform one of the following to mitigate XSS threats:

- Create a custom template. For instructions, refer to Adding a Template, on page 410.

  The following example shows a sample custom template:

  ```
  ##template properties
  name =ip_as_path;
  description = IP AS Path Custom Template;
  tags = ;
  userDefined = true;
  supportedPlatforms = All;
  templateType = POLICY;
  templateSubType = DEVICE;
  contentType = TEMPLATE_CLI;
  implements = ;
  dependencies = ;
  published = false;
  imports = ;
  ##
  ##template variables
  ##

  ##template content
  ip as-path access-list ORIGIN-ACL seq 10 permit "^$"
  ##
  ```

- Add a policy using this template for your switches from View/Edit Policies.
- Deploy the new policy to switches.

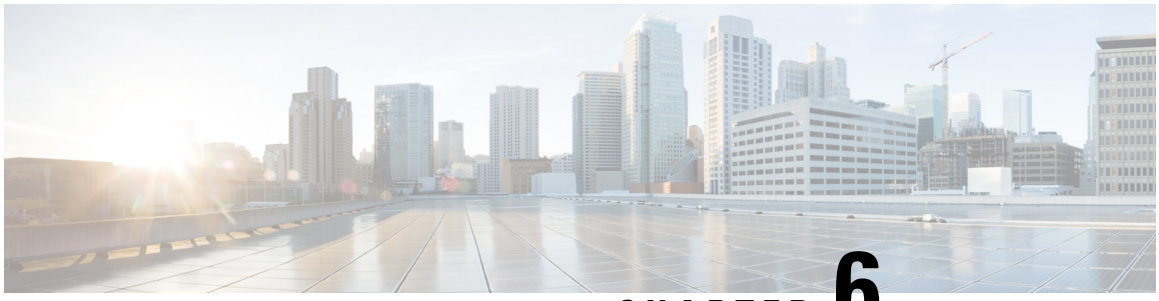

CHAPTER 6

Monitor

This chapter contains the following topics:

- Inventory, on page 437
- Monitoring Switch, on page 456
- Monitoring LAN, on page 459
- Endpoint Locator, on page 463
- Alarms, on page 463

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.

- **Mode** column displays the current mode of the switch. The switch can be in **Normal**, **Maintenance**, or **Migration** mode.
- **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.
- **Up Time** column displays the time period for which the switch is active.

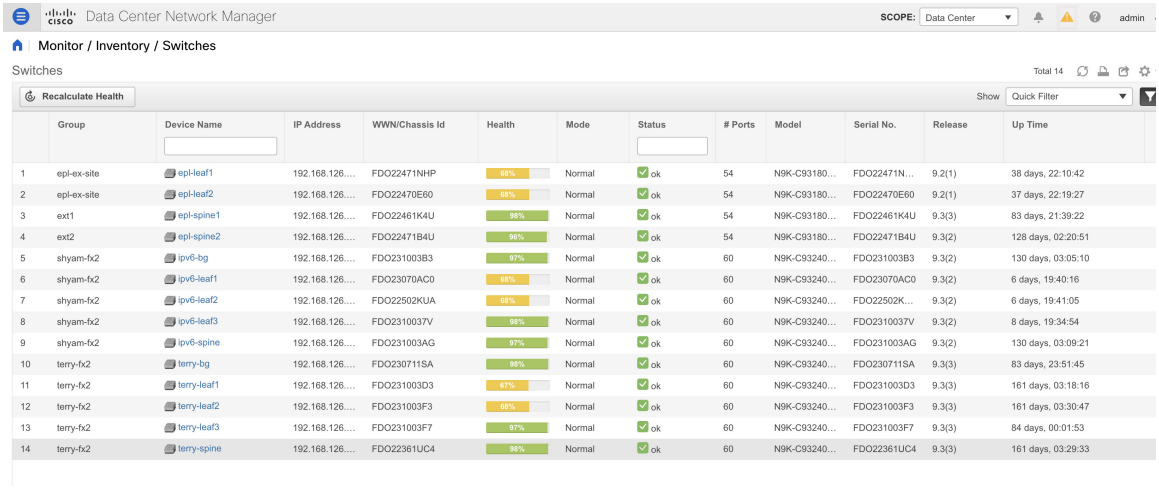

	Group	Device Name	IP Address	WWN/Chassis Id	Health	Mode	Status	# Ports	Model	Serial No.	Release	Up Time
1	epl-ex-site	epl-leaf1	192.168.126....	FDO22471NHP	68%	Normal	ok	54	N9K-C93180...	FDO22471N...	9.2(1)	38 days, 22:10:42
2	epl-ex-site	epl-leaf2	192.168.126....	FDO22470E60	68%	Normal	ok	54	N9K-C93180...	FDO22470E60	9.2(1)	37 days, 22:19:27
3	ext1	epl-spine1	192.168.126....	FDO22461K4U	98%	Normal	ok	54	N9K-C93180...	FDO22461K4U	9.3(3)	83 days, 21:39:22
4	ext2	epl-spine2	192.168.126....	FDO22471B4U	96%	Normal	ok	54	N9K-C93180...	FDO22471B4U	9.3(2)	128 days, 02:20:51
5	shyam-fx2	ipv6-bg	192.168.126....	FDO231003B3	97%	Normal	ok	60	N9K-C93240...	FDO231003B3	9.3(2)	130 days, 03:05:10
6	shyam-fx2	ipv6-leaf1	192.168.126....	FDO23070AC0	68%	Normal	ok	60	N9K-C93240...	FDO23070AC0	9.3(2)	6 days, 19:40:16
7	shyam-fx2	ipv6-leaf2	192.168.126....	FDO22502KUA	68%	Normal	ok	60	N9K-C93240...	FDO22502K...	9.3(2)	6 days, 19:41:05
8	shyam-fx2	ipv6-leaf3	192.168.126....	FDO2310037V	98%	Normal	ok	60	N9K-C93240...	FDO2310037V	9.3(2)	8 days, 19:34:54
9	shyam-fx2	ipv6-spine	192.168.126....	FDO231003AG	97%	Normal	ok	60	N9K-C93240...	FDO231003AG	9.3(2)	130 days, 03:09:21
10	terry-fx2	terry-bg	192.168.126....	FDO230711SA	98%	Normal	ok	60	N9K-C93240...	FDO230711SA	9.3(3)	83 days, 23:51:45
11	terry-fx2	terry-leaf1	192.168.126....	FDO231003D3	67%	Normal	ok	60	N9K-C93240...	FDO231003D3	9.3(3)	161 days, 03:18:16
12	terry-fx2	terry-leaf2	192.168.126....	FDO231003F3	68%	Normal	ok	60	N9K-C93240...	FDO231003F3	9.3(3)	161 days, 03:30:47
13	terry-fx2	terry-leaf3	192.168.126....	FDO231003F7	97%	Normal	ok	60	N9K-C93240...	FDO231003F7	9.3(3)	84 days, 00:01:53
14	terry-fx2	terry-spine	192.168.126....	FDO22361UC4	98%	Normal	ok	60	N9K-C93240...	FDO22361UC4	9.3(3)	161 days, 03:29:33

Step 3 Click **Health** to access the Health score window for a device. The Health score window includes health score calculation and health trend. The Overview tab displays the overall health score. All the modules, switch ports and alarms are taken into consideration while calculating the health score. Hover over the graph under Health Trend for detailed information on specific dates. Hover over the info icon next to Alarms to display the number of Critical, Major, Minor, and Warning alarms that have been generated.

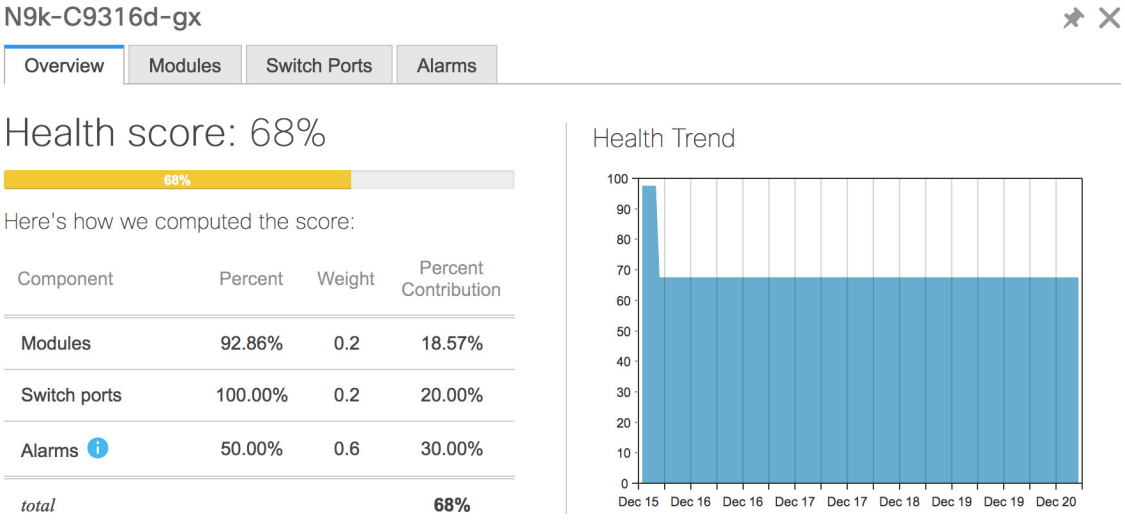

Click the **Modules** tab to display information about the various modules in the device. This tab displays information such as Name, Model name, Serial number, Status, Type, Slot, Hardware revision and Software revision.

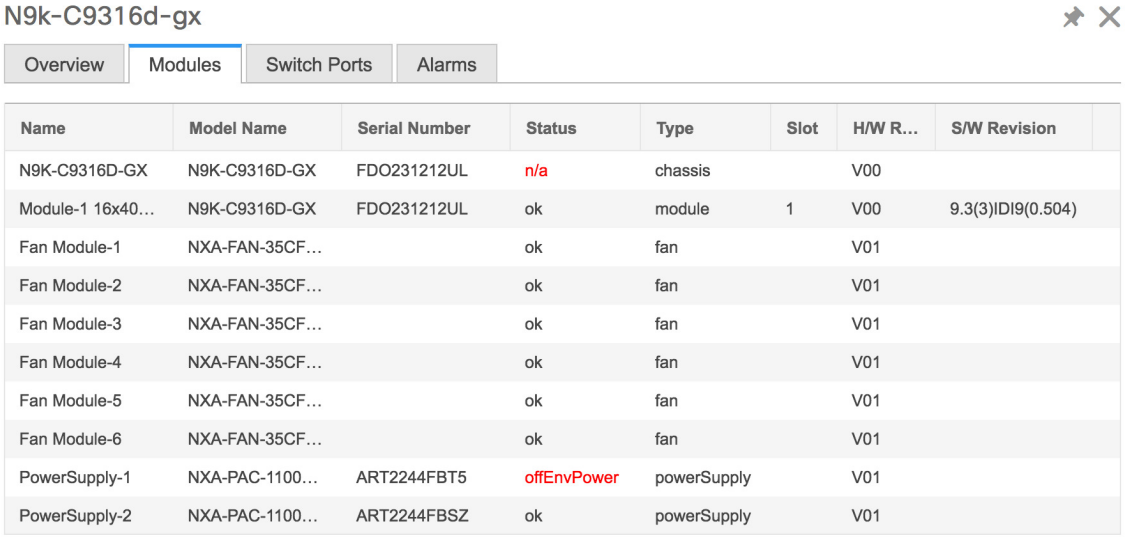

Click the **Switch Ports** tab to display information about the device ports. This tab displays information such as Name, Description, Status, Speed, and the device to which a port is connected .

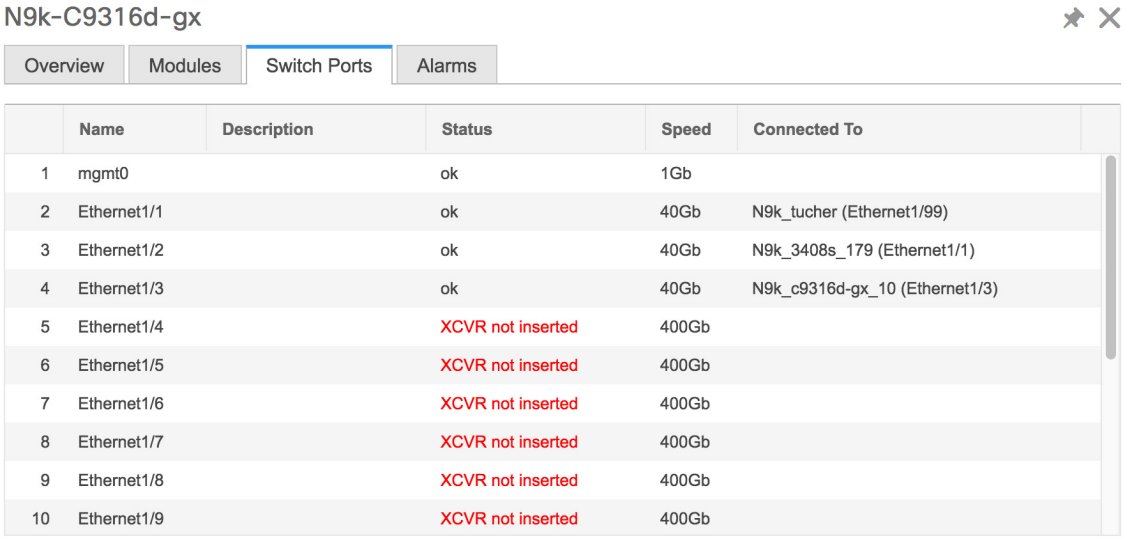

Click the **Alarms** tab to display information about the alarms that have been generated. This tab displays information such as alarm Severity, Message, Category, and the Policy that has been activated due to which the alarm is generated.

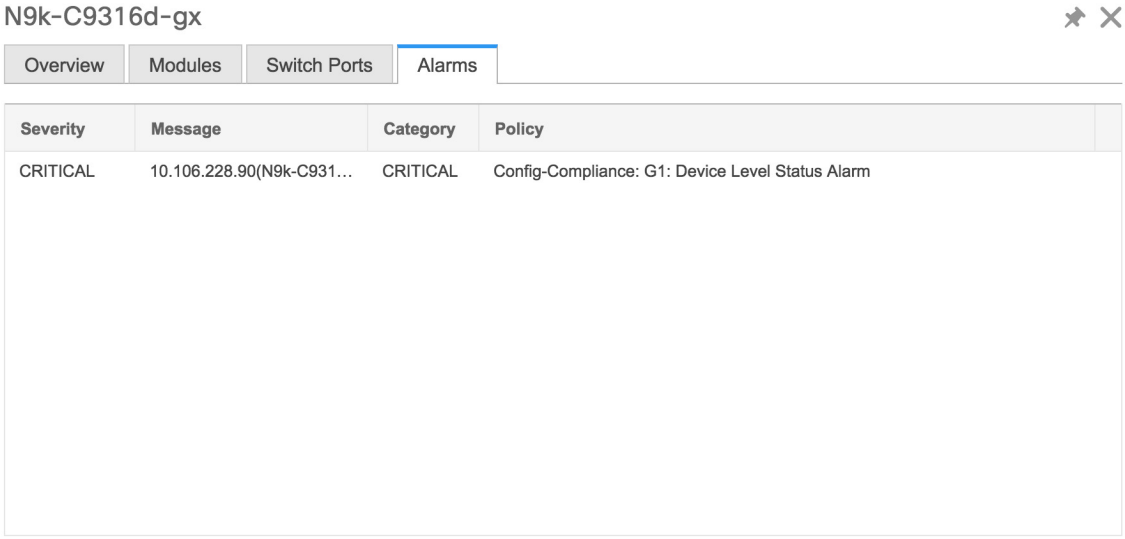

In the **Health** column, the switch health is calculated by the capacity manager based on the following parameters:

- Total number of modules
- Total number of modules impacted by the warning
- Total number of switch ports
- Total number of switch ports impacted by the warning

- Total number of critical severity alarms
- Total number of warning severity alarms
- Total number of major severity alarms
- Total number of minor severity alarms

Step 4 The value in the **Health** column is calculated based on the following:

- Percentage of modules impacted by warnings (Contributes 20% of the total health).
- Percentage of ports impacted by warnings (Contributes 20% of the total health).
- Percentage of alarms (Contributes 60% of the total health). The critical alarms contribute the highest value to this percentage followed by major alarms, minor alarms and warning alarms.

You may also have your own health calculation formula by implementing the common interface class: com.cisco.dcbu.sm.common.rif.HealthCalculatorRif.

The default Java class is defined as: health.calculator=com.cisco.dcbu.sm.common.util.HealthCalculatorAlarms.

- 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

Starting from Cisco DCNM 11.3(1) Release, you can view information about switch health along with the switch summary by clicking on a switch in the **Topology** window or by choosing **Control>Fabrics>Fabric Builder**, selecting a fabric and clicking on a switch in the **Fabric Builder** window.

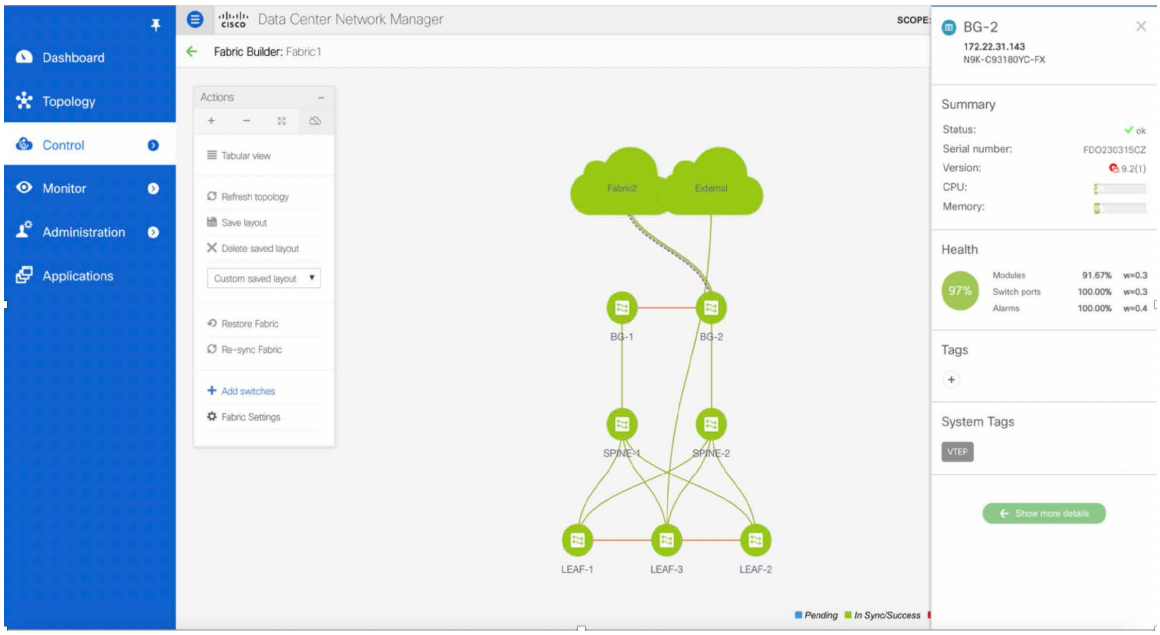

Viewing System Information

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 **Show Commands** to display the device show commands. The Device Show Commands page helps you to view commands and execute them.

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 7: 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 8: 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 9: 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 10: 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.

Field	Description
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 11: 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 **chassis**, **module**, **fan**, and **powerSupply**.
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.
IO FPGA	Specifies the IO field programmable gate arrays (FPGA) version.
MI FPGA	Specifies the MI field programmable gate arrays (FPGA) version.

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 12: FEX Operations

Field	Description
Show	Allows you to view various configuration details for the selected FEX ID. You can select the following from the drop-down list. • show_diagnostic • show_fex • show_fex_detail • show_fex_fabric • show_fex_inventory • show_fex_module The variables for respective show commands are displayed in the Variables area. Review the Variables and click **Execute**. The output appears in the **Output** area. You can create a show template for FEX. Select template type as SHOW and sub type as FEX.

Table 13: 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.

Field	Description
Pinning	An integer value that denotes the maximum pinning uplinks of the Fabric Extender that is active at a time.
State	Specifies the status of the FEX as associated with the Cisco Nexus Switch.
Model	Specifies the model of the FEX.
Serial No.	Specifies the configured serial number. **Note** 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.

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 14: 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.

Field	Description
Suspend	Allows you to suspend an active non-default VDC. Save the VDC running configuration to the startup configuration before suspending the VDC. Otherwise, you will lose the changes to the running configuration. **Note** You cannot suspend the default VDC. **Caution** Suspending a VDC disrupts all traffic on the VDC.
Rediscover	Allows you to resume a non-default VDC from the suspended state. The VDC resumes with the configuration that is saved in the startup configuration.
Show	Allows you to view the Interfaces and Resources that are allocated to the selected VDC. In the Interface tab, you can view the mode, admin-status, and operational status for each interface associated with the VDC. In the Resource tab, you can view the allocation of resources and current usage of these resources.

Table 15: Vdc Table Field and Description

Field	Description
Name	Displays the unique name for the VDC
Type	Species the type of VDC. The two types of VDCs are: • Ethernet • Storage
Status	Specifies the status of the VDC.
Resource Limit-Module Type	Displays the allocated resource limit and module type.

Field	Description
HA-Policy • Single Supervisor • Dual Supervisor	Specifies the action that the Cisco NX-OS software takes when an unrecoverable VDC fault occurs. 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: **Single supervisor module configuration:** • Bringdown—Puts the VDC in the failed state. To recover from the failed state, you must reload the physical device. • Reload—Reloads the supervisor module. • Restart—Takes down the VDC processes and interfaces and restarts it using the startup configuration. **Dual supervisor module configuration:** • Bringdown—Puts the VDC in the failed state. To recover from the failed state, you must reload the physical device. • Restart—Takes down the VDC processes and interfaces and restarts it using the startup configuration. • Switchover—Initiates a supervisor module switchover. 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.
Mac Address	Specifies the default VDC management MAC address.
Management Interface • IP Address Prefix • Status	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 16: Template Resource Limits

Resource	Minimum	Maximum
Global Default VDC Template Resource Limits		
Anycast Bundled		
IPv6 multicast route memory	8	8 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.

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 **Do not interpolate data**.

  Note Set the **pmchart.doInterpolate** property in the **Server Properties** window to false to use the **Do not interpolate data** option.

- 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.

Note To change traffic display unit from bytes to bits, From Cisco DCNM Web UI, choose **Administration > DCNM Server > Server Properties**, enter value as true for **pm.showTrafficUnitAsbit** property, and click **Apply Changes**.

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 **Do not interpolate data**.

  Note Set the **pmchart.doInterpolate** property in the **Server Properties** window to false to use the **Do not interpolate data** option.

- 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 17: vPC Performance, on page 461 displays the following vPC configuration details in the data grid view.

Table 17: vPC Performance

Column	Description
Search box	Enter any string to filter the entries in their respective column.
vPC ID	Displays vPC ID's configured device.
Domain ID	Displays the domain ID of the vPC peer switches.
Multi Chassis vPC EndPoints	Displays the multi-chassis vPC endpoints for each vPC ID under a vPC domain.
Primary vPC Peer - Device Name	Displays the vPC Primary device name.
Primary vPC Peer - Primary vPC Interface	Displays the primary vPC interface.
Primary vPC Peer - Capacity	Displays the capacity for the primary vPC peer.
Primary vPC Peer - Avg. Rx/sec	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 **Do not interpolate data**.

  Note Set the **pmchart.doInterpolate** property in the **Server Properties** window to false to use the **Do not interpolate data** option.

- 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.

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 is dependent on the scope selected by you from the **SCOPE** drop-down list.

- Endpoint Locator , on page 559
- Monitoring Endpoint Locator, on page 583

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 forward alarms to registered SNMP listeners in DCNM. From Cisco DCNM web UI, choose **Administration > DCNM Server > Server Properties**, enter an external port address in **alarm.trap.listener.address** field, click **Apply Changes,** and restart DCNM services.

Note Ensure that you select **Forwarding** check box in **Alarm Policy creation** dialog window to enable forwarding alarms to external SNMP listener.

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.

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 check boxes 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**

  The Regex definitions are simple expressions but not a complete regex. Variable regions of text are noted using $(LABEL) syntax. Each label represents a regex capture group (.+), which corresponds to one or more characters. The variable texts found in both raise and clear messages are used to associate the two messages. An Identifier is a sequence of one or more labels that appear in both messages. An Identifier is used to match a clear syslog message to the syslog message that raised the alarm. If the text appears only in one of the messages, it can be noted with a label and exclude it from the identifier.

  Example: A policy with "Value": "ID1-ID2",

"syslogRaise": "SVC-5-DOWN: $(ID1) module $(ID2) is down $(REASON)"

"syslogClear": "SVC-5-UP: $(ID1) module $(ID2) is up."

In the example, ID1 and ID2 labels can be marked as an identifier to find the alarm. This identifier will be found in corresponding syslog messages. Label “REASON” is in the raise but not in the clear message. This label can be excluded from the identifier, as it has no impact on the syslog message to clear the alarm.

Table 18: 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 19: 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 20: 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.

Enabling External Alarms

You can enable external alarms using one of the following methods:

- Using Cisco DCNM Web UI
  1. From Cisco DCNM Web UI, choose **Administration > DCNM Server > Server Properties**.
  2. Locate the **alarm.enable.external** property.
  3. Enter the value in the field as **true**.
- Using REST APIs
  1. Go the API documentation URL from your DCNM setup: https://<*DCNM-ip*>/api-docs
  2. Navigate to the **Alarms** section.
  3. Click **POST > rest/alarms/enabledisableextalarm**.
  4. Choose the **body** parameter value as **true** from the **Value** drop-down list.
  5. Click **Try it out!**.
- Using CLI
  1. Log into the DCNM server using SSH.
  2. Set the **alarm.enable.external** property to **true** in the server.properties file.

     The filepath is `/usr/local/cisco/dcm/fm/config/server.properties`.

Configuration Compliance Alarms

Starting from Cisco DCNM Release 11.3(1), the alarm policies and alarms under the External category are created by the applications running on DCNM. These External alarm policies are created by the applications and cannot be created or added via the DCNM Web UI.

Config-Compliance(CC) is a core application running on DCNM. CC registers and creates Alarms under the External Alarm category.

Config-Compliance : Alarm Policy

This External alarm category policy is activated on creation of the fabric and is enabled on all the devices in that fabric. The severity level of the policy is CRITICAL. If any device in the fabric moves from In-Sync to Out-of-Sync and the **Enable Alarms** checkbox is selected, a critical severity alarm is generated.

Choose **Monitor>Alarms>Policies** to display the default alarm policies. This alarm policy is not editable on the web UI. Click **Activate** or **Deactivate** to activate or deactivate the selected policy.

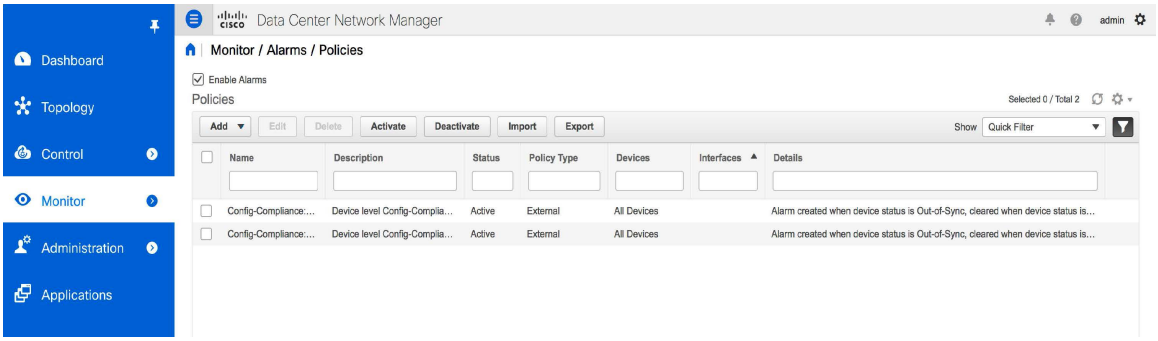

In case an alarm policy is deactivated using the DCNM Web UI, any alarms created or cleared for that policy will not be displayed in the **Monitor>Alarms>View** tab. To delete a policy, select the checkbox next to the policy and click **Delete**. However, we recommend not deleting a policy from the DCNM Web UI. If a policy is deleted, CC regenerates the policy on the next periodic run or when a Re-sync is triggered at the device level or fabric level under that fabric.

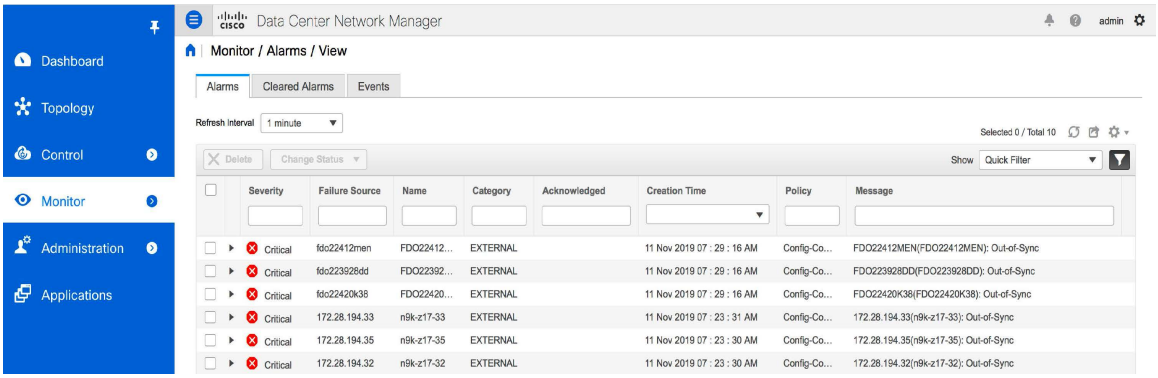

Click the arrow icon next to **Critical** to display detailed information about the alarm.

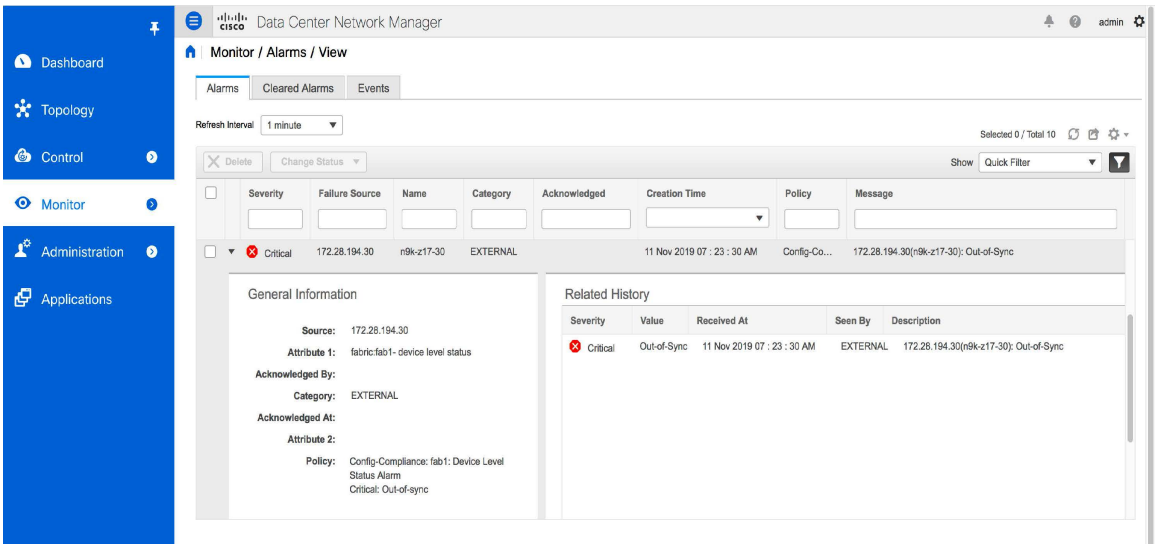

An Out-of-Sync status indicates that there is a difference between the intent defined for the device on DCNM and the running configuration on the device. An In-Sync status indicates that the intent defined for the device on DCNM matches the running configuration and CC found no differences between the configurations. For more details on computation of diff, refer *Configuration Compliance in DCNM*.

When a fabric is deleted, the alarm policy along with all the active alarms for the devices in that fabric are deleted.

Config-Compliance : Active Alarms

Consider a scenario in which CC is running on a fabric and a device in that fabric moves to Out-of-Sync status. This leads to the generation of a Critical severity alarm. Choose **Monitor->Alarms->View** to display the alarms. These alarms are active until the device moves from Out-of-Sync to In-Sync.

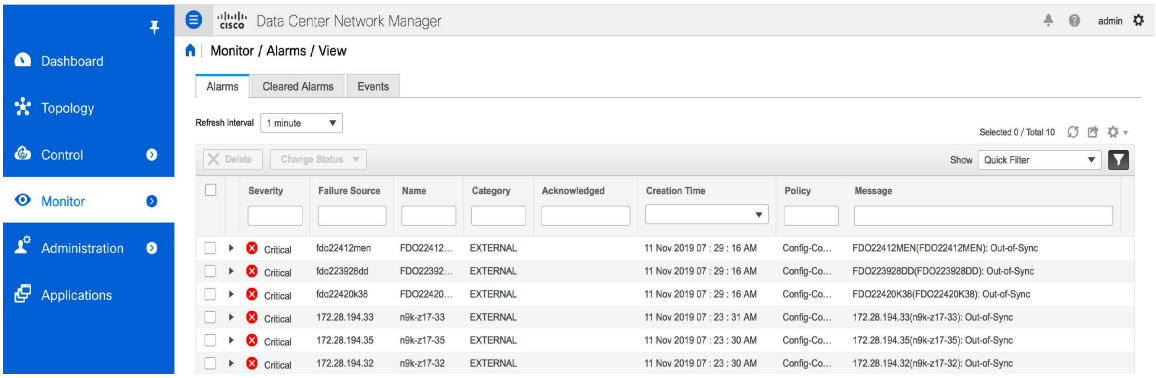

To clear active alarms, select the checkbox next to the alarm, click **Change Status** and select **Clear**. In case the same device moves to Out-of-Sync status again, the active alarm is re-created.

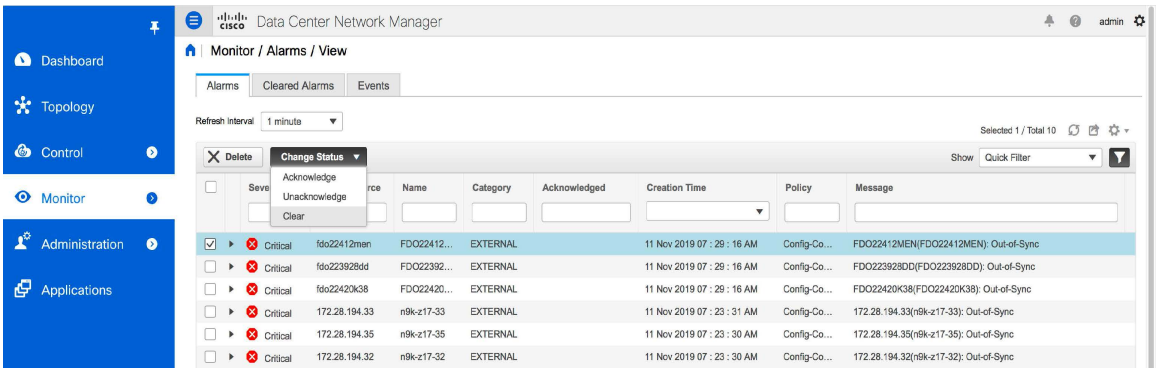

To delete active alarms, select the checkbox next to the alarm and click **Delete**. In case the same device moves to Out-of-Sync status again, a new active alarm is created.

Config-Compliance : Cleared Alarms

When the device that is in Out-of-Sync status moves to In-Sync status, the active alarm is cleared. To view the cleared alarms, select **Monitor>Alarms>View>Cleared Alarms**. The cleared alarms do not contribute to the overall device health score.

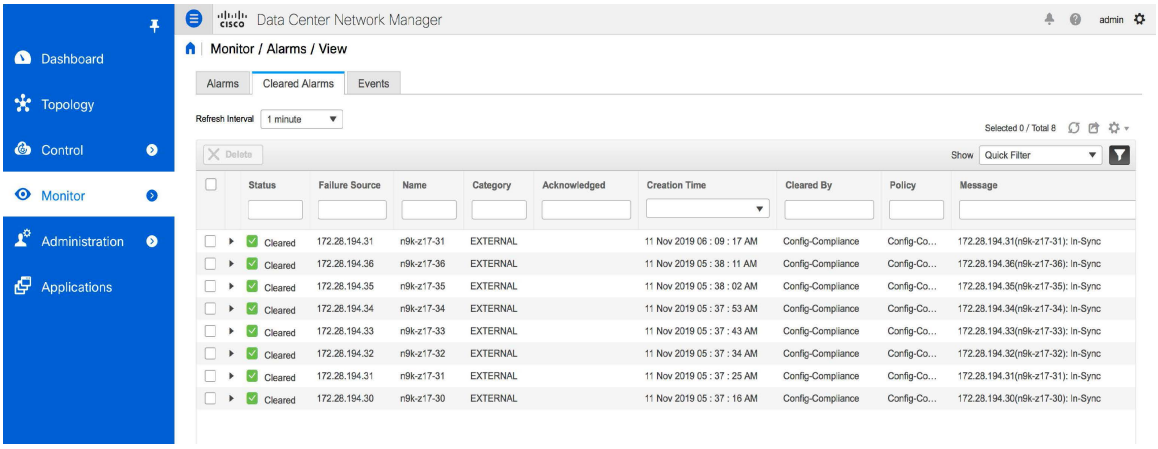

To delete a cleared alarm from the list of cleared alarms, choose **Monitor>Alarms>View>Cleared Alarms**, select the checkbox next to the alarm, and click **Delete**. This will delete the selected cleared alarms from the list.

Alarms are cleared when a switch moves from Out-of-Sync to In-Sync. Configuration compliance alarms also contribute to the overall Health score of a device.

For more information on Alarms and Policies, refer *Alarms*.

Endpoint Locator Alarms

Starting from Cisco DCNM Release 11.4(1), alarms are registered and created under the External alarm category by the Endpoint Locator (EPL).

Endpoint Locator: Alarm Policy

The EPL external alarm category policy is activated when EPL is enabled on a fabric. Alarms are raised for issues such as Duplicate IP addresses, Duplicate MAC addresses, Endpoints appearing on a VRF and Endpoints disappearing from a VRF, Endpoints moving within a fabric, loss of Route Reflector connectivity, and restoration of Route Reflector connectivity. Depending on the issue, the severity level of the alarm policy can be CRITICAL or MINOR.

Alarms are raised and categorized as CRITICAL for the following events:

- Route Reflector disconnection
- Detection of a duplicate IP address
- Detection of a duplicate MAC address

Alarms are raised and categorized as MINOR for the following events:

- Movement of an endpoint
- Appearance of a new VRF in a fabric
- Number of endpoints in a fabric goes down to 0
- Number of endpoints in a VRF goes down to 0
- Disappearance of all endpoints from a switch
- Connection of a Route Reflector (RR)

CRITICAL alarms are cleared automatically when the condition is corrected. For example, when the connectivity between DCNM and RR is lost, a CRITICAL alarm is generated. This alarm is automatically cleared when the connectivity between DCNM and RR is restored. Other MINOR alarms are automatically cleared after 30 minutes have passed since the alarm was generated.

Choose **Monitor>Alarms>Policies** to display the EPL alarm policies. These alarm policies are not editable on the web UI. Click **Activate** or **Deactivate** to activate or deactivate the selected policy.

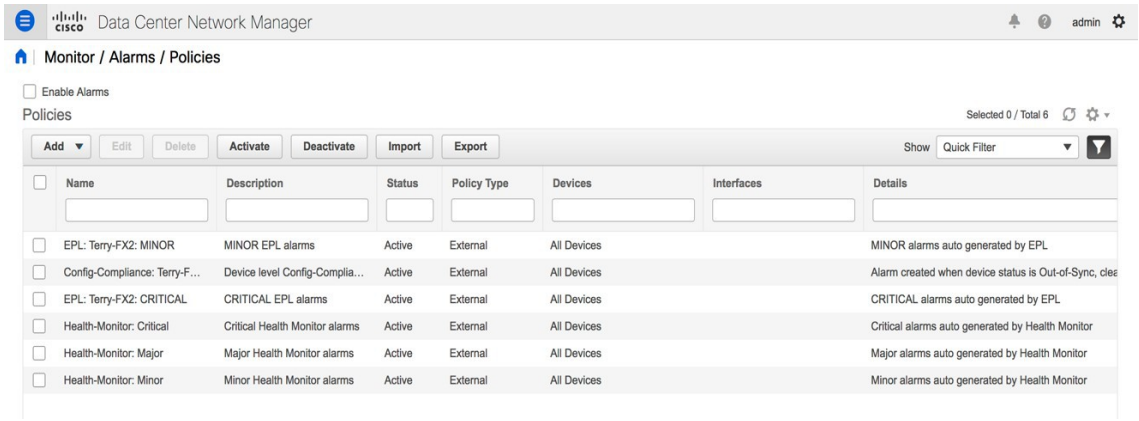

In case an alarm policy is deactivated using the DCNM Web UI, any alarms created or cleared for that policy will not be displayed in the **Monitor>Alarms>View** tab. To delete a policy, select the checkbox next to the policy and click **Delete**. However, we recommend not deleting a policy from the DCNM Web UI. When a fabric is deleted, the alarm policy along with all the active alarms for the devices in that fabric are deleted.

Endpoint Locator: Active Alarms

Choose **Monitor>Alarms>View** to display the active alarms.

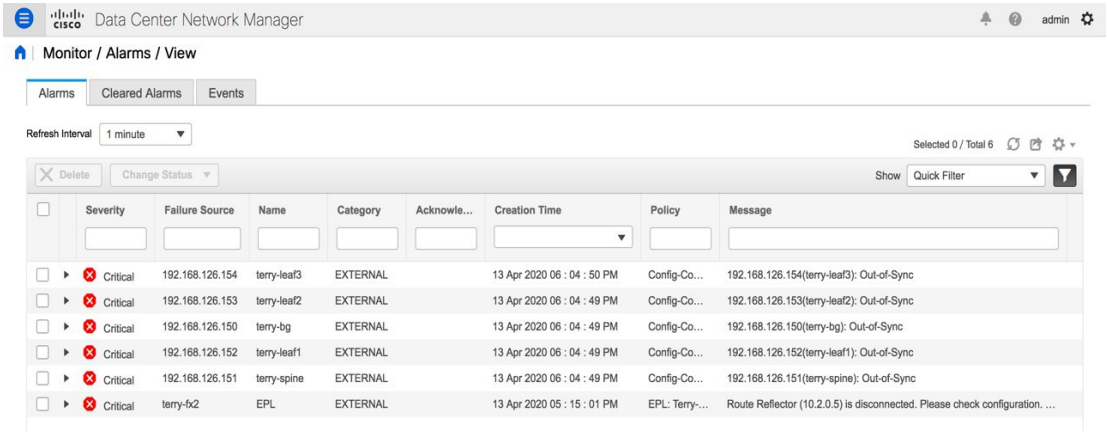

To clear active alarms, select the checkbox next to the alarm, click **Change Status** and select **Clear**.

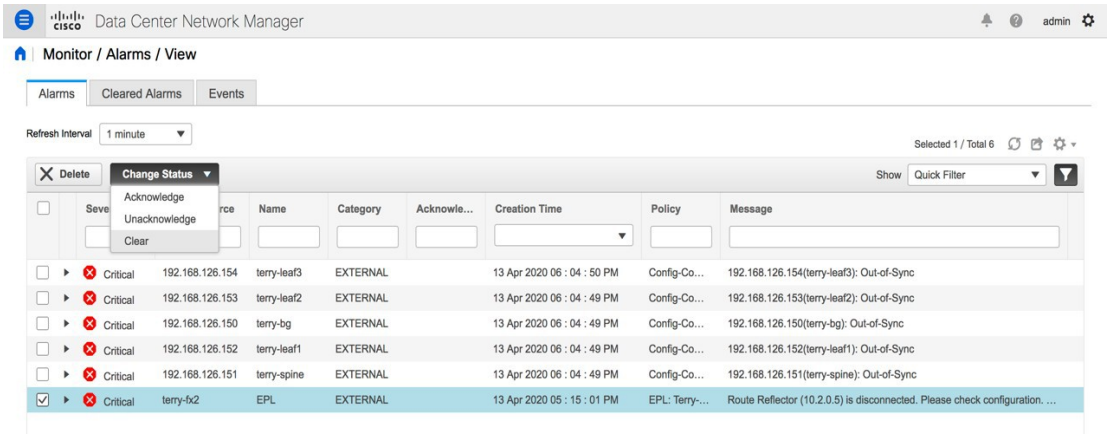

To delete active alarms, select the checkbox next to the alarm and click **Delete**.

Endpoint Locator: Cleared Alarms

To view the cleared alarms, select **Monitor>Alarms>View>Cleared Alarms**.

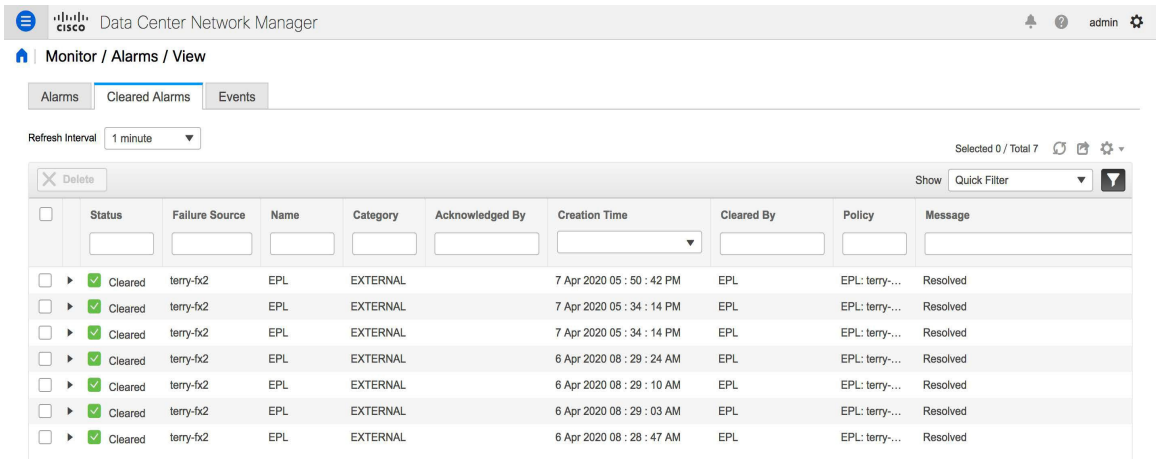

Click the arrow icon ▶ to display detailed information about the required alarm.

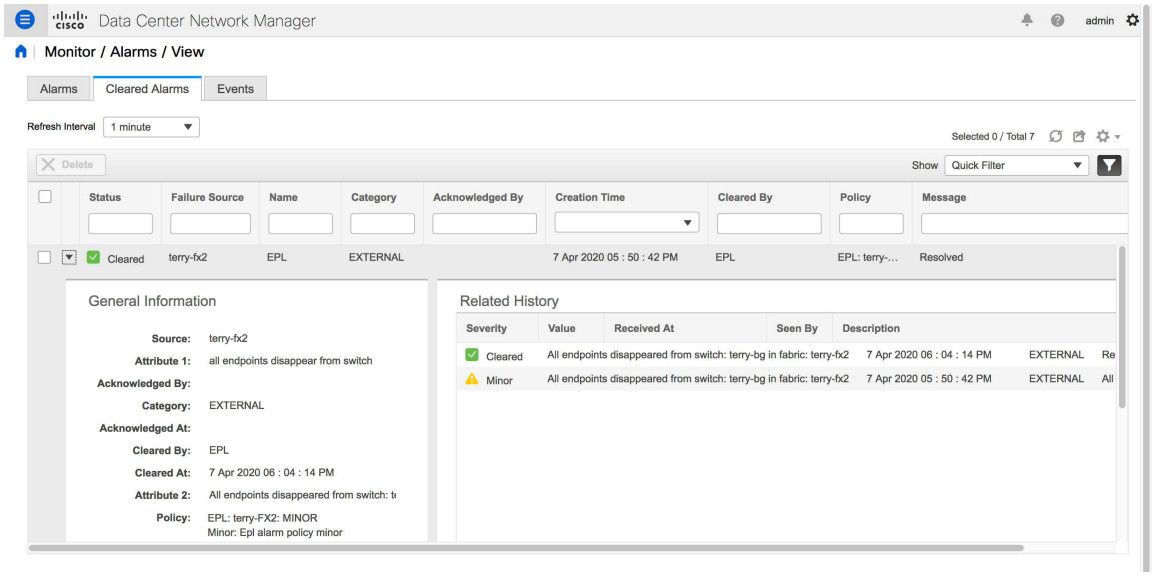

To delete a cleared alarm from the list of cleared alarms, select the checkbox next to the alarm and click **Delete**.

For more information on Alarms and Policies, refer Alarms.

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 >= 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 >= 80% and <90%

Alarms are raised and categorized as MINOR for the following events:

- CPU/Memory/Disk Utilization/ES JVM Heap Used Percentage >= 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.

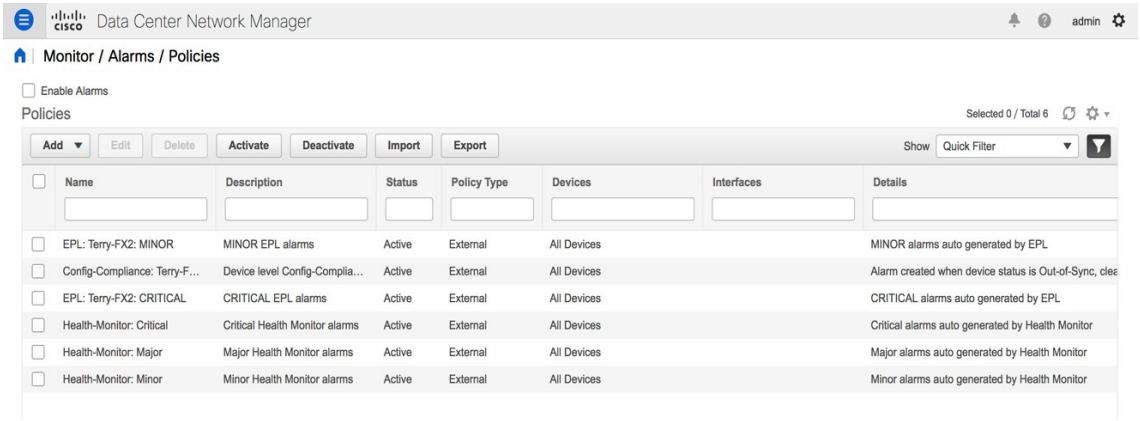

In case an alarm policy is deactivated using the GUI, any alarms created or cleared for that policy will not be displayed in the **Monitor>Alarms>View** tab. To delete a policy, select the checkbox next to the policy and click **Delete**. However, we recommend not deleting a policy from the GUI. When 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.

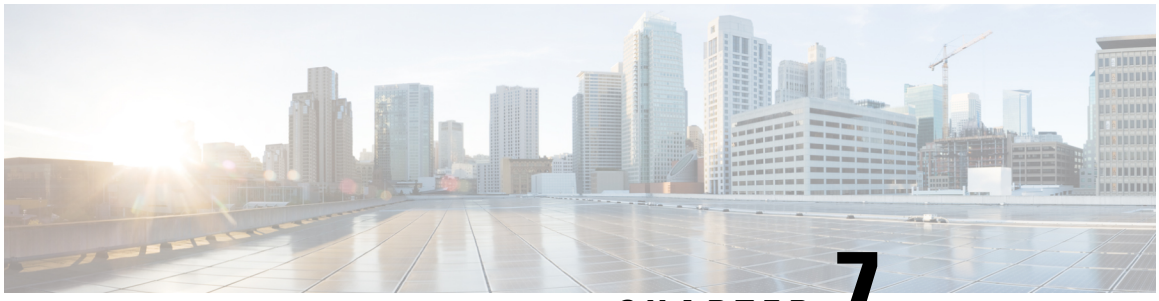

CHAPTER 7

Administration

This chapter contains the following topics:

- DCNM Server, on page 479
- Manage Licensing, on page 500
- Management Users, on page 515
- Performance Setup, on page 522
- Event Setup, on page 523
- Credentials Management, on page 527

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.

From Cisco DCNM Release 11.4(1), you can see the status of the following services as well:

Note The following services are available for OVA/ISO deployments only.

- NTPD server: NTPD service running on DCNM OVA, the IP address, and the port to which the service is bound.
- DHCP server: DHCP service running on DCNM OVA, the IP address, and the port to which the service is bound.
- SNMP traps
- Syslog Receiver

The DCNM servers for these services are as follows:

Service Name	DCNM Server
NTPD Server	0.0.0.0:123
DHCP Server	0.0.0.0:67
SNMP Traps	0.0.0.0:2162
Syslog Server	0.0.0.0:514

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.

Customization

From Cisco DCNM Release 11.3(1), you can modify the background image and message on the Web UI login page. This feature helps you to distinguish between the DCNM instances, when you have many instances running at the same time. You can also use a company-branded background on the login page. Click on Restore Defaults to reset the customizations to their original default values.

To remove the customizations and restore to the default values, click **Restore defaults**.

Login Image

This feature allows you to change the background image on the Cisco DCNM Web UI login page. If you have many instances of DCNM, this will help you identify the correct DCNM instance based on the background image.

To edit the default background image for your Cisco DCNM Web UI login page, perform the following steps:

1. Choose **Administration > DCNM Server > Customization**.
2. In the Login Image area, click **Add (+)** icon.

   Browse for the image that you need to upload from your local directory. You can choose any of the following format images: JPG, GIF, PNG, and SVG.
3. Select the image and click **Open**.

   A status message appears on the right-bottom corner.

```
Login image
Upload Successful
```

Note We recommend that you upload a scaled image for fast load times.

   The uploaded image is selected and applied as the background image.
4. To choose an existing image as login image, select the image and wait until you see the message on the right-bottom corner.
5. To revert to the default login image, click **Restore Defaults**.

Message of the day (MOTD)

This feature allows you to add a message to the Cisco DCNM Web UI login page. You can a list of messages that will rotate on the configured frequency. This feature allows you to convey important messages to the user on the login page.

To add or edit the message of the day on the Cisco DCNM Web UI login page, perform the following steps:

1. Choose **Administration > DCNM Server > Customization**.
2. In the **Message of the day (MOTD)** field, enter the message that must appear on the login page.
3. Click **Save**.

Default Fabric for Overlay Deployments

From Release 11.4(1), Cisco DCNM Customizations allows you to choose one of the valid Fabrics as default. This feature is available in the Cisco DCNM LAN Fabric deployment only.

To set a default fabric for all overlay deployments on the Cisco DCNM Web UI, perform the following steps:

Note Only a user with **network admin** role can use configure the default fabric.

1. Choose **Administration > DCNM Server > Customization**.
2. In the **Default Fabric for Overlay Deployments** drop-down list, select set a Fabric to set as a default for all the overlay deployments.
3. Click **Save** to set the fabric as default.

   A note appears in the right bottom of the window confirming that the default fabric is updated successfully.
4. To remove the default fabric, choose **--select as option** from the drop-down list and click **Save**.

Network Preferences

Earlier to Release 11.5(1), **appmgr update network-properties** command allows you to modify network properties. From Release 11.5(1), Cisco DCNM allows you to modify few network parameters from the Web UI. Modifying these overwrites the previously configured parameters.

Choose Cisco DCNM **Web UI > Admin > DCNM Server > Customization > Network Preferences** to modify the DNS, NTP, and the eth1/eth2 interfaces.

DNS

In the DNS field, enter the DNS IP address. You can also configure the DNS server using an IPv6 address. You can configure more than one DNS server. Use comma (,) as differentiator between the IP addresses.

Note If you’re using Network Insights applications, ensure that the DNS server is valid and reachable.

NTP

In the NTP field, enter the IP address of the NTP server. The value must be an IP or IPv6 address or RFC 1123 compliant name.

Routes

In-Band (eth2)

In the In-Band Network area, enter the IPv4 address and Gateway IPv4 Address for the in-band network. If DCNM is on the IPv6 network, configure the network by entering relevant IPv6 Address for IPv6 address and Gateway IPv6 Address.

Note When a Nexus Dashboard server is adding a Site from DCNM 11.5(1), it must reach the DCNM server over the Data Network. DCNM Data Network connectivity is defined to be over eth2 interface of the DCNM server; also known as Inband Connectivity interface in DCNM. When the eth2 connectivity of the DCNM with the Data Network Connectivity of the Nexus Dashboard is spanning multiple subnets, that is, when they are Layer3 Route connected, you must add routes in DCNM before adding the Site on ND. Enter the Routes to the ND Data Network over the In-band(eth2) inputs of the dashlet.

The In-Band Network provides reachability to the devices via the front-panel ports.

Out-of-Band (eth1)

In the Out-of-Band Network area, enter the IPv4 address and Gateway IPv4 Address. If DCNM is on the IPv6 network, configure the network by entering relevant IPv6 Address for IPv6 address and Gateway IPv6 Address.

Out-of-band management provides a connection to the device management ports (Typically mgmt0).

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. You can use the use **appmgr tech_support** command in the CLI to generate the techsupport file.

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 51st 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.

Native HA

Before you begin

Note Ensure that you clear your browser cache and cookies everytime after a Federation switchover or failover.

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/ pgsql-standby-backup.tgz. If the file exists, restore database data files:

```
rm -rf    data/*
tar -zxf    replication/ pgsql-standby-backup.tgz    data
/etc/init.d/postgresql-9.4 start
ps -ef | grep post
```

The active DCNM host will synchronize the two databases.

The TFTP server is not bound to the eth1 VIP address on the active host: The TFTP server should run on the active host (not on the standby host), and it should be bound to the eth1 VIP address. In some setups, the bind address is not the VIP address, as per the TFTP configuration file, and this could cause issues when switches try to use TFTP.

- Enter “grep bind /etc/xinetd.d/tftp” to check if the TFTP configuration file has the right bind address. If the displayed IP address is not the eth1 VIP address, then change the bind address to the VIP address. Repeat the procedure for the standby host. Update the bind address to the VIP address.
- Enter " " /etc/init.d/xinetd restart” on the active host to restart TFTP.

Note The TFTP server can be started or stopped with the “appmgr start/stop ha-apps” command.

Multi Site Manager

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 f Windows/RHEL installation type.

2. 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.

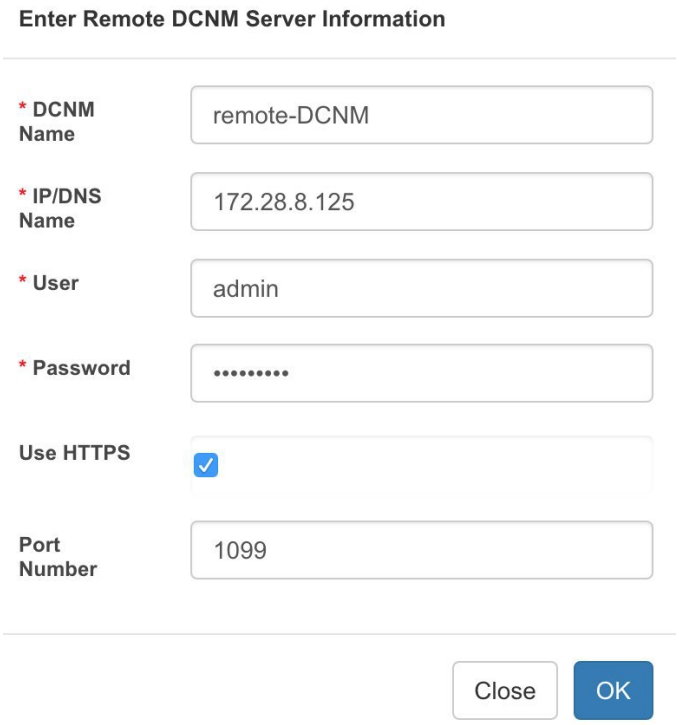

3. Click **OK**. After validation, the remote DCNM server is represented in the screen, next to the local DCNM server.

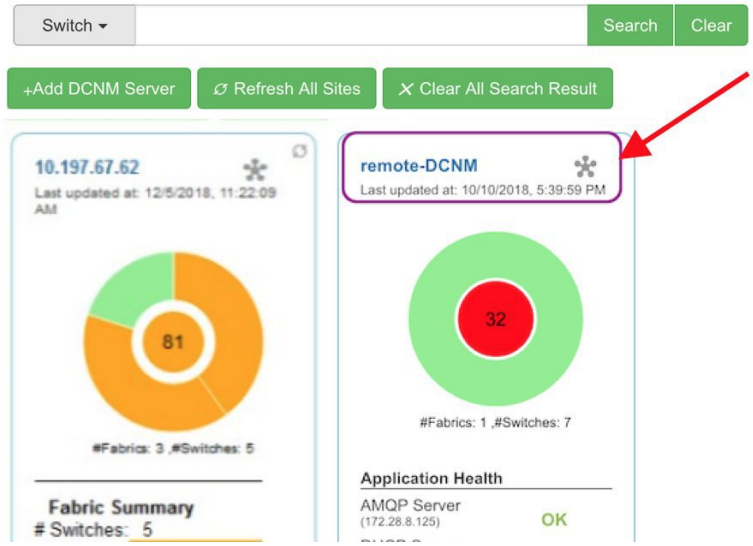

You can click **Refresh All Sites** to display updated information.

Retrieve Switch Information

1. Choose **Administration > DCNM Server > Multi Site Manager**. The Multi Site Manager screen comes up

2. 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.

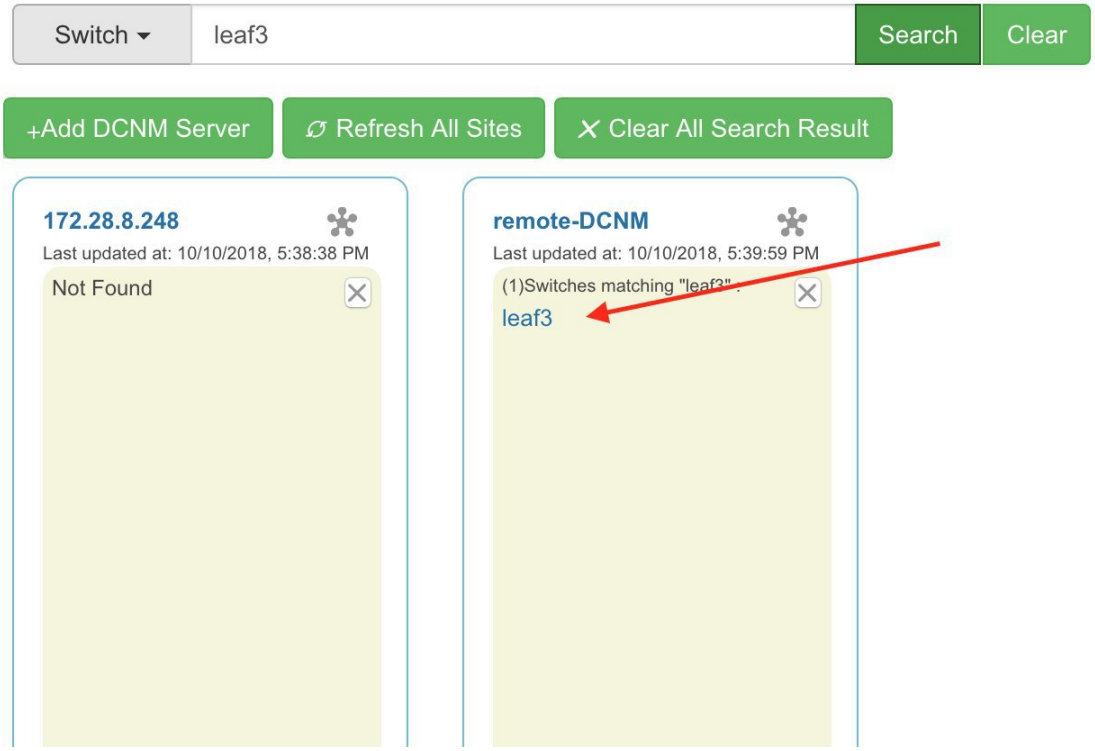

3. Click **leaf3** to view detailed switch information in an adjacent browser tab.

   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.

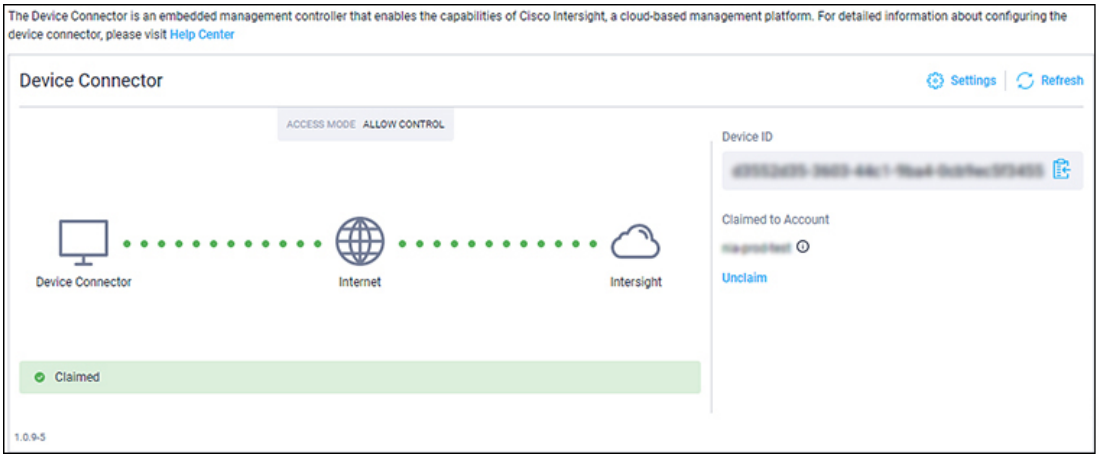

2. Click **Settings**.

   The **Settings - General** window appears.

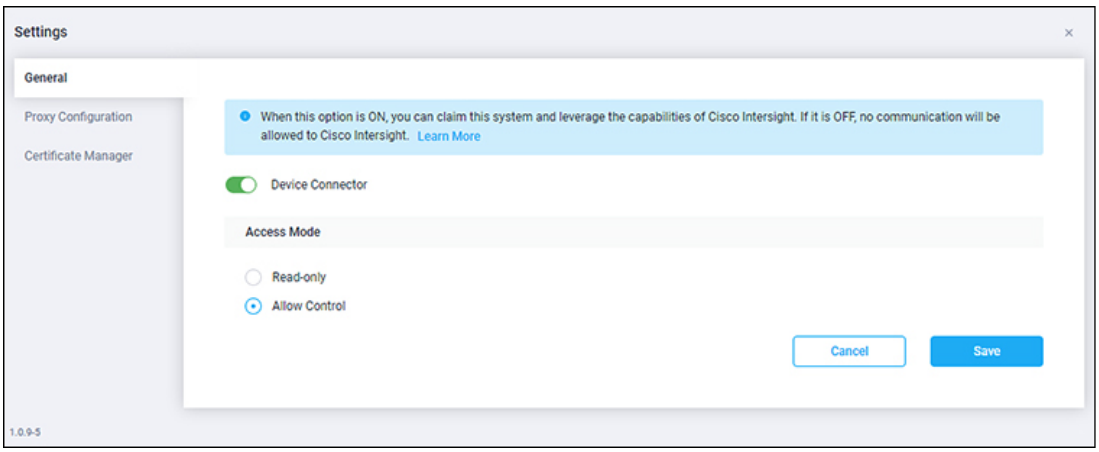

- **Device Connector (switch)**

  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.

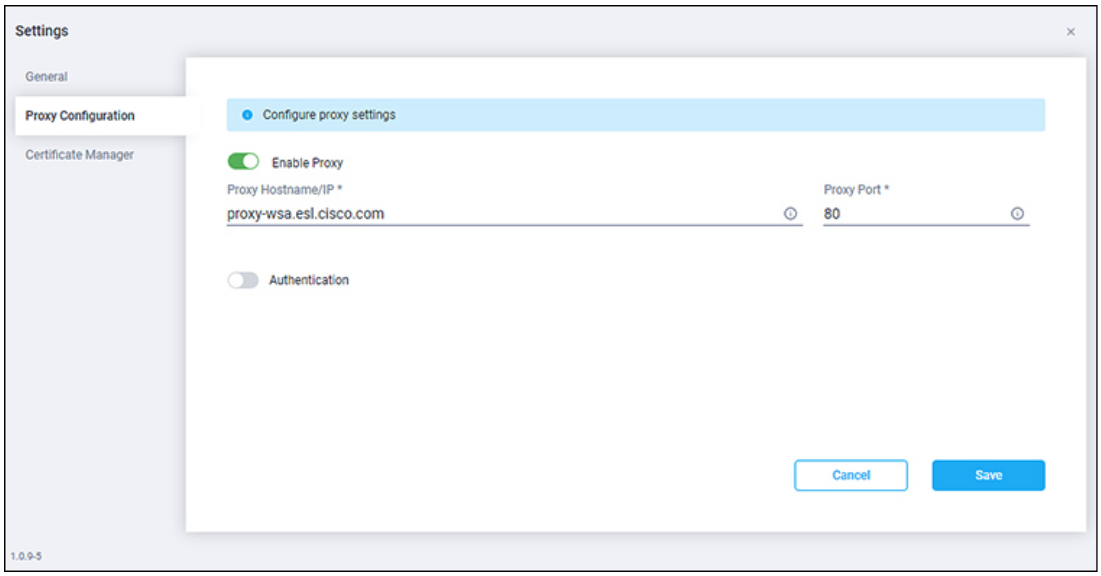

   - **Enable Proxy (switch)**

     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**.

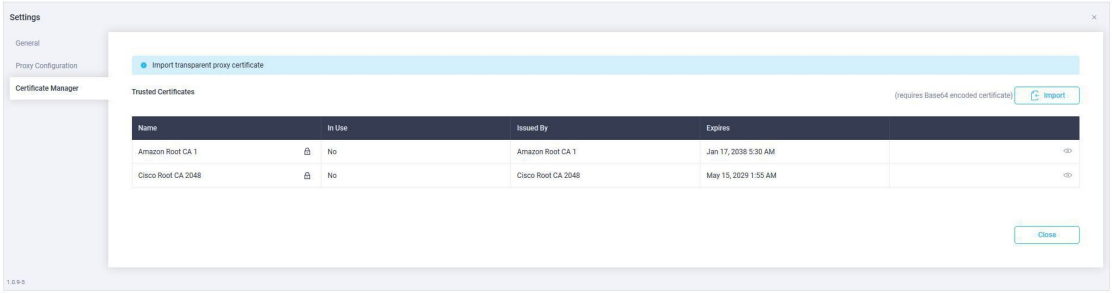

The trusted certificates appear in the table.

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.

NX-API Certificate Management for Switches

Cisco NX-OS switches require an SSL certificate to function in NX-API HTTPS mode. You can generate the SSL certificates and get it signed by your CA. You can install the certificates manually using CLI commands on switch console.

From Release 11.4(1), Cisco DCNM provides a Web UI framework to upload NX-API certificates to DCNM. Later, you can install the certificates on the switches that are managed by DCNM.

This feature is supported only on Cisco DCNM OVA/ISO deployments.

Note This feature is supported on switches running on Cisco NXOS version 9.2(3) or higher.

For each switch, the data center administrator generates an ASCII (base64) encoded certificate. This certificate comprises two files:

- `.key` file that contains the private key
- `.crt/.cer/.pem` file that contains the certificate

Cisco DCNM also supports a single certificate file that contains an embedded key file, that is, `.crt/.cer/.pem` file can also contain the contents of .key file.

DCNM doesn't support binary encoded certificates, that is, the certificates with `.der` extension are not supported. You can protect the key file with a password for encryption. Cisco DCNM does not mandate encryption; however, as this is stored on DCNM, we recommend that you encrypt the key file. DCNM supports AES encryption.

You can either choose CA-signed certificates or self-signed certificates. Cisco DCNM does not mandate the signing; however, the security guidelines suggest you use CA-signed certificates.

You can generate multiple certificates meant for multiple switches, to upload to DCNM. Ensure that you name the certificates appropriately, to help you choose the switch meant for that certificate.

You can upload one certificate and corresponding key file, or bulk upload multiple certificates and key files. After the upload is complete, you can view the upload list before installing these on the switches. If a certificate file that contains an embedded key file is uploaded, DCNM derives the key automatically.

Certificate and the key file must have the same filename. For example, if a certificate filename is `mycert.pem`, the key filename must be `mycert.key`. If the certificate and key pair filenames are not the same, then DCNM will not be able to install the certificate on the switch.

Cisco DCNM allows you to bulk install the certificates to the switches. Because bulk installation uses the same password, all encrypted keys must be encrypted with the same password. If the password is different for a key, you cannot install the certificate in bulk mode. Bulk mode installation allows you to install encrypted and unencrypted keys certificates together, but all encrypted keys must have the same password.

When you install a new certificate on the switch, it replaces the existing certificate and replaces it with the new certificate.

You can install the same certificate on multiple switches; however, you cannot use the bulk upload feature.

Note DCNM doesn't enforce the validity of certificates or options provided in it. It is up to you and the requirements on the switch to follow the convention. For example, if a certificate is generated for Switch-1 but it is installed on Switch-2, DCNM doesn't enforce it; switches may choose to accept or reject a certificate based on the parameters in the certificate.

On Cisco DCNM **Web UI > Administration > DCNM Server > NX API Certificates**, the following tables are displayed:

- **Certificate Installation Status table**: Displays the status of certificates last installed on the switches. It also displays the time when the certificates were updated previously.
- **Certificates Uploaded to DCNM table**: Displays the certificates uploaded on DCNM and any switch association.

  However, refer to the Certificate Installation Status table to see the certificate and switch association. Upload table is only meant for uploading certificates on DCNM and installing on the switches.

You can also watch the video that demonstrates how to use Switch NX-API SSL Certificate Management feature. See Video: Switch NX-API SSL Certificate Management.

Uploading the certificates on DCNM

To upload the certificates onto DCNM using the Cisco DCNM Web Client UI, perform the following steps:

Procedure

Step 1 Choose **Administration > DCNM Server > NX API Certificates**.

Step 2 In the **Certificates Uploaded to DCNM** area, click **Upload Certificates** to upload the appropriate license file.

Step 3 Browse your local directory and choose the certificate key pair that you must upload to DCNM.

You can choose certificates with extension .cer/.crt/.pem + .key file separately.

Cisco DCNM also allows you to upload a single certificate file that contains an embedded key file. The key file is automatically derived after upload.

Step 4 Click **Open** to upload the selected files to DCNM.

A successful upload message appears. The uploaded certificates are listed in the **Certificates Uploaded to DCNM** area.

In the **Certificate Installation Status** area, the certificate appears, with Status as **UPLOADED**.

If the certificate is uploaded without the key file, the status shows **KEY_MISSING**.

Installing Certificates on Switches

To install certificates on the switches using Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Administration > DCNM Server > NX API Certificates**.

Step 2 In the **Certificate Installation Status** area, for each certificate, click on the **Switch** column.

Step 3 From the drop-down list, select the switch to associate with the certificate.

Click **Save**.

Step 4 Select the certificate that you need to install and click **Install Certificates on Switch**.

You can select multiple certificates to perform a bulk install.

Step 5 In the **Bulk Certificate Install** window, upload the certificates to DCNM. Perform the following steps:

You can install a maximum of 20 certificates at the same instance, using the Bulk Install feature.

a) Choose the file transfer protocol to upload the certificate to DCNM.

You can choose either SCP or SFTP protocol to upload the certificates.

b) Check the VRF checkbox for the certificates to support the VRF configuration.

Enter the VRF name that the switch uses to reach DCNM. Generally, DCNM is reached via management VRF of switches, but it can be any VRF that is configured on the switch that is used to reach DCNM.

c) In the NX-API Certificate Credentials, enter the password which was used to encrypt the key while generating the certificates.

Leave this field empty, if the key uploaded along with the certificate is not encrypted.

Note that you can install unencrypted and encrypted keys and a certificate in a single bulk install; however, you must provide the key password used for encrypted keys.

d) Click **Install**.

A notification message appears to confirm if the certificate was successfully installed on the specific switch.

In the Certificate Installation Status area, the Status of certificate now shows **INSTALLED**.

Unlinking and Deleting certificates

After the certificates are installed on the switch, DCNM cannot uninstall the certificate from DCNM. However, you can always install a new certificate on the switch. The certificates that are not installed on the switches can be deleted. To delete the certificate installed on the switch, you must unlink the certificate from the switch, and then delete it from DCNM.

Note Unlinking the certificate from the switch does not delete the certificate on the switch. The certificate still exists on the switch. Cisco DCNM cannot delete the certificate on the Switch.

To delete certificates from DCNM repository, using the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Administration > DCNM Server > NX API Certificates**.

Step 2 In the **Certificate Installation Status** area, select the certificate(s) that you need to delete.

Step 3 Click **Clear** Certificates.

A confirmation message appears.

Step 4 Click **OK** to clear the selected certificates.

The status column shows UPLOADED. The Switch column shows NOT_INSTALLED.

Step 5 Select the certificate and click **Clear Certificates**.

The Certificate is removed from the Certificate Installation Status table.

Step 6 In the Certificates Uploaded to DCNM area, select the certificate that is now unlinked from the Switch.

Click **Delete Certificates**.

The certificate is deleted from DCNM.

Troubleshooting NX API Certificate Management

While installing a certificate, you can encounter errors. The following sections provide information about troubleshooting the NX-API Certificate Management for switches.

COPY_INSTALL_ERROR

Problem Statement: Error message COPY_INSTALL_ERROR

Reason Cisco DCNM cannot reach the switch.

Solution:

- Verify if the switch is reachable from Cisco DCNM. You can perform an SSH login and ping the switch to verify.
- Switch connects to DCNM through it's management interface. Verify if you can ping DCNM from the Switch console. If the switch requires VRF, very if the correct vrf is provided.
- If the certificate private key is encrypted, ensure that you provide the correct password.
- Verify is the correct key file is uploaded with the certificate. Ensure that the certificate file and the key file have the same filename.

CERT_KEY_NOT_FOUND

Problem Statement: Error message CERT_KEY_NOT_FOUND

Reason: Key file was not uploaded while uploading the certificate (.cer, .crt, .pem).

Solution:

- Ensure that the certificate (.cer, .crt, or .pem) file and its corresponding .key file has the same filename

  For example: If the certificate file name is mycert.crt, the key file must be mycert.key.
- DCNM identifies key file with certificate file name, and therefore, it is necessary to have the key file with same filename.
- Upload the certificate and key file with same filename, and install the certificate.

Backing up DCNM

From Cisco DCNM, Release 11.5(1), you can trigger scheduled DCNM backups from the Cisco DCNM Web UI. When you trigger a backup from the Web UI, the **appmgr backup** command is run. You can see the following information under the **Server Backup Jobs** tab in the **Backup** window.

Table 21: Server Backup Jobs Tab

Parameters	Description
Node	Specifies if the backup is active or standby. For standalone nodes, it will appear as a localpath. **Note** For HA cluster, one active node and one standby node is created. However, you can choose only the active node for an HA cluster.
Schedule	Specifies when the scheduled backup is triggered.
Local Path	Specifies the local path, where the backup is stored.
Remote Destination	Specifies the username, host IP, and the remote destination, where the backup is stored. It is empty if you do not save the backup in a remote location. **Note** A copy of the backup is also stored in the local path.
Log Path	Specifies the path where the log entries are stored. You can use this information to troubleshoot any issues.
Saved Backups	Specifies the number of versions of a backup. The default value is 5.

You can perform the following actions in the **Backup** window:

Creating a Backup

To create a backup from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Administration > DCNM Server > Backup**.

The **Backup** window appears, which has all the information under the **Server Backup Schedules** area.

Step 2 Click **Add**.

The **Create Backup Schedule** dialog box appears.

Step 3 Choose the time using the **Start At** drop-down list under the **Schedule** area.

Step 4 Choose the frequency of the backup.

The valid options are:

- **Daily**: Select this radio button if you want to trigger the backup everyday.

- **Weekly**: Select this radio button if you want to trigger the backup once a week. If you select this radio button, you get options to choose the day.

Step 5 Enter the number of backups you want to save in the **Max # of Saved Backups** field under the **Destination** area.

You can save upto 10 backups and the default value is 5.

Step 6 (Optional) Check the **Remote Destination** check box to save the backup in a remote location.

The following fields will be available after you check the **Remote Destination** check box.

Fields	Descriptions
User	Enter the username.
Password	Enter the password. **Note** You don't have to enter the password if you have enabled the key-less configuration between your DCNM and the remote host.
Host IP	Enter the host IP address which is connected to your DCNM.
Path	Enter the remote destination path where you want to save the backup.

Note
- The backup files are huge, with the size in gigabytes.
- A copy of the backup will always be saved in the local destination as well.

Step 7 Click **Create**.

The **Backup** window is populated even when you run the **appmgr backup** command using the CLI. You can also view the backups, which you scheduled from the Web UI, in the CLI using the **appmgr backup schedule show** command.

Modifying a Backup

To modify a backup from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Administration > DCNM Server > Backup**.

The **Backup** window appears, which has all the information under the **Server Backup Schedules** area.

Step 2 Click **Modify**.

The **Modify Backup Schedule** dialog box appears.

Step 3 Make the necessary changes.

Step 4 Click **Modify**.

Deleting a Backup

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

Procedure

Step 1 Choose **Administration > DCNM Server > Backup**.

The **Backup** window appears, which has all the information under the **Server Backup Schedules** area.

Step 2 Click **Delete**.

The confirmation dialog box appears.

Step 3 Click **Yes**.

Note If you run the **appmgr backup schedule none** command in the CLI, the backup is deleted. You can verify if the backup is deleted by refreshing the **Backup** window.

Job Execution Details

You can see the following information under the **Job Execution Details** tab in the **Backup** window.

Table 22: Server Backup Schedules Area

Parameters	Description
Node	Specifies if the node is active or standby. For standalone nodes, it will appear as a local node.
Backup File	Specifies the path, where the backup is stored.
Start Time	Specifies the time when the backup process started.
End Time	Specifies the time when the backup process ended.
Log File	Specifies the path where the log entries are stored. You can use this information to troubleshoot any issues.
Status	Specifies if the backup was a success or failed.
Error Message	Specifies error messages, if any, that appeared during the backup.

Manage Licensing

The Manage Licensing menu includes the following submenus:

Managing Licenses

You can view the existing Cisco DCNM licenses by choosing **Administration > Manage Licensing > DCNM**. 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. The total number of licenses for new installations are 50. However, the total number of licenses continues to be 500 for inline upgrade.
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.

Field	Description
License State	Displays the license state of the switch that can be one of the following: • Permanent • Eval • Unlicensed • Not Applicable • Expired • Invalid • Smart
License Type	Displays the license type of the switch that can be one of the following: • DCNM-Server • Switch • Smart • Honor • Switch-Smart
Expiration Date	Displays the expiry date of the license. **Note** Text under the **Expiration Date** 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**

To assign license to switches through POAP, refer to DCNM Licensing Guide.

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.

Honor License Mode

From Release 11.3(1), Cisco DCNM Eval license validity is extended from 30 days to 60 days. That implies, after 60 days. Every license has an expiry date attached to it. After the license expires, Cisco DCNM allows you to use all the licensed features. Switches remain in honor mode until the switch is licensed again or the user manually removes the license.

If there are switches in the Honor License mode, an error message appears after you logon to DCNM.

```
************************************
*Your licenses are out of compliance.
Your inventory contains switches that are unlicensed for DCNM Operation*

************************************
```

Go to **Administration > Manage Licensing > DCNM**, In the **Switches/VDCs** table, select the switch and click **Assign License** to renew the license.

Guidelines

- Switches that don’t have a license assigned to them is considered unlicensed. Unlicensed Switches aren’t allowed to use Licensed DCNM features.
- If a switch has an expired EVAL license, it will change from EVAL to Honor mode and the license features continues to be operational.
- You can’t assign expired EVAL licenses to the switches.
- Switches with switch-based honor license can’t be overwritten with any server-based license.
- When a license is assigned to a discovered switch and a valid license isn’t available, then an honor-based license with expiration date will be assigned to the switch.

Nag events for Honor-mode licenses

For every license in honor mode, an event is generated every seven days. A nag event informs the user "DCNM-SAN file license is in honor mode, need to assign/purchase a new license for this switch." Or "DCNM-LAN file license is in honor mode, need to assign/purchase a new license for this switch."

Additional popup notification appears when you logon to Cisco DCNM, to inform that “DCNM-SAN file license is in honor mode, need to assign/purchase a new license for this switch.”

Server-based honor license support

On the DCNM **Web UI > Administration > Manage Licensing > DCNM**, the **Licensed State** column displays **Honor** and **Expiration Date** column displays the date, time, and when the license expired and changed to the Honor mode.

Switches will remain in honor mode after reboot also. To change the license from honor mode, you must manually unassign the license or assign a new valid license to the switch.

The following image shows license page with a SAN switch in Honor mode.

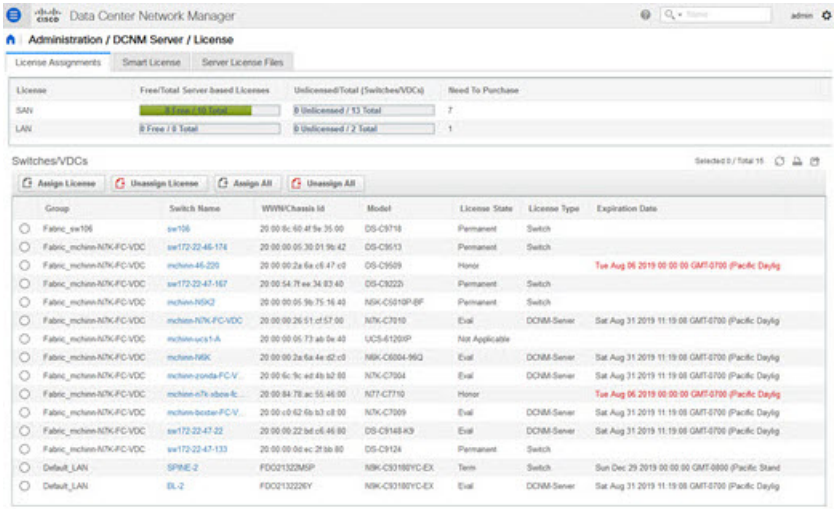

The following image shows license page with a LAN switch in Honor mode.

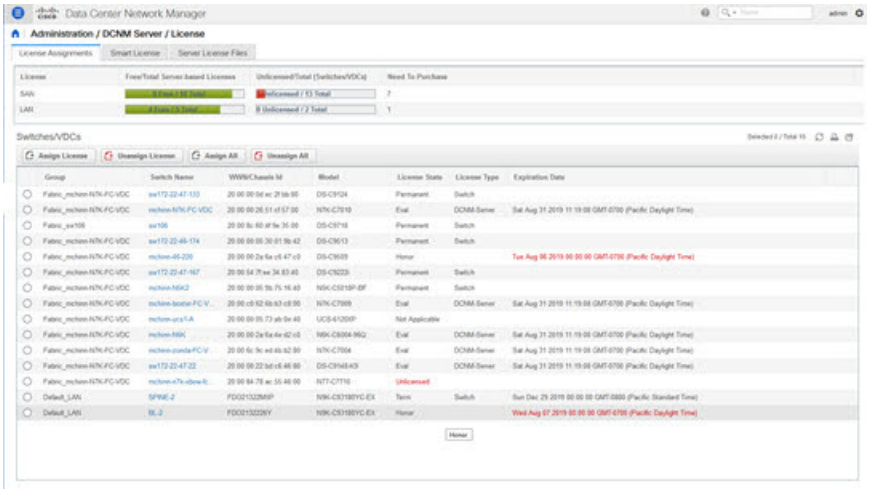

The following image shows the switch table displaying the honor mode of license and term.

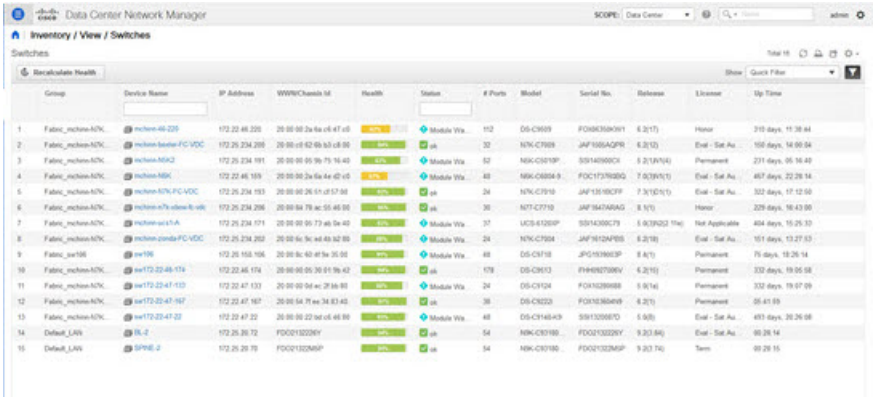

The following image shows Switch Dashboard with a LAN switch in Honor mode license.

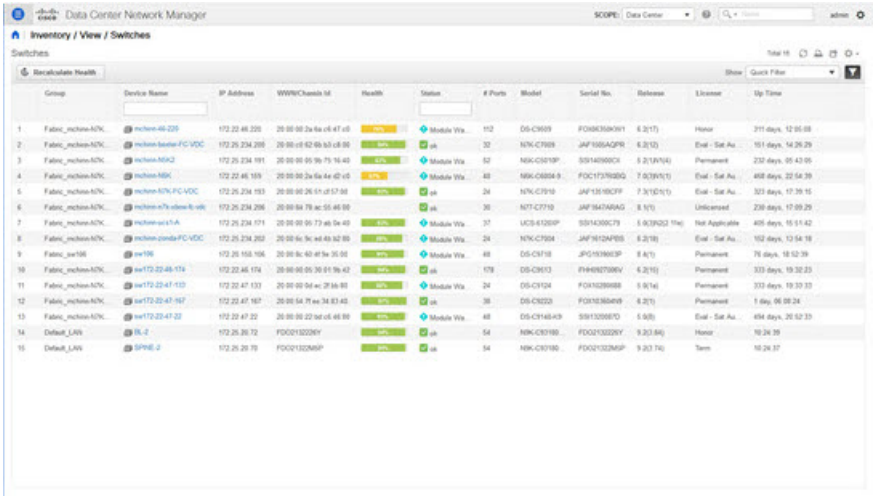

The following image shows Switch Dashboard with a SAN switch in Honor mode license.

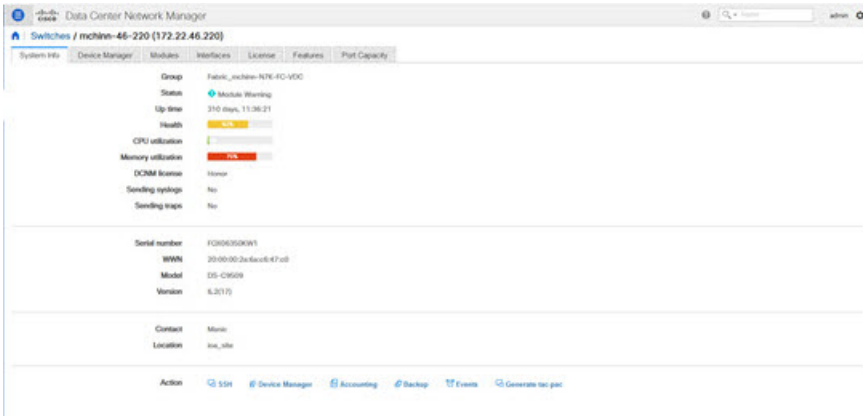

The following image shows the SAN Client License Agreement tab.

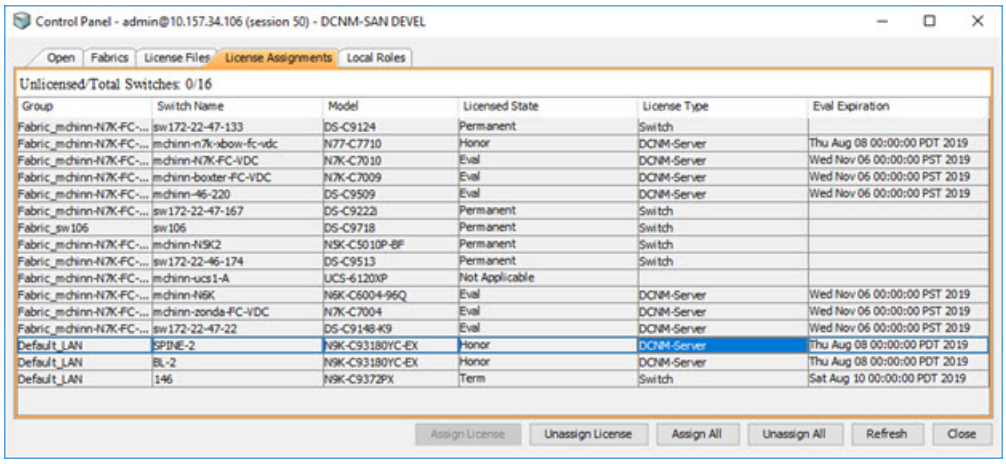

The following image shows the **SAN Client License** files tab.

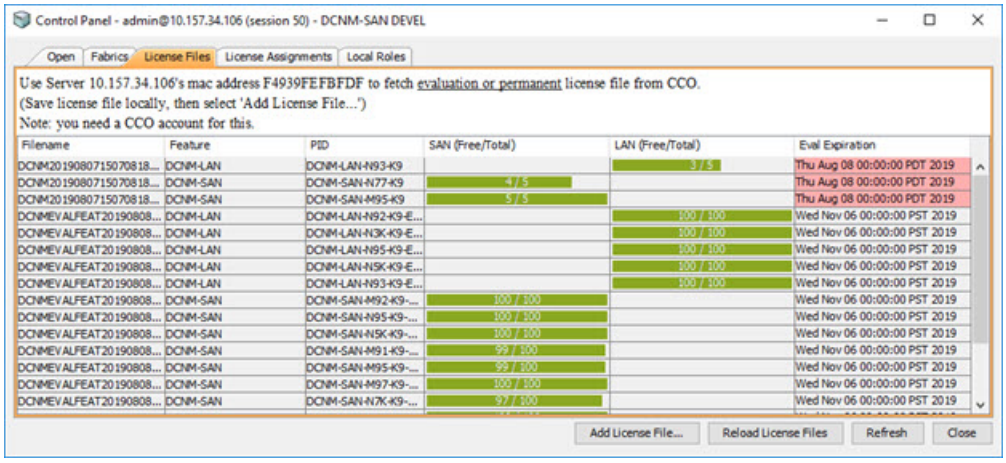

Note Switch-based honor licenses can’t be overwritten with server-based license files.

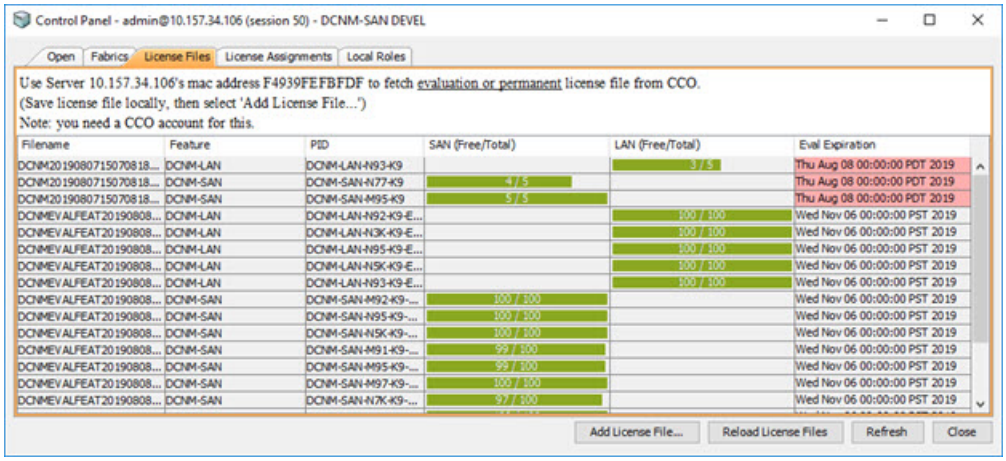

Smart License

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 > Manage Licensing > DCNM > Smart License**. You will see a brief introduction on Cisco smart licensing, a menu bar, and the **Switch Licenses** area.

Introduction to Smart Licensing

Cisco Smart Licensing is a flexible licensing model that provides you with an easier, faster, and more consistent way to purchase and manage software across the Cisco portfolio and across your organization. And it's secure – you control what users can access. With Smart Licensing you get:

- **Easy Activation**: Smart Licensing establishes a pool of software licenses that can be used across the entire organization—no more PAKs (Product Activation Keys).
- **Unified Management**: My Cisco Entitlements (MCE) provides a complete view into all of your Cisco products and services in an easy-to-use portal, so you always know what you have and what you are using.
- **License Flexibility**: Your software is not node-locked to your hardware, so you can easily use and transfer licenses as needed.

To use Smart Licensing, you must first set up a Smart Account on Cisco Software Central (https://software.cisco.com/software/csws/ws/platform/home).

For a more detailed overview on Cisco Licensing, go to https://www.cisco.com/c/en/us/buy/licensing/licensing-guide.html.

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.

Field	Description
Status	Specifies the status of the licenses used. Valid values are **Authorized** and **Out of Compliance**.
Description	Specifies the type and details of the license.
Last Updated	Specifies the timestamp when switch licenses were last updated.
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 > Manage Licensing > DCNM > 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 > Manage Licensing > DCNM > 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/oddce/services/DDCEService

- **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 > Manage Licensing > DCNM > 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 > Manage Licensing > DCNM > 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.

Switch Smart License

If the switch is pre-configured with a smart license, DCNM validates and assigns a switch smart license. To assign licenses to switch using the Cisco DCNM UI, choose **Administration > Manage Licensing >Assign License** or, **AssignAll**.

Note For switches in managed mode, switch smart license must be assigned through DCNM.

Note From Cisco NX-OS Release 9.3(6), switch smart license is supported.

To enable switch smart license on DCNM:

- Enable smart license feature on the switch, using freeform CLI configuration.
- Configure smart licensing on the switch, using **feature license smart** or **license smart enable** command on the switch.

- Push token of your device to smart account using license smart register **idtoken** command. Use **EXEC** option in DCNM to push token. For more details, refer to Running EXEC Mode Commands in DCNM.

For unlicensed switches, licenses are assigned based on this priority:

1. DCNM Smart License
2. DCNM Server License
3. DCNM Eval License

Server License Files

From Cisco DCNM Web UI, choose **Administration > Manage Licensing > DCNM > 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. **Note** Text in the **Expiration Date** 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 > Manage Licensing > DCNM** 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.

Switch Features—Bulk Install

From Release 11.3(1), Cisco DCNM allows you to upload multiple licenses at a single instance. DCNM parses the license files and extract the switch serial numbers. It maps the serial numbers in the license files with the discovered fabric to install the licenses on each switch. License files are moved to bootflash and installed.

To bulk install licenses to the switches on the Cisco DCNM Web Client UI, perform the following steps:

1. Choose **Administration > Manage Licensing > Switch features**.
2. In the Switch Licenses area, click **Upload License files** to upload the appropriate license file.

   The Bulk Switch License Install window appears.
3. In the Select file, click **Select License file(s)**.

   Navigate and choose the appropriate license file located in your local directory.

   Click **Open**.
4. Choose the file transfer protocol to copy the license file from the DCNM server to the switch.
   - Choose either **TFTP**, **SCP**, or **SFTP** protocol to upload the license file.

Note Not all protocols are supported for all platforms. TFTP is supported for Win/RHEL DCNM SAN installation only. However, SFTP/SCP supported for all installation types.

5. Check the **VRF** check box for the licenses to support VRF configuration.

   Enter the VRF name of one of their defined routes.
6. Check the **Overwrite file on Switch** checkbox, to overwrite the license file with the new uploaded license file.

Note The overwrite command copies the new file over the existing one in boot flash. If the previous license already installed, it won't override the installation.

7. In the DCNM Server credentials, enter the root username and password for the DCNM server.

   Enter the authentication credentials for access to DCNM. For DCNM Linux deployment, this is the username. For OVA\ISO deployments, use the credentials of the **sysadmin** user.
8. Click **Upload**.

   The License file is uploaded to the DCNM. The following information is extracted from the license file.

- Switch IP – IP Address of the switch to which this license is assigned.
- License File – filename of the license file
- Features List –list of features supported by the license file

9. Select the set of licenses that you want to upload and install on their respective switches. A license file is applicable for a single specific switch.
10. Click **Install Licenses**.

    The selected licenses are uploaded and installed on their respective switches. Status messages, including any issues or errors are updated for each file as it completes.
11. After the license matches with respective devices and installs, the **License Status** table displays the status.

Switch-based honor license support

On the DCNM **Web UI > Inventory > Switch > License**, the **Type** column displays "Unlicensed Honor License" and **Warnings** column displays **Honor started: …** with elapsed time since the license was changed to the Honor mode.

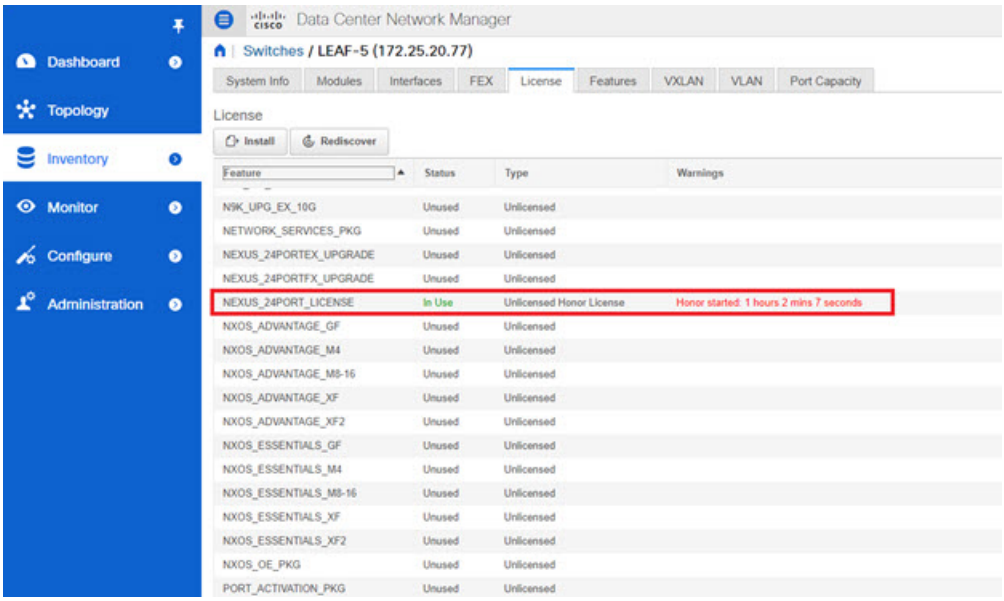

Note Switch-based honor licenses can't be overwritten with server-based license files.

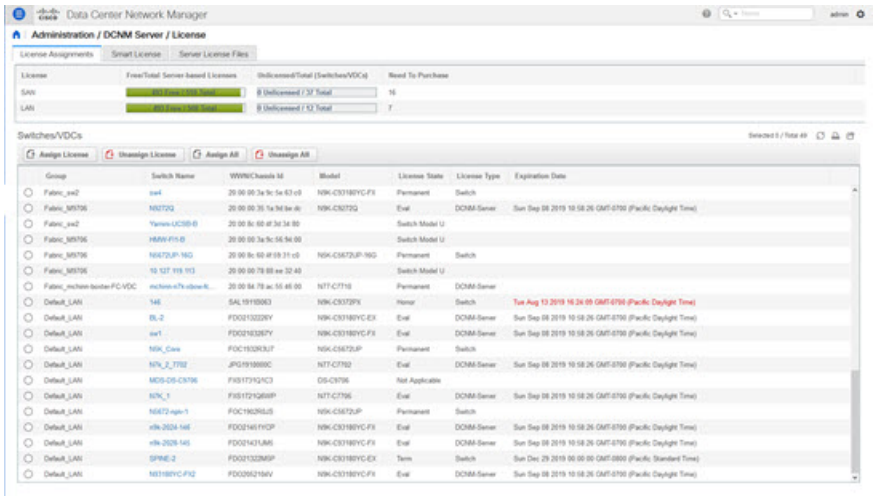

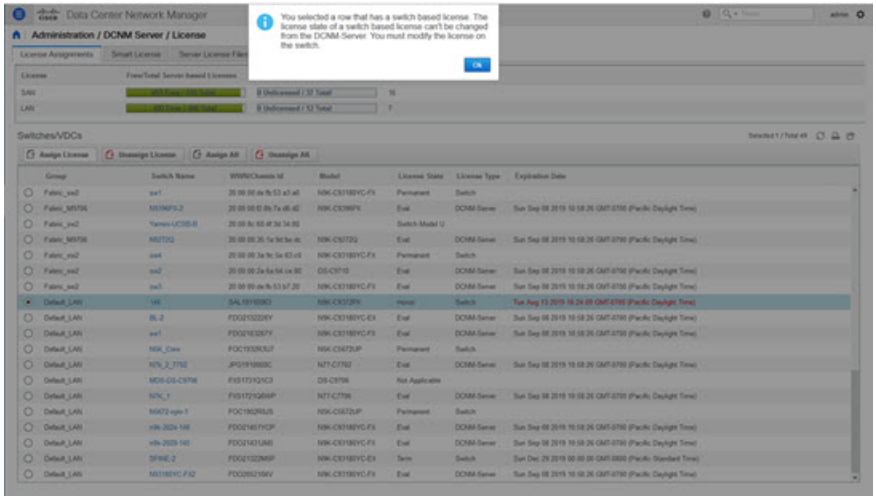

Application Licenses

From Release 11.3(1), you can manage licenses for applications on the Cisco DCNM. Choose **Web UI > Administration > Manage Licensing > Applications** to view the Application Licenses.

The Application Licenses tab displays the DCNM Applications with a summary of their unlicensed/total switches and if they are out of compliance. The PID Per Application Usage table displays the actual counts per PID given to the server from the Application Framework. The PIDs that need to be purchased for each application is also listed.

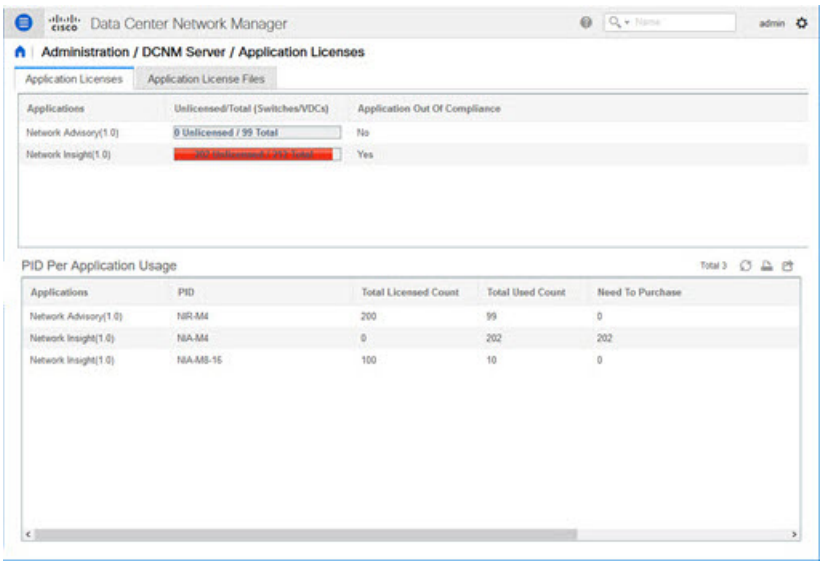

The Application License Files tab allows you to add license files for the applications. Click on Add license file to add license file from your local directory. The license filename, application name, PID, device count and expiration date details are extracted from the imported license file. If the license isn't permanent or is eval or term, the expiration date is also listed.

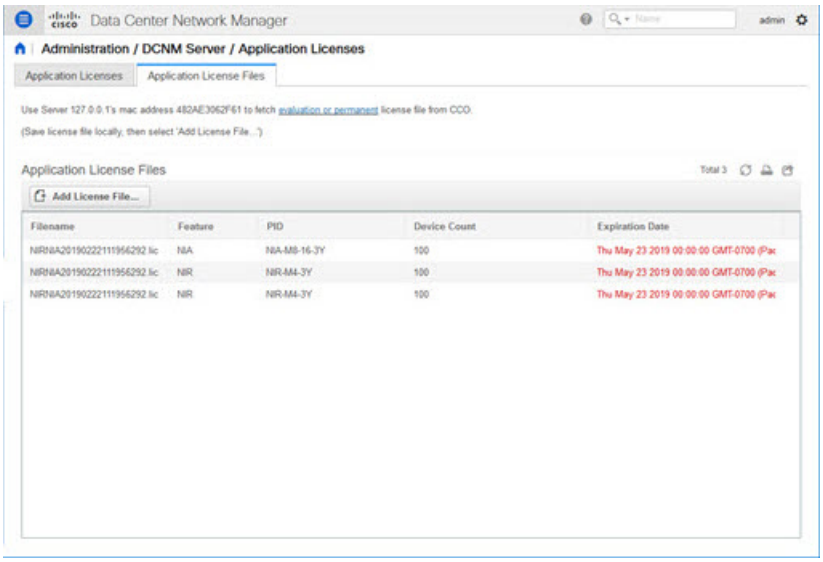

The following image shows a sample error message while uploading an application license file.

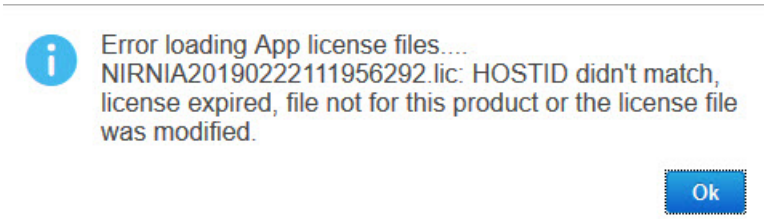

Management Users

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**.

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.

Note When using the DCNM AAA or Radius authentication, you should not specify the hash (#) symbol at the beginning of a secret key. Otherwise, DCNM will try to use **#** as encrypted, and it will fail.

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.

Note When using the DCNM AAA or Radius authentication, you should not specify the hash (#) symbol at the beginning of a secret key. Otherwise, DCNM will try to use **#** as encrypted, and it will fail.

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.

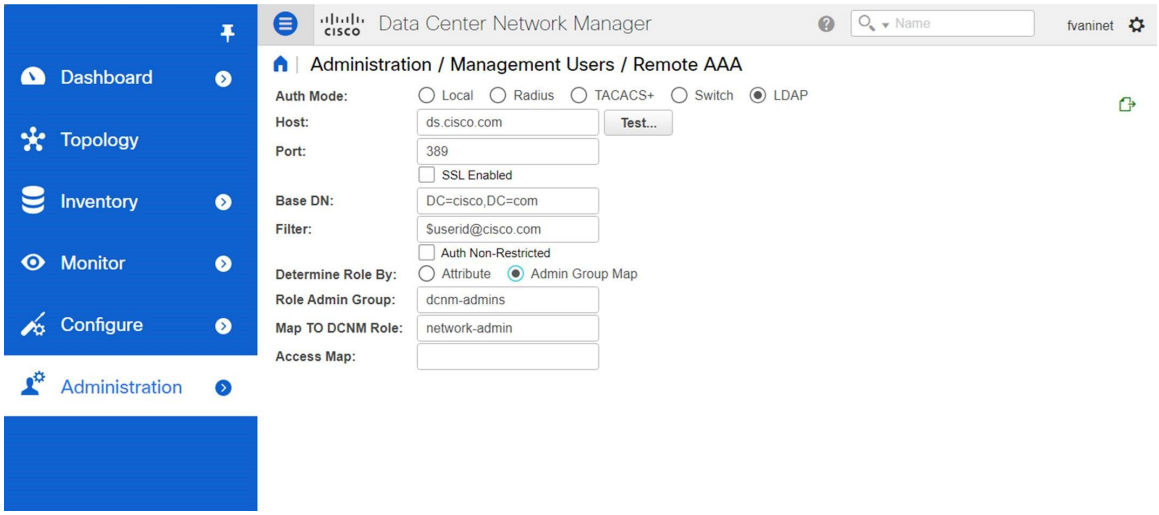

Step 2 In the **Host** field, enter either the IPv4 or IPv6 address.

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.

Note Cisco DCNM establishes a secured connection with the LDAP server using TLS. Cisco DCNM supports all versions of TLS. However, the specific version of TLS is determined by the LDAP server.

For example, if the LDAP server supports TLSv1.2 by default, DCNM will connect using TLSv1.2.

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:

```
ldapserver# 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.

From DCNM release 11.5(1), new user role **device-upg-admin** is added to perform operations only in Image Management window.

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**.

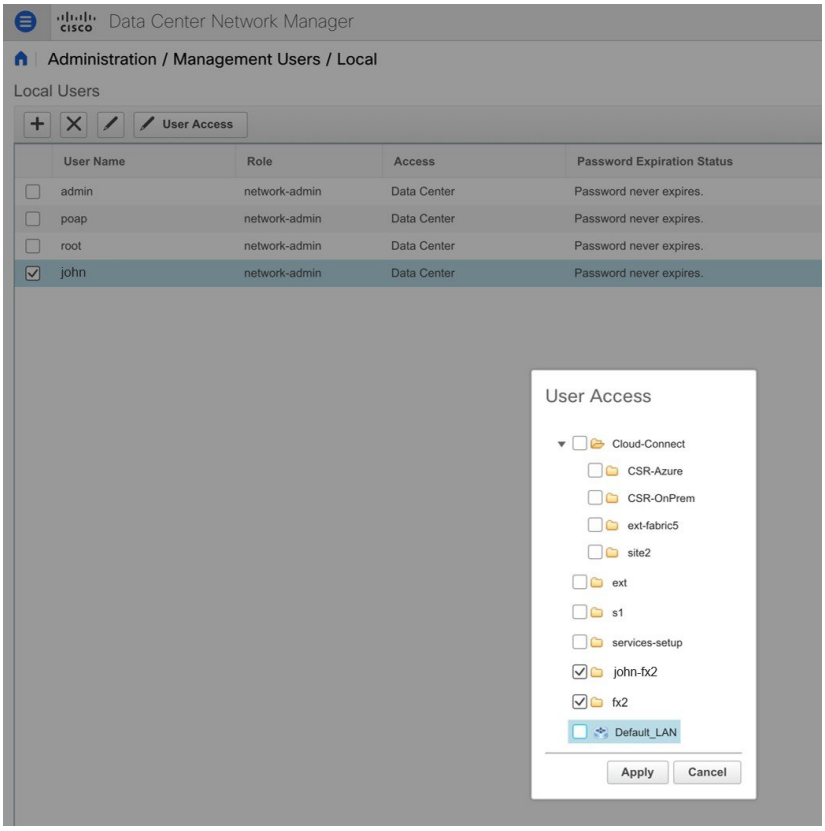

Note The **User Access** button grays out and the value under the **Access** column isn't **Data Center** if the user with the **network-admin** role doesn't have access to the entire data center. In that case, to create a new **network-admin** role user with access to the entire data center use the *addUser.sh/bat* script.

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 Select a value for **Performance Default Polling Interval** from the drop-down list. Valid values are **5 Mins, 10 Mins**, and **15 Mins** . The default value is **5 Mins**.

Step 4 Use the check boxes to select the types of LAN switches for which you want to collect performance data.

Step 5 Click **Apply** to save the configuration.

Step 6 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.

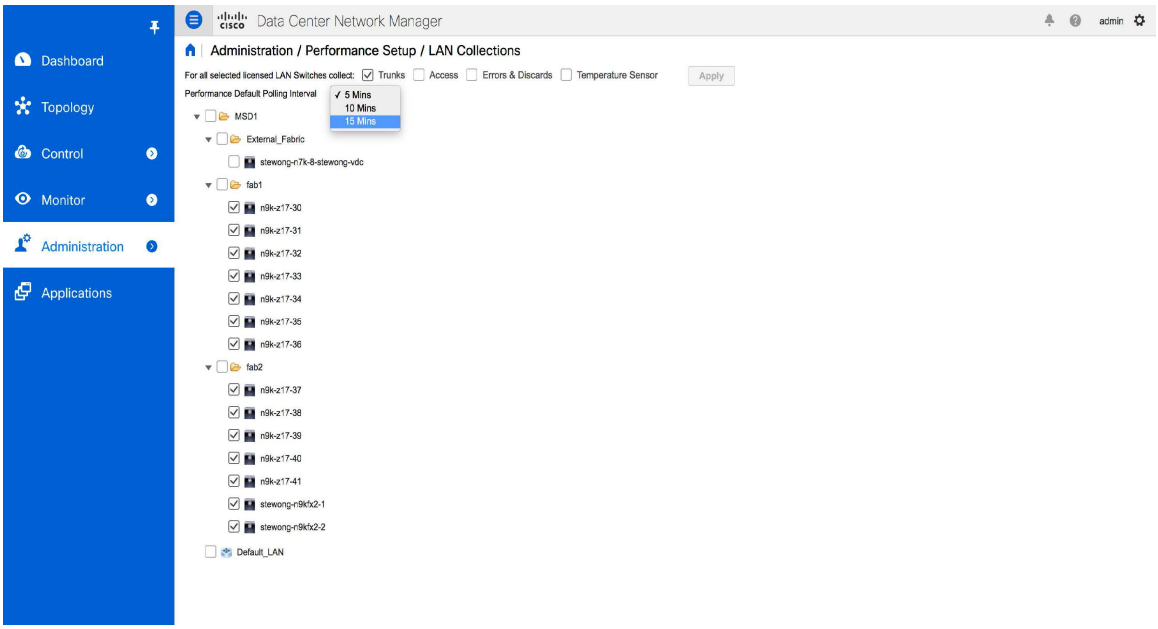

Event Setup

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.

Some SMTP servers may require addition of authentication parameters to emails that are sent from DCNM to the SMTP servers. Starting from Cisco DCNM Release 11.4(1), you can add authentication parameters to the emails that are sent by DCNM to any SMTP server that requires authentication. This feature can be configured by setting up the **SMTP>Authentication** properties in the **Administration>DCNM Server>Server Properties** window. Enter **true** in the **server.smtp.authenticate** field, enter the required username in the **server.smtp.username** field, and enter the required password in the **server.smtp.password** field.

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.

Step 5 In the **Event Count Filter**, add a filter for the event count to the event forwarder.

The forwarding stops forwarding an event if the event count exceeds the limit as specified in the event count filter. In this field, you can specify a count limit. Before an event can be forwarded, the Cisco DCNM checks if its occurrence exceeds the count limit. If it does, the event will not be forwarded.

Step 6 Select the **Snooze** checkbox and specify the **Start** date and time and the **End** date and time. Click **Apply** to save the configuration.

Step 7 Under the **Event Forwarder Rules** table, click the + icon to add an event forwarder rule.

You see the **Add Event Forwarder Rule** dialog box.

Step 8 In the **Forwarding Method**, choose either **E-mail** or **Trap**. If you choose **Trap**, a **Port** field is added to the dialog box.

Step 9 If you choose the **E-mail** forwarding method, enter the IP address in the **Email Address** field. If you choose the **Trap** method, enter the trap receiver IP address in the **Address** field and specify the port number.

You can either enter an IPv4 or IPv6 addresses or DNS server name in the **Address** field.

Step 10 For **Forwarding Scope**, choose the **Fabric/LAN** or **Port Groups** for notification.

Step 11 In the **Source** field, select **DCNM** or **Syslog**.

If you select **DCNM**, then:

a) From the **Type** drop-down list, choose an event type.
b) Check the **Storage Ports Only** check box to select only the storage ports.
c) From the **Minimum Severity** drop-down list, select the severity level of the messages to receive.
d) Click **Add** to add the notification.

If you select **Syslog**, then:

a) In the **Facility** list, select the syslog facility.
b) Specify the syslog **Type**.
c) In the **Description Regex** field, specify a description that matches with the event description.
d) From the **Minimum Severity** drop-down list, select the severity level of the messages to receive.
e) Click **Add** to add the notification.

Note The **Minimum Severity** option is available only if the **Event Type** is set to All.

The traps that are transmitted by Cisco DCNM correspond to the severity type. A text description is also provided with the severity type.

```
trap type(s) = 40990 (emergency)
40991 (alert)
40992 (critical)
40993 (error)
40994 (warning)
40995 (notice)
40996 (info)
40997 (debug)
textDescriptionOid = 1, 3, 6, 1, 4, 1, 9, 9, 40999, 1, 1, 3, 0
```

Removing Notification Forwarding

You can remove notification forwarding.

Procedure

Step 1 Choose **Administration > Event Setup > Forwarding**.

Step 2 Select the check box in front of the notification that you want to remove and click **Delete**.

Event Suppression

Cisco DCNM allows you to suppress the specified events that are based on the user-specified suppressor rules. Such events will not be displayed on the Cisco DCNM Web UI. The events will neither be persisted to DCNM database, nor forwarded via email or SNMP trap.

You can view, add, modify, and delete suppressor rules from the table. You can create a suppressor rule from the existing event table. Select a given event as the template, and invoke the rule dialog window. Event details are automatically ported from the selected event in the event table to the input fields of the rule creation dialog window.

Note You cannot suppress EMC Call Home events from the Cisco DCNM Web UI.

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.

Note After you enter appropriate credentials in **Password**, **Confirm Password** fields and click **Save**, the **Confirm Password** field is blank. A blank **Confirm Password** field implies that the password is saved successfully.

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 529
- Validate Credentials, on page 529
- Clear Switch Credentials, on page 529
- Credentials Management with Remote Access

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.

Credentials Management with Remote Access

DCNM allows you to authenticate users in different modes such as:

- Local Users - In this mode, you can use the Cisco DCNM Web UI to create a new user, assign a role, and provide access to one or more fabrics or groups for the user.
- Remote Users - In this mode, you can log in to DCNM. The DCNM server fetches information from the Remote Authentication server, for example, the Cisco Identity Services Engine (ISE), for AAA authentication. Cisco supports TACACS+, RADIUS, and LDAP options for remote authentication. For more information, see Remote AAA.

When you configure DCNM for remote authentication, the AAA server handles both authentication and authorization. DCNM forwards the entered user login and password to the AAA server to check for authentication. Post authentication, the AAA server returns the appropriate privileges/role assigned to the user through the **cisco-avpair** attribute. This attribute can contain the list of fabrics that a particular user can access. The supported roles for DCNM LAN deployments are as follows:

- network-admin
- network-operator
- network-stager
- access-admin
- device-upg-admin

Each role allows read and optional write privileges to resources of a certain category. For more information about DCNM roles, refer to Enhanced Role-based Access Control in Cisco DCNM.

Both device discovery credentials and LAN credentials provide write access to the devices, but they differ—as the write operation is performed only with LAN credentials. Device discovery credentials are associated with each device and entered only once, that is, when you import the device into DCNM. DCNM uses these credentials for periodic rediscovery using a mix of SSH and SNMPv3 access to the device. However, LAN credentials are configured for every user on a per-user basis. If a user with an appropriate role has access to DCNM, then that user can enter the LAN credentials to get write access to the devices. The write operations use the LAN credentials to access the device, which allows for an appropriate audit trail of the changes made in DCNM by every user and the resultant changes in the device.

When you configure DCNM using Remote Authentication Methods such as TACACS+ or RADIUS, the users can set their LAN credentials as follows:

- Regular AAA Remote Authentication
- AAA Remote Authentication Passthrough Mechanism
- AAA Remote Authentication Using DCNM Service Account

Regular AAA Remote Authentication

Post authentication, when a user with an appropriate role logs in to DCNM for the first time, DCNM prompts the user to enter the LAN credentials. As mentioned earlier, DCNM uses these credentials to provide write access to the devices. All users must follow this process. Consider that an internal business policy requires the users to change password every 3-6 months. Then all the users must update their passwords for device access in the DCNM **LAN Credentials** window. Also, they must update their passwords in the AAA server.

For example, let us consider a user named John, who has authentication on the ISE server.

1. John logs in to DCNM with his user credentials.
2. The ISE server authenticates the user credentials of John, and DCNM displays a message to enter his LAN switch credentials. DCNM uses these credentials to perform various configurations and write operations on the devices.

3. John enters his LAN switch credentials. DCNM uses the LAN switch credentials for all write operations triggered by John on all devices. However, John can also opt to enter LAN switch credentials on a per-device access basis. This per-device access option overrides the access provided by entering the default credentials.

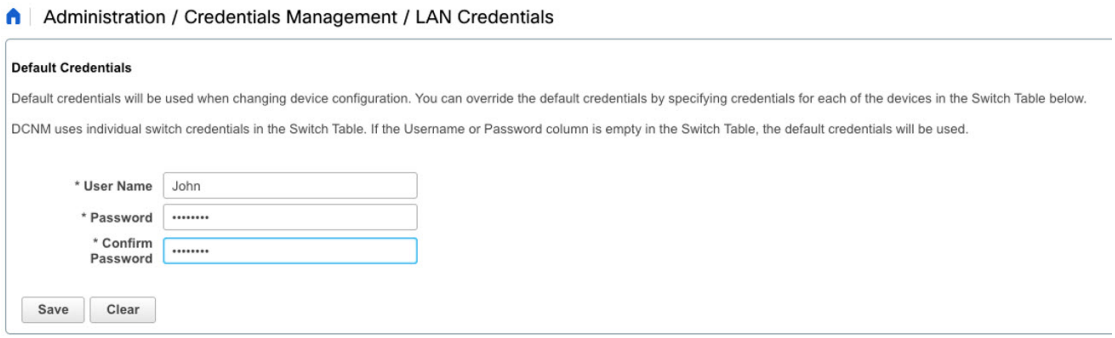

When John logs in to DCNM again, DCNM doesn't display any message to enter the LAN switch credentials as it has already captured his LAN switch credentials. John uses the same credentials to log in to DCNM and to the devices that he can access.

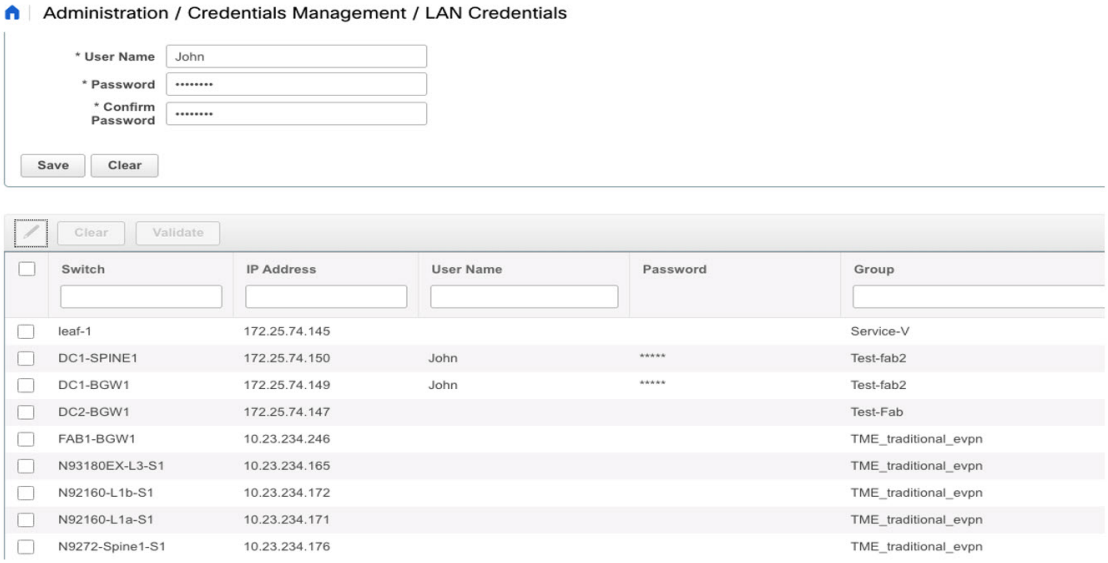

4. Now, consider that after a few months, the Corporate IT policy changes. Then John must update his password in the Remote AAA server, and also perform Step 3 to allow DCNM to update his LAN switch credentials.

Thus, in this mode, when John logs in to the DCNM Web GUI with his updated password, DCNM doesn't display any message to enter LAN credentials. However, John must update the password in LAN Credentials. Updating the password is necessary as it allows DCNM to inherit the newly updated password and perform write operations on the devices.

AAA Remote Authentication Passthrough Mechanism

In this mode, when a user enters the username and password to log in to DCNM, DCNM automatically copies the user credentials to the Default Credentials in the LAN switch credentials settings for that user. As a result, when the user logs in for the first time, DCNM doesn't display the message to enter the LAN switch credentials.

1. Use SSH to log in to DCNM as a sysadmin user.
2. Log in to the `/root/directory` using the **su** command.
3. Navigate to the `/usr/local/cisco/dcm/fm/conf/server.properties` file.
4. Add the following server property to the file and save the changes.

   dcnm.lanSwitch.sameUserAccount=true

```
[root@dcnm sysadmin]# cat /usr/local/cisco/dcm/fm/conf/server.properties | grep dcnm.lan
dcnm.lanSwitch.sameUserAccount=true
[root@dcnm sysadmin]#
```

5. Restart DCNM using the **service FMServer restart** command.
6. Now, John logs in to DCNM.
7. After successful authentication, DCNM doesn't display the message to update the LAN switch credentials, as it automatically copies this information to the LAN switch credentials.

8. Consider that after a few months, the Corporate IT policy changes. In this mode, John must update his password in the Remote AAA server. After that, when John logs in to DCNM, DCNM automatically copies the updated credentials to the Default LAN Credentials associated with the user John.

AAA Remote Authentication Using DCNM Service Account

Often, the customers prefer to track all the changes made from the DCNM controller with a common service account. In the following example, a user makes changes using the DCNM controller, which results in changes on the device. These changes are audit logged on the device, against a common service account. Thus, it is possible to distinguish the controller-triggered changes from other changes (also known as Out-of-Band changes) made by the user directly on the device. The Out-of-Band changes appear in the device accounting logs as made from the user account.

For example, create a service account with the name **Robot** on the remote AAA server. Using the corresponding credentials, the Robot user can log in to DCNM. The Robot user can enter the default LAN credentials to have write access to the devices. The DCNM network-admin enables a server property that automatically sets the default LAN credentials for all the users and inherits the default LAN credentials associated with Robot.

Therefore, when any user logs in to DCNM and makes any configuration changes, DCNM pushes the changes to the devices using the LAN credentials of Robot. The DCNM deployment history logs track the user who triggered the change and display the corresponding changes deployed from DCNM to the switch in the audit log with the user Robot.

To set up the service account on the DCNM, perform the following steps:

1. Use SSH to log in to DCNM as a sysadmin user.
2. Log in to the `/root/ directory` using the **su** command.
3. Navigate to the `/usr/local/cisco/dcm/fm/conf/server.properties` file.
4. Add the following server property to the file and save the changes.

   service.account=robot

   Note You can enable either an AAA passthrough account or a Service Account.

```
[root@dcnm sysadmin]# cat /usr/local/cisco/dcm/fm/conf/server.properties | grep robot
service.account=robot
[root@dcnm sysadmin]#
```

5. Restart DCNM using the **service FMServer restart** command.
6. Now, John logs in to DCNM.
7. After successful authentication, DCNM doesn't display the message to update the LAN switch credentials. However, when John navigates to the **LAN Credentials** page, DCNM displays a message stating that the Service Account is enabled in DCNM and, hence, all LAN credentials will be inherited from the service account.

Service Account Configuration Audit

The following workflow example allows for verification of the configuration audit while using the DCNM service account feature. However, you must have completed the Service Account Activation procedure.

1. John creates a test loopback on a device.

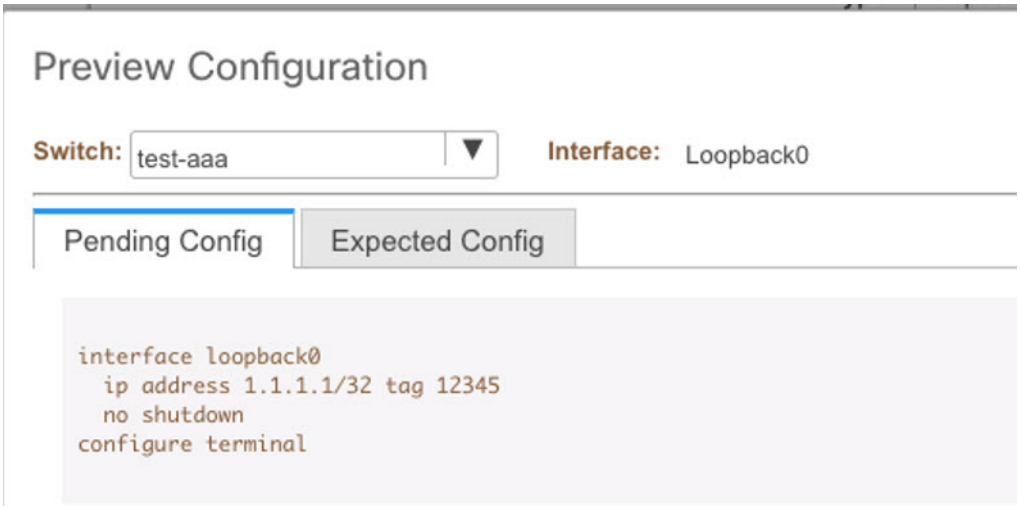

2. John deploys the configuration using DCNM.
3. The DCNM Deployment history confirms that John made the recent configuration change.

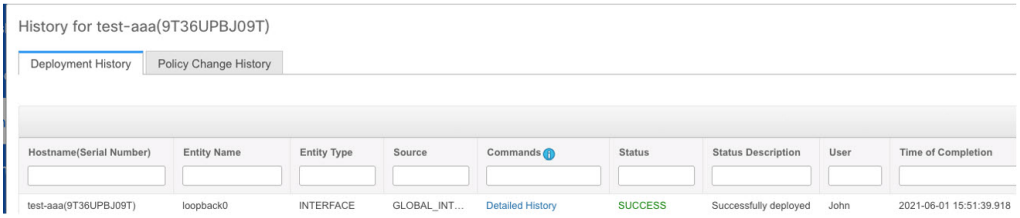

4. The accounting logs of the device indicate that the DCNM Service Account (that is, Robot, in this example) has triggered the changes on the NX-OS device.

```
Tue Jun  1 22:50:04 2021:type=update:id=172.25.74.142@pts/5:user=robot:cmd=terminal length 0 (SUCCESS)
Tue Jun  1 22:50:04 2021:type=update:id=172.25.74.142@pts/5:user=robot:cmd=terminal session-timeout 30 (SUCCESS)
Tue Jun  1 22:50:04 2021:type=update:id=172.25.74.142@pts/5:user=robot:cmd=terminal dont-ask (SUCCESS)
Tue Jun  1 22:50:04 2021:type=update:id=172.25.74.142@pts/5:user=robot:cmd=terminal width 511 (SUCCESS)
Tue Jun  1 22:50:05 2021:type=update:id=172.25.74.142@pts/5:user=robot:cmd=configure terminal ; interface loopback0 (REDIRECT)
Tue Jun  1 22:50:05 2021:type=update:id=172.25.74.142@pts/5:user=robot:cmd=configure terminal ; interface loopback0 (SUCCESS)
Tue Jun  1 22:50:05 2021:type=update:id=172.25.74.142@pts/5:user=robot:cmd=configure terminal ; interface loopback0 ; ip address 1.1.1.1/32 tag 12345
(REDIRECT)
Tue Jun  1 22:50:05 2021:type=update:id=172.25.74.142@pts/5:user=robot:cmd=configure terminal ; interface loopback0 ; ip address 1.1.1.1/32 tag 12345
(SUCCESS)
Tue Jun  1 22:50:06 2021:type=update:id=172.25.74.142@pts/5:user=robot:cmd=configure terminal ; interface loopback0 ; no shutdown (REDIRECT)
Tue Jun  1 22:50:06 2021:type=update:id=172.25.74.142@pts/5:user=robot:cmd=configure terminal ; interface loopback0 ; no shutdown (SUCCESS)
Tue Jun  1 22:50:06 2021:type=stop:id=172.25.74.142@pts/5:user=robot:cmd=shell terminated because the ssh session closed
test-aaa#
```

PART I

Applications

- Applications Framework, on page 539
- Endpoint Locator, on page 559
- IPAM Integrator, on page 595
- Health Monitor, on page 601
- PTP Monitoring, on page 609
- Programmable Reports, on page 613
- ServiceNow Integration, on page 629

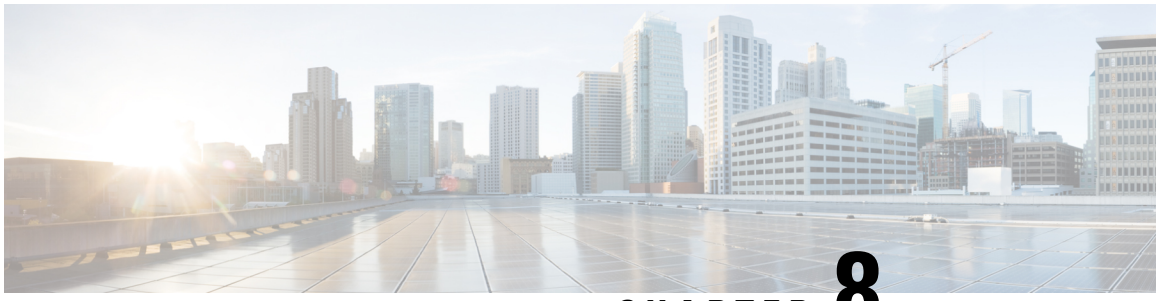

CHAPTER 8

Applications Framework

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 539
- Cisco DCNM in Clustered Mode, on page 540
- Installing and Deploying Applications, on page 551
- Application Framework User Interface, on page 554
- Catalog, on page 555
- Compute, on page 555
- Preferences, on page 557
- Failure Scenario, on page 557

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 551.

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 if 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.

Refer to *Installing Cisco DCNM Compute Node* in the *Cisco DCNM Installation and Upgrade Guide for LAN Fabric Deployment*.

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.

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 (the maximum qualified scale is 350 switches starting from Cisco DCNM 11.3(1) Release), 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)

Table 23: Upto 80 Switches

Node	CPU Deployment Mode	CPU	Memory	Storage	Network
DCNM	OVA/ISO	16 vCPUs	32G	500G HDD	3xNIC
Computes	NA	—	—	—	—

Table 24: 81–350 Switches

Node	CPU Deployment Mode	CPU	Memory	Storage	Network
DCNM	OVA/ISO	16 vCPUs	32G	500G HDD	3xNIC
DCNM on Applications Service Engine (SE)	ISO	16 vCPUs	32G	500G HDD	3xNIC
Computes	OVA/ISO	16 vCPUs	64G	500G HDD	3xNIC

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.

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.

Refer to Installing Cisco DCNM Compute Node in the *Cisco DCNM Installation and Upgrade Guide for LAN Fabric Deployment, Release 11.5(1)*.

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 2: Security Settings for vSwitch Port-Group

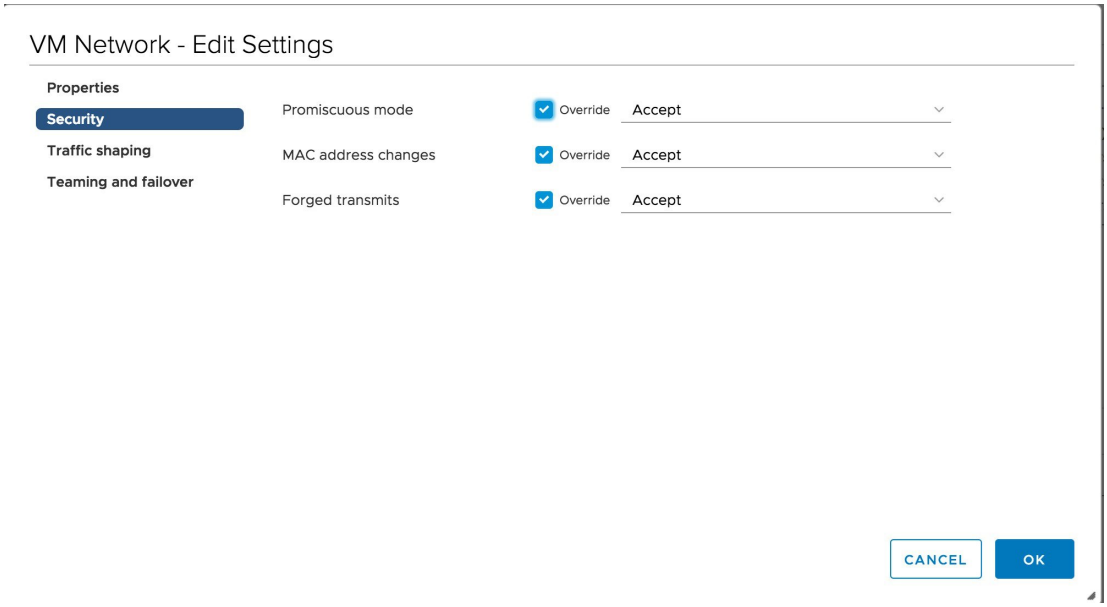

Figure 3: Security Settings for DVSwitch Port-Group

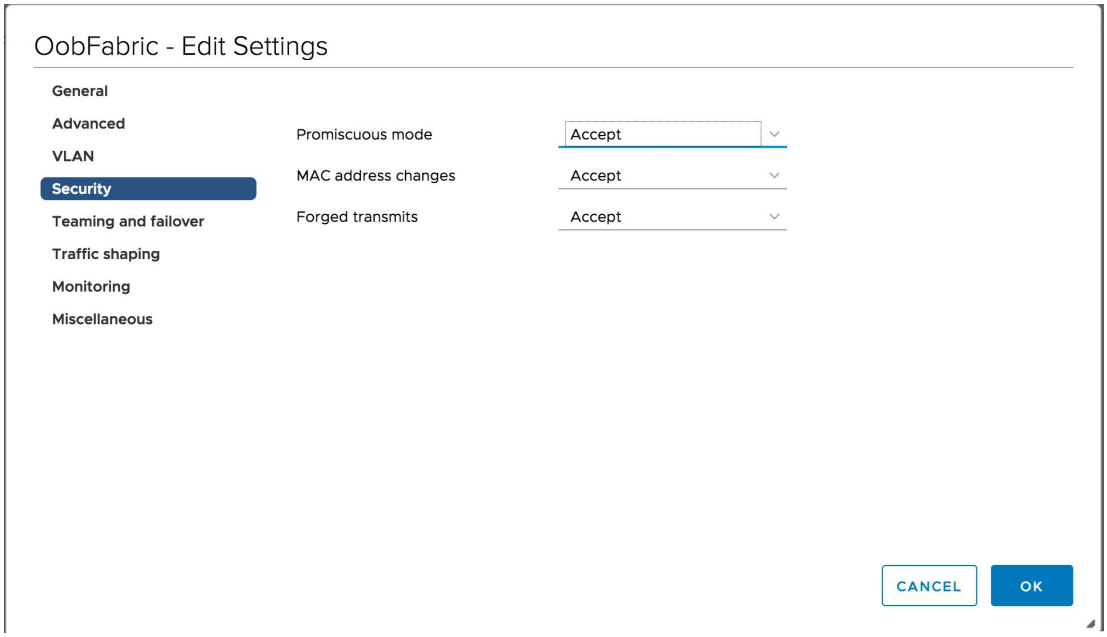

Note Ensure that you repeat this procedure on all the hosts, where a Compute OVA is running.

Enabling the Compute Cluster

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.

Note If you enable clustered mode while installing Cisco DCNM, you don't need to enable cluster. The compute nodes will be discovered on Cisco DCNM **Web UI > Applications > Compute**. Go to Compute, on page 555 to form a cluster.

If you did not enable clustered mode while installation, 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.

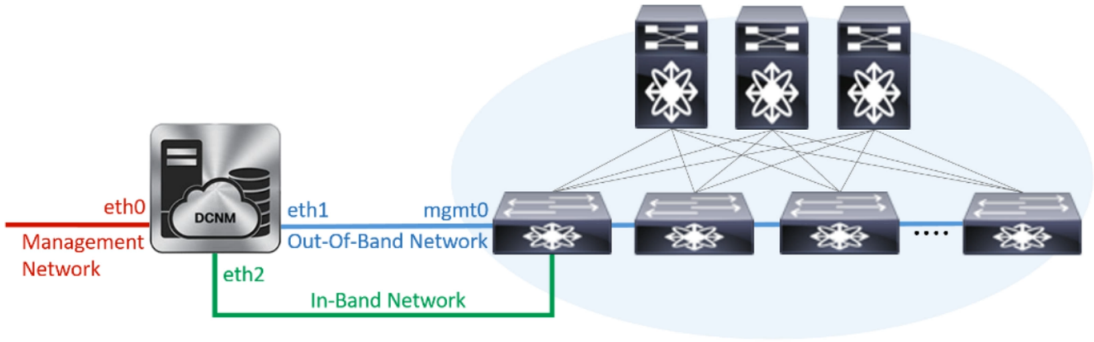

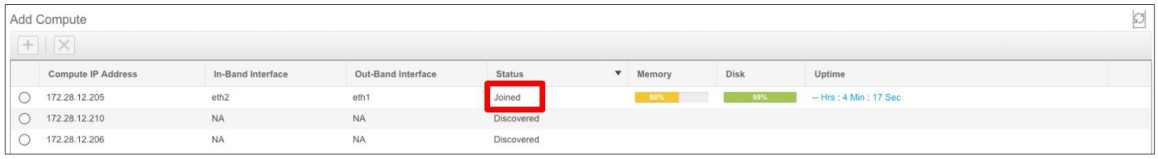

The other two computes are Discovered automatically, and is displayed on the Cisco DCNM **Web UI** > **Applications** > **Compute**.

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 546.

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>***--compute***<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.

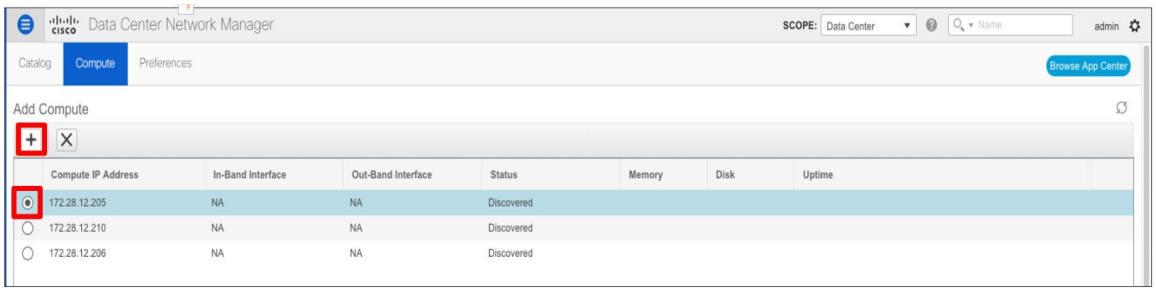

- While using Compute, ensure that Cisco DCNM GUI shows nodes as Joined.
- Offline indicates some connectivity issues, therefore no applications are running on Offline Computes.
- Failed indicates that the compute node could not join the cluster.
- Health indicates the amount of free memory and disk on the Compute node. The Health Monitor 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.

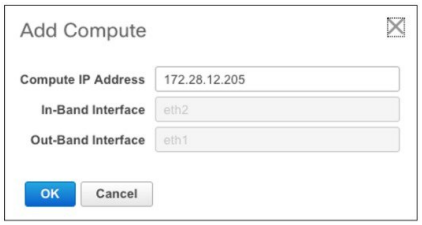

Step 4 Click **OK**.

The Status for that Compute IP changes to **Joining**.

Wait until the Compute IP status shows **Joined**.

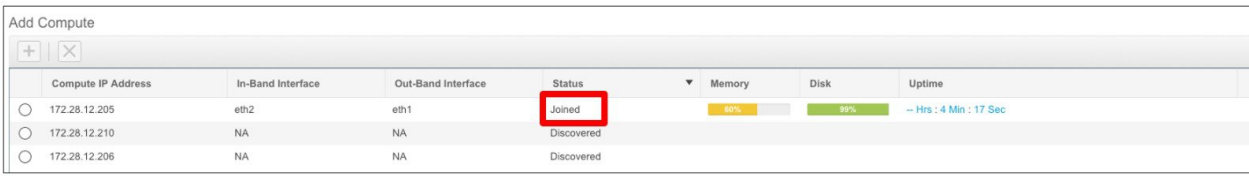

Step 5 Repeat the above steps to add the remaining compute node.

All the Computes appear as **Joined**.

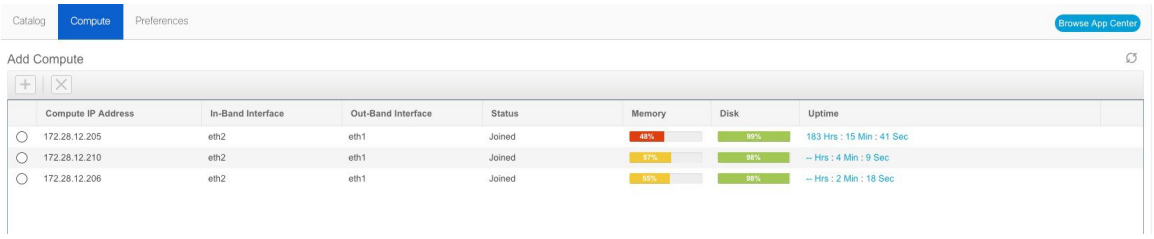

Note When you install compute as a virtual machine on the VMware platform, vSwitch or DV switch port groups associated eth1 and eth2 must allow for packets that are associated with Mac address other than eth1 and eth2 to be forwarded.

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 548 to Step 7, on page 548 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 3 to 5 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.

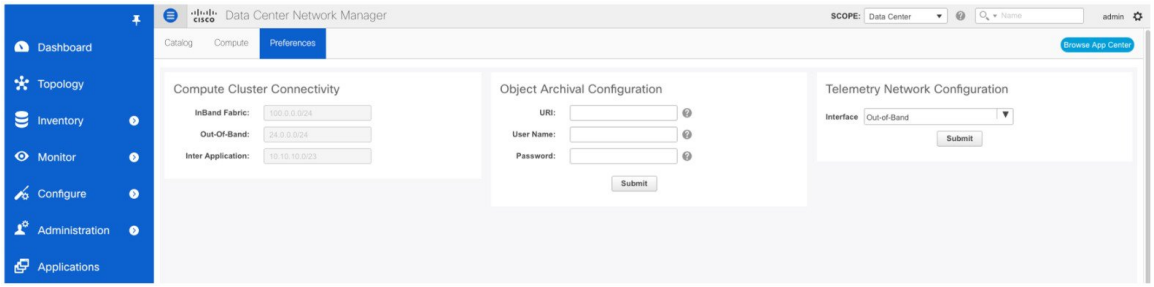

Compute Cluster Connectivity

The fields show the IP address that is used to configure the connectivity interfaces for the cluster node. The IP addresses for in-band fabric, out-of-band fabric, and Inter-Application are displayed.

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.

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
- LAN Classic fabric in Monitored or Managed Mode (Applicable for DCNM 11.4(1) or later)

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
```

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.

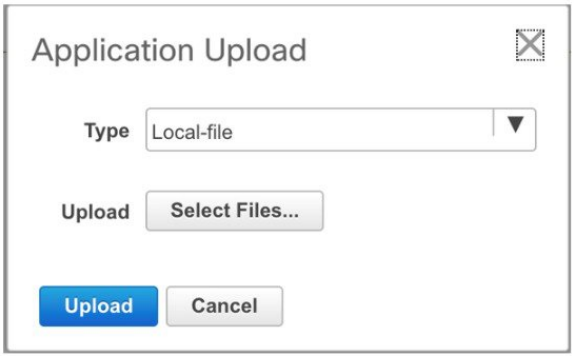

From the Type drop-down list, select one of the following:

- 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 is 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.

The green icon on the left-top corner indicates that the application is launched successfully and is operational.

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.

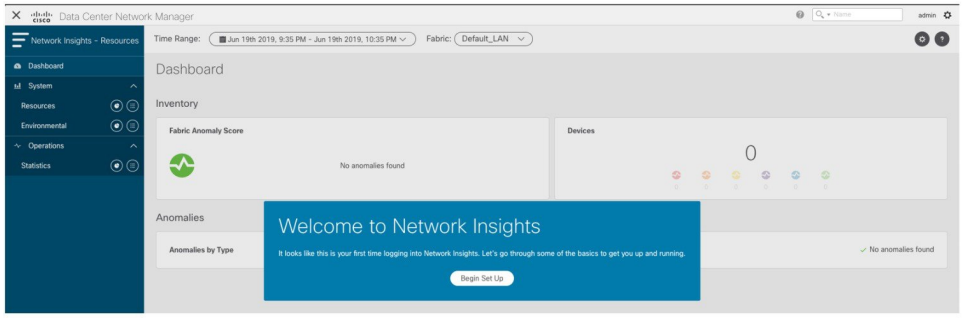

To check the services that are running go back to **Applications > Catalog**. Click the gear icon on the left-bottom on the application icon to view the Application Specifications. The Info tab displays the running container information and the Specs tab displays the configuration as shown in the figures below.

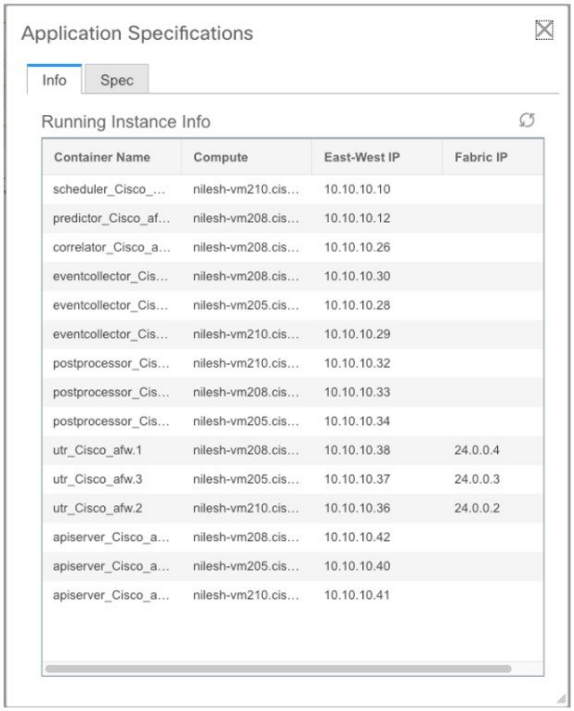

Container Name	Compute	East-West IP	Fabric IP
scheduler_Cisco_...	nilesh-vm210.cis...	10.10.10.10	
predictor_Cisco_af...	nilesh-vm208.cis...	10.10.10.12	
correlator_Cisco_a...	nilesh-vm208.cis...	10.10.10.26	
eventcollector_Cis...	nilesh-vm208.cis...	10.10.10.30	
eventcollector_Cis...	nilesh-vm205.cis...	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 554.

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.
4. Click **Stop** to stop the application from streaming data from Cisco DCNM.

   The Green icon disappears after the application is successfully stopped.
5. After you stop the application, click the **Waste Basket** icon to remove the application from the Catalog.

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 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 555.

  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 549.

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:

- Health Monitor (2.1)
- PTP Monitoring (1.1)
- Kibana (2.0)
- Programmable report (1.1.0)
- Elastic Service (1.1)
- Compliance (4.0.0)
- Debug Tools (2.1)
- IPAM Integrator (1.0)
- Endpoint Locator (2.1)
- Kubernetes Visualizer (1.1)
- vmmplugin (4.1)

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 551.

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.

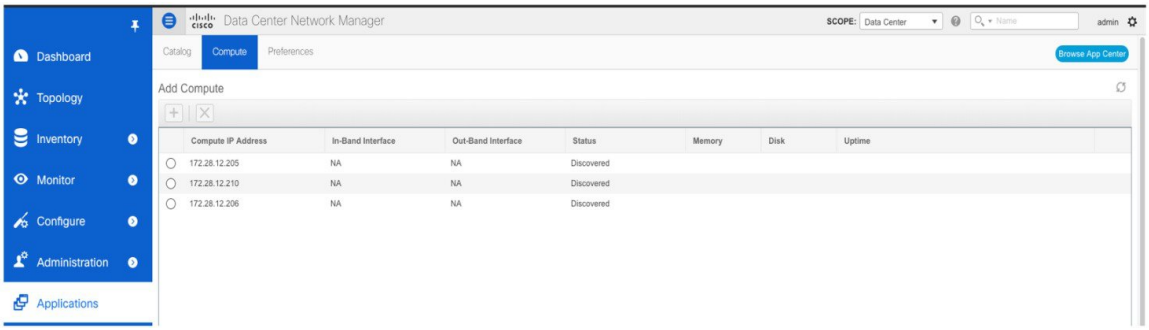

Note If the NTP server for compute nodes is not synchronized with the NTP server for DCNM Servers (Active and Standby) and Computes, you cannot configure a cluster.

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 25: 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. • Joined • Discovered • Failed • Offline
Memory	Specifies the memory that is consumed by the node.
Disk	Specifies the disk space that is consumed on the compute node.

Field	Description
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 546.

To configure or modify the Cluster Connectivity preferences, see Preferences, on page 549.

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.

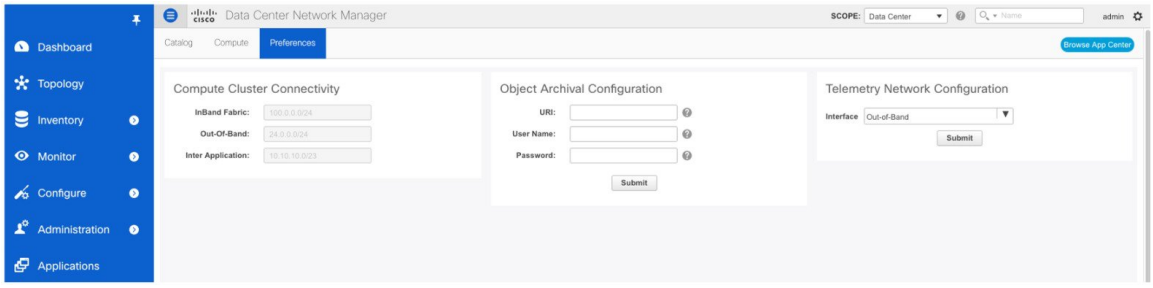

Compute Cluster Connectivity

The fields show the IP address that is used to configure the connectivity interfaces for the cluster node. The IP addresses for in-band fabric, out-of-band fabric, and Inter-Application are displayed.

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.

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.

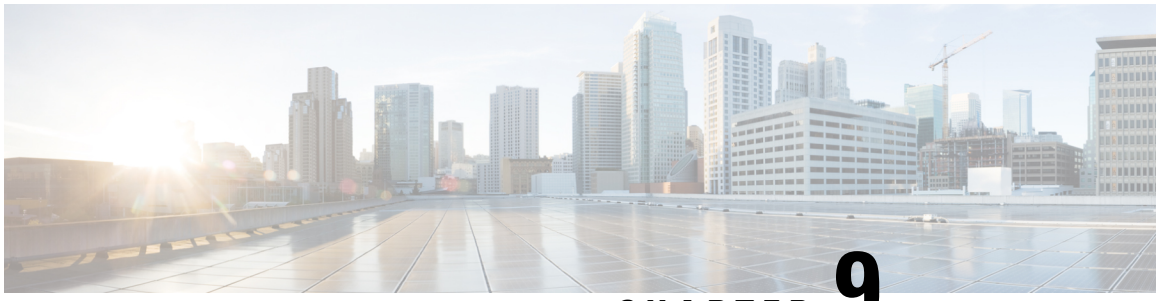

CHAPTER 9

Endpoint Locator

- Endpoint Locator , on page 559
- Monitoring Endpoint Locator, on page 583

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. Starting from Cisco DCNM Release 11.3(1), the EPL feature is also capable of displaying MAC-Only endpoints. By default, MAC-Only endpoints are not displayed. 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). Starting from Cisco DCNM Release 11.3(1), EPL is also capable of displaying MAC-Only endpoints.Select the **Process MAC-Only Advertisements** checkbox while configuring EPL to enable processing of EVPN Route-type 2 advertisements having a MAC address only. L2VNI:MAC is the unique endpoint identifier for all such endpoints. EPL can now track endpoints in Layer-2 only network deployments where the Layer-3 gateway is on a firewall, load-balancer, or other such nodes.

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).
- Starting from Cisco DCNM Release 11.3(1), you can enable the EPL feature for upto 4 fabrics. This is supported only in clustered mode.
- Starting from Cisco DCNM Release 11.3(1), EPL is supported on Multi-Site Domain (MSD).
- Starting from Cisco DCNM Release 11.3(1), IPv6 underlay is supported.
- Support for high availability
- Support for endpoint data that is stored for up to 180 days, amounting to a maximum of 100 GB storage space.
- Support for optional flush of the endpoint data in order to start afresh.
- Supported scale: 50K unique endpoints per fabric. A maximum of 4 fabrics is supported. However, the maximum total number of endpoints across all fabrics should not exceed 100K.

  Starting from Cisco DCNM Release 11.4(1), if the total number of endpoints across all fabrics exceeds 100K, an alarm is generated and is listed under the **Alarms** icon at the top right of the window. This icon starts flashing whenever a new alarm is generated.

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

Configuring Endpoint Locator

The DCNM OVA or the ISO installation comes with three interfaces:

- eth0 interface for external access
- eth1 interface for fabric management (Out-of-band or OOB)
- eth2 interface for in-band network connectivity

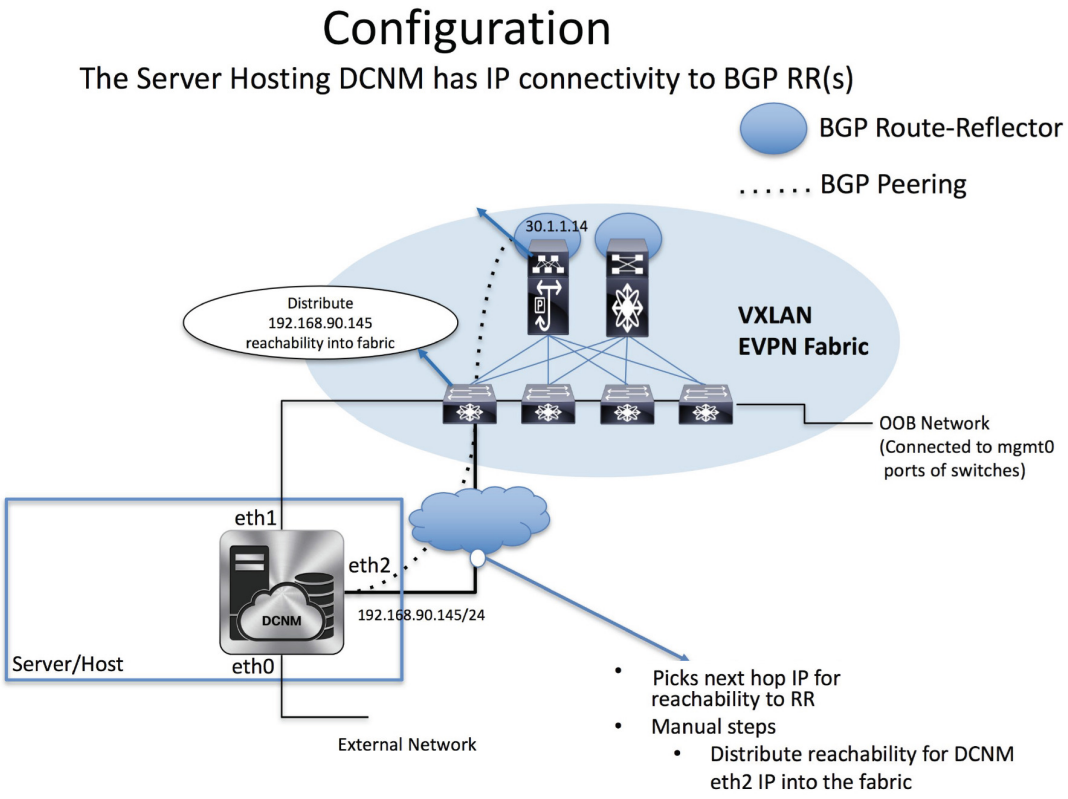

The eth1 interface provides reachability to the devices via the mgmt0 interface either Layer-2 or Layer-3 adjacent. This allows DCNM to manage and monitor these devices including POAP. EPL requires BGP peering between the DCNM and the Route-Reflector. Since the BGP process on Nexus devices typically runs on the default VRF, in-band IP connectivity from the DCNM to the fabric is required. For this purpose, the eth2 interface can be configured using the **appmgr setup inbandappmgr update network-properties** 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), run the **appmgr setup inbandappmgr update network-properties** command again. Refer Editing Network Properties Post DCNM Installation to run the **appmgr setup inbandappmgr update network-properties** command.

Note The setup of eth2 interface on the DCNM is a prerequisite of any application that requires the in-band connectivity to the devices within fabric. This includes EPL and Network Insights Resources (NIR).

Note For configuring EPL in standalone mode, you must add a single neighbor to EPL. DCNM eth2 IP address is EPL IP.

On the fabric side, for a standalone DCNM deployment, if the DCNM eth2 port is directly connected to one of the front-end interfaces on a leaf, then that interface can be configured using the **epl_routed_intf** template. An example scenario of how this can be done when IS-IS or OSPF is employed as the IGP in the fabric, is depicted below:

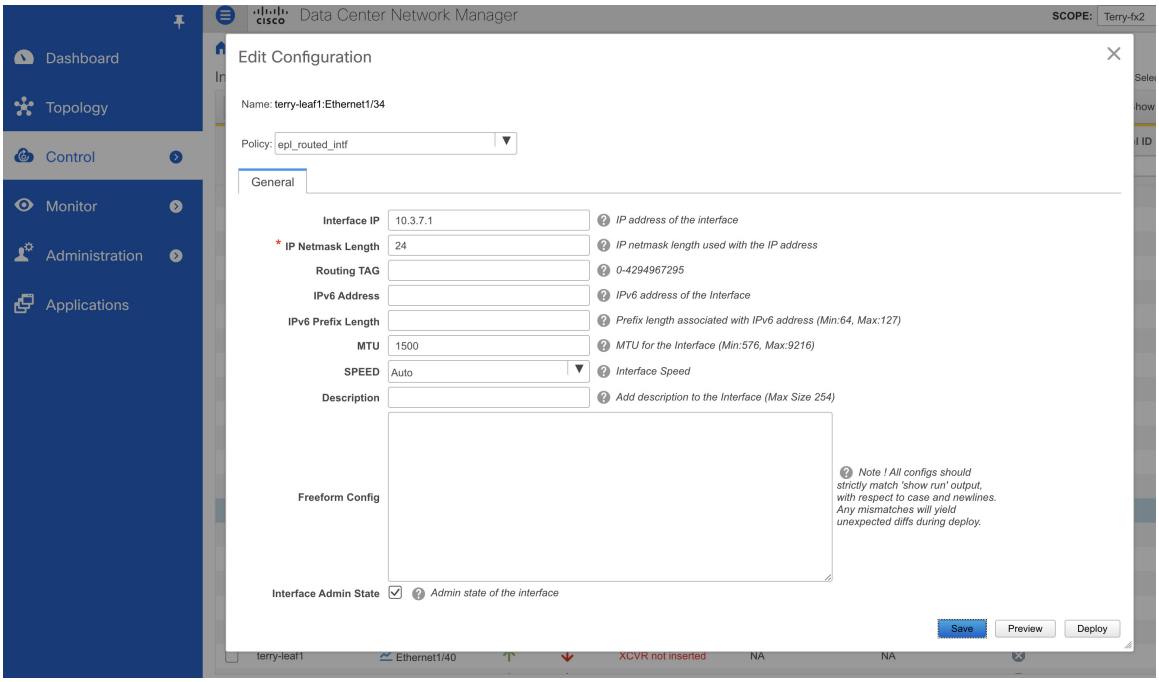

However, for redundancy purposes, it is always advisable to have the server on which the DCNM is installed to be dual-homed or dual-attached. With the OVA DCNM deployment, the server can be connected to the switches via a port-channel. This provides link-level redundancy. To also have node-level redundancy on the network side, the server may be attached to a vPC pair of Leaf switches. In this scenario, the switches must be configured such that the HSRP VIP serves as the default gateway of the eth2 interface on the DCNM. The following image depicts an example scenario configuration:

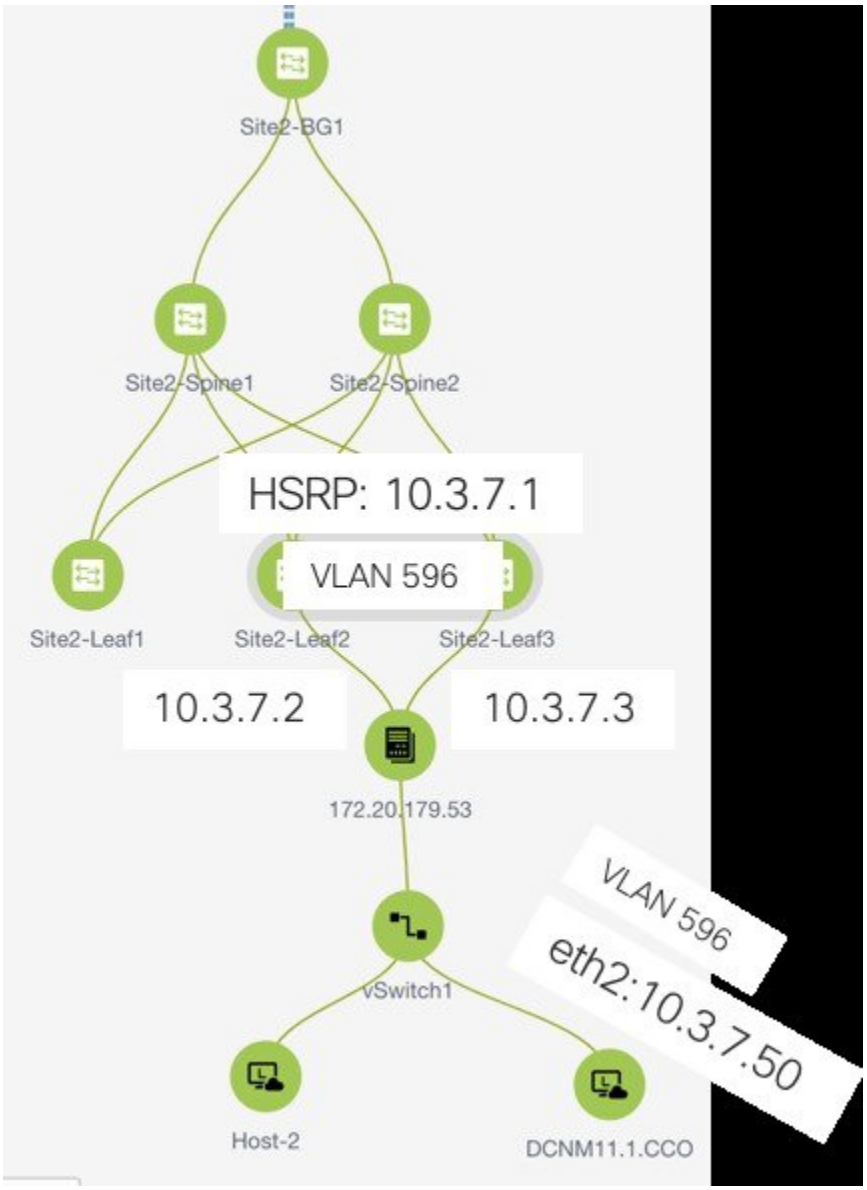

In this example, the server with the DCNM VM is dual-attached to a vPC pair of switches that are named Site2-Leaf2 and Site2-Leaf3 respectively. VLAN 596 associated with the IP subnet 10.3.7.0/24 is employed for in-band connectivity. You can configure the vPC host port toward the server using the **interface vpc trunk host** policy as shown is the following image:

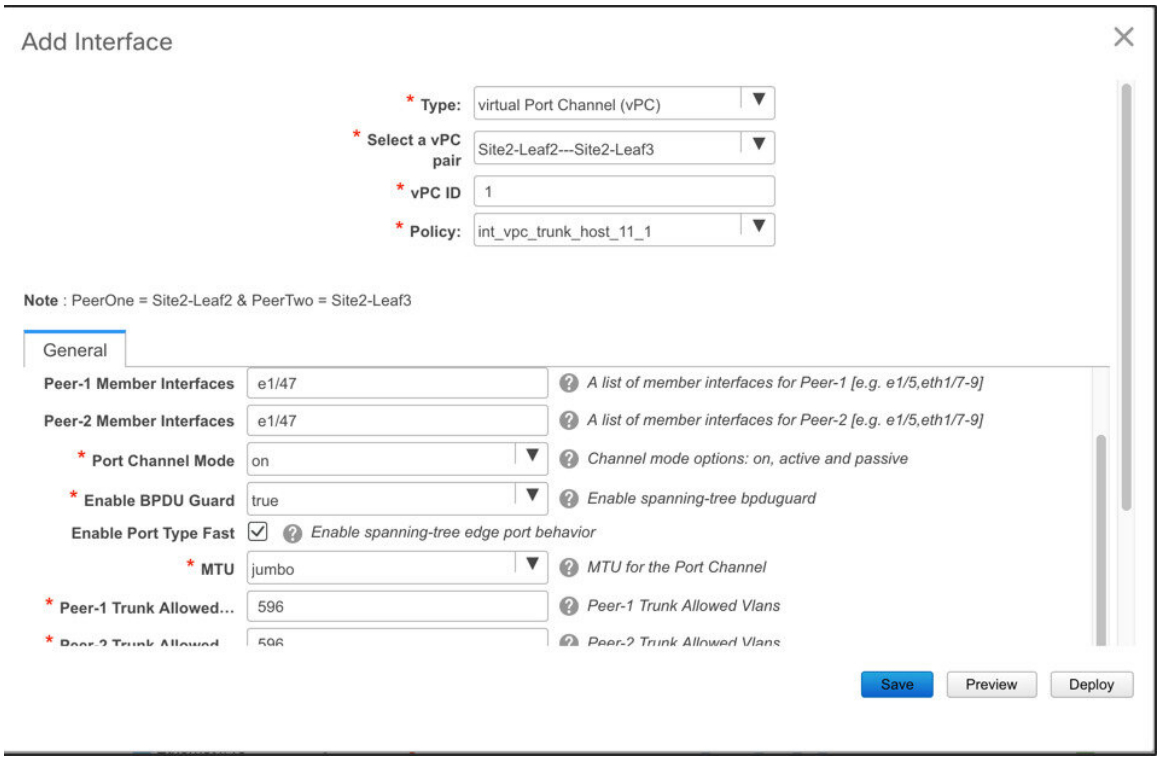

For the HSRP configuration on Site2-Leaf2, the **switch_freeform** policy may be employed as shown in the following image:

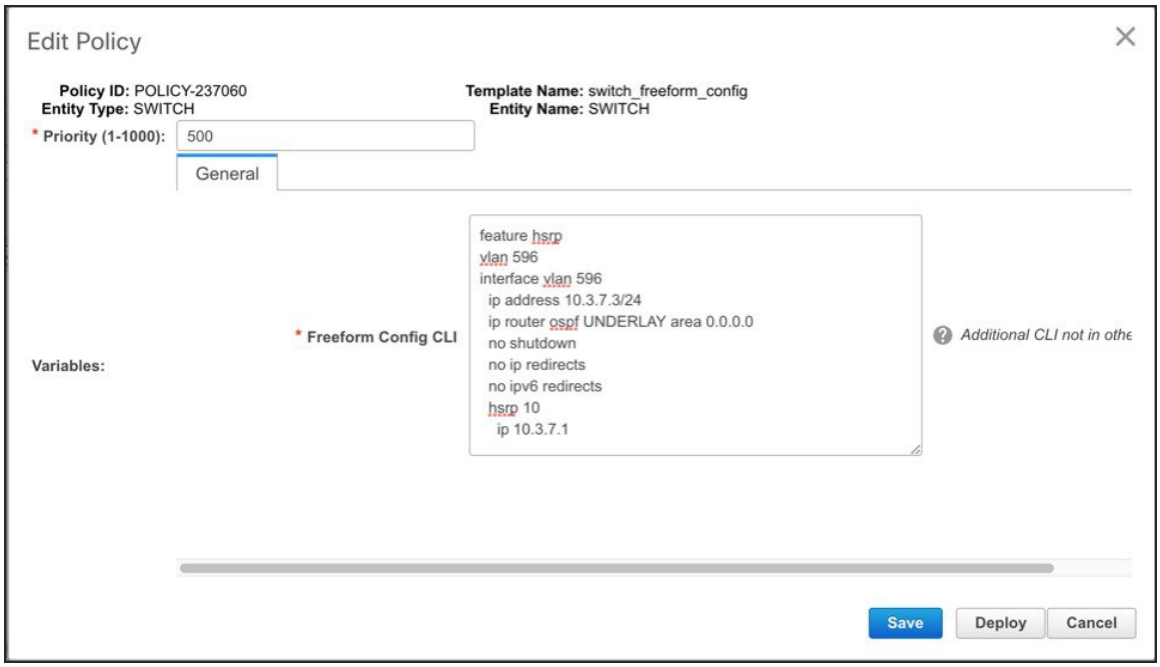

You can deploy a similar configuration on Site2-Leaf3 while using IP address 10.3.7.2/24 for SVI 596. This establishes an in-band connectivity from the DCNM to the fabrics over the eth2 interface with the default gateway set to 10.3.7.1.

After you establish the in-band connectivity between the physical or virtual DCNM and the fabric, you can establish BGP peering.

During the EPL configuration, 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.

To configure Endpoint Locator from the Cisco DCNM Web UI, choose **Control > Endpoint Locator > Configure**. The **Endpoint Locator** window appears.

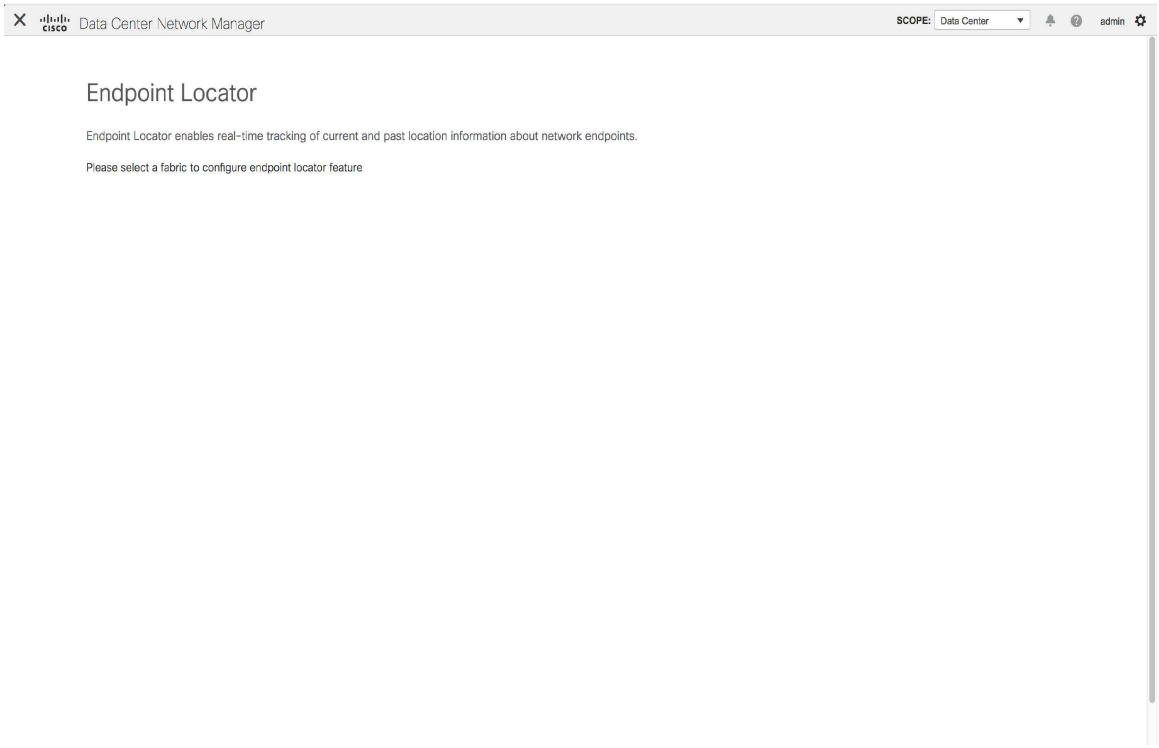

Select a fabric from the **Scope** drop-down list on which the endpoint locator feature should be enabled to track endpoint activity. You can enable EPL for one fabric at a time.

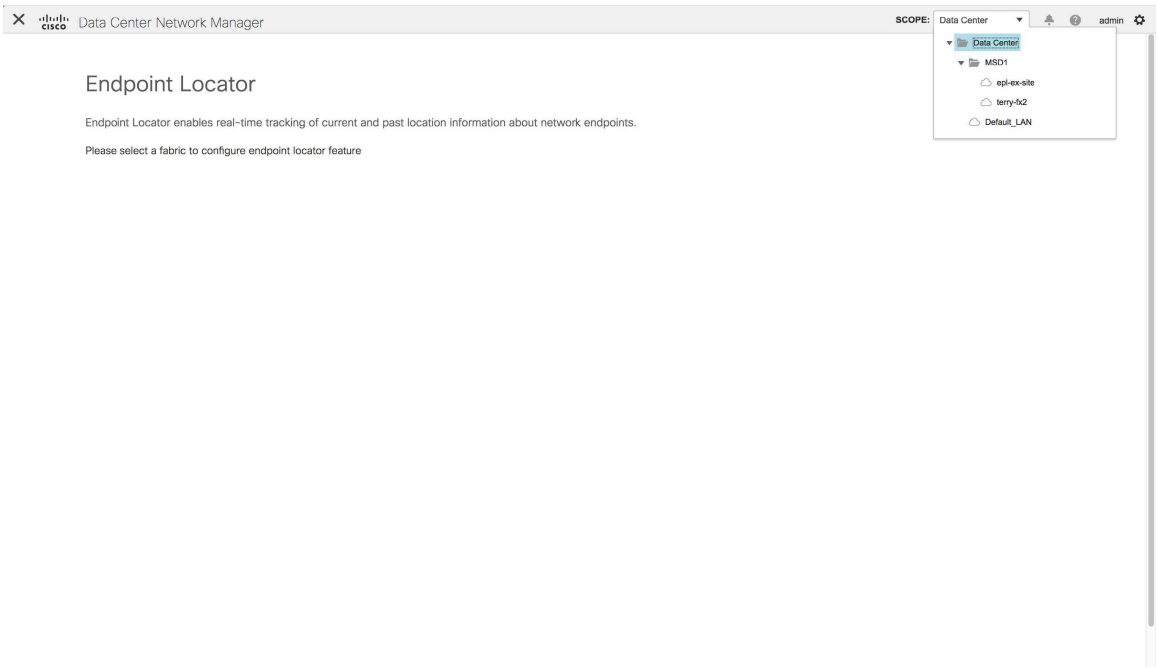

Select the switches on the fabric hosting the RRs from the drop-down list. Cisco DCNM will peer with the RRs.

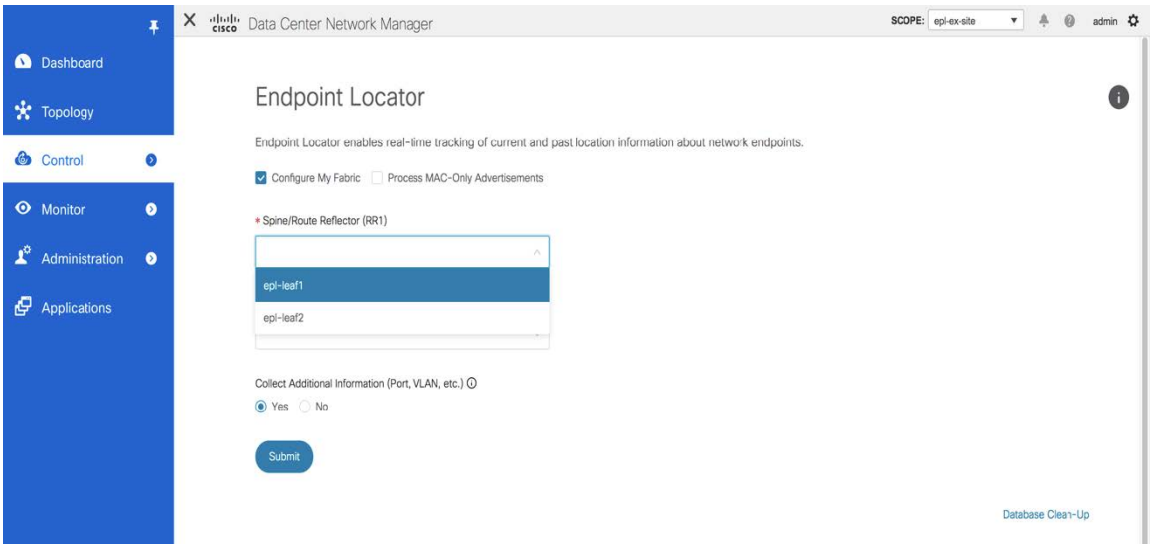

By default, the **Configure My Fabric** option is selected. This knob controls whether BGP configuration will be pushed to the selected spines/RRs as part of the enablement of the EPL feature. If the spine/RR needs to be configured manually with a custom policy for the EPL BGP neighborship, then this option should be unchecked. For external fabrics that are only monitored and not configured by DCNM, this option is greyed out as these fabrics are not configured by DCNM.

Select the **Process MAC-Only Advertisements** option to enable processing of MAC-Only advertisements while configuring the EPL feature.

Note If EPL is enabled on a fabric with or without selecting the **Process Mac-Only Advertisements** checkbox and you want to toggle this selection later, then you have to first disable EPL and then click **Database Clean-up** to delete endpoint data before re-enabling EPL with the desired **Process Mac-Only Advertisements** setting.

Select **Yes** under **Collect Additional Information** to enable collection of additional information such as PORT, VLAN, VRF etc. while enabling the EPL feature. To gather additional information, NX-API must be supported and enabled on the switches, ToRs, and leafs. If the **No** option is selected, this information will not be collected and reported by EPL.

Note For all fabrics except external fabrics, NX-API is enabled by default. For external fabrics, you have to enable NX-API in the external fabric settings by selecting the **Enable NX-API** checkbox in the **Advanced** tab of the External_Fabric_11_1 fabric template.

You can also watch the video that demonstrates how to configure EPL using Cisco DCNM. See Configuring Endpoint Locator.

Starting from Cisco DCNM Release 11.4(1), click the **i** icon to view a template of the configuration that is pushed to the switches while enabling EPL. This configuration can be copied and pasted on spines or border gateway devices to enable EPL on external monitored fabrics.

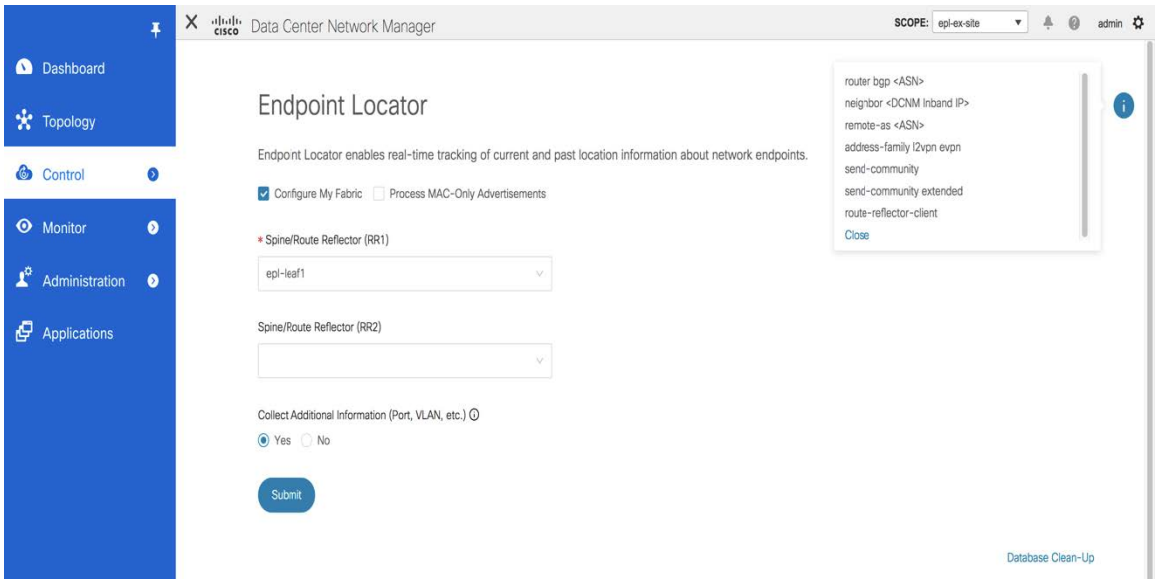

Once the appropriate selections are made and various inputs have been reviewed, click **Submit** to enable EPL. If there are any errors while you enable EPL, the enable process aborts and the appropriate error message is displayed. Otherwise, EPL is successfully enabled.

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. 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.

Fore more information about the EPL dashboard, refer Monitoring Endpoint Locator.

Enabling High Availability

Consider a scenario in which EPL is enabled on a DCNM deployment that is in non-HA mode and then, DCNM is moved to HA-mode. In such scenarios, the **Enable HA** toggle appears on the **Endpoint Locator** window. Toggle the **Enable HA** knob to enable high availability sync between primary and secondary DCNM.

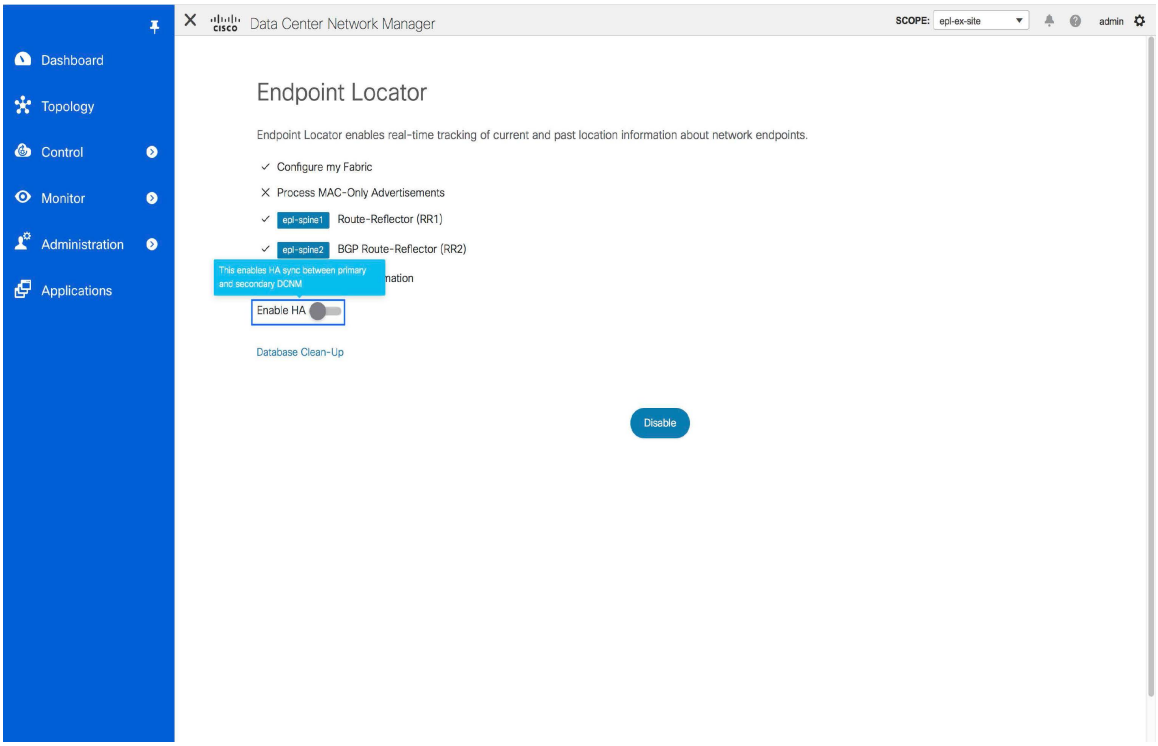

To enable high availability sync from the Cisco DCNM Web UI, perform the following steps:

Procedure

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

Step 2 Toggle the **Enable HA** button.

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. Starting from Cisco DCNM Release 11.4(1), you can flush the endpoint database even if you have not re-enabled the EPL feature on a fabric on which the EPL feature was previously disabled.

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 **Database Clean-Up**.

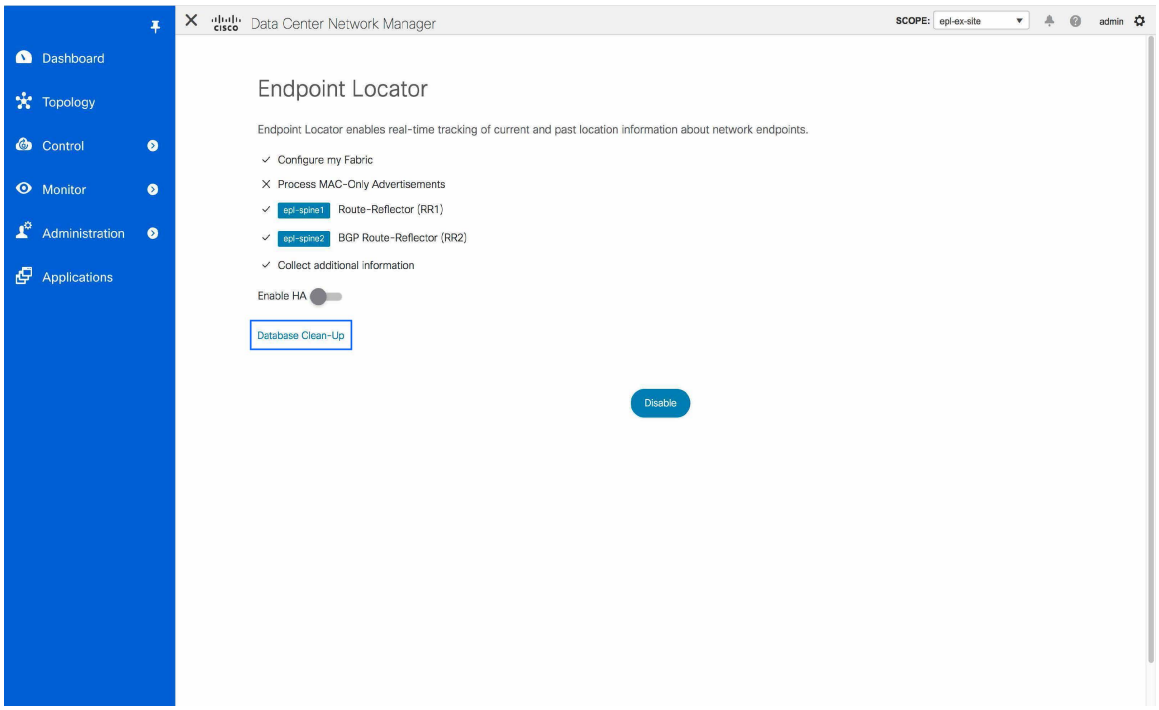

A warning is displayed with a message indicating that all the endpoint information that is stored in the database will be flushed.

Step 2 Click **Delete** to continue or **Cancel** to abort.

Configuring Endpoint Locator in DCNM High Availability Mode

Note For configuring EPL in native HA mode, you must add 2 neighbors to EPL. EPL IP being DCNM Primary eth2 and DCNM Secondary eth2 address respectively.

For production deployments, a native HA pair of DCNM nodes is recommended. Since the DCNM active and standby nodes need to be Layer-2 adjacent, their respective eth2 interfaces should be part of the same IP subnet or vlan. In addition, both DCNM nodes should be configured with the same eth2 gateway. The recommended option is to connect the DCNM active and standby nodes to a vPC pair of nexus switches (they may be leafs) so that there is enough fault-tolerance in case of single link failure, single device or a single DCNM node failure.

The following example shows a sample output for the **appmgr update network-properties** command for a Cisco DCNM Native HA Appliance. In this example, 1.1.1.2 is the primary eth2 interface IP address, 1.1.1.3 is the standby eth2 interface IP address, 1.1.1.1 is the default gateway and 1.1.1.4 is the virtual IP (VIP) for inband.

On Cisco DCNM Primary appliance:

```
appmgr update network-properties session start
appmgr update network-properties set ipv4 eth2 1.1.1.2 255.255.255.0 1.1.1.1
appmgr update network-properties set ipv4 peer2 1.1.1.3
appmgr update network-properties set ipv4 vip2 1.1.1.4 255.255.255.0
appmgr update network-properties session apply
appmgr update ssh-peer-trust
```

On Cisco DCNM Secondary appliance:

```
appmgr update network-properties session start
appmgr update network-properties set ipv4 eth2 1.1.1.3 255.255.255.0 1.1.1.1
appmgr update network-properties set ipv4 peer2 1.1.1.2
appmgr update network-properties set ipv4 vip2 1.1.1.4 255.255.255.0
appmgr update network-properties session apply
appmgr update ssh-peer-trust
```

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

Procedure

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

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

Step 2 Select a fabric from the **SCOPE** dropdown list to configure endpoint locator in DCNM HA mode.

Step 3 Select the Route-Reflectors (RRs) from the drop-down lists.

Step 4 Select **Yes** under **Collect Additional Information** to enable collection of additional information such as PORT, VLAN, VRF etc. while enabling the EPL feature. If the No option is selected, this information will not be collected and reported by EPL.

Step 5 Click **Submit**.

What to do next

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

Configuring Endpoint Locator in DCNM Cluster Mode

Note For configuring EPL in cluster mode, you must add a single neighbor to EPL. DCNM EPL container Inband IP address is EPL IP.

With the DCNM cluster mode deployment, in addition to the DCNM nodes, an additional 3 compute nodes are present in the deployment. For information about deploying applications in cluster mode, see *Cisco DCNM in Clustered Mode*.

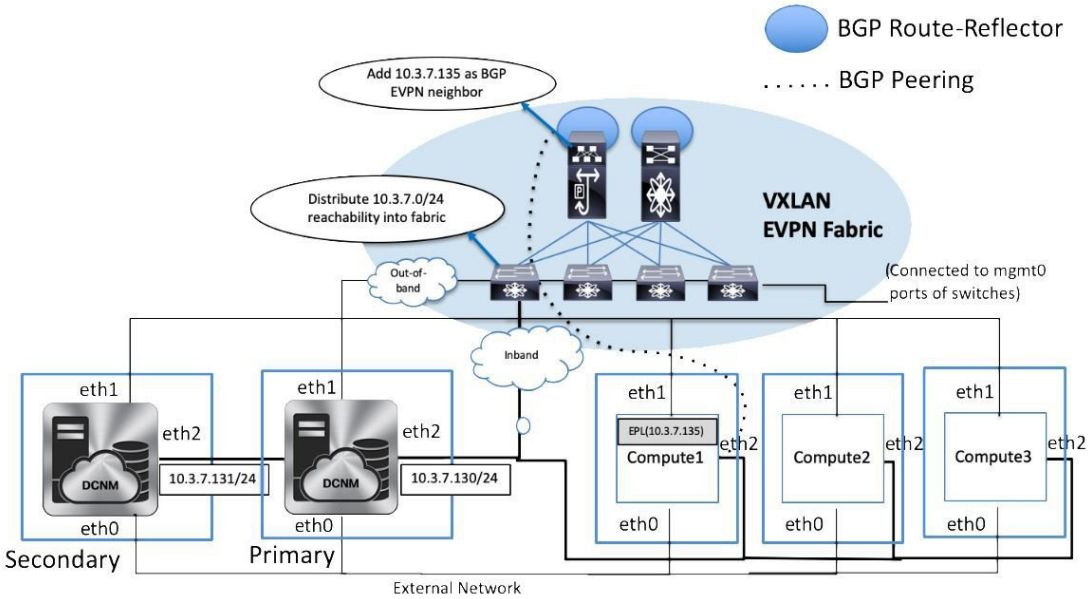

In DCNM Cluster mode, all applications including EPL run on the compute nodes. The DCNM application framework takes care of the complete life cycle management of all applications that run on the compute nodes. The EPL instance runs as a container that has its own IP address allocated out of the inband pool assigned to the compute nodes. This IP address will be in the same IP subnet as the one allocated to the eth2 or inband interface. Using this IP address, the EPL instance forms a BGP peering with the spines/RRs when the EPL feature is enabled. If a compute node hosting the EPL instance will go down, the EPL instance will be automatically respawned on one of the remaining 2 compute nodes. All IP addresses and other properties associated with the EPL instance are retained.

The Layer-2 adjacency requirement of the compute nodes dictates that the compute node eth2 interfaces should be part of the same IP subnet as the DCNM nodes. Again, in this case, connecting the compute nodes to the same vPC pair of switches is the recommended deployment option. Note that for cluster mode DCNM OVA setups, ensure that promiscuous mode is enabled in the port group corresponding to eth2 interface in order to establish inband connectivity as depicted below:

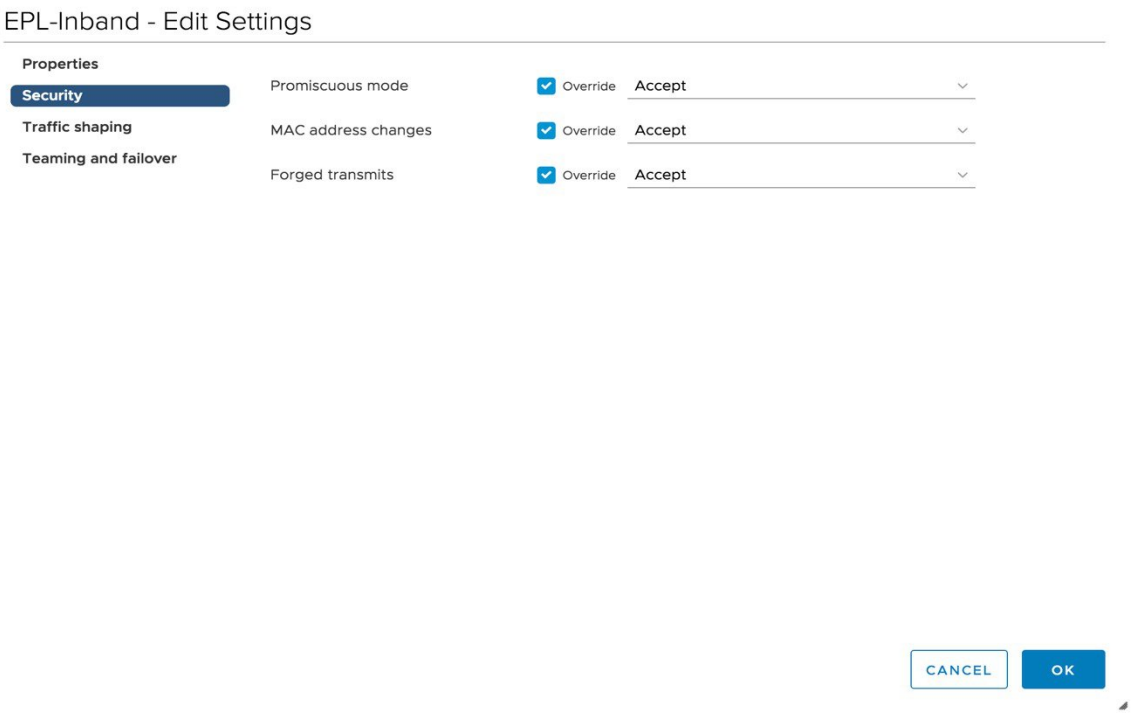

The enablement of the EPL feature for DCNM cluster mode is identical to that in the non-cluster mode. The main difference is that on the spine/RRs, only a single BGP neighborship is required that points to the IP address allocated to the EPL instance. Recall that for the DCNM native HA deployment in the non-cluster mode, all spines/RRs always had 2 configured BGP neighbors, one pointing to the DCNM primary eth2 interface and other one pointing to the DCNM secondary eth2 interface. However, only one neighbor would be active at any given time.

Configuring Endpoint Locator for External Fabrics

In addition to Easy fabrics, DCNM Release 11.2(1) allows you to enable EPL for VXLAN EVPN fabrics comprising of switches that are imported into the external fabric. The external fabric can be in managed mode or monitored mode, based on the selection of **Fabric Monitor Mode** flag in the **External Fabric** Settings. For external fabrics that are only monitored and not configured by DCNM, this flag is disabled. Therefore, you must configure BGP sessions on the Spine(s) via OOB or using the CLI. To check the sample template, click icon to view the configurations required while enabling EPL.

In case the **Fabric Monitor Mode** checkbox in the External Fabric settings is unchecked, then EPL can still configure the spines/RRs with the default **Configure my fabric** option. However, disabling EPL would wipe out the router bgp config block on the spines/RRs. To prevent this, the BGP policies must be manually created and pushed onto the selected spines/RRs.

Configuring Endpoint Locator for eBGP EVPN Fabrics

From Cisco DCNM Release 11.2(1), you can enable EPL for VXLAN EVPN fabrics, where eBGP is employed as the underlay routing protocol. Note that with an eBGP EVPN fabric deployment, there is no traditional RR similar to iBGP. The reachability of the in-band subnet must be advertised to the spines that behave as Route

Servers. To configure EPL for eBGP EVPN fabrics from the Cisco DCNM Web UI, perform the following steps:

Procedure

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

Select the fabric to configure eBGP on or create eBGP fabric with the **Easy_Fabric_eBGP** template.

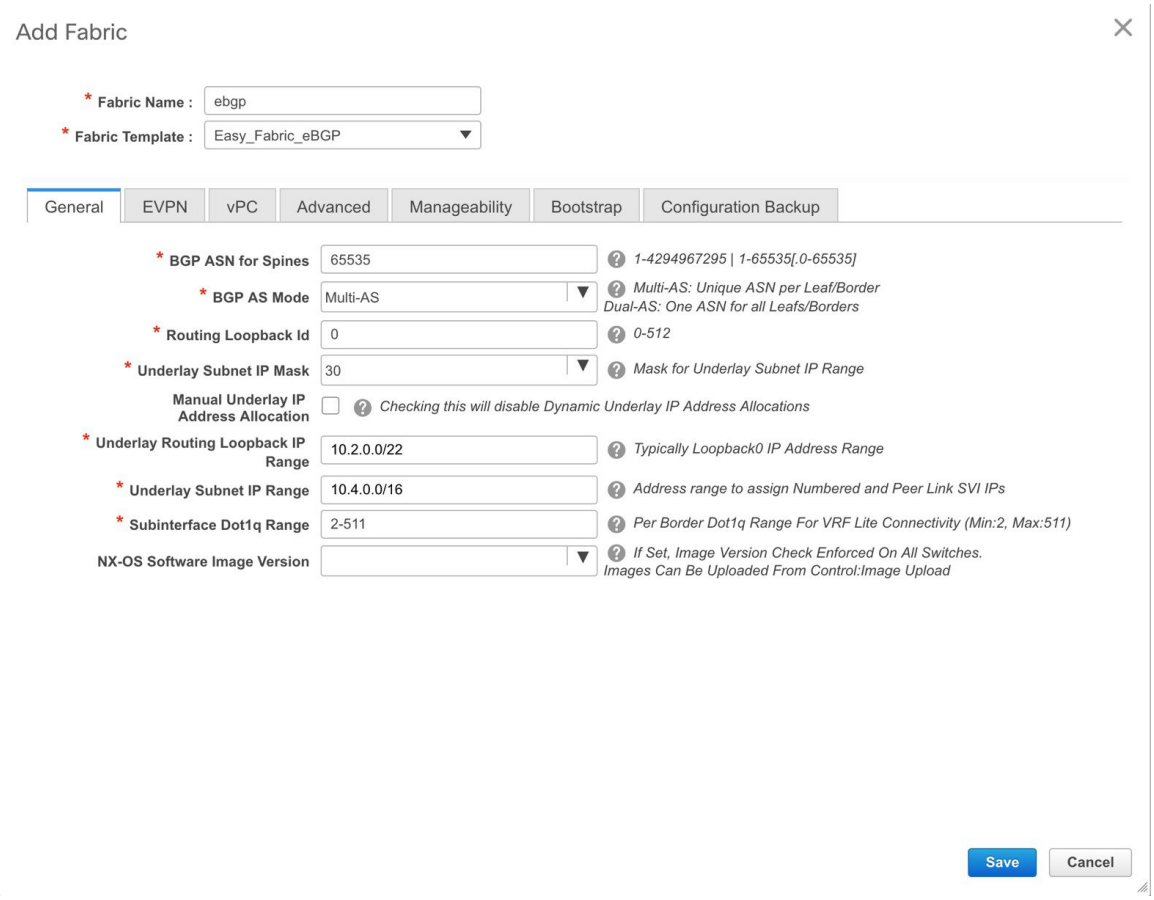

Step 2 Use the **leaf_bgp_asn** policy to configure unique ASNs on all leaves.

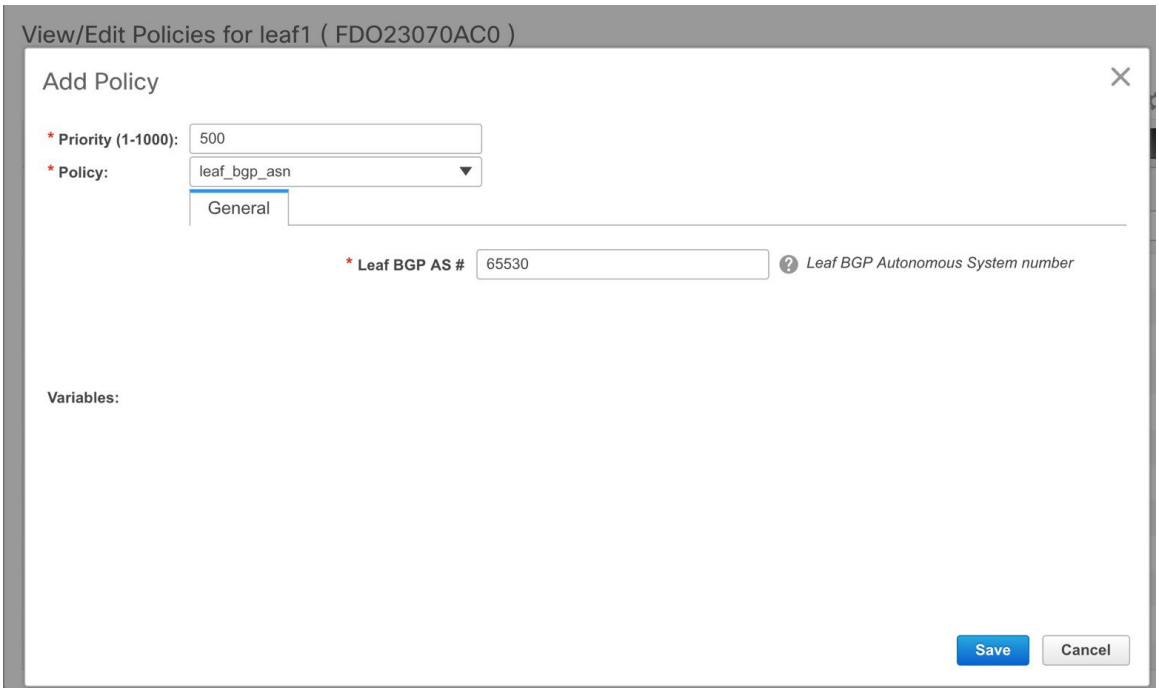

Step 3 Add the **ebgp_overlay_leaf_all_neighbor** policy to each leaf.

Fill **Spine IP List** with the spines' BGP interface IP addresses, typically the loopback0 IP addresses.

Fill **BGP Update-Source Interface** with the leaf's BGP interface, typically loopback0.

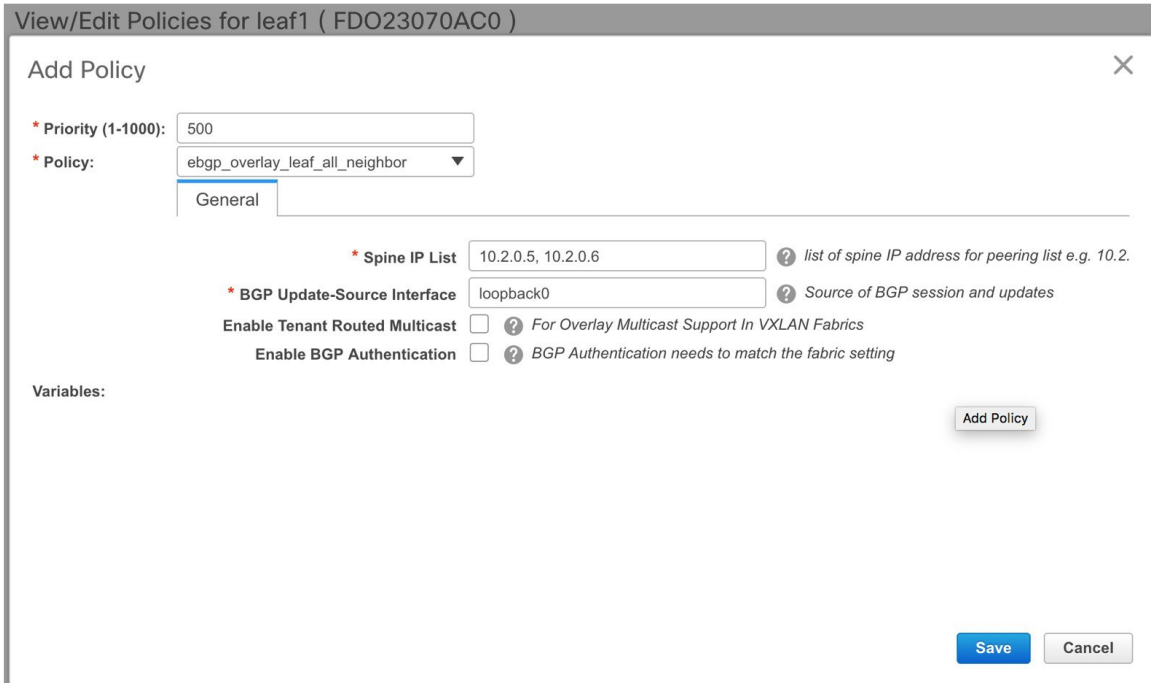

Step 4 Add the **ebgp_overlay_spine_all_neighbor** policy to each spine.

Fill **Leaf IP List** with the leaves' BGP interface IPs, typically the loopback0 IPs.

Fill **Leaf BGP ASN** with the leaves' ASNs in the same order as in **Leaf IP List**.

Fill **BGP Update-Source Interface** with the spine's BGP interface, typically loopback0.

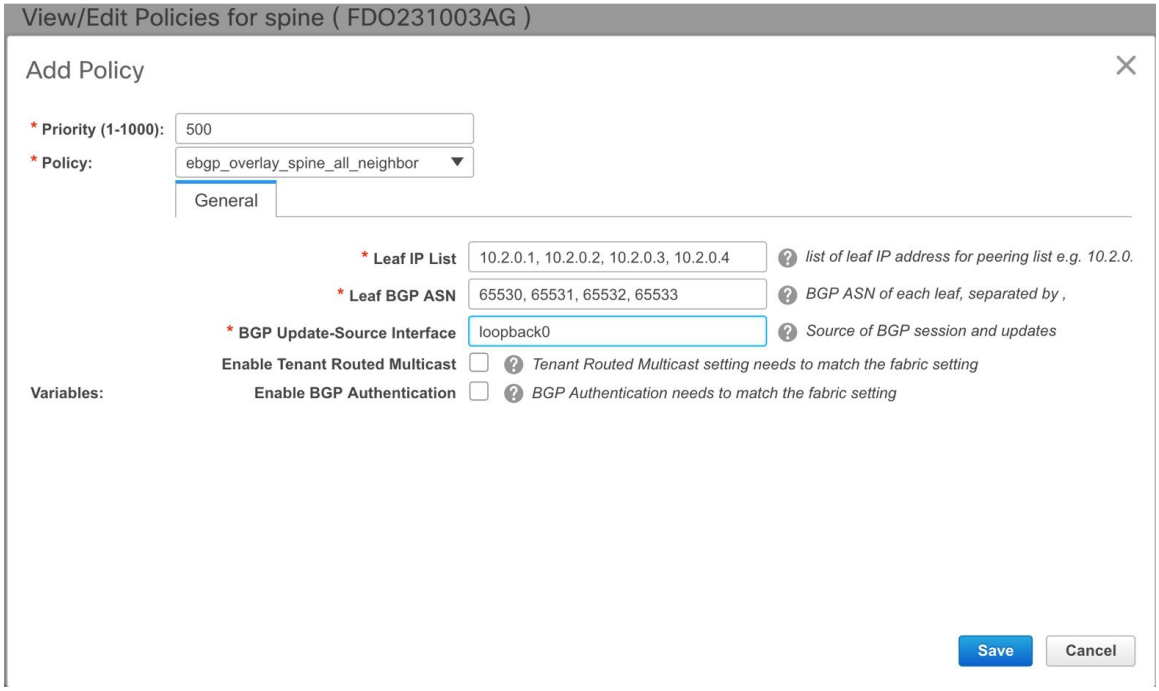

After the in-band connectivity is established, the enablement of the EPL feature remains identical to what is listed so far. EPL becomes a iBGP neighbor to the Route Servers running on the spines.

EPL Connectivity Options

Sample topologies for the varios EPL connectivity options are as given below.

Cisco DCNM supports the following web browsers:

DCNM Cluster Mode: Physical Server to VM Mapping

We recommend a minimum of 3 physical servers, or a maximum of 5 physical servers in which each DCNM and compute is located on an individual physical server.

Figure 4: A minimum of 3 physical servers

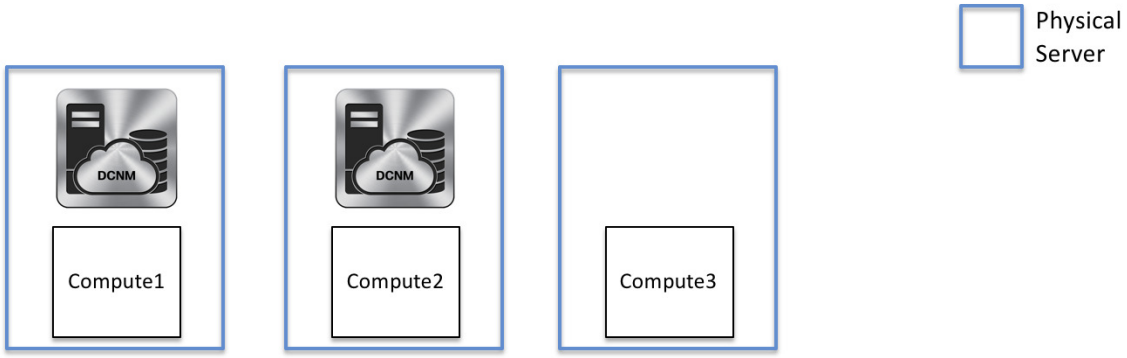

Figure 5: A maximum of 5 physical servers

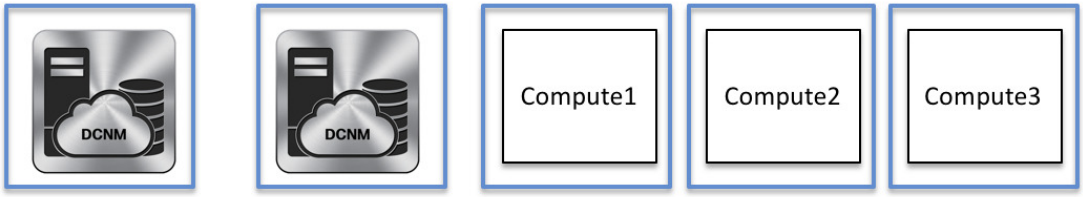

DCNM/Compute VM Physical Connectivity

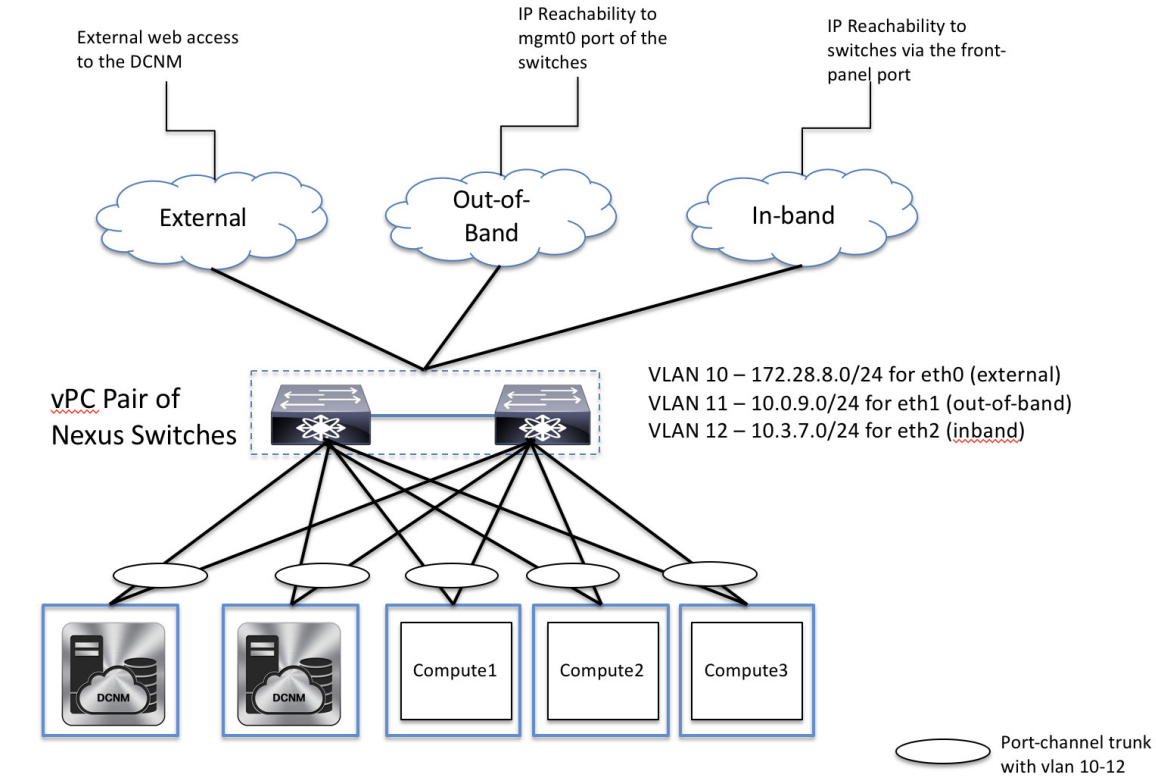

DCNM Cluster Mode

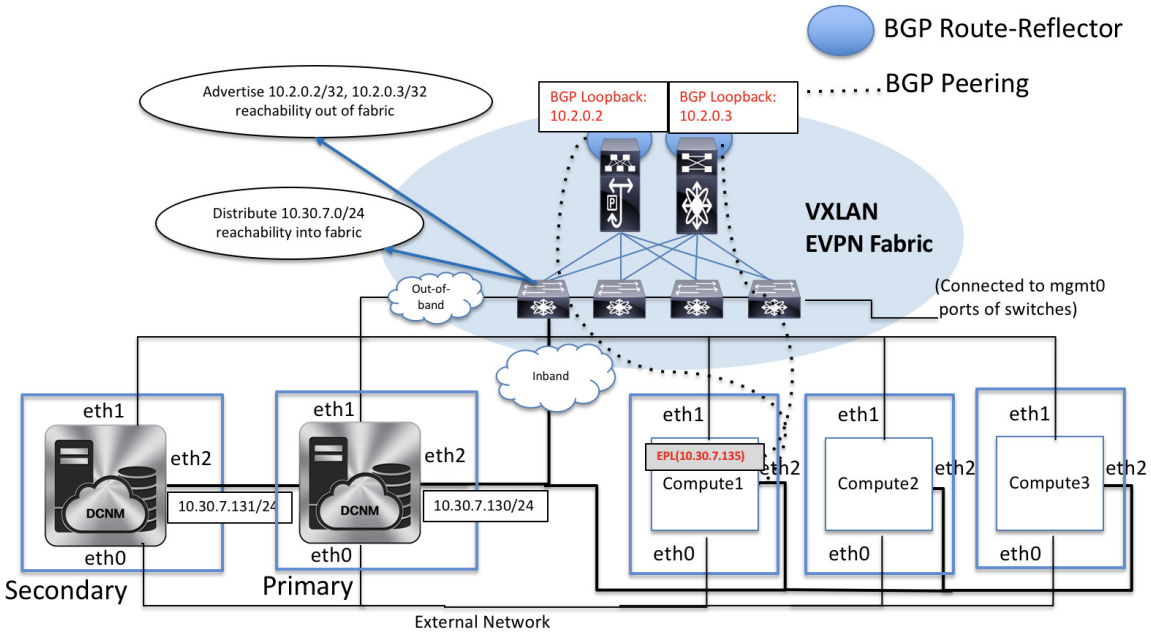

DCNM Multi-Fabric Connectivity

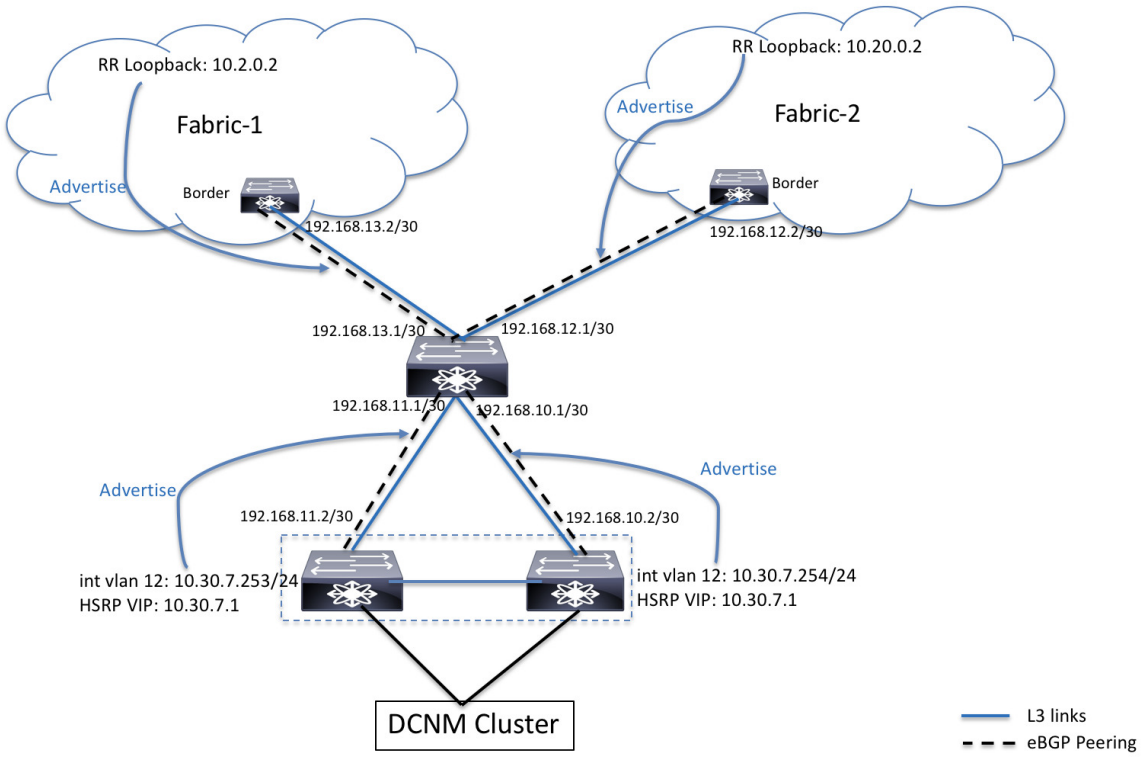

EPL Connectivity for Native HA

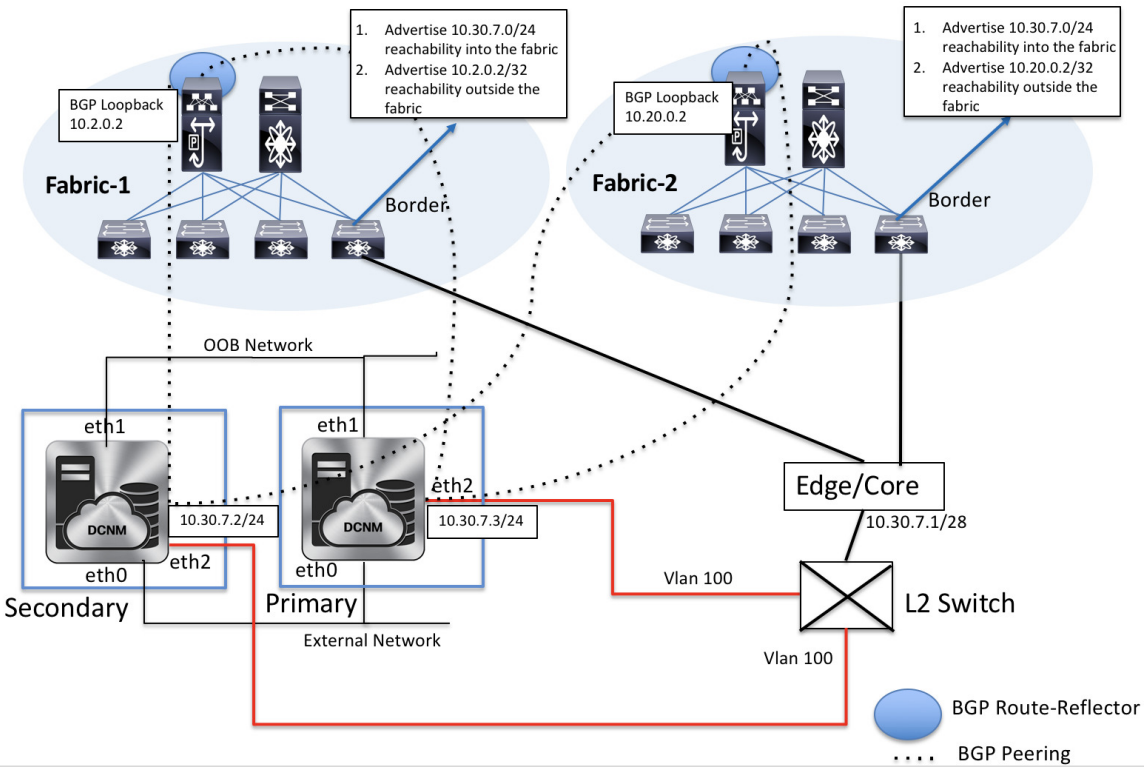

Disabling Endpoint Locator

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. Select the required fabric from the **SCOPE** dropdown list. The fabric configuration details are then displayed for the selected fabric.

Step 2 Click **Disable**.

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 log that provides details on what occurred when the EPL feature is enabled or disabled, is present in the file epl.log at the location: /usr/local/cisco/dcm/fm/logs/epl.log. The following example provides a snapshot of the epl.log that shows the EPL configuration progress for a fabric.

```
2019.12.05 12:18:23  INFO  [epl] Found DCNM Active Inband IP: 192.168.94.55/24
2019.12.05 12:18:23  INFO  [epl] Running script: [sudo, /sbin/appmgr, setup, inband-route,
 --host, 11.2.0.4]
2019.12.05 12:18:23  INFO  [epl] Getting EPL configure progress for fabric 4
2019.12.05 12:18:23  INFO  [epl] EPL Progress 2
2019.12.05 12:18:23  INFO  [epl] [sudo, /sbin/appmgr, setup, inband-route, --host, 11.2.0.4]
 command executed, any errors? No
2019.12.05 12:18:23  INFO  [epl] Received response:
2019.12.05 12:18:23  INFO  [epl] Validating host route input
2019.12.05 12:18:23  INFO  [epl] Done configuring host route
2019.12.05 12:18:23  INFO  [epl] Done.
2019.12.05 12:18:23  INFO  [epl] Running script: [sudo, /sbin/appmgr, setup, inband-route,
 --host, 11.2.0.5]
2019.12.05 12:18:23  INFO  [epl] [sudo, /sbin/appmgr, setup, inband-route, --host, 11.2.0.5]
 command executed, any errors? No
2019.12.05 12:18:23  INFO  [epl] Received response:
2019.12.05 12:18:23  INFO  [epl] Validating host route input
2019.12.05 12:18:23  INFO  [epl] Done configuring host route
2019.12.05 12:18:23  INFO  [epl] Done.
2019.12.05 12:18:23  INFO  [epl] Running command: sudo /sbin/appmgr show inband
2019.12.05 12:18:24  INFO  [epl] Received response: Physical IP=192.168.94.55/24
Inband GW=192.168.94.1
No IPv6 Inband GW found

2019.12.05 12:18:26  INFO  [epl] Call:
http://localhost:35000/afw/apps?imagetag=cisco:epl:2.0&fabricid=epl-ex-site, Received
response:
[illegible]
2019.12.05 12:18:26  INFO  [epl] Epl started on AFW
```

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 100 MB.

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.

The endpoint data displayed on the dashboard may be slightly inaccurate in a large-scale setup. An approximately 1% accuracy tradeoff is made at higher endpoint counts for performance. If the dashboard greatly differs from what is expected, the validity can be checked with a verifier script that is packaged in DCNM. As root, run the epl-rt-2.py script in /root/packaged-files/scripts/. This script needs the RR/spine IP and the associated username and password. Note that the /root/packaged-files/scripts/ directory is read only, so the script needs to be run outside that directory. For example, to run the script for a spine with IP 10.2.0.5, username admin, and password cisco123, run **/root/packaged-files/scripts/epl-rt-2.py -s 10.2.0.5 -u admin -p cisco123** while the working directory is /root/. If the EPL dashboard still does not display expected numbers and the epl-rt-2.py script output differs significantly from the dashboard, please contact tech support.

In cluster mode, BGP is not established between the spines/RRs and DCNM. Check that the **Promiscuous mode** setting for the port group corresponding to the eth2 DCNM interface is set to **Accept**. If a connection is still not established, perform the following steps to check the connectivity between DCNM's BGP client and the spine/RR:

1. Open a shell on the active DCNM and run the following commands:
   a. *docker service ls*

      *Note the ID for the EPL service

   b. *docker service ps $ID*

      *Note the NODE field

   c. *afw compute list -b*

      *Note the HostIp matching the HostName (NODE) from before. This is the compute that the EPL service is currently running on.

2. Open a shell on the compute noted from Step 1 - c and run the following commands:
   a. *docker container ls*

      Note the CONTAINER ID for EPL. If there are multiple EPL containers check the container name to see which one corresponds to which fabric. The naming scheme is epl_cisco_$FabricName_afw.

   b. *docker container inspect $CONTAINER_ID*

      *Note the value of SandboxKey

   c. *nsenter --net=$SandboxKey*

      This command enters the network namespace of the EPL container. Now network commands such as ifconfig, ip, and ping will act as if they're being ran from inside the container until "exit" is issued in the shell.

3. Try pinging the spine/RR. Make sure that the Inband IP Pool that the DCNM cluster is configured with does not conflict with any switch loopback IPs.

EPL with ISE Policies

Consider a scenario in which AAA configurations are configured on switches running Cisco NX-OS Release 9.3(4) or earlier releases. A sample AAA switch configuration is given below.

```
feature tacacs+
tacacs-server host ISE_ACS_IP_ADDDRESS 5 key 7 "Fewhg12345"
aaa group server tacacs+ AAA_TACACS
       server ISE_ACS_IP_ADDDRESS
       use-vrf management
       source-interface mgmt0
aaa authentication login default group AAA_TACACS local
aaa authentication login console local
aaa authorization config-commands default group AAA_TACACS local
aaa authorization commands default group AAA_TACACS local
aaa accounting default group AAA_TACACS
aaa authentication login error-enable
```

The ISE server is configured such that the **guestshell**, **run guestshell**, and **show** commands, are permitted to reach the discovery account or policies that are created in the ISE. The permitted commands are set in the **TACACS Command Sets** window under the **Policy Elements** tab in the ISE.

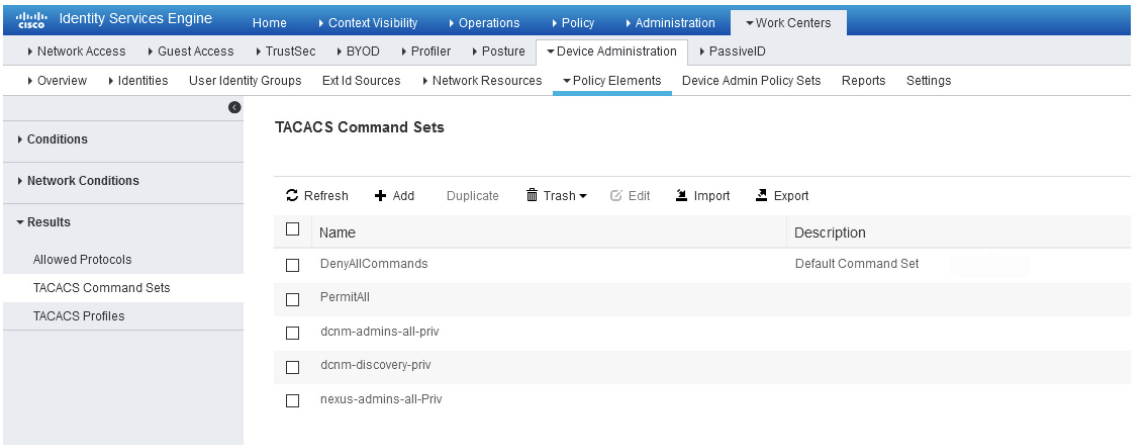

The eth0 IP of DCNM and the subnet for the fabric devices are also allowed. This is configured in the **Device Admin Policy Sets** window under the **Device Administration** tab.

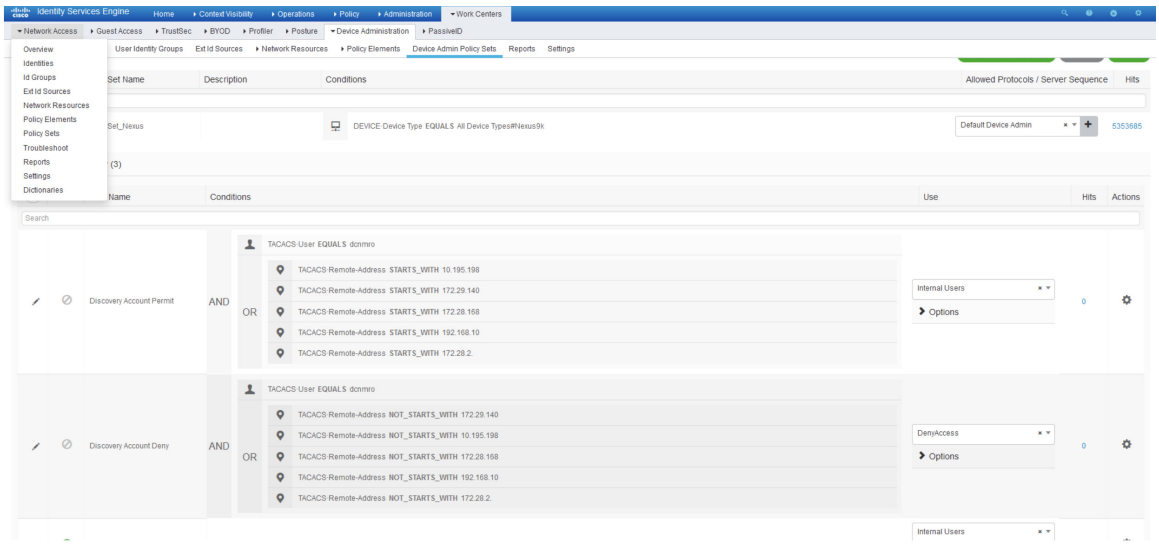

Now, DCNM is configured to use the discovery account to run all the **show** commands that are required for the Endpoint Locator feature. However, due to an issue with the switch NXAPI, AAA validation fails as the

requestor IP is not populated in the remote AAA authorization requests. Since the **show** commands are not seen as originating from an IP address, the commands are blocked which prevents the EPL dashboard from displaying the required endpoint information.

As a workaround, we recommend relaxing AAA rules and allowing requests from "blank" senders. To allow requests from “blank” senders, click the icon under the **Status** column for both **Discovery Account Permit** and **Discovery Account Deny** in the **Device Admin Policy Sets** window, choose **Disabled**, and click **Save**.

Also, this issue is not seen on switches running Cisco NXOS Release 9.3(5) and later releases.

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.

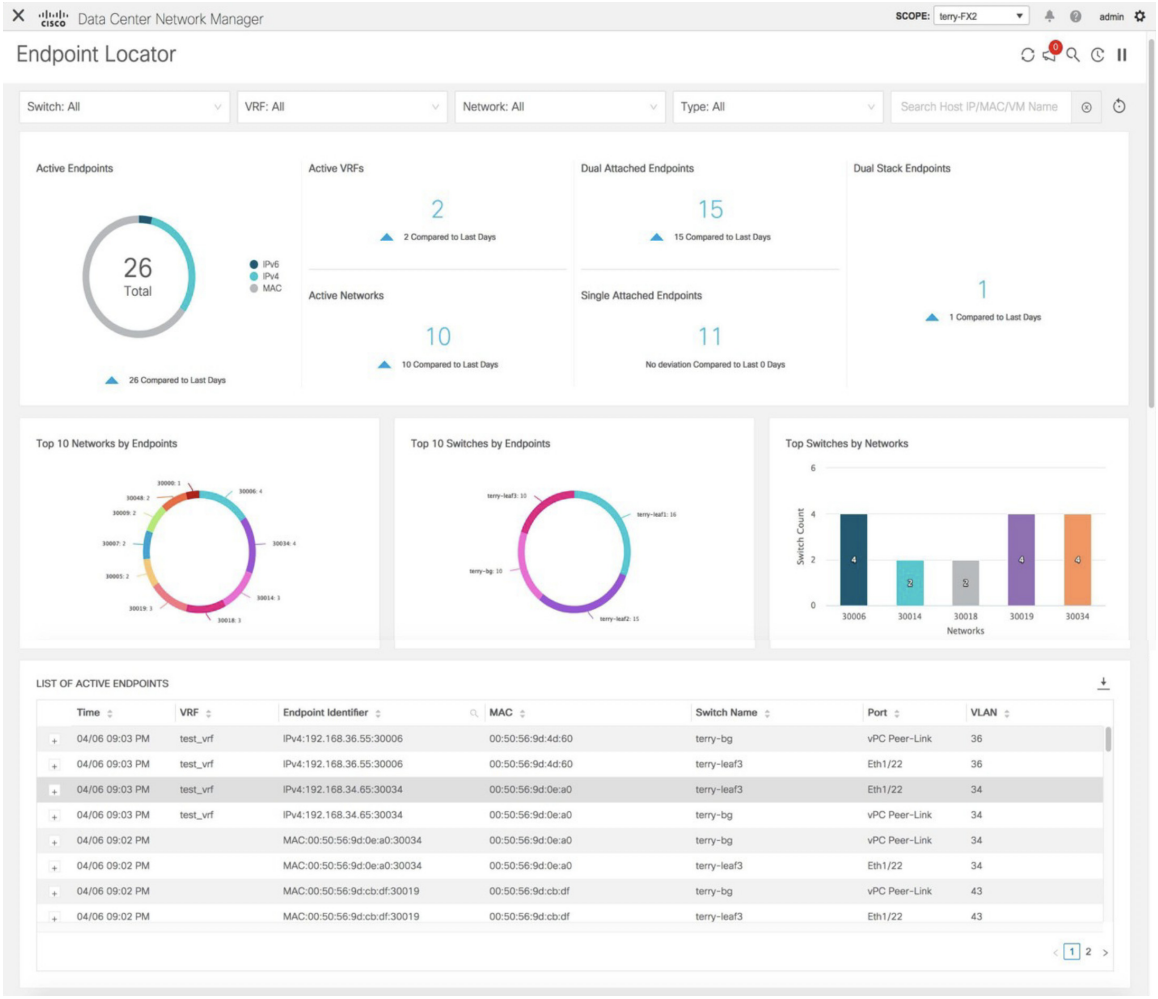

Note Due to an increase in scale from Cisco DCNM Release 11.3(1), the system may take some time to collect endpoint data and display it on the dashboard. Also, on bulk addition or removal of endpoints, the endpoint information displayed on the EPL dashboard takes a few minutes to refresh and display 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. Starting from Cisco DCNM Release 11.4(1), the name of the network is also displayed in the **Network** drop-down list. By default, the selected option is **All** for these fields. You can also display endpoint data for a specific device by entering the host IP address, MAC address, or the name of the virtual machine in the **Search Host IP/MAC/VM Name** field.

Note You can initiate a search by using the available options from the dropdown lists or by using the **Search Host IP/MAC/VM Name** field. You cannot initiate a search by using a combination of dropdown lists and the search field.

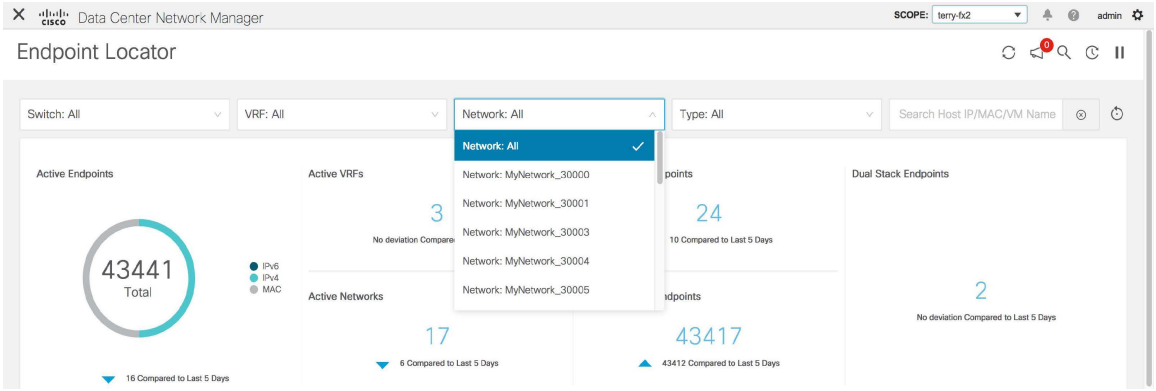

You can reset the filters to the default options by clicking the **Reset Filters** icon.

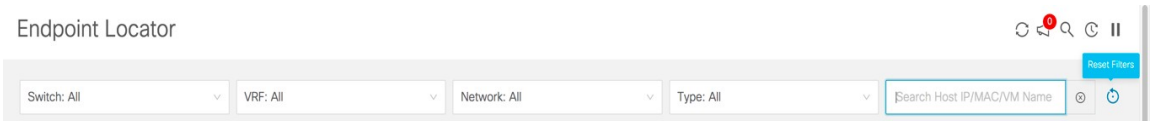

The 'top pane' of the window displays the number of active endpoints, active VRFs, active networks, dual attached endpoints, single attached endpoints and dual stacked endpoints, for the selected scope. Support for displaying the number of dual attached endpoints, single attached endpoints and dual stacked 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.

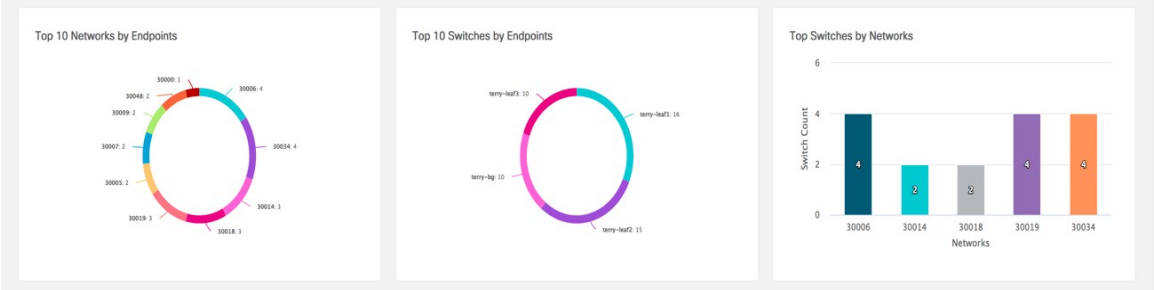

The 'bottom pane' of the window displays the list of active endpoints.

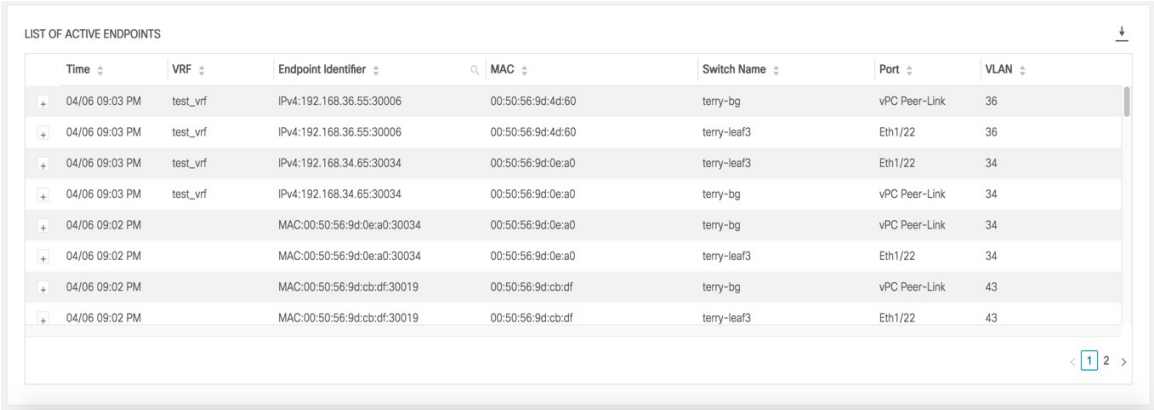

Click + to display more information for a specific endpoint. If a virtual machine has been configured, the name of the VM is displayed in the **Node Name** field. Note that it can take up to 15 minutes for the name of the VM to be reflected in the EPL dashboard. Until then, the EPL dashboard displays **No DATA** in the **Node Name** field.

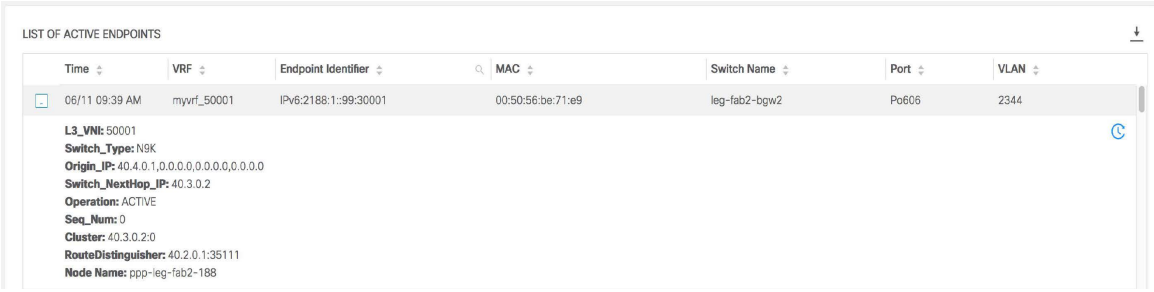

Click the **Host Life** icon to display the **Endpoint Life** window for that endpoint.

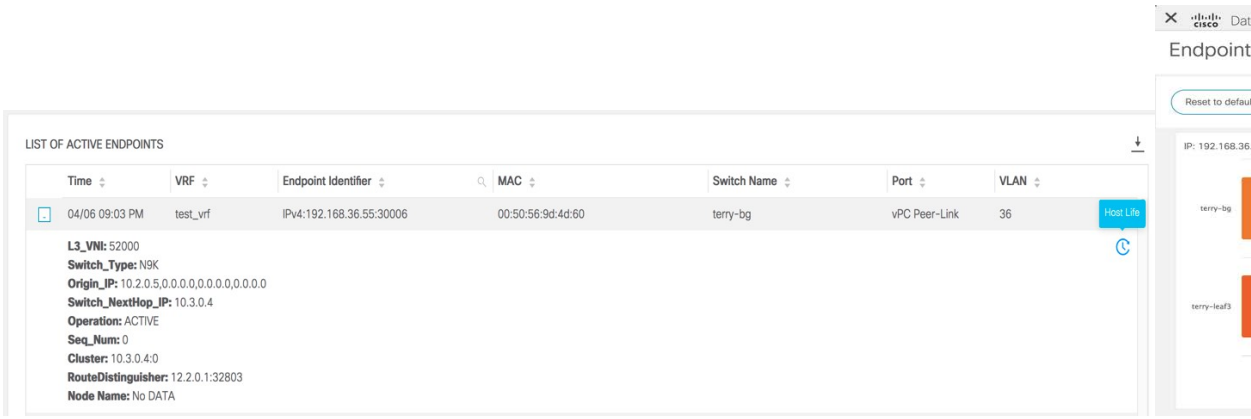

Click the search icon in the **Endpoint Identifier** column to search for specific IP addresses.

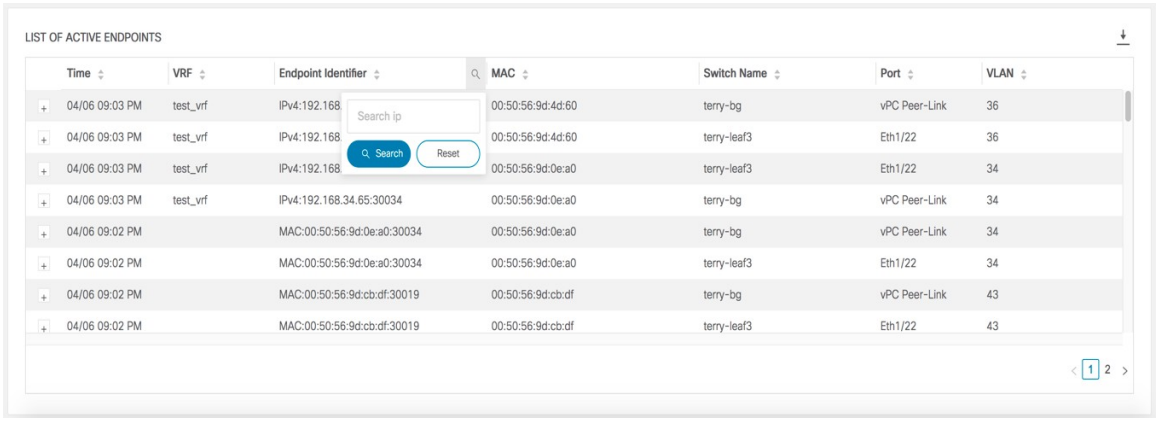

In certain scenarios, the datapoint database may go out-of-sync and information, such as the number of endpoints, may not be displayed correctly due to network issues such as -

- Endpoint moves under the same switch between ports and the port information needs some time to be updated.
- An orphan endpoint is attached to the second VPC switch and is no longer an orphan endpoint.
- NX-API not enabled initially and then enabled at a later point in time.
- NX-API failing initially due to misconfiguration.
- Change in Route Reflector (RR).
- Management IPs of the switches are updated.

In such cases, clicking the **Resync** icon leads to the dashboard syncing to the data currently in the RR. However, historical data is preserved. We recommend not clicking **Resync** multiple times as this is a compute-intense activity.

Click the **Notifications** icon to display a list of the most recent notifications.

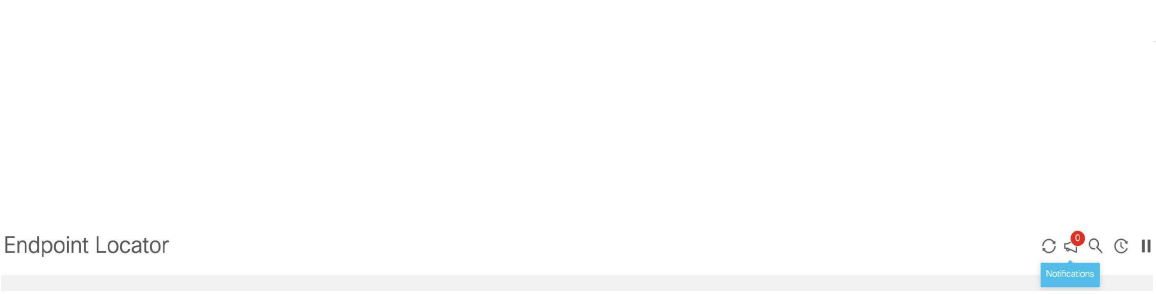

Information such as the time at which the notification was generated, the description of the notification, severity level, and the name of the node is displayed.

Notifications are generated for events such as duplicate IP addresses, duplicate MAC-Only addresses, VRF disappears from a fabric, all endpoints disappear from a switch, endpoint moves, endpoints on a fabric going to zero, when endpoints are attached to a switch, when a new VRF is detected, and when the RR BGP connectivity status changes. The RR connected status indicates that the DCNM can connect to the RR through BGP (DCNM and RR are BGP neighbors). The RR disconnected status indicates that the RR is disconnected and the underlying BGP is not functioning. Click the download icon to download the list of notifications as a CSV file.

Starting from Cisco DCNM Release 11.4(1), an alarm is generated if there are any endpoint-related anomalies. For more information on endpoint alarms, refer Endpoint Locator Alarms.

Click the **Pause** icon to temporarily stop the near real-time collection and display of data.

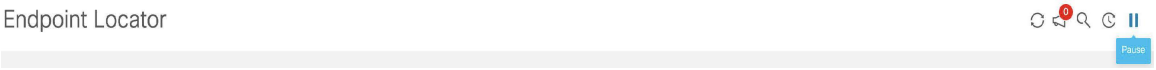

Consider a scenario in which EPL is first enabled and the **Process MAC-Only Advertisements** checkbox is selected. Then, EPL is disabled and enabled again without selecting the **Process MAC-Only Advertisements** checkbox. As the cache data in elasticsearch is not deleted on disabling of EPL, the MAC 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.

Hover over the graph at specific points to display more information. The points in the graph are plotted at 30-minute intervals. You can also display the graph for a specific requirement by clicking the color-coded points at the bottom of each graph. For example, click on all color-coded points other than **active (IPv4)** in the Active Endpoints window displayed above such that only **active (IPv4)** is highlighted and the other points are not highlighted. In such a scenario, only the active IPv4 endpoints are displayed on the graph. You can also hover over the color-coded points at the bottom of the graph to display the graph for a specific requirement. For example, hover over **active (IPv4)** to display only the active IPv4 endpoints on the graph.

Click on any point in the graph to display a window that has detailed information about that point of time. For example, click on a specific point in the **Active Endpoints** graph to display the **Endpoints** window. This window has information about the endpoints along with the name of the switch and the VRF associated with the endpoint. Click the download icon at the top right of the **Endpoints** window to download the data as a CSV file.

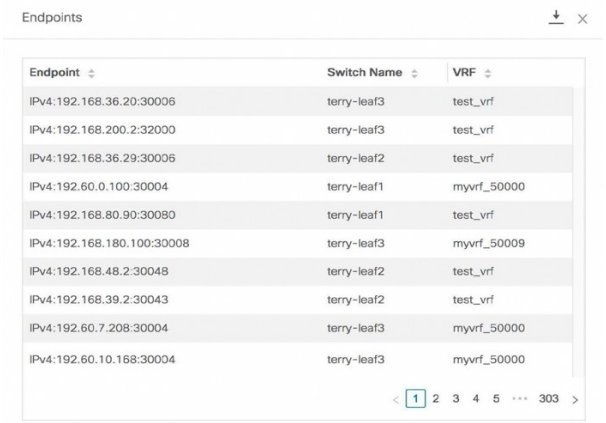

Endpoint Snapshots

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.

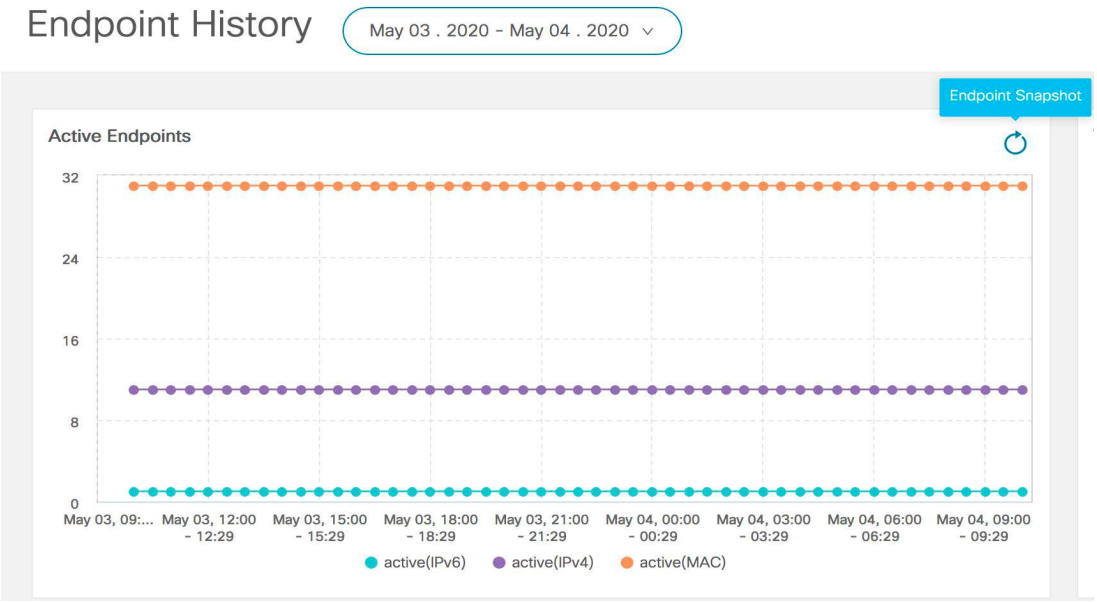

By default, endpoint snapshot comparison data for the previous hour is displayed.

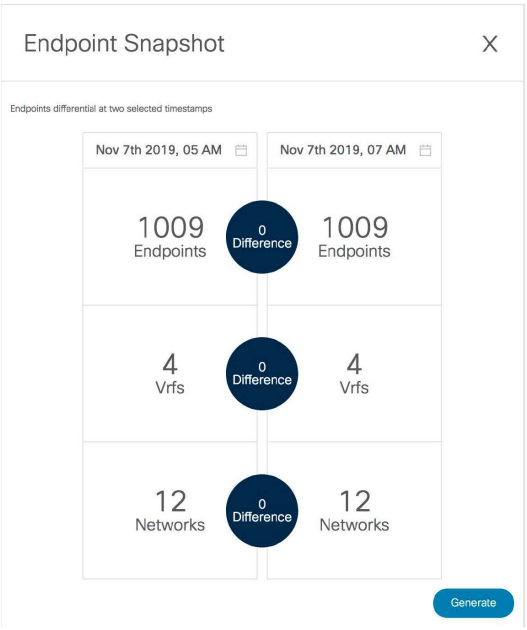

To compare endpoint snapshots at specific points in time, select two points in time, say T1 and T2, and click **Generate**.

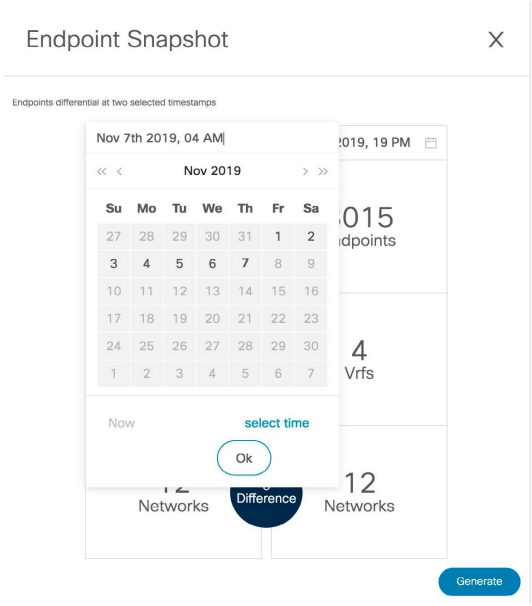

A comparison of the endpoints, VRFs, and networks at the selected points in time are displayed. Click each tile to download more information about the endpoints, VRFs, or networks. Click the **Difference** icon to download details about the differences in data for the specified time interval. Snapshots are stored for a maximum of three months and then discarded.

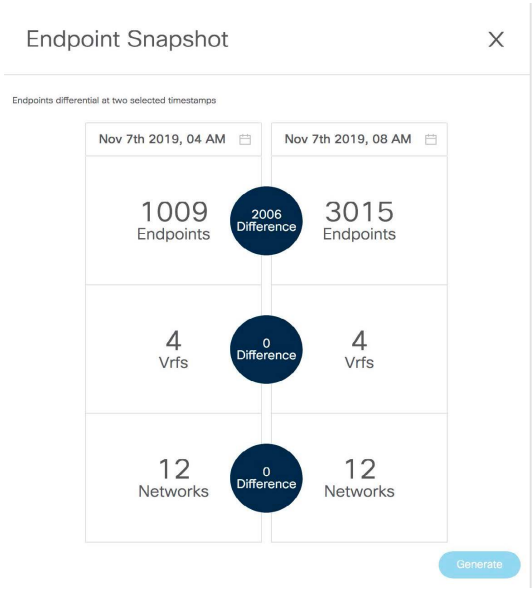

Endpoint Search

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.

The results displayed here are dependent on the fields listed under **Selected fields** located in the menu on the left. You can add any field listed under **Available fields** to **Selected fields** to initiate a search using the required fields.

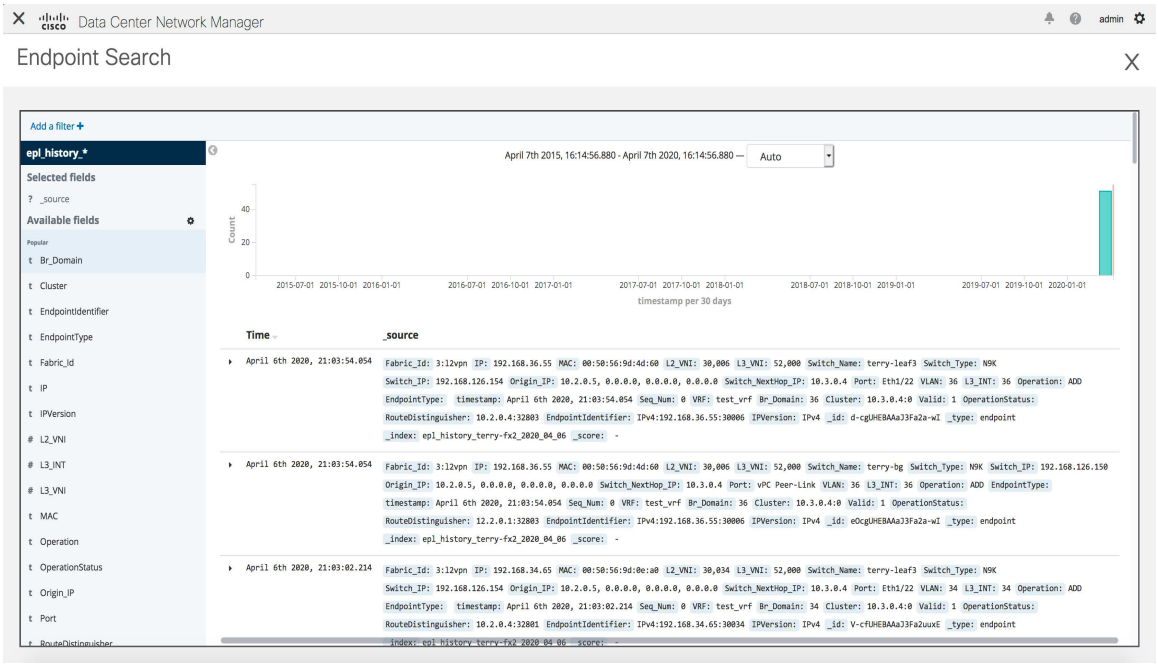

Endpoint Life

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**.

Initiate a search by using an IPv4 or IPv6 address to display the **Endpoint Life** graph for IPv4/IPv6 endpoints. Initiate a search by using a MAC address to display the **Endpoint Life** graph for MAC-Only endpoints.

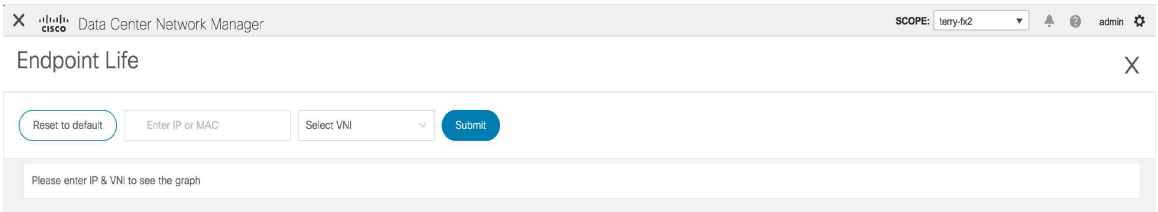

The window that is displayed is essentially the endpoint life of a specific endpoint. The bar that is orange in color represents the active endpoint on that switch. If the endpoint is viewed as active by the network, it will have a band here. If an endpoint is dual-homed, then there will be two horizontal 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.

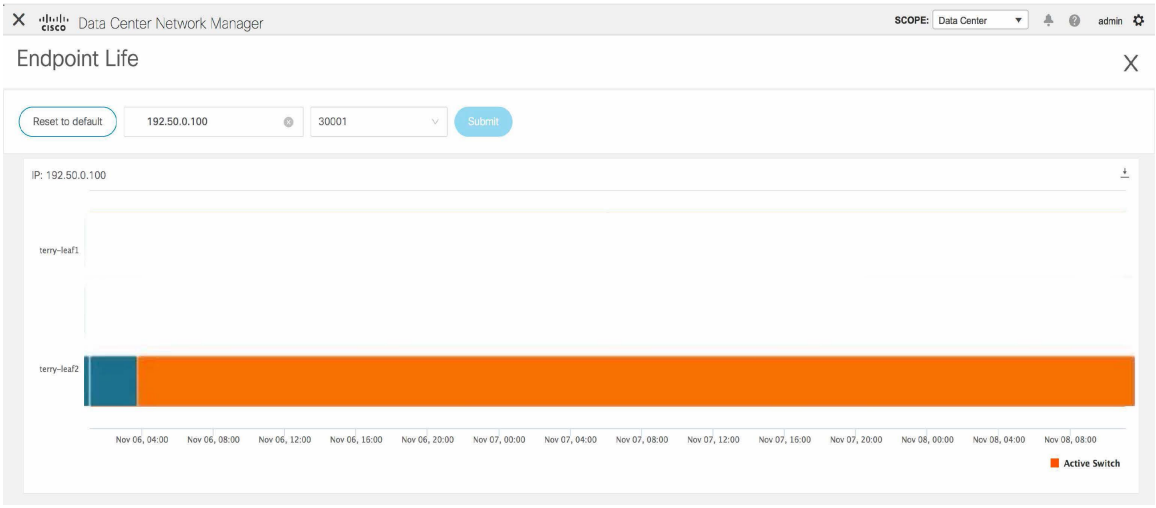

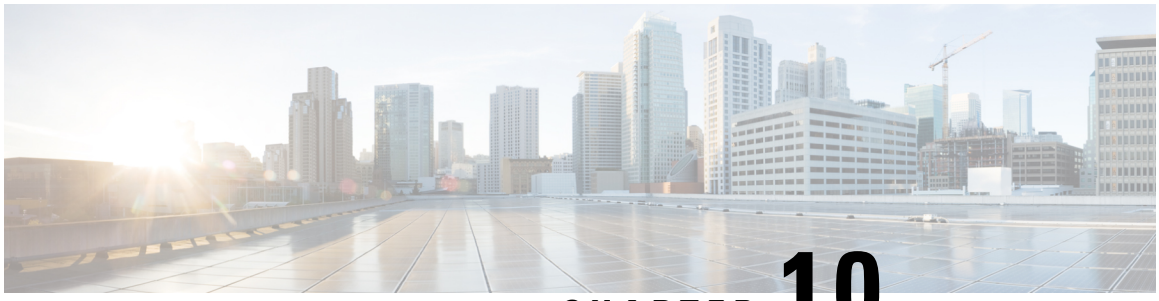

CHAPTER 10

IPAM Integrator

- Catalog, on page 595

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:

- Health Monitor (2.1)
- PTP Monitoring (1.1)
- Kibana (2.0)
- Programmable report (1.1.0)
- Elastic Service (1.1)
- Compliance (4.0.0)
- Debug Tools (2.1)
- IPAM Integrator (1.0)
- Endpoint Locator (2.1)
- Kubernetes Visualizer (1.1)
- vmmplugin (4.1)

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 551.

IPAM Integrator

From Cisco DCNM Release 11.4(1), you can use the IPAM Integrator application to view the IP allocation in IPAM server and relevant networks defined in DCNM. In DCNM 11.4(1), IPAM integration is with Infoblox.

The IPAM Integrator application in DCNM 11.4(1) allows read-only access to the IPAM and DCNM servers. Currently, IPv4 overlay DHCP is supported. In read-only access mode, IPAM records are retrieved and mapped to DCNM networks in Easy Fabric and eBGP VXLAN fabric. You can also choose to sync up records on-demand between DCNM and IPAM server. A Infoblox user who has the API permission and at least IPv4 network read permission of IPAM will be able to view the retrieved Infoblox records.

In addition to the matched subnets that exist on both IPAM server and DCNM, the IPAM Integrator application lists the subnets with conflicting netmask for review.

You can also watch the video that demonstrates how to use the IPAM Integrator application to view the IP allocation in IPAM server and the relevant networks defined in DCNM. See Video: Using IPAM Integrator in Cisco DCNM.

Accessing IPAM Integrator

This procedure shows how to access IPAM integrator application.

Procedure

Step 1 Navigate to **Applications > Catalog**.

Step 2 Click the IPAM Integrator application icon to access the application. If application is not yet started, this action starts the application before accessing the GUI.

Step 3 Provide the required access details in the **Access Authentication** window.

Note You can provide the access details of an Infoblox server or an Infoblox grid manager.

- **IPAM User Name** – Specifies the user name for the IPAM server. The Infoblox user has to be granted API permission for the application to retrieve data from Infoblox server via API.
- **Password** – Specifies the password for the IPAM server with respect to the username.
- **IPAM Server IP Address** – Specifies the IP address of the IPAM server.
- **Poll Interval (minutes)** – Specifies the time in minutes that determines how often you want the data to be retrieved from Cisco DCNM and IPAM server. The default value is 15 minutes. The range of the polling value is 2–60 minutes.

Step 4 Click **Create**.

Step 5 After you access IPAM, you can remove or modify the access details using the **Settings** icon. You can also edit the poll interval using **Edit**.

Note Only the DCNM users with the **admin** role can add, update, and delete the access setting. Also, only Infoblox user who has been granted with API permission and at least IPv4 network read access of IPAM permission is able to view the retrieved Infoblox network records.

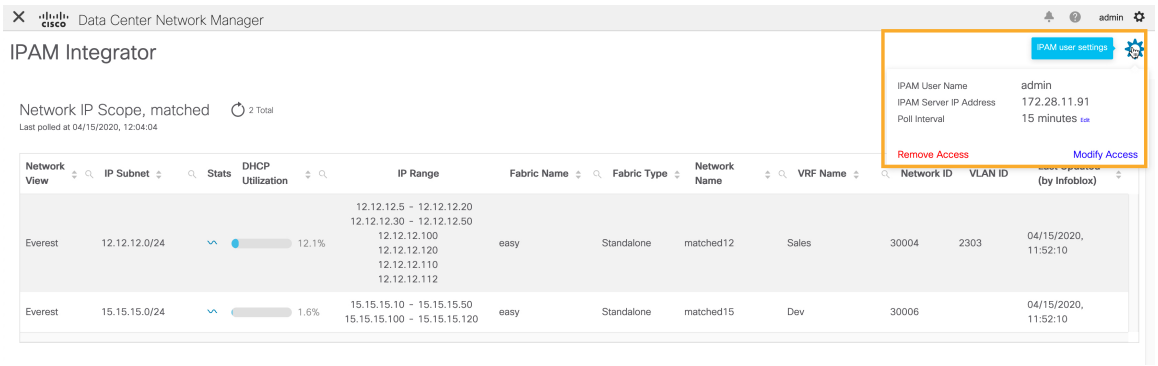

Viewing Network IP Scope

Network IP Scope is the landing page after you access the IPAM Integrator application.

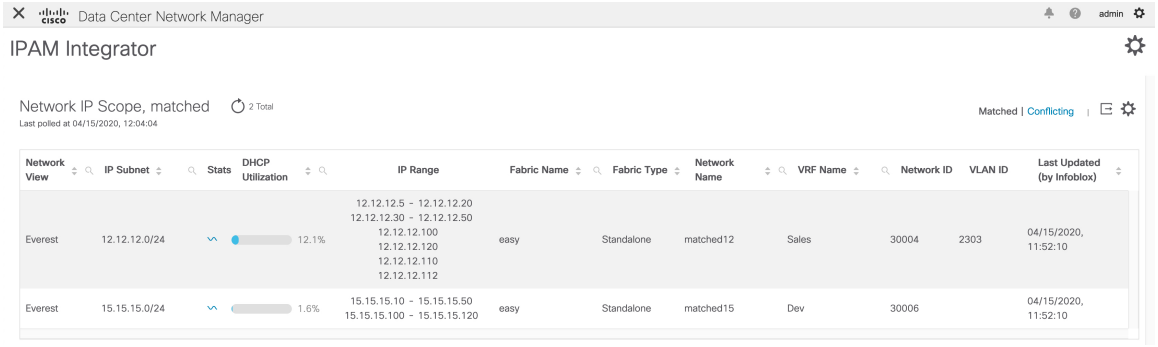

The following table describes the fields retrieved from the IPAM server.

Field	Description
Network View	Specifies the network view, which is a single routing domain with its own networks and shared networks on the Infoblox server.
IP Subnet	Specifies the IP subnet defined in the IPAM server. A subnet, or subnetwork, is a segmented piece of a larger network. More specifically, subnets are a logical partition of an IP network into multiple, smaller network segments.
Stats	Click the icon under the **Stats** column to view the statistics for the utilization of the IP subnet. For more information, see Viewing Statistics for the Subnet Utilization, on page 598.
DHCP Utilization	Specifies the utilization percentage of a network in terms of the IP addresses that are leased out. Hover over the percentage value to view the number of allocated IPs and their details. In the Infoblox server, it takes time to calculate the DHCP utilization. The IPAM utilization is calculated approximately every 15 minutes on the Infoblox server, and the latest value will be reflected on the IPAM Integrator app after that.

IP Range	Specifies the IP range for the network. Hover over a range to view the enabled DHCP range, the reserved DHCP range, and the fixed addresses for a network.

The following table describes the fields retrieved from DCNM.

Field	Description
Fabric Name	Specifies the name of the fabric.
Fabric Type	Specifies the type of the fabric. It can be multi-site deployment (MSD), or a standalone easy fabric or an eBGP VXLAN fabric.
Network Name	Specifies the name of the network.
VRF Name	Specifies the name of the VRF.
Network ID	Specifies the network ID.
VLAN ID	Specifies the VLAN ID.
Last Updated (by Infoblox)	Specifies the date and time when the data was last updated by Infoblox. **Note** The date and time of the last poll are displayed under the **Network IP Scope** title.

Click **Export** to export the data in a .csv file.

For each field, you can sort the values by clicking the arrow icons, and search by clicking the **search** icon and entering a value.

Click the **Settings** icon above the fields to remove or add the fields to be displayed.

The polling of data is based on the following criteria:

- Poll interval value that the user configured initially in the **Access Authentication** window. It specifies how often you want the data to be retrieved from Cisco DCNM and IPAM.
- User can click the **Refresh** icon to receive instantaneous data from DCNM and IPAM server.
- DCNM Web UI automatically refreshes every 2 minutes and displays data retrieved from DCNM and Infoblox server.

  For example, if the poll interval is 15 minutes and user doesn't refresh (on-demand) the data during this 15-minute duration, the DCNM Web UI displays the same polling data after every 2-minute refresh until 15 minutes. After 15 minutes, new data is polled from DCNM and IPAM, and saved in the database. This new data is fetched by DCNM after a total of 16 minutes.

Viewing Statistics for the Subnet Utilization

Click the icon under the **Stats** column to view the statistics for the utilization of the IP subnet over a time.

From the **Time Range** drop-down list, select the time for which you want to view the statistics. These stats include utilization of subnet such as DHCP utilization, total allocations, dynamic allocations, and static allocations.

Viewing IP Allocation for Hosts

Click the IP range value under the **IP Range** column to view the IP allocation for each host.

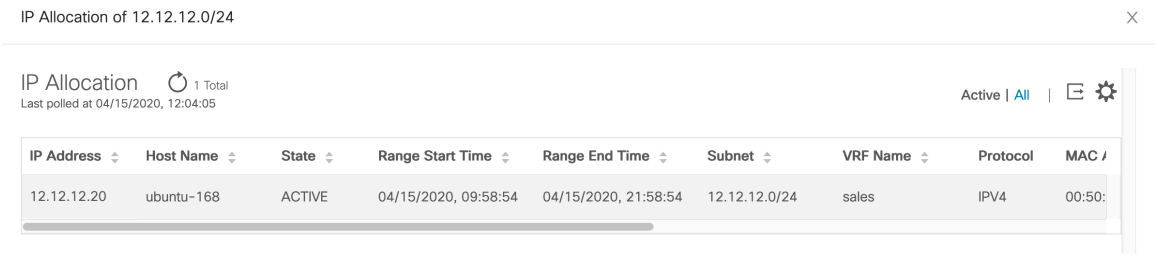

The following fields are displayed for each host in the **IP Allocation** window. The data for these fields is retrieved from the IPAM server.

- IP Address
- Host Name
- State of the host, that is, active or free
- Range start time and end time
- Subnet
- VRF Name

- Protocol version
- MAC address
- DHCP server info such as IP address and server name
- Last requested by the host

For each field, you can sort the values by clicking the arrow icons, and search by clicking the **search** icon and entering a value.

By default, information about only active hosts are displayed. Click the **All** value to view information about all hosts retrieved from the IPAM server. Click **Export** to export the data in a .csv file.

Hosts that were recently freed show as "FREE" in the **All** tab. The all the originally free hosts won't be shown as FREE. Only the hosts that were recently freed appear in this tab.

Click the **Settings** (gear) icon on the right-side to remove or add the fields to be displayed.

Viewing Conflicting Networks

IPAM Integrator detects conflicting networks defined in IPAM server and DCNM. You can view this info by clicking **Conflicting** in the **Network IP Scope** window.

For example, if one network is a subset of another, the conflicting IP addresses of the network are displayed under **Conflicting**.

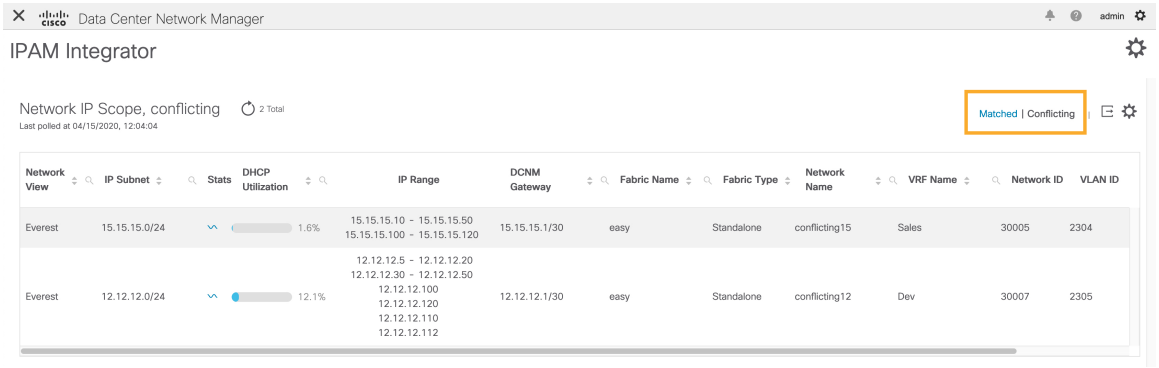

The data is displayed similar to how the **Matched** data is displayed. You can click the IP range value under the **IP Range** column to view the IP allocation for each host.

Note that this table also lists the DCNM Gateway for the conflicting IP scopes in addition to the subnet information from the IPAM server.

For each field, you can sort the values by clicking the arrow icons, and search by clicking the **search** icon and entering a value.

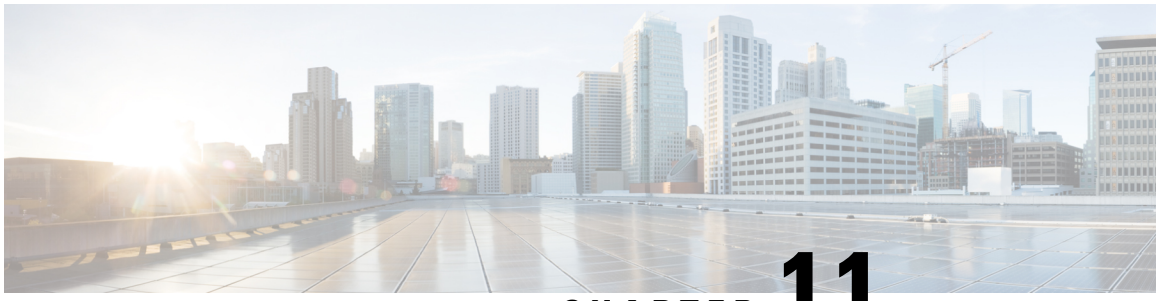

CHAPTER 11

Health Monitor

- Catalog, on page 601

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:

- Health Monitor (2.1)
- PTP Monitoring (1.1)
- Kibana (2.0)
- Programmable report (1.1.0)
- Elastic Service (1.1)
- Compliance (4.0.0)
- Debug Tools (2.1)
- IPAM Integrator (1.0)
- Endpoint Locator (2.1)
- Kubernetes Visualizer (1.1)
- vmmplugin (4.1)

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 551.

Health Monitor

The Health Monitor helps you to monitor the infrastructure health and status. You can monitor the Alerts, Service Utilization, and Compute Utilization using the Health Monitor application. When you install or upgrade to 11.2(1), the Health Monitor application is installed and operational, by default.

To launch the Health Monitor app, on the Cisco DCNM Web UI, choose **Applications**. On the Catalog tab, click on **Health Monitor** to launch the application.

Note Health Monitor 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 Health Monitor 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.

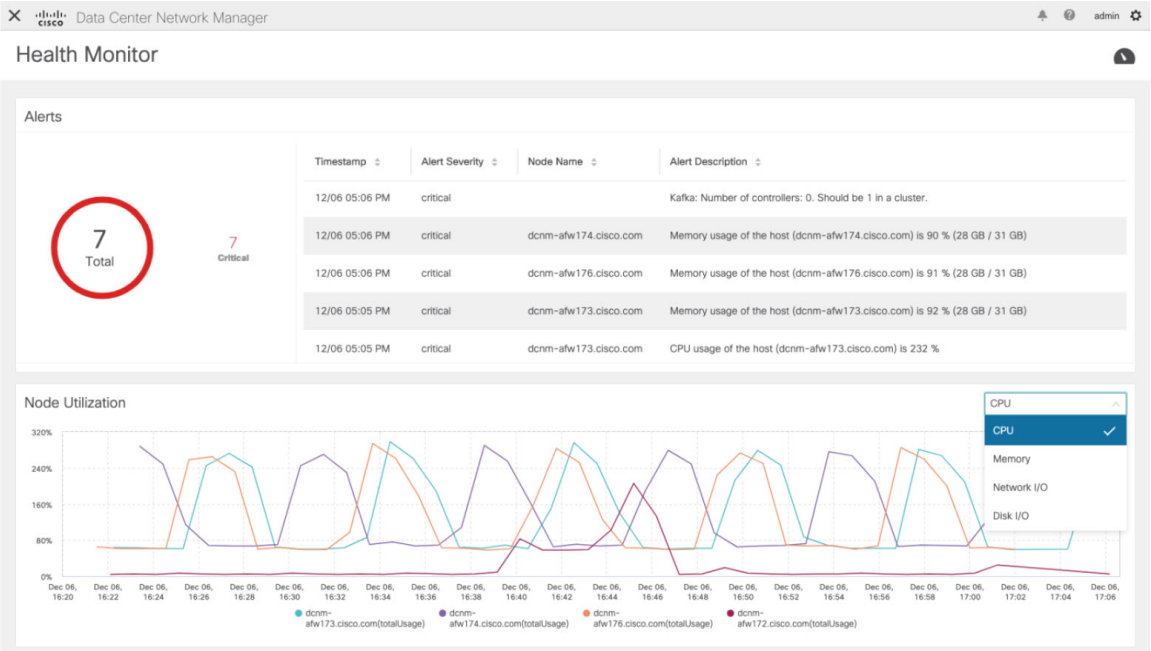

In the List View, alerts are displayed in tabular format with the following categories:

- **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.
- **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.

Health Monitor generates alerts for the following metrics:

- CPU utilization >= 65 %
- Memory utilization >= 65 %
- Disk utilization >= 65 %
- Elasticsearch cluster status: Red/yellow
- Elasticsearch unassigned shards > 0
- Elasticsearch JVM heap used >= 65 %
- Kafka partitions without leader: Controller offline partitions count > 0
- Kafka controllers count: Controller active controller count != 1

- Kafka partition leader: Controller unclear leader elections count > 0

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.

Guidelines and Limitations for Health Monitor in Service Utilization

- The CPU utilization for applications without a CPU limit, like Kafka, ElasticSearch, FMserver, and so on, may show 100% utilization in the graphs. 100% utilization is because this application uses one or more cores.
- The following alerts are triggered for the CPU utilization of applications:
  - Minor alert: 200-400 %

    Major alert: 400-600%

    Critical: > 600%

- The transient message for Kafka controller counts appears as a severe alert sometimes. You can ignore the alert if it clears within two minutes after refresh.
- The **Disk I/O** and **Memory Utilization** metrics are not available for Kafka and Elastic Service.
- The **Network I/O** metric is not available for **DCNM: FMServer** and **DCNM: Postgres**.
- The metrics does not auto-refresh. Navigate between different windows using the options in the drop-down list to refresh the metrics. Additionally, you can change the time range to refresh the metrics for a selected period.
- There might be duplicate alerts for the same feature.

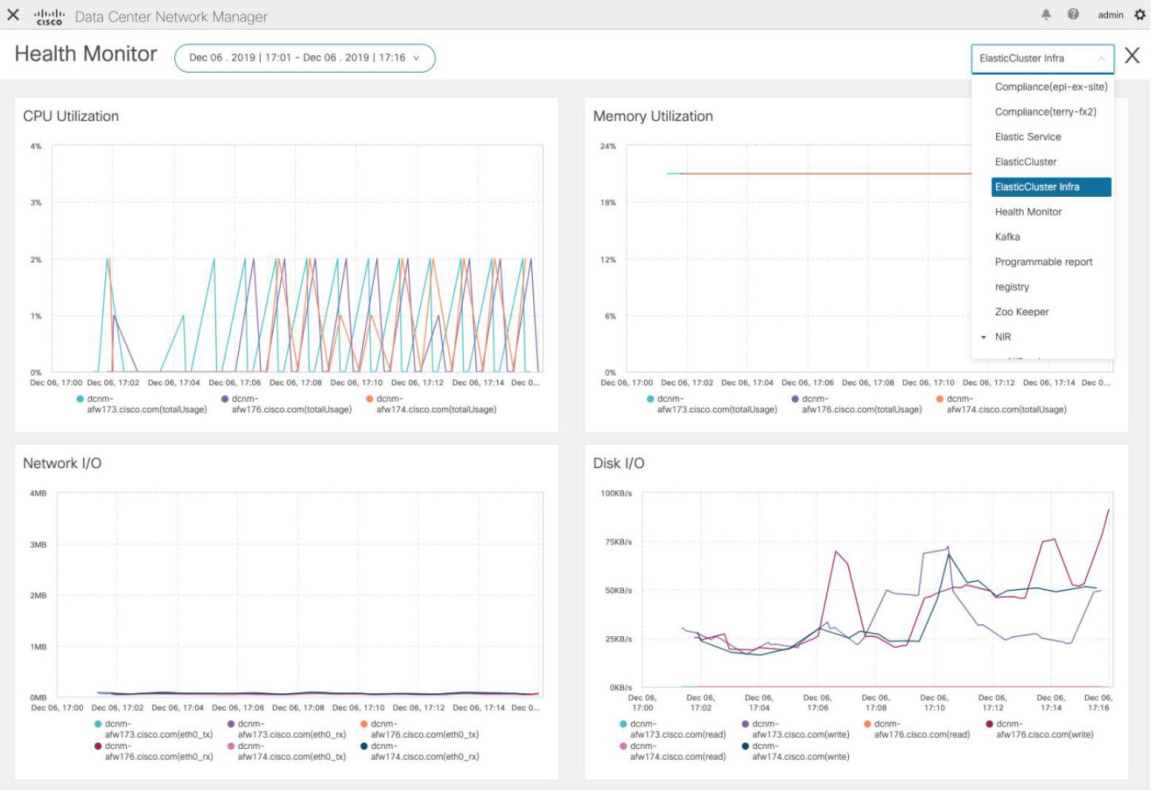

The following additional metrics are collected for Elastic Cluster:

- Get latency: Latency for getting a single record by id
- Cluster heap used: Heap memory used by the cluster
- Unassigned shards: Count of unassigned shards
- Node heap used percentage: Percentage heap memory used by the node
- Search latency: Latency for getting a collection of records

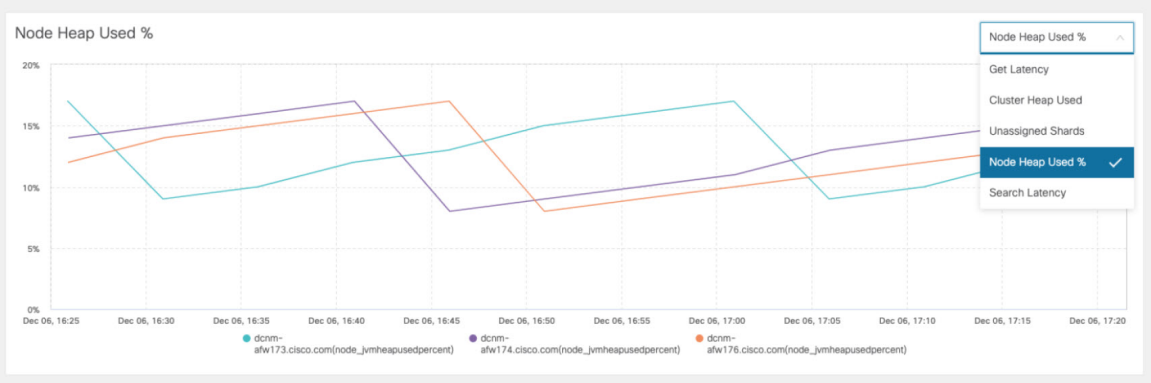

The following additional metrics are collected for Kafka broker:

- CPU: CPU utilization of broker

- Memory: Memory utilization of broker
- Heap memory: Heap memory utilized by broker
- Total time: Network produce, network fetch follower, network fetch consumer time
- Purgatory size: Server fetch purgatory size, server produce purgatory size of broker
- Data in: Bytes in for the broker
- Data out: Bytes out for the broker
- Messages in: Messages received by the broker
- Fetch request: Total fetch requests for the broker
- ISR: In-sync-replicas expands and shrinks for the broker

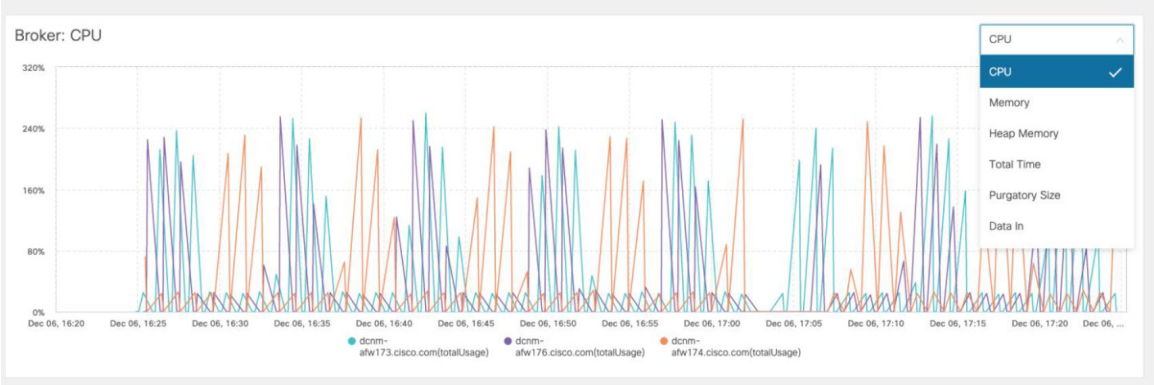

The following additional metrics are collected for top 5 Kafka topics:

- Data in: Bytes in for the topic
- Data out: Bytes out for the topic
- Messages in: Message in count for topic
- Messages out: Message out count for topic
- Lag: Lag per topic

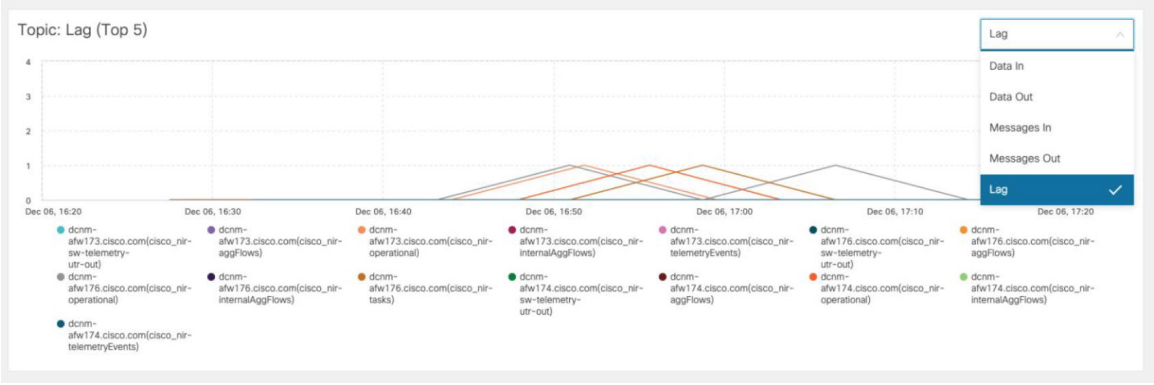

Compute Utilization

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.

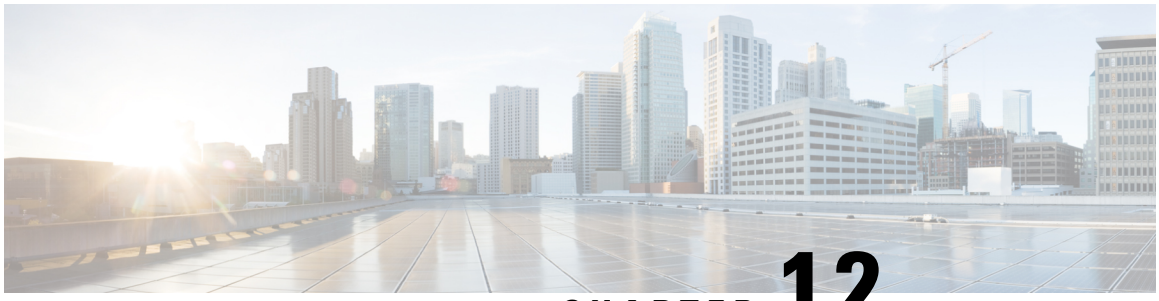

CHAPTER 12

PTP Monitoring

- Catalog, on page 609

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:

- Health Monitor (2.1)
- PTP Monitoring (1.1)
- Kibana (2.0)
- Programmable report (1.1.0)
- Elastic Service (1.1)
- Compliance (4.0.0)
- Debug Tools (2.1)
- IPAM Integrator (1.0)
- Endpoint Locator (2.1)
- Kubernetes Visualizer (1.1)
- vmmplugin (4.1)

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 551.

PTP Monitoring

This section explains the preview functionality of the Precision Time Protocol (PTP) monitoring. PTP is a time synchronization protocol for nodes that are distributed across a network. On a local area network, it achieves clock accuracy in the sub-nanosecond range, making it suitable for measurement and control systems.

In DCNM, PTP Monitoring can be installed as an application. From the DCNM Web UI, navigate to **Applications** and click **PTP Monitoring**. This application works in the IPFM mode only.

In the **PTP Management** window, you can view PTP related information based on the switch selected from the **Select a switch** drop-down list. You can click the **Telemetry Switch Sync Status** link to check whether the switches are in sync. The **Sync Status** column displays the status of the switches.

The following tabs are displayed in this window:

- **Correction & Mean Path Delay**
- **Clock & Port Status**

Note The PTP related info is displayed for the switch group that you select from the **SCOPE** drop-down list.

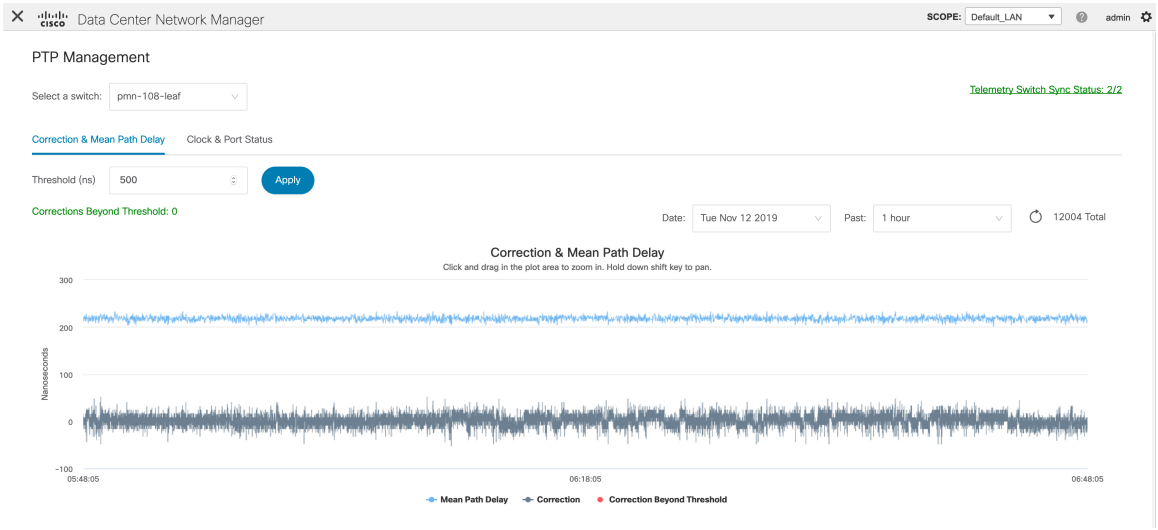

Correction and Mean Path Delay

The **Correction & Mean Path Delay** tab displays a graph showing the PTP operational statistics: mean path delay, correction, and correction beyond threshold. You can click and drag in the plot area to zoom in and hold the **shift** key to pan. Click the **Reset zoom** button to reset zoom.

By default, the graph is displayed for the threshold value of 500 nanoseconds (ns). You can also display data based on a specific threshold value. In the **Threshold (ns)** field, enter the required value in nanoseconds and click **Apply**. Note that the threshold value is persistent in the DCNM settings, and it is used to generate PTP correction threshold AMQP notifications.

From the **Date** drop-down list, you can select the appropriate date to view the data. The PTP data is stored up to the last seven (7) days. The default value for the stored data is 7 days. To change this value, navigate

to **Administration > DCNM Server > Server Properties** and set the updated value for the **pmn.elasticsearch.history.days** property.

From the **Past** drop-down list, you can also select a timeframe over which the data has to be displayed. The values in the **Past** drop-down list are 1, 6, 12, and 24 hours.

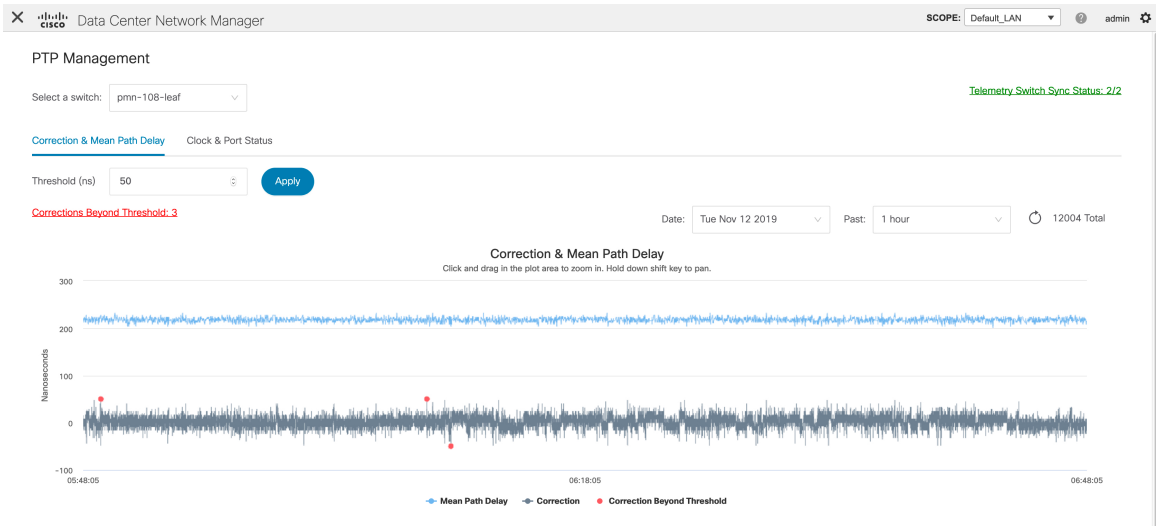

Note that you can click the legends in the graph to hide or display statistics.

If there are any corrections, you can view them in a tabular format by clicking the **Corrections Beyond Threshold** link.

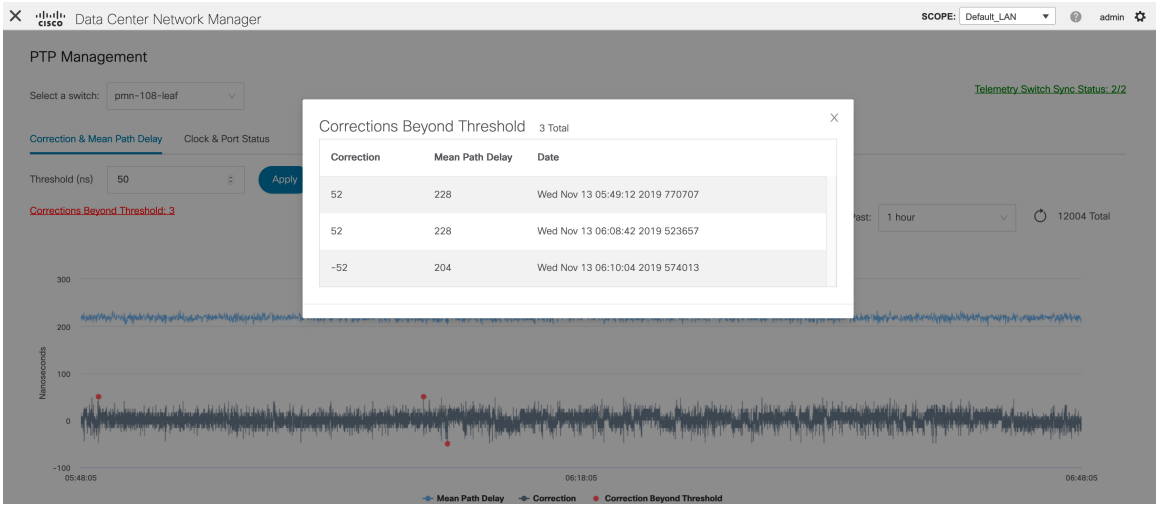

Clock and Port Status

The **Clock & Port Status** tab displays status for Parent Clock, Grandmaster Clock, and ports.

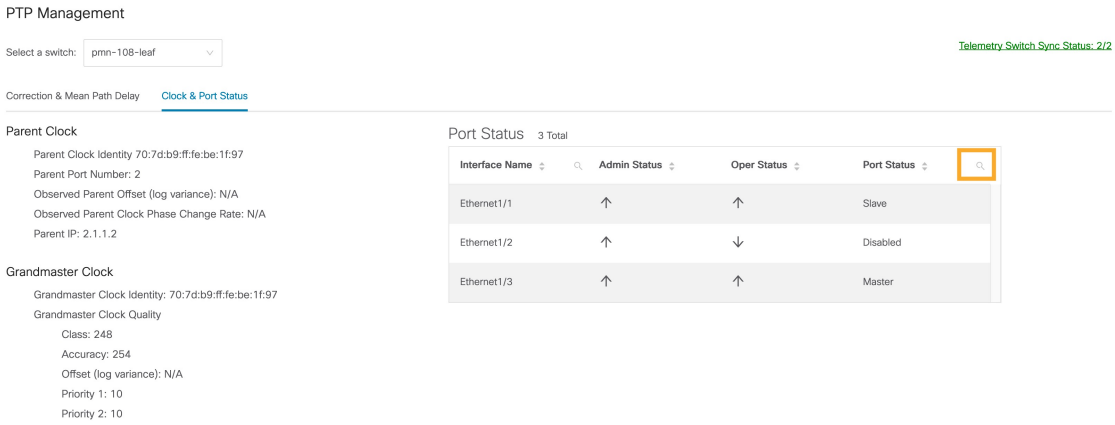

The **Port Status** table displays the status of the ports and the peer ports. Click the **Search** icon, and enter the port status, and click **Search** to filter the port status.

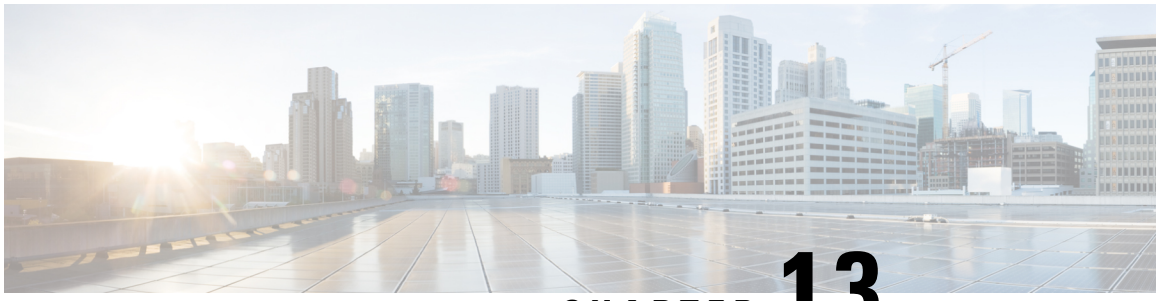

CHAPTER 13

Programmable Reports

- Catalog, on page 613

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:

- Health Monitor (2.1)
- PTP Monitoring (1.1)
- Kibana (2.0)
- Programmable report (1.1.0)
- Elastic Service (1.1)
- Compliance (4.0.0)
- Debug Tools (2.1)
- IPAM Integrator (1.0)
- Endpoint Locator (2.1)
- Kubernetes Visualizer (1.1)
- vmmplugin (4.1)

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 551.

Programmable Report

The Programmable Report application enables the generation of reports using Python 2.7 scripts. Report jobs are run to generate reports. Each report job can generate multiple reports. You can schedule the report to run for a specific device or fabric. These reports are analyzed to obtain detailed information about the devices.

The REPORT template type is used to support the Programmable Report feature. This template has two template subtypes- UPGRADE and GENERIC. For more information on the REPORT template, refer Report Template. A python SDK is provided to simplify report generation. This SDK is bundled with DCNM and provides APIs to generate reports. For more information on APIs, refer Report Python Library.

RBAC support

- An admin or a network operator can create a report..
- A network operator can view reports created by other admins and operators.
- A network operator cannot delete/edit/rerun any reports created by an admin and other network operators.
- An admin can view and delete any report irrespective of the user creating them.
- An admin cannot edit any report created by other user, including network operator, due to fabric and device association.

Note A Jython template supports a maximum file size of 100k bytes. In case any report template exceeds this size, Jython execution may fail.

To launch the Programmable Report app, on the Cisco DCNM Web UI, choose **Applications**. On the **Catalog** tab, click **Programmable report** to launch the application. The Report window is displayed. This window has two tabs, **User Defined** and **Internal**. The **User Defined** tab displays the report jobs which are created by a user. For information on creating a report job, refer Creating a Report Job. On any window, click the **Home** icon at the top left of the screen to return to this **Report** window. All operations such as creating, deleting, editing, re-running, displaying history, and downloading of report job information, are supported for the jobs displayed in this tab.

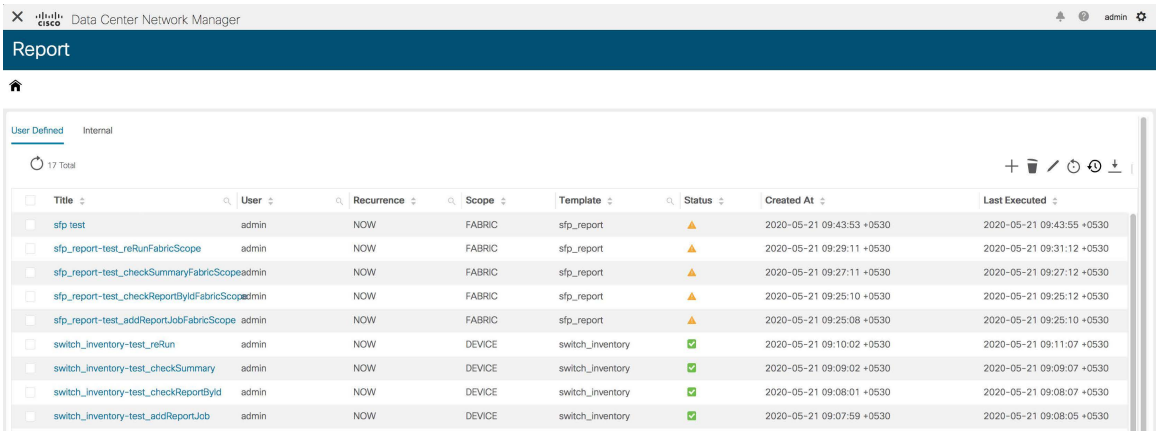

Field	Description
Title	Specifies the title of the report job.

Field	Description
User	Specifies the user who has initiated the report generation.
Recurrence	Specifies the frequency at which the reports are generated.
Scope	Specifies the scope of the report. The report can be generated for a device or a fabric.
Template	Specifies the name of the template.
Status	Specifies the status of the report. The status messages are as follows- *Success - Report is generated successfully. *Scheduled - A report generating schedule is set. *Running - A report job is running. *Failed - Report execution failed for one or more selected switches/fabrics or an issue occurred during running of the report job. *Unknown - Job state could not be identified.
Created At	Specifies the time at which the report was created.
Last Executed	Specifies the time at which the report was last generated.
Start Date	Specifies the date at which the report generation is scheduled to start from.
End Date	Specifies the date at which the report generation is scheduled to end at.

The **Internal** tab displays the report jobs that have been created by DCNM. For example, Pre and Post-ISSU report jobs that are created by the ISSU wizard are considered as internal jobs. However, you can only view report job information and history for the report jobs in this tab. You cannot delete a report job displayed under this tab as any deletion may affect the working of a DCNM feature that depends on this report.

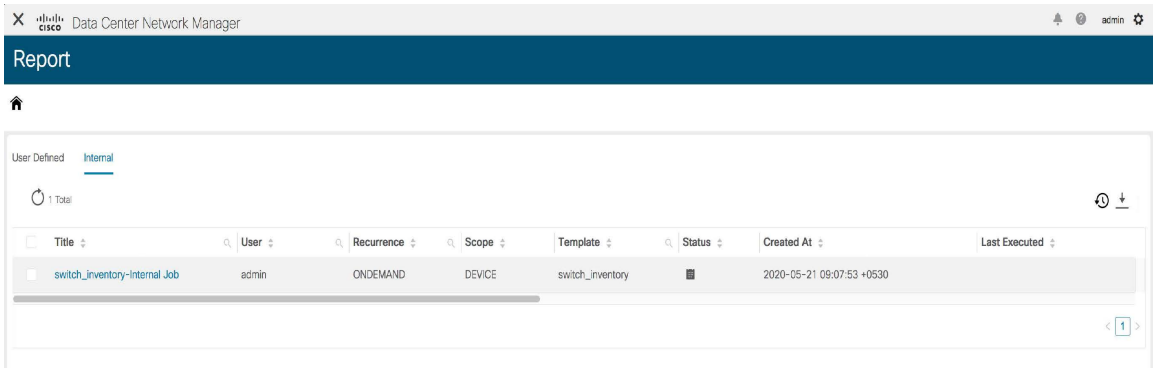

You can also watch the video that demonstrates how to use the Programmable Report application in Cisco DCNM. See Programmable Report.

Creating a Report Job

To create a report job, perform the following steps:

Procedure

Step 1 Click **Create Report** icon.

The Create Report window is displayed.

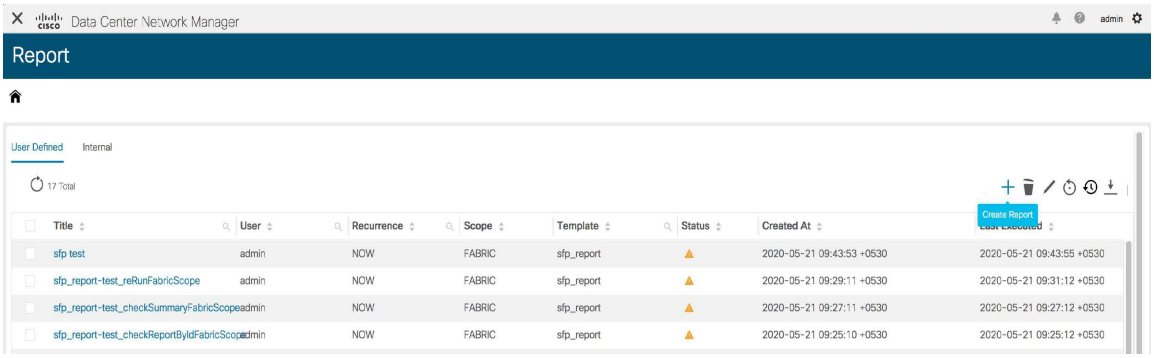

Step 2 Enter a name for the report job in the **Name** field.

Step 3 Specify the frequency at which the report job should be run by selecting the required button next to Recurrence. For this use-case, select

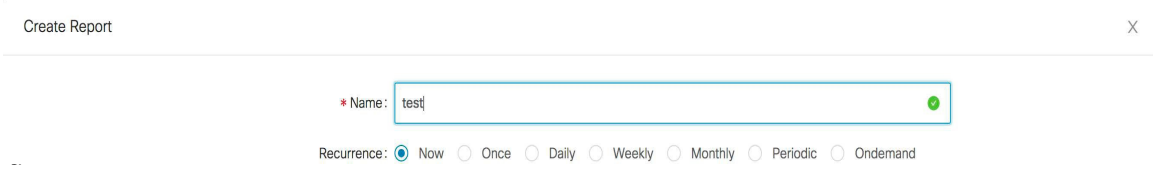

The Recurrence options are as given below:

*Now - The report is generated now.

*Once - The report is generated once at the specified time.

*Daily - The report is generated daily at a specified time between the Start Date and End Date.

*Weekly - The report is generated once a week at a specified time between the Start Date and End Date.

*Monthly - The report is generated once a month at a specified time between the Start Date and End Date.

*Periodic - The report is generated periodically in a time period between the specified Start Date and End Date. The interval of time between the reports can be specified in minutes or hour(s).

Note When you are creating a **Periodic NVE VNI** Counters report, the report generation frequency has to be set to 60 minutes or more. If the frequency is less than 60 minutes, an error message is displayed.

*Ondemand - The report is generated on demand. This report can only be generated by clicking the **Rerun** icon in the Report window.

Note The **Start Date and Time** and **End Date and Time** are displayed in a 24-hour format.

Step 4 Specify the devices or fabrics on which the report job should be run to generate the

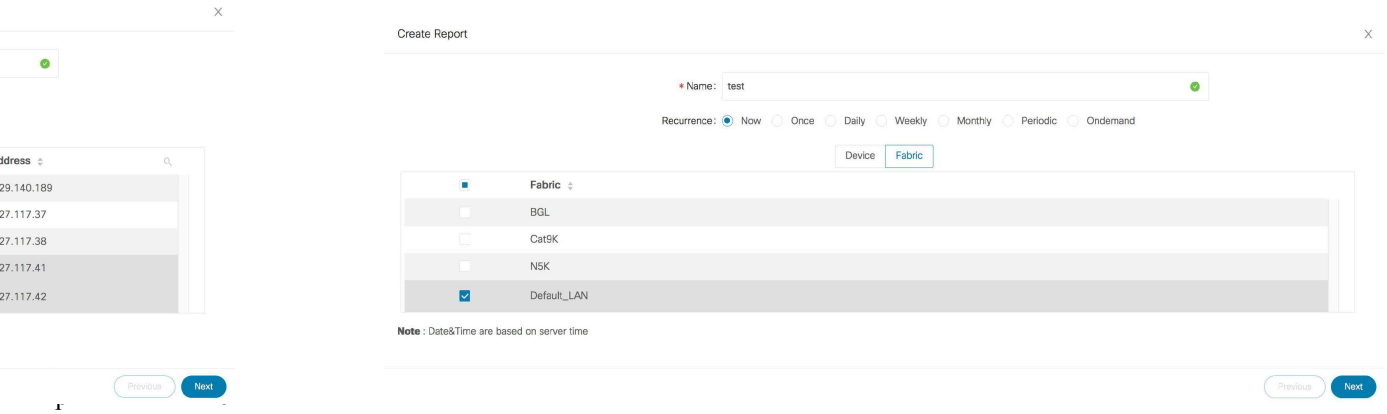

Step 5 Click **Next**. Select a template from the **Template** drop-down list in the **Create Report** window. Each report template has a device or fabric tag associated with it.

The following pre-defined templates are available:

Device scope-

*switch_inventory

Fabric scope-

*fabric_nve_vni_counter

*fabric_resources

*sfp_report

In addition to the templates listed above, any other templates that have been created by you will also be listed here. For more information on default templates and creating customized templates, refer Template Library.

Templates are listed based on the associated tags. If you select the **Device** scope, the templates with the device tag are listed in the drop-down list. If you select the **Fabric** scope, the templates with the fabric tag are listed in the drop-down list.

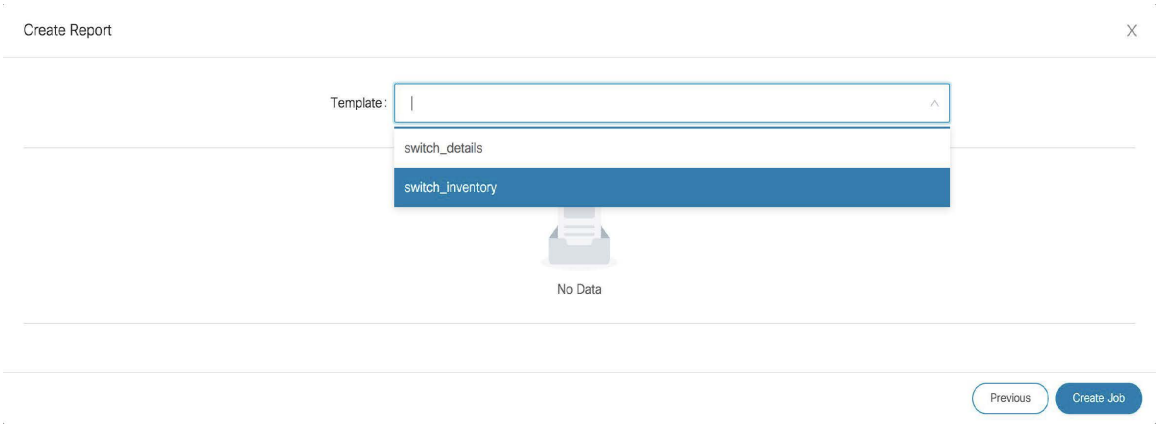

Step 6 Click **Create Job**. A pop-up comes up stating that the Job report has been created successfully. Click the **Refresh** icon in case the newly created job report is not displayed in the list. When you hover over the Status

column for the new report, the status is **Running** which indicates that a report is currently being generated. The status will change to a green tick indicating **Success** after the report has been successfully generated.

Viewing a Report Job

Click a report title from the list of displayed report jobs to view the required information.

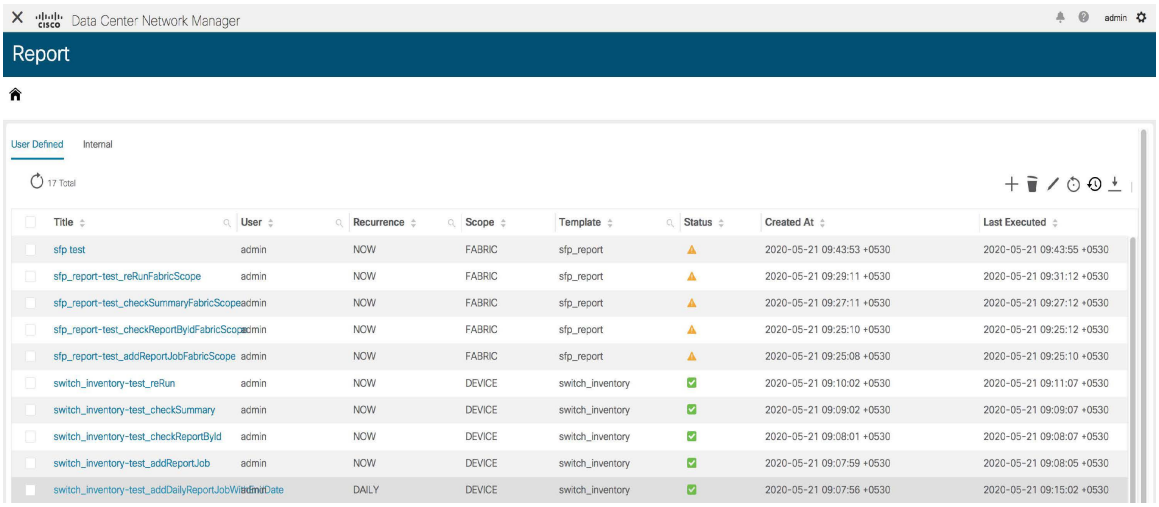

The **Report** window is displayed. The number of errors, warnings, info, and success messages that are color coded and displayed in this window depends on the report details. Errors are displayed in red, Warnings in yellow, Info in blue and Success in green. The summary is not considered for the generation of these numbers.

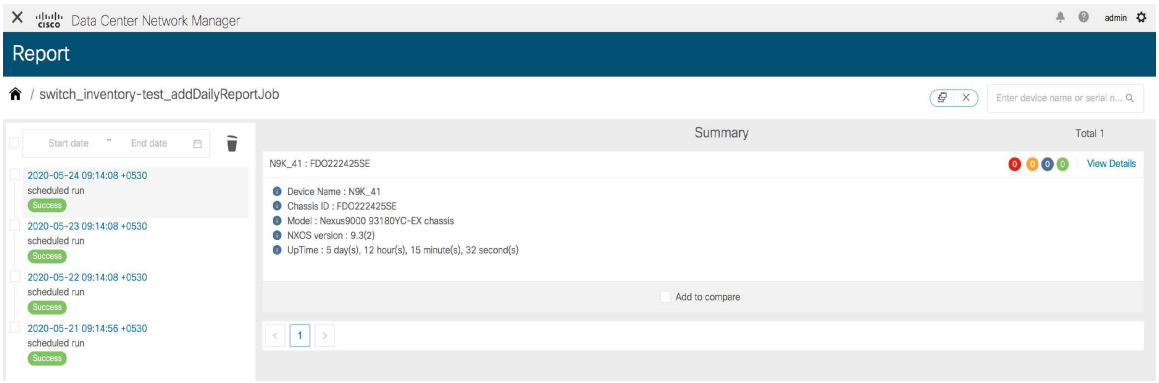

Reports can be generated for multiple devices. You can also see a timeline on the left of the window indicating the time at which a report was generated. You can click an item in this timeline to display the report that was generated at that moment. You can also display the reports that are generated in a specific time-frame by selecting a **Start date** and **End date**.

Click **View Details** to display detailed information.

The **Details** tab displays a report summary along with other relevant information based on the type of report template.

The report detail is logically grouped into sections. Each section is displayed separately with a collapsible widget. The number of errors, warnings, info, and success messages generated in the report is color-coded and is displayed on the top right part of the window.

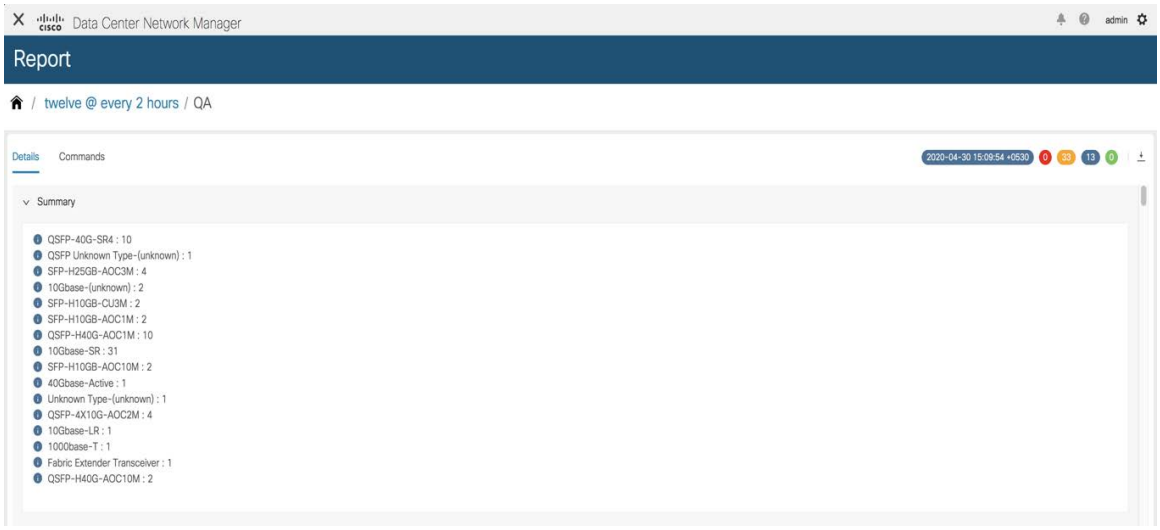

Click the **Commands** tab to display the commands that have been executed to generate the report. The commands are displayed based on the templates and the API that is used to run the commands on the device. For example, in the switch_inventory template, the **show version**, **show inventory** and **show license usage** commands are run to retrieve information. Note that the commands are displayed only if the **show_and_store** API is used to run the commands on the device.

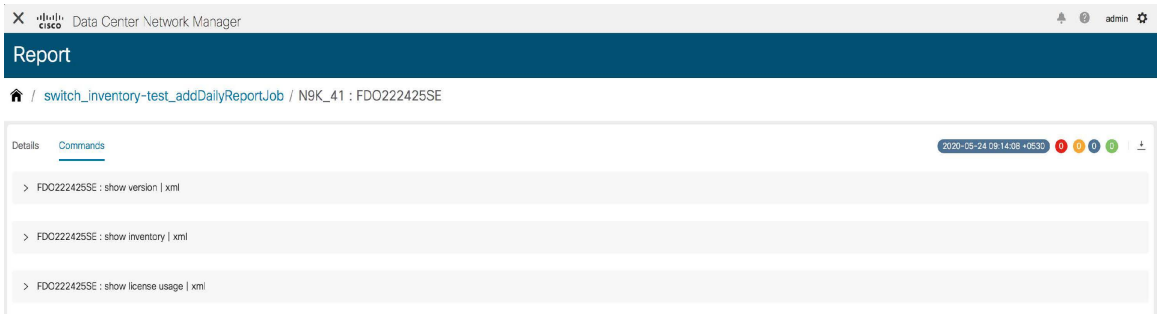

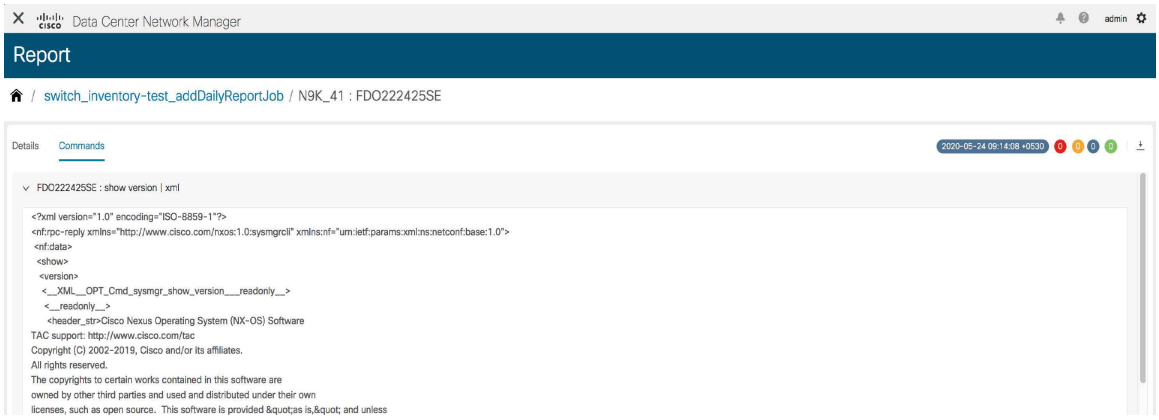

Downloading Report Information

To download report information, click the **Export Report** icon in the Details tab. To download the commands along with the outputs, click the **Export Report** icon in the **Commands** tab.

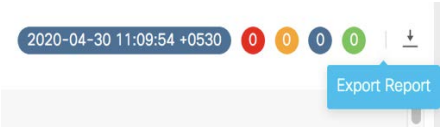

Detailed information about the report is then displayed in a separate window.

SFP Report

Summary

ERRORS	0
WARNINGS	0
SUCCESS	0
INFO	0

10Gbase-(unknown)	10
SFP-H10GB-CU3M	2
10Gbase-SR	12

Device-Level SFP count

warnings	0
title	"Device-Level SFP count"
success	1
errors	0
info	0

SFP count per device

Device	Device level SFP Count
N5648-38	13
N5596-37	11

Device Level: N5648-38

warnings	0
title	"Device Level: N5648-38"
success	1
errors	0
info	0

Interface SFP type

lencu	part	serial_number	interface	type	speed	cdp	device_name	name	model
2	AFBR-7IER02Z-CS1	SSI15470HJ5	Eth1/10	10Gbase-(unknown)	1000,10000	N/A	N5648-38	CISCO-AVAGO	N5K-C5548UP-SUP
N/A	SFBR-7702SDZ-CS5	SSI15470HJ5	Eth1/11	10Gbase-SR	1000,10000	N5596-37(FOX1816G0S9)@Ethernet1/11	N5648-38	CISCO-AVAGO	N5K-C5548UP-SUP
N/A	FTLX8571D3BCL-C2	SSI15470HJ5	Eth1/13	10Gbase-SR	1000,10000	N/A	N5648-38	CISCO-FINISAR	N5K-C5548UP-SUP
N/A	FTLX8571D3BCL-CS	SSI15470HJ5	Eth1/15	10Gbase-SR	1000,10000	N/A	N5648-38	CISCO-FINISAR	N5K-C5548UP-SUP
N/A	FTLX8571D3BCL-C2	SSI15470HJ5	Eth1/17	10Gbase-SR	1000,10000	N/A	N5648-38	CISCO-FINISAR	N5K-C5548UP-SUP
3	74752-9520	SSI15470HJ5	Eth1/21	SFP-H10GB-CU3M	1000,10000	N5596-37(FOX1816G0S9)@Ethernet1/21	N5648-38	CISCO-MOLEX	N5K-C5548UP-SUP

Deleting a Report

To delete a report, select a report on the report timeline that has to be deleted and click the **Delete** icon.

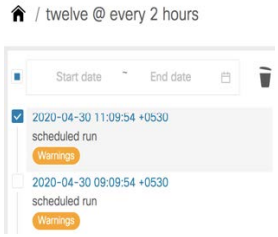

A pop-up window is displayed asking for confirmation to delete the report. Click **Yes** to delete the report.

Comparing Reports

You can compare any two reports that have been generated from the same report job.

To compare reports, perform the following steps:

Procedure

Step 1 Select the **Add to compare** checkboxes in a **Report** window. The number of reports that have been selected for comparison is indicated by the red colored number icon next to the X at the top right of the

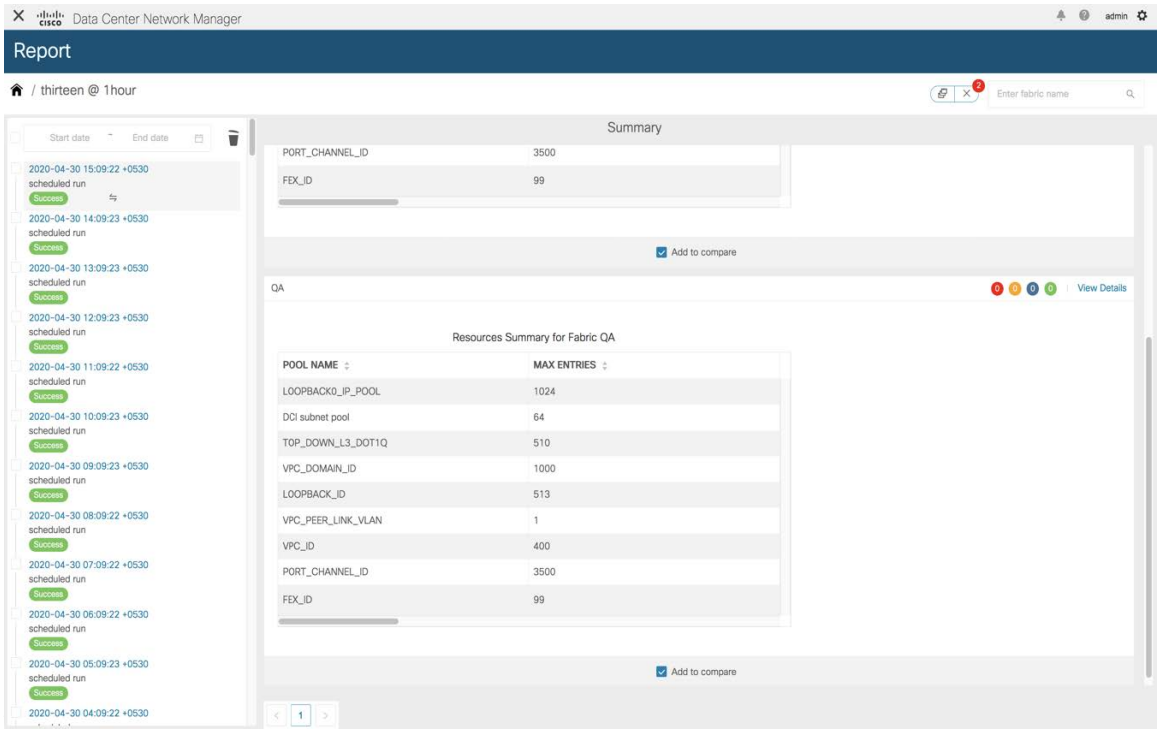

Step 2 Click **Compare Reports** icon at the top right of the

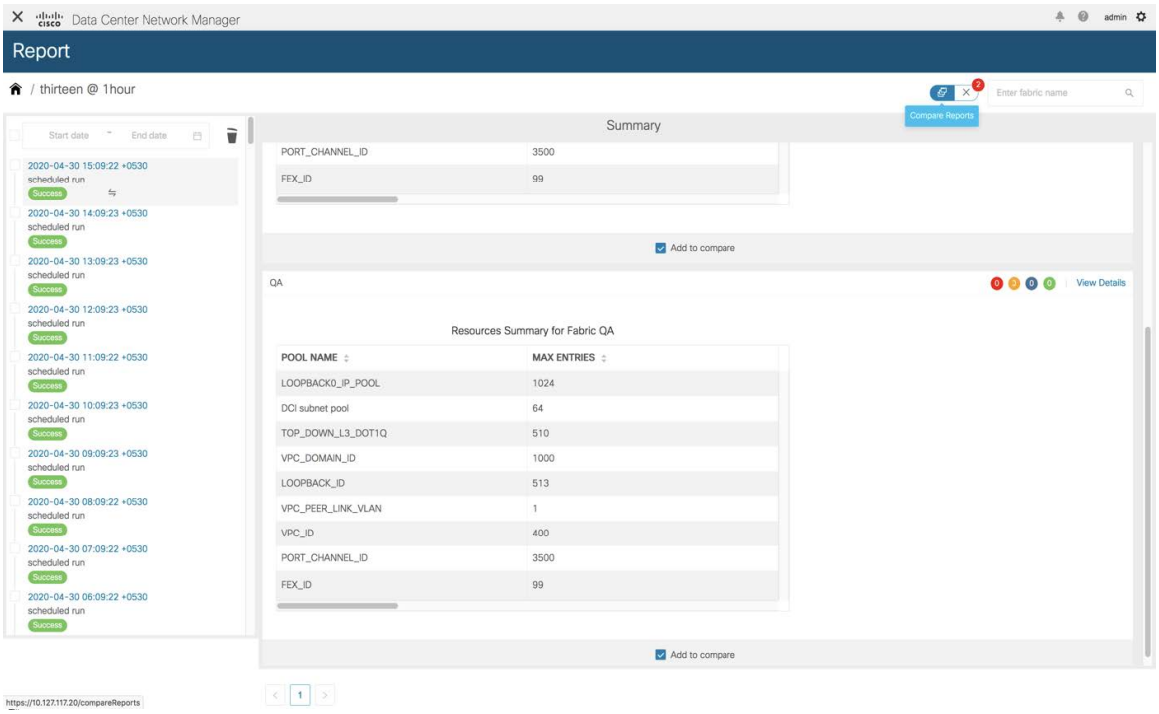

Step 3 The **Compare Reports** window is displayed with a side-by-side

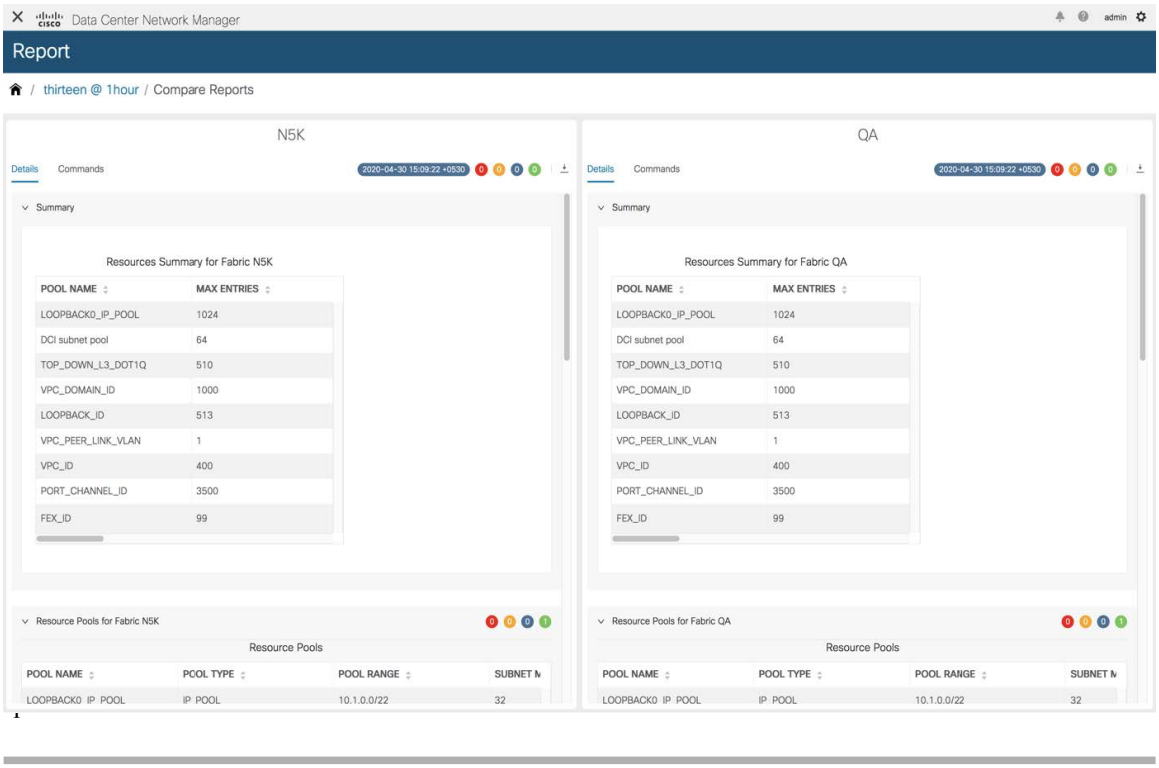

Deleting a Report Job

To delete a report job, select the checkbox next to the report job that has to be deleted and click the **Delete Report** icon.

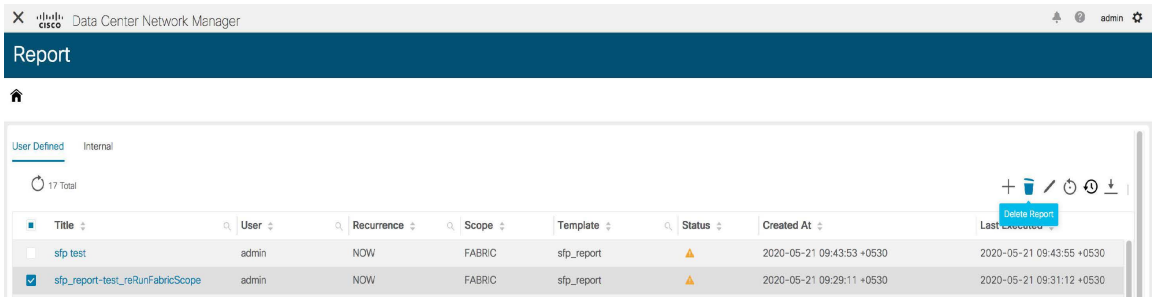

Deleting a report job also deletes all the reports that have been generated by that report job.

Editing a Report Job

To edit a report job, perform the following steps:

Procedure

Step 1 Select the checkbox next to the report that has to be edited and click the **Edit Report** icon. Note that only one job can be edited at a

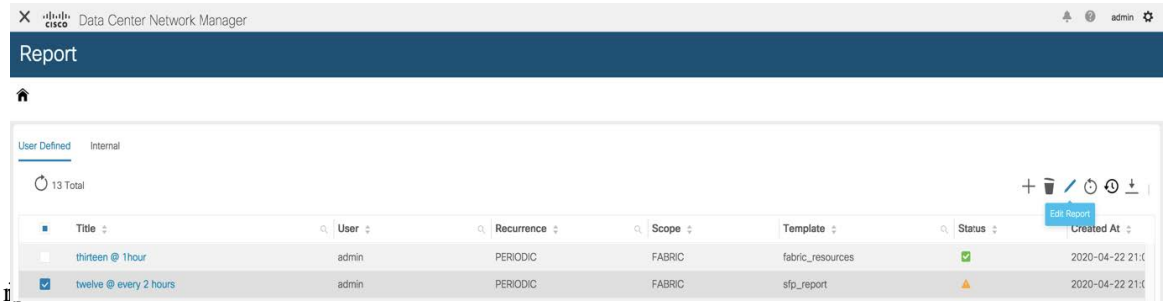

Step 2 The **Create Report** window is displayed. You can edit the Start Date & Time, End Date & Time, Period, and Device or Fabric selected. After editing the required parameters, click

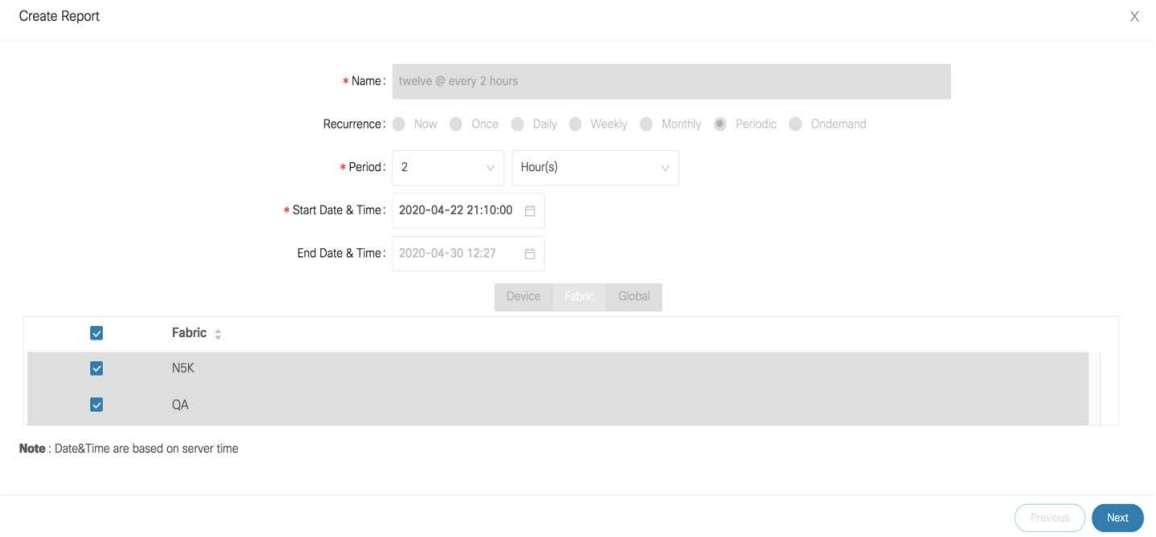

Step 3 Click **Update**

A pop-up window is displayed indicating that the report job has been updated successfully.

Rerunning a Report Job

On the **Report** window, select the checkbox next to the report that has to be generated again and click the **Rerun Report** icon to run a report job again. A pop-up window is displayed indicating that the report job has

been run again

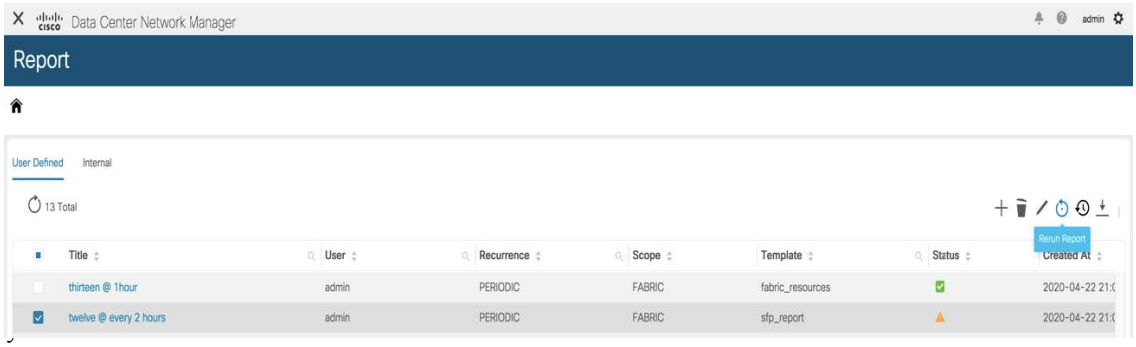

You can use the re-run option to generate a report before the scheduled execution time. In case of an **Ondemand** job, you need to click the **Rerun Report** icon to generate the report.

Displaying Report Job History

To display report job history, select the checkbox next to the report job for which the history has to be displayed, and click the **History** icon. The **Job History** window is displayed. You can view the last 100 entries per report

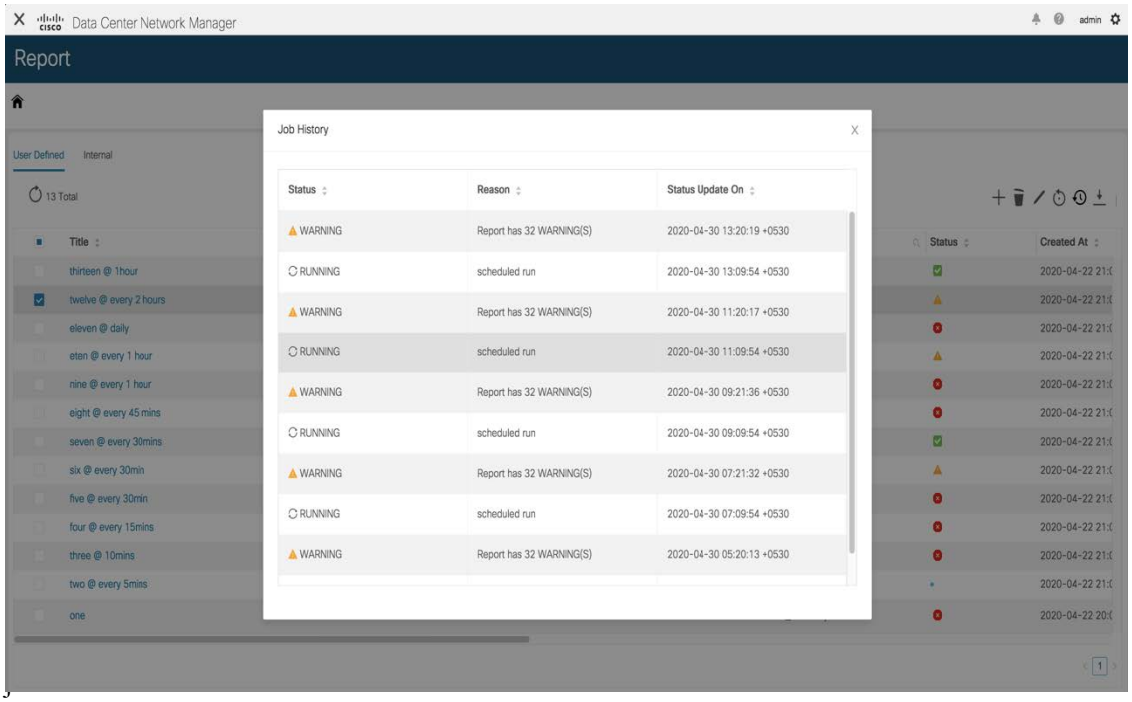

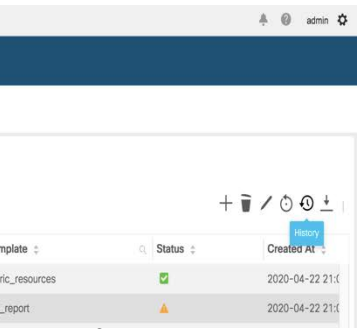

Downloading Report Job Information

To download report job information as a CSV file, click the **Download** icon, select the location to which the CSV file should be downloaded to, and click **OK**. A CSV file with information about the report job is

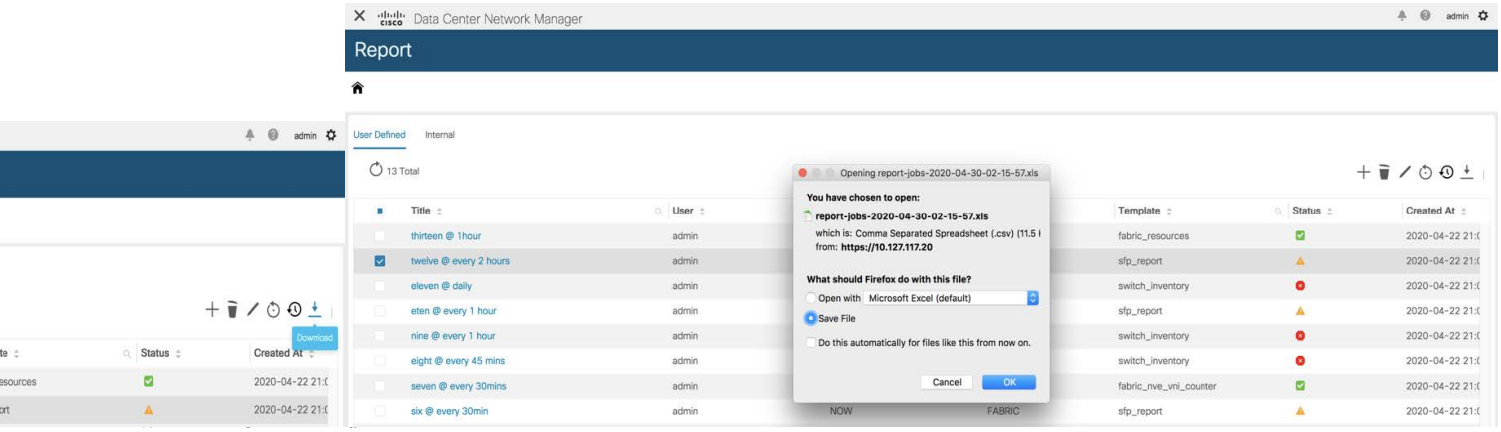

Report Purging

Purging reports leads to creation of space for new reports. A separate report index is created for each report job as each report job can have a different recurrence interval. So, the number of reports that are generated in a given period of time may be different for different report jobs. Separate report indexes help in purging the reports easily. Each report index has a maximum size limit of 100MB and a maximum report count of 500. If the limits are exceeded, the older reports are deleted to make space for new reports.

At any instant, only 70% of the maximum threshold value is retained. For example, consider a scenario in which the maximum report index size is 100MB. On purging, only reports that take up to 70MB space on an average are retained. This provides space for new reports that are generated. The maximum report count of 500 also has a threshold value of 70%.

- To modify the limits and the threshold percentage values, use the following REST API:

  URL: *appcenter/Cisco/preport/integrated/admin/reportconfig*

  METHOD: POST

  You can configure one or more threshold attributes and values.

  ```
  {
  <threshold attributes>: <value>
  }
  ```

  The <threshold attributes> are given below-

  MAX_REPORT_SIZE: Maximum size of the report index in KB

  MAX_USAGE_PERCENTAGE: Maximum percentage of “MAX_REPORT_SIZE” to retain

  MAX_NUMBER_OF_REPORTS: Maximum number of reports

  MAX_NUMBER_OF_REPORTS_PERCENTAGE: Maximum percentage of

  “MAX_NUMBER_OF_REPORTS” to retain

  MAX_HISTORY_SIZE: Maximum size of history in KB

  MAX_HISTORY_PERCENTAGE: Maximum Percentage of “MAX_HISTORY_SIZE” to retain

MAX_NUMBER_OF_HISTORY: Maximum number of historical reports to retain

MAX_NUMBER_OF_HISTORY_PERCENTAGE: Maximum percentage of

“MAX_NUMBER_OF_HISTORY” to retain

Enter integers for the <*value*> attribute.

- *To retrieve the currently configured limits, use the following API:

  URL: *appcenter/Cisco/preport/integrated/admin/reportconfig*

  METHOD: GET

- To retrieve current utilization statistics, use the following API:

  URL: *appcenter/Cisco/preport/integrated /admin/index/stats*

  METHOD: GET

- Reports are purged once a day at 12 AM. You can also initiate a purge. To initiate a purge, use the following-

  REST API:

  URL: *appcenter/Cisco/preport/integrated/admin/purge/report*

  METHOD: POST

- Report execution history is stored in a single index for all the jobs and purged at even hours. The maximum index limit for report execution history is 1000 and the maximum size allowed is 500MB. To change these limits, use the following REST API:

  URL: *appcenter/Cisco/preport/integrated/admin/reportconfig*

  METHOD: POST

- To initiate a purge of the report execution history, use the following REST API:

  URL: *appcenter/Cisco/preport/integrated/admin/purge/history*

  METHOD: POST

- To initiate a purge of both reports and report execution history, use the following REST API:

  URL: *appcenter/Cisco/preport/integrated/admin/purge*

  METHOD: POST

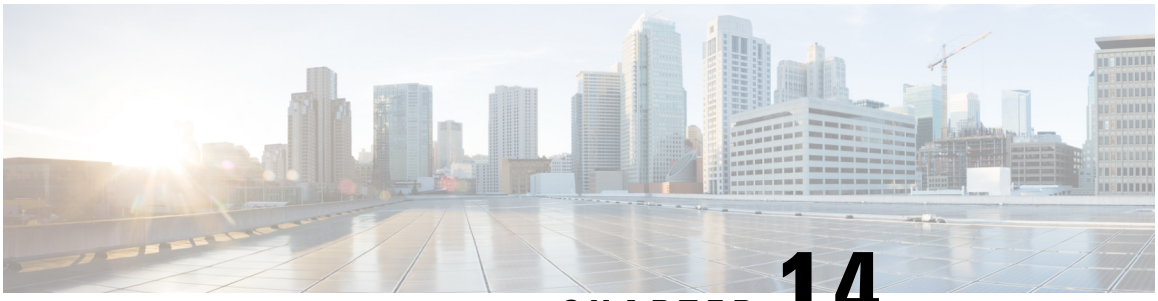

CHAPTER 14

ServiceNow Integration

• DCNM Integration with ServiceNow, on page 629

DCNM Integration with ServiceNow

ServiceNow offers applications for IT Service Management (ITSM) and IT Operations Management (ITOM). There are four primary modules - inventory discovery, incident management, event management & change management workflows. Starting from Cisco DCNM Release 11.3(1), we provide Cisco DCNM integration with ServiceNow. This enables you to integrate end-user IT data with the ServiceNow platform. The integration provides a default set of ServiceNow custom tables which are populated with configuration data.

To utilize this functionality, install the DCNM application in the ServiceNow customer instance and provide the DCNM mid-server details. Information or data regarding switch details, port details, and alarms, is retrieved to the ServiceNow Configuration Management Database (CMDB) tables. By default, data is retrieved every 15 minutes and displayed.

Details about the switches and ports of each switch are collected from the DCNM inventory. The alarms are collected by polling DCNM. Alarms are then filtered and categorized based on their type, such as, CPU, MEMORY, POWER, LINKSTATE, EXTERNAL, ICMP, SNMP, and SSH. The alarms are then stored in an Events table. These events are then used to generate incidents for the CPU, MEMORY, SNMP, and SSH categories. The source, description, severity and category of each alarm is stored. However, when an alarm ceases to exist in DCNM, the incident that was raised for it is not updated or cleared on the DCNM ServiceNow application. When polling of alarms is initiated for the first time, the alarms that were raised in the last seven days are pulled in from DCNM.

The DCNM application on ServiceNow runs scheduled scripts and connects with the mid-server which in turn connects with DCNM to retrieve data. DCNM sends the requested data to the mid-server which then passes on the data to the DCNM application on ServiceNow. The tables in the DCNM instance on ServiceNow are then populated with this retrieved data.

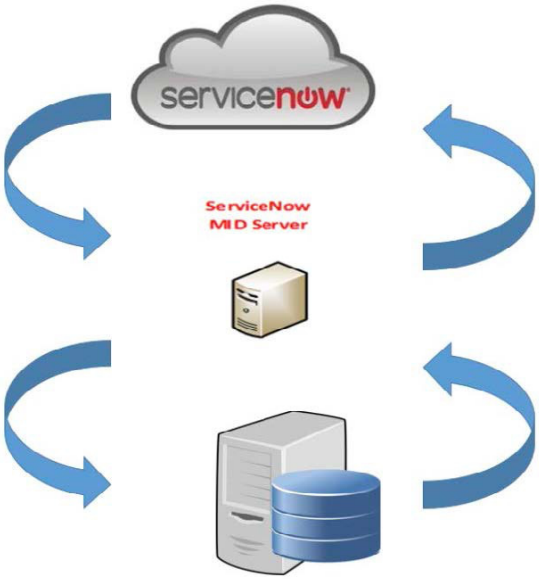

Guidelines and Limitations of DCNM Integration with ServiceNow

- In the ServiceNow Cisco DCNM Application version 1.0, details about only one MID server can be added in the **Cisco DCNM>Properties** table. Starting from Cisco DCNM Application version 1.1, multiple MID servers can be added in the **Cisco DCNM>Properties** table. This means that data can be retrieved from multiple DCNM setups at the same time. In the ServiceNow GUI, data from each DCNM is distinguished by the DCNM IP address.

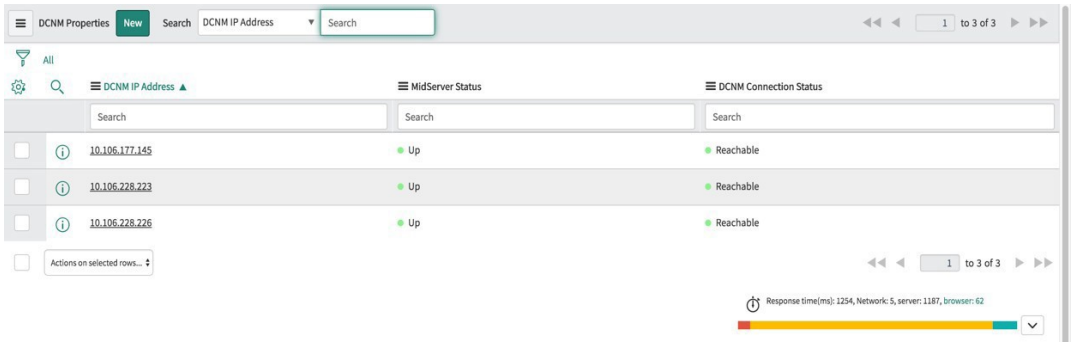

- Scheduled scripts to retrieve data are run only after insertion of a server record in the **Cisco DCNM>Properties** table.
- In case the mid-server IP Address and credentials in the **Cisco DCNM>Properties** table are changed, the data that was imported using the previous mid-server is deleted from the application scope tables. However, data that was imported to the ServiceNow CMDB (global scope) remains and is not deleted.
- To ensure optimal performance in the ServiceNow database, each entry is matched with the switch database ID and IP Address ensuring that there is no duplication of entries.
- Entries in the cmdb_ci_ip_switch table have to be manually deleted in case a new server is added in the **Cisco DCNM>Properties** table.

Installing and Configuring the Cisco DCNM Application on ServiceNow

Procedure

Step 1 Log in to https://dcnm1.service-now.com. Select **System Applications** > **Applications**. Install the **Cisco DCNM Application** from the **All Apps** tab.

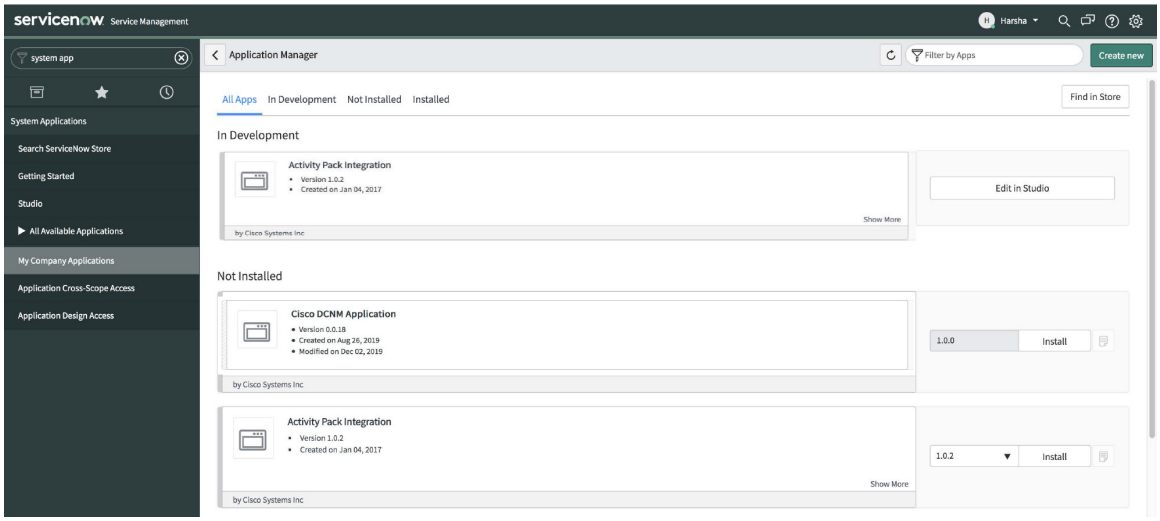

Step 2 After installation is complete, verify that the Cisco DCNM Properties and Dashboard tabs are appearing in the application.

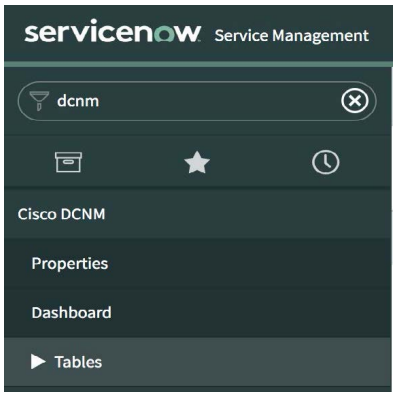

Step 3 Choose **MID Servers** and click the MID Server that is used for DCNM integration.

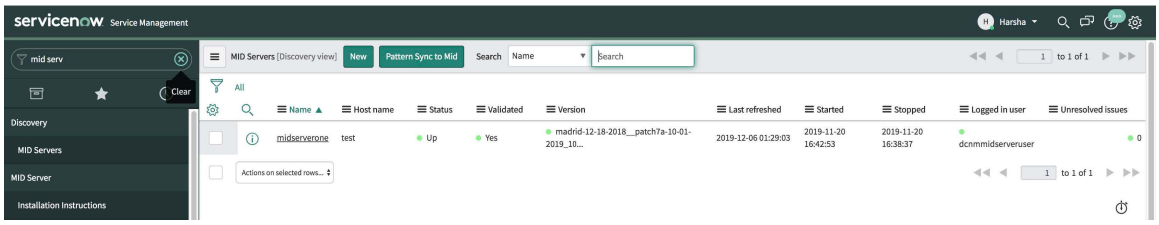

Step 4 Scroll down and click the **Properties** tab. Click **New** and add the property given below in the **MID Server Property New record** window. Click **Submit**.

Name	Type	Value
glide.http.outbound.max_timeout.enabled	True/false	False

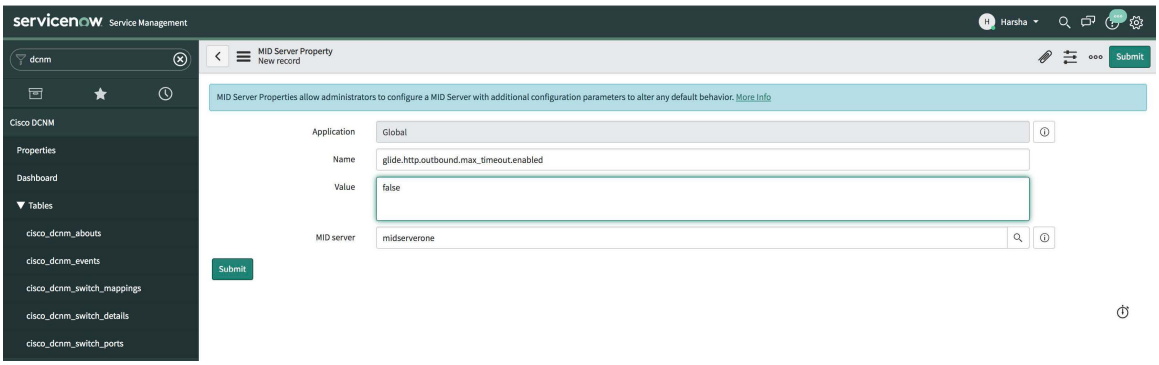

Step 5 Now, select the **Configuration Parameters** tab.

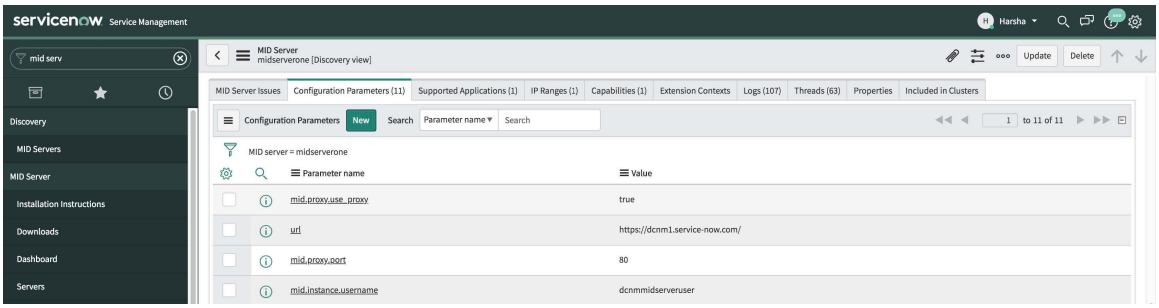

Step 6 In the **Configuration Parameters** tab, click **New**. Enter the required details in the fields.

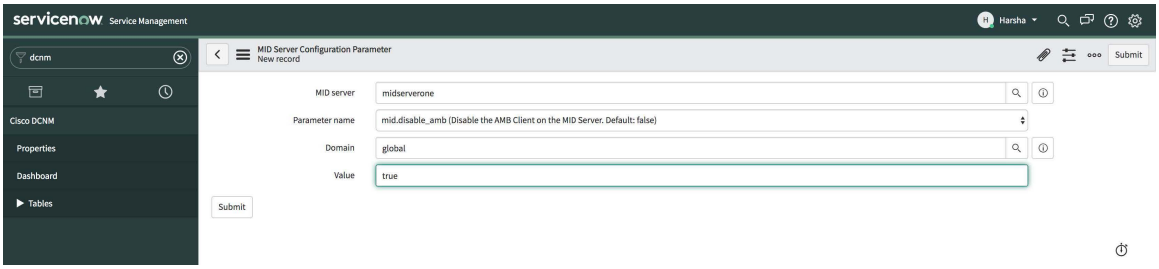

Step 7 Click **Submit** to set up the MID Server.

Step 8 Choose **Cisco DCNM > Properties**. Click **New Server**. Enter the required parameters.

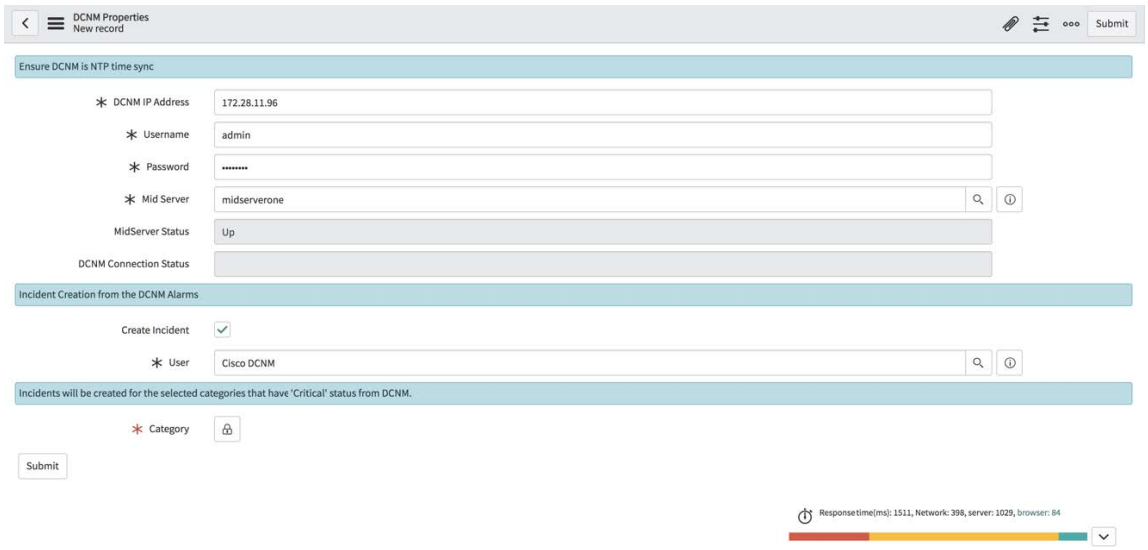

DCNM IP Address - IP Address of the DCNM.

Username - Enter the username used to log in to DCNM.

Password - Enter the password used to log in to DCNM.

Note Access should be provided only for DCNM admins.

Mid server - Specify the name of the mid server to be used. The name is auto-populated as you type. You can also click the search icon next to this field to bring the MID Servers window. You can then select a MID Server from the list that is displayed.

MidServer Status - Indicates whether the MID server is up or down.

DCNM Connection Status - Indicates whether the DCNM IP address that has been provided is reachable or not to retrieve data. This status field is populated when you click **Submit** after you have entered the required information. **Reachable** is displayed on successful communication with DCNM, and **Unreachable**, in case the connection is unsuccessful.

Create Incident - Select this checkbox in case you need incidents to be raised automatically for alarm events.

User - Create a new user and add the user name in this field. The Caller field in the incidents that are created is populated with this user name. This field is auto-populated as you type. You can also click the search icon next to this field to bring the Users window. You can then select a user from the list that is displayed.

Category - Click the lock icon to create incidents automatically for specific categories only.

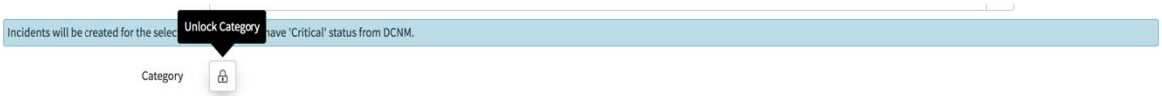

Select the required category for which incidents have to be created from the drop-down list below the **Category** window. The available categories for creation of incidents are CPU, DEVICE_ACCESS_SNMP,DEVICE_ACCESS_SSH, and MEMORY. Refer the following table for more information on this.

Table 26: Events & Incidents

Category	Data Collection in ServiceNow	Incident Raised	Incident Rule	ServiceNow Incident details
CPU	Yes	Yes	DCNM Alarm severity = 'Critical'	Priority = 2 Urgency = 2 Impact = 2
Memory	Yes	Yes	DCNM Alarm severity = 'Critical'	Priority = 2 Urgency = 2 Impact = 2
Power	Yes	No	NA	NA
Linkstate	Yes	No	NA	NA
ICMP	Yes	No	NA	NA
SNMP	Yes	Yes	DCNM Alarm severity = 'Critical'	Priority = 2 Urgency = 2 Impact = 2
SSH	Yes	Yes	DCNM Alarm severity = 'Critical'	Priority = 2 Urgency = 2 Impact = 2

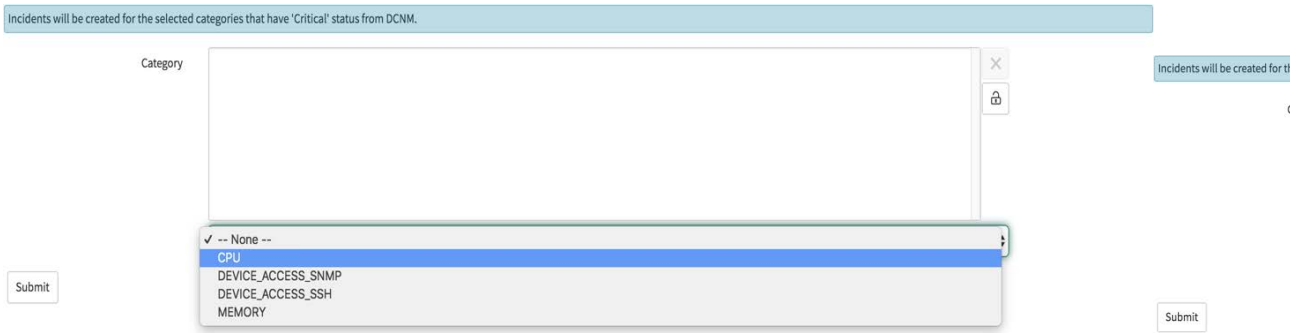

Now, click **Submit**.

Viewing the Dashboard

Choose **Cisco DCNM>Dashboard** to display the dashboard. The **DCNM IP Address**, the **DCNM Connection Status** and the **MidServer Status** are displayed at the top of the dashboard.

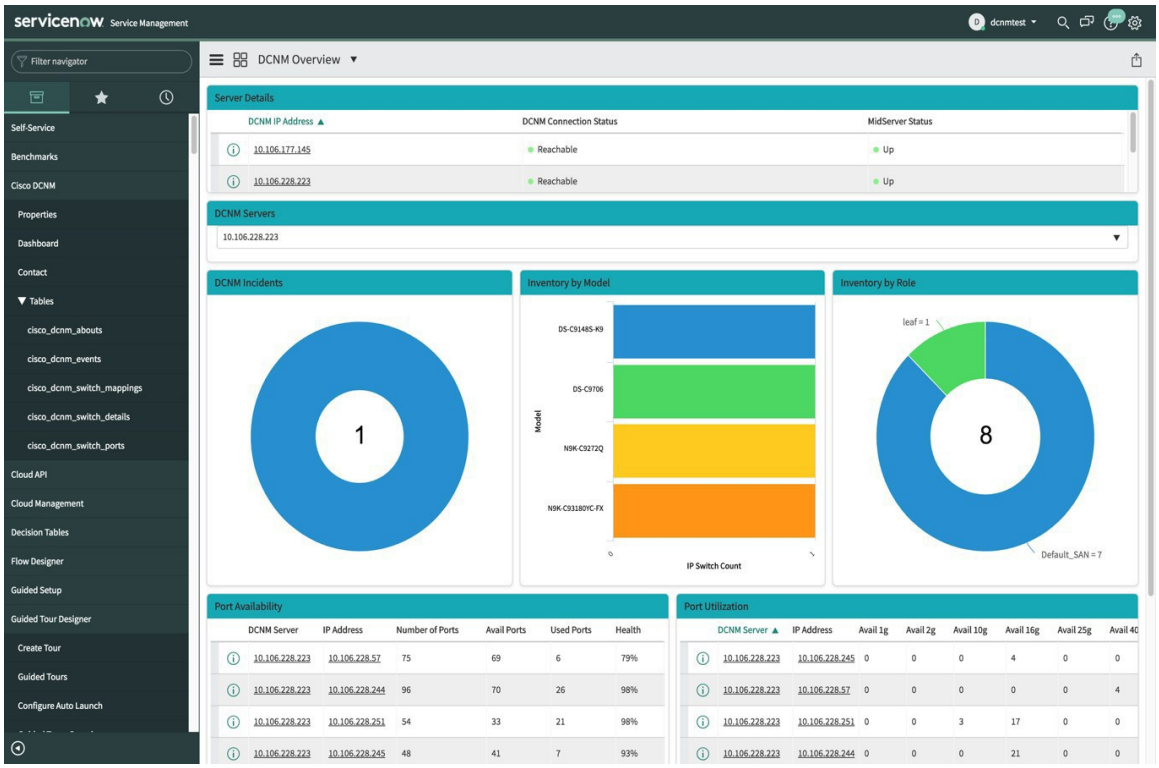

The **DCNM Servers** section displays the IP address of the DCNM server from which the data is being retrieved and displayed. Click the dropdown list to select any other DCNM server as per your requirement.

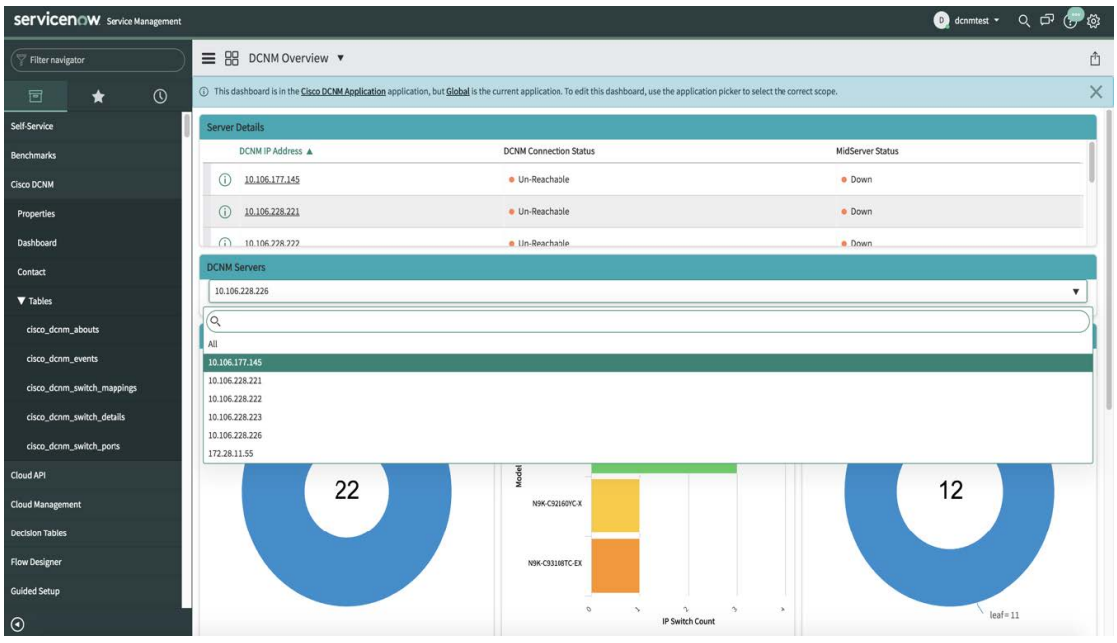

Click **All** to retrieve and display data from all the DCNM Servers that are displayed in the dropdown list. When the **All** option is selected, the number of incidents that are displayed in the DCNM Incidents donut are color-coded and displayed based on the different DCNM server IP addresses. The Inventory by Model and

Inventory by Role donuts also display data from all the DCNM servers. The Port Availability and Port Utilization donuts display data along with the DCNM Server that each IP address belongs to.

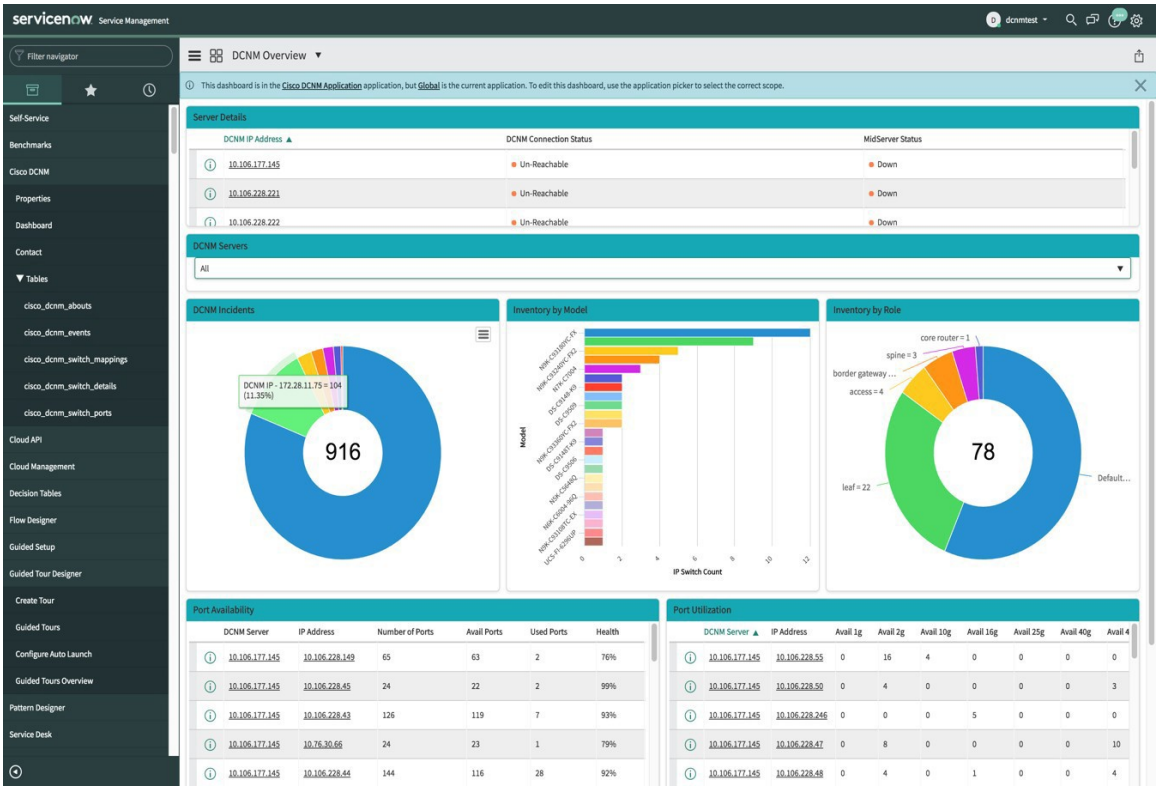

DCNM Incidents - This displays the number of incidents that have been raised based on the alarms retrieved from DCNM. Click the donut for more details about the

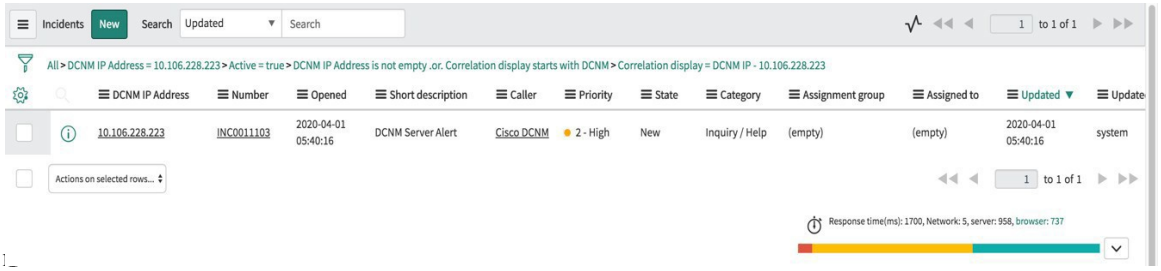

Inventory by Model - This displays the number and type of switches present in DCNM. Each band represents a device model. Click a band for more

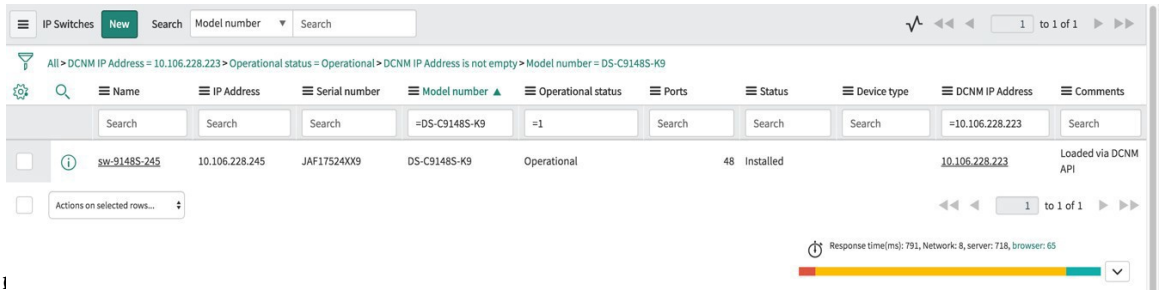

Inventory by Role - This displays the number and types of switch roles present in DCNM. Click the required section to display the number of roles that are operational and click on that pictorial representation to display more details about the roles.

Note The number that is displayed in the Inventory by Role donut does not change in case switches are removed from DCNM. The switches that are removed are displayed as Non Operational and there is no change in the number that is displayed in the donut.

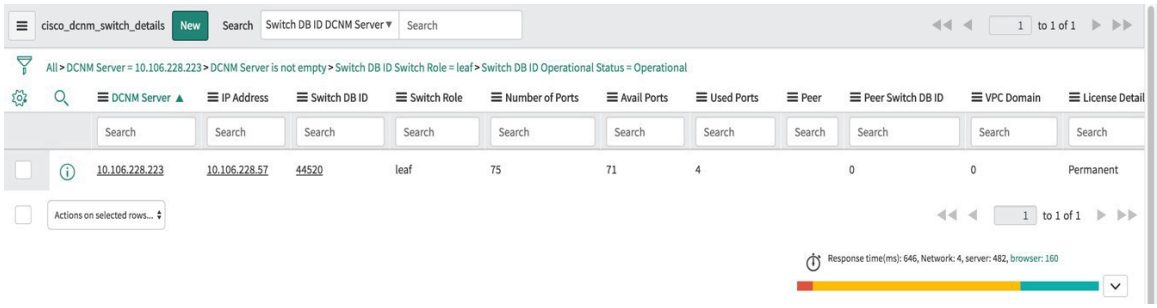

Port Availability - This displays information about port availability. The DCNM server and IP address along with the total number of ports, available ports, used ports and health of the switch is displayed. Click an IP address to display more

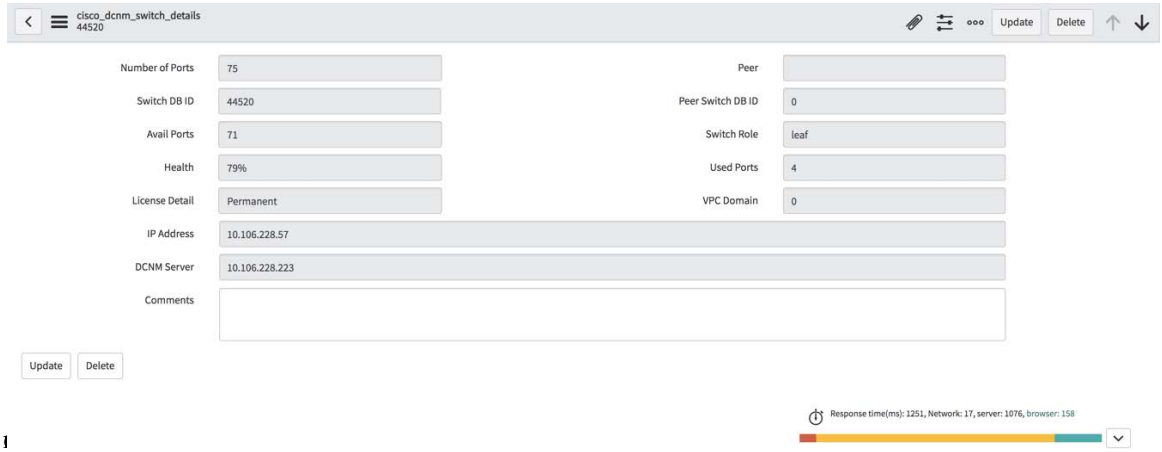

Port Utilization - This displays information about port utilization based on each IP address. The number of ports having 1G, 2G, 4G, 8G, 10G, 16G, 25G, 32G, 40G, and 100G availability, are displayed. Click an IP

address to display more

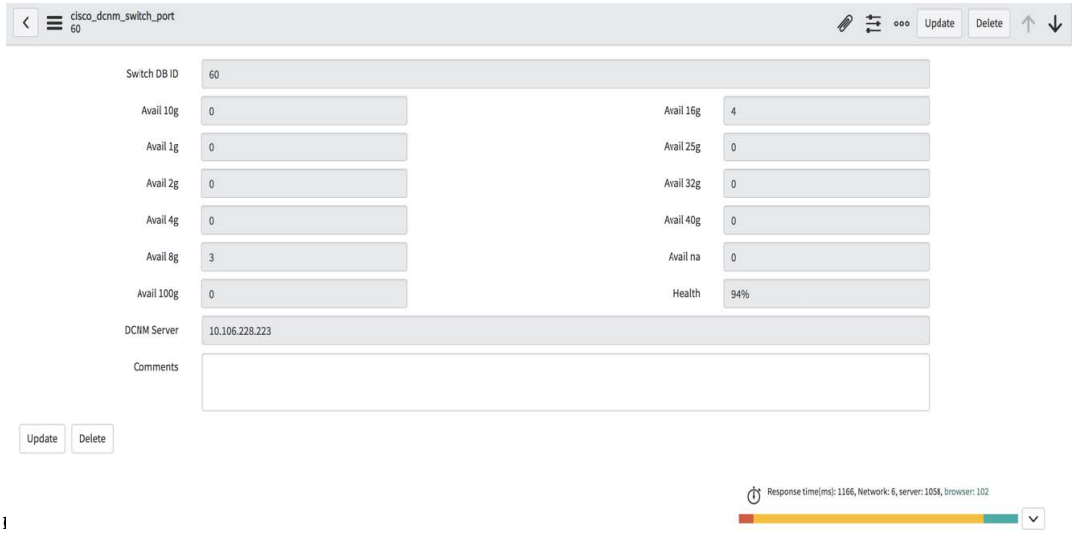

Contact Us

Choose **Cisco DCNM>Contact** to display an email address and a telephone number that can be used to contact Cisco Systems for any queries.

Troubleshooting DCNM Integration with ServiceNow

In case data is not being retrieved in the ServiceNow table:

- Check if the MID server is up or down.
- Check for information entries in system logs with the source “x_caci_cisco_dcnm”.
- Check the login credentials added in Cisco DCNM Properties.
- Consider a scenario in which data is being displayed on the ServiceNow dashboard for the selected DCNM server and then you want to display data for another DCNM server. In such a scenario, the ServiceNow dashboard may take some time to load data from the other DCNM server due to a delay in refreshing the cache. To refresh the data manually, click the **Refresh** icon that appears on the top right corner of the individual tiles when you hover the cursor over the tiles.

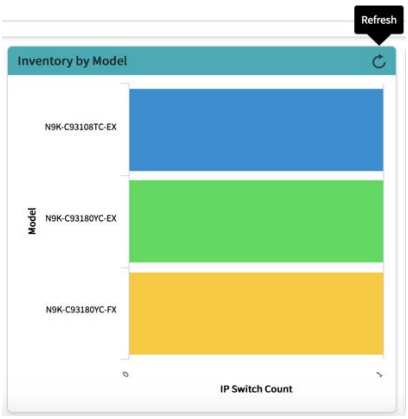

You can also refresh the whole dashboard by clicking on the **Dashboard Controls** icon and then clicking **Refresh** to load the reports correctly.

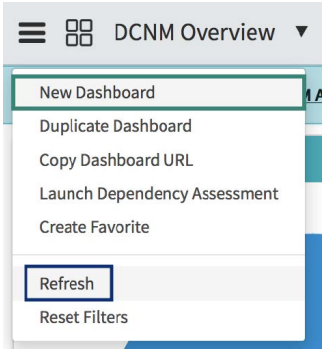

For more information on DCNM application integration with ServiceNow, click here.

PART II

Easy Provisioning of VXLAN BGP EVPN Fabrics

- Managing a Greenfield VXLAN BGP EVPN Fabric, on page 643
- Managing a Brownfield VXLAN BGP EVPN Fabric, on page 695
- Configuring a VXLANv6 Fabric, on page 757
- Auto-Provisioning ToR Switches Attached to VXLAN VTEPs, on page 763

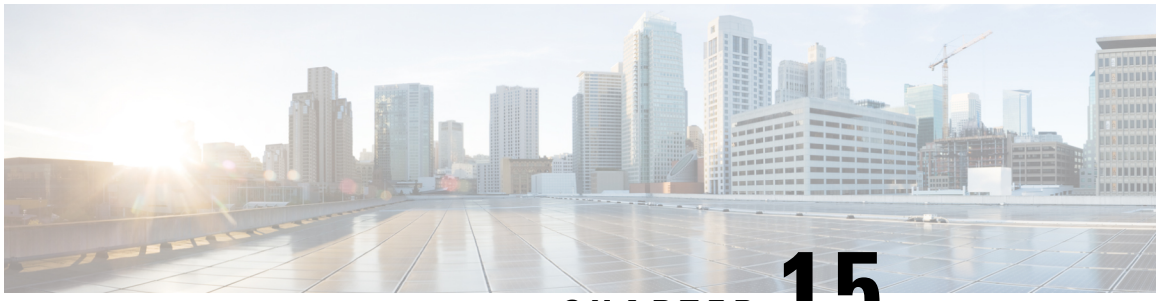

CHAPTER 15

Managing a Greenfield VXLAN BGP EVPN Fabric

This chapter describes how to manage a greenfield VXLAN BGP EVPN fabric.

- VXLAN BGP EVPN Fabrics Provisioning, on page 643
- Creating a New VXLAN BGP EVPN Fabric, on page 646
- Adding Switches to a Fabric, on page 665
- VXLAN EVPN Deployment with eBGP EVPN, on page 678

VXLAN BGP EVPN Fabrics Provisioning

DCNM 11 introduces an enhanced “Easy” fabric workflow for unified underlay and overlay provisioning of VXLAN BGP EVPN configuration on Nexus 9000 and 3000 series of switches. The configuration of the fabric is achieved via a powerful, flexible, and customizable template-based framework. Using minimal user inputs, an entire fabric can be brought up with Cisco recommended best practice configurations, in a short period of time. The set of parameters exposed in the Fabric Settings allow users to tailor the fabric to their preferred underlay provisioning options.

Border devices in a fabric typically provide external connectivity via peering with appropriate edge/core/WAN routers. These edge/core routers may either be managed or monitored by DCNM. These devices are placed in a special fabric called the External Fabric. The same DCNM controller can manage multiple VXLAN BGP EVPN fabrics while also offering easy provisioning and management of Layer-2 and Layer-3 DCI underlay and overlay configuration among these fabrics using a special construct called a Multi-Site Domain (MSD) fabric.

Note that in this document the terms switch and device are used interchangeably.

The DCNM GUI functions for creating and deploying VXLAN BGP EVPN 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).

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]).

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 (ST-FEX), Active-Active FEX (AA-FEX), loopback, subinterface, etc.
- Create breakout and unbreakout ports.
- Shut down and bring up interfaces.
- Rediscover ports and view interface configuration history.

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.

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

Provisioning of configuration on service leafs to which L4-7 service appliances may be attached. For more information, see *L4-L7 Service Basic Workflow*.

This chapter mostly covers configuration provisioning for a single VXLAN BGP EVPN fabric. EVPN Multi-Site provisioning for Layer-2/Layer-3 DCI across multiple fabrics using the MSD fabric, is documented in a separate chapter. The deployment details of how overlay Networks and VRFs can be easily provisioned from the DCNM, is covered under Creating and Deploying Networks and VRFs .

Guidelines for VXLAN BGP EVPN Fabrics Provisioning

- For any switch to be successfully imported into DCNM, the user specified for discovery/import, should have the following permissions:
  - SSH access to the switch
  - Ability to perform SNMPv3 queries
  - Ability to run the **show** commands including show run, show interfaces, etc.
- The switch discovery user need not have the ability to make any configuration changes on the switches. It is primarily used for read access.

- 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.

- LAN credentials are required to be set of any user that needs to be perform any write access to the device. LAN credentials need to be set on the DCNM, on a per user per device basis. When a user imports a device into the Easy Fabric, and LAN credentials are not set for that device, DCNM moves this device to a migration mode. Once the user sets the appropriate LAN credentials for that device, a subsequent Save & Deploy will retrigger the device import process.

- The **Save & Deploy** button triggers the intent regeneration for the entire fabric as well as a configuration compliance check for all the switches within the fabric. This button is required but not limited to the following cases:

  - A switch or a link is added, or any change in the topology
  - A change in the fabric settings that must be shared across the fabric
  - A switch is removed or deleted
  - A new vPC pairing or unpairing is done
  - A change in the role for a device

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

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

- 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
  ```

  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*.

  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.

- 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.

This procedure contains descriptions for the IPv4 underlay. For information about IPv6 underlay, see IPv6 Underlay Support for Easy Fabric, on page 114.

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

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

   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 appear.

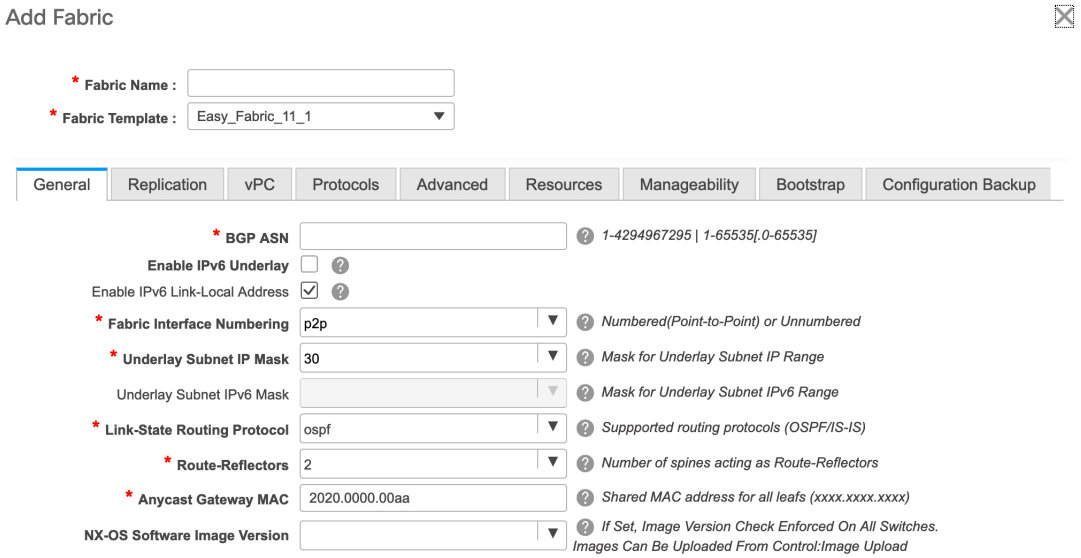

The tabs and their fields in the screen are explained in the subsequent points. The overlay and underlay network parameters are included in these tabs.

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

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

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

   Enable IPv6 Underlay: Enable the IPv6 underlay feature. For information, see IPv6 Underlay Support for Easy Fabric, on page 114.

   Enable IPv6 Link-Local Address: Enables the IPv6 Link-Local address.

   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.

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

   Route-Reflectors (RRs) – 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 two spine devices designated as RRs.

   Decreasing the count - When you reduce four route reflectors to two, 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 & Deploy** in the fabric topology window.

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, and save the Fabric Settings, the system checks that all the switches within the fabric have the selected version. If some devices do not run the image, a warning is prompted to perform an In-Service Software Upgrade (ISSU) to the specified image. The warning is also accompanied with a Resolve button. This takes the user to the image management screen with the mismatched switches auto selected for device upgrade/downgrade to the specified NX-OS image specified in Fabric Settings. Till, all devices run the specified image, the deployment process is incomplete.

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

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

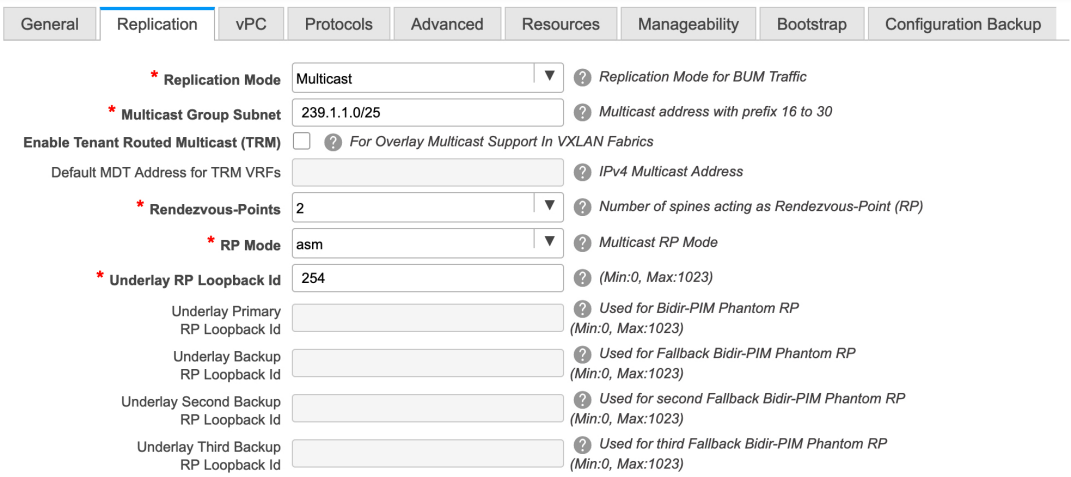

Replication Mode : The mode of replication that is used in the fabric for BUM (Broadcast, Unknown Unicast, Multicast) traffic. The choices are Ingress Replication or Multicast. When you choose Ingress replication, the multicast related fields get disabled.

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. A 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 check box to enable Tenant Routed Multicast (TRM) that allows overlay multicast traffic to be supported over EVPN/MVPN in the VXLAN BGP EVPN fabric.

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**.

For more information, see Overview of Tenant Routed Multicast, on page 196.

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 sof release 9.2(1) onwards.

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

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

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

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

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

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

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

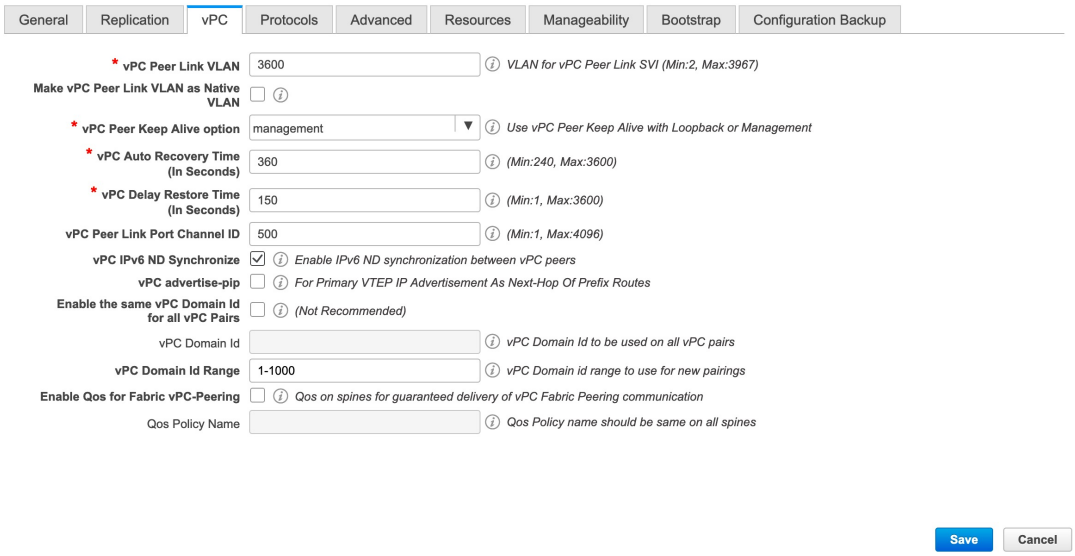

vPC Peer Link VLAN – VLAN used for the vPC peer link SVI.

Make vPC Peer Link VLAN as Native VLAN - Enables vPC peer link VLAN as Native VLAN.

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 ID - 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.

You can enable the advertise PIP feature on a specific vPC as well. For more information, see Advertising PIP on vPC, on page 240.

Enable the same vPC Domain Id for all vPC Pairs: Enable the same vPC Domain ID for all vPC pairs. When you select this field, the **vPC Domain Id** field is editable.

vPC Domain Id - Specifies the vPC domain ID to be used on all vPC pairs.

vPC Domain Id Range - Specifies the vPC Domain Id range to use for new pairings.

Enable Qos for Fabric vPC-Peering - Enable QoS on spines for guaranteed delivery of vPC Fabric Peering communication. For more information, see QoS for Fabric vPC-Peering, on page 233.

Note QoS for vPC fabric peering and queuing policies options in fabric settings are mutually exclusive.

Qos Policy Name - Specifies QoS policy name that should be same on all fabric vPC peering spines. The default name is **spine_qos_for_fabric_vpc_peering**.

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

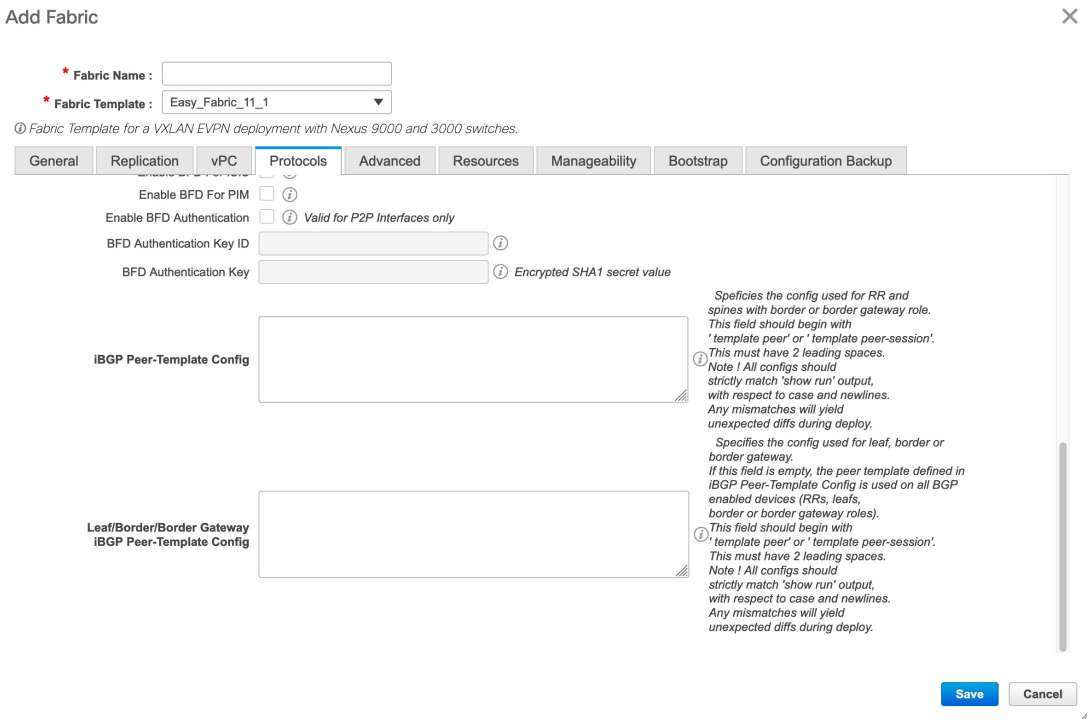

Underlay Routing Loopback Id - The loopback interface ID is populated as 0 since loopback0 is usually used for fabric underlay IGP peering purposes.

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

Underlay 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 **Underlay Routing Protoc** 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. Log in to the switch, retrieve the encrypted key and enter it in this field. Refer, *Retrieving the Authentication Key* section for details.

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, such as CiscoisisAuth.

IS-IS Authentication Key ID - The Key ID is populated.

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

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

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 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. Log in to the switch, retrieve the encrypted key and enter it in BGP Authentication Key field. Refer the Retrieving the Authentication Key section for details.

Enable PIM Hello Authentication - Enables the PIM hello authentication.

PIM Hello Authentication Key - Specifies the PIM hello authentication key.

Enable BFD: Select the check box to enable **feature bfd** on all switches in the fabric. This feature is valid only on IPv4 underlay and the scope is within a fabric.

From Cisco DCNM Release 11.3(1), BFD within a fabric is supported natively. The BFD feature is disabled by default in the Fabric Settings. If enabled, BFD is enabled for the underlay protocols with the default settings. Any custom required BFD configurations must be deployed via the per switch freeform or per interface freeform policies.

The following config is pushed after you select the **Enable BFD** check box:

```
feature bfd
```

Note After you upgrade from DCNM Release 11.2(1) with BFD enabled to DCNM Release 11.3(1), the follo configurations are pushed on all P2P fabric interfaces:

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

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 be 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 set to **unnumbered**. The BFD authentication fields will be grayed out automatically. BFD authentication valid for only for P2P interfaces.

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 Encrypted BFD Authentication Key, on page 249.

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
```

Until Cisco DCNM Release 11.3(1), iBGP peer template for iBGP definition on the leafs or border role devices and BGP RRs were same. From DCNM Release 11.4(1), the following fields can be used to specify different configurations:

- **iBGP Peer-Template Config** – Specifies the config used for RR and spines with border role.
- **Leaf/Border/Border Gateway iBGP Peer-Template Config** – Specifies the config used for leaf, border, or border gateway. If this field is empty, the peer template defined in **iBGP Peer-Template Config** is used on all BGP enabled devices (RRs, leafs, border, or border gateway roles).

In brownfield migration, if the spine and leaf use different peer template names, both **iBGP Peer-Template Config** and **Leaf/Border/Border Gateway iBGP Peer-Template Config** fields need to be set according to the switch config. If spine and leaf use the same peer template name and content (except for the “route-reflector-client” CLI), only **iBGP Peer-Template Config** field in fabric setting needs to be set. If the fabric settings on iBGP peer templates do not match the existing switch configuration, an error message is generated and the migration will not proceed.

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

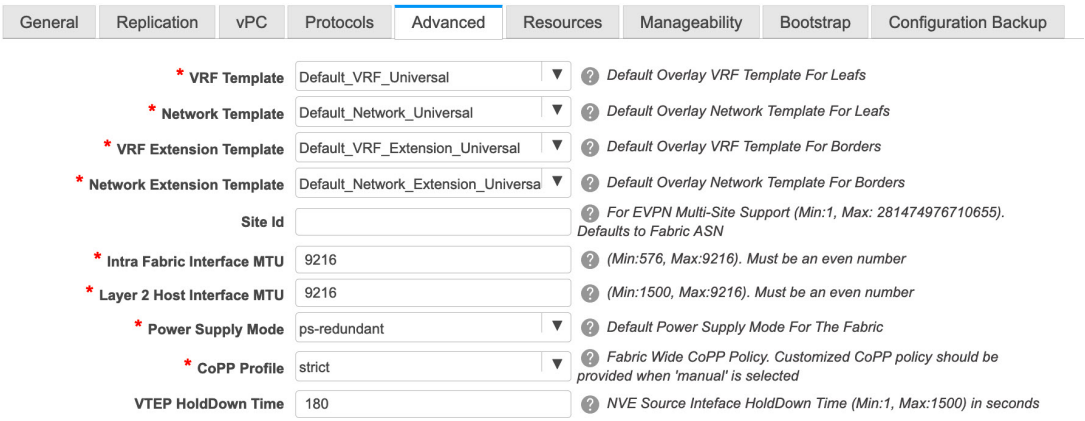

VRF Template and **VRF Extension Template**: Specifies the VRF template for creating VRFs, and the VRF extension template for enabling VRF extension to other fabrics.

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.

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

VTEP HoldDown Time - Specifies the NVE source interface hold down time.

Brownfield Overlay Network Name Format: Enter the format to be used to build the overlay network name during a brownfield import or migration. The network name should not contain any white spaces or special characters except underscore (_) and hyphen (-). The network name must not be changed once the brownfield migration has been initiated. See the *Creating Networks for the Standalone Fabric* section for the naming convention of the network name. The syntax is **[<string> | $$VLAN_ID$$] $$VNI$$ [<string>| $$VLAN_ID$$]** and the default value is **Auto_Net_VNI$$VNI$$_VLAN$$VLAN_ID$$**. When you create networks, the name is generated according to the syntax you specify. The following table describes the variables in the syntax.

Variables	Description
$$VNI$$	Specifies the network VNI ID found in the switch configuration. This is a mandatory keyword required to create unique network names.
$$VLAN_ID$$	Specifies the VLAN ID associated with the network. VLAN ID is specific to switches, hence DCNM picks the VLAN ID from one of the switches, where the network is found, randomly and use it in the name. We recommend not to use this unless the VLAN ID is consistent across the fabric for the VNI.
<string>	This variable is optional and you can enter any number of alphanumeric characters that meet the network name guidelines.

Example overlay network name: Site_VNI12345_VLAN1234

Note Ignore this field for greenfield deployments. The **Brownfield Overlay Network Name Format** applies fo the following brownfield imports:

- CLI-based overlays
- Configuration profile-based overlay where the configuration profiles were created in Cisco DCNM Release

  10.4(2).

Enable CDP for Bootstrapped Switch - Enables CDP on management (mgmt0) interface for bootstrapped switch. By default, for bootstrapped switches, CDP is disabled on the mgmt0 interface.

Enable VXLAN OAM - Enables the VXLAM OAM functionality for devices in the fabric. 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 check box to enable feature dhcp and associated configurations globally on all switches in the fabric. This is a pre-requisite for support of DHCP for overlay networks that are part of the tenant VRFs.

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

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

Enable NX-API on HTTP on Port - Specifies enabling of NX-API on HTTP. Enable this check box and the **Enable NX-API** check box to use HTTP. This check box is checked by default. If you uncheck this check box, the applications that use NX-API and supported by Cisco DCNM, such as Endpoint Locator (EPL), Layer 4-Layer 7 services (L4-L7 services), VXLAN OAM, and so on, start using the HTTPS instead of HTTP.

Note If you check the **Enable NX-API** check box and the **Enable NX-API on HTTP** check box, applicatio HTTP.

Enable Policy-Based Routing (PBR) - Select this check box to enable routing of packets based on the specified policy. Starting with Cisco NX-OS Release 7.0(3)I7(1) and later releases, this feature works on Cisco Nexus 9000 Series switches with Nexus 9000 Cloud Scale (Tahoe) ASICs. This feature is used along with the Layer 4-Layer 7 service workflow. For information on Layer 4-Layer 7 service, refer the *Layer 4-Layer 7 Service* chapter.

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.

Enable AAA IP Authorization - Enables AAA IP authorization, when IP Authorization is enabled in the remote authentication server. This is required to support DCNM in scenarios where customers have strict control of which IP addresses can have access to the switches.

Enable DCNM as Trap Host - Select this check box to enable DCNM as a SNMP trap destination. Typically, for a native HA DCNM deployment, the eth1 VIP IP address will be configured as SNMP trap destination on the switches. By default, this check box is enabled.

Greenfield Cleanup Option – Enable the switch cleanup option for switches imported into DCNM with Preserve-Config=No, without a switch reload. This option is typically recommended only for the fabric environments with Cisco Nexus 9000v Switches to improve on the switch clean up time. The recommended option for Greenfield deployment is to employ Bootstrap or switch cleanup with a reboot. In other words, this option should be unchecked.

Enable Precision Time Protocol (PTP): Enables PTP across a fabric. When you select this check box, PTP is enabled globally and on core-facing interfaces. Additionally, the **PTP Source Loopback Id** and **PTP Domain Id** fields are editable. For more information, see Precision Time Protocol for Easy Fabric, on page 100.

PTP Source Loopback Id: Specifies the loopback interface ID Loopback that is used as the Source IP Address for all PTP packets. The valid values range from 0 to 1023. The PTP loopback ID cannot be the same as RP, Phantom RP, NVE, or MPLS loopback ID. Otherwise, an error will be generated. The PTP loopback ID can be the same as BGP loopback or user-defined loopback which is created from DCNM.

If the PTP loopback ID is not found during **Save & Deploy**, the following error is generated:

Loopback interface to use for PTP source IP is not found. Create PTP loopback interface on all the devices to enable PTP feature.

PTP Domain Id: Specifies the PTP domain ID on a single network. The valid values range from 0 to 127.

Enable MPLS Handoff: Select the check box to enable the MPLS Handoff feature. For more information, see the *Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - MPLS SR and LDP Handoff* chapter.

Underlay MPLS Loopback Id: Specifies the underlay MPLS loopback ID. The default value is 101.

Enable TCAM Allocation: TCAM commands are automatically generated for VXLAN and vPC Fabric Peering when enabled.

Enable Default Queuing Policies: Check this check box to apply QoS policies on all the switches in this fabric. To remove the QoS policies that you applied on all the switches, uncheck this check box, update all the configurations to remove the references to the policies, and save and deploy. From Cisco DCNM Release 11.3(1), pre-defined QoS configurations are included that can be used for various Cisco Nexus 9000 Series Switches. When you check this check box, the appropriate QoS configurations are pushed to the switches in the fabric. The system queuing is updated when configurations are deployed to the switches. You can perform the interface marking with defined queuing policies, if required, by adding the required configuration to the per interface freeform block.

From Cisco DCNM Release 11.4(1), the DSCP mapping for QoS 5 has changed from 40 to 46 in the policy template. For DCNM 11.3(1) deployments that have been upgraded to 11.4(1), you will see the diffs that need to be deployed.

Review the actual queuing policies by opening the policy file in the template editor. From Cisco DCNM Web UI, choose **Control > Template Library**. Search for the queuing policies by the policy file name, for example, **queuing_policy_default_8q_cloudscale**. Choose the file and click the **Modify/View template** icon to edit the policy.

See the *Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide* for platform specific details.

N9K Cloud Scale Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all Cisco Nexus 9200 Series Switches and the Cisco Nexus 9000 Series Switches that ends with EX, FX, and FX2 in the fabric. The valid values are **queuing_policy_default_4q_cloudscale** and **queuing_policy_default_8q_cloudscale**. Use the **queuing_policy_default_4q_cloudscale** policy for FEXes. You can change from the **queuing_policy_default_4q_cloudscale** policy to the **queuing_policy_default_8q_cloudscale** policy only when FEXes are offline.

N9K R-Series Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all Cisco Nexus switches that ends with R in the fabric. The valid value is **queuing_policy_default_r_series**.

Other N9K Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all other switches in the fabric other than the switches mentioned in the above two options. The valid value is **queuing_policy_default_other**.

Enable MACsec - Enables MACsec for the fabric. For more information, see MACsec Support in Easy Fabric and eBGP Fabric, on page 194.

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. For more information, refer Enabling Freeform Configurations on Fabric Switches .

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

Intra-fabric Links Additional Config - Add CLIs that should be added to the intra-fabric links.

8. Click the **Resources** tab.

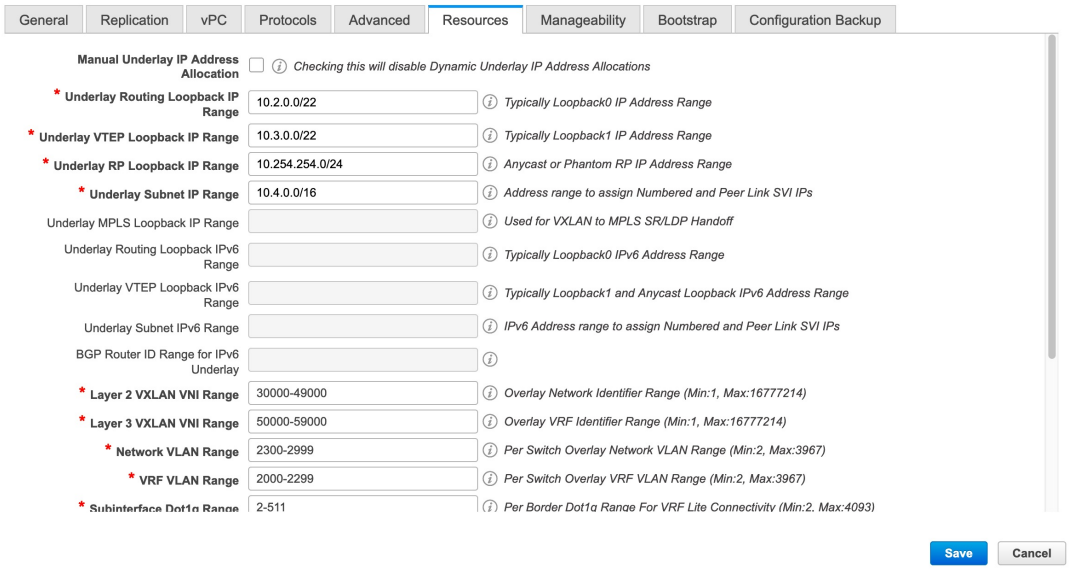

Manual Underlay IP Address Allocation – *Do not* select this check box if you are transitioning your VXLAN fabric management to DCNM.

- 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.

Underlay MPLS Loopback IP Range: Specifies the underlay MPLS loopback IP address range.

For eBGP between Border of Easy A and Easy B, Underlay routing loopback and Underlay MPLS loopback IP range must be a unique range. It should not overlap with IP ranges of the other fabrics, else VPNv4 peering will not come up.

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.

The VRF Lite Subnet IP Range field specifies resources reserved for IP address used for VRF LITE when VRF LITE IFCs are auto-created. If you select Back2BackOnly, ToExternalOnly, or Back2Back&ToExternal then VRF LITE IFCs are auto-created.

Auto Deploy Both - This check box is applicable for the symmetric VRF Lite deployment. When you select this check box, it would set the auto deploy flag to true for auto-created IFCs to turn on symmetric VRF Lite configuration.

This check box can be selected or deselected when the **VRF Lite Deployment** field is not set to Manual. In the case, a user explicitly unchecks the auto-deploy field for any auto-created IFCs, then the user input is always given the priority. This flag only affects the new auto-created IFC and it does not affect the existing IFCs.

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 upda one range of values at a time. If you want to update more than one range of values, do it in separate instance 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**.

Service Network VLAN Range - Specifies a VLAN range in the Service Network VLAN Range field. This is a per switch overlay service network VLAN range. The minimum allowed value is 2 and the maximum allowed value is 3967.

Route Map Sequence Number Range - Specifies the route map sequence number range. The minimum allowed value is 1 and the maximum allowed value is 65534.

9. Click the **Manageability** tab.

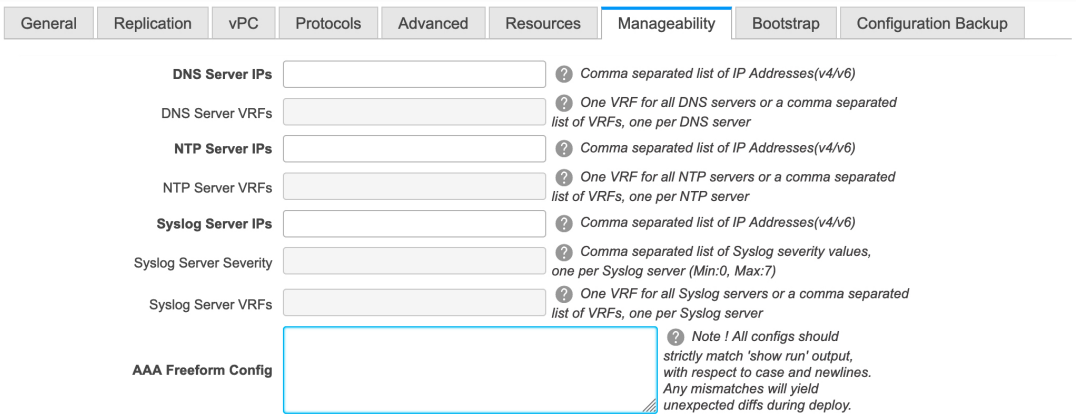

The fields in this tab are:

DNS Server IPs - Specifies the comma separated list of IP addresses (v4/v6) of the DNS servers.

DNS Server VRFs - Specifies one VRF for all DNS servers or a comma separated list of VRFs, one per DNS server.

NTP Server IPs - Specifies comma separated list of IP addresses (v4/v6) of the NTP server.

NTP Server VRFs - Specifies one VRF for all NTP servers or a comma separated list of VRFs, one per NTP server.

Syslog Server IPs – Specifies the comma separated list of IP addresses (v4/v6) IP address of the syslog servers, if used.

Syslog Server Severity – Specifies the comma separated list of syslog severity values, one per syslog server. The minimum value is 0 and the maximum value is 7. To specify a higher severity, enter a higher number.

Syslog Server VRFs – Specifies one VRF for all syslog servers or a comma separated list of VRFs, one per syslog server.

AAA Freeform Config – Specifies the AAA freeform configurations.

If AAA configurations are specified in the fabric settings, **switch_freeform** PTI with source as **UNDERLAY_AAA** and description as **AAA Configurations** will be created.

10. Click the **Bootstrap** tab.

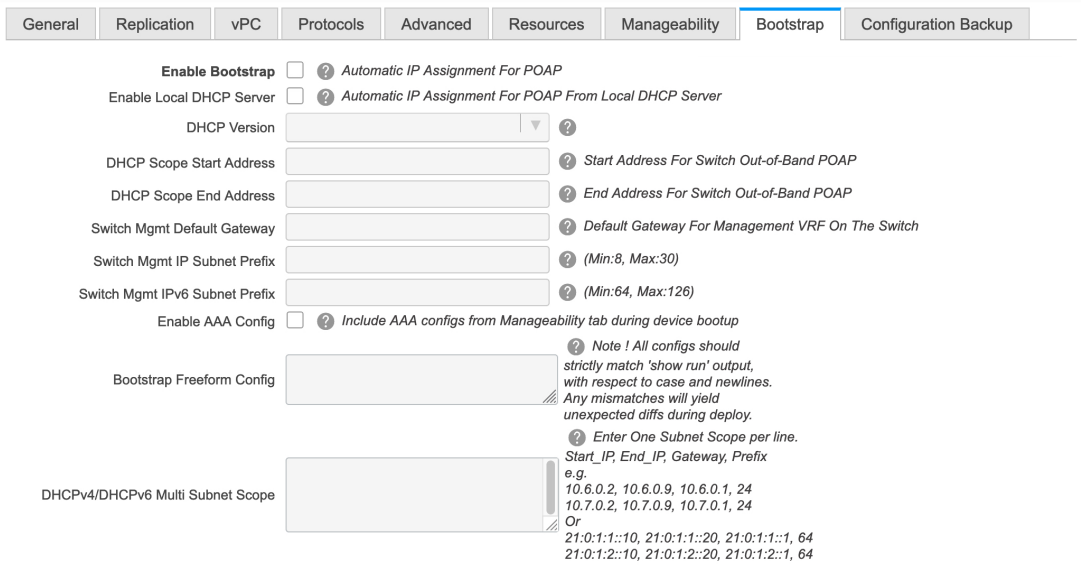

Enable Bootstrap - Select this check box to enable the bootstrap feature. Bootstrap allows easy day-0 import and bring-up of new devices into an existing fabric. Bootstrap leverages the NX-OS POAP functionality.

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 Mgmt Default Gateway** and **Switch Mgmt IP Subnet Prefix** fields.
- Local DHCP Server: Enable the **Local DHCP Server** check box 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 Version – Select DHCPv4 or DHCPv6 from this drop-down list. When you select DHCPv4, the **Switch Mgmt IPv6 Subnet Prefix** field is disabled. If you select DHCPv6, the **Switch Mgmt IP Subnet Prefix** is disabled.

Note Cisco Nexus 9000 and 3000 Series Switches support IPv6 POAP only when switches are either Layer-2 adjacent (eth1 or out-of-band subnet must be a /64) or they are L3 adjacent residing in some IPv6 /64 subn Subnet prefixes other than /64 are not supported.

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

Switch Mgmt IP 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.

Switch Mgmt IPv6 Subnet Prefix - Specifies the IPv6 prefix for the Mgmt0 interface on the switch. The prefix should be between 112 and 126. This field is editable if you enable IPv6 for DHCP.

Enable AAA Config – Select this check box to include AAA configurations from the Manageability tab as part of the device startup config post bootstrap.

Bootstrap Freeform Config - (Optional) Enter additional commands as needed. For example, if you require some additional configurations to be pushed to the device and be available post device bootstrap, they can be captured in this field, to save the desired intent. After the devices boot up, they will contain the configuration 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 347.

DHCPv4/DHCPv6 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

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

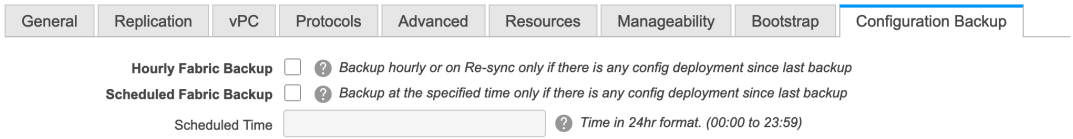

Hourly Fabric Backup: Select the check box to enable an hourly backup of fabric configurations and the intent.

The hourly backups are triggered during the first 10 minutes of the hour.

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.

The backup process is initiated after you click **Save**.

The scheduled backups are triggered exactly at the time you specify with a delay of up to two minutes. The scheduled backups are triggered regardless of the configuration deployment status.

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 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.

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.

(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.
- **Backup Now**: You can initiate a fabric backup manually by clicking **Backup Now**. Enter a name for the tag and click **OK**. Regardless of the settings you choose under the **Configuration Backup** tab in the **Fabric Settings** dialog box, you can initiate a backup using this option.

- **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 and interfaces. For more information, see Pre-provisioning a Device , on page 86 and Pre-provisioning an Ethernet Interface, on page 90.

Note When DCNM discovers a switch with the hostname containing the period character (.), it is treated as a domain-name and truncated. Only the text prior to the period character (.) is considered as a hostname. For example:

- If hostname is **leaf.it.vxlan.bgp.org1-XYZ**, DCNM shows only **leaf**
- If hostname is **leaf-itvxlan.bgp.org1-XYZ**, DCNM shows only **leafit-vxlan**

Discovering New Switches

1. When a new Cisco NX-OS device is powered on, typically that device has no startup configuration or any configuration state for that matter. Consequently, it powers on with NX-OS and post initialization, goes into a POAP loop. The device starts sending out DHCP requests on all the interfaces that are up including the mgmt0 interface.

2. As long as there is IP reachability between the device and the DCNM, the DHCP request from the device, will be forwarded to the DCNM. For easy day-0 device bring-up, the bootstrap options should be enabled in the **Fabric Settings** as mentioned earlier.

3. With bootstrap enabled for the fabric, the DHCP request coming from the device will be serviced by the DCNM. The temporary IP address allocated to the device by the DCNM will be employed to learn basic information about the switch including the device model, device NX-OS version, etc.

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

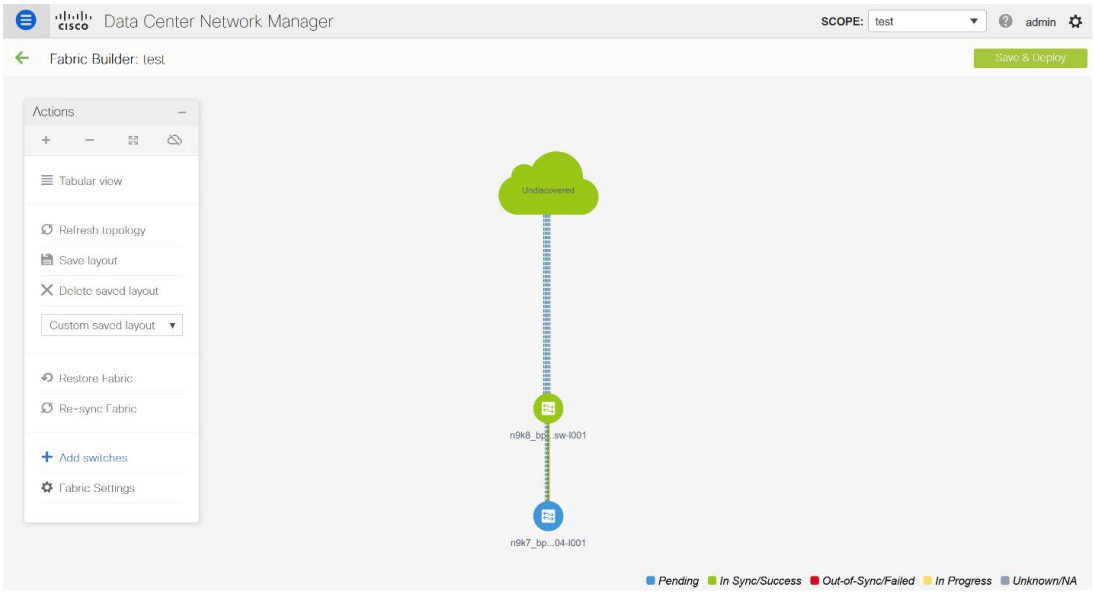

Go to the fabric topology window and click the **Add switches** option from the **Actions** panel. The Inventory Management window comes up.

5. Click the **POAP** tab.

As mentioned earlier, DCNM retrieves the serial number, model number, and version from the device 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

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

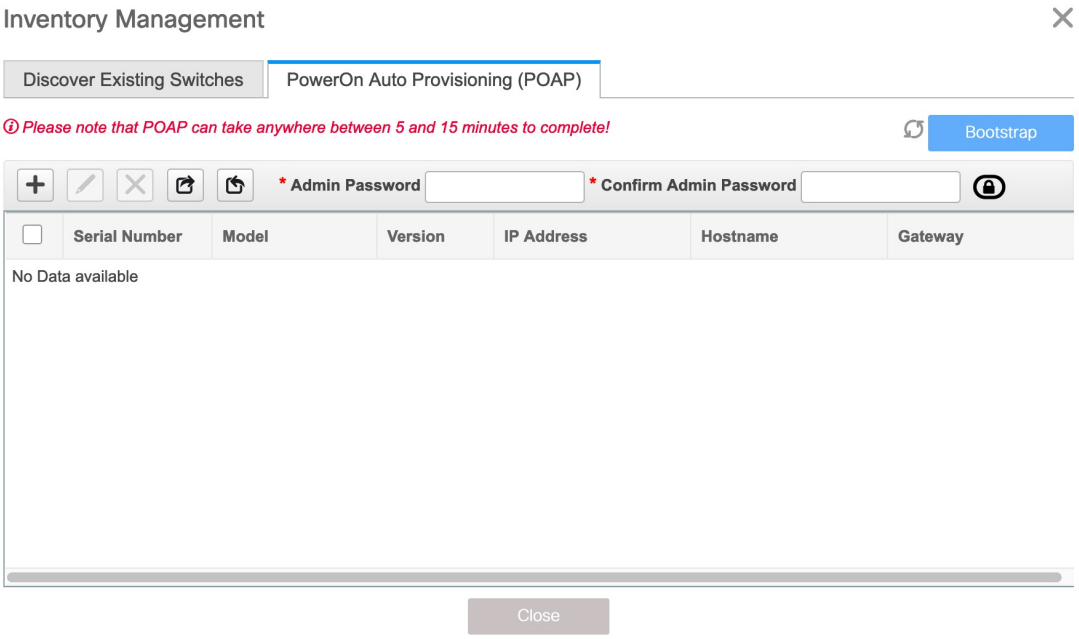

Select the checkbox next to the switch and enter the switch credentials: IP address and host name.

Based on the IP address of your device, you can either add the IPv4 or IPv6 address in the **IP Address** field.

Beginning with Release 11.2(1), you can provision devices in advance. To pre-provision devices, refer to Pre-provisioning a Device , on page 86.

6. 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 authentic that is, RADIUS or TACACS credentials for discovery only.

7. (Optional) Use discovery credentials for discovering switches.

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

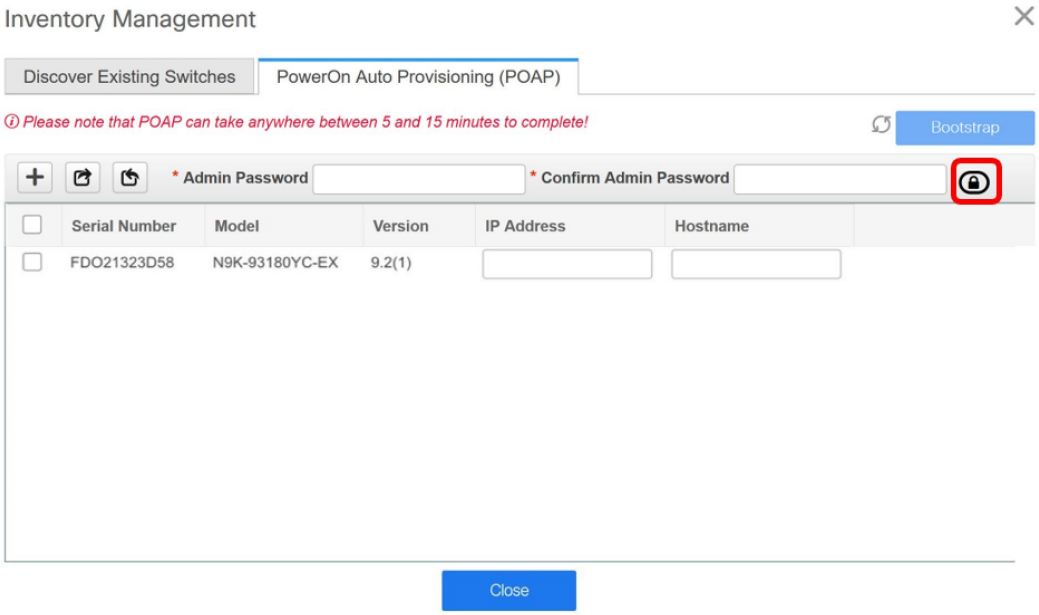

b. In the **Discovery Credentials** window, enter the discovery credentials such as discovery username and password.

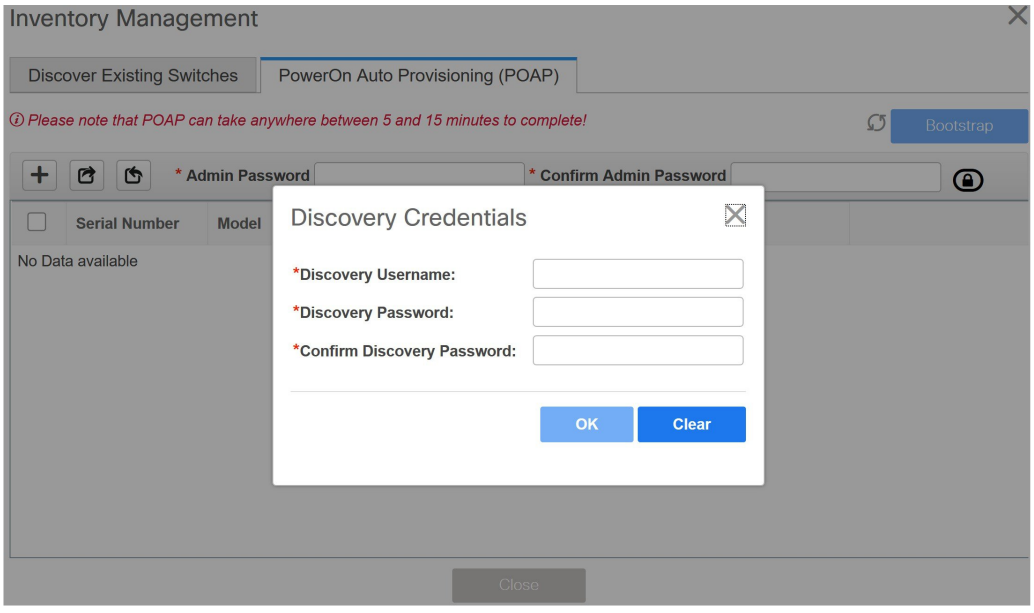

Click **OK** to save the discovery credentials.

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

8. 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.

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

10. After the added switch completes POAP, the fabric builder topology page is refreshed with the added switch thereby depicting its discovered physical connections. Set the appropriate role for the switch followed by a Save & Deploy operation at the fabric level. The Fabric Settings, switch role, the topology etc. are evaluated by the Fabric Builder and the appropriate intended configuration for the switch is generated as part of the Save operation. The pending configuration will provide a list of the configurations that need to be deployed to the new switch in order to bring it IN-SYNC with the intent.

Note For any changes on the fabric that results in the Out-of-Sync, then you must deploy the changes. The pr 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 n update the AAA server password on each switch. Else, switch discovery fails.

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

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

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

14. 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.

15. 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. 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 pairing/un-pairing 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:

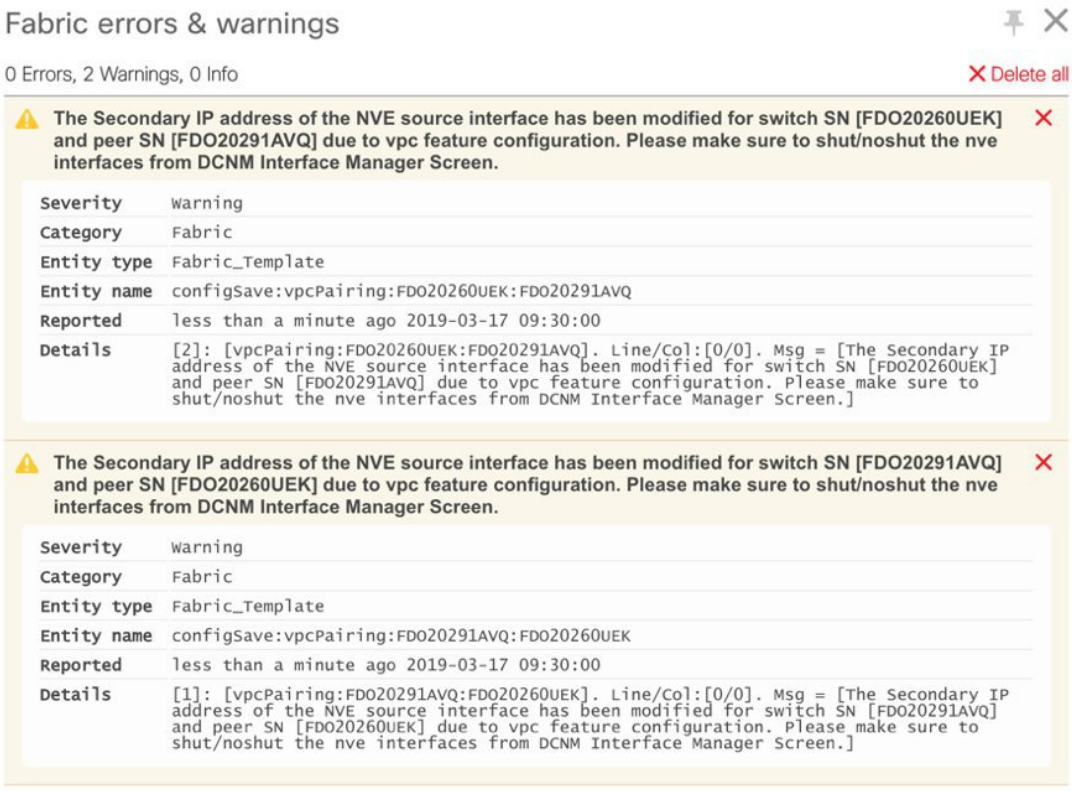

To resolve, go to the Control > Interfaces screen and deploy the Shutdown operation on the nve interface followed by a No Shutdown configuration. This is depicted in the figure below where the up arrow corresponds to a No Shutdown operation while a down arrow corresponds to a Shutdown operation.

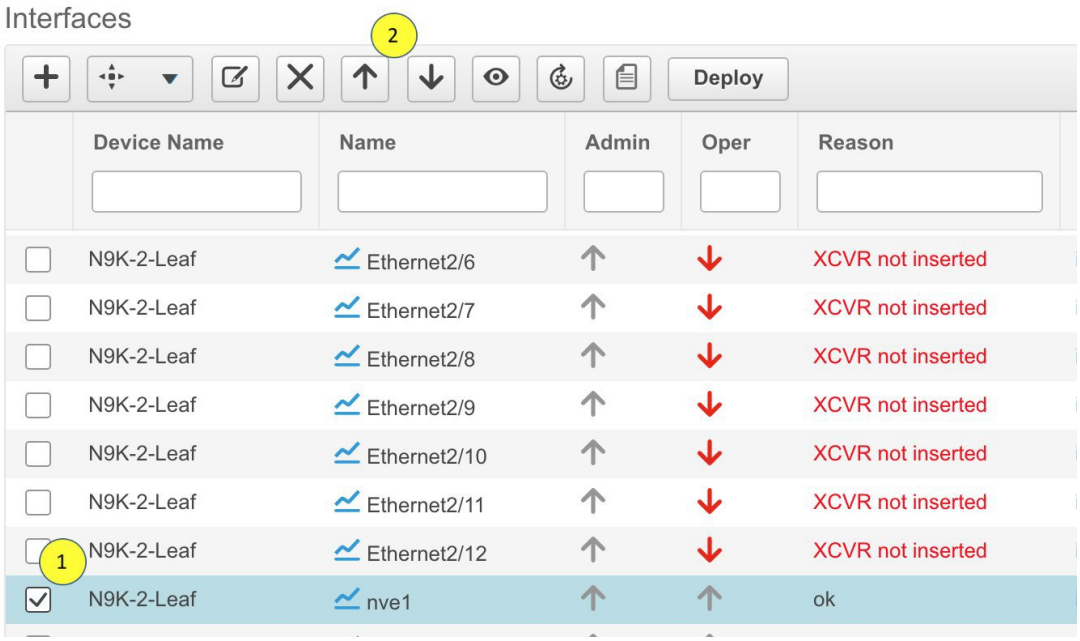

You can right-click the switch to view various options:

- **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 205.

- **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 and policy change history.

  The **Policy Change History** tab lists the history of policies along with the users who made the changes like add, update, or delete.

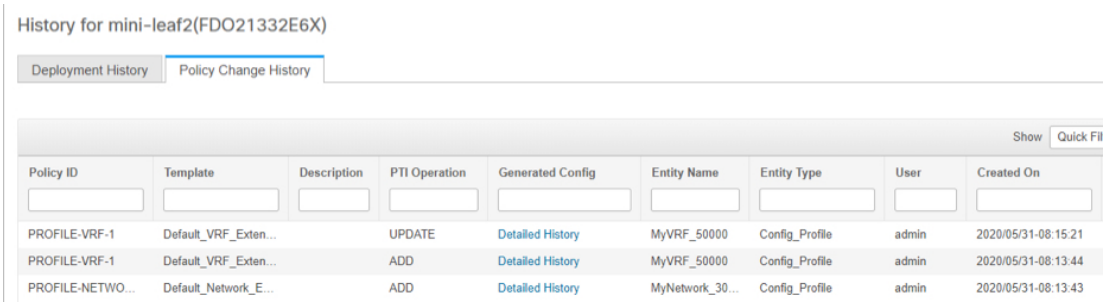

Under the **Policy Change History** tab, for a policy, click **Detailed History** under the **Generated Config** column to view the generated config before and after.

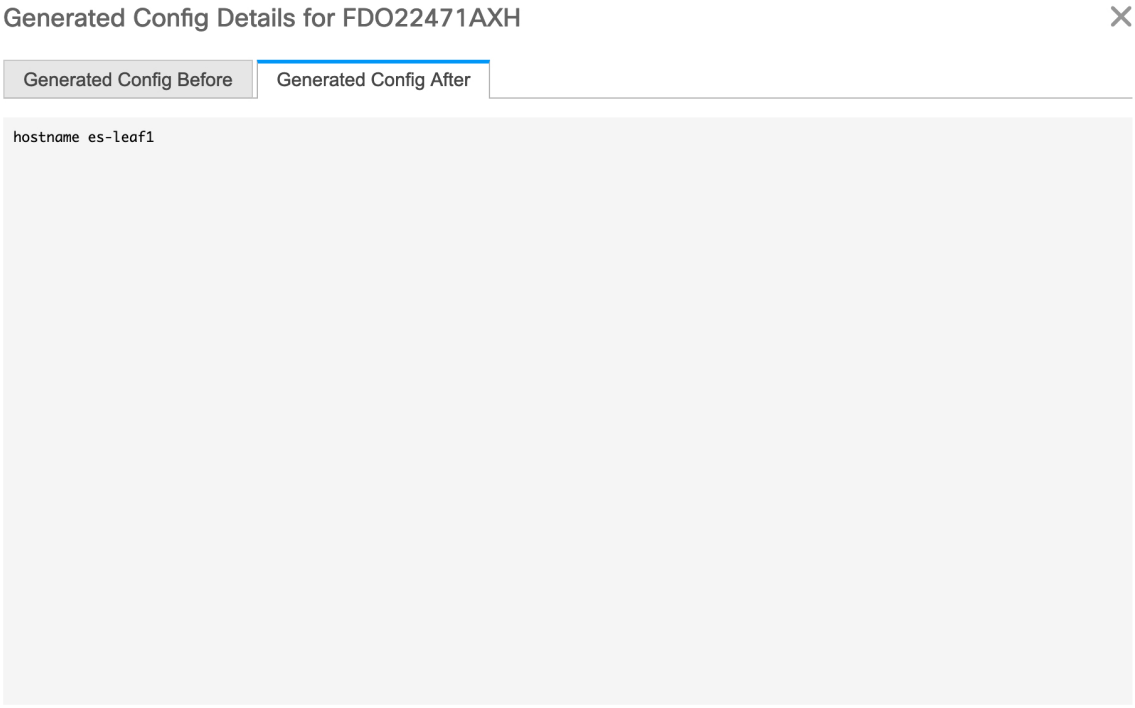

The following table provides the summary of generated config before and after for Policy Template Instances (PTIs).

PTI Operations	Generated Config Before	Generated Config After
Add	Empty	Contains the config
Update	Contains config before changes	Contains config after changes
Mark-Delete	Contains the config to be removed.	Contains the config to be removed with colour change.
Delete	Contains the config	Empty

Note When a policy or profile template is applied, an instance is created for each application of the template, which is known as Policy Template Instance or PTI.

- **Preview Config** - View the pending configuration and the side-by-side comparison of the running and expected configuration.
- **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 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 networks and deploy them on the switches. [Refer Creating and Deploying Networks and VRFs].

Discovering Existing Switches

1. After clicking on **Add Switches**, use the **Discover Existing Switches** tab to add one or more existing switches into the fabric. In this case, a switch with known credentials and a pre-provisioned IP address, is added to the fabric. The IP address (Seed IP), administrator username, and password (**Username** and **Password** fields) of the switch are provided as the input by a user. The **Preserve Config** knob is set to **yes** by default. This is the option that a user would select for a brownfield import of a device into the fabric. For a greenfield import where the device configuration will be cleaned up as part of the import process, the user should set the **Preserve Config** knob to **no**.

Note Easy_Fabric_eBGP does not support brownfield import of a device into the fabric.

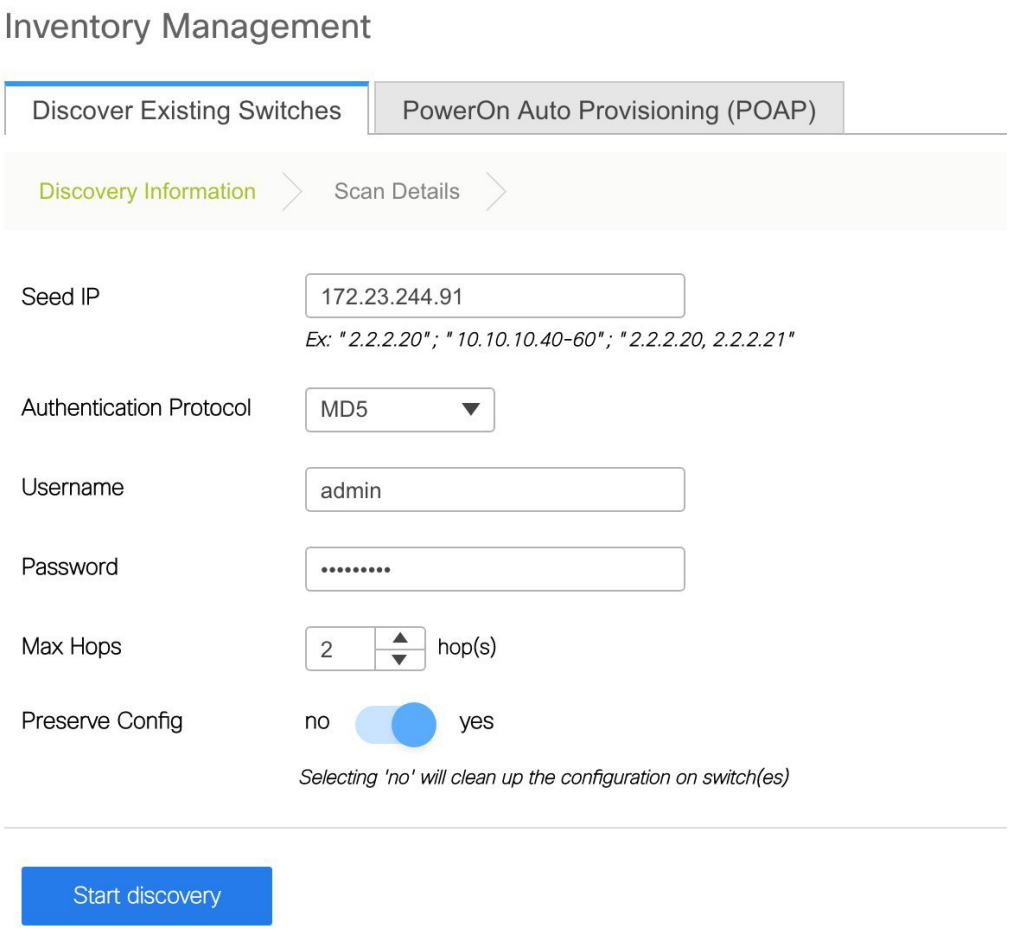

2. Click **Start discovery**. The **Scan Details** window comes up shortly. Since the **Max Hops** field was populated with 2 (by default), the switch with the specified IP address (leaf-91) and switches two hops from that switch, are populated in the **Scan Details**result.

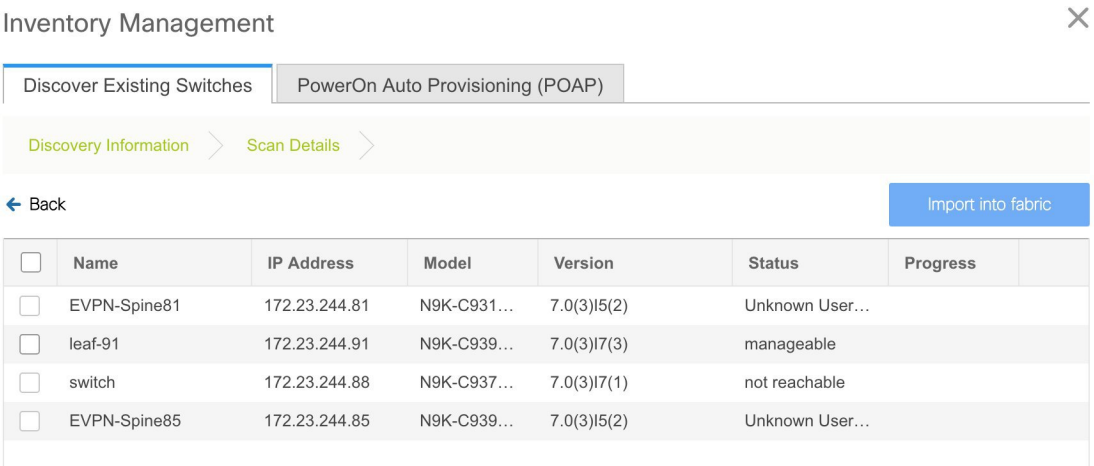

3. If the DCNM was able to perform a successful shallow discovery to a switch, the status will show up as **Manageable**. Select the check box next to the appropriate switch(es) and click **Import into fabric**.

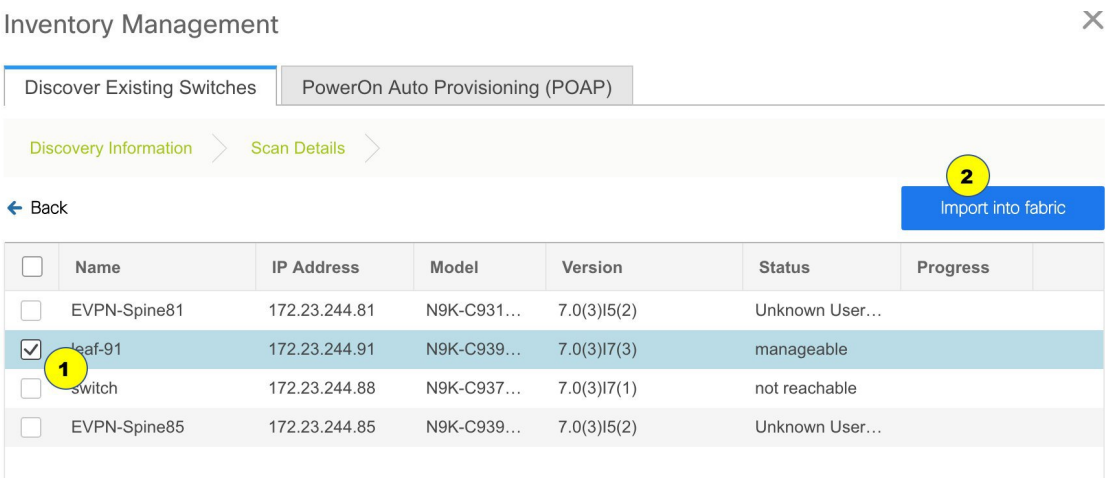

Though this example describes the discovery of one switch, multiple switches can be discovered at once.

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 c error message comes up.

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

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.

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 devices, assign an appropriate role to each device. For this purpose, tight click the device, and use the **Set role** option to set the appropriate role. Alternatively, the tabular view may be employed to assign the same role to multiple devices at one go.

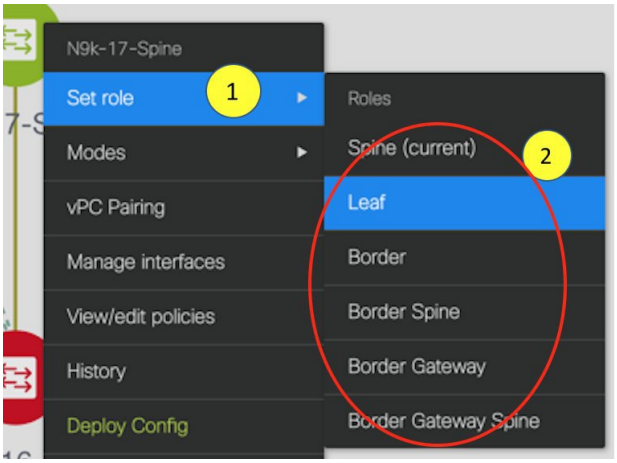

If you choose the Hierarchical layout for display (in the Actions panel), the topology automatically gets aligned as per role assignment, with the leaf devices at the bottom, the spine devices connected on top of them, and the border devices at the top.

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

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

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

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

   The template and interface configurations form the underlay network configuration on the switches. Also, freeform CLIs that were entered as part of fabric settings (leaf and spine switch freeform configurations

entered in the Advanced tab) are deployed. For more details on freeform configurations, refer Enabling Freeform Configurations on Fabric Switches .

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

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

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

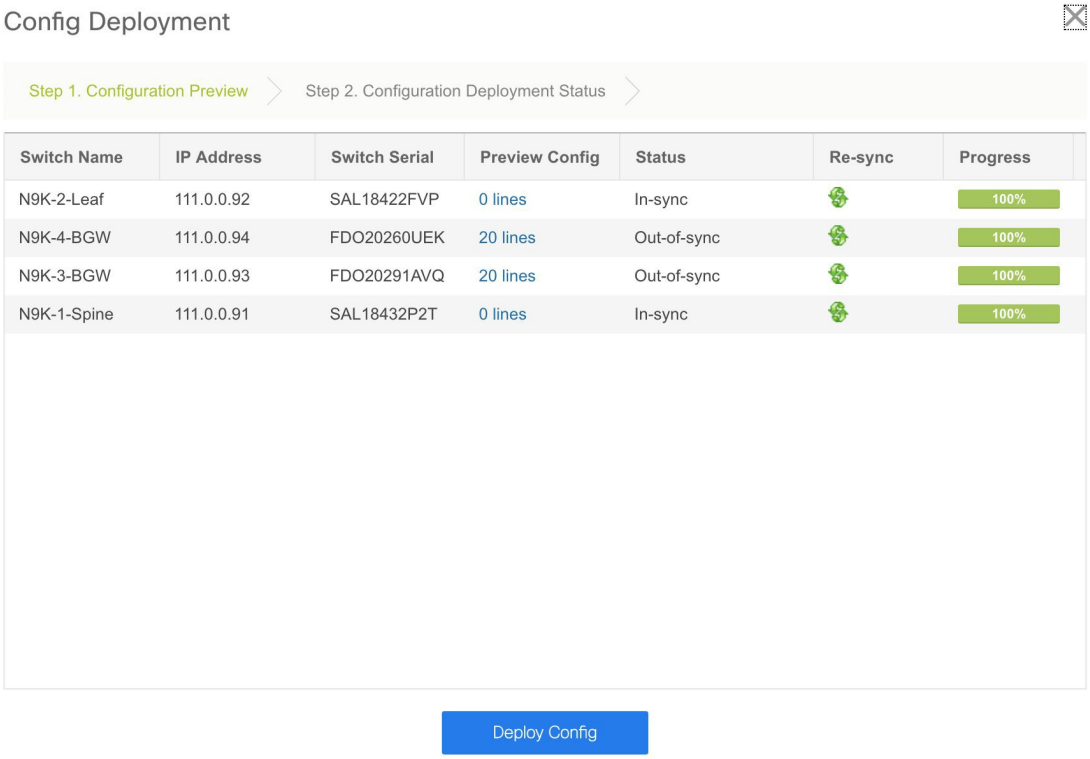

If the status is out-of-sync, it suggests that there is inconsistency between the DCNM and configuration on the device.

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

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

The **PendingConfig** tab displays the pending configurations for successful deployment.

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

In DCNM 11, 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 are 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. The pending state indicates that there is a pending deployment or pending recomputation. You can click on the switch and review the pending deployments using **Preview** or **Deploy Config** options, or click **Save & Deploy** to recompute the state of the switch.

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

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.

From Cisco NX-OS Release 11.4(1), if you uncheck the **FEX** check box in the **Topology** window, FEX devices are hidden in the **Fabric Builder** topology window as well. To view FEX in **Fabric Builder**, you need to check this check box. This option is applicable for all fabrics and it is saved per session or until you log out of DCNM. If you log out and log in to DCNM, the FEX option is reset to default, that is, enabled by default. For more information, see Show Panel, on page 24.

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

VXLAN EVPN Deployment with eBGP EVPN

Creating a eBGP New VXLAN EVPN with 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.

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.

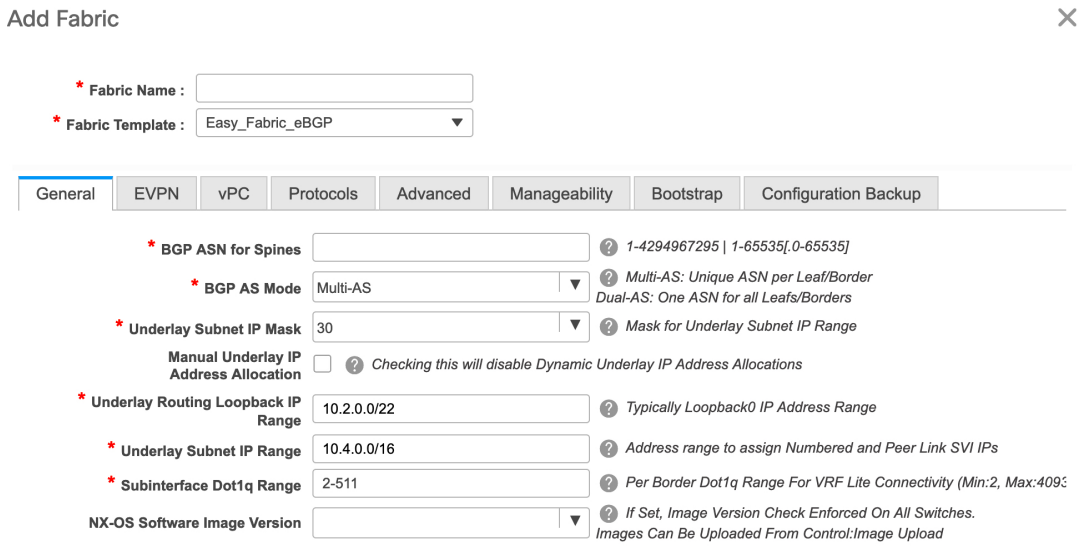

3. The **General** tab is displayed by default. The fields in this tab are:

   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.

   Manual Underlay IP Address Allocation – Select this check box to disable Dynamic Underlay IP Address Allocations.

   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:

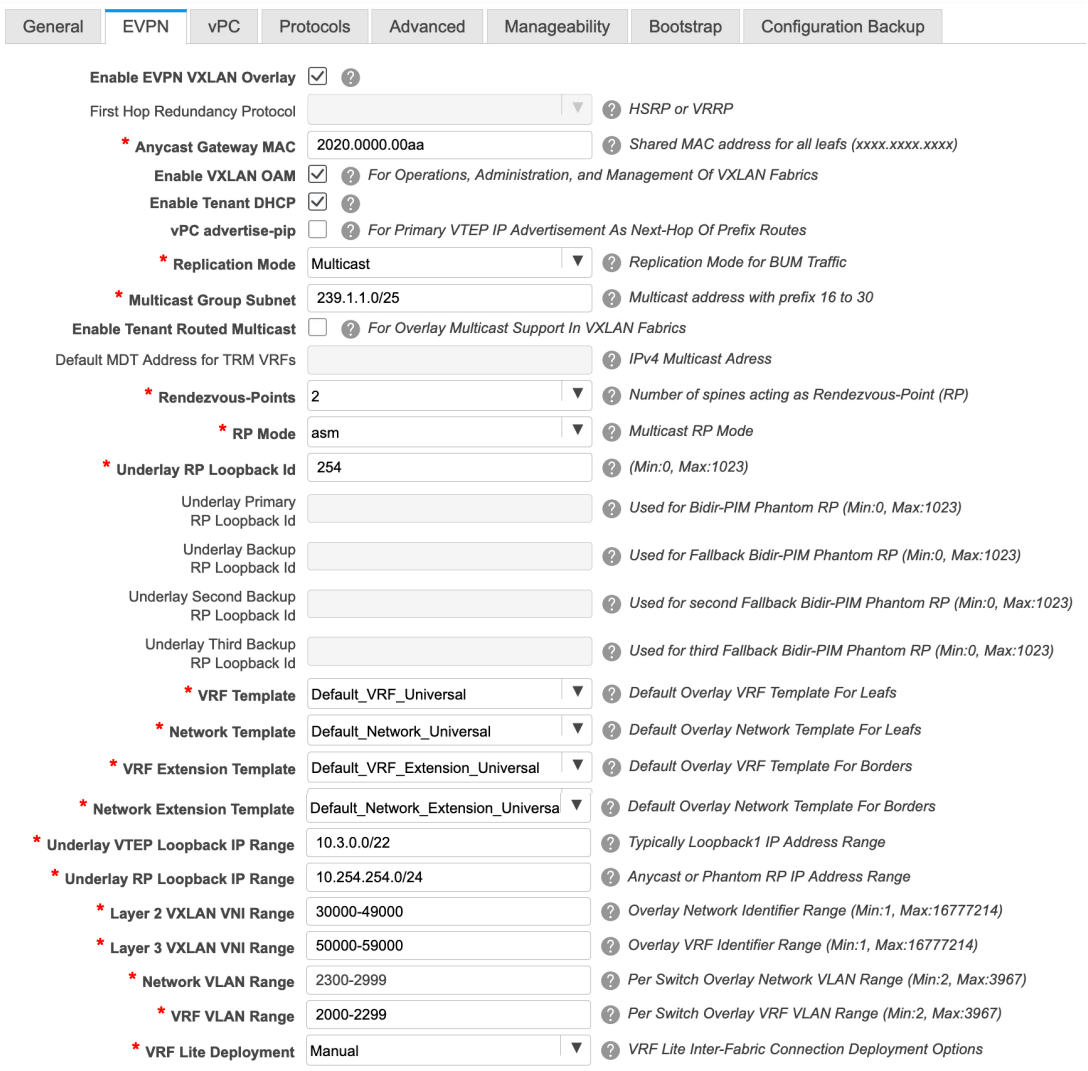

Enable EVPN VXLAN Overlay: Enables the VXLAN overlay provisioning for the fabric.

You can convert a routed fabric to a VXLAN enabled fabirc 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 ore 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. In a Routed Fabric, you can create and deploy networks. For more information, see Overview of Networks in a Routed Fabric, on page 988.

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.

If a network or a VRF is created in a fabric, you cannot switch between VXLAN EVPN mode and Routed Fabric mode by selecting the **Enable EVPN VXLAN Overlay** check box. You need to delete these networks or VRFs to change the fabric setting.

Note that **Routed_Network_Universal Template** is only applicable to a Routed Fabric. When you convert the routed fabric to EVPN VXLAN fabric, set the network template and network extension template to the ones defined for EVPN VXLAN: **Default_Network_Universal** and **Default_Network_Universal**. If you have a customized template for EVPN VXLAN fabric, you can also choose to use it.

First Hop Redundancy Protocol: Specifies the FHRP protocol. Choose either **hsrp** or **vrrp**. This field is only applicable to a Routed Fabric.

Note

- After a network has been created, you cannot change this fabric setting. You should delete all netw and then change the FHRP setting.
- The rest of the fields in the EVPN tab section are only applicable if you enable the EVPN VXLA Overlay.

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

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 free-form configurations to enable OAM an 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.

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 softwa release 9.2(1) onwards.

Underlay RP Loopback ID: The loopback ID used for the rendezvous point (RP), for multicast protocol peering purposes in the fabric underlay. The 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**: Specify the VRF template for creating VRFs, and the VRF extension template for enabling VRF extension to other fabrics.

Network Template and **Network Extension Template**: Specify 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**: Specify the VXLAN VNI IDs for the fabric.

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

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

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

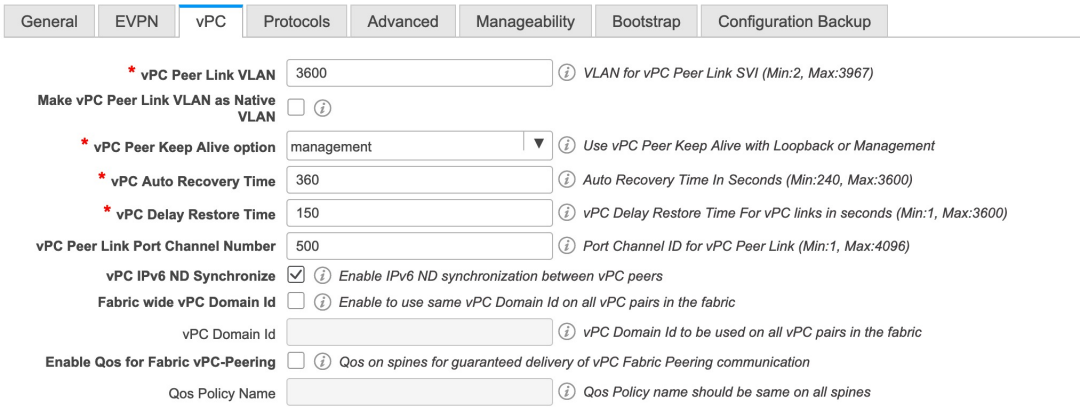

vPC Peer Link VLAN: VLAN used for the vPC peer link SVI.

Make vPC Peer Link VLAN as Native VLAN - Enables vPC peer link VLAN as Native VLAN.

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.

Fabric wide vPC Domain Id: Enables the usage of same vPC Domain Id on all vPC pairs in the fabric. When you select this field, the **vPC Domain Id** field is editable.

vPC Domain Id - Specifies the vPC domain ID to be used on all vPC pairs.

Enable Qos for Fabric vPC-Peering - Enable QoS on spines for guaranteed delivery of vPC Fabric Peering communication.

Qos Policy Name - Specifies QoS policy name that should be same on all spines.

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

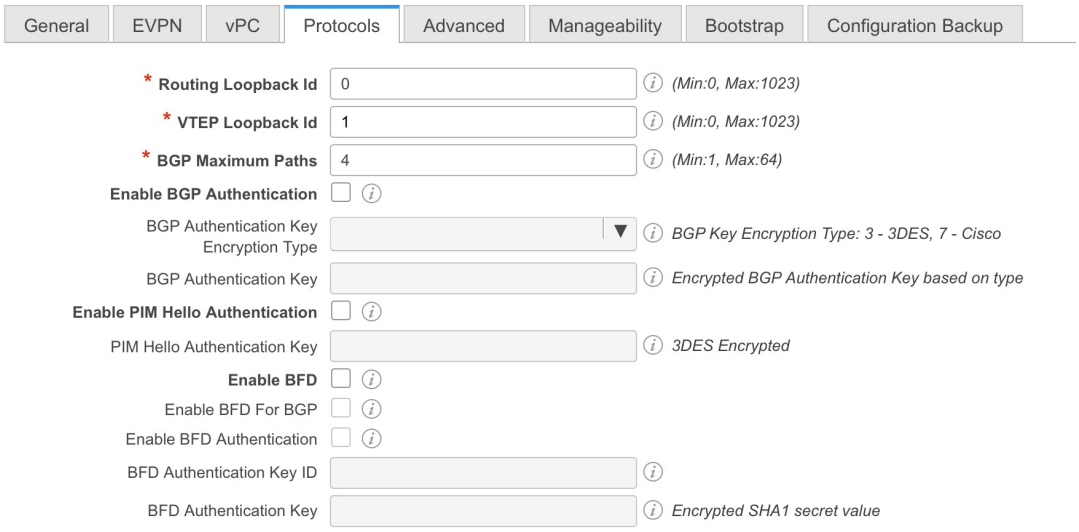

Routing Loopback Id - The loopback interface ID is populated as 0 by default. It is used as the BGP router ID.

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

BGP Maximum Paths - Specifies the BGP maximum paths.

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 BG Authentication Key field. Refer the Retrieving the Authentication Key section for details.

Enable PIM Hello Authentication: Enables the PIM hello authentication.

PIM Hello Authentication Key: Specifies the PIM hello authentication key.

Enable BFD: Select the check box to enable **feature bfd** on all switches in the fabric. This feature is valid only on IPv4 underlay and the scope is within a fabric.

From Cisco DCNM Release 11.3(1), BFD within a fabric is supported natively. The BFD feature is disabled by default in the Fabric Settings. If enabled, BFD is enabled for the underlay protocols with the default settings. Any custom required BFD configurations must be deployed via the per switch freeform or per interface freeform policies.

The following config is pushed after you select the **Enable BFD** check box:

```
feature bfd
```

Note After you upgrade from DCNM Release 11.2(1) with BFD enabled to DCNM Release 11.3(1), the follo configs are pushed on all P2P fabric interfaces:

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

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 BGP: Select the check box to enable BFD for the BGP neighbor. This option is disabled by default.

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.

BFD Authentication Key ID: Specifies the BFD authentication key ID for the interface authentication.

BFD Authentication Key: Specifies the BFD authentication key.

For information about how to retrieve the BFD authentication parameters, see *Retrieving the Encrypted BFD Authentication Key, in Cisco DCNM LAN Fabric Configuration Guide*.

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

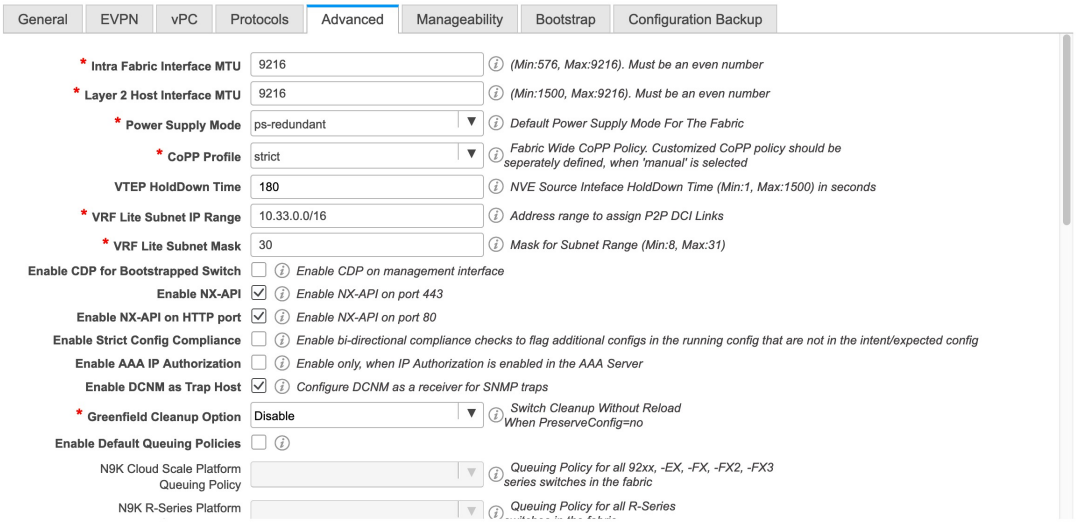

Intra Fabric Interface MTU - Specifies the MTU for the intra fabric interface. This value should be an even number.

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

VTEP HoldDown Time - Specifies the NVE source interface hold down time.

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

Enable CDP for Bootstrapped Switch - Select the check box to enable CDP for bootstrapped switch.

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

Enable NX-API on HTTP - Specifies enabling of NX-API on HTTP. Enable this check box and the **Enable NX-API** check box to use HTTP. This check box is checked by default. If you uncheck this check box, the applications that use NX-API and supported by Cisco DCNM, such as Endpoint Locator (EPL), Layer 4-Layer 7 services (L4-L7 services), VXLAN OAM, and so on, start using the HTTPS instead of HTTP.

Note If you check the **Enable NX-API** check box and the **Enable NX-API on HTTP** check box, applications u HTTP.

Enable Strict Config Compliance - Enable the Strict Config Compliance feature by selecting this check box.

For Strict Configuration Compliance, see *Enhanced Monitoring and Monitoring Fabrics Guide*.

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

Enable DCNM as Trap Host - Select this check box to enable DCNM as a trap host.

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 Default Queuing Policies: Check this check box to apply QoS policies on all the switches in this fabric. To remove the QoS policies that you applied on all the switches, uncheck this check box, update all the configurations to remove the references to the policies, and save and deploy. From Cisco DCNM Release 11.3(1), pre-defined QoS configurations are included that can be used for various Cisco Nexus 9000 Series Switches. When you check this check box, the appropriate QoS configurations are pushed to the switches in the fabric. The system queuing is updated when configurations are deployed to the switches. You can perform the interface marking with defined queuing policies, if required, by adding the required configuration to the per interface freeform block.

Review the actual queuing policies by opening the policy file in the template editor. From Cisco DCNM Web UI, choose **Control > Template Library**. Search for the queuing policies by the policy file name, for example, **queuing_policy_default_8q_cloudscale**. Choose the file and click the **Modify/View template** icon to edit the policy.

See the *Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide* for platform specific details.

N9K Cloud Scale Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all Cisco Nexus 9200 Series Switches and the Cisco Nexus 9000 Series Switches that ends with EX, FX, and FX2 in the fabric. The valid values are **queuing_policy_default_4q_cloudscale** and

queuing_policy_default_8q_cloudscale. Use the **queuing_policy_default_4q_cloudscale** policy for FEXes. You can change from the **queuing_policy_default_4q_cloudscale** policy to the **queuing_policy_default_8q_cloudscale** policy only when FEXes are offline.

N9K R-Series Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all Cisco Nexus switches that ends with R in the fabric. The valid value is **queuing_policy_default_r_series**.

Other N9K Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all other switches in the fabric other than the switches mentioned in the above two options. The valid value is **queuing_policy_default_other**.

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.

Intra-fabric Links Additional Config - Add CLIs that should be added to the intra-fabric links.

8. Click the **Manageability** tab.

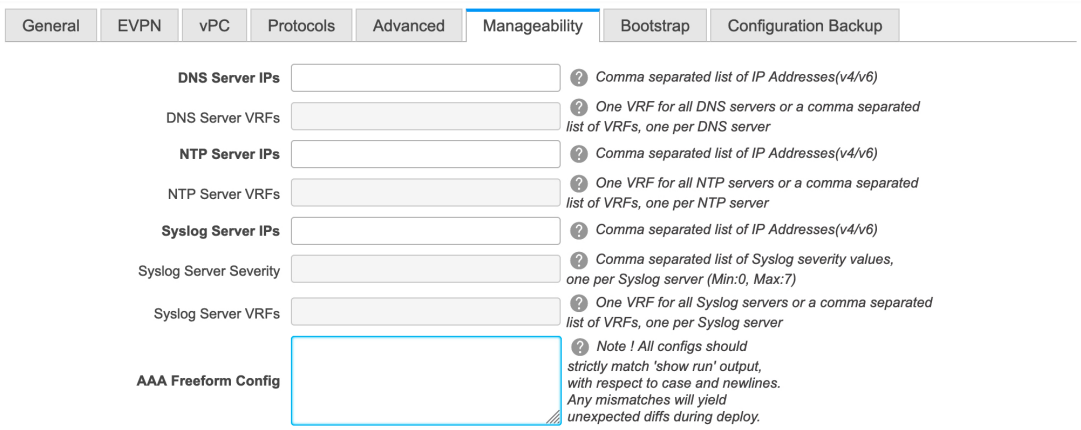

The fields in this tab are:

DNS Server IPs - Specifies the comma separated list of IP addresses (v4/v6) of the DNS servers.

DNS Server VRFs - Specifies one VRF for all DNS servers or a comma separated list of VRFs, one per DNS server.

NTP Server IPs - Specifies comma separated list of IP addresses (v4/v6) of the NTP server.

NTP Server VRFs - Specifies one VRF for all NTP servers or a comma separated list of VRFs, one per NTP server.

Syslog Server IPs – Specifies the comma separated list of IP addresses (v4/v6) IP address of the syslog servers, if used.

Syslog Server Severity – Specifies the comma separated list of syslog severity values, one per syslog server. The minimum value is 0 and the maximum value is 7. To specify a higher severity, enter a higher number.

Syslog Server VRFs – Specifies one VRF for all syslog servers or a comma separated list of VRFs, one per syslog server.

AAA Freeform Config – Specifies the AAA freeform configs.

If AAA configs are specified in the fabric settings, **switch_freeform** PTI with source as **UNDERLAY_AAA** and description as "**AAA Configurations**" will be created.

9. Click the **Bootstrap** tab.

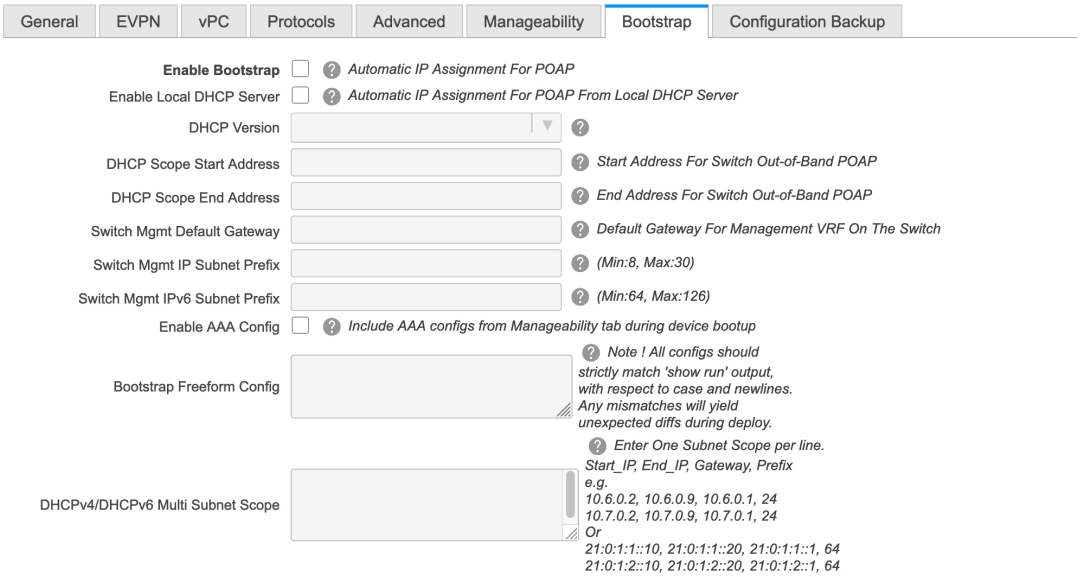

Enable Bootstrap - Select this check box to enable the bootstrap feature.

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

- External DHCP Server: Enter information about the external DHCP server in the **Switch Mgmt Default Gateway** and **Switch Mgmt IP 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 Version – Select DHCPv4 or DHCPv6 from this drop-down list. When you select DHCPv4, the **Switch Mgmt IPv6 Subnet Prefix** field is disabled. If you select DHCPv6, the **Switch Mgmt IP Subnet Prefix** is disabled.

Note Cisco DCNM IPv6 POAP is not supported with Cisco Nexus 7000 Series Switches. Cisco Nexus 9000 an 3000 Series Switches support IPv6 POAP only when switches are either L2 adjacent (eth1 or out-of-band subnet must be a /64) or they are L3 adjacent residing in some IPv6 /64 subnet. Subnet prefixes other than /64 are not supported.

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 Mgmt Default Gateway: Specifies the default gateway for the management VRF on the switch.

Switch Mgmt IP 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.

Switch Mgmt IPv6 Subnet Prefix - Specifies the IPv6 prefix for the Mgmt0 interface on the switch. The prefix should be between 112 and 126. This field is editable if you enable IPv6 for DHCP.

Enable AAA Config – Select this check box to include AAA configs from the Manageability tab during device bootup.

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*.

DHCPv4/DHCPv6 Multi Subnet Scope - Specifies the field to enter one subnet scope per line. This field is editable after you check the **Enable Local DHCP Server** check box.

The format of the scope should be defined as:

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

For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24

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

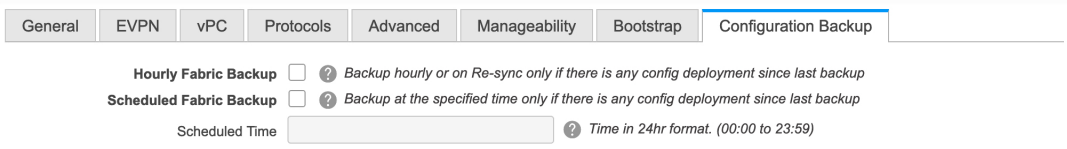

Hourly Fabric Backup: Select the check box to enable an hourly backup of fabric configurations and the intent.

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.

The backup process is initiated after you click **Save**.

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

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

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

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

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

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

VXLAN Fabric With eBGP Underlay – Pointers

- Deploy the leaf overlay and underlay policies on all leaf switches at once, since they have a common AS number.
- Brownfield migration is not supported for eBGP fabrics.
- 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 swtch leaf_bsp_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*.

Deploying Fabric Underlay eBGP Policies

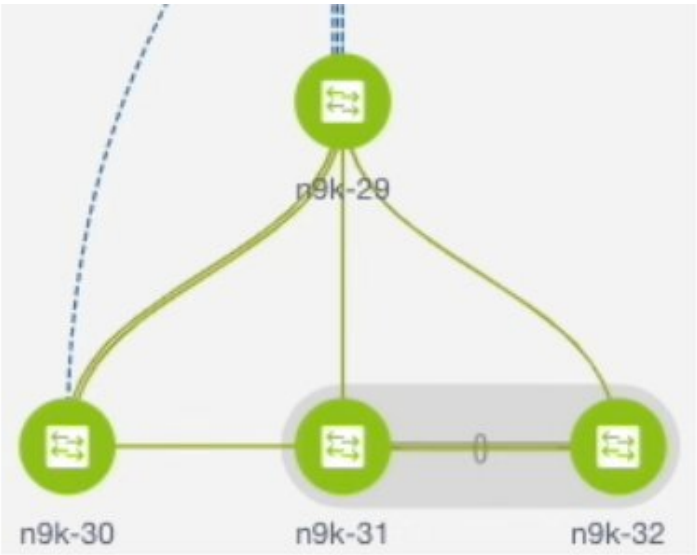

The topology shows a VXLAN fabric enabled with eBGP for the underlay. In DCNM, a fabric with the **Easy_Fabric_eBGP** template is created. One spine switch (n9k-29) and three leaf switches (n9k-30, and vPC switch pair n9k-31 and n9k-32) are imported to it.

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.

DCNM Release 11.5(3a) introduces **adv_static_prefix_bgp** template that allows you to create a static route for the specified IPv4 prefix and advertise it to eBGP peers.

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 mod 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), a you have assigned a BGP AS number to all of them, as explained in the earlier steps.

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

Deploying Fabric Overlay eBGP Policies

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

Deploying Spine Switch Overlay Policies

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

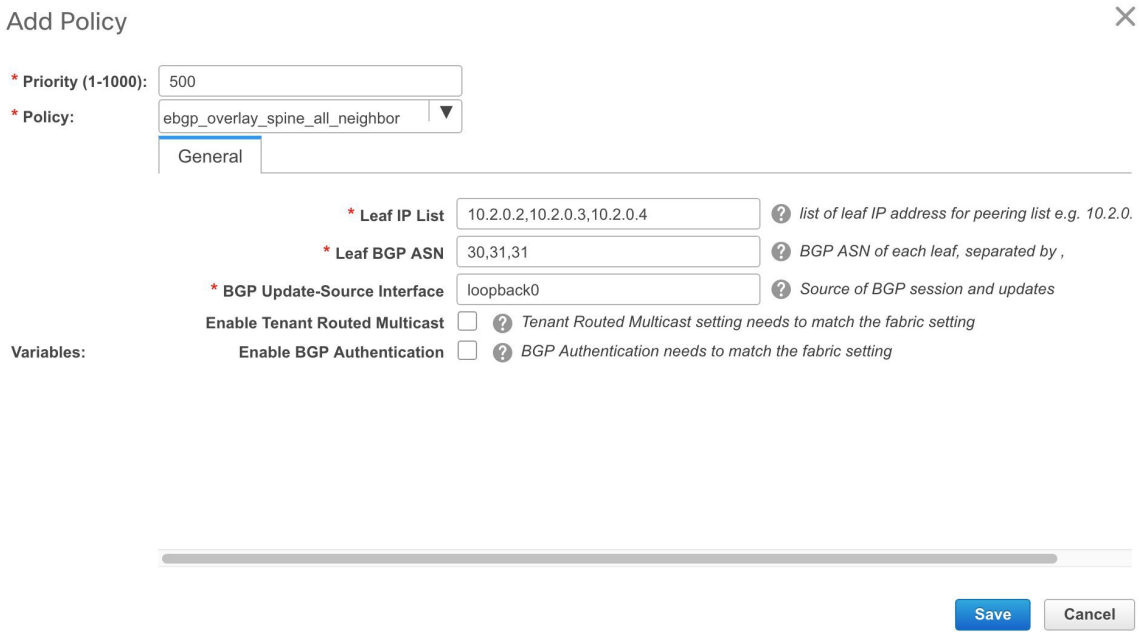

The fields on the screen are:

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

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

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

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

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

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

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

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

Deploying Leaf Switch Overlay Policies

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

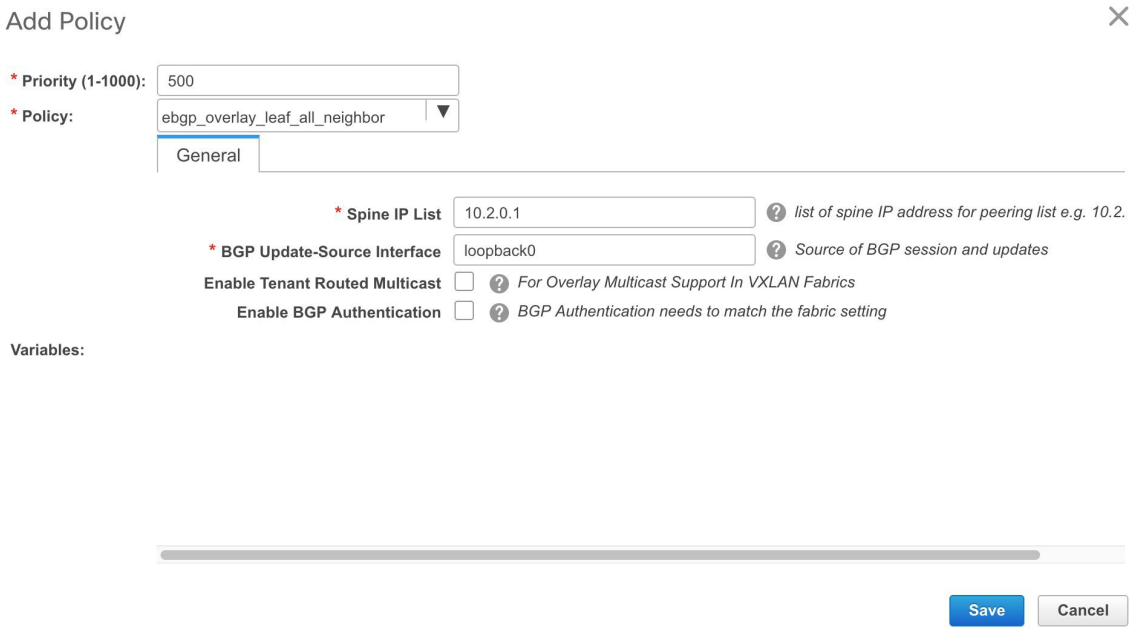

The fields on the screen are:

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

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

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

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

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

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

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

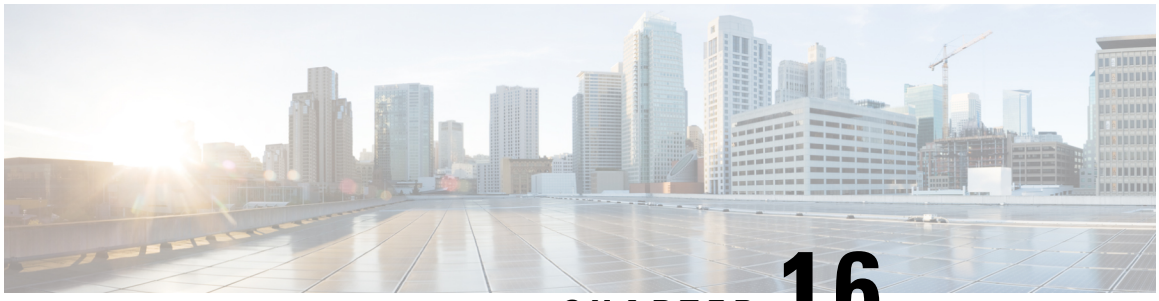

CHAPTER 16

Managing a Brownfield VXLAN BGP EVPN Fabric

This chapter explains how to migrate a Brownfield fabric into Cisco DCNM.

- Overview, on page 695
- Prerequisites, on page 696
- Guidelines and Limitations, on page 696
- Fabric Topology Overview, on page 698
- DCNM Brownfield Deployment Tasks, on page 699
- Verifying the Existing VXLAN BGP EVPN Fabric, on page 699
- Creating a VXLAN BGP EVPN Fabric, on page 702
- Adding Switches and Transitioning VXLAN Fabric Management to DCNM, on page 716
- Verifying the Import of the VXLAN BGP EVPN Fabric, on page 729
- Configuration Profiles Support for Brownfield Migration, on page 737
- Migrating a Bottom-Up VXLAN Fabric to DCNM, on page 737
- 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 746
- 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 750
- Changing a Brownfield Imported BIDIR Configuration, on page 752
- Manually Adding PIM-BIDIR Configuration for Leaf or Spine Post Brownfield Migration , on page 753
- Migrating an MSD Fabric with Border Gateway Switches , on page 753

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.
- 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.

- From Cisco DCNM Release 11.5(1), the brownfield import in DCNM supports the simplified NX-OS VXLAN EVPN configuration CLIs. For more information, see Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide, Release 9.3(x).
- The following features are unsupported.
  - eBGP underlay
  - Layer 3 port channel
- 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.

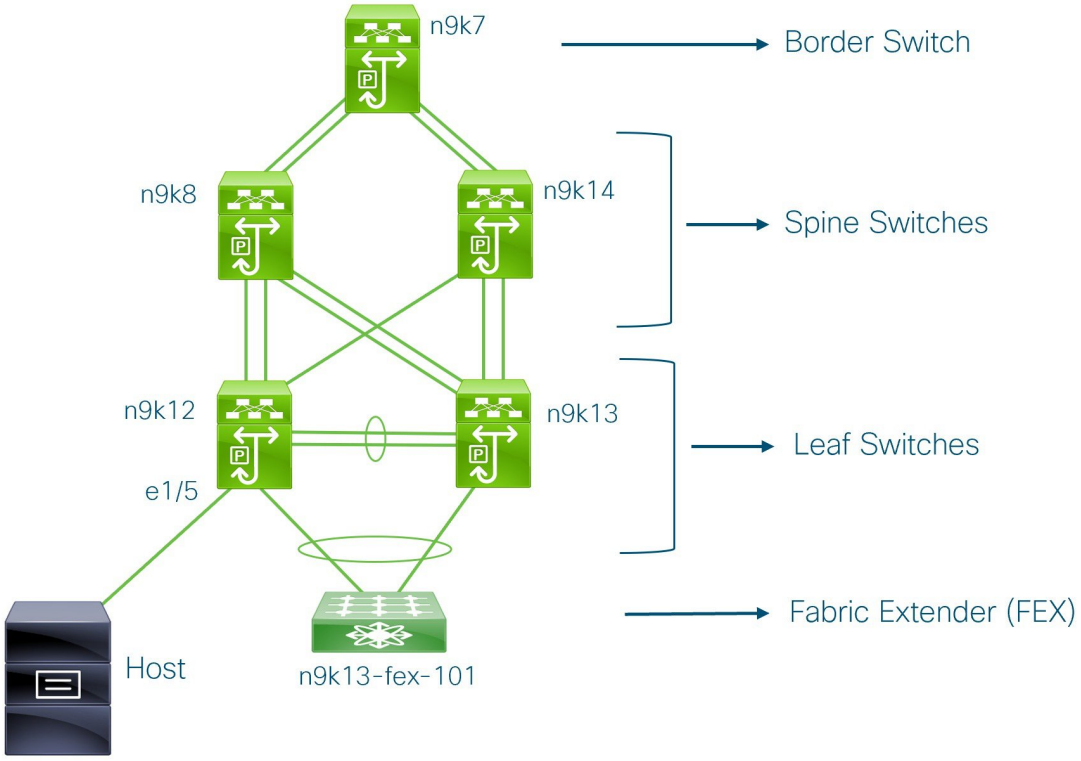

You can see that there is a border switch, two spine switches, two leaf switches, and a Fabric Extender or FEX.

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 699
2. Creating a VXLAN BGP EVPN Fabric, on page 702
3. Adding Switches and Transitioning VXLAN Fabric Management to DCNM, on page 716
4. Verifying the Import of the VXLAN BGP EVPN Fabric, on page 729

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 l2vpn 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:

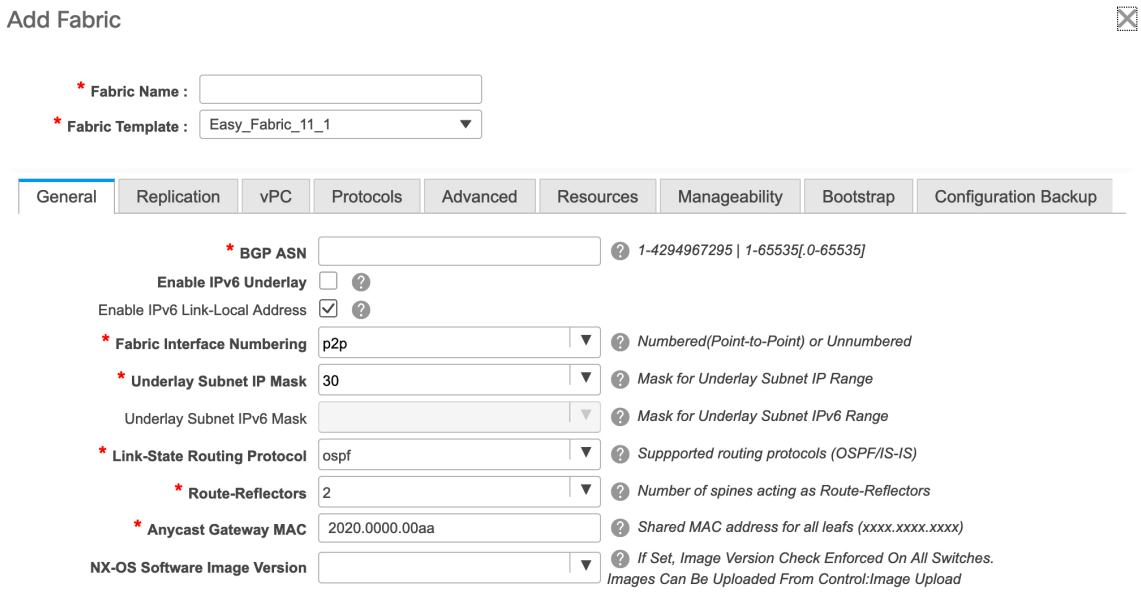

BGP ASN: Enter the BGP AS number the fabric is associated with.

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.

For information about IPv6 underlay, see *Configuring a VXLANv6 Fabric*.

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.

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.

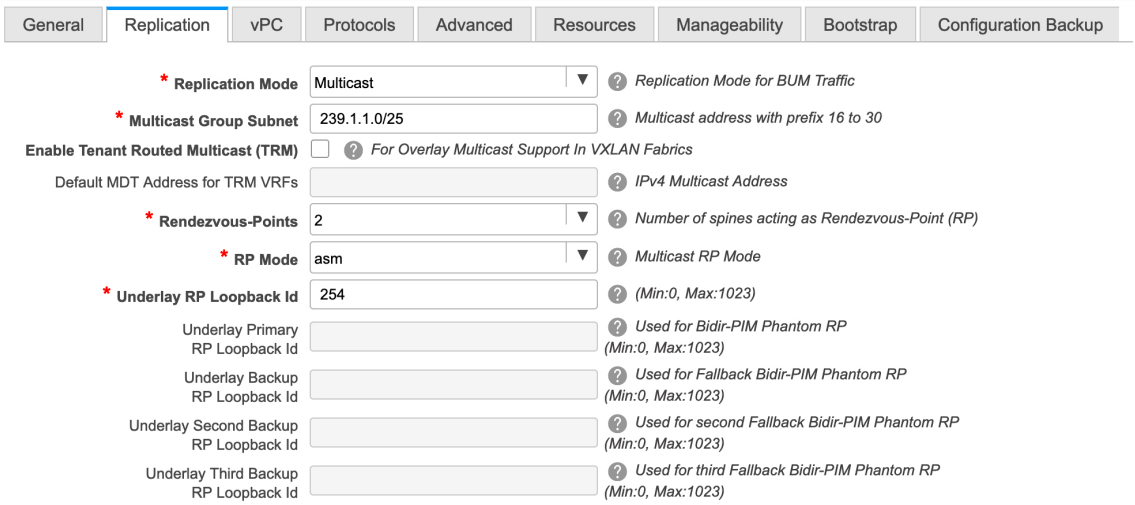

Replication Mode: The mode of replication that is used in the existing fabric, Ingress Replication, or Multicast.

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.

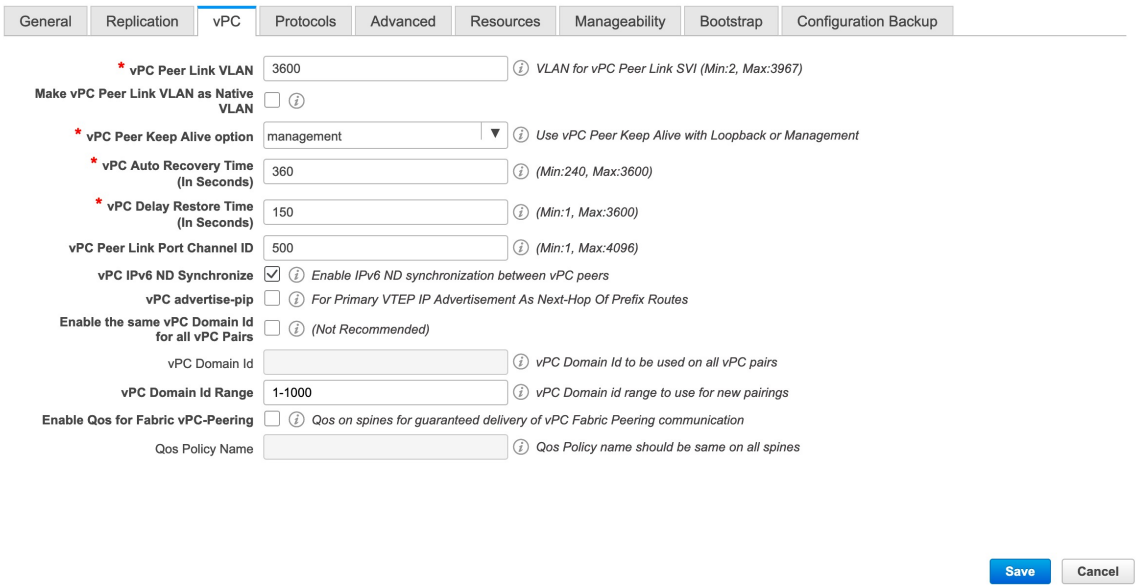

vPC Peer Link VLAN - Enter the VLAN ID used for the vPC peer link SVI in the existing fabric.

Make vPC Peer Link VLAN as Native VLAN - Enables vPC peer link VLAN as Native VLAN.

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 ID - 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.

Enable the same vPC Domain Id for all vPC Pairs: Enable the same vPC Domain ID for all vPC pairs. When you select this field, the **vPC Domain Id** field is editable.

vPC Domain Id - Specifies the vPC domain ID to be used on all vPC pairs.

vPC Domain Id Range - Specifies the vPC Domain Id range to use for new pairings.

Enable Qos for Fabric vPC-Peering - Enable QoS on spines for guaranteed delivery of vPC Fabric Peering communication. For more information, see QoS for Fabric vPC-Peering, on page 233.

Note QoS for vPC fabric peering and queuing policies options in fabric settings are mutually exclusive.

Qos Policy Name - Specifies QoS policy name that should be same on all fabric vPC peering spines. The default name is **spine_qos_for_fabric_vpc_peering**.

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 PIM Hello Authentication - Enables the PIM hello authentication.

PIM Hello Authentication Key - Specifies the PIM hello authentication key.

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.

From Cisco DCNM Release 11.3(1), BFD within a fabric is supported natively. The BFD feature is disabled by default in the Fabric Settings. If enabled, BFD is enabled for the underlay protocols with the default settings. Any custom required BFD configurations must be deployed via the per switch freeform or per interface freeform policies.

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 be 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 is valid for only for P2P interfaces.
- After you upgrade from DCNM Release 11.2(1) with BFD enabled to DCNM Release 11.3(1), the following configs are pushed to the switch:

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

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.

Until Cisco DCNM Release 11.3(1), iBGP peer template for iBGP definition on the leafs or border role devices and BGP RRs were same. From DCNM Release 11.4(1), the following fields can be used to specify different configurations:

- **iBGP Peer-Template Config** – Specifies the config used for RR and spines with border role.
- **Leaf/Border/Border Gateway iBGP Peer-Template Config** – Specifies the config used for leaf, border, or border gateway. If this field is empty, the peer template defined in **iBGP Peer-Template Config** is used on all BGP enabled devices (RRs, leafs, border, or border gateway roles).

In brownfield migration, if the spine and leaf use different peer template names, both the **iBGP Peer-Template Config** and **Leaf/Border/Border Gateway iBGP Peer-Template Config** fields need to be set according to the switch config. If spine and leaf use the same peer template name and content (except for the

"route-reflector-client" CLI), only **iBGP Peer-Template Config** field in fabric setting needs to be set. If the fabric settings on iBGP peer templates do not match the existing switch configuration, an error message is generated and the migration will not proceed.

Step 7 Click the **Advanced** tab. Most of the fields are auto generated.

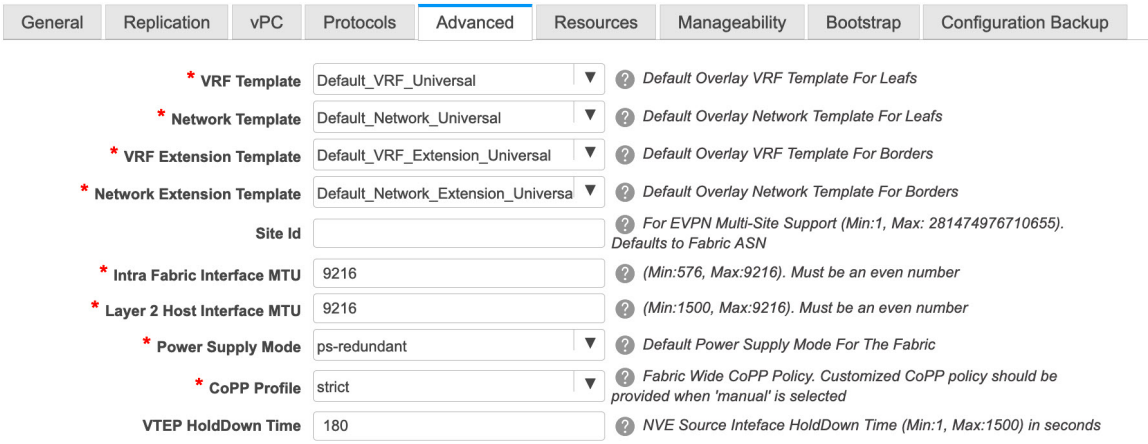

VRF Template and **VRF Extension Template**: Specifies the VRF template for creating VRFs, and the VRF extension template for enabling VRF extension to other fabrics.

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.

Brownfield Overlay Network Name Format: Enter the format to be used to build the overlay network name during a brownfield import or migration. The network name should not contain any white spaces or special characters except underscore (_) and hyphen (-). The network name must not be changed once the brownfield migration has been initiated. See the *Creating Networks for the Standalone Fabric* section for the naming convention of the network name. The syntax is **[<string> | $$VLAN_ID$$] $$VNI$$ [<string>| $$VLAN_ID$$]** and the default value is **Auto_Net_VNI$$VNI$$_VLAN$$VLAN_ID$$**. When you create networks, the name is generated according to the syntax you specify. The following table describes the variables in the syntax.

Variables	Description
$$VNI$$	Specifies the network VNI ID found in the switch configuration. This is a mandatory keyword required to create unique network names.
$$VLAN_ID$$	Specifies the VLAN ID associated with the network. VLAN ID is specific to switches, hence DCNM will pick the VLAN ID from one of the switches, where the network is found, randomly and use it in the name. We recommend not to use this unless the VLAN ID is consistent across the fabric for the VNI.
<string>	This variable is optional and you can enter any number of alphanumeric characters that meet the network name guidelines.

Example overlay network name: Site_VNI12345_VLAN1234

Note Ignore this field for greenfield deployments. The **Brownfield Overlay Network Name Format** applies for the following brownfield imports:

- CLI-based overlays
- Configuration profile-based overlay where the configuration profiles were created in Cisco DCNM Release

  10.4(2).

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

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

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

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

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 Precision Time Protocol (PTP): Enables PTP across a fabric. When you select this check box, PTP is enabled globally and on core-facing interfaces. Additionally, the **PTP Source Loopback Id** and **PTP Domain Id** fields are editable. For more information, see *Precision Time Protocol for Easy Fabric in Cisco DCNM LAN Fabric Configuration Guide*.

PTP Source Loopback Id: Specifies the loopback interface ID Loopback that is used as the Source IP Address for all PTP packets. The valid values range from 0 to 1023. The PTP loopback ID cannot be the same as RP, Phantom RP, NVE, or MPLS loopback ID. Otherwise, an error will be generated. The PTP loopback ID can be the same as BGP loopback or user-defined loopback which is created from DCNM.

If the PTP loopback ID is not found during **Save & Deploy**, the following error is generated:

Loopback interface to use for PTP source IP is not found. Please create PTP loopback interface on all the devices to enable PTP feature.

PTP Domain Id: Specifies the PTP domain ID on a single network. The valid values range from 0 to 127.

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.

Enable TCAM Allocation: TCAM commands are automatically generated for VXLAN and vPC Fabric Peering when enabled.

Enable Default Queuing Policies: Check this check box to apply QoS policies on all the switches in this fabric. To remove the QoS policies that you applied on all the switches, uncheck this check box, update all the configurations to remove the references to the policies, and save and deploy. From Cisco DCNM Release 11.3(1), pre-defined QoS configurations are included that can be used for various Cisco Nexus 9000 Series Switches. When you check this check box, the appropriate QoS configurations are pushed to the switches in the fabric. The system queuing is updated when configurations are deployed to the switches. You can perform the interface marking with defined queuing policies, if required, by adding the required configuration to the per interface freeform block.

Review the actual queuing policies by opening the policy file in the template editor. From Cisco DCNM Web UI, choose **Control > Template Library**. Search for the queuing policies by the policy file name, for example, **queuing_policy_default_8q_cloudscale**. Choose the file and click the **Modify/View template** icon to edit the policy.

See the *Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide* for platform specific details.

N9K Cloud Scale Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all Cisco Nexus 9200 Series Switches and the Cisco Nexus 9000 Series Switches that ends with EX, FX, and FX2 in the fabric. The valid values are **queuing_policy_default_4q_cloudscale** and **queuing_policy_default_8q_cloudscale**. Use the **queuing_policy_default_4q_cloudscale** policy for FEXes.

You can change from the **queuing_policy_default_4q_cloudscale** policy to the **queuing_policy_default_8q_cloudscale** policy only when FEXes are offline.

N9K R-Series Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all Cisco Nexus switches that ends with R in the fabric. The valid value is **queuing_policy_default_r_series**.

Other N9K Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all other switches in the fabric other than the switches mentioned in the above two options. The valid value is **queuing_policy_default_other**.

Enable MACsec - Enables MACsec for the fabric. For more information, see MACsec Support in Easy Fabric and eBGP Fabric, on page 194.

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.

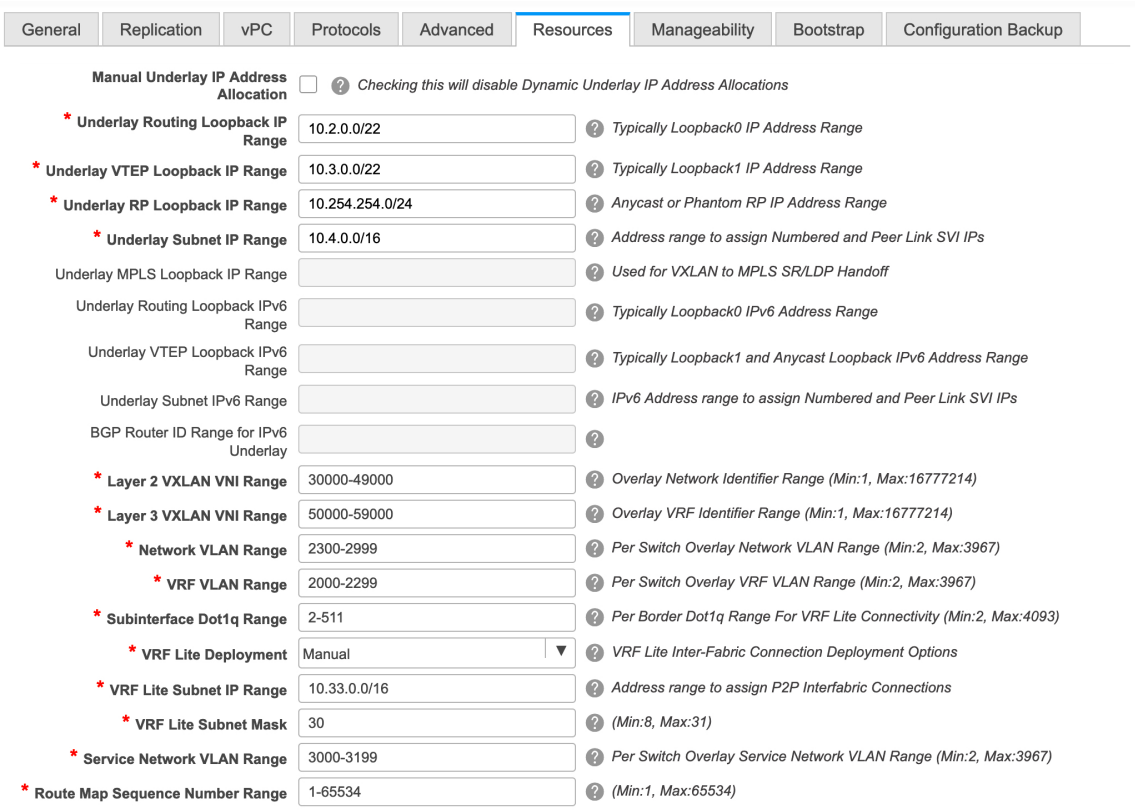

Manual Underlay IP Address Allocation – *Do not* select this check box if you are transitioning your VXLAN fabric management to DCNM.

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.

The VRF Lite Subnet IP Range field specifies resources reserved for IP address used for VRF LITE when VRF LITE IFCs are auto-created. If you select Back2BackOnly, ToExternalOnly, or Back2Back&ToExternal then VRF LITE IFCs are auto-created.

Auto Deploy Both - This check box is applicable for the symmetric VRF Lite deployment. When you select this check box, it would set the auto deploy flag to true for auto-created IFCs to turn on symmetric VRF Lite configuration.

This check box can be selected or deselected when the **VRF Lite Deployment** field is not set to Manual. In the case, a user explicitly unchecks the auto-deploy field for any auto-created IFCs, then the user input is always given the priority. This flag only affects the new auto-created IFC and it does not affect the existing IFCs.

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**.

Service Network VLAN Range - Specifies a VLAN range in the Service Network VLAN Range field. This is a per switch overlay service network VLAN range. The minimum allowed value is 2 and the maximum allowed value is 3967.

Route Map Sequence Number Range - Specifies the route map sequence number range. The minimum allowed value is 1 and the maximum allowed value is 65534.

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.

Note If AAA configs are not specified in the fabric settings, **switch_freeform** PTI with source as **UNDERLAY_AAA** and description as **DCNM Extra AAA Configurations** will be created.

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.

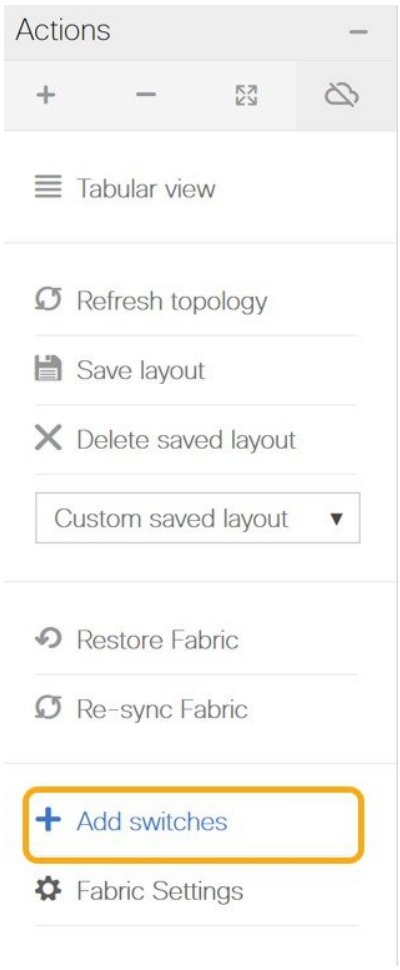

Step 2 Under the **Discover Existing Switches** tab, enter the IP address of the switch in the **Seed IP** field. Enter the username and password of the switches that you want to discover.

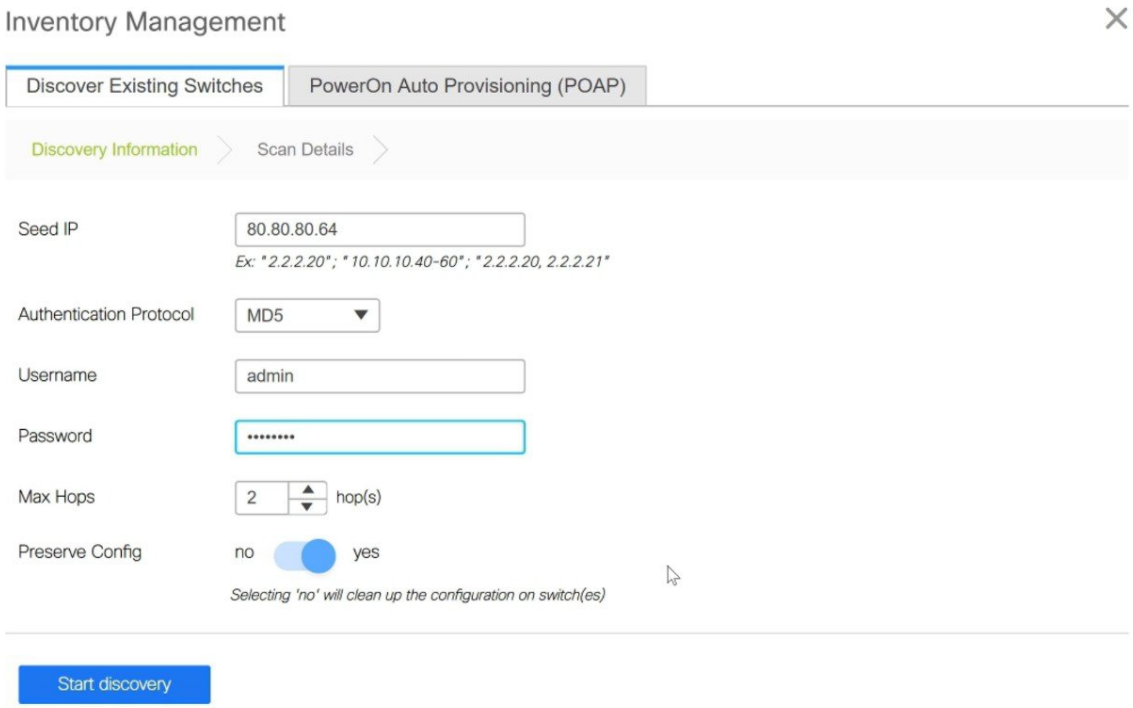

By default, the value in the **Max Hops** field is **2**. The switch with the specified IP address and the switches that are 2 hops from it will be populated after the discovery is complete.

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**.

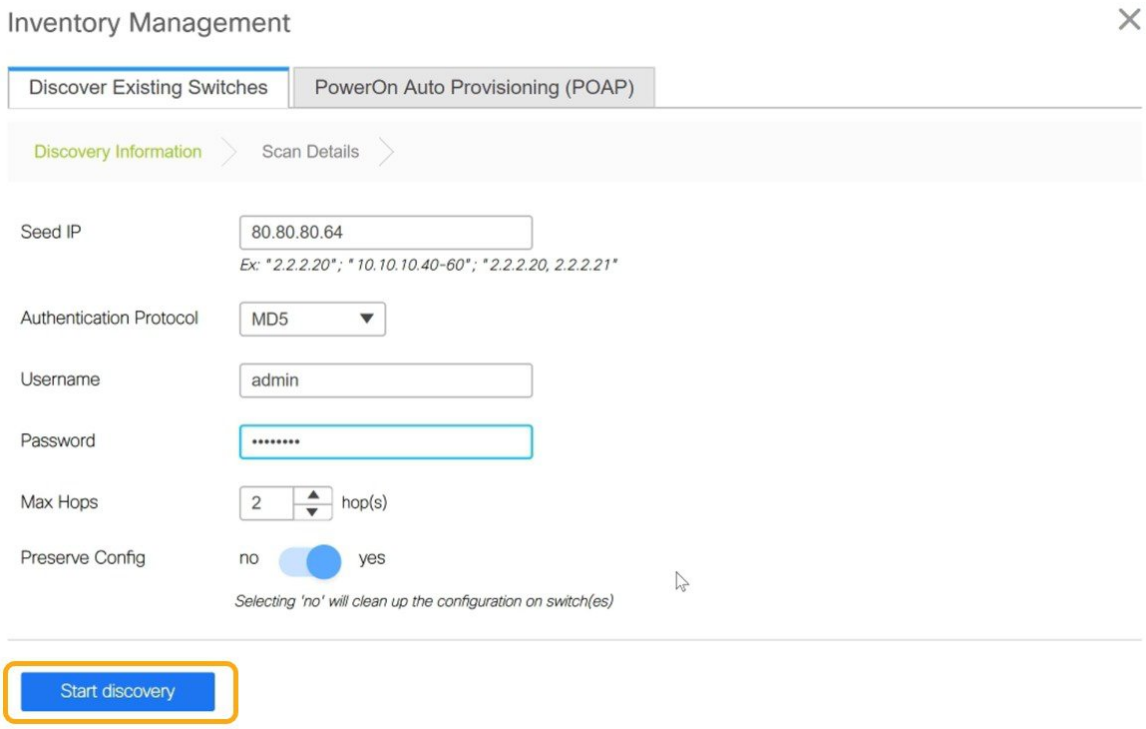

The switch with the specified IP address and switches up to two hops away (depending on the setting of Max Hops) from it are populated in the Scan Details section.

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**.

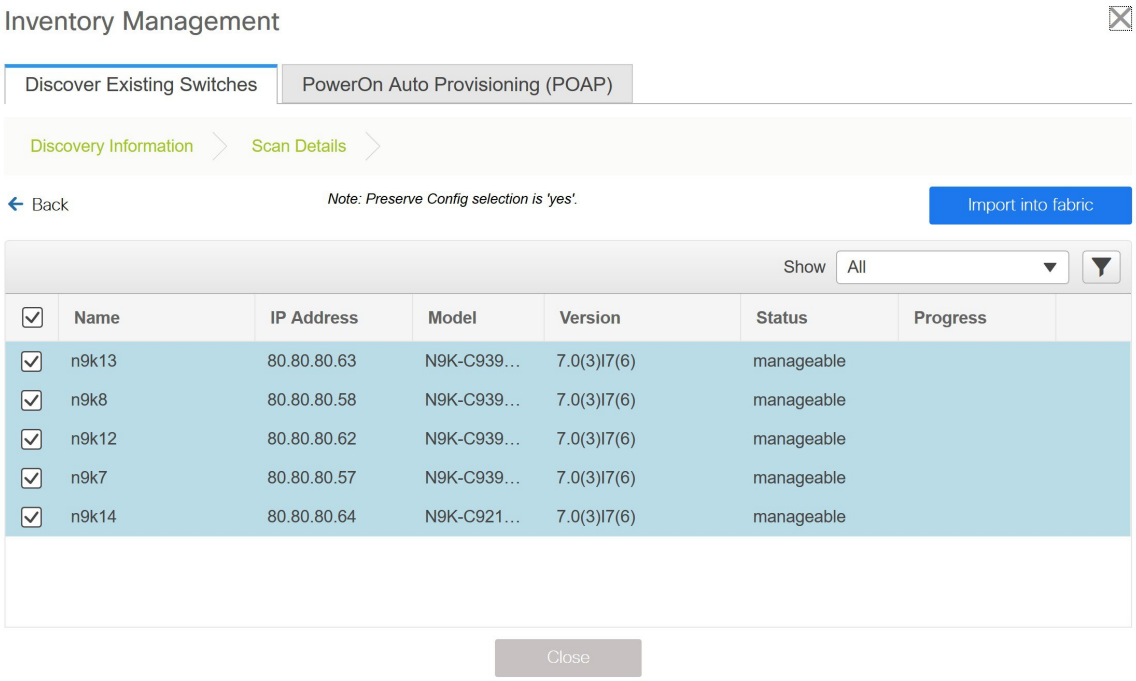

Step 5 Click **Import into fabric**.

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**.

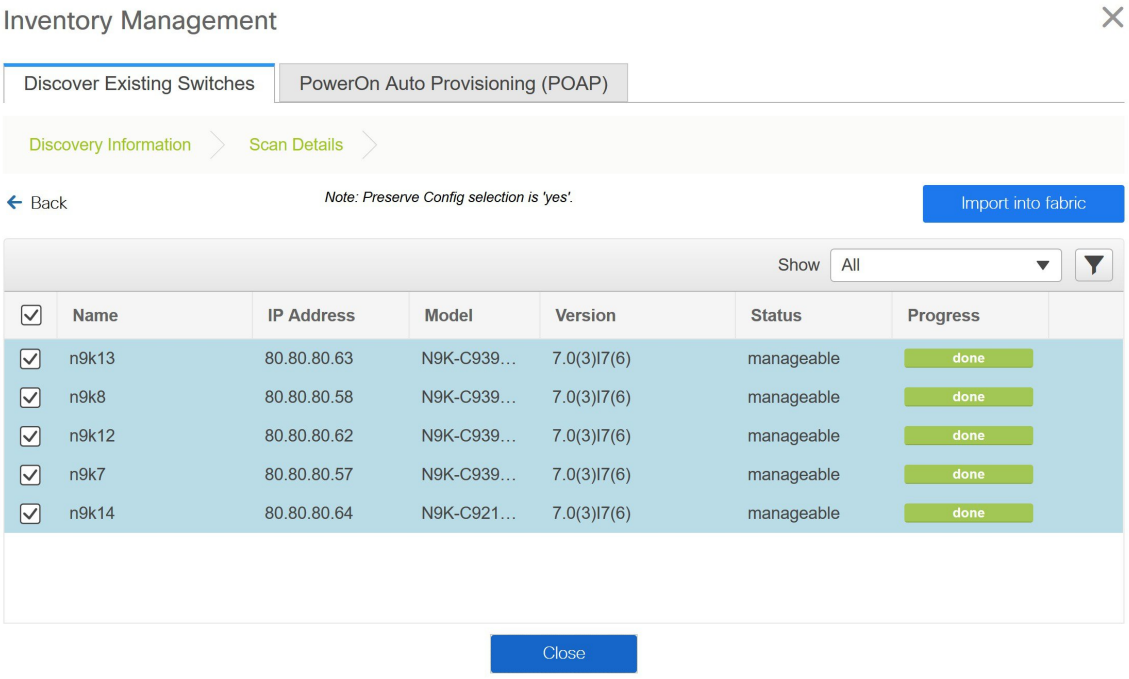

After closing the window, the fabric topology window comes up again. The switch is in Migration Mode now, and the Migration mode label is displayed on the switch icons.

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.

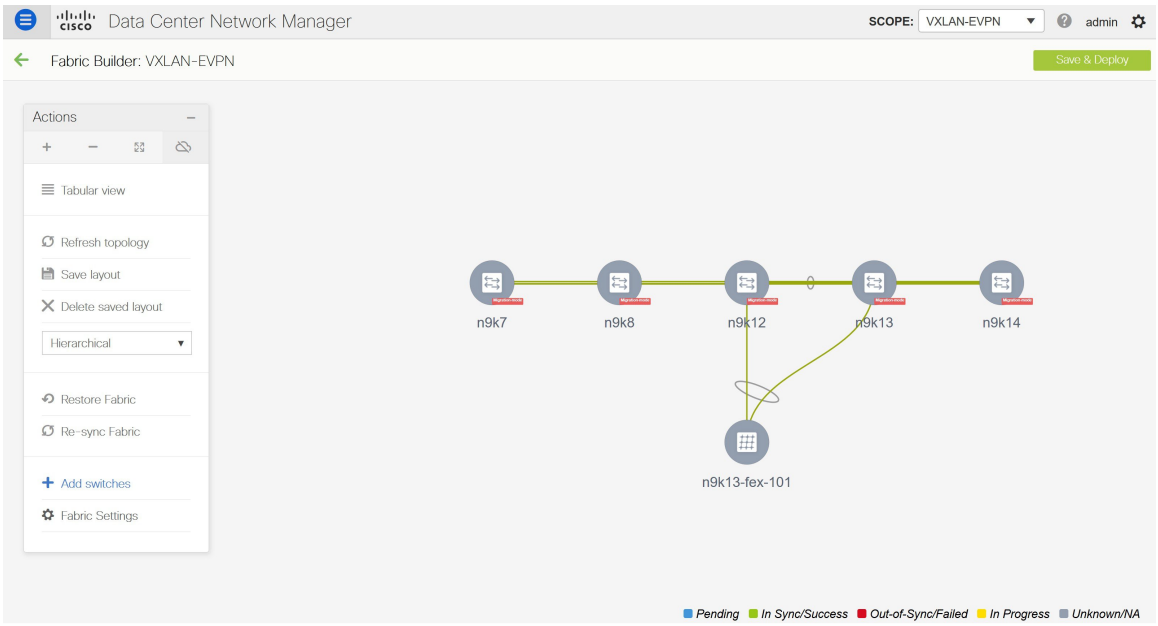

The switch discovery process might fail for a few switches, and the Discovery Error message is displayed. However, such switches are still displayed in the fabric topology. You should remove such switches from the fabric (Right-click the switch icon and click **Discovery** > **Remove** from fabric), 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.

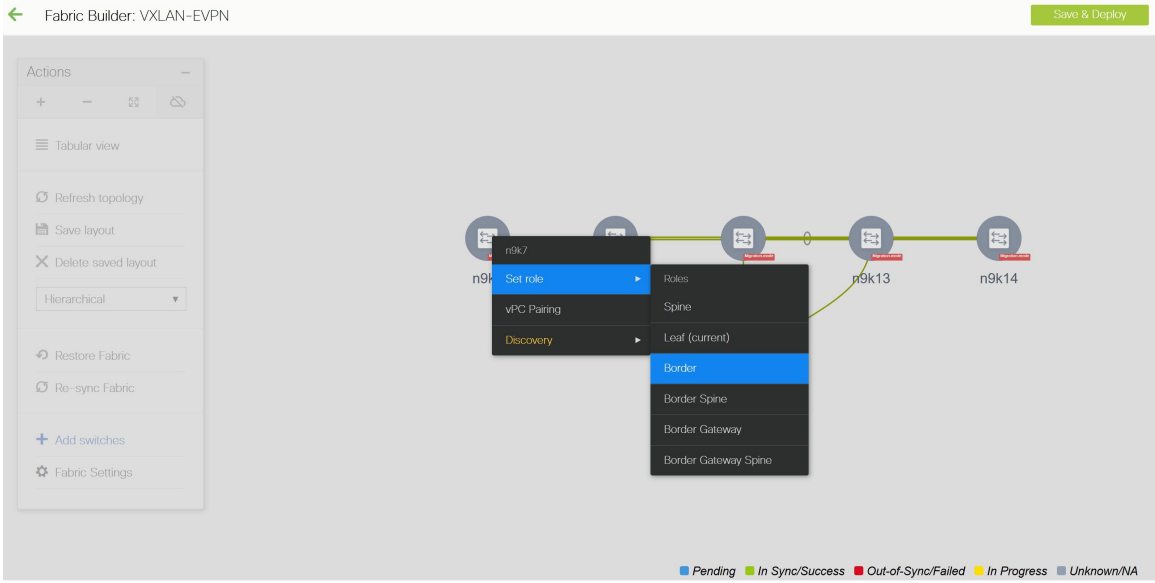

Similarly, set the **Spine** role for the **n9k-14** and **n9k-8** spine switches.

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 109.

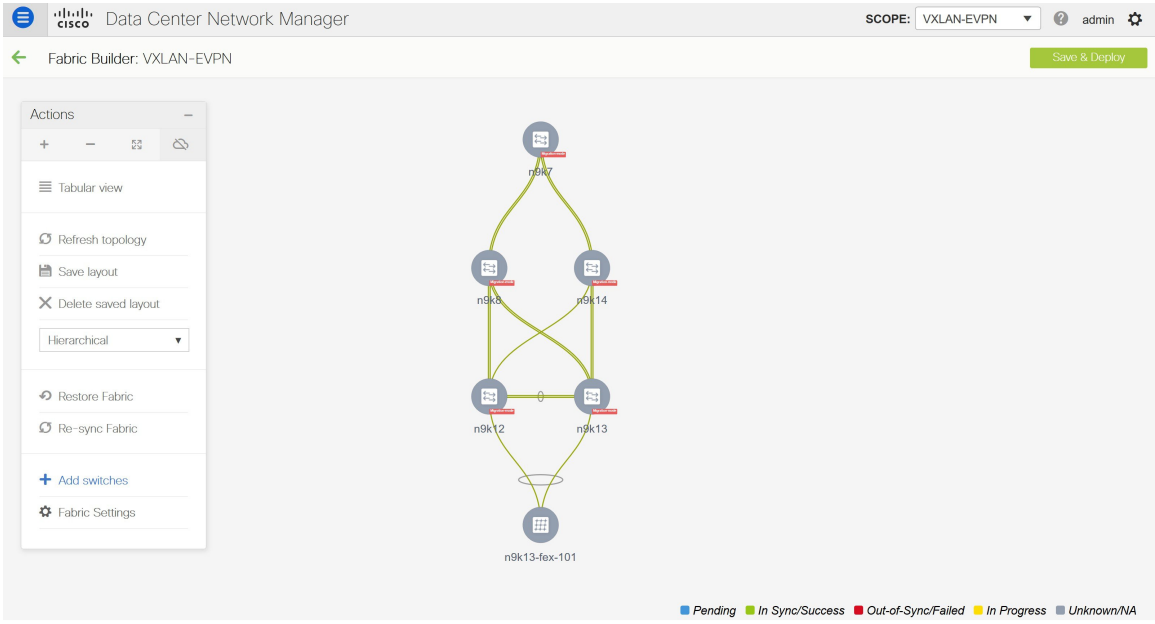

vPC Pairing - The vPC pairing must be done for switches where the Layer 3 vPC peer-keep alive is used. The vPC configuration is automatically picked up from the switches when the vPC peer keep alive is established through the management option. This pairing reflects in the GUI only after the migration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

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

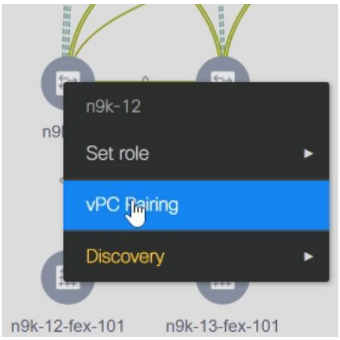

The Select vPC peer screen comes up. It lists potential vPC peer switches.

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

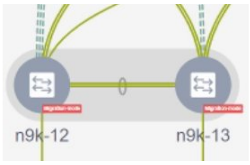

Note Check if you have added all switches from the current fabric. If you have missed adding switches, add them now. Once you are certain that you have imported all existing switches, move to the next step, the Save and Deploy option.

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 the *Control* chapter.
- 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.

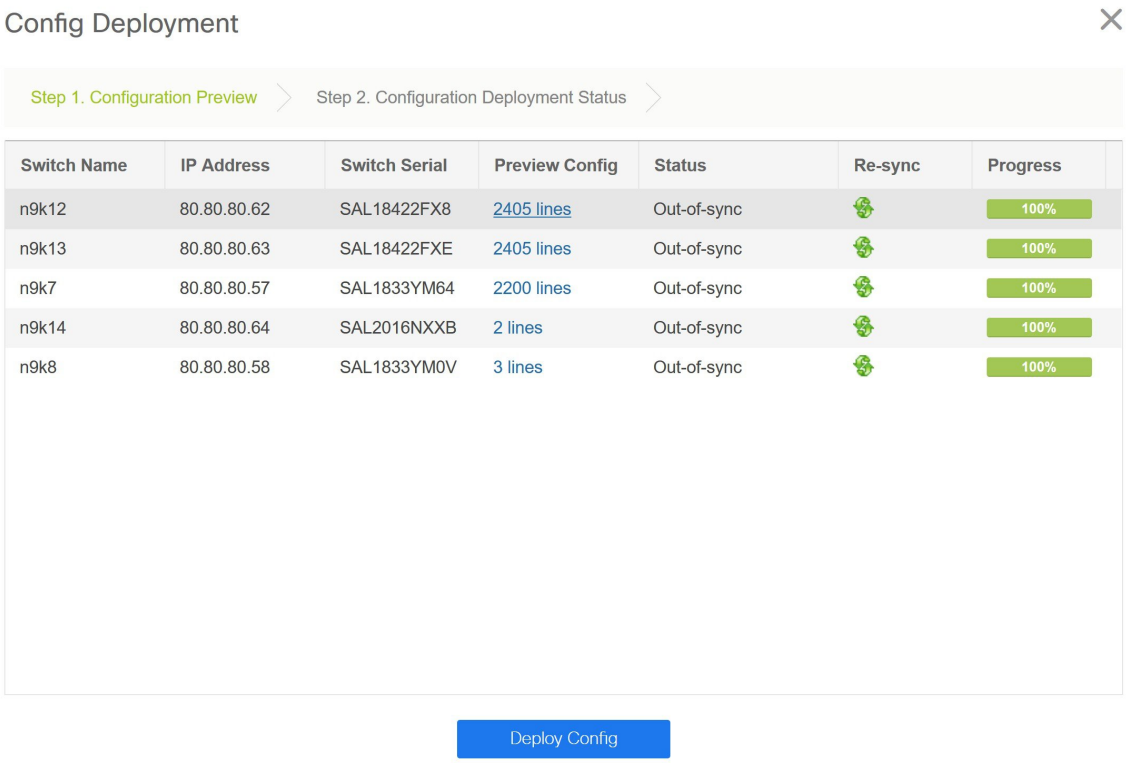

We strongly recommend that you preview the configuration before proceeding to deploy it on the switches. Click the Preview Config column entry. The Config Preview screen comes up. It lists the pending configurations on the switch.

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 a 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.

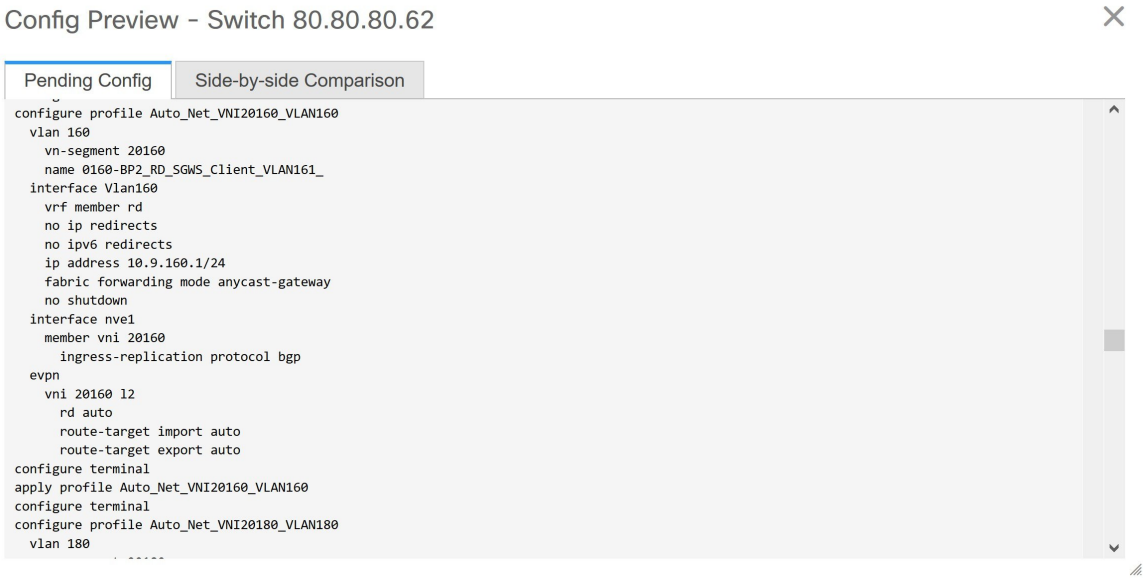

As part of the import process, after the configuration profiles are applied, the original CLI based configuration references will be removed from the switches. These are the 'no' CLIs that will be seen towards the end of the diffs. The VXLAN configurations on the switches will be persisted in the configuration profiles. In the following image, you can see that the configurations will be removed, specifically, **no vlan 160**.

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.

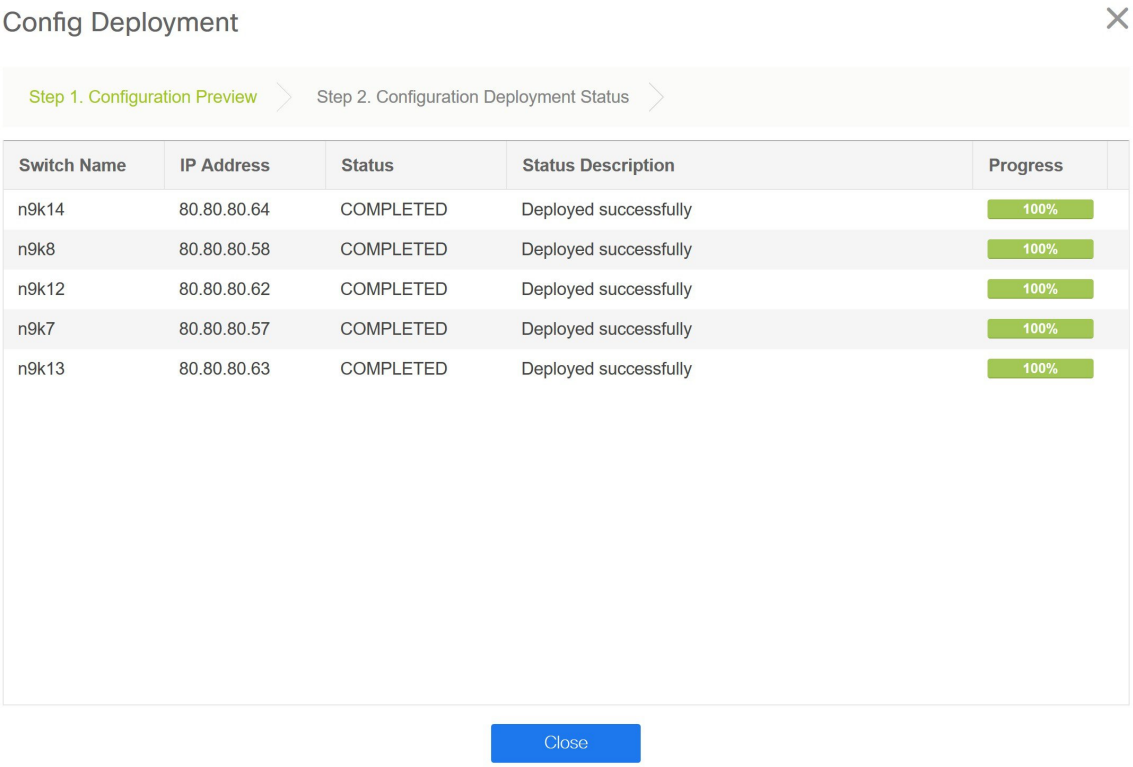

If the **Status** column displays **FAILED**, investigate the reason for failure to address the issue.

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.

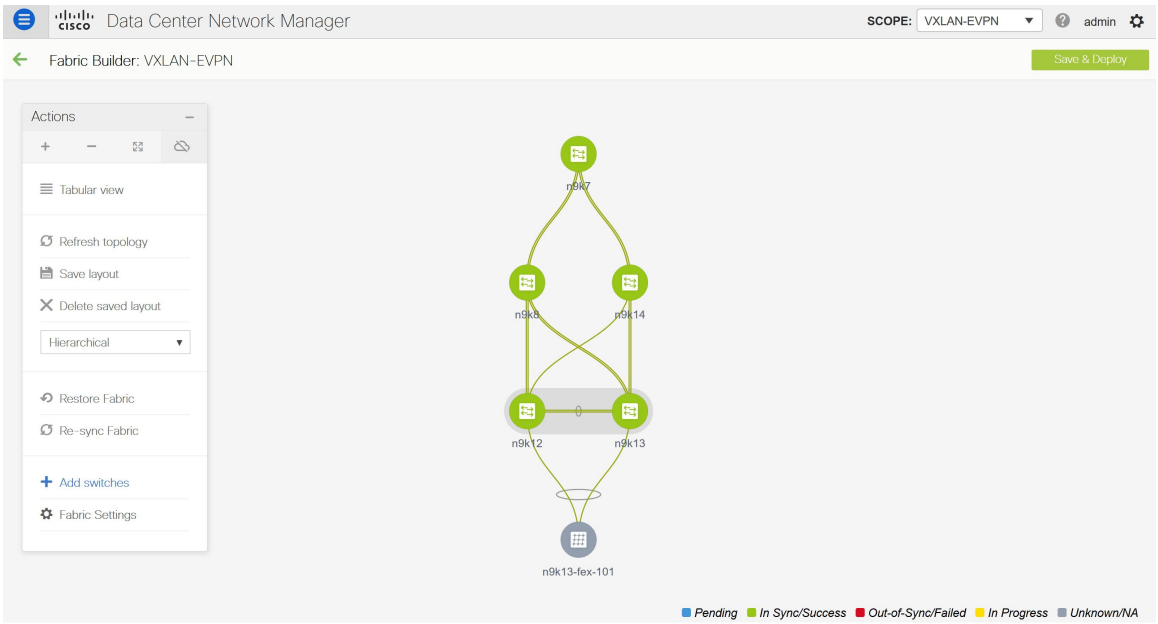

Post-transitioning of VXLAN fabric management to DCNM - This completes the transitioning process of VXLAN fabric management to DCNM. Now, you can add new switches and provision overlay networks for your fabric. For details, refer the respective section in the Fabrics topic in the configuration guide.

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.
- **Backup Now**: You can initiate a fabric backup manually by clicking **Backup Now**. Enter a name for the tag and click **OK**. Regardless of the settings you choose under the **Configuration Backup** tab in the **Fabric Settings** dialog box, you can initiate a backup using this option.

- **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.

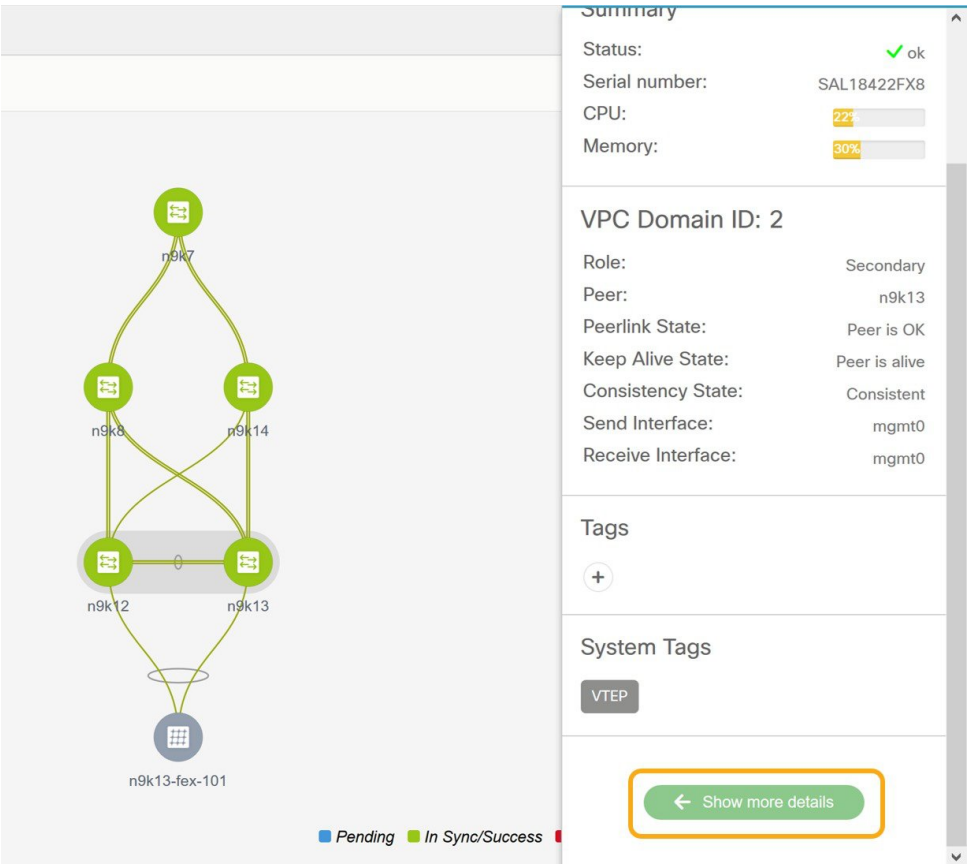

Step 2 Click the **VXLAN** tab.

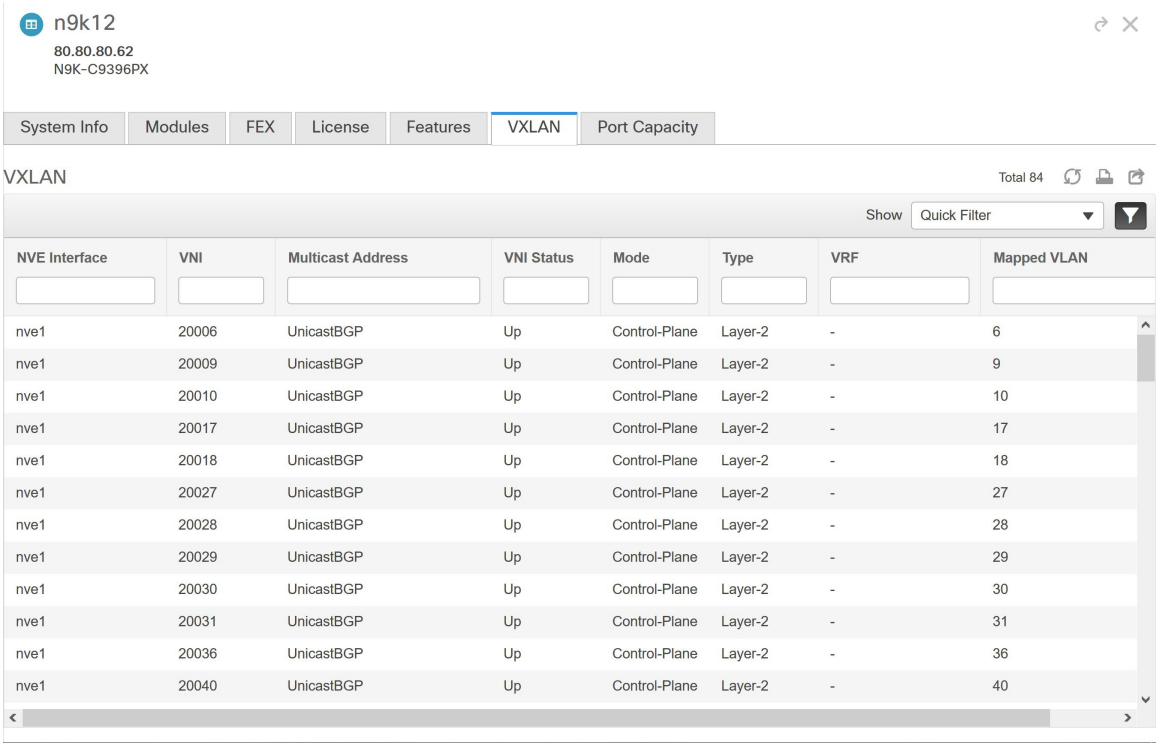

You can see that all the VXLANs have been migrated successfully.

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.

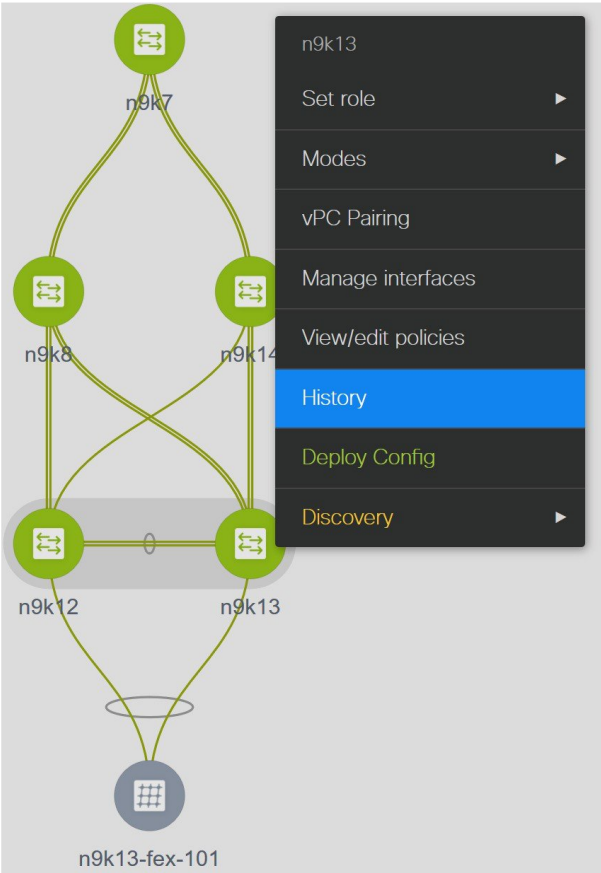

Step 4 Click the **Success** hyperlink under the **Status** column to view the commands pushed by DCNM.

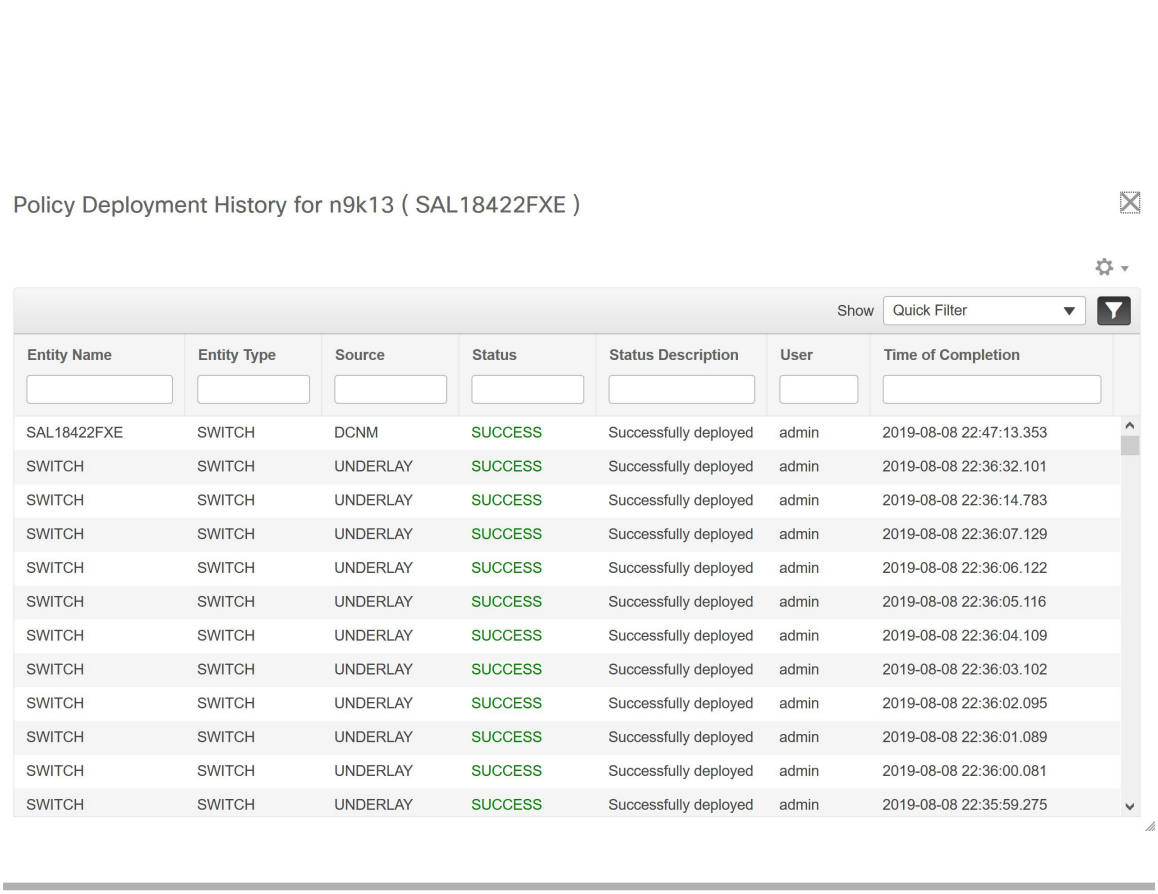

Verifying Resources

DCNM has a resource manager that tracks all the resources used in a fabric. Navigate to **Control > Management > Resources** in the left menu.

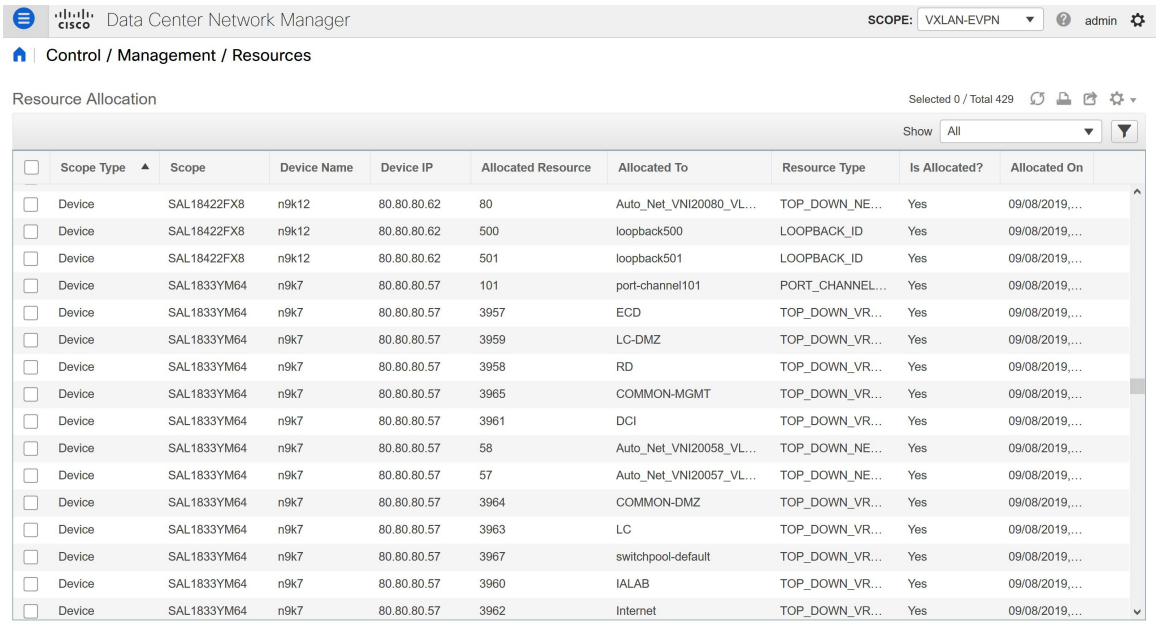

The resources that are being utilized by the VXLAN EVPN fabric such as VLAN IDs, port channel IDs, point to point IP addresses, and loopback IDs are displayed in this window.

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.

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.

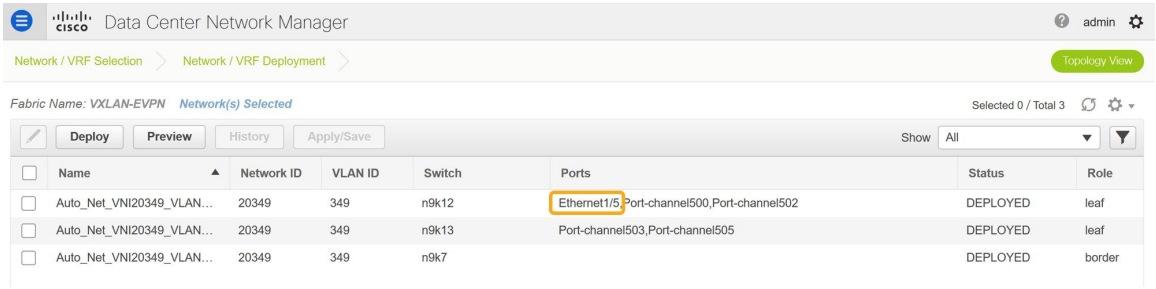

Let us verify the overlay network associated with this interface.

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**.

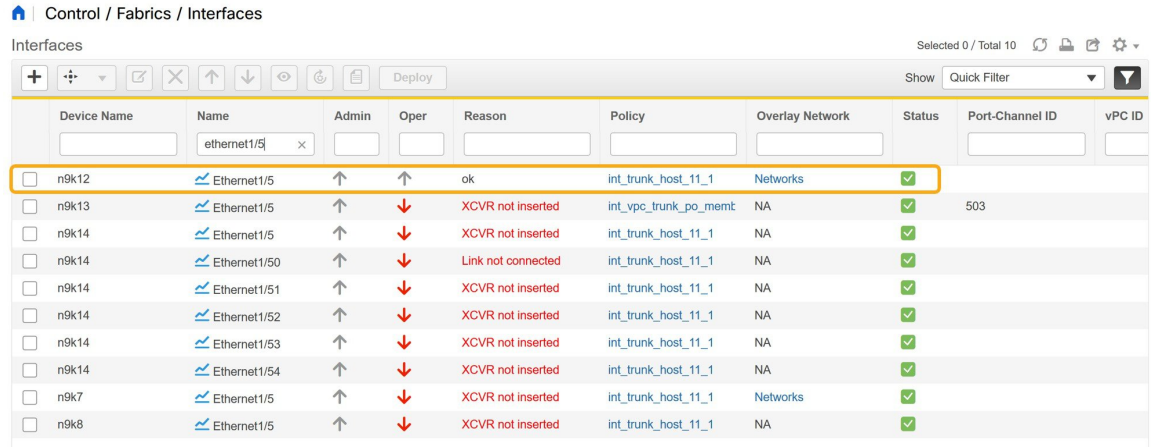

This interface is attached to the host through the **n9k-12 switch**.

Step 7 In the **Overlay Networks** column, click **Networks** that corresponds to the n9k-12 switch and the Ethernet 1/5 interface.

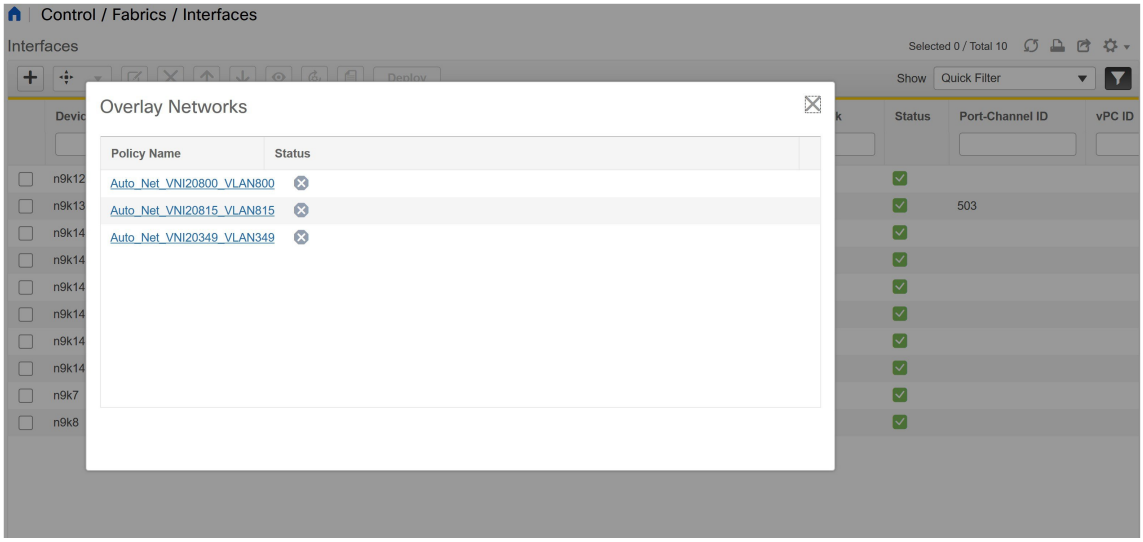

These are the networks that are attached to the **Ethernet 1/5** interface.

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.

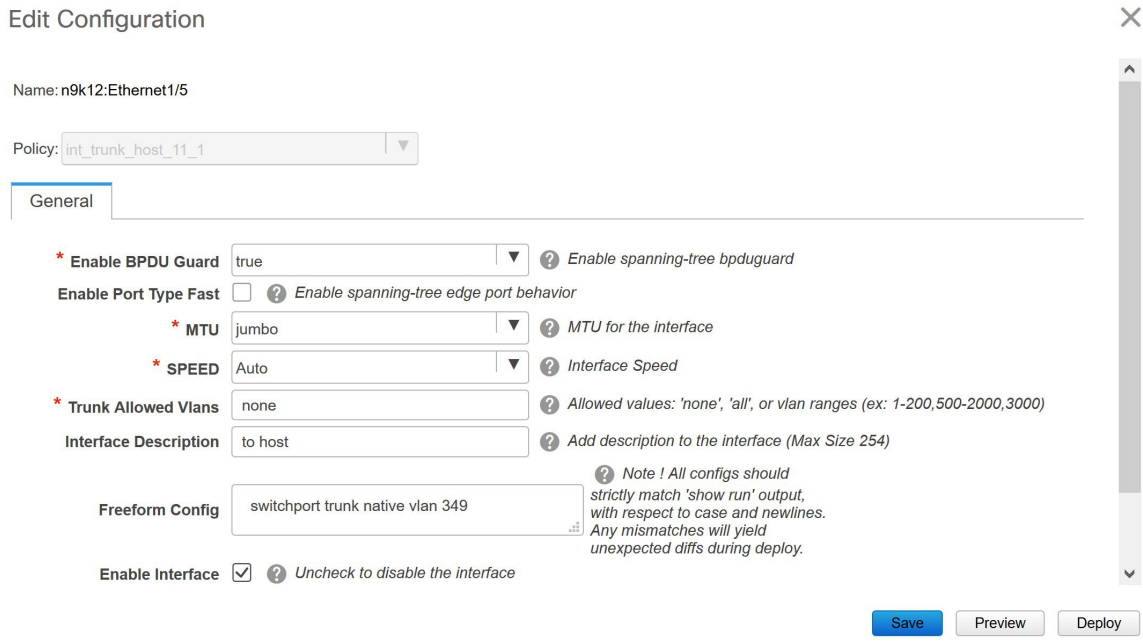

You can see that all the settings for this interface have been successfully imported, including the BPDU guard settings and the interface description.

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.

Configuration Profiles Support for Brownfield Migration

Cisco DCNM Release 11.3(1) supports the Brownfield import of fabrics with VXLAN overlay provisioned with configuration profiles. This import process recreates the overlay configuration intent based on the configuration profiles. The underlay migration is performed with the usual Brownfield migration.

This feature can be used to recover your existing Easy fabric when a DCNM backup is not available to be restored. In this case, you must install the latest DCNM release, create a fabric, and then import the switches into the fabric.

Note that this feature is not recommended for the DNCM upgrade. For more information, see *Cisco DCNM Installation and Upgrade Guide for LAN Fabric Deployment*.

The following are the guidelines for the support of configuration profiles:

- The Brownfield migration of configuration profiles is supported for the **Easy_Fabric_11_1** template.
- The configuration profiles on the switches must be a subset of the default overlay **Universal** profiles. If extra configuration lines are present that are not part of the **Universal** profiles, unwanted profile refreshes will be seen. In this case, after you click **Save & Deploy**, review the diffs using the **Side-by-side Comparison** feature and deploy the changes.
- Brownfield migration with switches having a combination of VXLAN overlay configuration profiles and regular CLIs is not supported. If this condition is detected, an error is generated, and migration is aborted. All the overlays must be with either configuration profiles or regular CLIs only.

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.

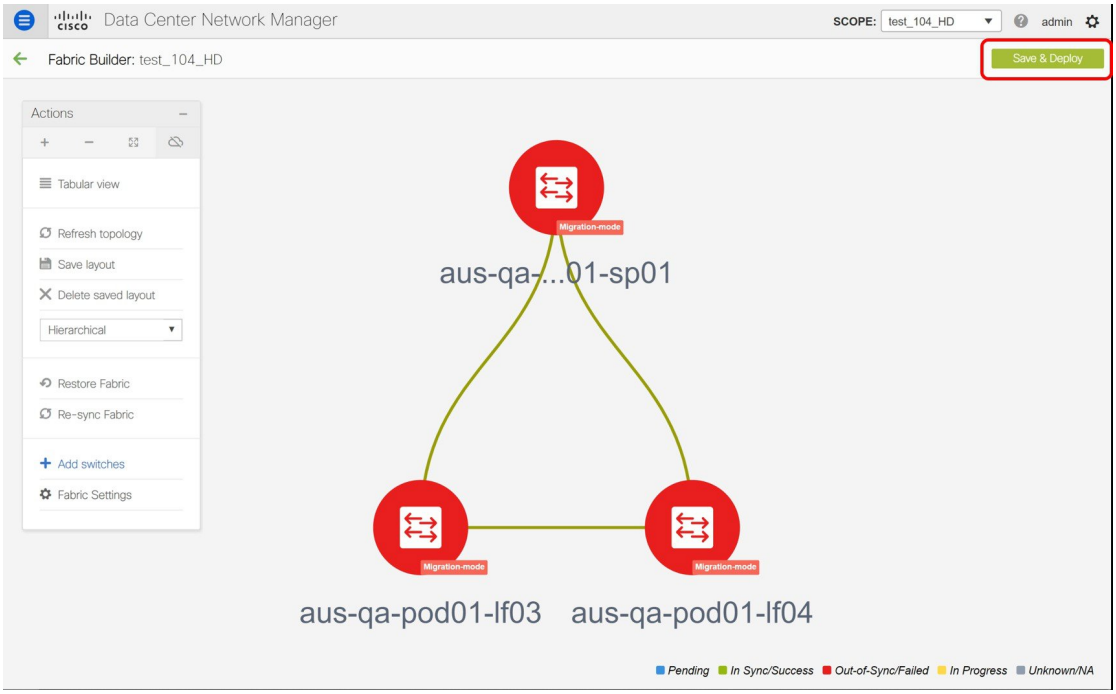

If the added switches contain bottom-up configurations, an error is displayed saying – Reloading switch to clean up bottom up config. Please wait for switch to come online and try **Save & Deploy**.

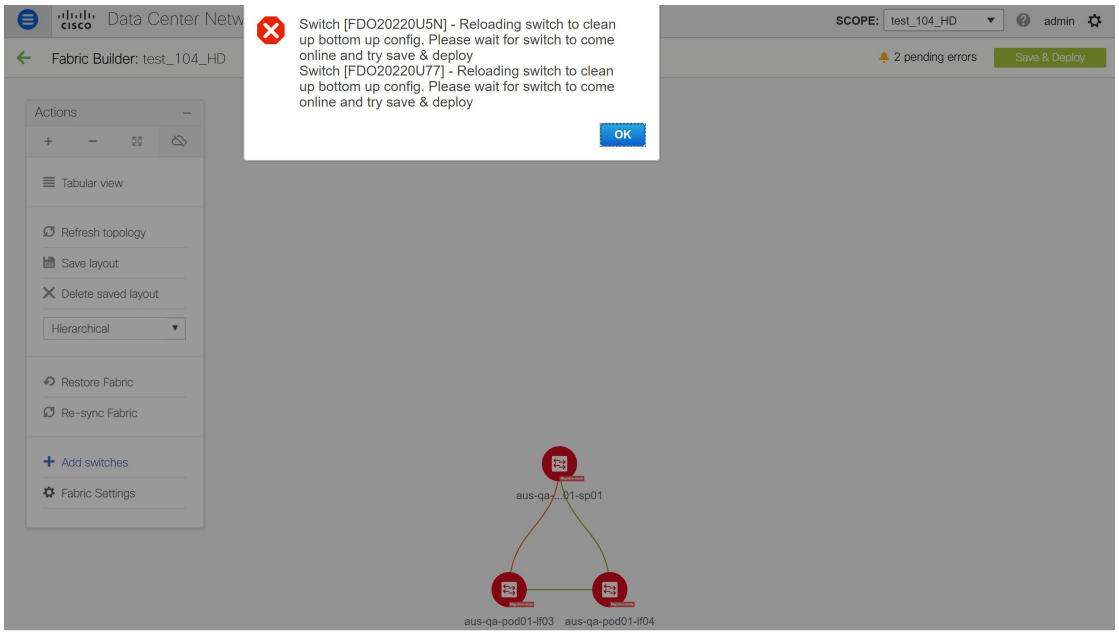

4. Wait for the switches to complete the reload operation. Click **Tabular view** under the **Actions** menu to view the status of the switches.

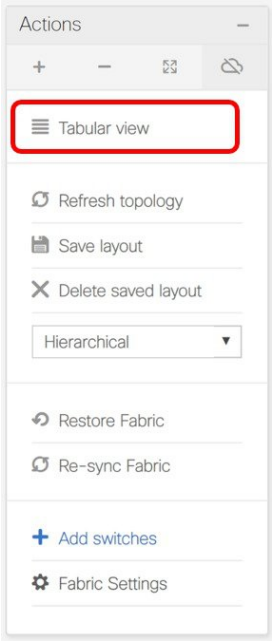

5. (Optional) Rediscovery of the reloaded switches occurs every 5 minutes. If you want to manually rediscover switches, select the switches and click the **Rediscover switch** icon.

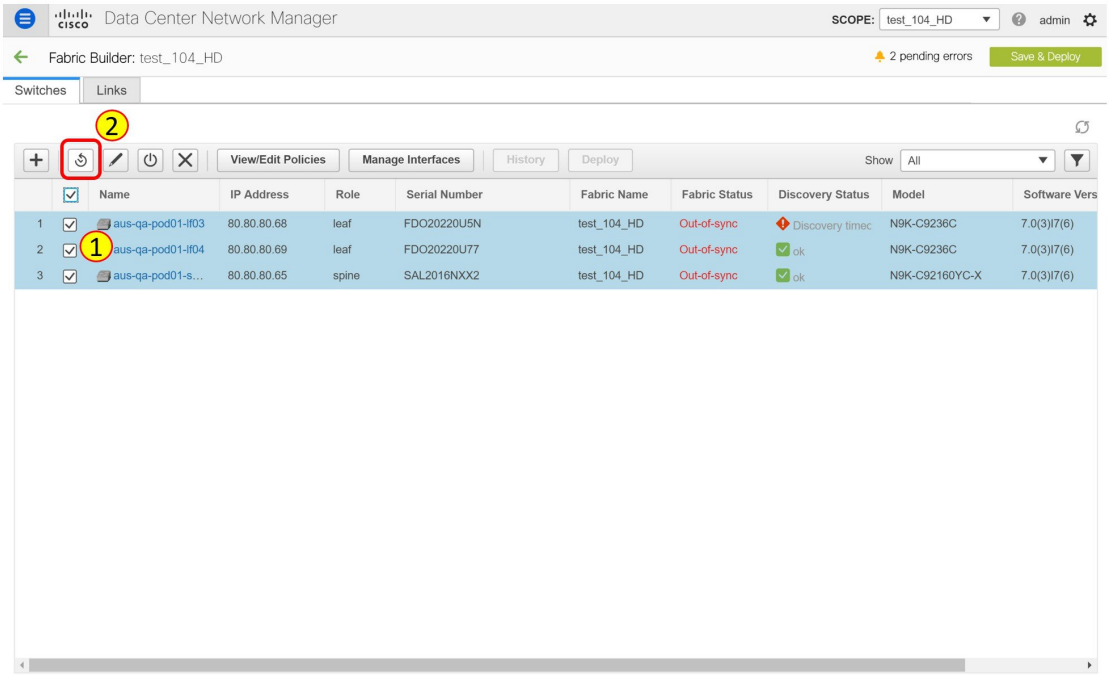

Note Click the **Refresh** icon to refresh the **Fabric Builder** window and see the updated discovery status of switche

6. 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**.

Note When a switch is in **Unreachable** discovery status, the last available information of the switch is retained other columns.

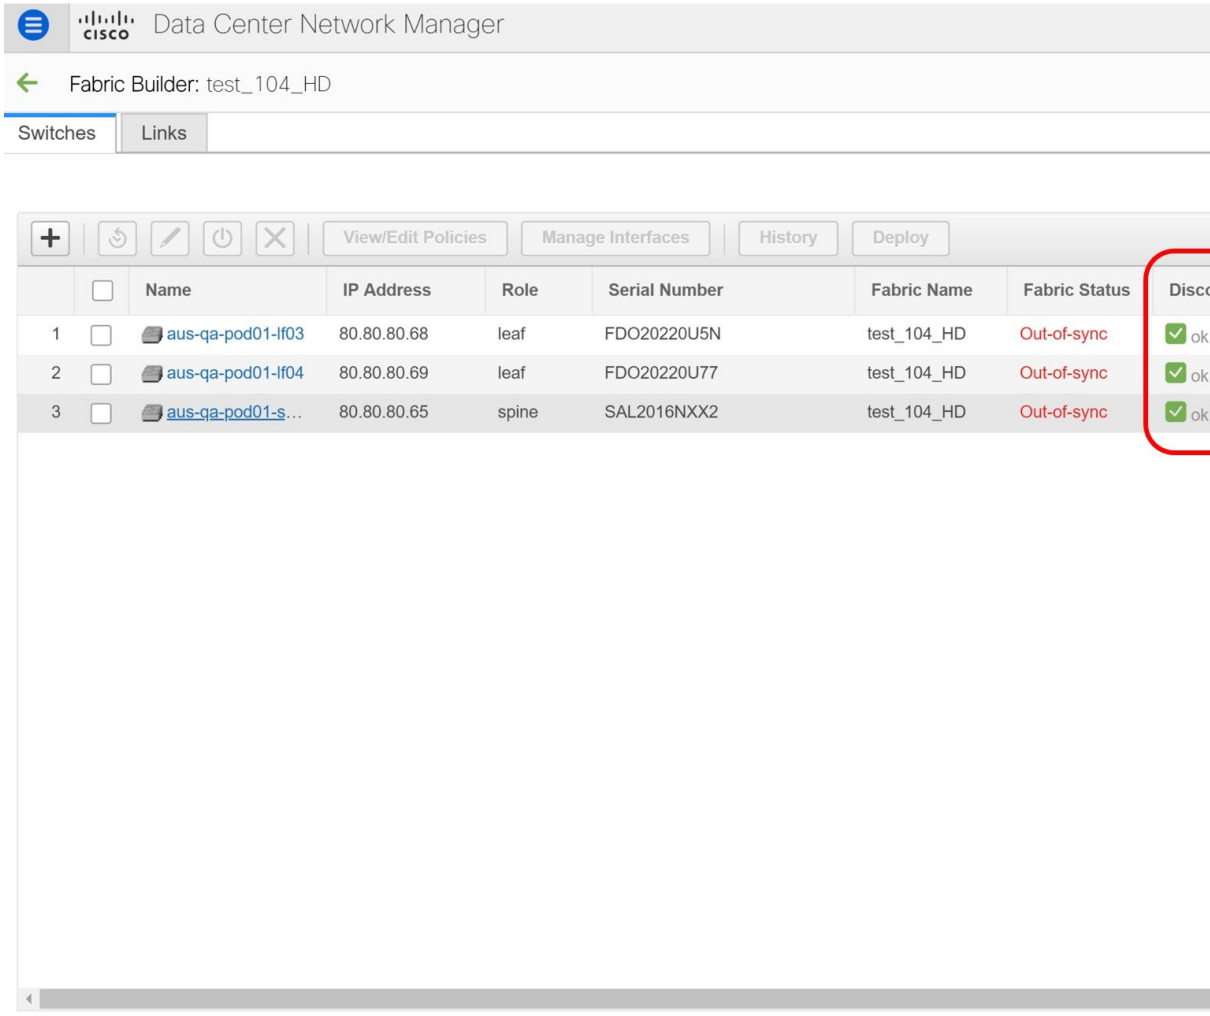

7. Click **Save & Deploy** again to sync configurations between the switches and DCNM.

   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.

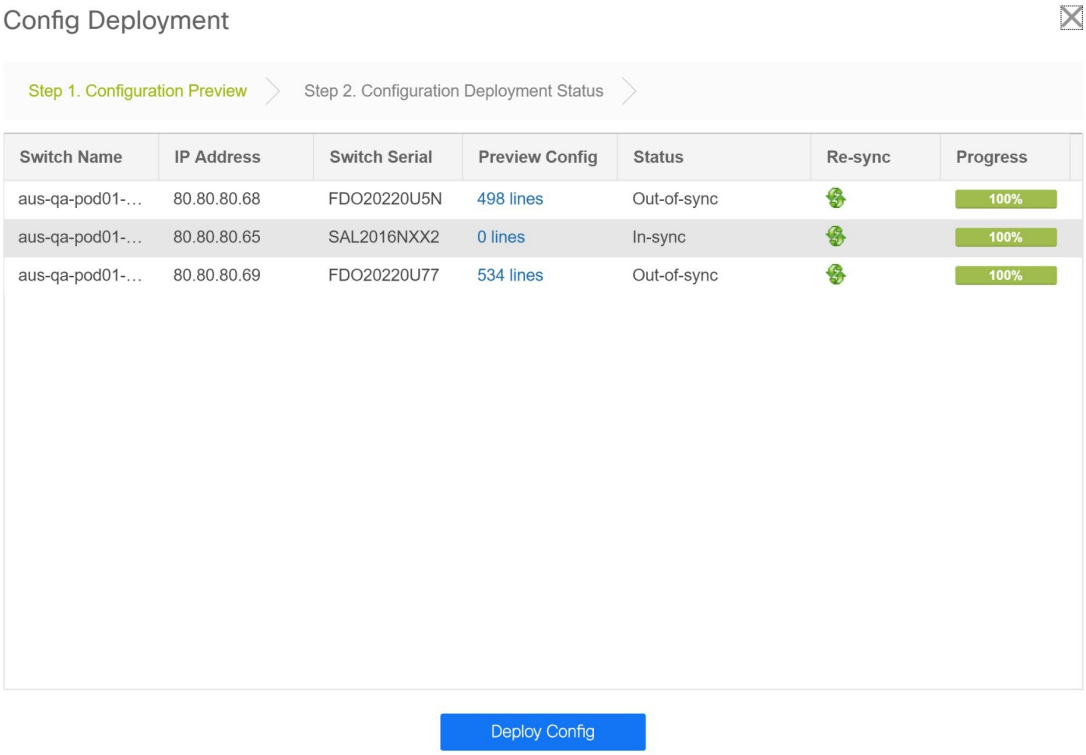

The **Preview Config** column is updated with entries denoting a specific number of lines.

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.

Close the **Config Preview** window.

8. 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.

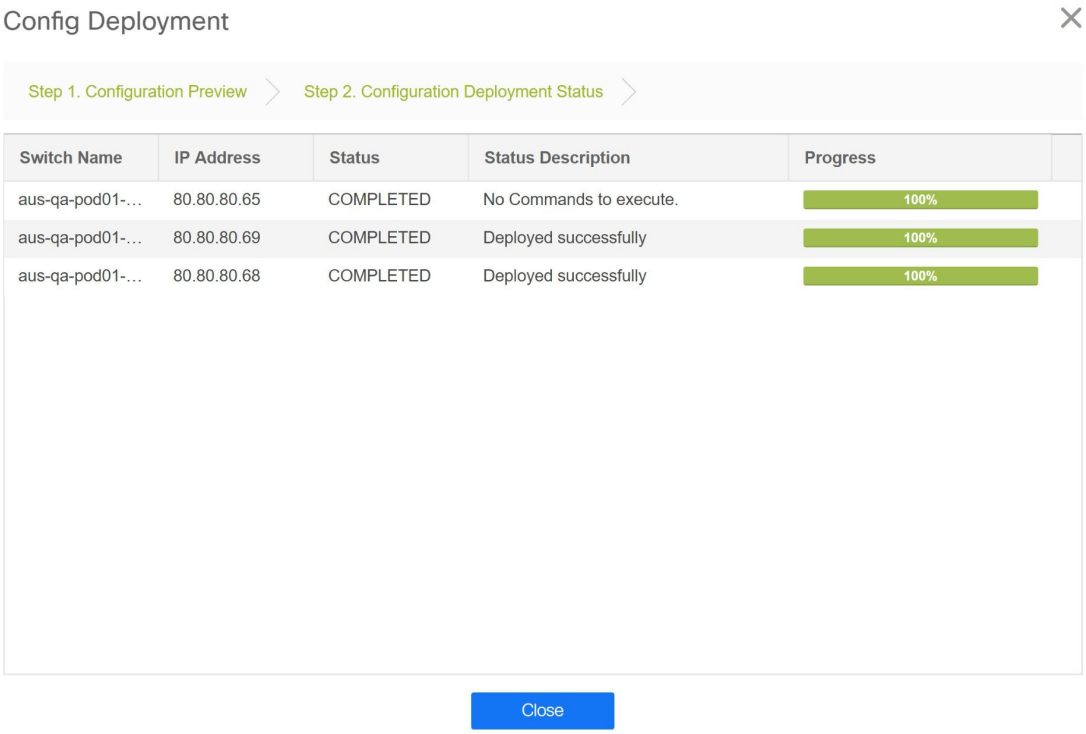

The progress bar shows 100% for each switch. After correct provisioning and successful configuration compliance, close the **Config Deployment** window.

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.

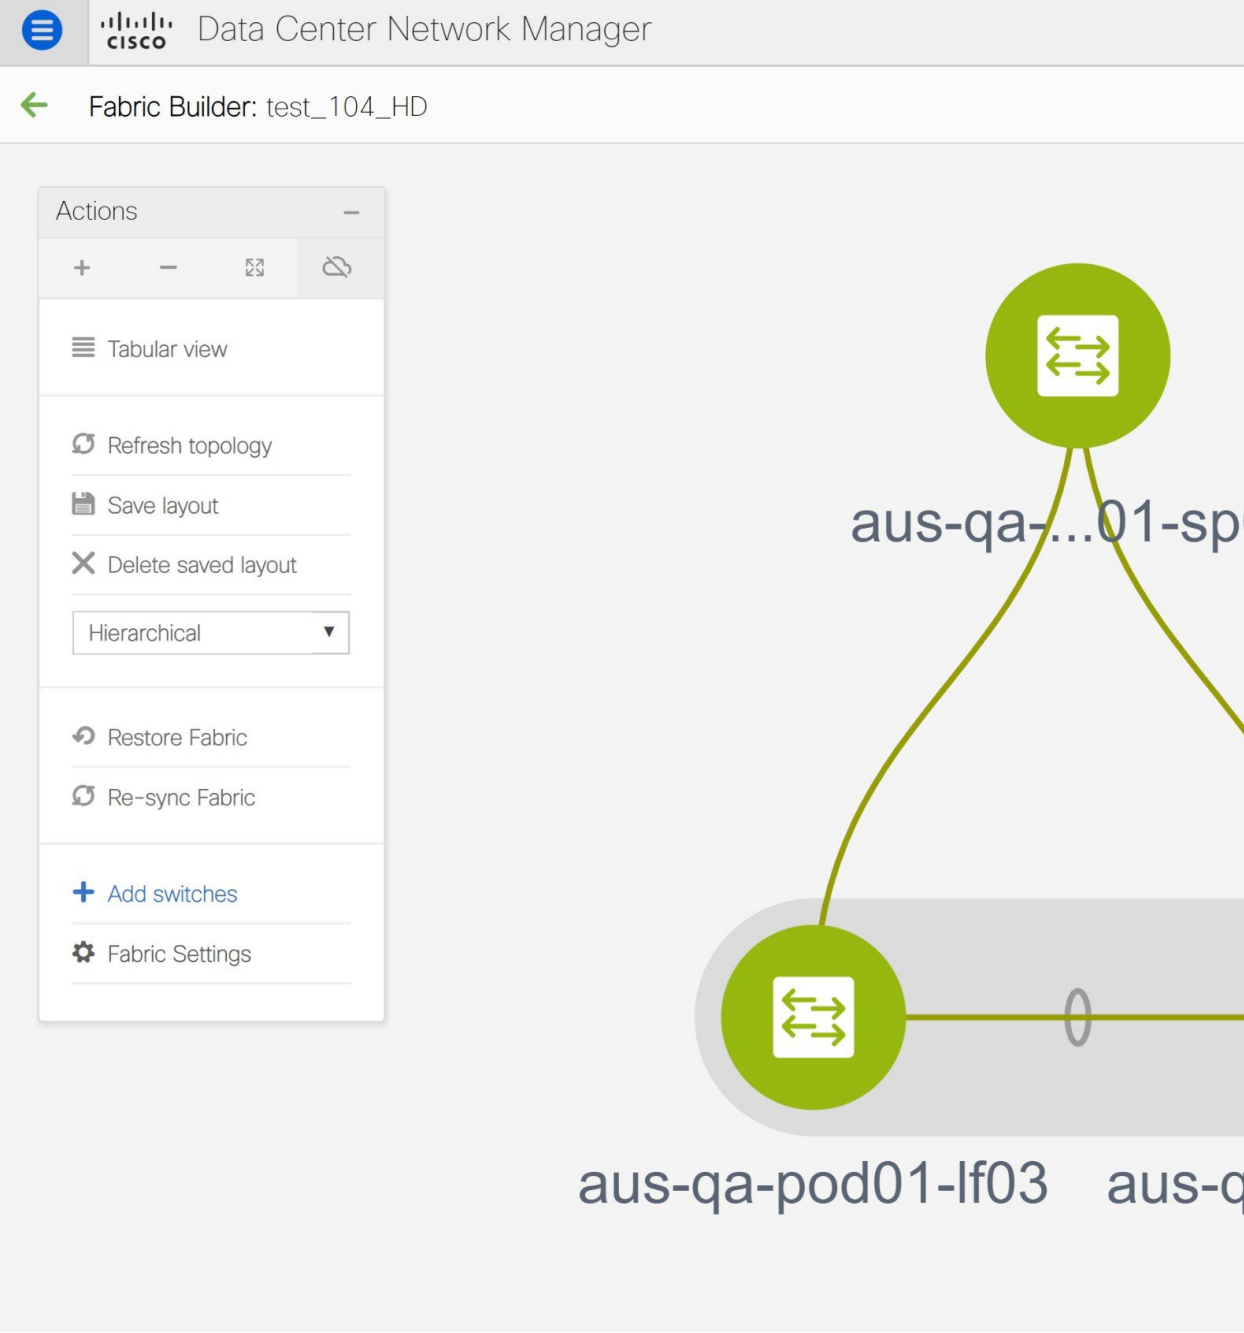

This completes the migration process of bottom-up VXLAN fabric to DCNM.

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**.

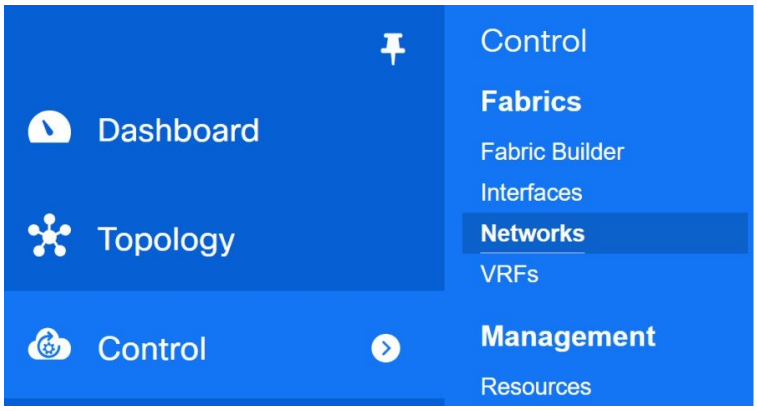

2. Select the fabric from the **SCOPE** drop-down list in the **Networks** window.

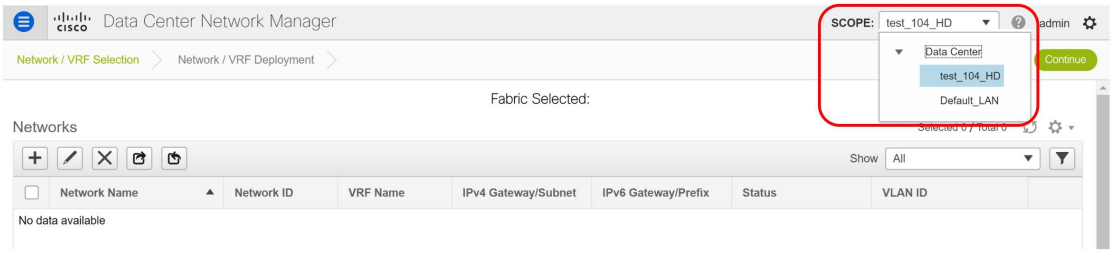

3. Check the networks that are migrated from the bottom-up VXLAN fabric and their deployment status.

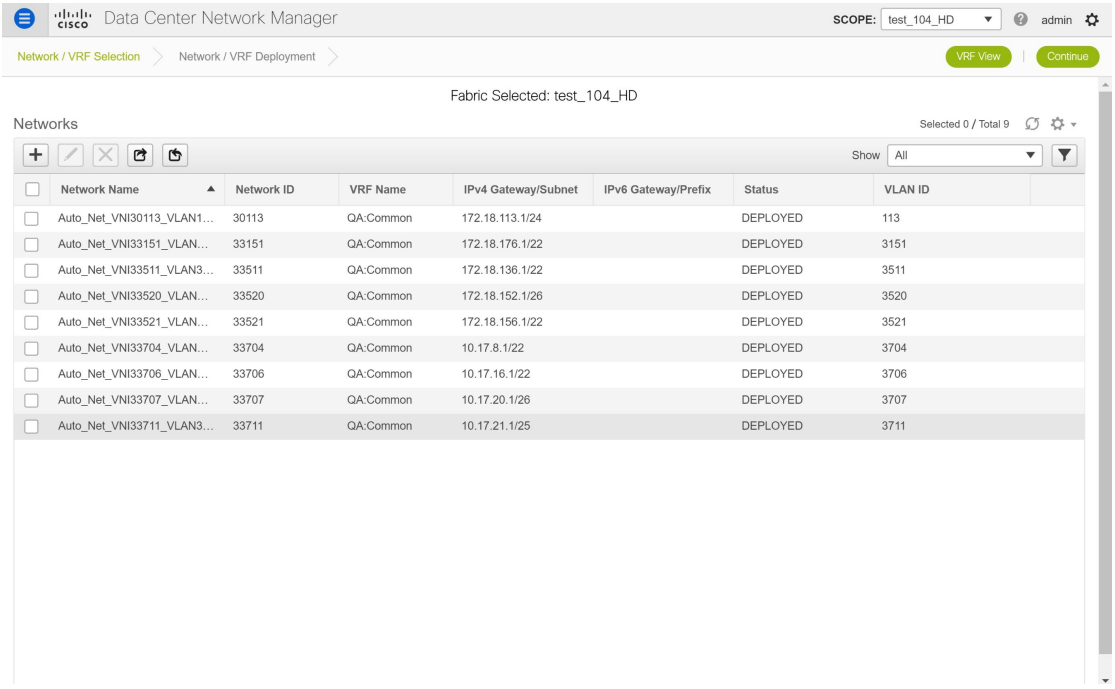

Resolving Config Compliance Error on Switches with Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) Images

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.

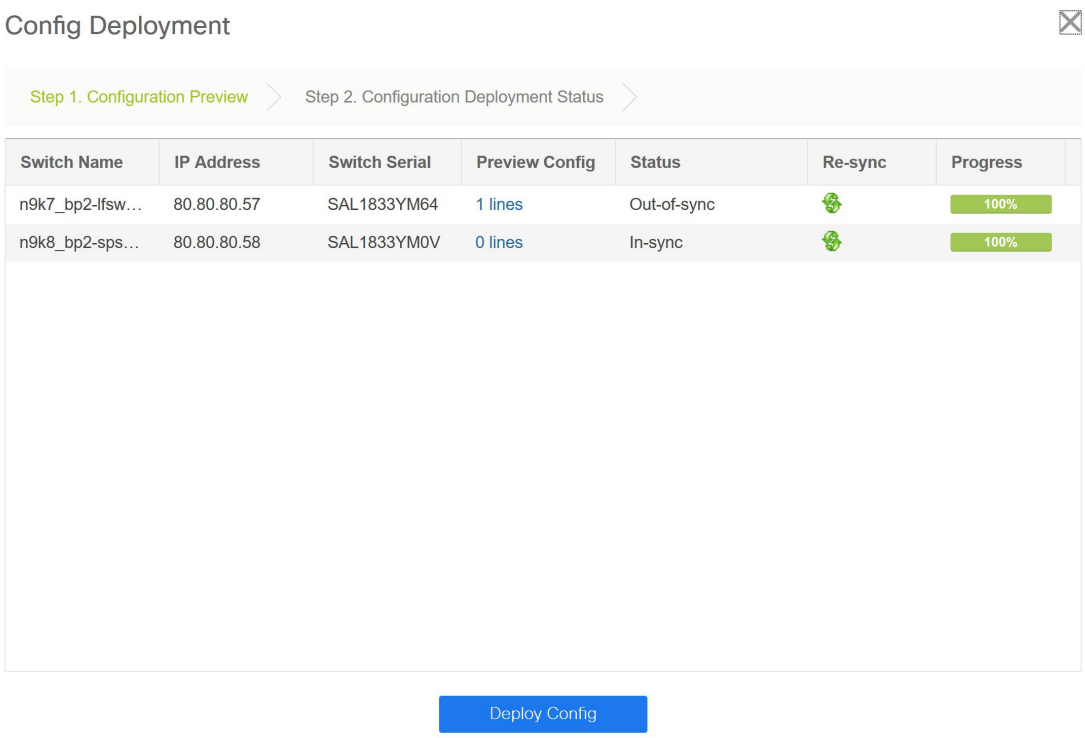

4. (Optional) Click the entry showing **1 lines** under the **Preview Config** column.

   You can see the TCAM command under the **Pending Config** tab in the **Config Preview** window.

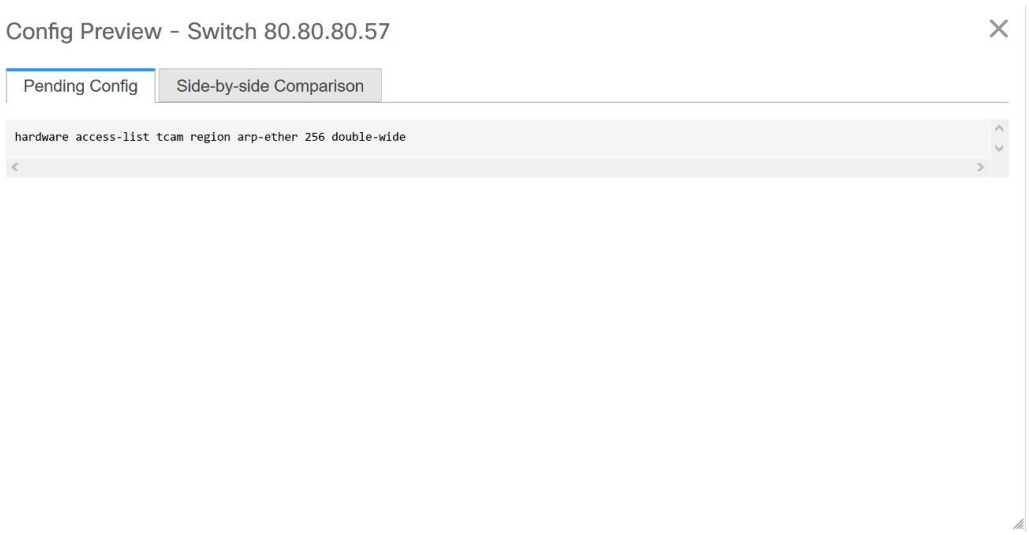

Close the **Config Preview** window.

5. Right-click a switch and click **View/Edit Policies**.

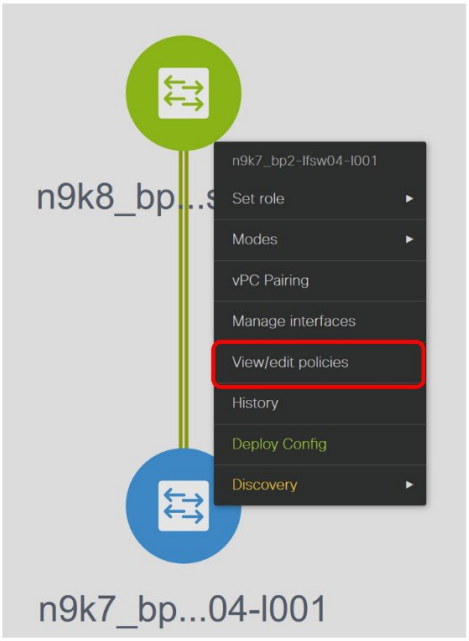

6. Search for the **tcam_pre_config_vxlan** policy in the **Template** search field.
7. Select the **tcam_pre_config_vxlan** policy and click the **Delete** icon to delete the policy.

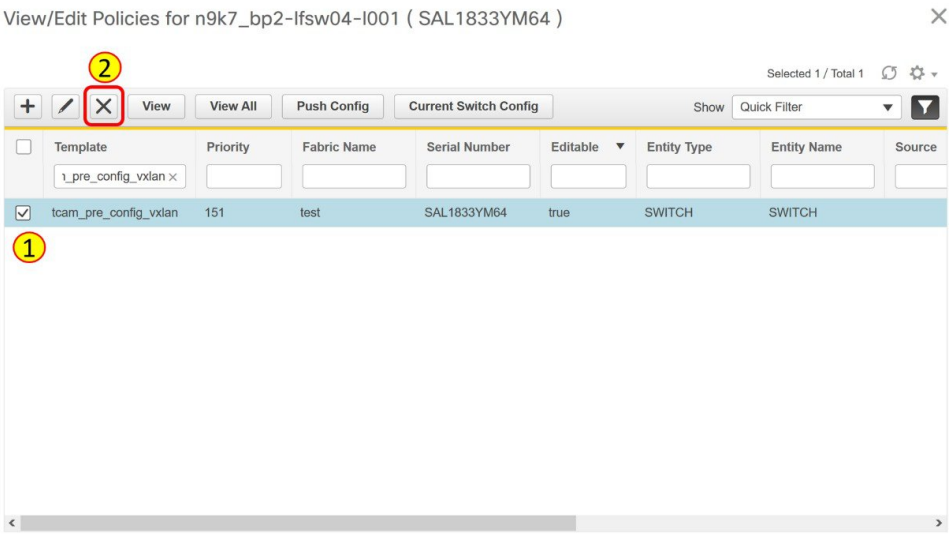

Close the **View/Edit Policies** window.

8. (Optional) Click **Save & Deploy** to verify whether there are any pending configs.

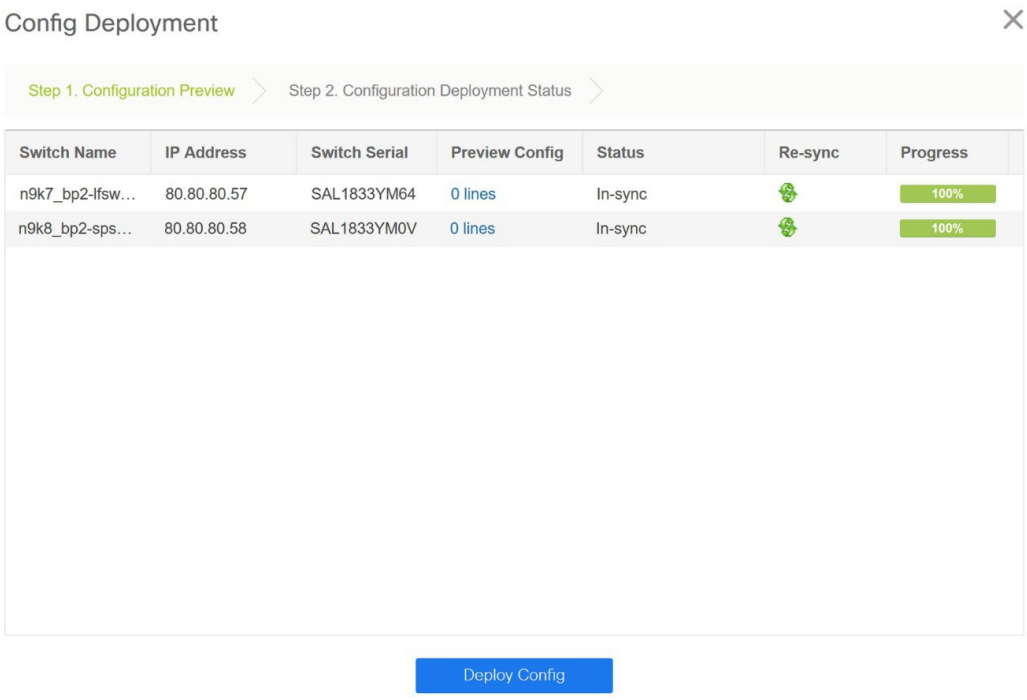

Resolving Config Compliance Error on Switches for RMA, and Write Erase and Reload Operations

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.

1. 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.

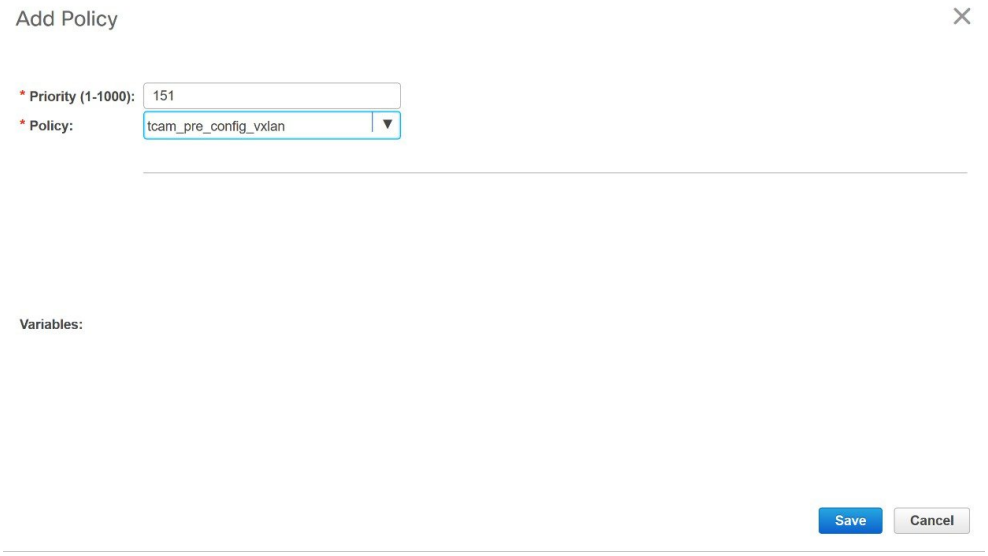

6. Click **Save**.
7. Complete the RMA or Write Erase and Reload operation.

   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.

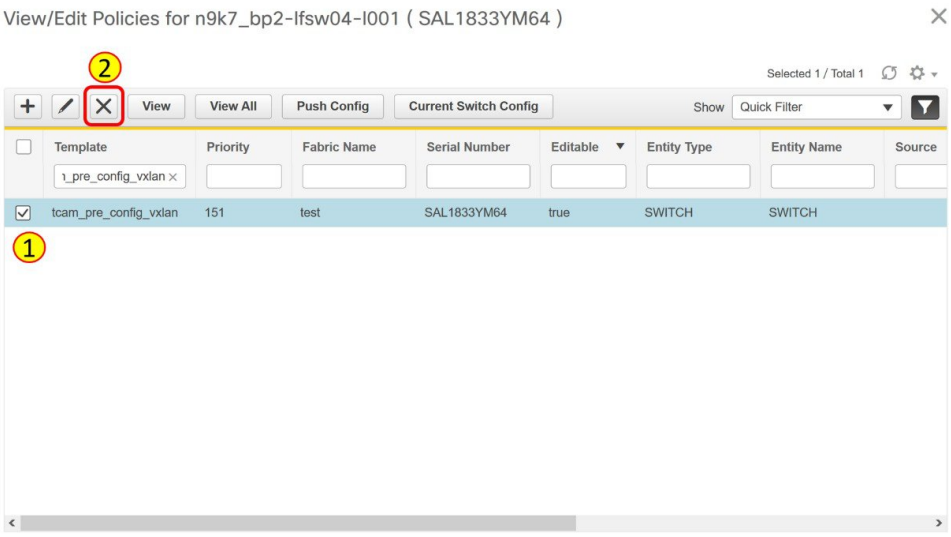

Close the **View/Edit Policies** window.

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.

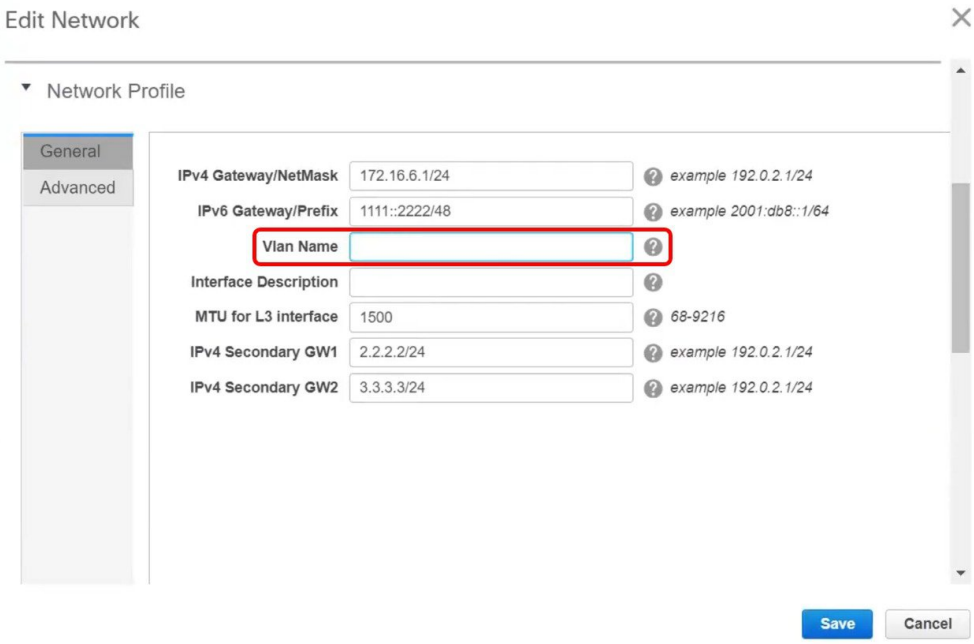

In the **Edit Network** window, the **Vlan Name** field is empty because DCNM has not captured this info in the overlay profile. Instead, the VLAN name is captured in the freeform config associated with the overlay network or VRF.

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.

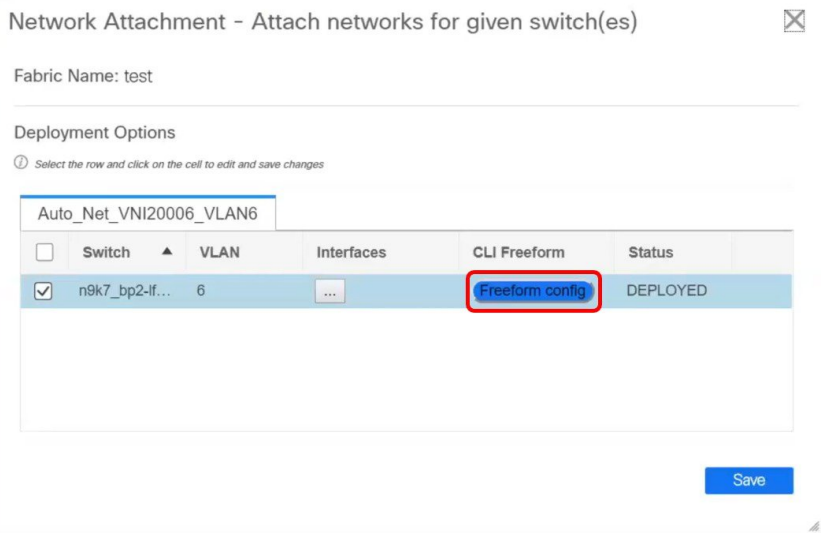

Step 7 Verify the VLAN name in the **Free Form Config** window.

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
.
.
.
```

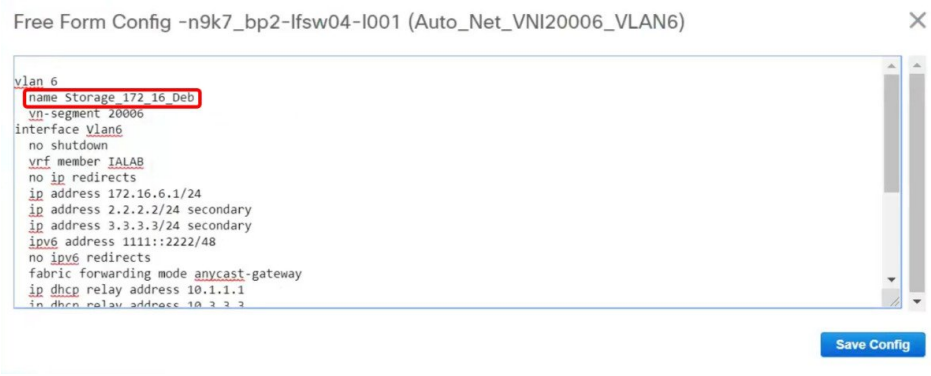

Step 9 Click **Save** in the **Network Attachment** window.

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:

- Uncheck all **Auto** IFC creation related fabric settings. Review the settings and ensure they are unchecked as follows:
  - Easy_Fabric_11_1 fabric

  - MSD_Fabric_11_1 fabric

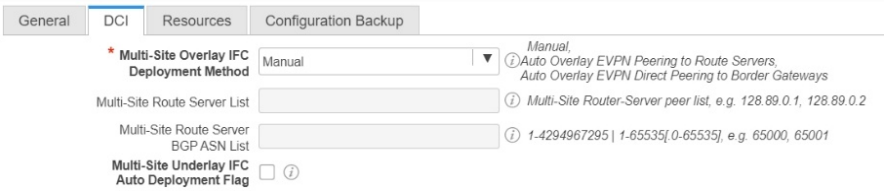

- Underlay Multisite peering: The eBGP peering and corresponding routed interfaces for underlay extensions between sites are captured in **switch_freeform** and **routed_inerfaces**, and optionally in the **interface_freeform** configs. This configuration includes all the global configs for multisite. 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. 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 g added successfully to the MSD. Errors will be reported if any mismatches are found. These must be fixed 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.

Note Additional interface configurations must be added to the Source/Destination interface freeform fields in th 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.

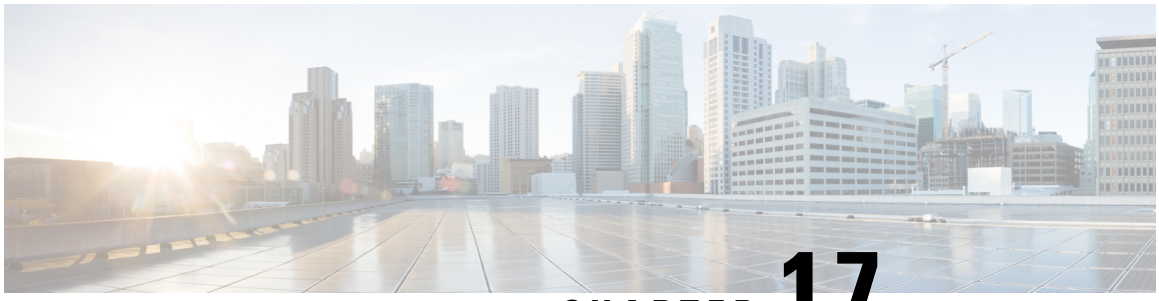

CHAPTER 17

Configuring a VXLANv6 Fabric

This chapter describes how to configure a VXLAN fabric with IPv6 underlay.

- Overview, on page 757
- Creating a VXLAN Fabric with IPv6 Underlay, on page 758

Overview

From Cisco DCNM Release 11.3(1), you can create an Easy fabric with IPv6 only underlay. The IPv6 underlay is supported only for the **Easy_Fabric_11_1** template. In the IPv6 underlay fabric, intra-fabric links, routing loopback, vPC peer link SVI, and NVE loopback interface for VTEP are configured with IPv6 addresses. EVPN BGP neighbor peering is also established using IPv6 addressing.

The following guidelines are applicable for IPv6 underlay:

- IPv6 underlay is supported for the Cisco Nexus 9000 Series switches with Cisco NX-OS Release 9.3(1) or higher.
- VXLANv6 is only supported Cisco Nexus 9332C, Cisco Nexus C9364C, and Cisco Nexus modules that end with EX, FX, FX2, FX3, or FXP.

Note VXLANv6 is defined as a VXLAN fabric with IPv6 underlay.

- In VXLANv6, the platforms supported on spine are all Nexus 9000 Series and Nexus 3000 Series platforms.
- The overlay routing protocol supported for the IPv6 fabric is BGP EVPN.
- vPC with physical multichassis EtherChannel trunk (MCT) feature is supported for the IPv6 underlay network in DCNM. The vPC peer keep-alive can be loopback or management with IPv4 or IPv6 address.
- 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.
- DHCPv6 is supported for the IPv6 underlay network.
- The following features are not supported for VXLAN IPv6 underlay:

- Multicast underlay
- Tenant Routed Multicast (TRM)
- ISIS, OSPF, and BGP authentication
- VXLAN Multi-Site
- Dual stack underlay
- vPC Fabric Peering
- DCI SR-MPLS or MPLS-LDP handoff
- BFD
- Super Spine switch roles
- NGOAM

Creating a VXLAN Fabric with IPv6 Underlay

This procedure shows how to create a VXLAN BGP EVPN fabric with IPv6 underlay. Only the fields for creating a VXLAN fabric with IPv6 underlay are documented. For information about the remaining fields, see Creating a New VXLAN BGP EVPN Fabric.

Procedure

Step 1 Navigate to **Control** > **Fabric Builder**.

Step 2 In the **Fabric Builder** window, click **Create Fabric**.

The **Add Fabric** window appears.

- **Fabric Name** - Enter the name of the fabric.
- **Fabric Template** - From the drop-down list, choose the **Easy_Fabric_11_1** fabric template.

Step 3 Enter the relevant values under the **General** tab.

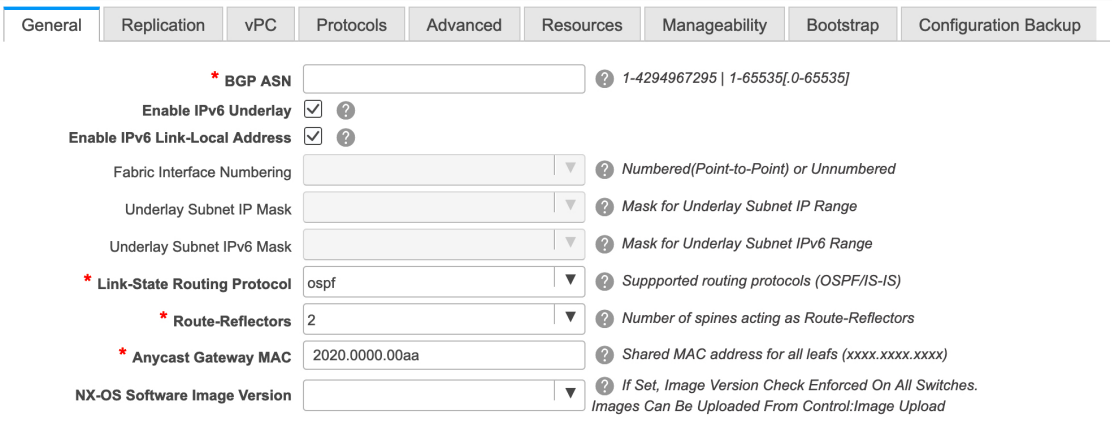

BGP ASN: Enter the BGP AS number for the fabric. You can enter either the 2 byte BGP ASN or 4 byte BGP ASN.

Enable IPv6 Underlay:Select this check box to enable the IPv6 underlay feature.

Enable Link-Local Address:Select this check box to use the link local addresses in the fabric between leaf-spine and spine-border interfaces. If you select this check box, the **Underlay Subnet IPv6 Mask** field is not editable. By default, the **Enable Link-Local Address** field is enabled.

IPv6 underlay supports only the **p2p** networks. Therefore, the **Fabric Interface Numbering** drop-down list field is disabled.

Underlay Subnet IPv6 Mask: Specifies the subnet mask for the fabric interface IPv6 addresses.

Link-State Routing Protocol: The IGP used in the fabric, that is, OSPFv3 or IS-IS for VXLANv6.

Step 4 Click the **Replication** tab.

IPv6 underlay supports only the ingress replication mode.

All the fields under this tab are disabled.

Step 5 Click the **vPC** tab.

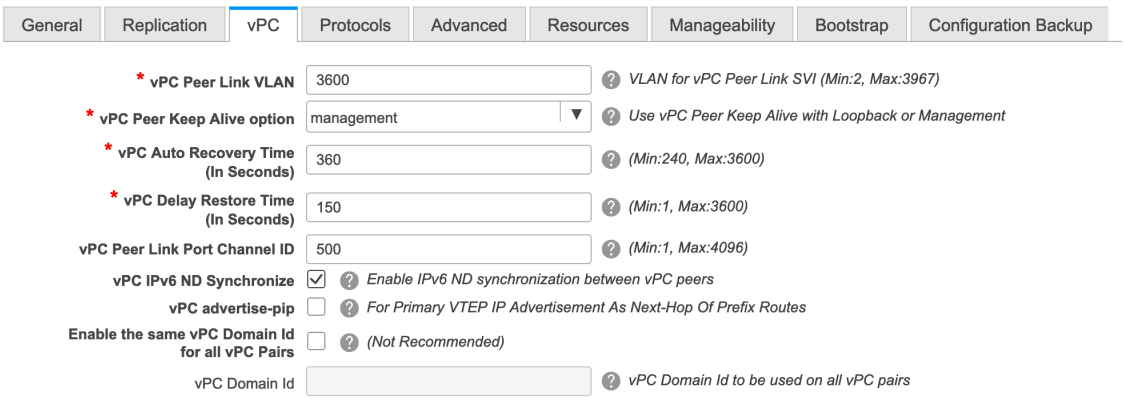

vPC Peer Keep Alive option – Choose the management or loopback option. If you want to use IP addresses assigned to the management port and the management VRF, choose management. If you use IP addresses assigned to loopback interfaces (and a non-management VRF), choose loopback. Both the options are supported for IPv6 underlay.

Step 6 Click the **Protocols** tab.

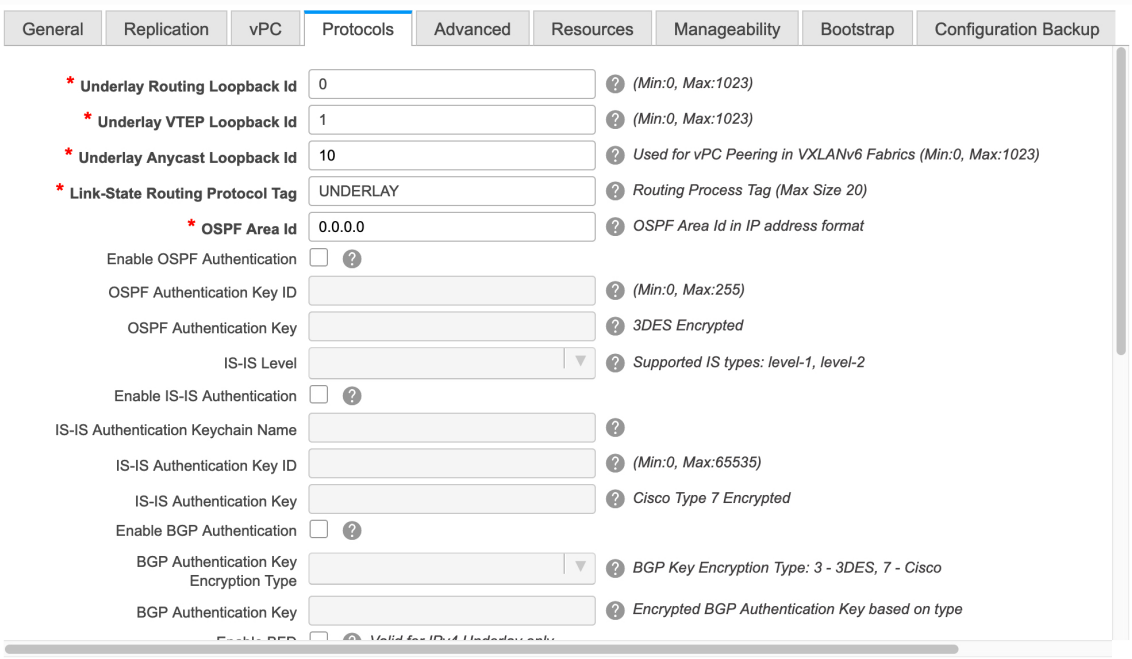

Underlay Anycast Loopback Id: Specifies the underlay anycast loopback ID for IPv6 underlay. Since an IPv6 address cannot be configured as secondary, an additional loopback interface is allocated on each vPC device. Its IPv6 address will be used as the VIP.

Step 7 Click the **Resources** tab.

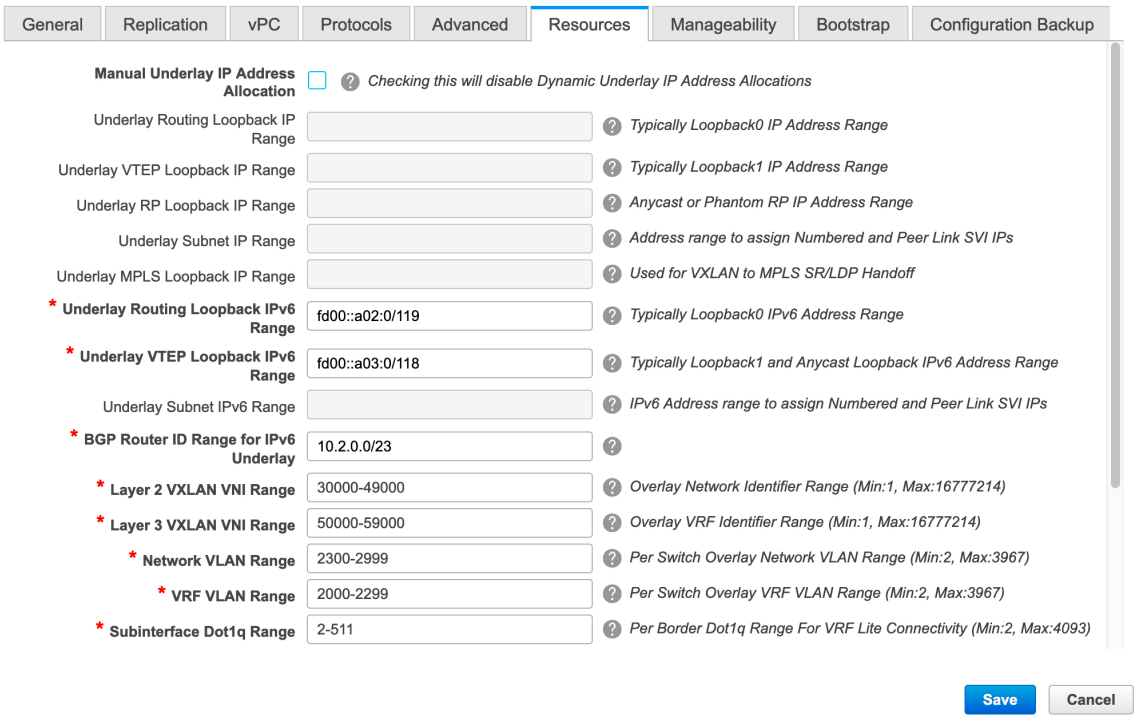

Manual Underlay IP Address Allocation: Select this check box to manually allocate underlay IP addresses. The dynamic underlay IP addresses fields are disabled.

Underlay Routing Loopback IPv6 Range: Specifies loopback IPv6 addresses for the protocol peering.

Underlay VTEP Loopback IPv6 Range: Specifies loopback IPv6 addresses for VTEPs. The IPv6 address for anycast will be assigned from this range.

Underlay Subnet IPv6 Range: Specifies the IPv6 address range that is used for assigning IP addresses for numbered and peer link SVIs. To edit this field, you need to unselect the **Enable Link-Local Address** check box under the **General** tab.

Underlay BGP Router ID Range: Specifies the address range to assign the BGP Router IDs.

Step 8 Click the **Bootstrap** tab.

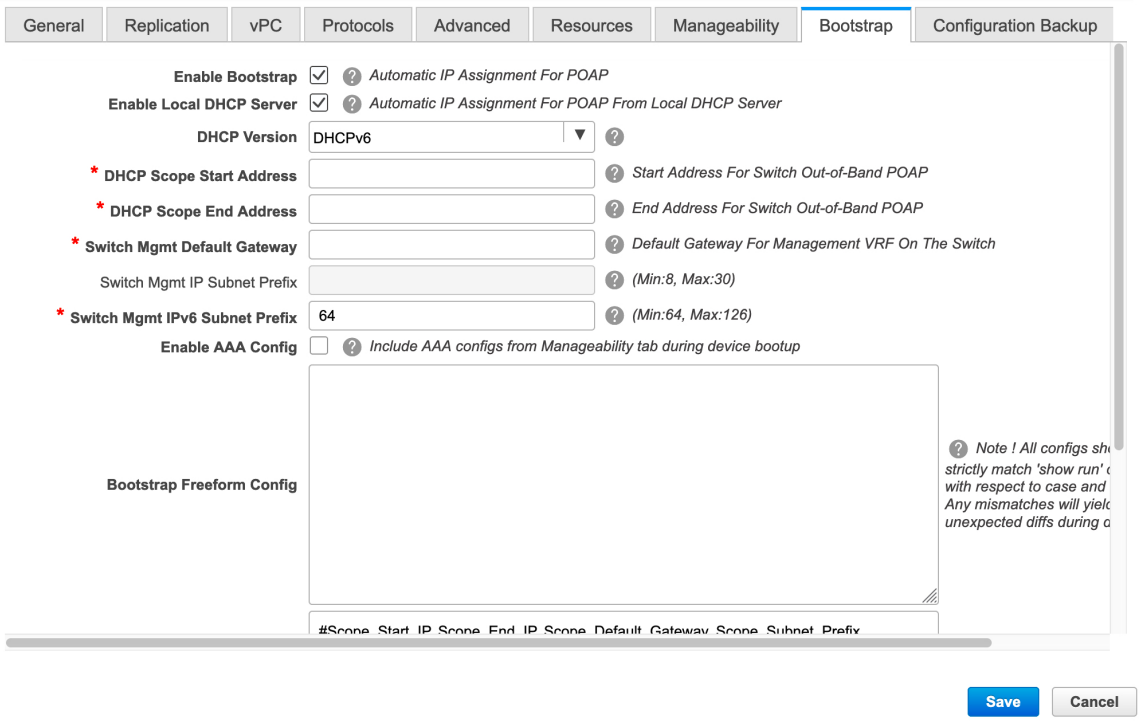

DHCP Version – Select DHCPv4 or DHCPv6 from this drop-down list. When you select DHCPv4, the **Switch Mgmt IPv6 Subnet Prefix** field is disabled. If you select DHCPv6, the **Switch Mgmt IP Subnet Prefix** is disabled.

Switch Mgmt IPv6 Subnet Prefix - Specifies the IPv6 prefix for the Mgmt0 interface on the switch. The prefix can be between 64 and 126. This field is editable if you enable IPv6 for DHCP.

For information about the remaining tabs and fields, see Creating a New VXLAN BGP EVPN Fabric.

What to do next

Adding Switches to a Fabric

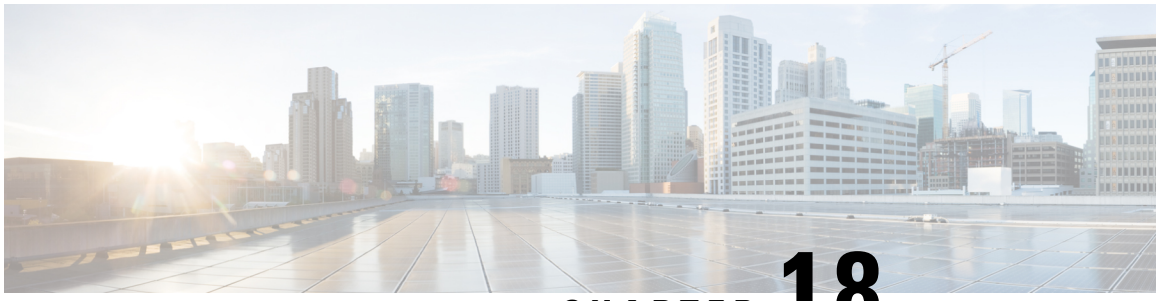

CHAPTER 18

Auto-Provisioning ToR Switches Attached to VXLAN VTEPs

This chapter describes how to configure the Top-of-Rack (ToR) switches and deploy networks in DCNM.

- Overview, on page 763
- Supported Topologies for ToR Switches, on page 763
- Configuring ToR Switches, on page 769
- Deploying Networks on ToR Switches, on page 775

Overview

From Cisco DCNM 11.3(1), support for the Top-of-Rack (ToR) switches is added in Cisco DCNM. You can add the Layer 2 ToR switches in an external fabric, and they can be connected to the Leaf switches in the Easy Fabric. Typically, the Leaf and ToR devices are connected with back-to-back vPC connection. For more information, see Supported Topologies for ToR Switches.

You can also watch the video that shows how to configure the ToR switches and deploy networks on these switches using Cisco DCNM. See Configuring ToR Switches.

Supported Topologies for ToR Switches

The following topologies with ToR switches are supported in DCNM:

Note Cisco Nexus 7000 Series Switches do not support the **ToR** switch role in Cisco DCNM.

- ToR switches with back to back vPC connection to the leaf switches.

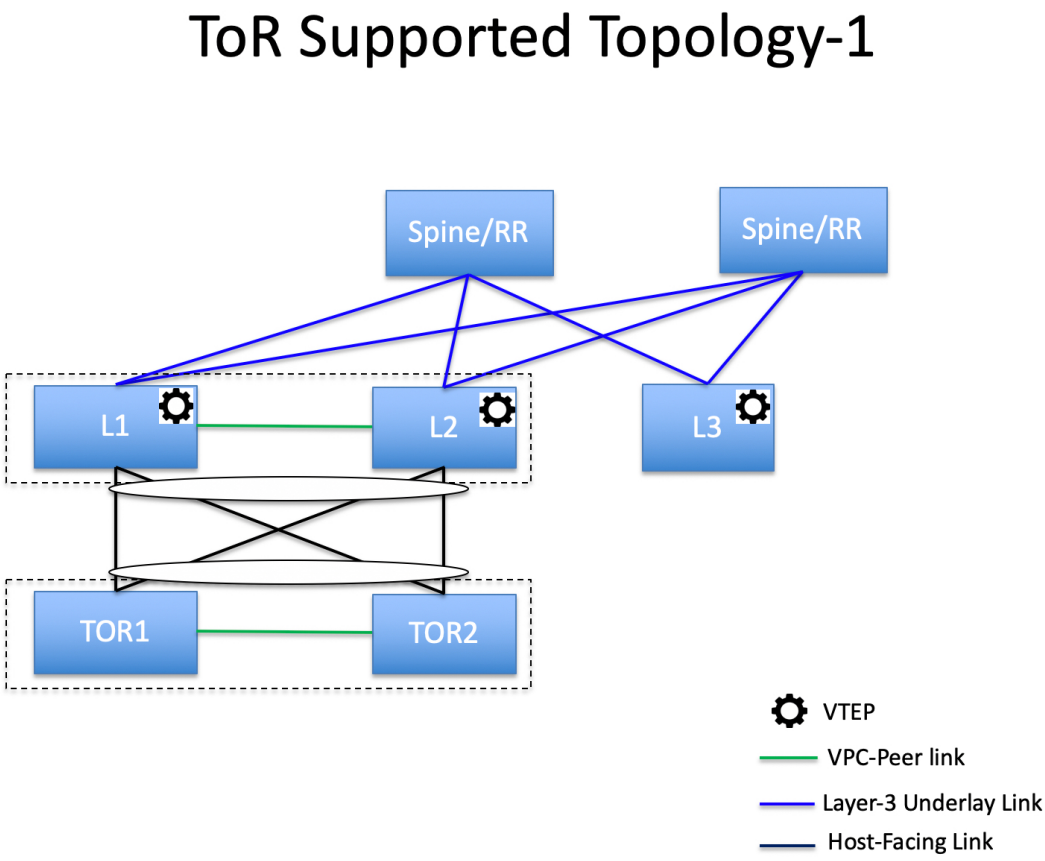

- ToR switches with port channels connected to both the leaf switches. The L1 and L2 switches are connected as a vPC pair.

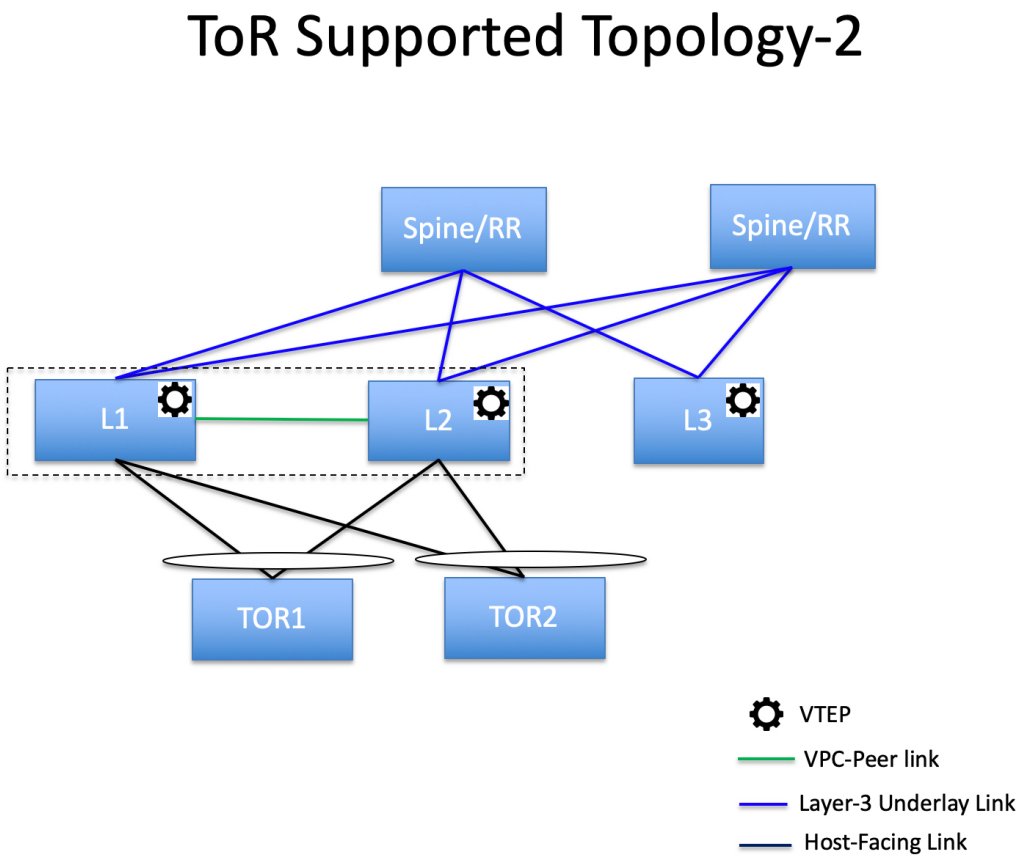

- ToR switches with port channels directly connected to the leaf switches. The L1 and L2 switches are connected as a vPC pair.

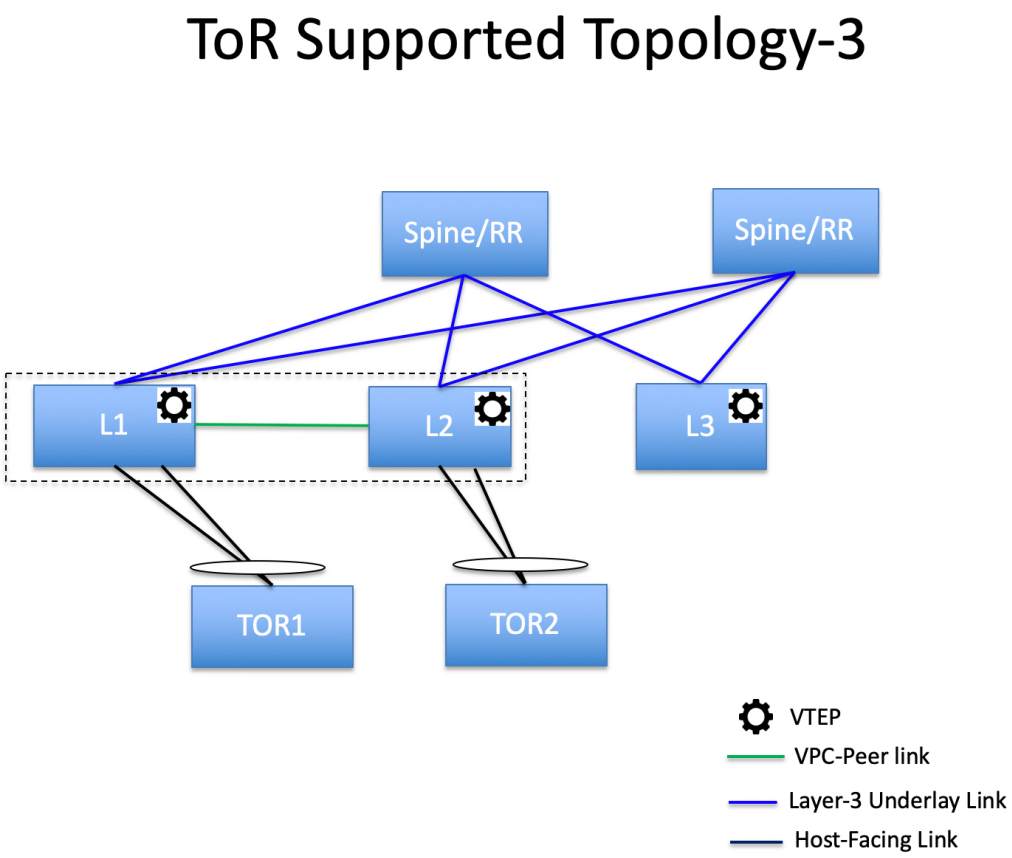

- ToR switches with port channels directly connected to the leaf switches. vPC pairs are not configured for the leaf or ToR switches.

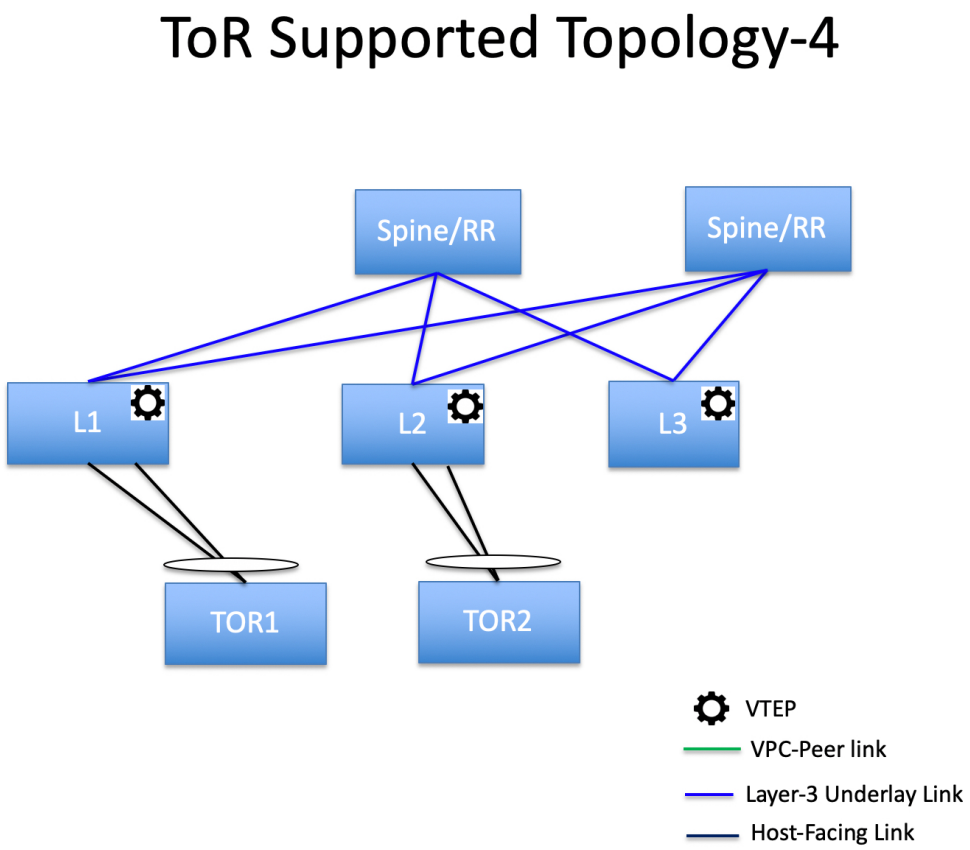

- ToR switches directly connected to the leaf switches. vPC pairs are not configured for the leaf or ToR switches.

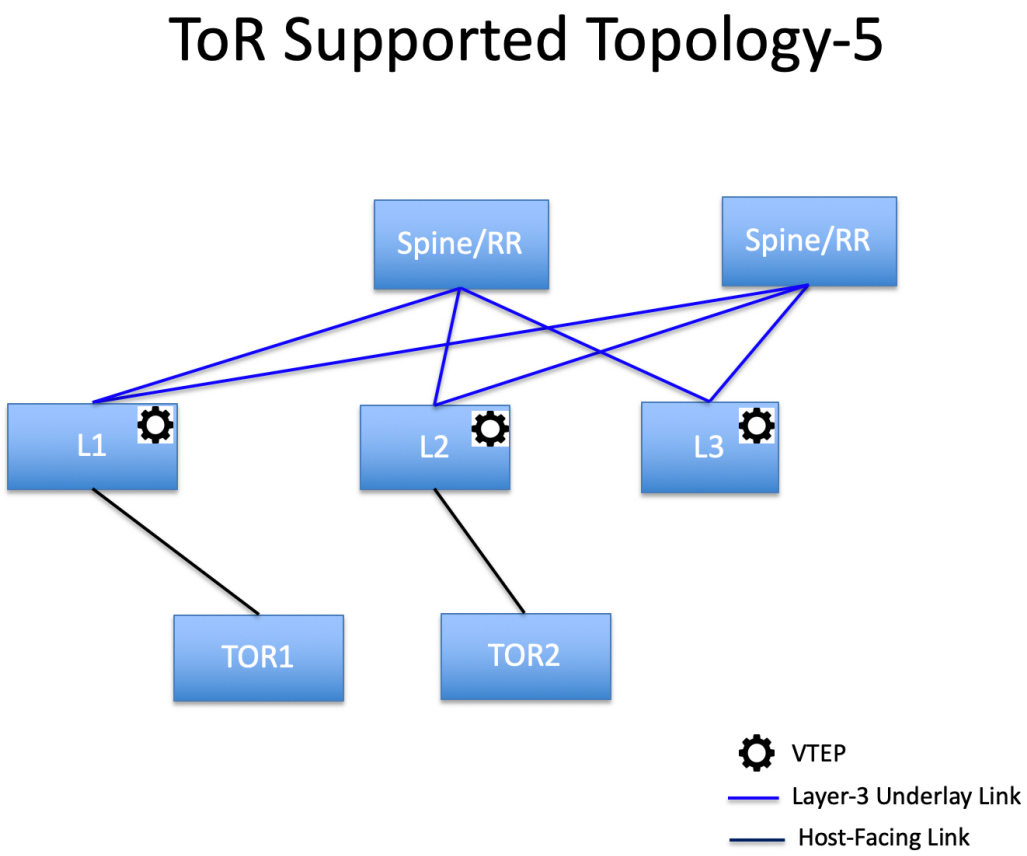

The following topology with ToR switches is not supported in DCNM:

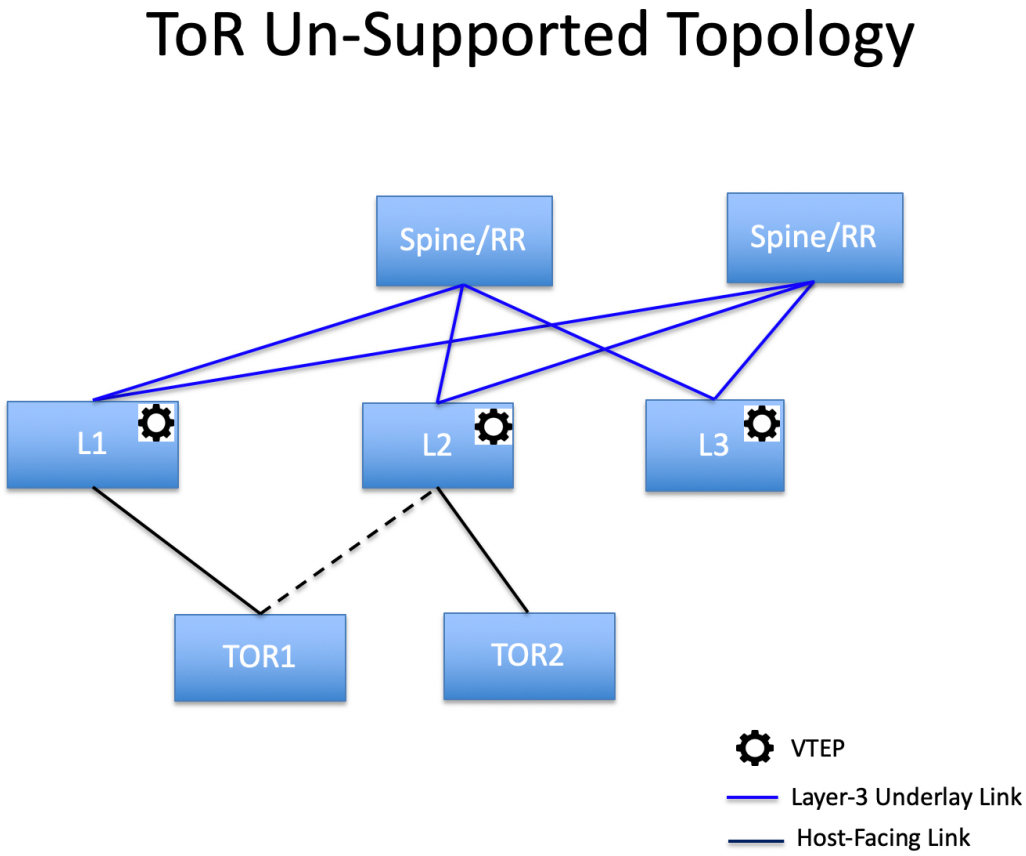

Configuring ToR Switches

Before you begin, make sure you have an Easy Fabric or create and deploy a new fabric. For more information, see *Creating a New VXLAN BGP EVPN Fabric* in *Cisco DCNM LAN Fabric Configuration Guide*.

Note DCNM supports the trunk_host policies for the ToR switches. Make sure ToR has vPC policies, port channel, and trunk host. These policies are used to connect the ToR switches in the external fabric to the Leaf switches in the Easy Fabric.

Procedure

Step 1 Create an external fabric and add two ToR switches. For more information, see *Creating an External Fabric* in *Cisco DCNM LAN Fabric Configuration Guide*.

The number of ToR switches can be more than two. This procedure shows how to configure ToR switches as shown in the ToR Topology-1, where ToR switches are connected using vPC. The following are the different scenarios for connecting the ToR switches:

- If vPC is not configured on the ToR switches, then vPC policies need to be applied on ToR facing interfaces if uplinks of these ToR switches are connected to vPC leaf switches.
- If ToR switches are connected to leaf using port-channel, then port-channel policies need to be applied on the ToR interfaces connected to the leaf switches.
- If ToR switches are connected to leaf switches as standalone, the trunk policies need to be applied on the TOR interfaces.

Note
- While creating the external fabric, make sure that the **Fabric Monitor Mode** check box is not selected.
- The two ToR switches must be connected and have same switch role.

After adding the ToR switches, make sure that the role for the ToR switches is selected as ToR.

Step 2 Right-click a ToR switch and select **vPC Pairing**.

Select the second ToR switch as a vPC Peer.

Step 3 Under vPC Pair Template, enter all the relevant details for a vPC connection between both the ToR switches. For more information about fields and their descriptions, see *Creating a vPC Setup in the External Fabric* in *Cisco DCNM LAN Fabric Configuration Guide*.

Note The Step 2 and 3 are required since this example shows the ToR configuration for Topology-1. For Topology 2, 3, 4, and 5, the steps 2 and 3 are not required.

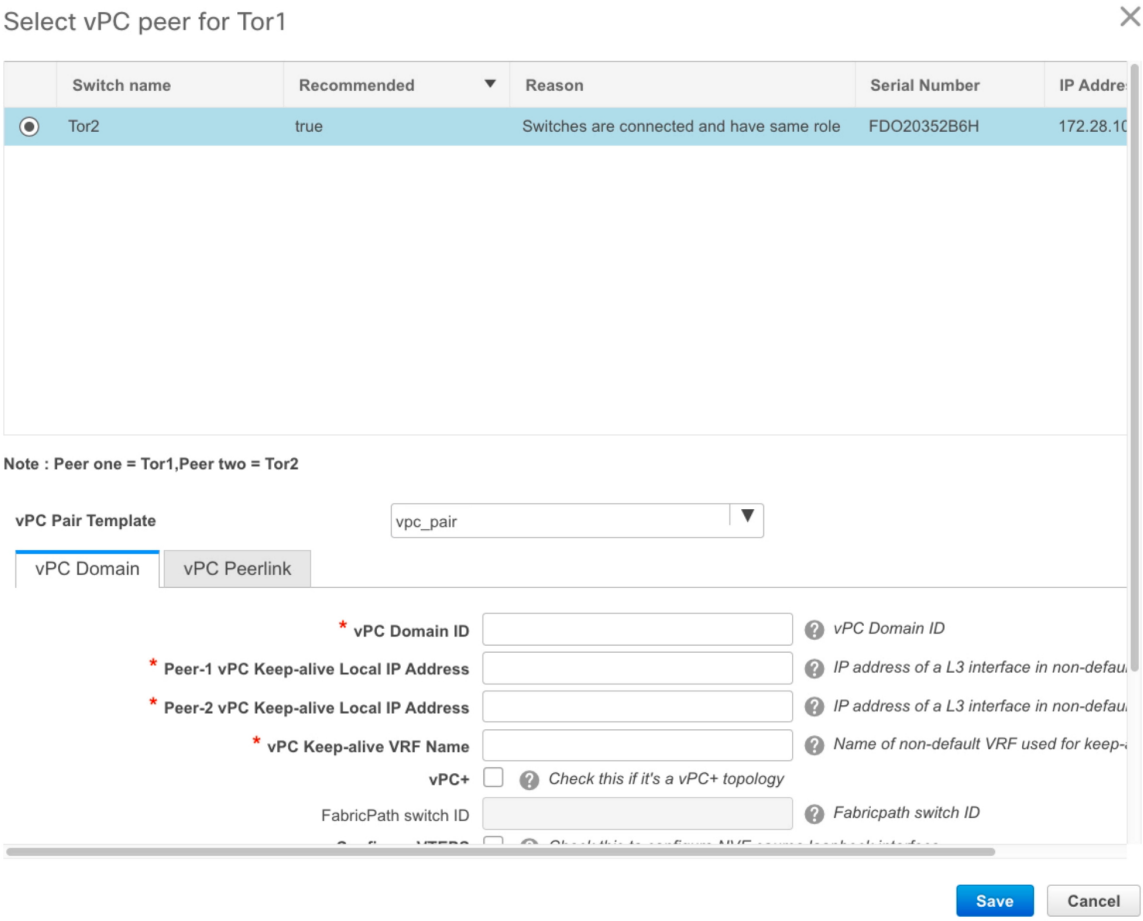

Step 4 Click **Save & Deploy**, and then click **Deploy Config**.

Step 5 After the progress bar shows 100% in the **Config Deployment** window, click **Close**.

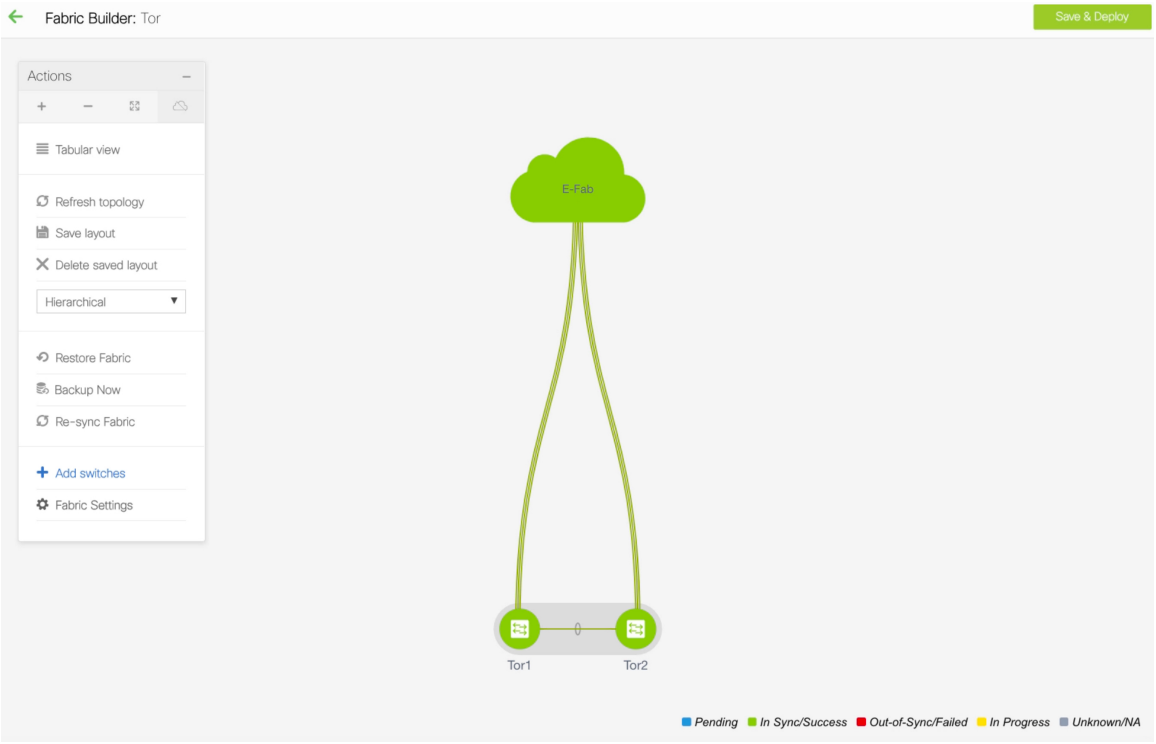

Step 6 Create an MSD fabric.

While creating the MSD fabric, under the General tab, select the ToR Auto-deploy Flag check box. This action enables automatic deployment of the networks and VRFs in the Easy Fabric to the ToR switches in the External Fabric when you click Save & Deploy in the MSD fabric. For more information, see Deploying Networks on ToR Switches.

For information about the remaining tabs and fields, see *Creating an MSD Fabric* in *Cisco DCNM LAN Fabric Configuration Guide*.

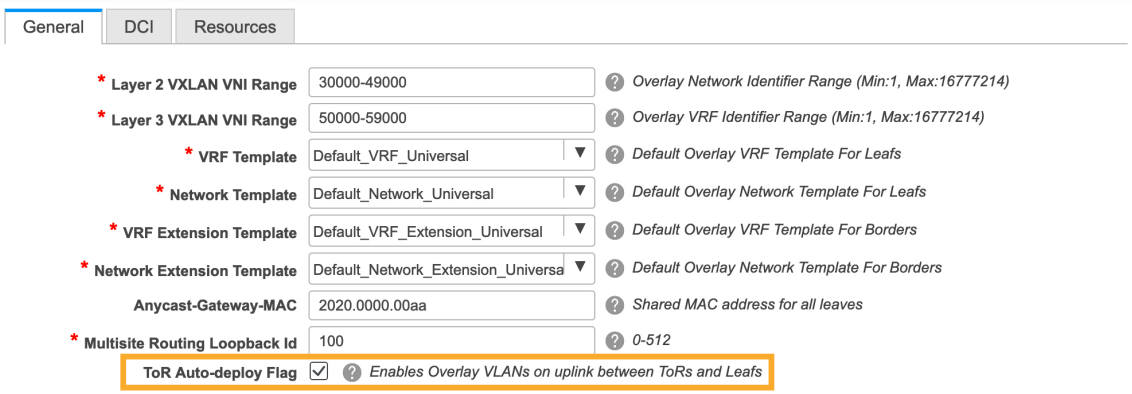

Step 7 Click **Move Fabric** in the **Action** panel. In the **Move Fabric** window, select the Easy Fabric and click **Add**.

Similarly, move the external fabric that contains the ToR switches to the MSD fabric.

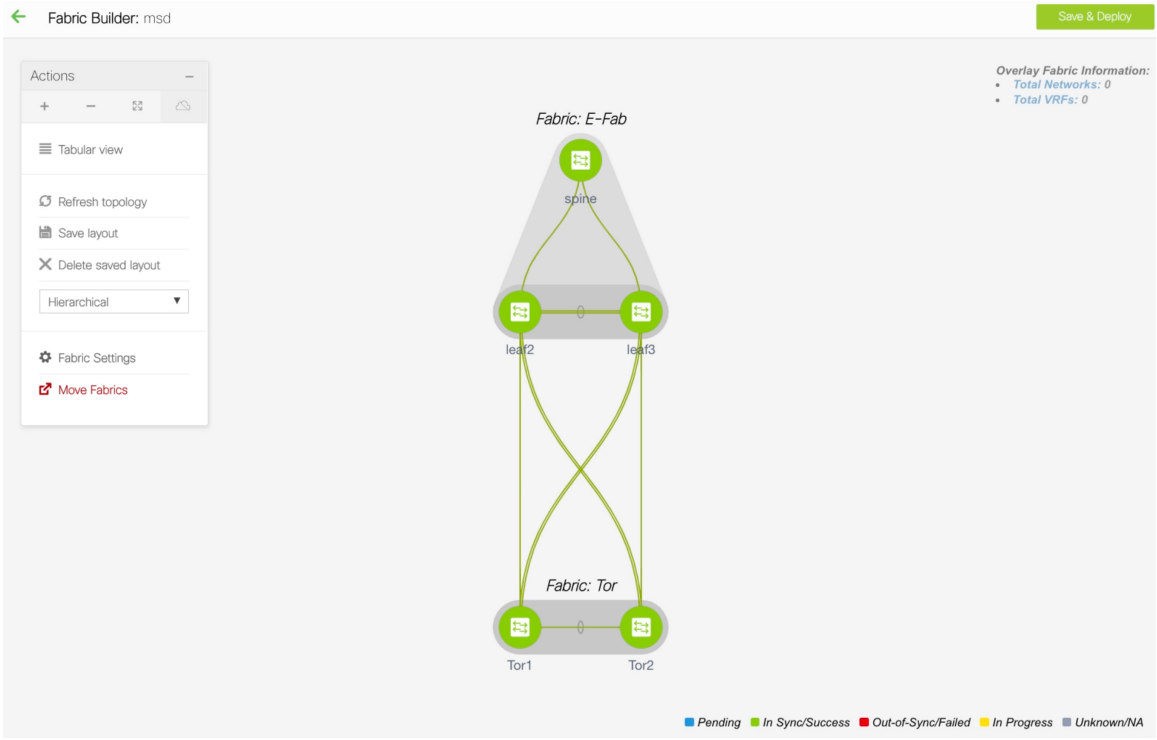

Step 8 Click the **Back** icon and click the Easy fabric containing the leaf switches.

Step 9 You need to create a vPC between the leaf and ToR switches. Right-click a leaf switch and select **Manage Interfaces**.

Step 10 In the **Manage Interfaces** window, click the **Add** icon to create a vPC.

Enter all the relevant details in the **Add Interface** window and click **Save**.

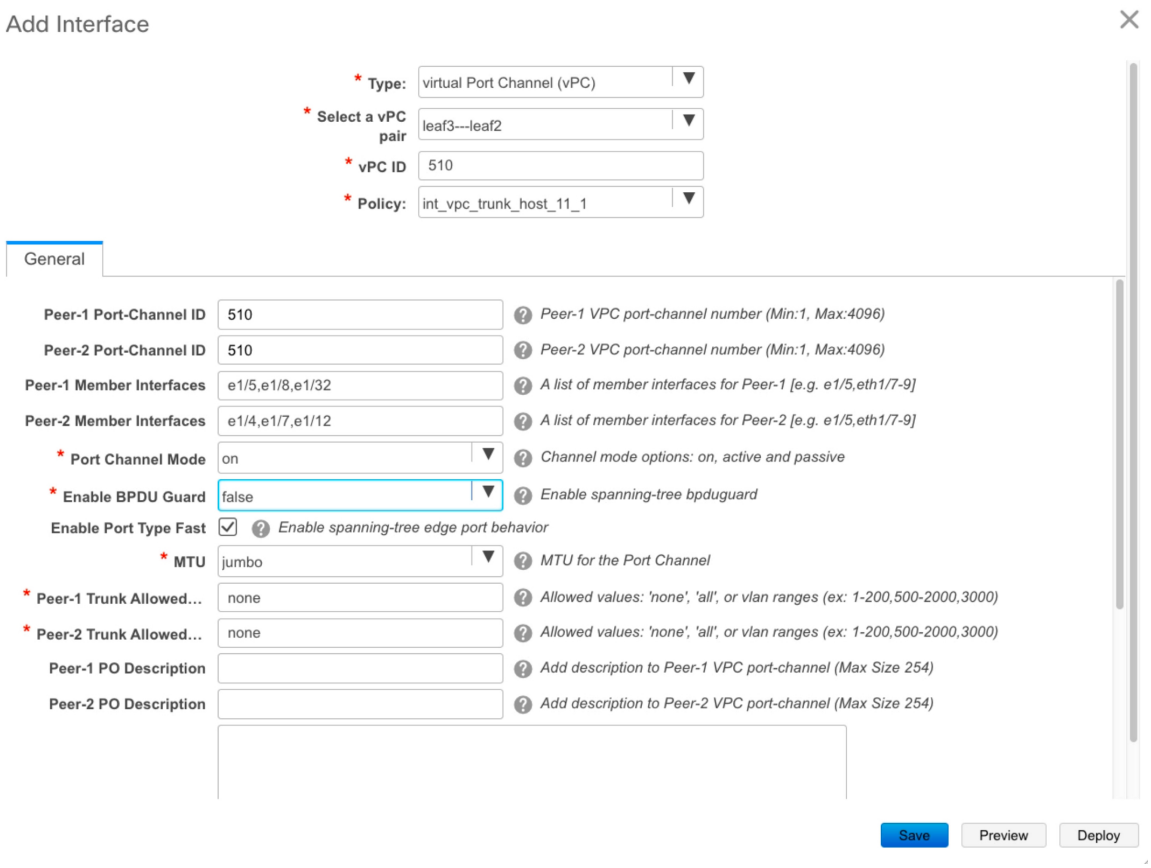

For more information about the fields in this window, see Adding Interfaces.

After saving all the information, click **Deploy**.

Similarly, follow the Steps 9 and 10 to create a vPC in the ToR switch as well.

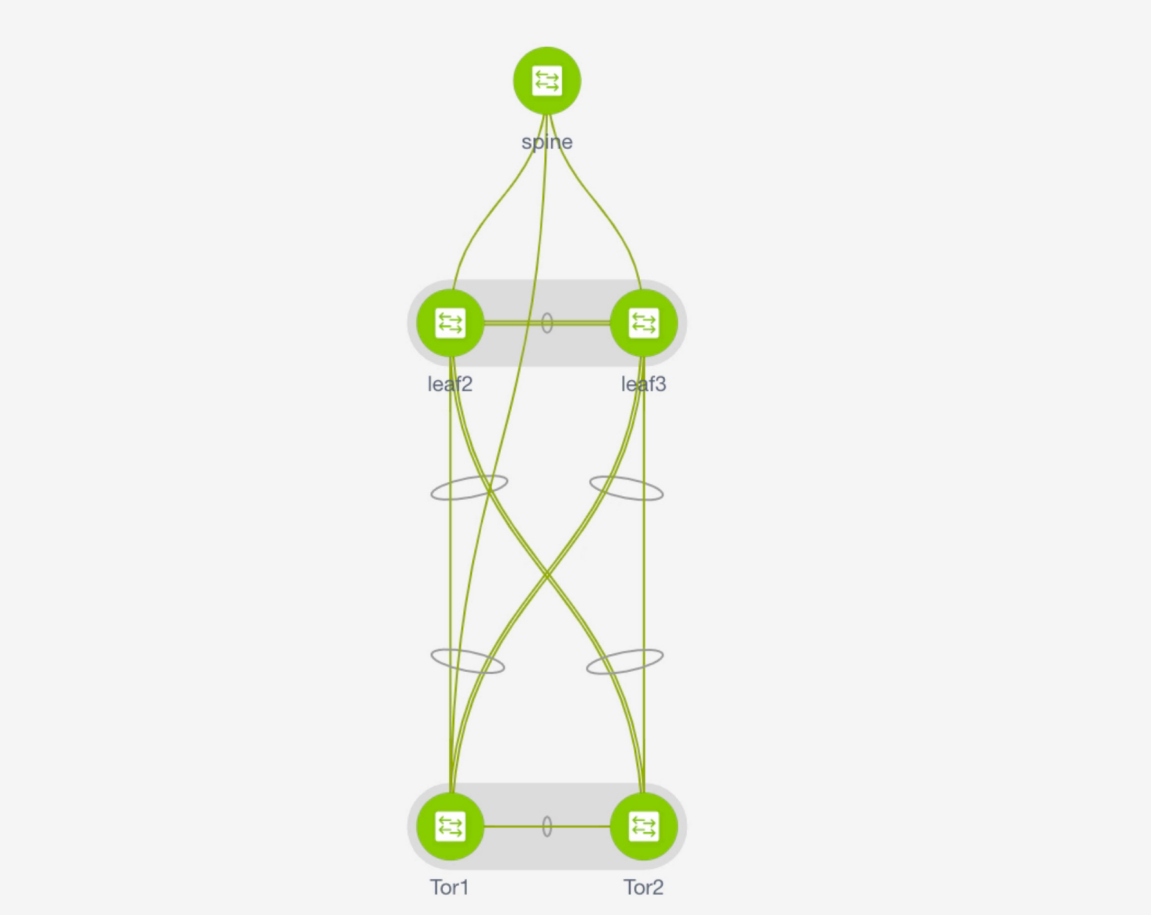

Deploying Networks on ToR Switches

To deploy networks on ToR switches in the external fabrics, you need to deploy them on the switches in the Easy Fabric through MSD. These switches should be connected to the ToR switches.

Procedure

Step 1 Navigate to **Control** > **Networks**.

Step 2 In the **Networks** window, from the **SCOPE** drop-down list, select the MSD fabric.

Step 3 Select the networks that you want to deploy or create a new network. For information about creating a network, see *Creating Networks for the Standalone Fabric* in the *Cisco DCNM LAN Fabric Configuration Guide*.

Click **Continue**.

Step 4 In the **Network Deployment** window, select the **Multi-select** check box and drag the cursor over the leaf switches in the Easy Fabric.

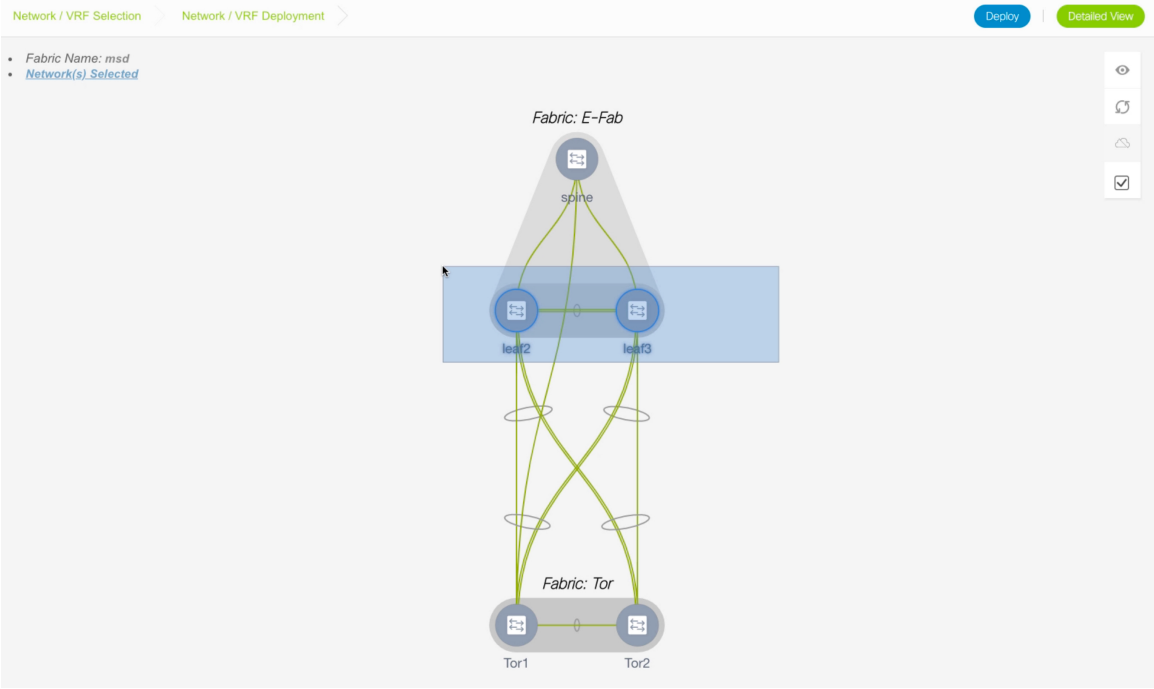

Step 5 In the **Network Attachment** window, click **...** in the **Interfaces** column.

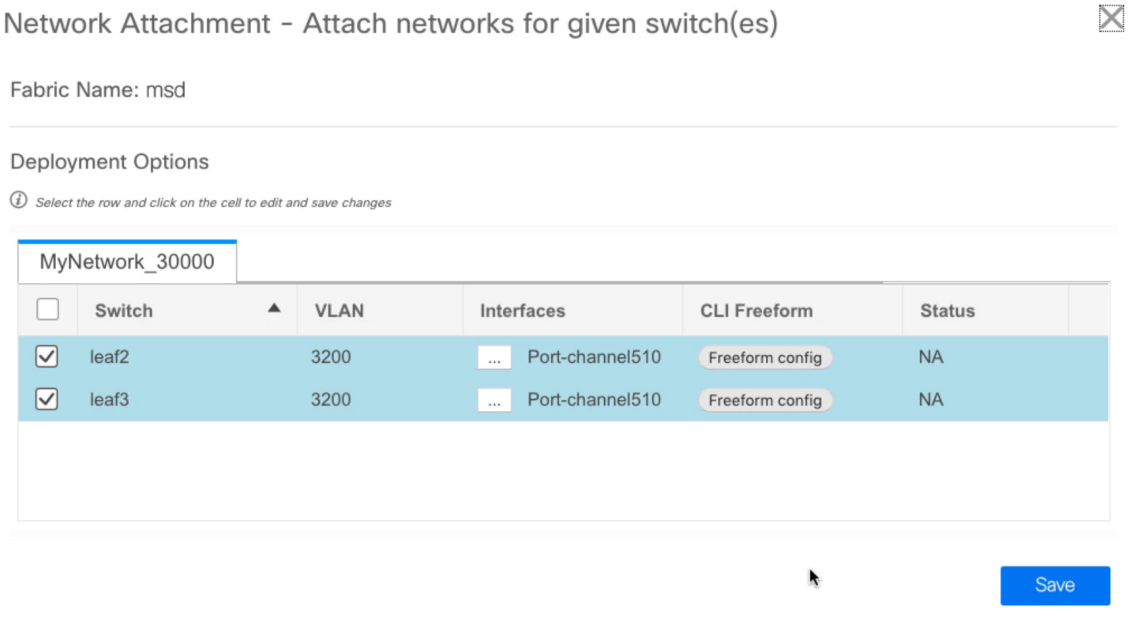

The **Interfaces** window lists interfaces or port channels. You can select interfaces/port channels to associate them with the selected network. These port channels connect the leaf switches to the ToR switches. The networks will be deployed on these port channels.

Click **Save** and close this window.

Step 6 Click **Deploy**.

Now the VLANs are deployed on the leaf switches.

Step 7 Navigate to **Control > Fabric Builder**.

Step 8 Click the MSD fabric and click **Save & Deploy**.

The networks created and deployed on the leaf switches in the Easy Fabric are also deployed on the ToR switches in the external fabric. This step allows the same VLANs to be configured on the ToR switches that are deployed on the leaf switches in the Step 6.

Note If VLANs are created on the ToR switches manually using the freeform configs, they are not modified.

PART III

External/WAN Layer 3 Connectivity for VXLAN BGP EVPN Fabrics

- VRF Lite, on page 781
- MPLS SR and LDP Handoff, on page 813

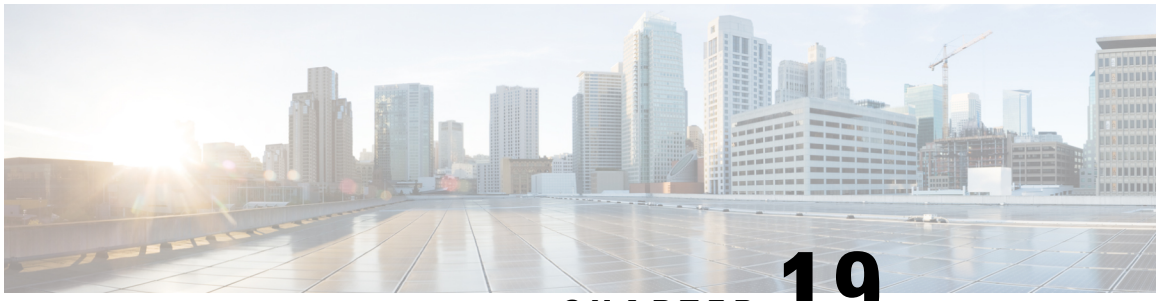

CHAPTER 19

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 DCNM, 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 and Guidelines, on page 781
- Sample Scenarios, on page 784
- VRF Lite Through the DCNM GUI – From a BGW Device to a Nexus 7000 Series Edge Router , on page 785
- VRF Lite Through the DCNM GUI – From a BGW Device To a Non-Nexus Device , on page 797
- Automatic VRF Lite (IFC) Configuration, on page 804
- Deleting VRF Lite IFCs, on page 808
- Additional References, on page 810
- Appendix , on page 810

Prerequisites and Guidelines

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**.

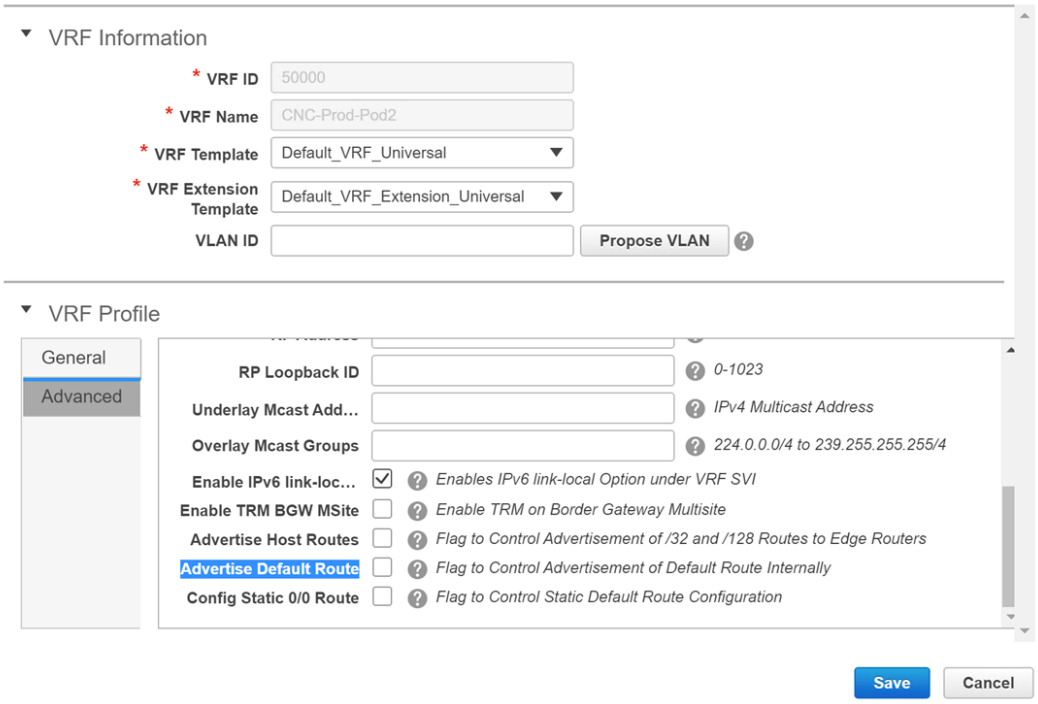

The following options apply only when VRF Lite connectivity is enabled on the border devices. By default, following Cisco best practices, DCNM uses eBGP over sub-interfaces for VRF Lite, Option-A peering. In other words, for each VRF Lite Inter-fabric connection (IFC), there is a per VRF per peer eBGP 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.

Guideline

In a DCNM Release 10.4(2) setup where VRF-Lite IFCs are created, the required default prefix-lists or route-maps configs are added on the switch. When this DCNM Release 10.4(2) setup is upgraded to any of the DCNM 11.x releases, VRF-Lite related RPM configs might be saved as part of the switch_freeform policy.

The following route-map config is part of this switch_freeform:

```
route-map EXTCON-RMAP-FILTER-V6 deny 20
match ip address prefix-list host-route-v6
```

When this setup is upgraded from DCNM Release 11.x to 11.3(1), the route-map config is corrected with the following config:

```
route-map EXTCON-RMAP-FILTER-V6 deny 20
match ipv6 address prefix-list host-route-v6
```

Since RPM configs are saved in DCNM 11.x as switch_freeform, you need to manually delete the ip prefix-list match config in the switch_freeformpolicy so that ipv6 match config is successful on the switch.

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

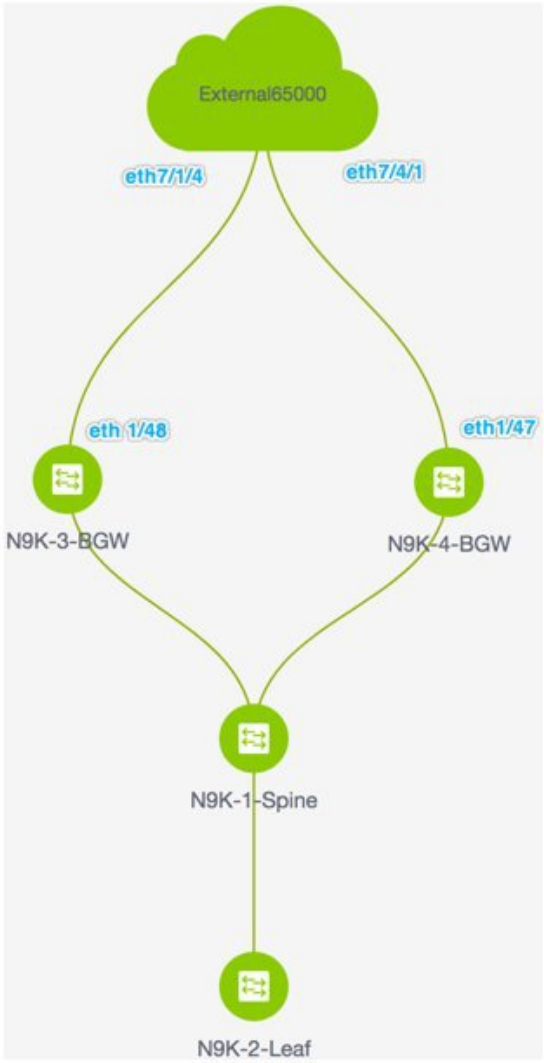

- The topology displays the VXLAN BGP EVPN fabric **Easy7200** connected to the external fabric **External65000** (the cloud icon). The BGWs of the VXLAN fabric are connected to the edge router **n7k1-Edge1** (not visible in the image) in the external fabric.
- The BGWs are special devices that allow 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. VRF Lite can be enabled on physical Ethernet interface or layer 3 port-channel. Sub-interface over physical interface or layer 3 port-channel interface is created by DCNM at the VRF extension moment for each VRF lite link the VRF is extended over.

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 a fabric overlay between two loopback interfaces). When you add a physical connection between devices, t 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.

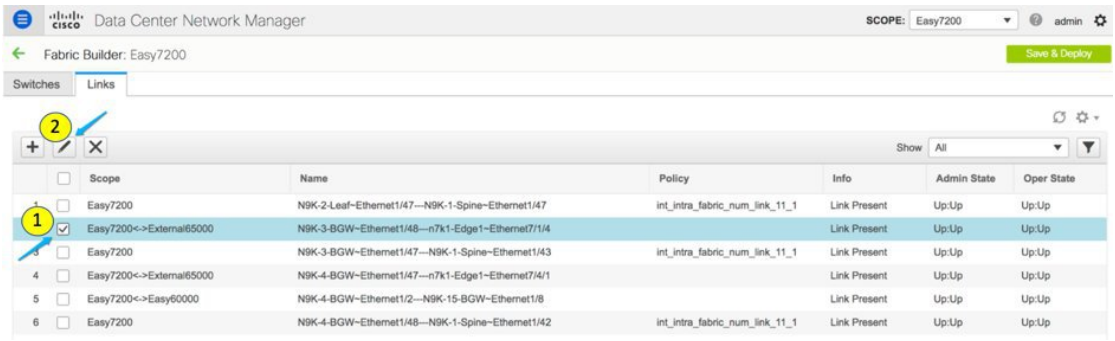

The fields are:

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.

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.

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.

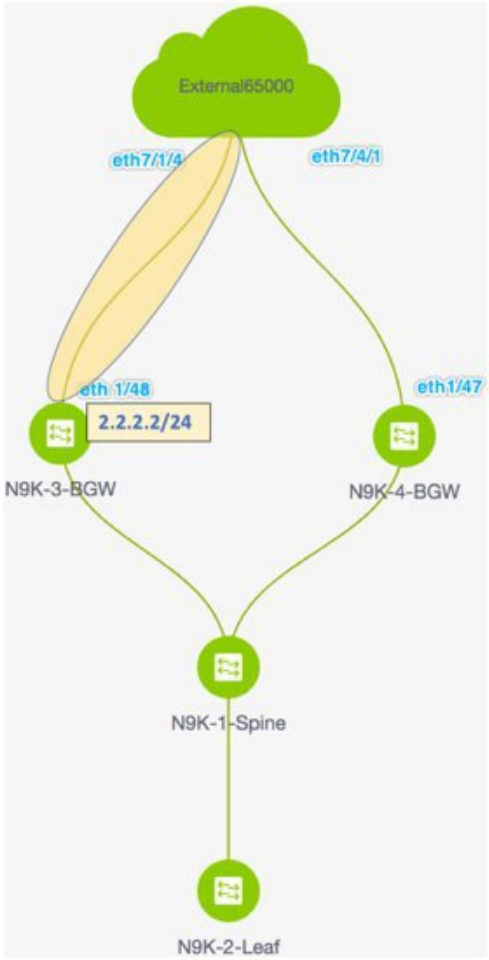

6. Similarly, create an IFC from **eth 1/47** on **N9K-4-BGW** towards **eth 7/4/1** on **n7k1-Edge1**. An entry is seen in the **Links** screen. A representation of the topology is shown below.

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.

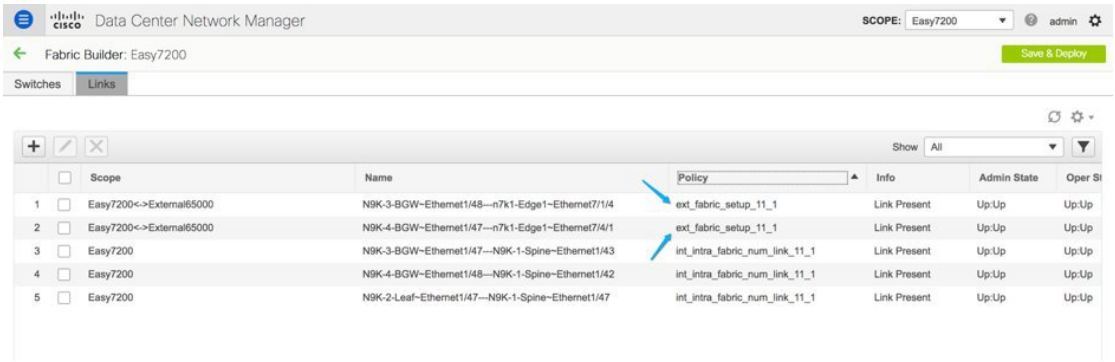

8. Go to the **Scope** drop down box at the top right part of the screen and choose **External65000**. The external fabric **Links** screen is displayed. You can see that the two IFCs created from **Easy7200** to **External65000** is displayed here.

Note When you create an IFC or edit its setting in the VXLAN fabric, the corresponding entry is automatica created in the connected external fabric.

9. Click **Save and Deploy** to save the IFCs creation on **External65000**.

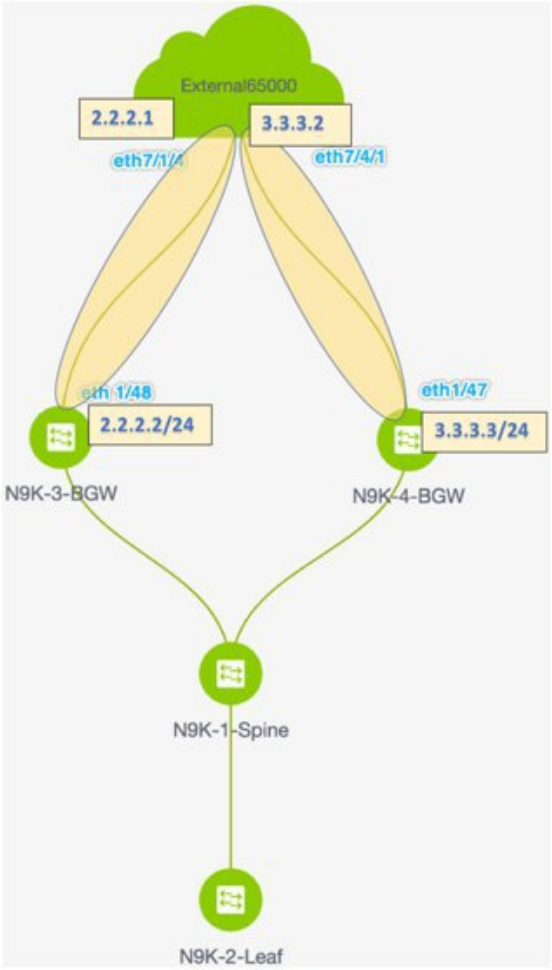

Base configurations – For VRF Lite to function, appropriate route maps and policies that apply to VRFs have to be deployed on the border devices **N9K-3-BGW** and **N9K-4-BGW**. You do not need to manually enable the base configurations. They are automatically deployed via a default template **ext_base_border_vrflite_11_1**.

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.

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.

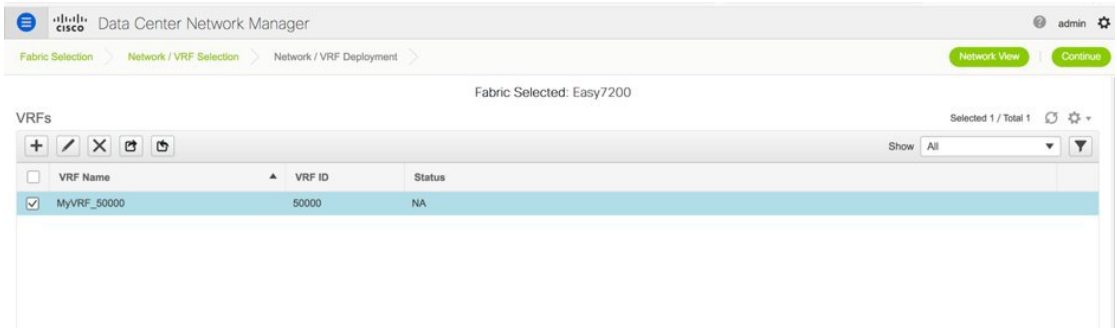

The **Easy7200** fabric topology comes up.

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.

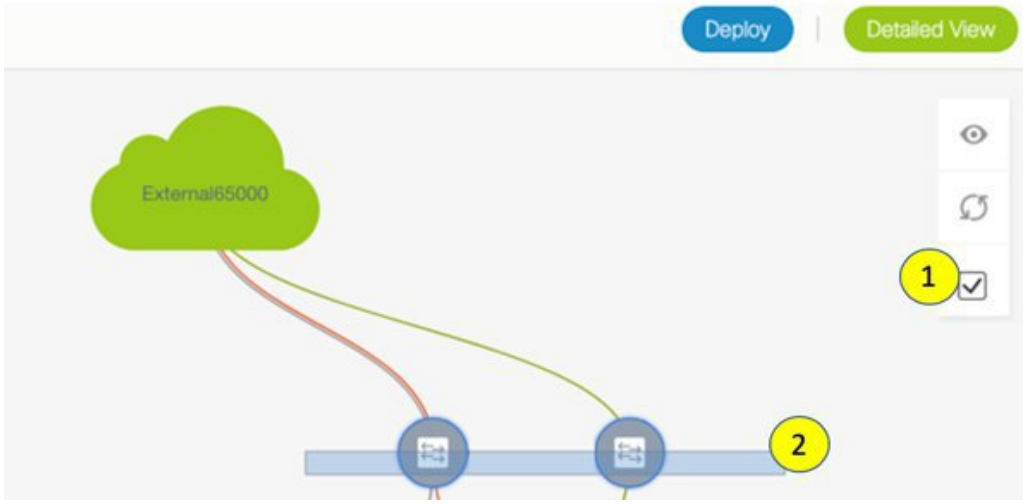

The **VRF Extension Attachment** screen comes up. Each row represents a switch and each tab a VRF. Update settings for each tab as explained.

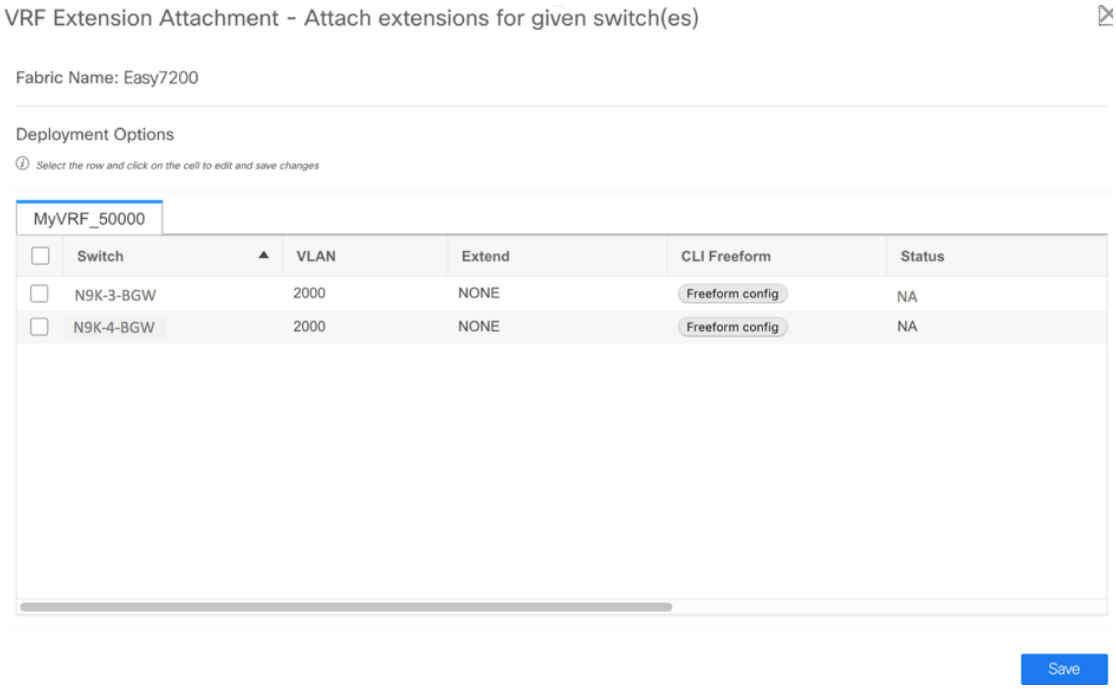

In the **Extend** column, click on **NONE** and choose the **VRF_LITE** option from the drop down box. Do this for the second row too.

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.

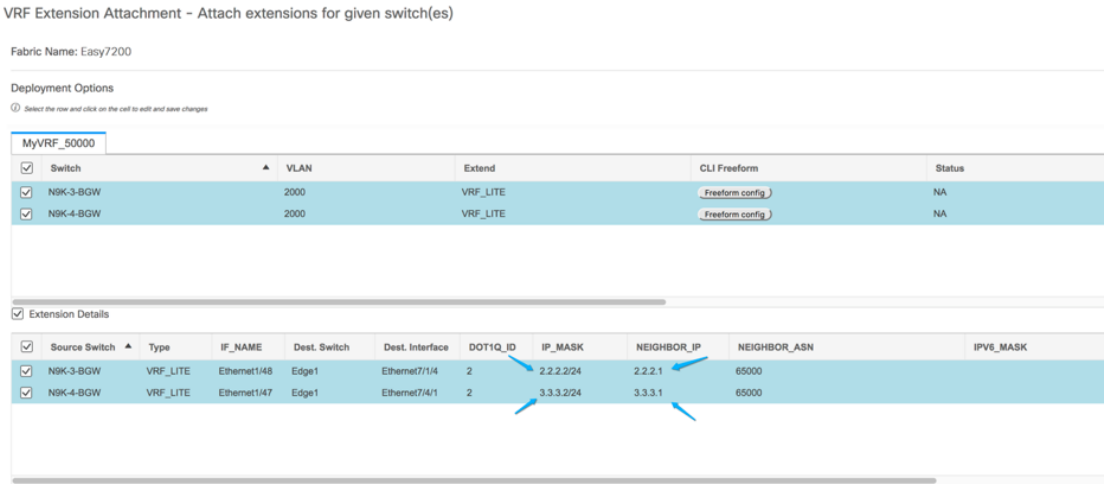

DCNM will create one sub-interface for each VRF-LITE link above using the values in DOT1Q_IP, IP_MASK and NEIGHBOR_IP fields. The IP_MASK and NEIGHBOR_IP fields for each VRF LITE extension are filled with the **IP Address/Mask** and **BGP Neighbor IP** values entered in VRF LITE link creation. The IP_MASK and NEIGHBOR_IP fields, along with the DOT1Q_ID field, can be overwritten.

IPV6_MASK and NEIGHBOR_IPV6 fields can be optionally entered if IPv6 eBGP session over the sub-interface is needed.

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.

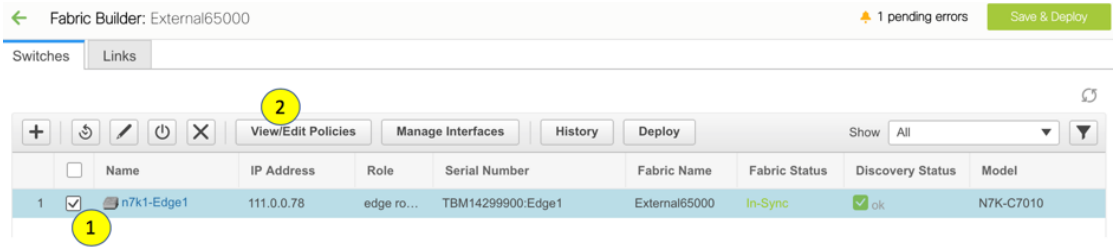

The **View/Edit Policies** screen comes up.

3. 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, wh applicable.

This defines a policy from the edge router towards **N9K-3-BGW**.

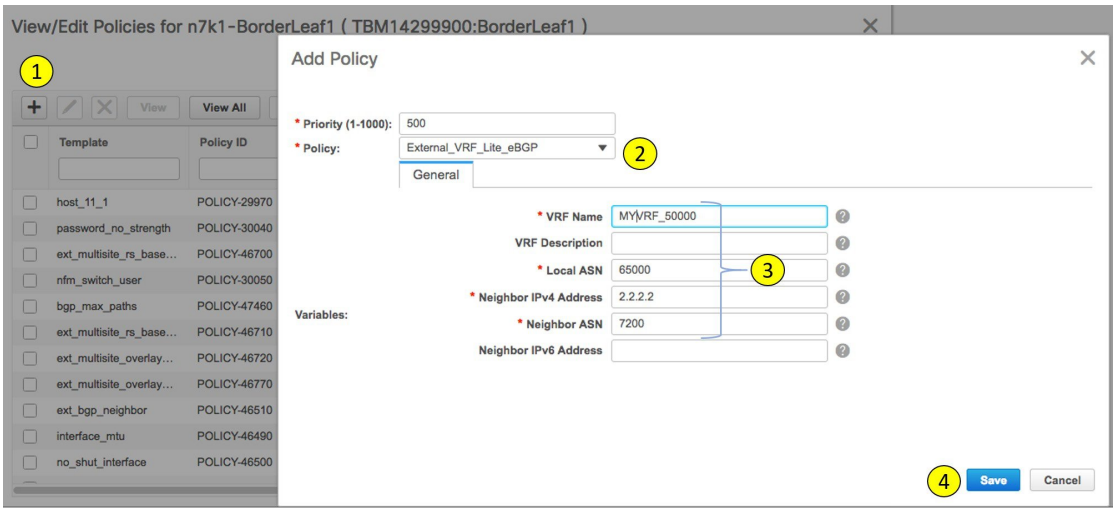

4. As per the earlier steps, create a policy for the VRF extension towards **N9K-4-BGW**. The **Neighbor IPv4 Address** field for the second extension is updated with 3.3.3.3.

Sub interface policy on Edge Router

1. In the **External65000** fabric topology screen, click **Tabular view**.

   The **Switches | Links** screen comes up.

2. Select the switch checkbox and click the **Manage Interfaces** button.

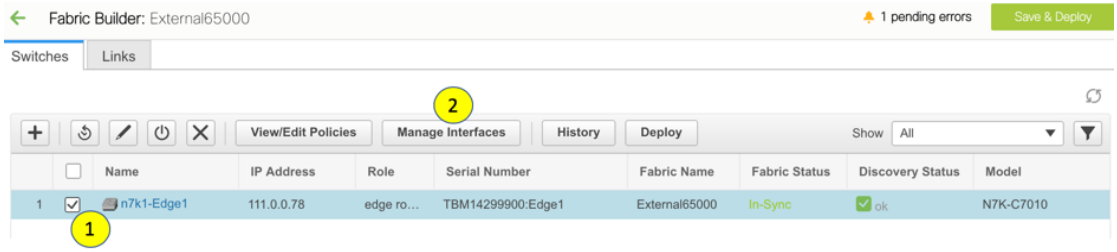

The **Manage Interfaces** screen comes up.

3. 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.

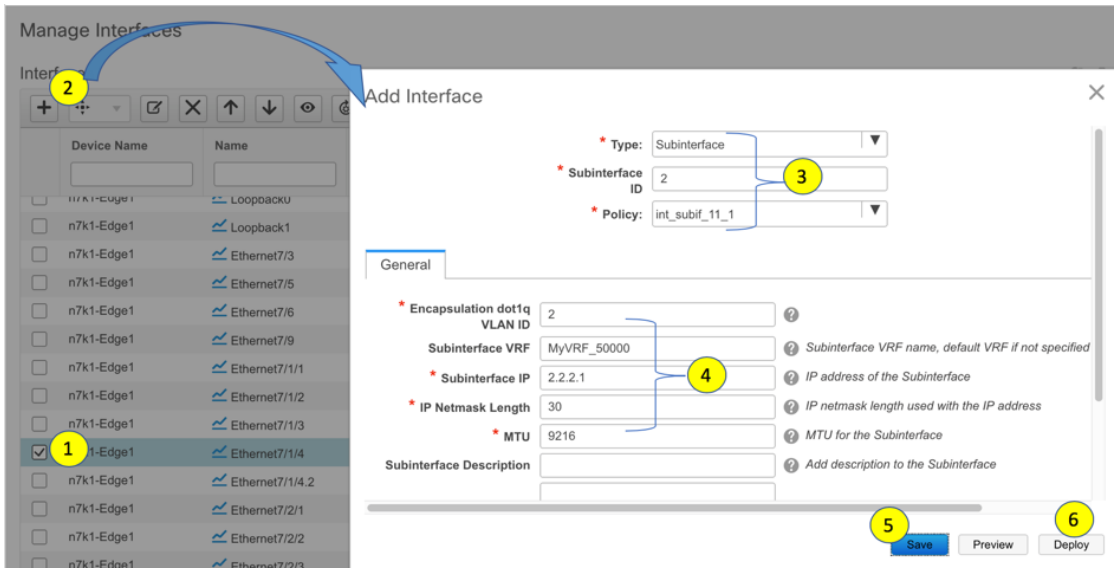

The example shows a break out port on the Cisco Nexus 7000 Series switch. This breakout must be performed using the DCNM breakout policy (the template name is **breakout_interface**). If this is not done, the subinterface deletion is blocked by DCNM.

4. Click **Save** to save the settings, and **Deploy** to deploy the settings onto the switch.

5. 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:

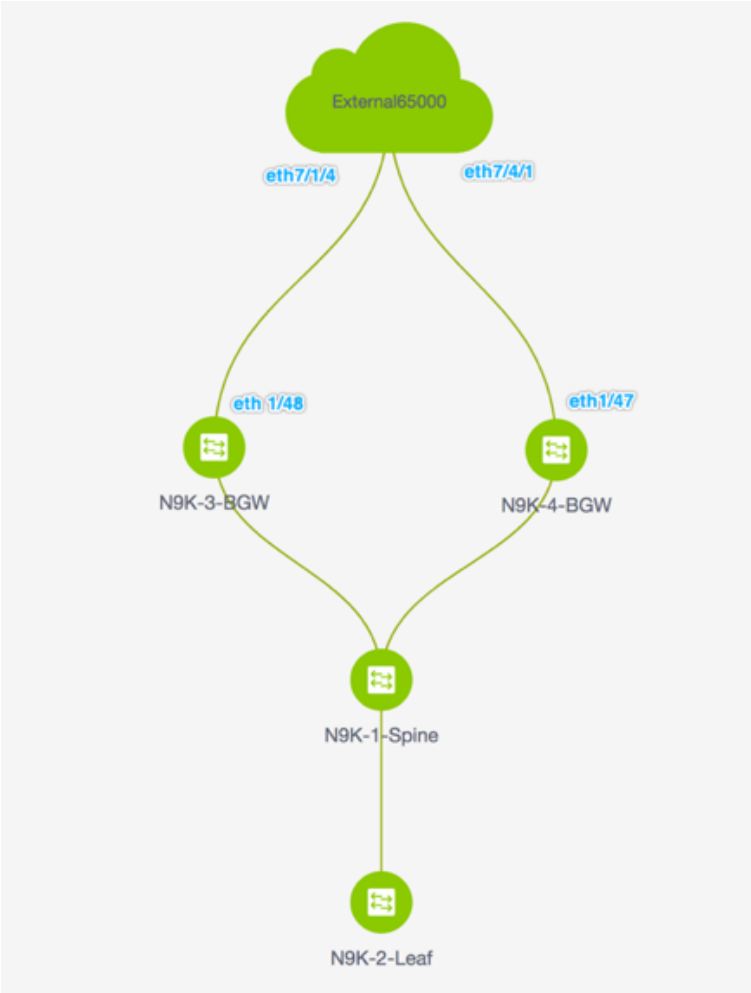

The steps are:

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.

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** cha 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:

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.

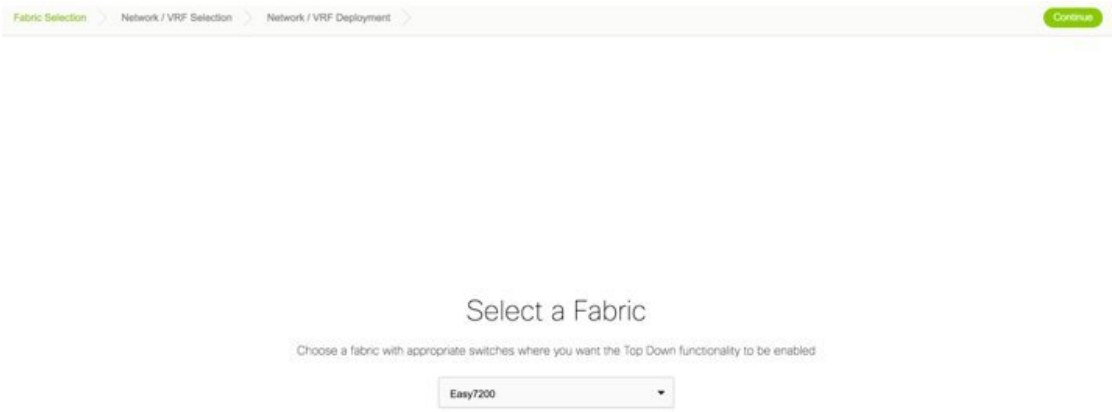

   The **Networks** screen comes up.
4. Click **VRFs** at the top right part of the screen. The **VRFs** screen comes up.
5. Select the VRF that you want to deploy (**MyVRF_5000** in this case) and click **Continue** at the top right part of the screen.

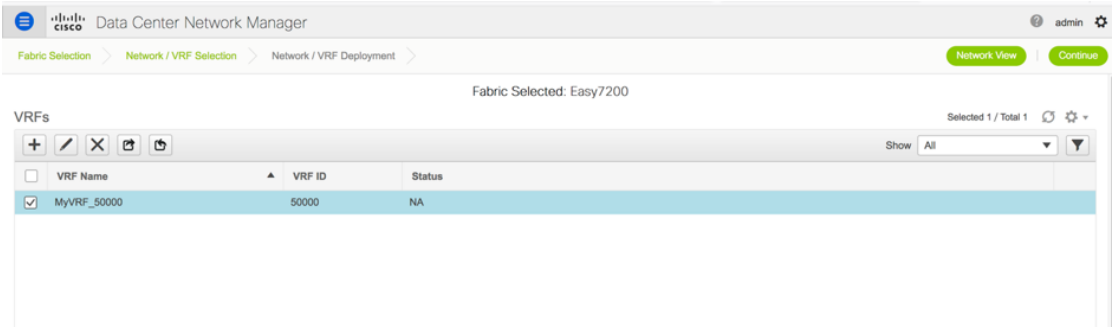

   The Easy7200 fabric topology comes up.
6. Double-click the **N9K-3-BGW** icon on which you want to deploy the VRF and VRF extension configuration.

   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.

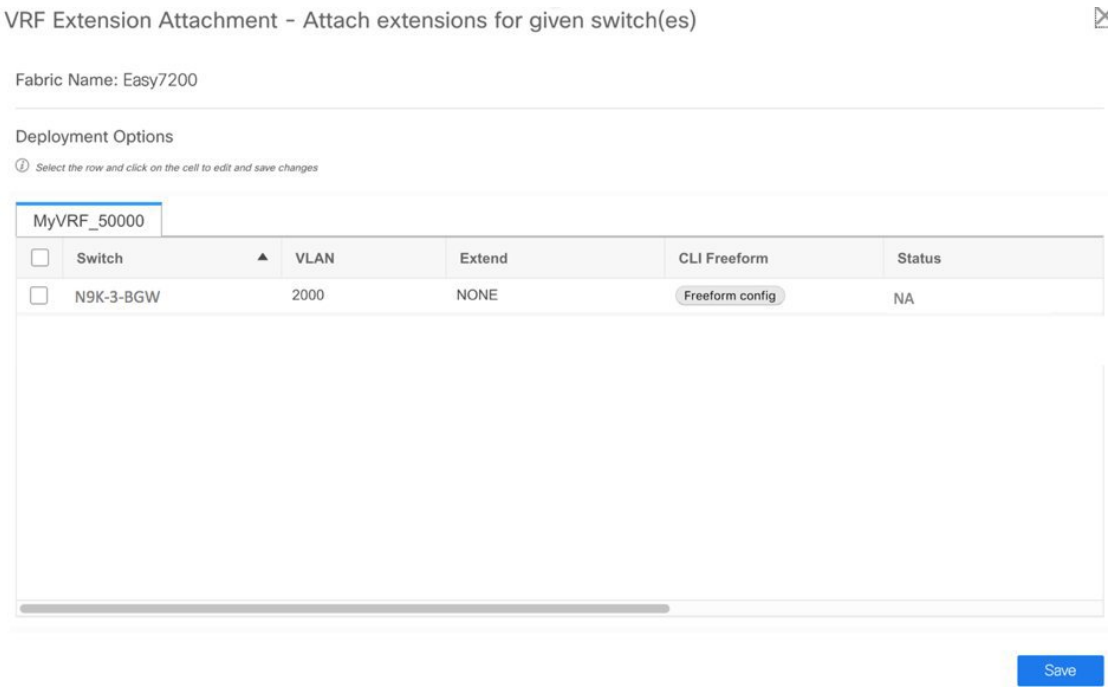

In the **Extend** column, click on **NONE**. A drop down box appears. Choose the **VRF_LITE** option, and click outside the row.

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.

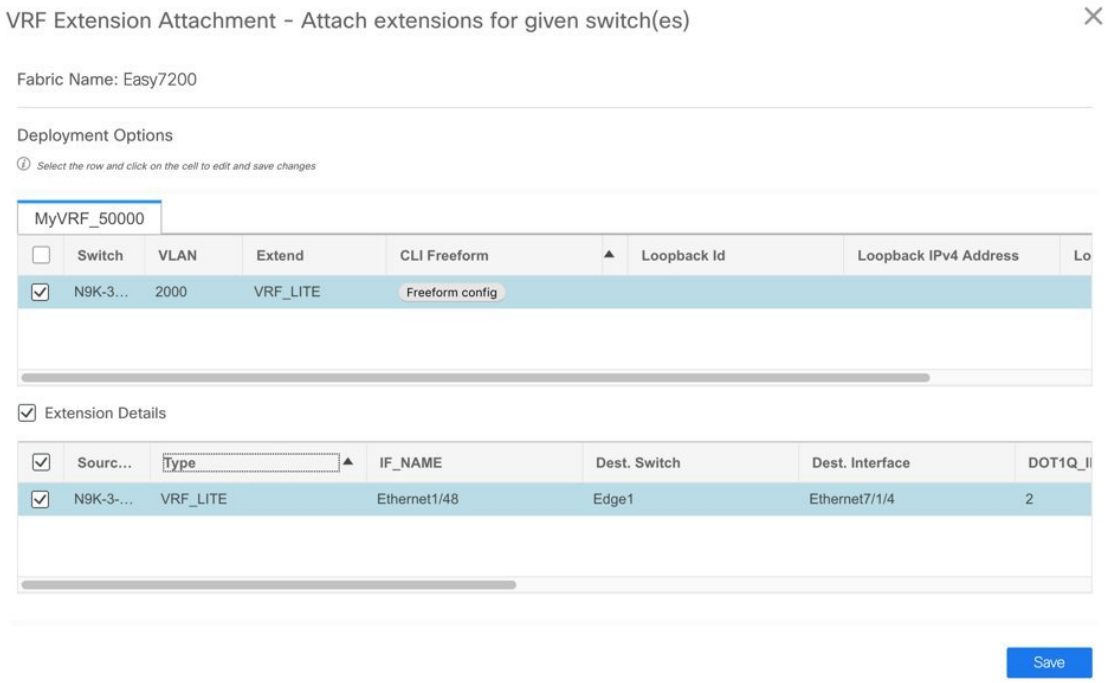

Click **Save** at the bottom right part of the screen.

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.

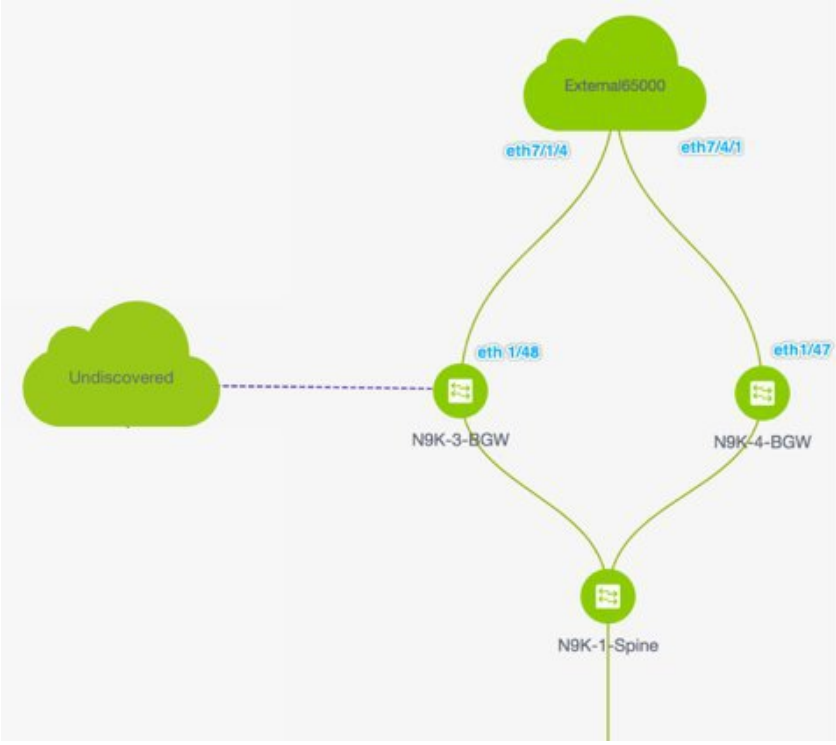

Automatic VRF Lite (IFC) Configuration

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.

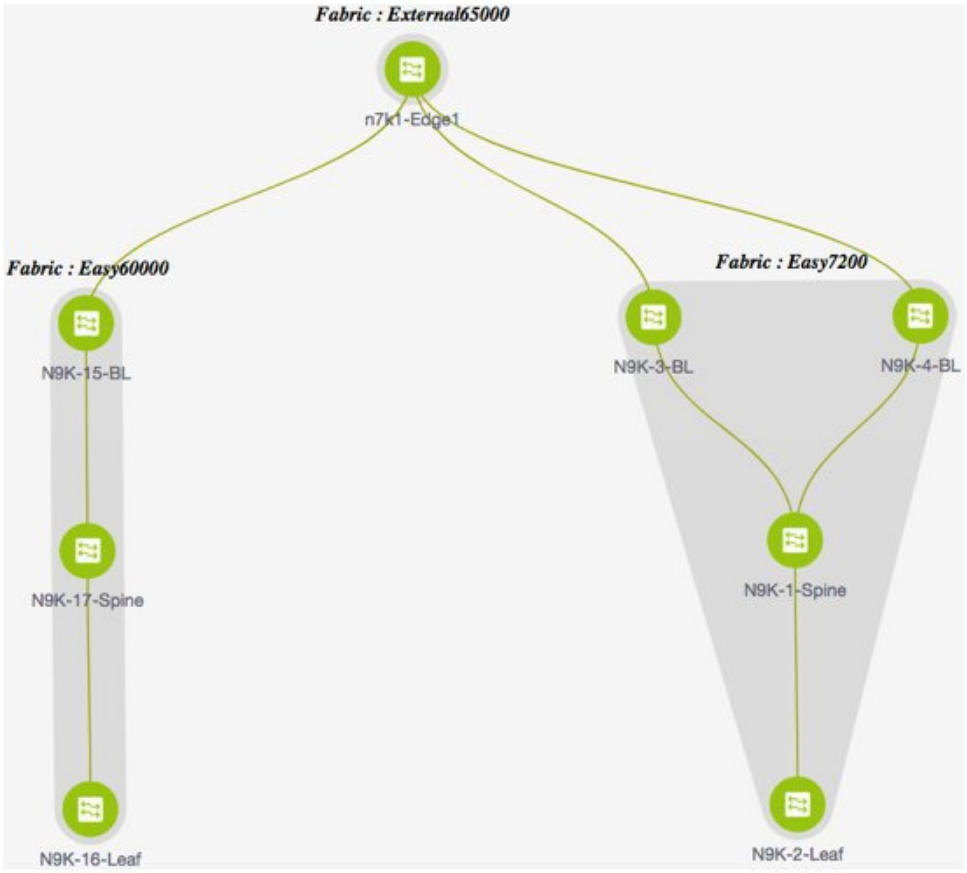

- The topology displays VXLAN BGP EVPN fabrics **Easy60000** (at the left) and **Easy7200** (at the right) and external fabric **External65000** (at the top). The border leaf of one VXLAN fabric is connected to the border leaf of the other through the edge router **n7k1-Edge1** in the external fabric.
- 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

- The Auto IFC is supported for the Cisco Nexus devices only.
- From Cisco DCNM Release 11.4(1), you can configure a Cisco ASR 1000 Series routers and Cisco Catalyst 9000 Series switches as edge routers, set up a VRF-lite IFC, and connect it as a border device with an easy fabric.

- If the device in the External fabric is non-Nexus, then IFC must be created manually.
- 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. **Back2Back&ToExternal** - 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.

Auto Deploy Both: This check box is applicable for the symmetric VRF Lite deployment. When you check this check box, the **Auto Deploy Flag** is set to true for auto-created IFCs to turn on symmetric VRF Lite configuration. You can check or uncheck this check box when the **VRF Lite Deployment** field is not set to **Manual**. The value you choose takes priority. This flag only affects the new auto-created IFC and it does not affect the existing IFCs.

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.

- Check the **Auto Deploy Flag** check box in the **Link Management** dialog box. Checking this check box enables VRF lite deployment, including VRF Lite sub-interface and BGP peering configuration, on both ends of the link for managed devices.

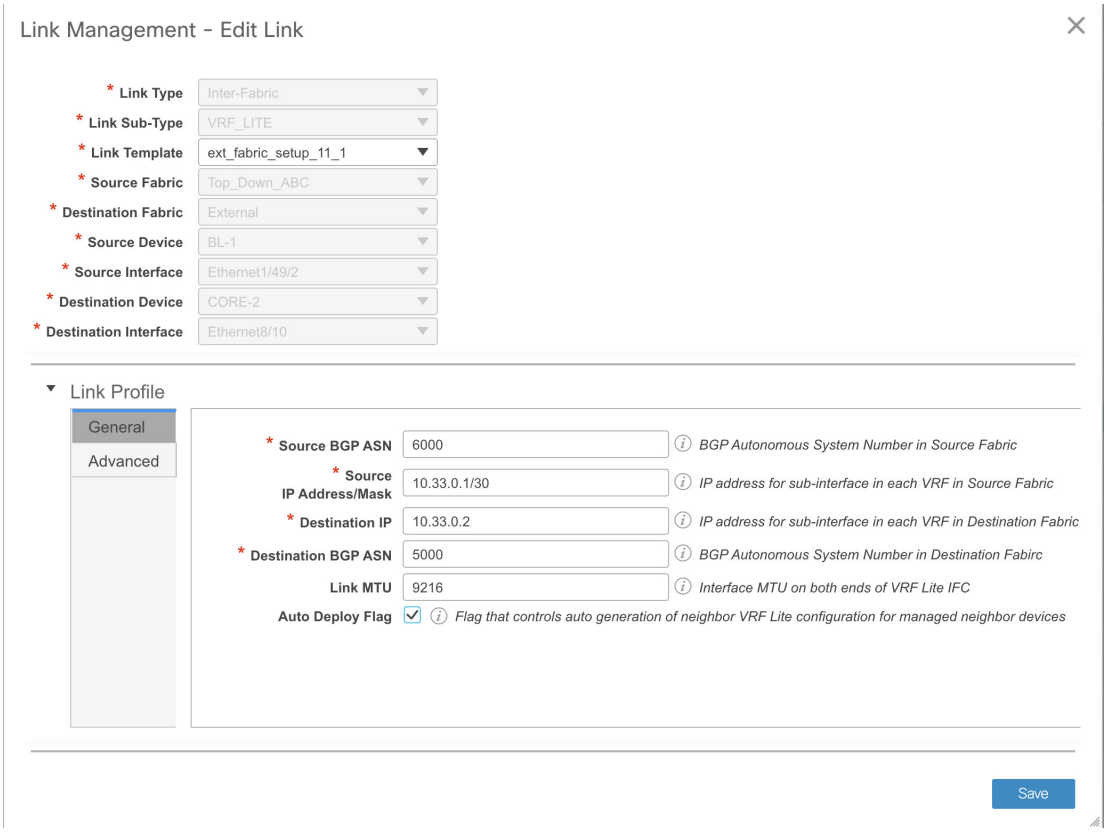

- When you extend the VRF lite in a consecutive scenario, the VRF must be present in the peer fabric and the VRF name must be same. An error message appears, if the VRF is not present in the peer fabric and if you try to extend the VRF lite.
- When you extend the VRF lite between an easy fabric and an external fabric, the VRF name can be same as the source fabric, default, or another VRF name. Enter the VRF name used in the external fabric in the **PEER_VRF_NAME** field. The child PTIs for the subinterface, the VRF creation and BGP peering on the external fabric has the non-empty source. Hence, the policies cannot be edited or deleted from the **View/Edit policies** window.
- Deploy configurations in both the fabrics. Perform **Save & Deploy** on the external fabrics to deploy the configurations. The easy fabric configuration can be deployed either from the topdown VRFs page or from the **Fabric Builder** window.

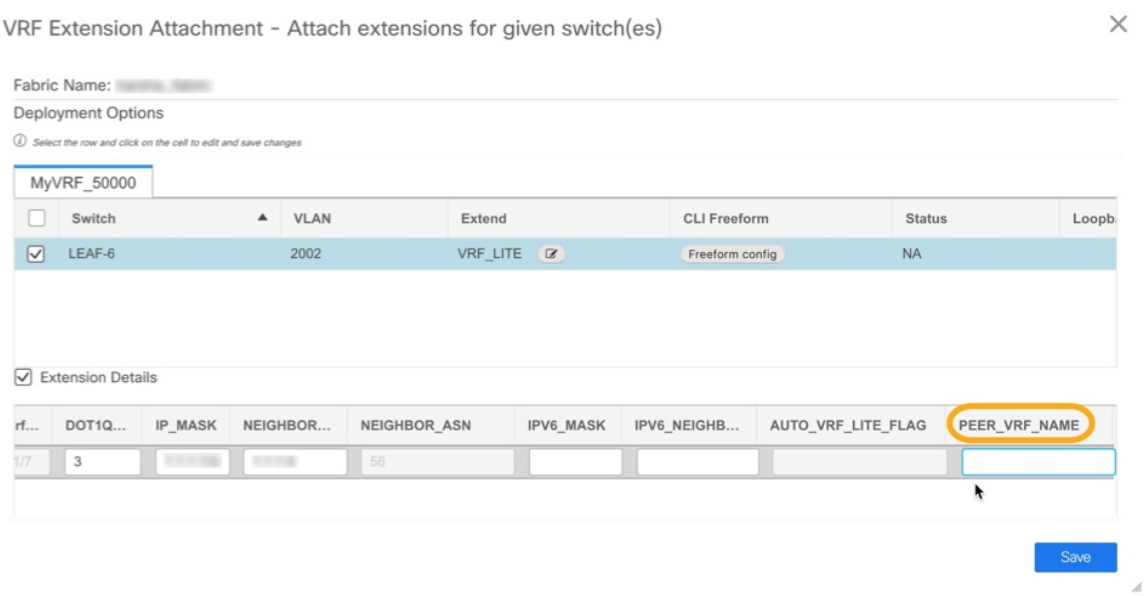

VRF Lite Extension over SVI

VRF Lite extension is achieved by creating a VRF LITE IFC over physical or port-channel interface and selecting VRF LITE IFCs on vrf attachment to switches of border role. In 11.5(3a), you can also have VRF LITE extension over SVI interface using **Ext_VRF_Lite_SVI** template. You can create a policy using this template to generate SVI configuration and IPv4 eBGP peering configuration. The policy must be attached to both directly connected switches (managed by the same DCNM or two different DCNMs) to configure eBGP session. The interfaces connecting the two switches need to be trunk or access port.

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.

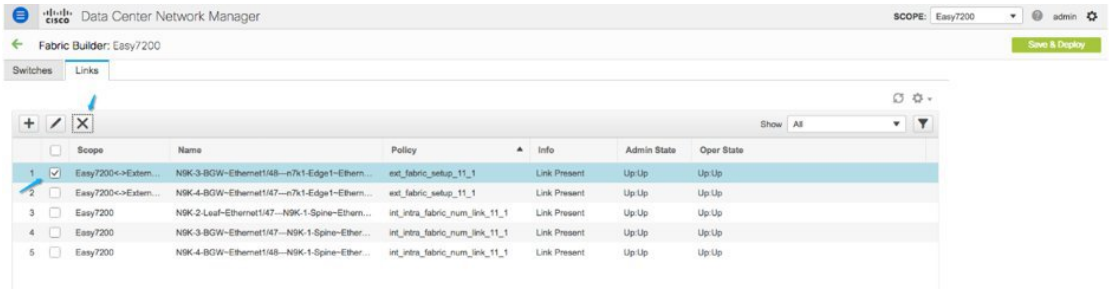

3. Click OK to confirm deletion.
4. Execute the Save and Deploy option in the fabric to reset the VRF Lite policy.

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 the click the delete button.

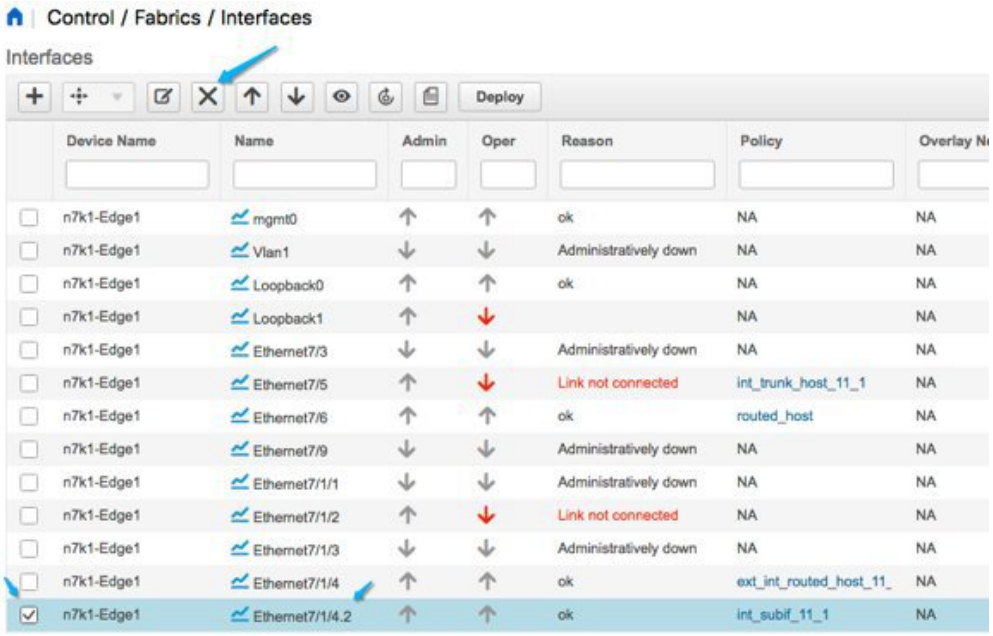

Deleting the eBGP policy

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.

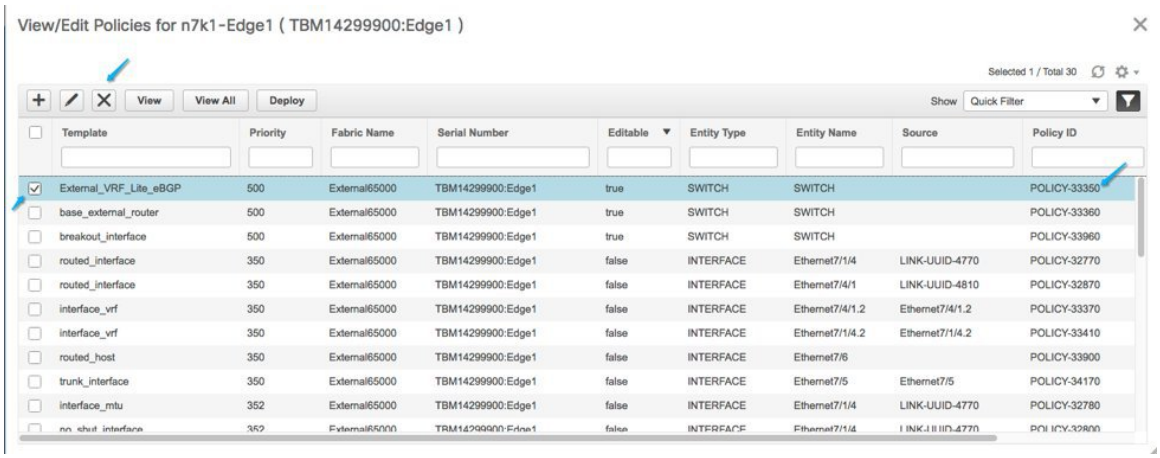

Deleting IFCs Created By Automatic VRF Lite creation

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
Cisco Programmable Fabric with VXLAN BGP EVPN Configuration Guide	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-redist-subnet
            maximum-paths ibgp 2
            network 0.0.0.0/0
        address-family ipv6 unicast
            advertise l2vpn evpn
            redistribute direct route-map fabric-rmap-redist-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
```

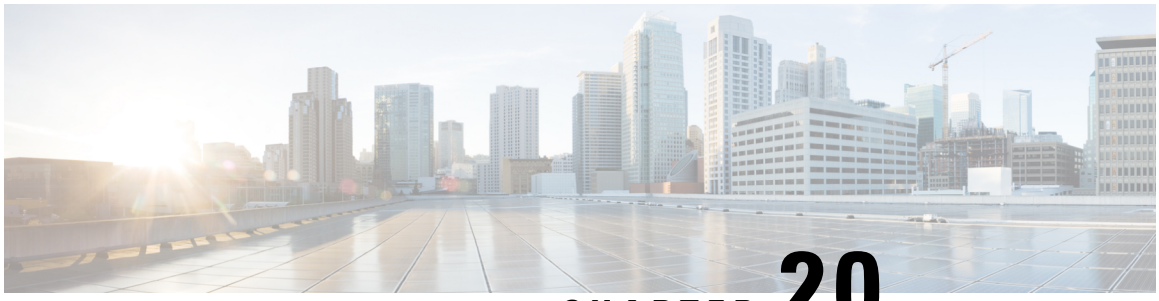

CHAPTER 20

MPLS SR and LDP Handoff

This chapter describes how to configure the MPLS handoff features.

- Overview of VXLAN EVPN to SR-MPLS and MPLS LDP Interconnection, on page 813
- VXLAN MPLS Topology, on page 815
- Configuration Tasks for VXLAN MPLS Handoff , on page 817
- Editing Fabric Settings for MPLS Handoff, on page 817
- Creating an Underlay Inter-Fabric Connection , on page 820
- Creating an Overlay Inter-Fabric Connection, on page 823
- Deploying VRFs, on page 825
- Changing the Routing Protocol and MPLS Settings, on page 827

Overview of VXLAN EVPN to SR-MPLS and MPLS LDP Interconnection

From Cisco DCNM Release 11.3(1), the following handoff features are supported:

- VXLAN to SR-MPLS
- VXLAN to MPLS LDP

These features are provided on the border devices, that is, border leaf, border spine, and border super spine in the VXLAN fabric using the **Easy_Fabric_11_1** template. Note that the devices should be running Cisco NX-OS Release 9.3(1) or later. These DCI handoff approaches are the one box DCI solution where no extra Provider Edge (PE) device is needed in the external fabric.

Note If the switch is running a Cisco NX-OS Release 7.0(3)I7(X), enabling the MPLS handoff feature causes the switch to remove the NVE related config-profile CLIs when the switch is reloaded.

In the DCNM DCI MPLS handoff feature, the underlay routing protocol to connect a border device to an external fabric is ISIS or OSPF, and the overlay protocol is eBGP. The N-S traffic between the VXLAN fabric and external fabric running SR-MPLS or MPLS LDP is supported. Though, you can use DCNM for connecting two Data Center VXLAN fabrics via SR-MPLS or MPLS LDP.

Supported Platforms and Configurations

The following table provides information about the supported platforms:

Feature	Supported Platforms
VXLAN to SR-MPLS	Cisco Nexus 9300-FX2, Jericho+ based Nexus 9000 R-Series, and Nexus 3600 R-Series switches
VXLAN to MPLS LDP	Jericho+ based Cisco Nexus 9000 R-series and Cisco Nexus 3600 R-series switches

The following features aren't supported as they aren't supported on a switch:

- Coexisting of MPLS LDP and SR-MPLS interconnections
- vPC

The VXLAN to SR-MPLS handoff feature comprises the following configurations:

- Base SR-MPLS feature configuration.
- Underlay configuration between the DCI handoff device and the device in the external fabric for the underlay connectivity. DCNM supports ISIS or OSPF as the routing protocol for the underlay connectivity.
- Overlay configuration between a DCI handoff device and a core or edge router in the external fabric, or another border device in another fabric. The connectivity is established through eBGP.
- VRF profile

The VXLAN to MPLS LDP handoff feature comprises the following configurations:

- Base MPLS LDP feature configuration.
- Underlay configuration between the DCI handoff device and the device in the external fabric for the underlay connectivity. DCNM supports ISIS or OSPF as the routing protocol for the underlay connectivity.
- Overlay configuration between a DCI handoff device and a core or edge router in the external fabric, or another border device in another fabric. The connectivity is established through eBGP.
- VRF profile

Inter-Fabric Connections for MPLS Handoff

The following two inter-fabric connection links are introduced:

- **VXLAN_MPLS_UNDERLAY** for underlay configuration: This link corresponds to each physical link or Layer 3 port channel between the border and the external device (or a P router in MPLS or SR-MPLS). A border device can have multiple inter-fabric connection links as there could be multiple links connected to one or more external devices.
- **VXLAN_MPLS_OVERLAY** for eBGP overlay configuration: This link corresponds to the virtual link between a DCI handoff device and a core or edge router in the external fabric, or another border device in another fabric. This inter-fabric connection link can only be created on border devices which meet the image and platform requirement. A border device can have multiple of this type of IFC link as it could communicate to multiple core or edge routers.

These inter-fabric connections can be manually created by using the DCNM Web UI or REST API. Note that the automatic creation of these inter-fabric connections isn't supported.

VXLAN MPLS Topology

MPLS-SR Topology

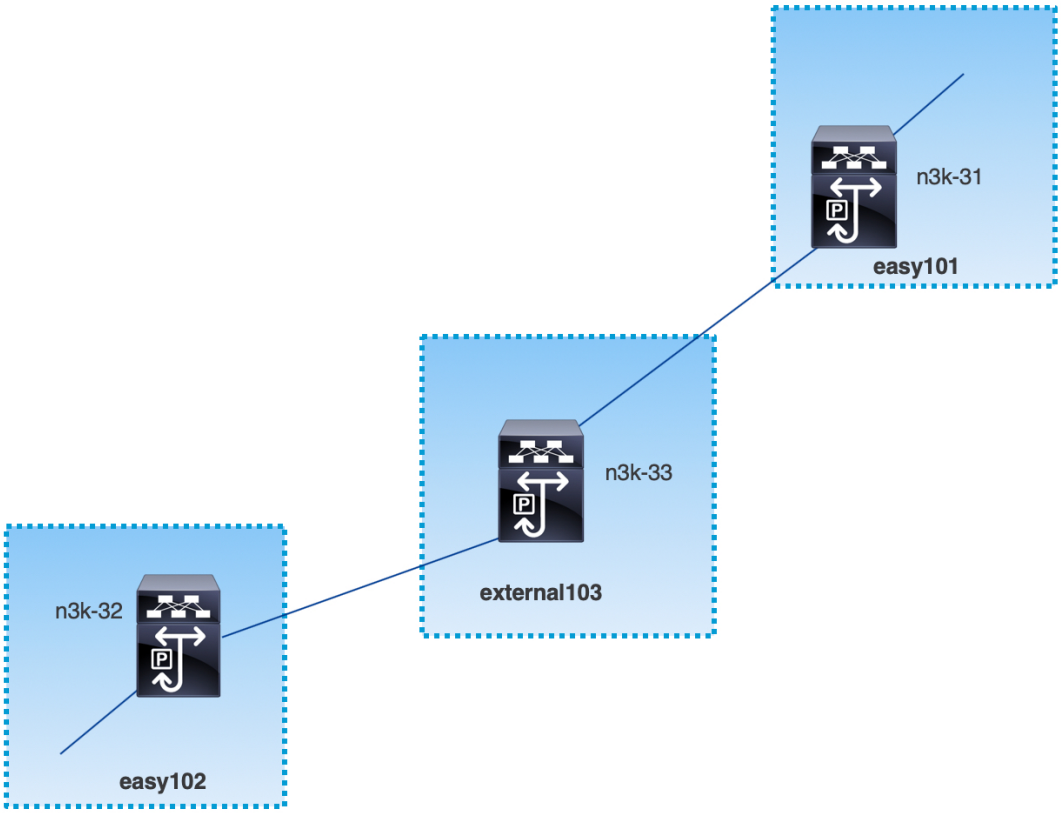

MPLS-LDP Topology

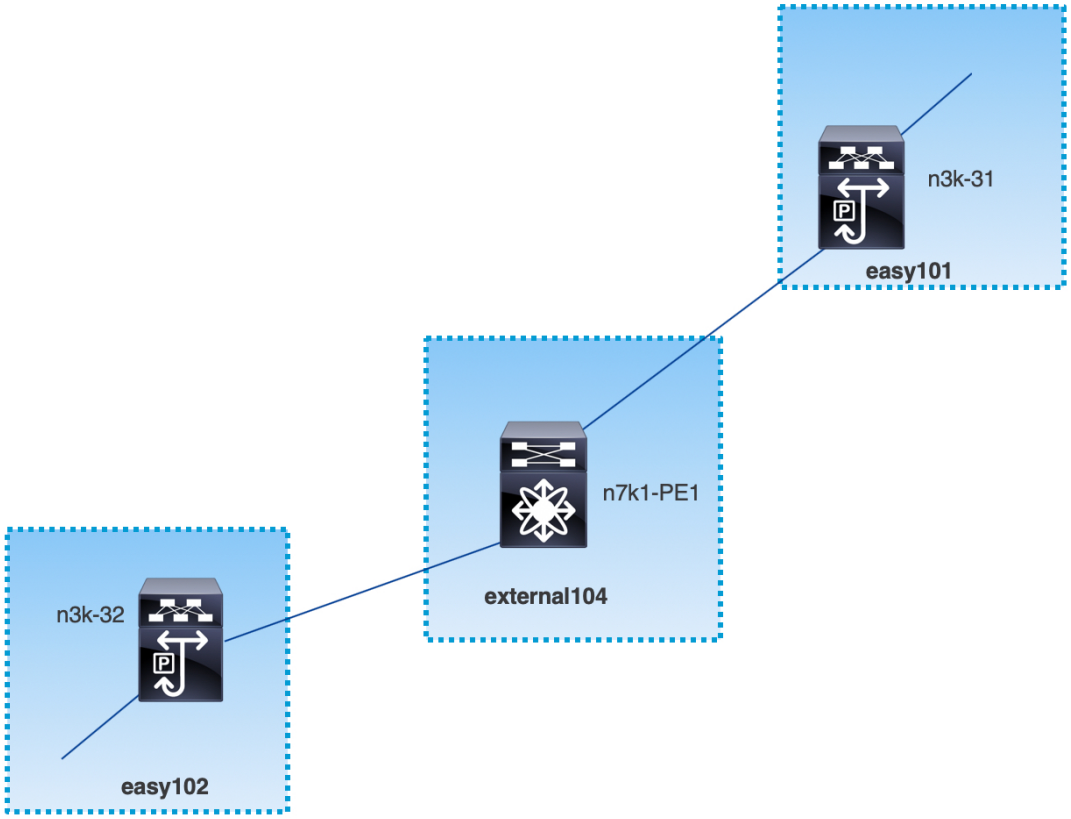

This topology shows only the border devices in the Easy Fabric and the core or edge router in the external fabric.

- The fabrics that are using the **Easy_Fabric_11_1** template are:
  - **easy101**
  - **easy102**
- The fabrics that are using the **External_Fabric_11_1** template are:
  - **external103**
  - **external104**
- The external fabric **external103** is running the MPLS SR protocol.
- The external fabric **external104** is running the MPLS LDP protocol.
- **n3k-31** and **n3k-32** are border devices performing VXLAN to MPLS handoff.
- **n7k-PE1** only supports MPLS LDP.
- **n3k-33** supports SR-MPLS.

Configuration Tasks for VXLAN MPLS Handoff

The following tasks are involved in configuring the MPLS handoff features:

1. Editing the fabric settings to enable MPLS handoff.
2. Creating an underlay inter-fabric connection link between the fabrics.

   Specify whether you're using MPLS SR or LDP in the inter-fabric connection link settings.
3. Creating an overlay inter-fabric connection link between the fabrics.
4. Deploying a VRF for VXLAN to MPLS interconnection.

Editing Fabric Settings for MPLS Handoff

This section shows how to edit the fabric settings for the easy fabric and the external fabric to enable the MPLS handoff feature.

Editing Easy Fabric Settings

Procedure

Step 1 Navigate to **Control > Fabric Builder**.

Step 2 Click the **Edit Fabric** icon to edit the fabric settings.

Step 3 Click the **Advanced** tab.

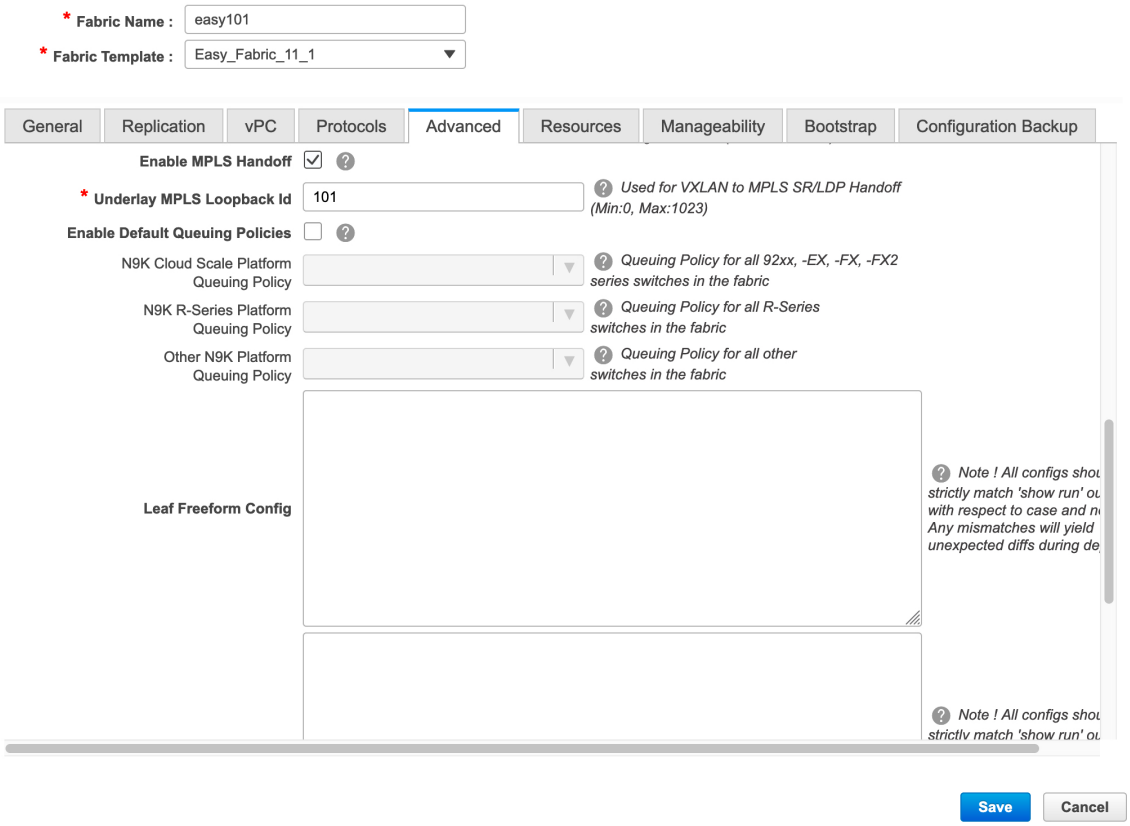

Enable MPLS Handoff: Select the check box to enable the MPLS Handoff feature.

Note: For the brownfield import, 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.

Step 4 Click the **Resources** tab.

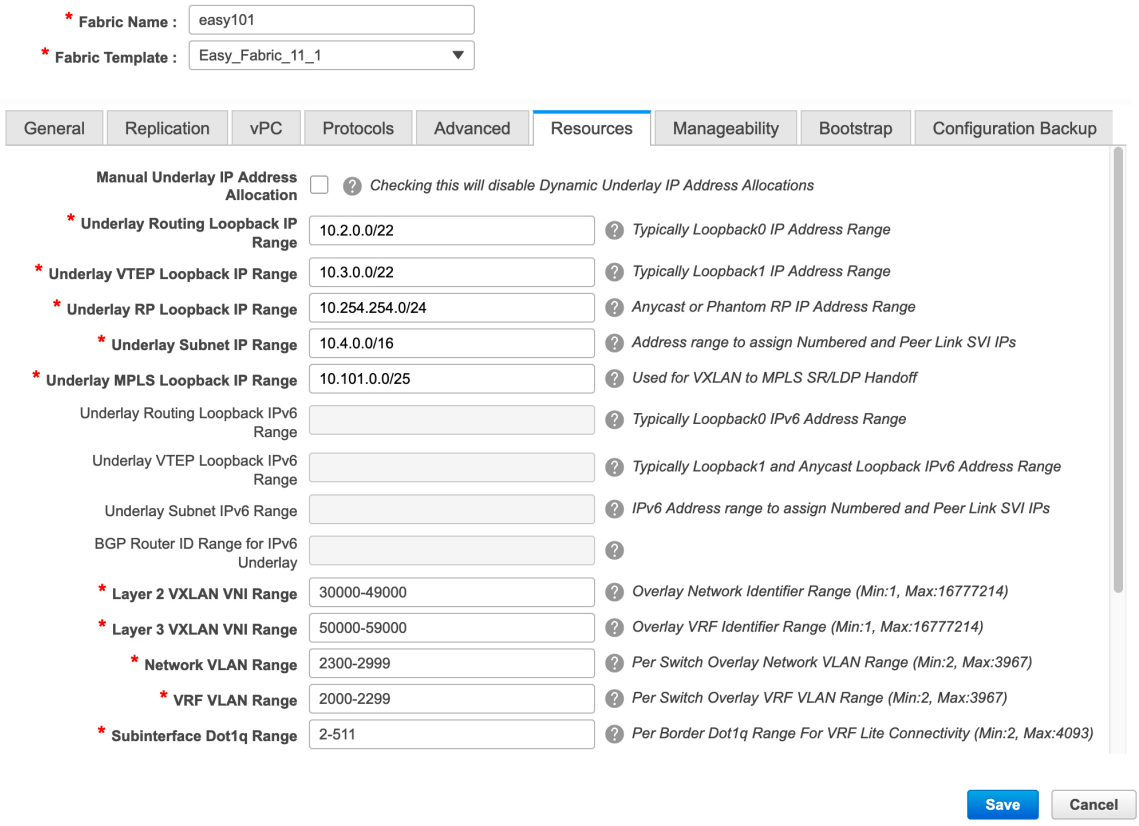

Underlay MPLS Loopback IP Range: Specifies the underlay MPLS loopback IP address range.

For eBGP between Border of Easy A and Easy B, Underlay routing loopback and Underlay MPLS loopback IP range must be a unique range. It should not overlap with IP ranges of the other fabrics, else VPNv4 peering will not come up.

Step 5 Click **Save & Deploy** to configure the MPLS feature on each border device in the fabric.

For more information about remaining fields, see Creating a New VXLAN BGP EVPN Fabric.

Editing External Fabric Settings

Procedure

Step 1 Navigate to **Control > Fabric Builder**.

Step 2 Click the **Edit Fabric** icon to edit the fabric settings.

Step 3 (Optional) Under the **General** tab, uncheck the **Fabric Monitor Mode** check box.

Step 4 Click the **Advanced** tab.

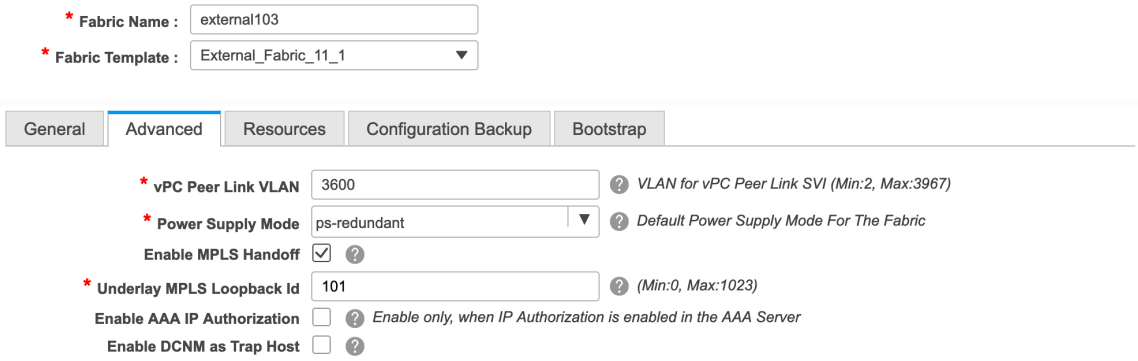

Enable MPLS Handoff: Select the check box to enable the MPLS Handoff feature.

Underlay MPLS Loopback Id: Specifies the underlay MPLS loopback ID. The default value is 101.

Step 5 Click the **Resources** tab.

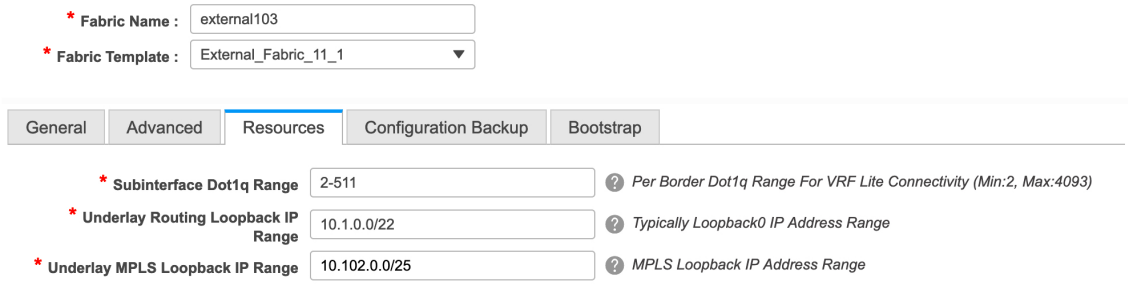

Underlay MPLS Loopback IP Range: Specifies the underlay MPLS SR or LDP loopback IP address range. Note that the IP range should be unique, that is, it should not overlap with IP ranges of the other fabrics.

Step 6 Click **Save & Deploy** to configure the MPLS feature on each edge or core router in the fabric.

For more information about remaining fields, see Creating an External Fabric.

Creating an Underlay Inter-Fabric Connection

This procedure shows how to create an underlay inter-fabric connection link.

Procedure

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

Step 2 Choose a VXLAN fabric from which you want to create an underlay inter-fabric connection to MPLS.

Step 3 Click **Tabular view** in the **Actions** panel that is displayed at the upper left of the window.

Step 4 Click the **Links** tab.

Step 5 Check the existing links that are already discovered for the fabric.

In this example, the link from **easy101** to **external103** is already discovered.

Step 6 Select the existing discovered link and click the **Update Link** icon.

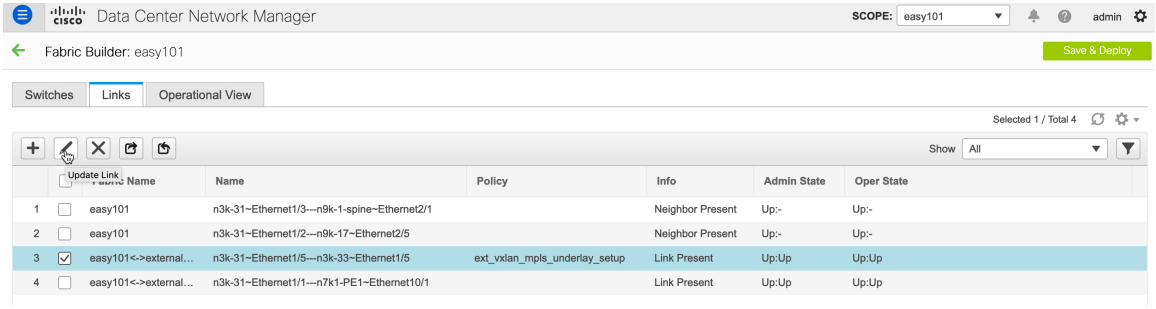

If a link isn't discovered, click the **Add Link** icon and provide all the details for adding an inter-fabric link.

Step 7 In the **Link Management - Edit Link** window, **Link Type** should be **Inter-Fabric**. Choose **VXLAN_MPLS_UNDERLAY** from the **Link Sub-Type** drop-down list and choose **ext_vxlan_mpls_underlay_setup** from the **Link Template** drop-down list.

Step 8 Under **Link Profile**, provide all the required information for the **General** tab.

MPLS-SR Configuration Example for Inter Fabric Link

MPLS-LDP Configuration Example for Inter Fabric Link

IP address/Mask: Specifies the IP address with mask for the source interface.

Neighbor IP: Specifies the IP address of the destination interface.

MPLS Fabric: Specifies whether the external fabric is running SR or LDP.

Note MPLS SR and LDP can't coexist on a single device.

Source SR Index: Specifies a unique SID index for the source border. This field is disabled if you choose **LDP** in the **MPLS Fabric** field.

Destination SR Index: Specifies a unique SID index for the destination border. This field is disabled if you choose LDP for the **MPLS Fabric** field.

SR Global Block Range: Specifies the SR global block range. You need to have the same global block range across the fabrics. The default range is from 16000 to 23999. This field is disabled if you choose LDP for the **MPLS Fabric** field.

DCI Routing Protocol: Specifies the routing protocol used on the DCI MPLS underlay link. You can choose either **is-is** or **ospf**.

OSPF Area ID: Specifies the OSPF area ID if you choose OSPF as the routing protocol.

DCI Routing Tag: Specifies the DCI routing tag used for the DCI routing protocol.

Step 9 Click **Save**.

Step 10 Click **Save & Deploy** to deploy the updated configurations.

Step 11 In the **Config Deployment** window, click **Deploy Config**.

Step 12 Navigate to the destination fabric from the **Fabric Builder** window and perform a **Save & Deploy**, that is, perform steps 10 and 11.

Creating an Overlay Inter-Fabric Connection

This procedure shows how to create an overlay inter-fabric connection after the underlay inter-fabric connection is created. The overlay inter-fabric connection is the same for MPLS SR and LDP because the overlay connection uses eBGP.

Procedure

Step 1 Click the **Add Link** icon.

Step 2 In the **Link Management - Add Link** window, specify all the details.

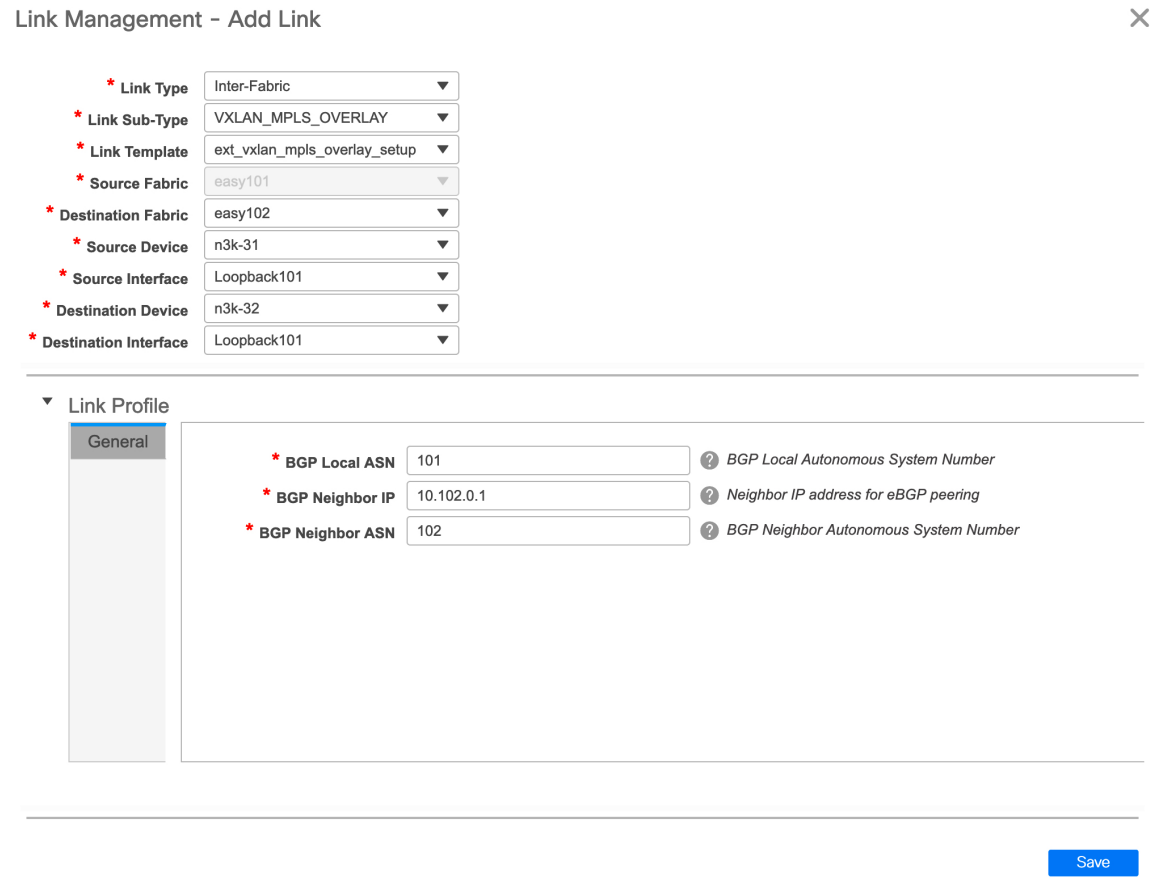

Link Type: Choose **Inter-Fabric**.

Link-Sub Type: Choose **VXLAN_MPLS_OVERLAY** from the drop-down list.

Link Template: Choose **ext_vxlan_mpls_overlay_setup** from the drop-down list.

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 the MPLS loopback interface. The IP address of the loopback interface will be used for overlay eBGP peering.

Destination Device and **Destination Interface**: Choose the destination device and a loopback interface that connects to the source device.

General tab in the **Link Profile** section.

BGP Local ASN: In this field, the AS number of the source device is autopopulated.

BGP Neighbor IP: Fill up this field with the IP address of the loopback interface at the destination device for eBGP peering.

BGP Neighbor ASN: In this field, the AS number of the destination device is autopopulated.

Step 3 Click **Save**.

Step 4 Click **Save & Deploy** to deploy the updated configurations.

Step 5 In the **Config Deployment** window, click **Deploy Config**.

Step 6 Navigate to the destination fabric from the **Fabric Builder** window and perform a **Save & Deploy**, that is, perform steps 4 and 5.

Note If there is only one MPLS overlay IFC link on the switch, you can remove it only when there's no VRF attached to either end of the MPLS overlay link.

Deploying VRFs

This procedure shows how to deploy VRFs for VXLAN to MPLS interconnection.

Note When you use the 4 byte ASN and auto route target is configured, the route target that is automatically generated is 23456:VNI. If two different VRFs in two different fabrics have the same VNI value, the route-target of the two VRFs would be the same due to auto route target and the value 23456 is always constant. For two fabrics connected via VXLAN MPLS handoff, this could result in unintended route exchange. Therefore, for security reasons, if you want to disable auto route target, you can disable it by customizing the network template and network extension template.

Procedure

Step 1 Navigate to **Control > Fabrics > VRFs**.

Step 2 In the **VRFs** window, click the **Add** icon to create a VRF. For more information, see Creating VRFs for the Standalone Fabric.

Step 3 Select the newly added VRF and click **Continue**.

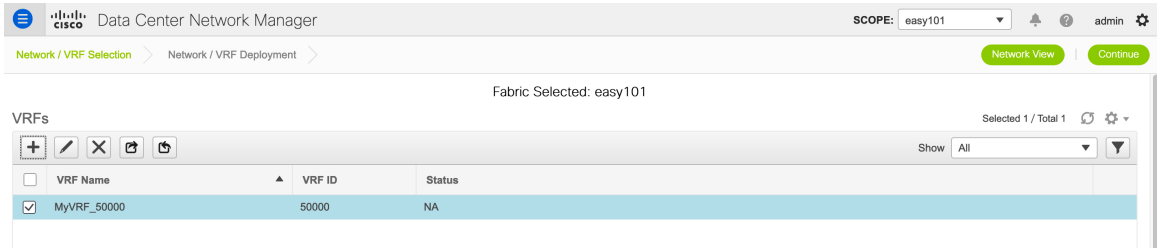

Step 4 In the **VRF Deployment** window, you can see the topology of the fabric. Select a border device to attach a VRF to the border device where the MPLS LDP IFC link is created.

In this example, **n3k-31** is the border device in the **easy101** fabric.

Step 5 In the **VRF Extension Attachment** window, select the VRF and click the **Freeform config** button under the CLI Freeform column.

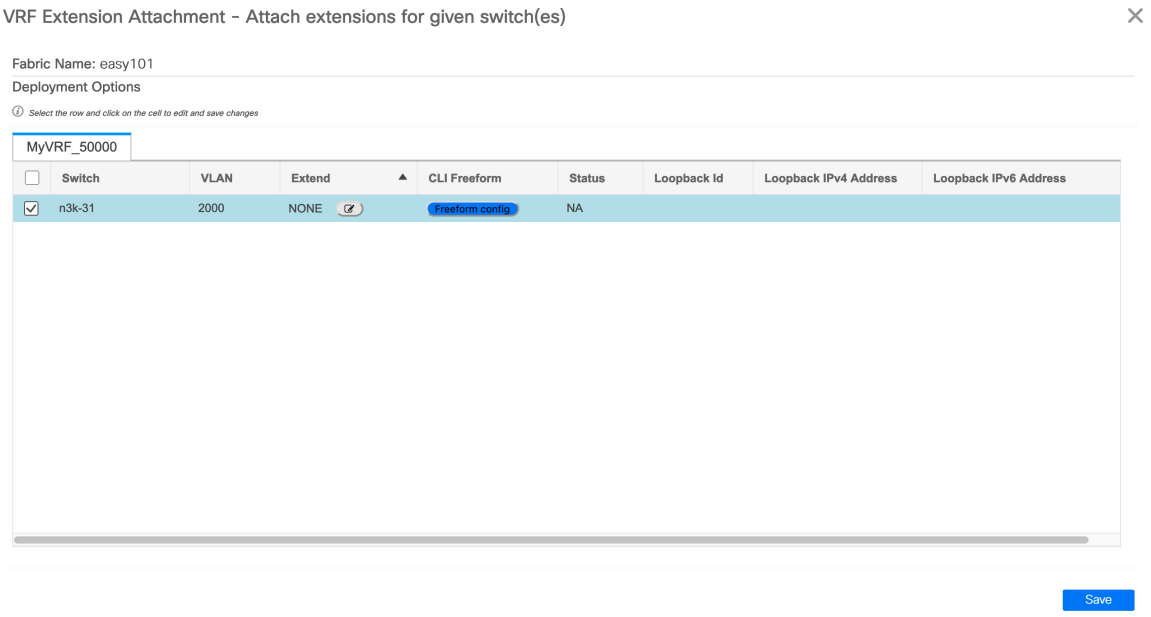

Step 6 Add the following freeform config manually to the VRF:

```
vrf context $$VRF_NAME$$
  address-family ipv4 unicast
    route-target import $$REMOTE_PE_RT$$
  address-family ipv6 unicast
    route-target import $$REMOTE_PE_RT$$
```

In the freeform config, *REMOTE_PE_RT* refers to the neighbor's BGP ASN and VNI number in the **ASN:VNI** format if the neighbor is a border device in Easy Fabric managed by DCNM.

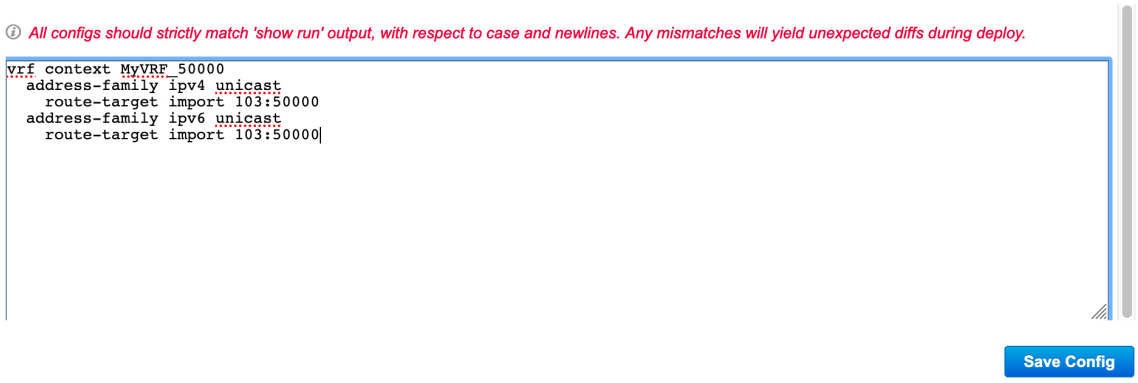

Step 7 Click **Save Config**.

Step 8 (Optional) Enter the Loopback Id and Loopback IPv4 Address and IPv6 address for the border device.

Step 9 Click **Save**.

Step 10 (Optional) Click the **Preview** icon in the **VRF Deployment** window to preview the configuration that will be deployed.

Step 11 Click **Deploy**.

Perform the same task from Step 3 to Step 11 in the destination fabric if the neighbor is a border device in Easy Fabric managed by DCNM.

Changing the Routing Protocol and MPLS Settings

This procedure shows how to change the routing protocol of a device from using IS-IS to OSPF, or from using MPLS SR to LDP for underlay IFC.

Note MPLS SR and LDP cannot co-exist on a device, and using both IS-IS and OSPF for MPLS handoff on the same device is not supported.

Procedure

Step 1 Remove all the MPLS underlay and overlay IFCs from the device that needs the change of DCI routing protocol or MPLS fabric.

Step 2 Click **Save & Deploy** for each fabric that is involved in the removal of the IFCs.

This step deletes all global MPLS SR/LDP configurations and the MPLS loopback interface that was previously created.

Step 3 Create a new IFC using the preferred DCI routing protocol and MPLS settings. For more information, see Creating an Underlay Inter-Fabric Connection , on page 820.

PART IV

Layer-2/Layer-3 DCI with VXLAN EVPN Multi-Site

- Auto-Provisioning Border Gateways with Multi-Site Domains, on page 831

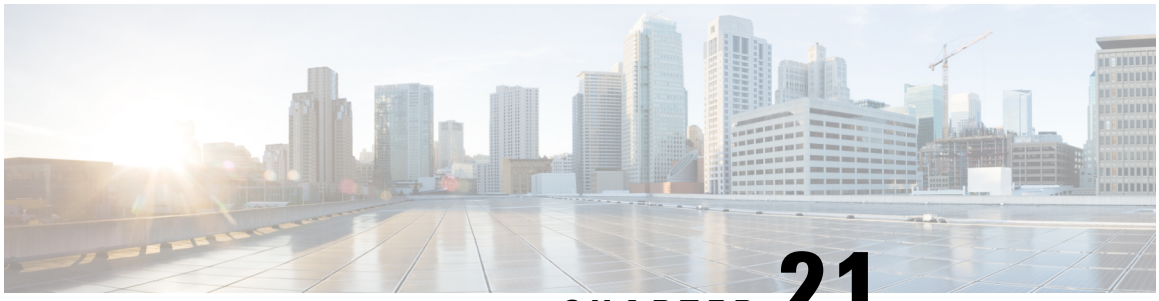

CHAPTER 21

Auto-Provisioning Border Gateways with Multi-Site Domains

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 831
- Prerequisites , on page 832
- Limitations, on page 833
- Save & Deploy Operation in the MSD Fabric , on page 833
- EVPN Multi-Site Configuration , on page 835
- Viewing, Editing and Deleting Multi-Site Overlays , on page 845
- Deleting Multi-Site IFCs, on page 846
- Creating and Deploying Networks and VRFs in the MSD Fabric , on page 847
- Deploying a Legacy Site BGW (vPC-BGWs), on page 850
- Additional References, on page 854
- Appendix , on page 854

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
`evpn multisite border-gateway 7200` `  delay-restore time 300`	7200 is the site ID of the member fabric Easy7200. 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.
`interface nve1` `multisite border-gateway interface loopback100`	The loopback interface 100 is the configuration set in the MSD fabric settings. Once a loopback ID is chosen and **Save and Deploy** is executed, the loopback ID cannot be changed. To modify the role of the BGW in the MSD fabric, perform the following steps: **1.** In the easy fabric, modify the role of the BGW to leaf or border. **2.** Save and deploy the changes. This will remove the loopback 100 from the switch **3.** Change role back to BGW, and do a save and deploy. **4.** In the MSD fabric, change the loopback ID setting to a desired value, and do a save and deploy.
`interface ethernet1/47` `  evpn multisite fabric-tracking`	The **evpn multisite fabric-tracking** command is configured on all ports on a Border Gateway that have a connection to a switch with a Spine role. In case of a Border Gateway Spine role, all ports facing the leaf switch have this command configured

`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`
	The Multi-Site loopback interface. This is configured on all Border Gateway (Spines).
	All BGWs in the same fabric get the same IP address. Each fabric gets its own unique IP address.
	It is not possible to change this address or ID without first changing role of the BGW.

Configuration	Description
`route-map rmap-redist-direct permit 10` `  match tag 54321`	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.

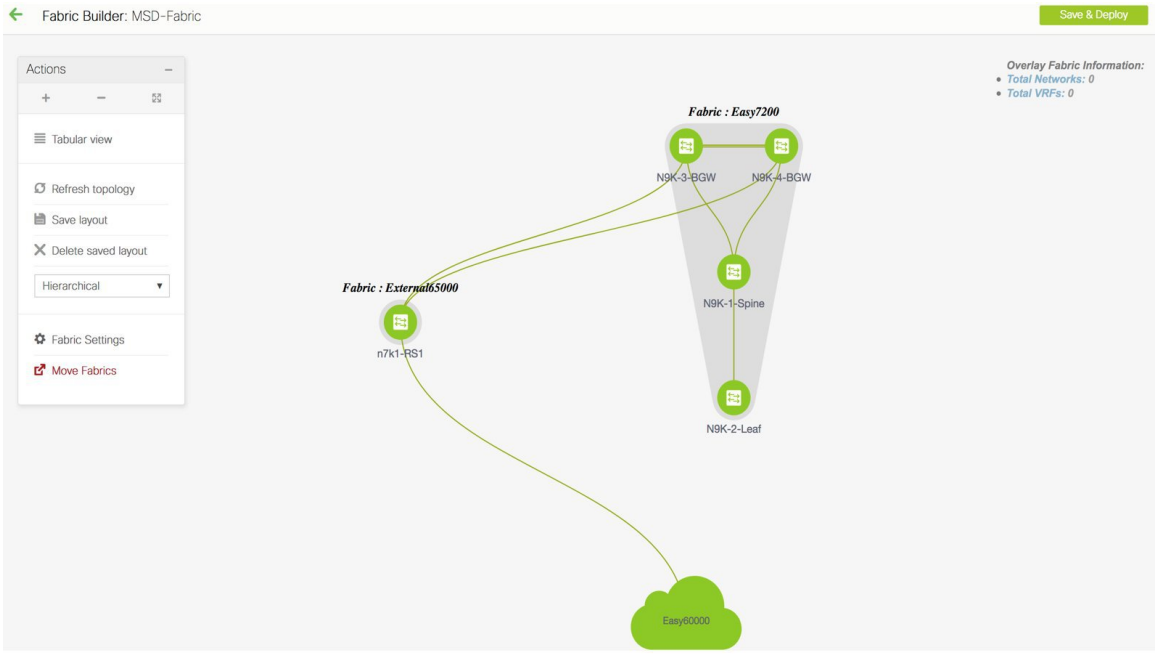

VXLAN BGP EVPN intra-fabric configurations, including network and VRF configurations are provisioned on the switches through DCNM software, 11.1(1) release. However, server traffic between the fabrics is only possible through the following configurations:

- 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 release 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. Ent
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 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.
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.

2. From the **Actions** panel at the left, click **Tabular view**.

   The **Switches | Links** window appears.

3. Click the **Links** tab and click **+**.

   The **Add Link** window appears.

4. Fill in the fields.

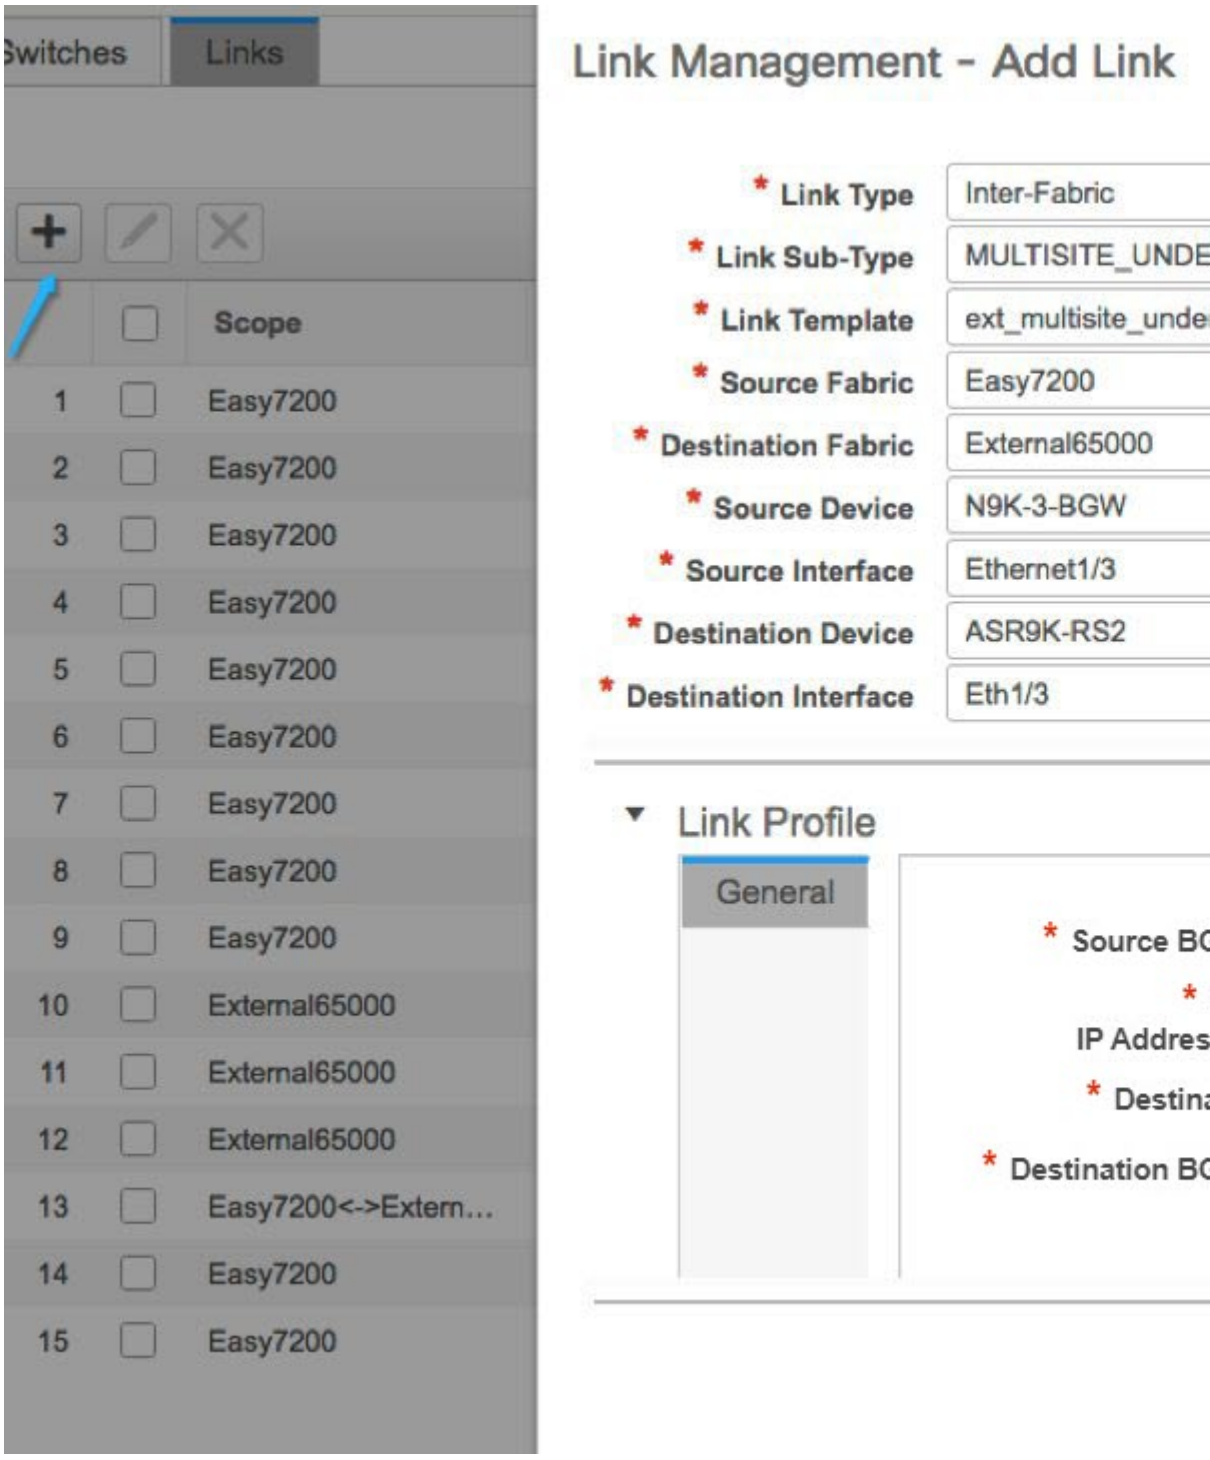

Link Type – Choose **Inter-Fabric**.

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.

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. 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 Route_Server (such as for this example) or Direct to BGW. By default, the Manual option is selected.

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 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 Route_Server is selected in the Deploy Border Gateway Method field in the above image.

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.

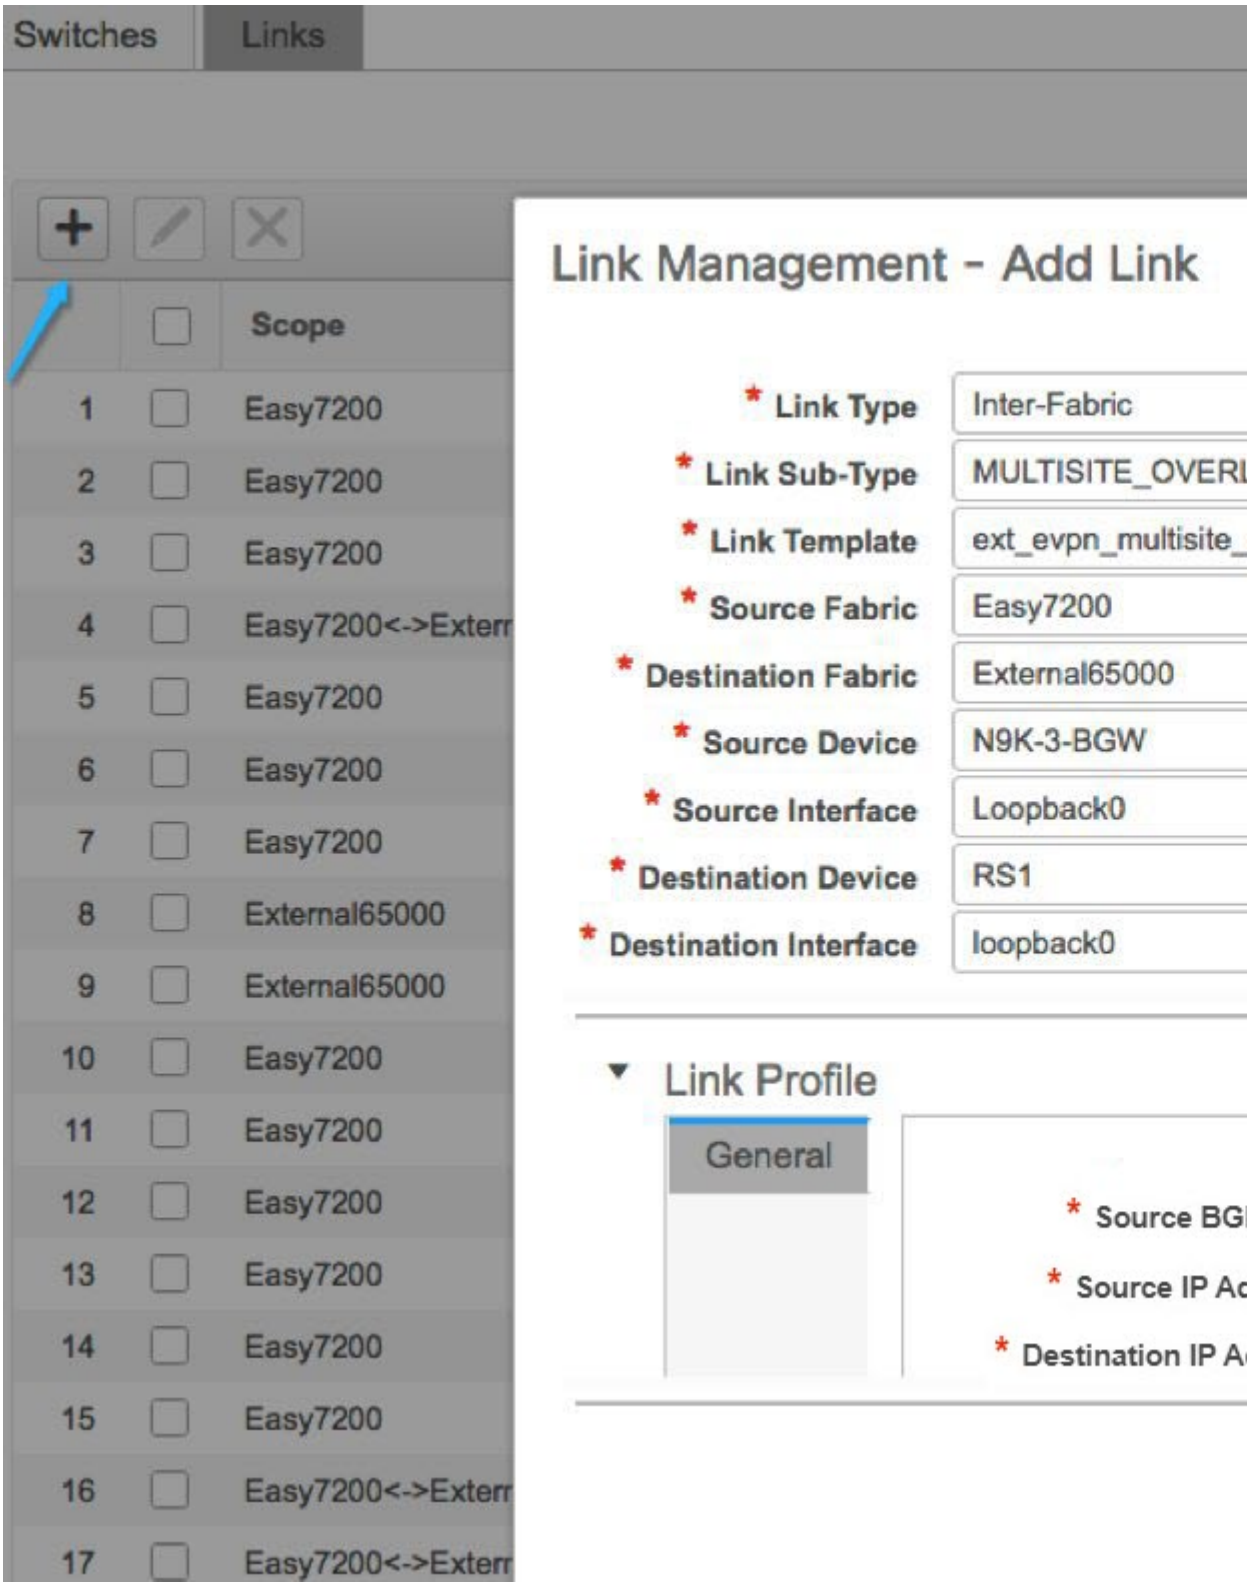

Link Type – Choose **Inter-Fabric**.

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.

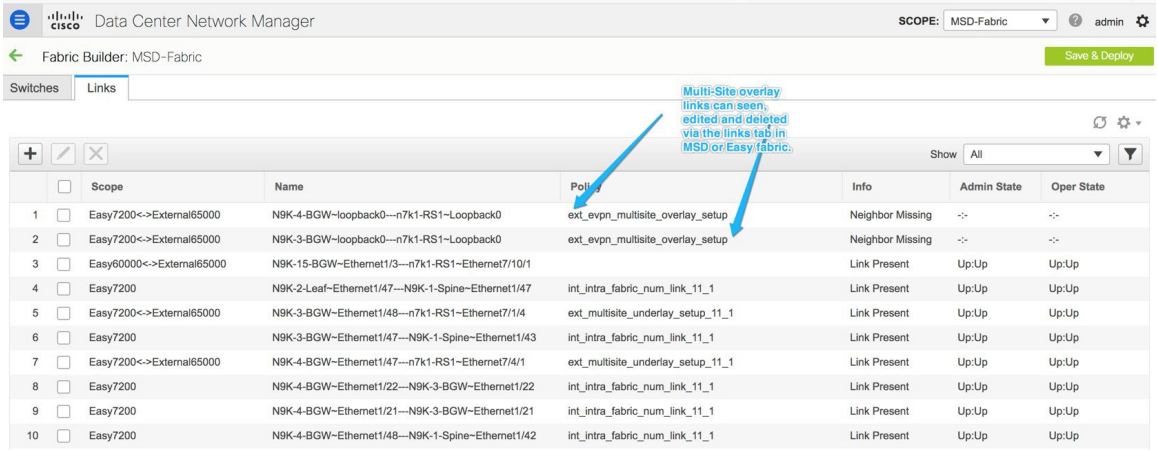

	Scope	Name	Policy	Info	Admin State	Oper State
1	Easy7200<->External65000	N9K-4-BGW~loopback0---n7k1-RS1~Loopback0	ext_evpn_multisite_overlay_setup	Neighbor Missing	-:-	-:-
2	Easy7200<->External65000	N9K-3-BGW~loopback0---n7k1-RS1~Loopback0	ext_evpn_multisite_overlay_setup	Neighbor Missing	-:-	-:-
3	Easy60000<->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.

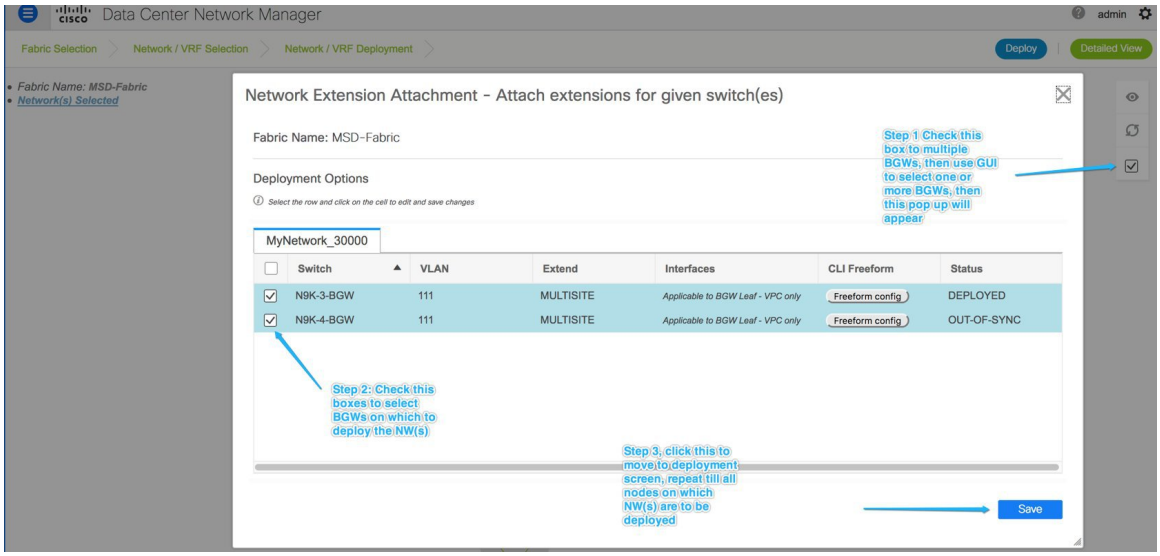

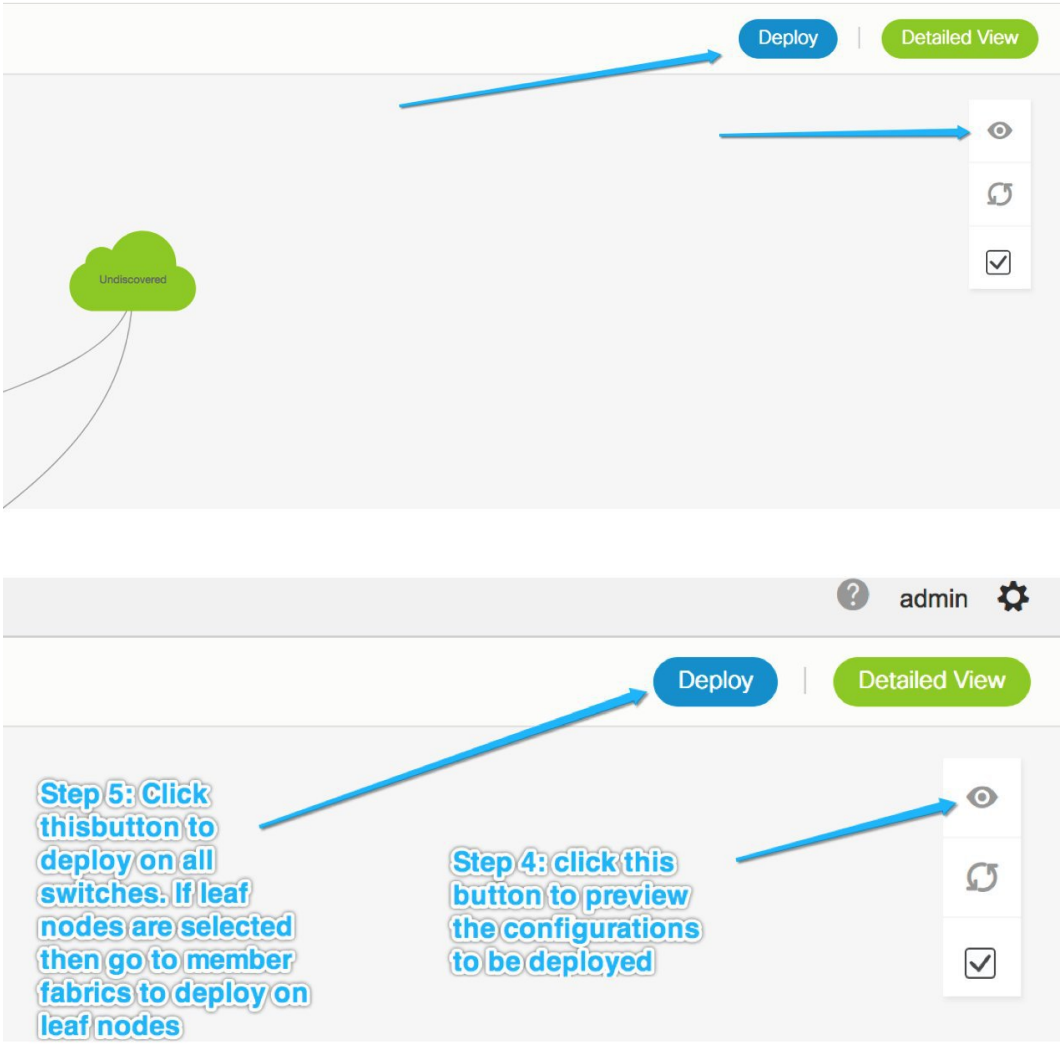

Deploying Networks with a Layer 3 Gateway on a BGW

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 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.

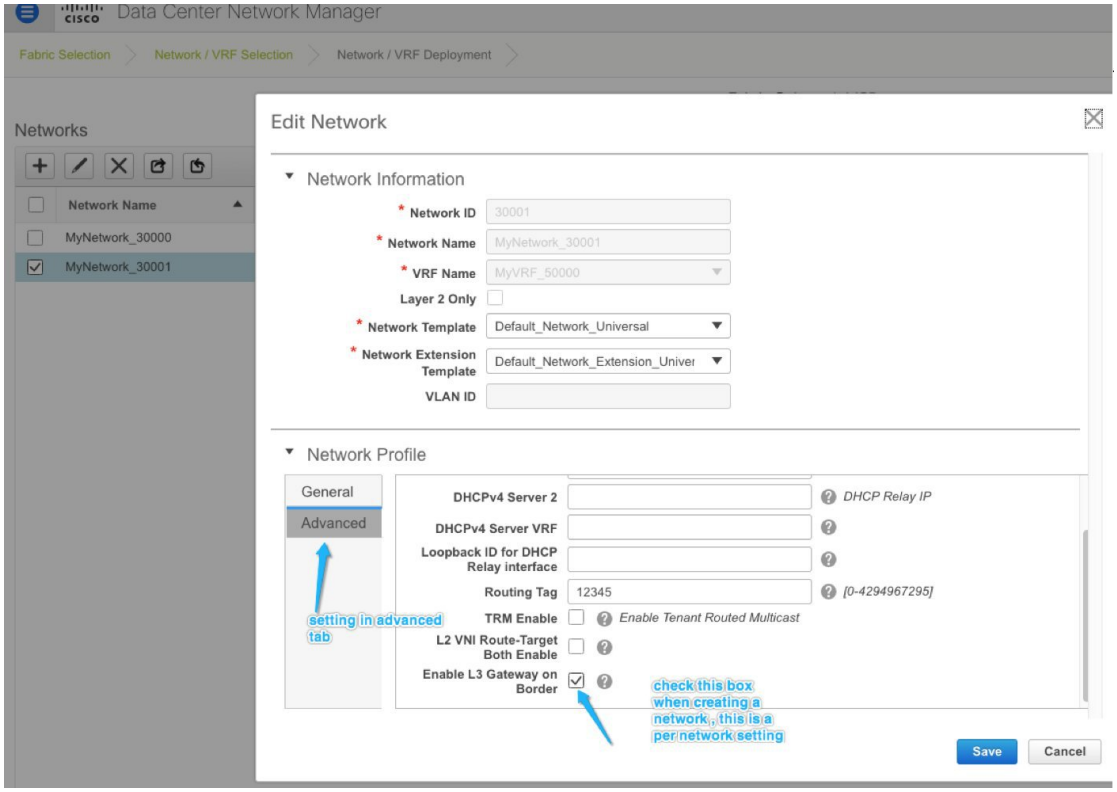

2. When deploying the network on the vPC pair of BGWs, select the interface(s) in the Interfaces column. Only vPC port channel interfaces are the candidate interfaces.

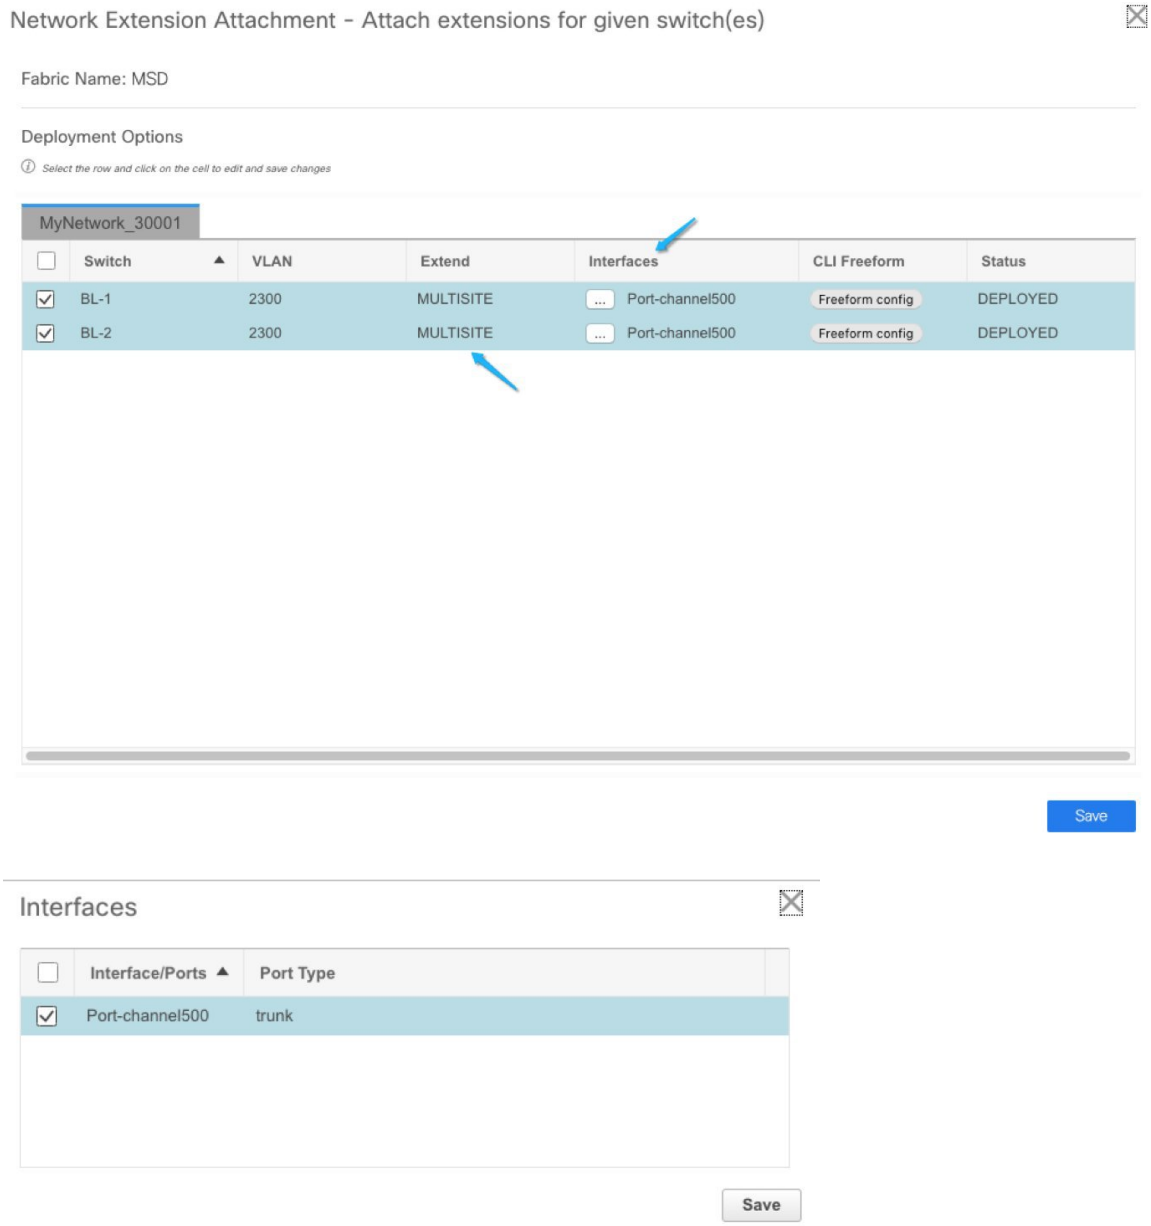

Deploying a Legacy Site BGW (vPC-BGWs)

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.

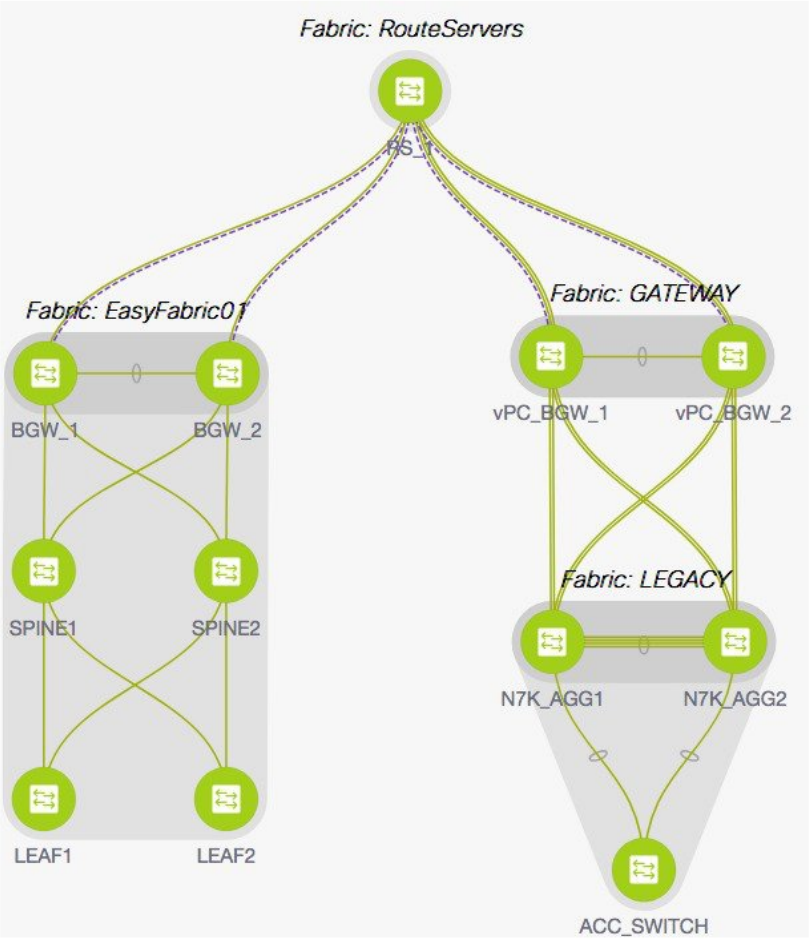

This topology contains the following five fabrics:

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.

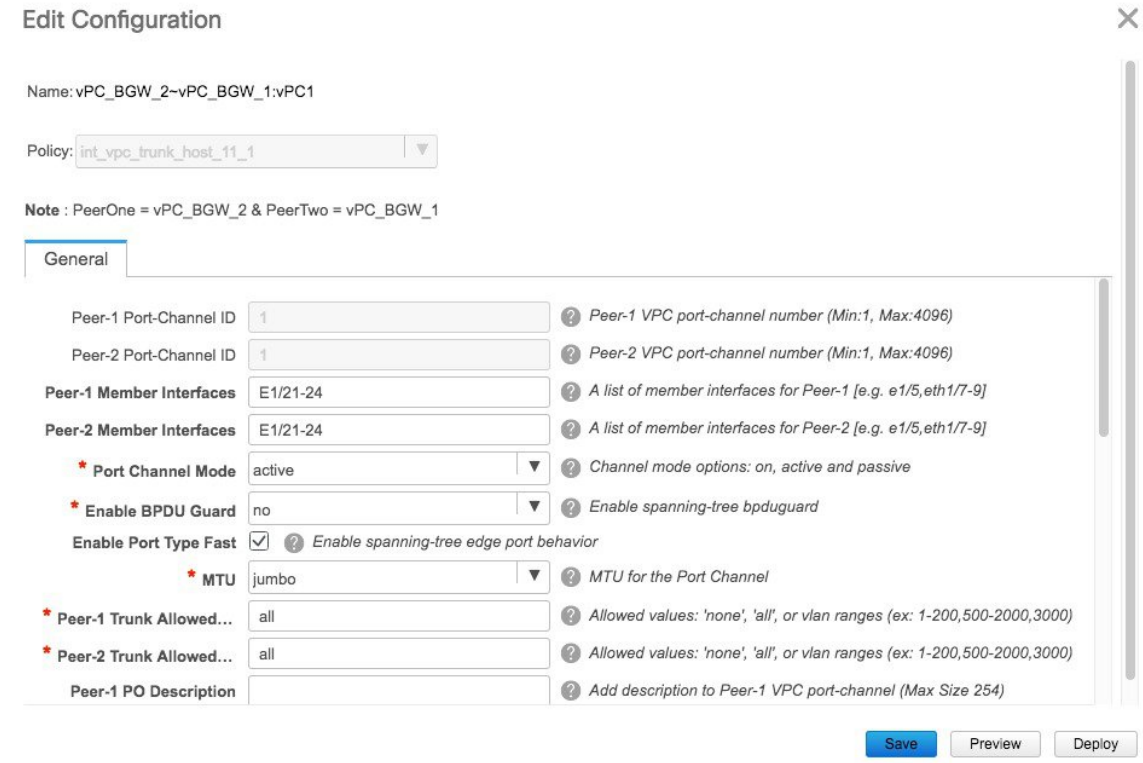

After entering all the information, click **Preview** to preview the configurations that are deployed, and then click **Deploy**.

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.

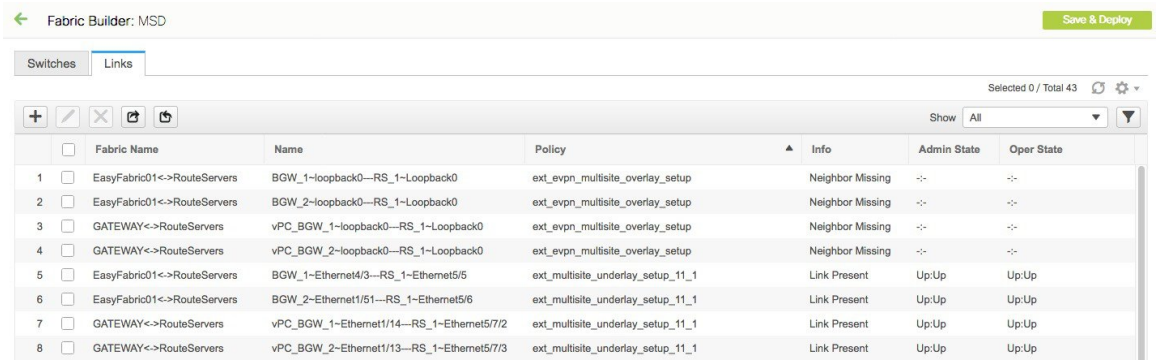

	Fabric Name	Name	Policy	Info	Admin State	Oper State
1	EasyFabric01<->RouteServers	BGW_1~loopback0---RS_1~Loopback0	ext_evpn_multisite_overlay_setup	Neighbor Missing	-:-	-:-
2	EasyFabric01<->RouteServers	BGW_2~loopback0---RS_1~Loopback0	ext_evpn_multisite_overlay_setup	Neighbor Missing	-:-	-:-
3	GATEWAY<->RouteServers	vPC_BGW_1~loopback0---RS_1~Loopback0	ext_evpn_multisite_overlay_setup	Neighbor Missing	-:-	-:-
4	GATEWAY<->RouteServers	vPC_BGW_2~loopback0---RS_1~Loopback0	ext_evpn_multisite_overlay_setup	Neighbor Missing	-:-	-:-
5	EasyFabric01<->RouteServers	BGW_1~Ethernet4/3---RS_1~Ethernet5/5	ext_multisite_underlay_setup_11_1	Link Present	Up:Up	Up:Up
6	EasyFabric01<->RouteServers	BGW_2~Ethernet1/51---RS_1~Ethernet5/6	ext_multisite_underlay_setup_11_1	Link Present	Up:Up	Up:Up
7	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	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
VXLAN EVPN Multi-Site Design and Deployment White Paper	This document explains Multi-Site design and deployment in detail.
Configuring VXLAN EVPN Multi-Site	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.

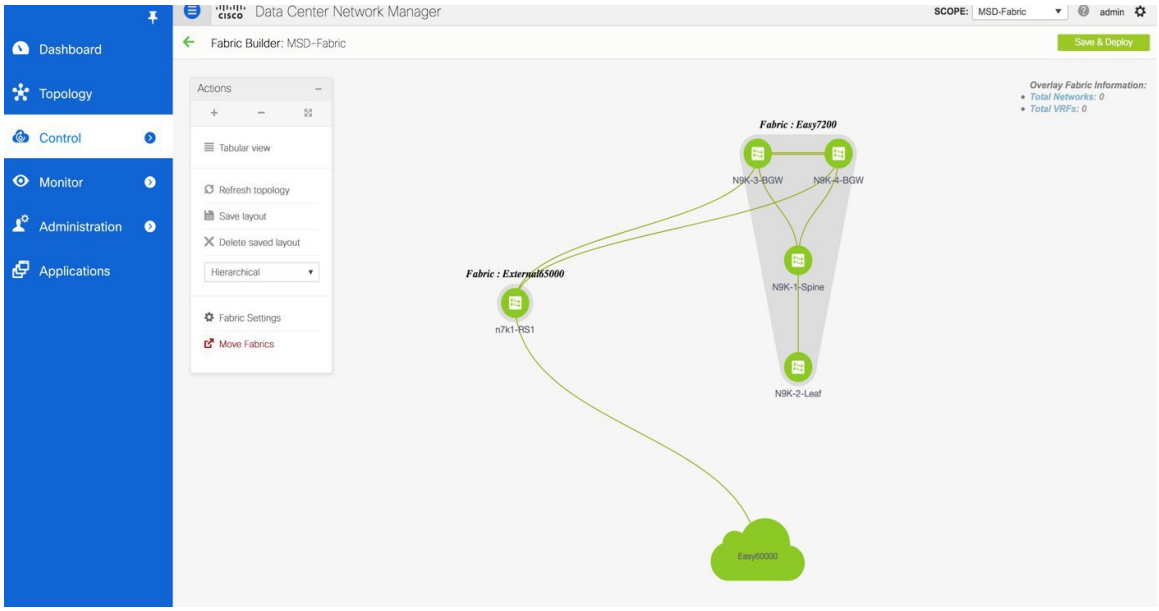

N9K-BGW-3	N9K-BGW-4
router bgp 7200 neighbor 10.4.0.17 remote-as 7200 update-source ethernet1/22 address-family ipv4 unicast next-hop-self	router bgp 7200 neighbor 10.4.0.18 remote-as 7200 update-source Ethernet1/22 address-family ipv4 unicast next-hop-self
router bgp 7200 neighbor 10.4.0.13 remote-as 7200 update-source ethernet1/21 address-family ipv4 unicast next-hop-self	router bgp 7200 neighbor 10.4.0.14 remote-as 7200 update-source Ethernet1/21 address-family ipv4 unicast next-hop-self
interface ethernet1/22 evpn multisite dci-tracking no switchport ip address 10.4.0.18/30 description connected-to-N9K-4-BGW--Ethernet1/22	interface Ethernet1/22 evpn multisite dci-tracking no switchport ip address 10.4.0.17/30 description connected-to-N9K-3-BGW-Ethernet1/22

N9K-BGW-3	N9K-BGW-4

interface ethernet1/21
  evpn multisite dci-tracking
  no switchport
  ip address 10.4.0.14/30
  description
connected-to-N9K-4-BGW-Ethernet1/21
	interface Ethernet1/21
	  evpn multisite dci-tracking
	  no switchport
	  ip address 10.4.0.13/30
	  description
	connected-to-N9K-3-BGW-Ethernet1/21

Changing loopback0 Policy to Modify IP Address

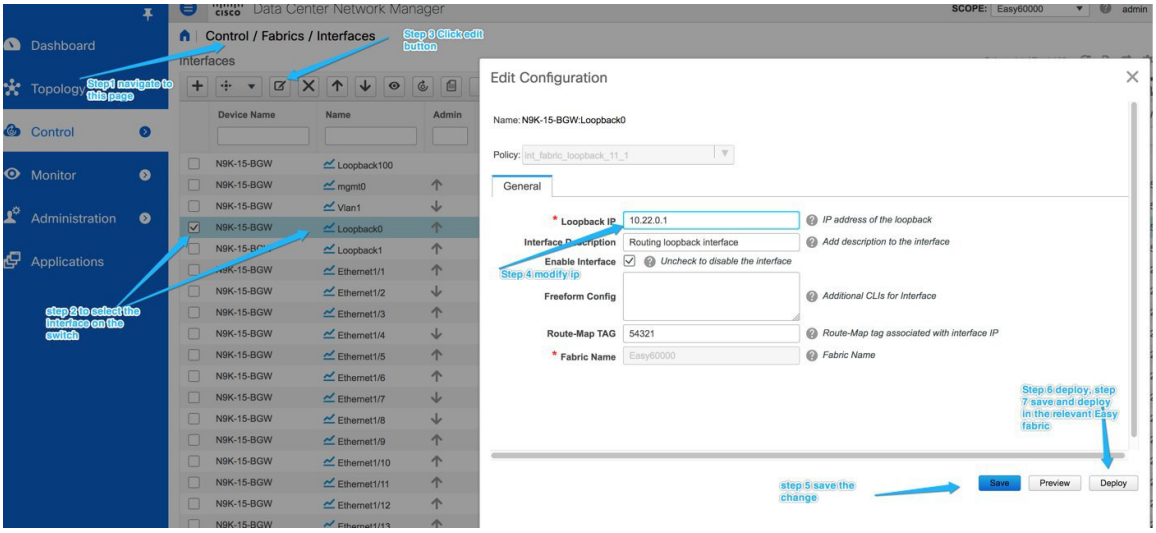

Route Server Configuration

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
route-map unchanged permit 10 set ip next-hop unchanged	–

Configuration	Description

router bgp 65000
  address-family ipv4 unicast
  network /32
	The network command to redistribute the BGP peering address of RS1 to the eBGP underlay sessions so that BGWs know how to reach RS.
	If operator is using a different method to distribute the route server peering address to BGW, then this is not needed

interface ethernet1/22
  evpn multisite dci-tracking
  no switchport
  ip address 10.4.0.18/30
  description
connected-to-N9K-4-BGW--Ethernet1/22
	interface Ethernet1/22
	  evpn multisite dci-tracking
	  no switchport
	  ip address 10.4.0.17/30
	  description
	connected-to-N9K-3-BGW-Ethernet1/22

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
	The knob in the external fabric controls if send community is sent in the form shown here, or as send-community both.
	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.

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

router bgp 7200
  neighbor
  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
	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

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
router bgp 7200 neighbor 10.10.1.6 remote-as 65000 update-source ethernet1/47 address-family ipv4 unicast next-hop-self	router bgp 65000 neighbor 10.10.1.5 remote-as 7200 update-source ethernet7/4/1 address-family ipv4 unicast next-hop-self
interface ethernet1/47 mtu 9216 no shutdown no switchport ip address 10.10.1.5/30 tag 54321 evpn multisite dci-tracking	interface ethernet7/4/1 mtu 9216 no shutdown no switchport ip address 10.10.1.6/30 tag 54321

The tag 54321 attached to the IP address is not required for correct functioning and will be removed in subsequent releases. It is benign.

PART V

Network Provisioning for L4-Layer7 Services

- **L4-L7 Service Basic Workflow**, on page 861
- **L4-L7 Service Use Cases**, on page 903

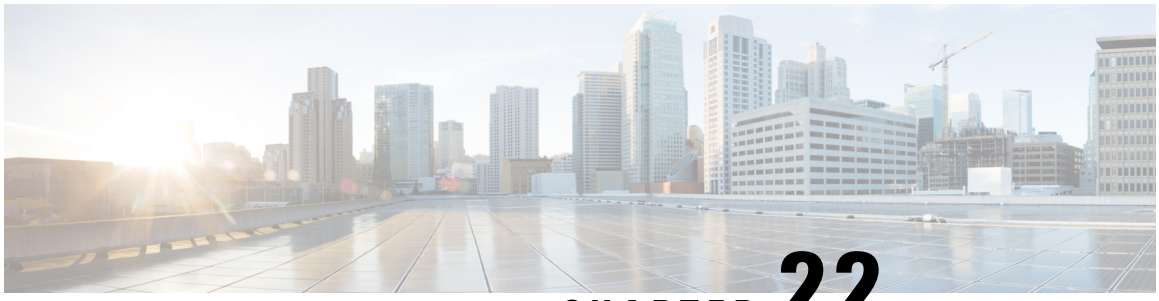

CHAPTER 22

L4-L7 Service Basic Workflow

- Layer 4-Layer 7 Service, on page 861

Layer 4-Layer 7 Service

Cisco DCNM Release 11.3(1) introduces the ability to insert Layer 4-Layer 7 (L4-L7) service devices in a data center fabric, and also enables selectively redirecting traffic to these service devices. You can add a service node, create route peering between the service node and the service leaf switch, and then selectively redirect traffic to these service nodes.

You can also watch a video that demonstrates how to orchestrate a L4-L7 Service Appliance with a VXLAN Fabric in a data center managed by Cisco DCNM. This demo covers provisioning, defining of service policies, and monitoring of redirected flows. For information, see Video: Service Redirection in Cisco DCNM.

Service Node

You have to create an external fabric and specify that a service node resides in that external fabric during service node creation. DCNM does not auto-detect or discover any service node. You also have to specify the service node name, type, and form factor. The name of the service node has to be unique within a fabric. The service node is attached to a leaf, border leaf, border spine, or a border super spine. Starting from Cisco DCNM Release 11.4(1), the service node can be attached to a vPC border gateway also. DCNM does not define a new switch role for a service leaf.

DCNM manages the switches that are attached to a service node. DCNM also manages the interfaces of these attached switches. Ensure that the interfaces to which the service node is attached to are in trunk mode and do not belong to any interface group. The L4-L7 service will not change its mode. In case the attached switches are forming a vPC pair, the name of the attached switch is a combination of both switches.

Route Peering

Route peering creates service networks. DCNM supports both static route and eBGP-based dynamic route peering options. After you specify the service network and select the peering policy for the tenant, DCNM automatically creates the service network under the specified tenant. Note that the terms, tenant and VRF, will be used interchangeably in this guide. If you select a route peering and click **Deploy** in the **Service Nodes** window, the L4-L7 service deploys the corresponding service network and VRF configuration to the leaf that is attached to the service node. Click **Preview** to review both the peering and service network configuration.

The automatically created service network will also be listed on the **Control > Fabrics > Networks** window. You can view and edit the corresponding config parameters in the **Networks** window. However, you cannot

delete the service network. Deletion of service networks is handled automatically during the service route peering deletion process. There can be multiple route peerings defined per tenant/VRF.

Service Policy

From DCNM 11.5(1), you can define service policies with any or arbitrary network and associate it with L3 routed interface on border switches. For more information, see PBR Support on WAN Interfaces of Border Switches. The L4-L7 service does not create any VRF or network other than the service networks that are defined during route peering. When you define the service policy between the created networks, the source and destination network can be a subnet, an individual IP address or the networks that are defined in the Control > Fabrics > Networks window. For intra-tenant firewall, one-arm and two-arm load balancer, the L4-L7 service in DCNM uses Policy-Based Routing (PBR) for service insertion. The inter-tenant firewall does not have a service policy. You only need to create a service node and route peering for inter-tenant firewall.

As the source and destination network can be attached or deployed independent of service policy deployment, the tenant/ VRF-related service policy configuration is only attached or pushed to the switch that is attached to the service node, and the source and destination network is updated with the service policy-related configuration. You can preview and confirm the generated configuration. By default, the service policy is defined but is not enabled or attached. You have to enable or attach the service policy to activate it.

The service configuration that is related to the source and destination network will be auto-processed when the source and destination networks are to be attached, or auto-updated in case the networks are already attached or deployed. By default, DCNM will collect statistics every 5 minutes and store it in ElasticSearch for aggregation and analysis. Click the graph line under **Stats** in the **Service Policy** tab of the **Service Nodes** window to view the historical time-based statistics. By default, the statistics are stored for a maximum of 7 days.

The service insertion is effective only on the flows to be created. There is no impact on any existing flows. Deletion of a network is not allowed in case an enabled service policy is associated with that network.

The L4-L7 service integration is built on top of the easy fabric policy enforcement. Use the fabric builder to create a VXLAN EVPN fabric and then import Cisco Nexus 9000 Series switches into the fabric with pre-defined fabric policies.

MSD Support

Starting from Cisco DCNM Release 11.4(1), this feature supports Multi-Site Domains (MSD). Select the MSD member fabric from the DCNM fabric scope selector, create a service node (for example, firewall, or load balancer), attach the service node to the switch in the selected MSD member fabric, define the route peering and service policies, and deploy relevant configurations on the selected MSD member fabric. For more information on the procedure to configure Layer 4-Layer 7 service, refer Configuring Layer 4-Layer 7 Service, on page 868.

RBAC Support

Starting from Cisco DCNM Release 11.4(1), the Layer 4-Layer 7 Service supports Role-Based Access Control (RBAC) along with fabric access mode.

The admin, stager, and operator, are pre-defined roles in DCNM. The table given below lists the various operations that each role can perform.

L4-L7 Service Operation	Service Node	Route Peering	Service Policy
Create/Update/Delete/Import	admin	admin, stager	admin, stager

L4-L7 Service Operation	Service Node	Route Peering	Service Policy
List/Export	admin, stager, operator	admin, stager, operator	admin, stager, operator
Attach/Detach	NA	admin, stager	admin, stager
Deploy	NA	admin (blocked if fabric is in fabric monitor or read-only mode)	admin (blocked if fabric is in fabric monitor or read-only mode)
Preview/Deployment History	NA	admin, stager, operator	admin, stager, operator

Note If a fabric is in fabric monitor or read-only mode, an admin cannot deploy the route peering or service policy. Also, the icon to delete the service node is not displayed if the external fabric where the service node is located is in fabric Monitor Mode. Remove the fabric from the fabric Monitor Mode to display the icon to delete the service node. This icon will be shown only to users with admin role access.

The Layer 4-Layer 7 Service windows are displayed based on the logged-in user's role and reflect the actions that the user is allowed to perform. Example screenshots of the Service Nodes window for an admin, stager, and operator role are as given below:

Figure 6: Admin role

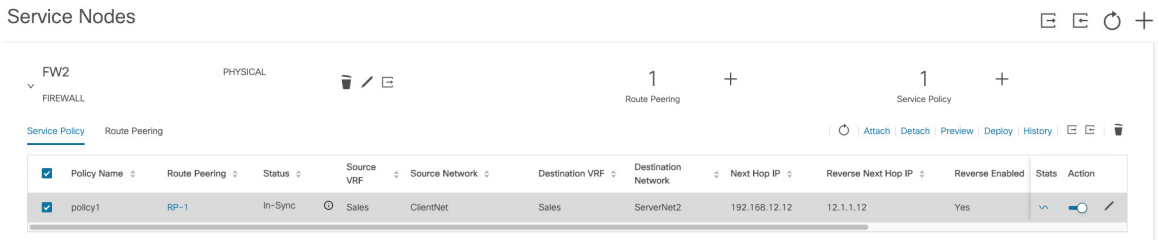

Figure 7: Stager role

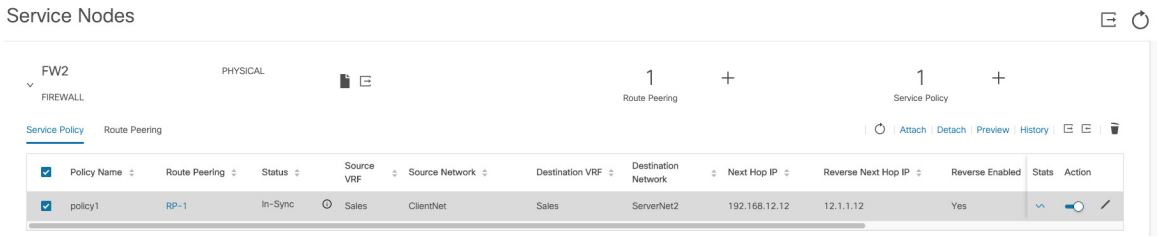

Figure 8: Operator role

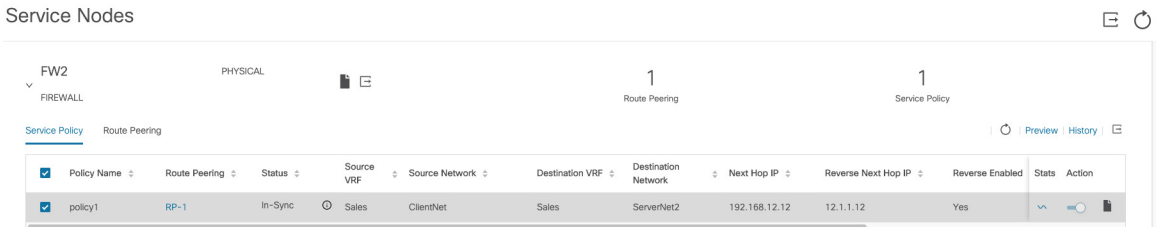

PBR Support on WAN Interfaces of Border Switches

In Cisco DCNM Release 11.4(1) and earlier releases, you have to manually associate a service policy with a specific switch interface by using a freeform configuration template to specify ‘any’ source or destination network during service policy creation. Starting from Cisco DCNM Release 11.5(1), you can specify an arbitrary network, that has not been defined in the top-down configuration, as a source or destination network in the service policy. This helps in streamlining policy enforcement for north-south traffic. The DCNM UI lists out routed Layer-3 interfaces of all border switches, standalone or vPC, that have a VRF association. You can then choose the required interface that has to be associated with the defined policy. The border switches include border leaf, border spine, border super spine and border gateway. There can be multiple interface associations. For example, multiple L3 interfaces, subinterfaces, and port-channels, can be selected for one border switch. You can also select multiple border switches for interface association. DCNM filters out the subinterfaces of Layer 3 port-channel as PBR is not supported with Layer 3 port-channel subinterfaces. For information, see NX-OS Unicast Routing Configuration Guide.

Depending on the policy direction, the border switch and interface association for ‘any’ or arbitrary network may not be needed. For example, for a forwarding policy, the border switch and interface input or route-map association is not needed for ‘any’ or arbitrary destination network. For a reversed policy, the border switch and interface or route-map association is not needed for ‘any’ or arbitrary source network.

When the policy with ‘any’ or arbitrary network is attached, the policy related CLIs are generated and associated with the selected L3 routed interfaces of the border switches. The deployment of that policy pushes the CLIs to the selected border switches. The deployment history will include the corresponding entries and can be quickly accessed using VRF filtering. The service policy stats diagram includes the PBR stats of route maps that are associated with the selected L3 routed interfaces of the border switches.

Static Route

On Cisco DCNM Release 11.4(1) and earlier releases, static routes are deployed only on the service leaf switches when static route peering is used. Starting from Cisco DCNM Release 11.5(1), the Layer 4-Layer 7 Service pushes static routes on all VTEPs, including service leaf switches, where the VRF being referenced in the static route is attached. This expedites service node failover with static routes.

Guidelines and Limitations for Layer 4-Layer 7 Service

- L4-L7 service in DCNM does not manage or provision service nodes, such as firewall and load balancer.
- The L4-L7 service feature is supported only on the VXLAN BGP EVPN fabrics with the **Easy_Fabric_11_1** template.
- The service policies defined in this feature leverage Policy-Based Routing (PBR). Refer Nexus 9000 Series NX-OS Unicast Routing Configuration Guide for PBR related configuration, constraints, and so on.

- This feature supports Cisco Nexus 9300-EX and 9300-FX platform switches as leaf, border leaf, border spine, border super spine, and border gateway switches.
- Configurations involving intra-tenant and inter-tenant firewall for L3 networks, and one-arm and two-arm deployed load balancers, are supported.
- The existing DCNM topology view is also leveraged to display redirected flows associated with the switches that the service node is attached to, and to locate specific redirected flows.
- From Cisco DCNM Release 11.5(1), one-arm Virtual Network Function is supported.
- From Cisco DCNM Release 11.5(1), Layer 4-Layer 7 Service REST APIs are accessible via DCNM packaged REST API documentation. For more information, refer Cisco DCNM REST API Reference Guide, Release 11.5(1).
- Load sharing is not supported.
- This feature creates, updates, and deletes the service network, as required. Service networks cannot be created or deleted from the **Control** > **Fabrics** > **Networks** window.

Types of Layer 4–Layer 7 Service Devices

The L4-L7 service in Cisco DCNM supports any vendors service node attachments. Typical service node types that are deployed in a data center are Firewalls, Load Balancers, and other Layer-4 to Layer-7 products.

Examples of supported Firewall vendors are Cisco Systems, Palo Alto Networks, Fortinet, Check Point Software Technologies, and others.

Examples of supported Load Balancer vendors are F5 Networks, Citrix Systems, A10 Networks, and others.

Note that these example lists are meant to serve as examples and not intended to be **exhaustive** lists. The L4-L7 service attachment is generic and applies to any vendors service node.

Configuring Fabric Settings for Layer 4-Layer 7 Service

Certain fabric settings have to be configured to enable L4-L7 service functionality. To configure these settings, click **Fabric Settings** under **Actions** in the **Fabric Builder** window.

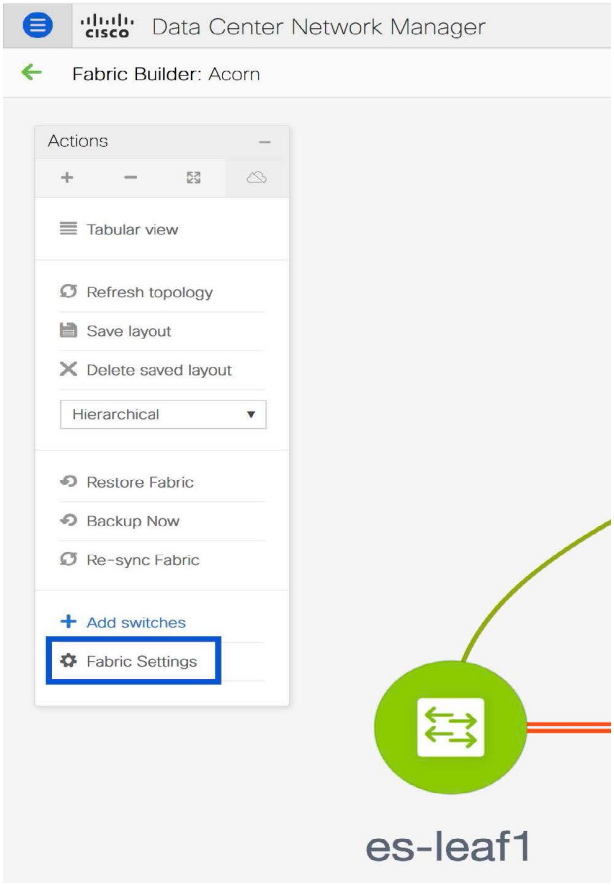

The **Edit Fabric** window is displayed. Click **Advanced**. Select the **Enable Policy-Based Routing (PBR)** checkbox to enable routing of packets based on the specified policy.

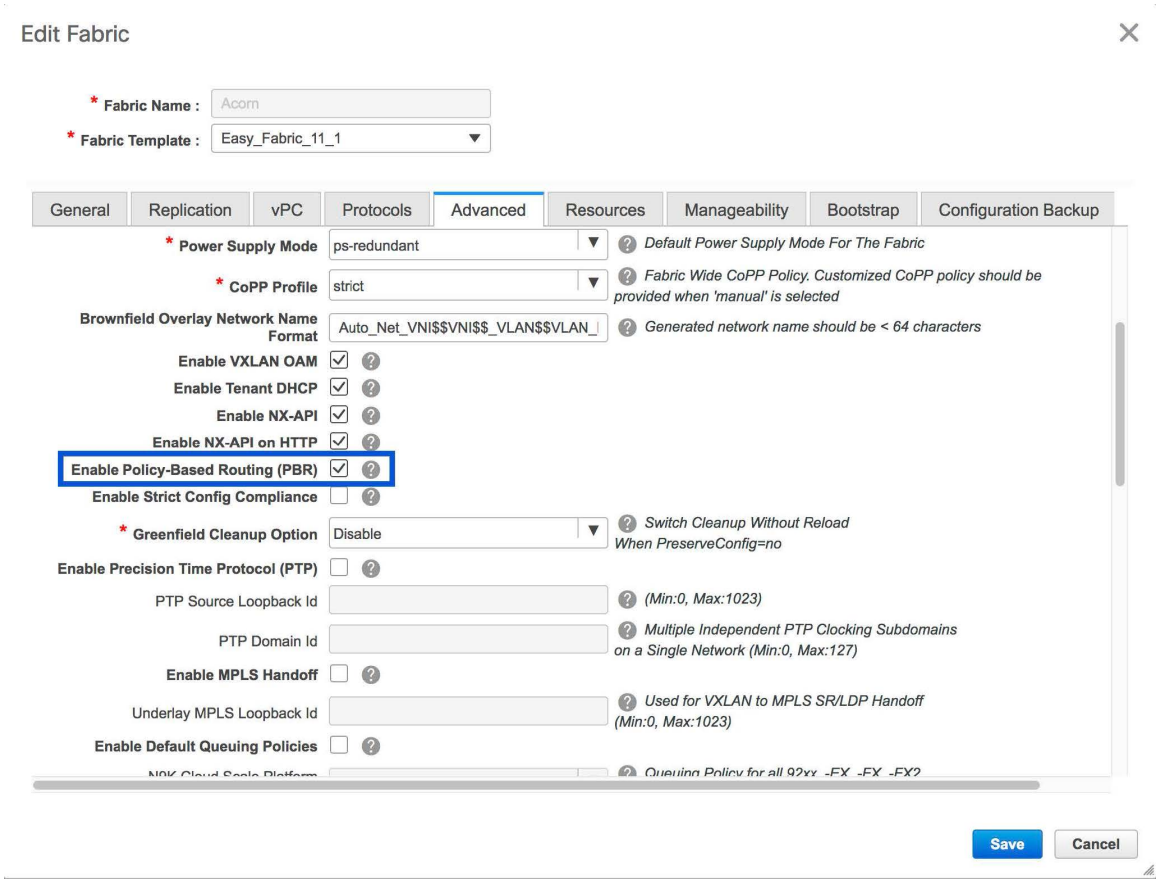

Now, click **Resources**. Specify a VLAN range in the **Service Network VLAN Range** field. This is a per switch overlay service network VLAN range. The minimum allowed value is 2 and the maximum allowed value is 3967. Also, specify a value for the **Route Map Sequence Number Range** field. The minimum allowed value is 1 and the maximum allowed value is 65535. Click **Save and Deploy** to deploy the updated configuration.

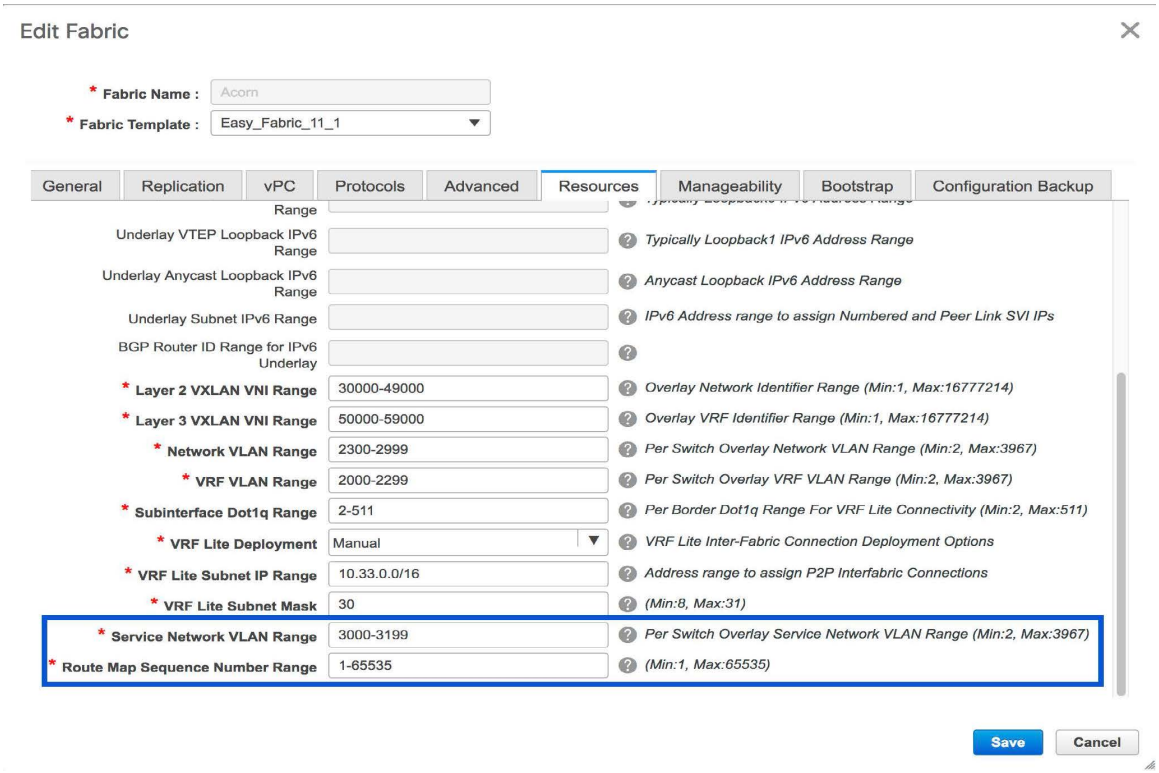

Configuring Layer 4-Layer 7 Service

To launch the L4-L7 Service, or the Elastic Service, on the Cisco DCNM Web UI, choose **Control>Fabrics>Services**.

The **Service Nodes** window is displayed. Select a valid switch fabric to display or define the service nodes, route peerings, and service policies, in that fabric.

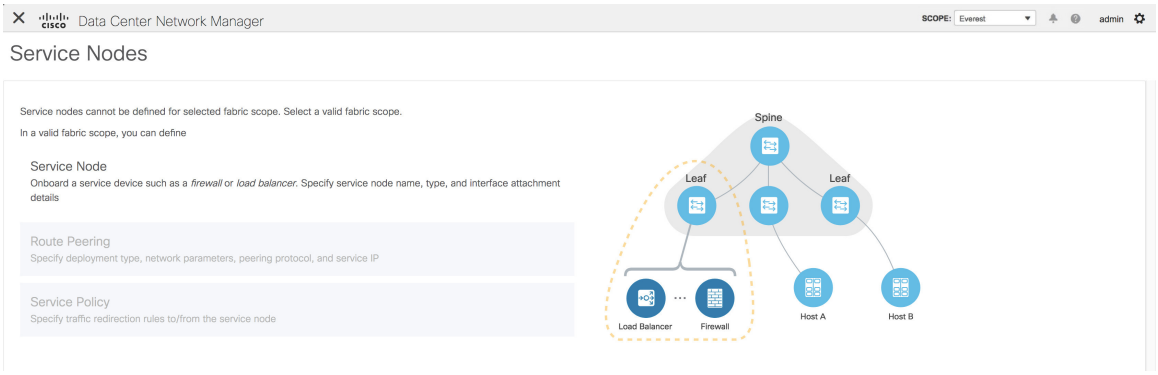

Note From Cisco DCNM Release 11.5(1), service nodes, route peering, and service policies updated within the last 15 minutes are highlighted.

The L4-L7 service configuration procedure consists of the following steps:

Create Service Node

To create a service node, click the + icon at the top right of the **Service Nodes** window to display the **New Service Nodes** window.

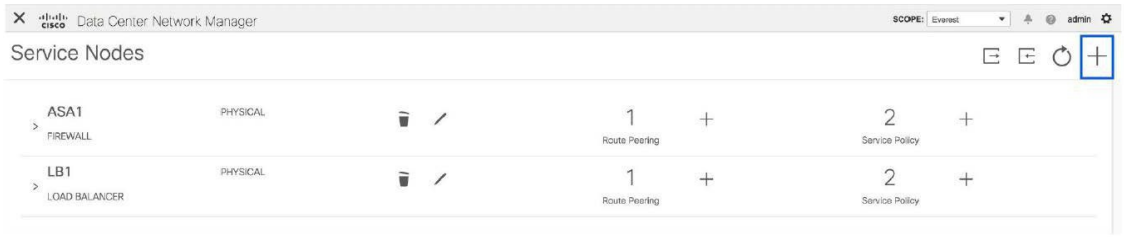

The **New Service Nodes** window has three steps, **Create Service Node, Create Route Peering** and **Create Service Policy**.

The **Create Service Node** window has two sections - Create Service Node and Switch Attachment, followed by a **Link Template** drop-down list. You can select service_link_trunk, service_link_port_channel_trunk and service_link_vpc from this drop-down list..

Figure 9: Example: Link Template - service_link_trunk

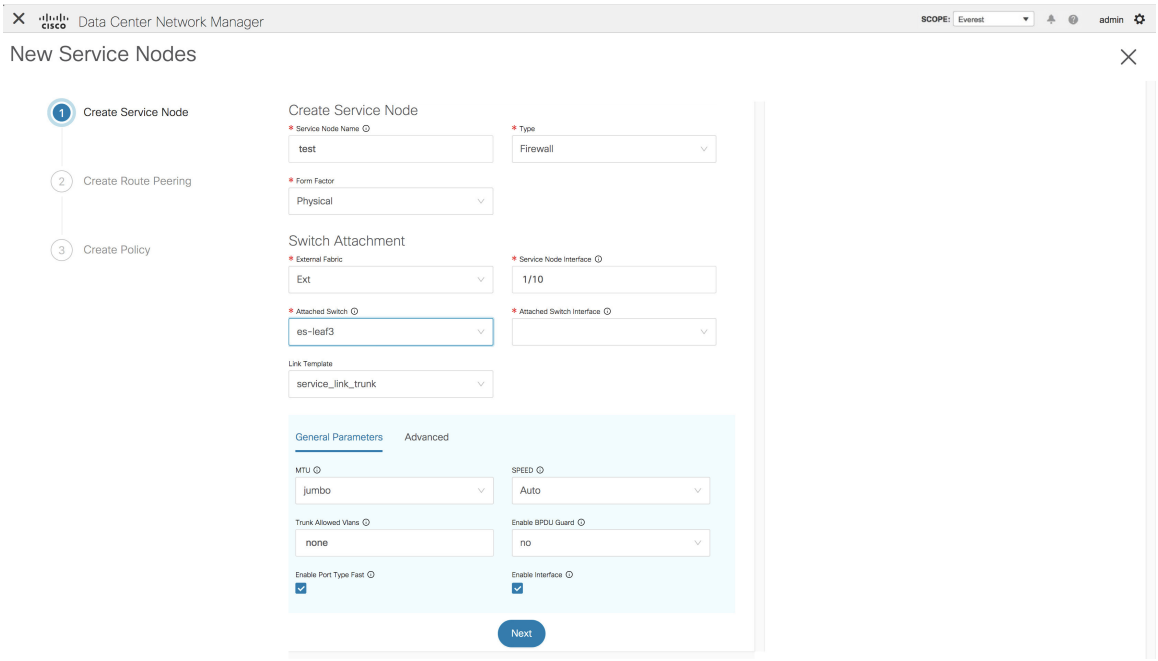

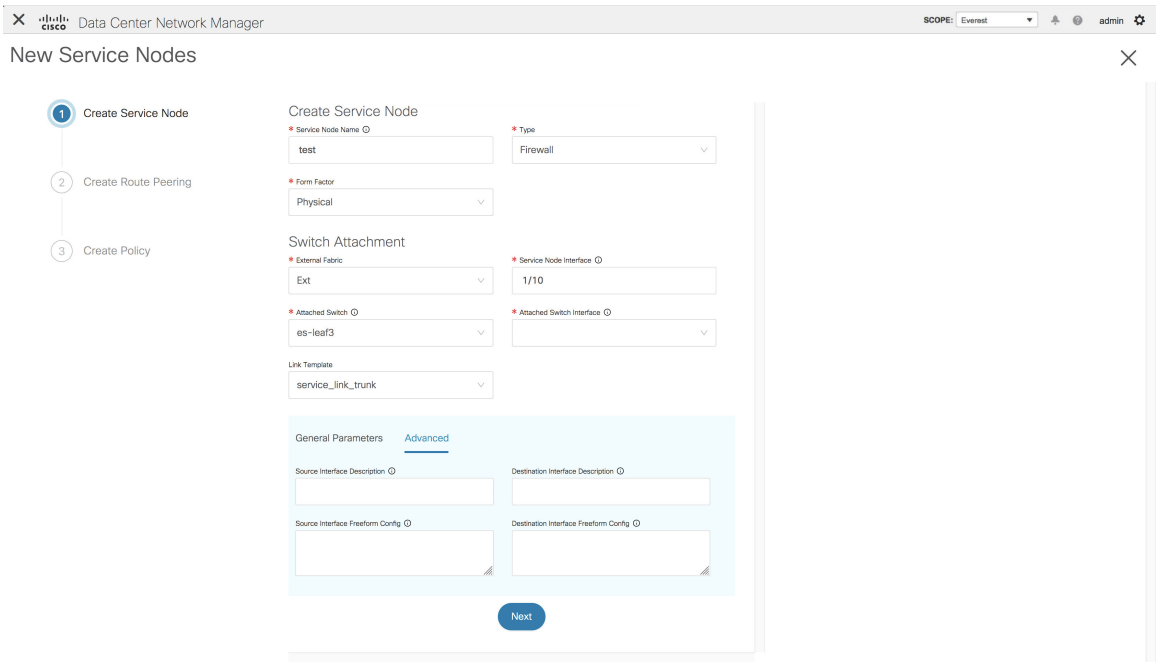

Figure 10: Example: Link Template - service_link_port_channel_trunk

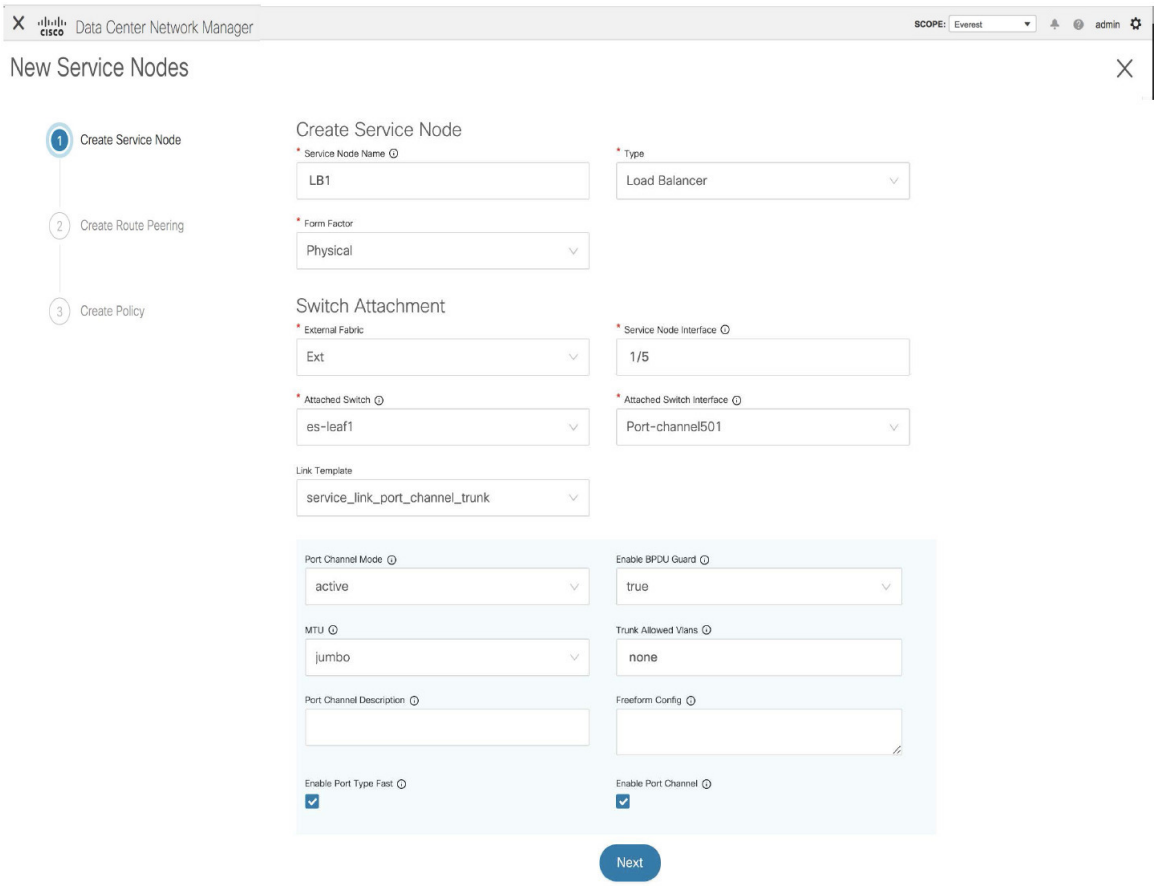

Figure 11: Example: Link Template - service_link_vpc

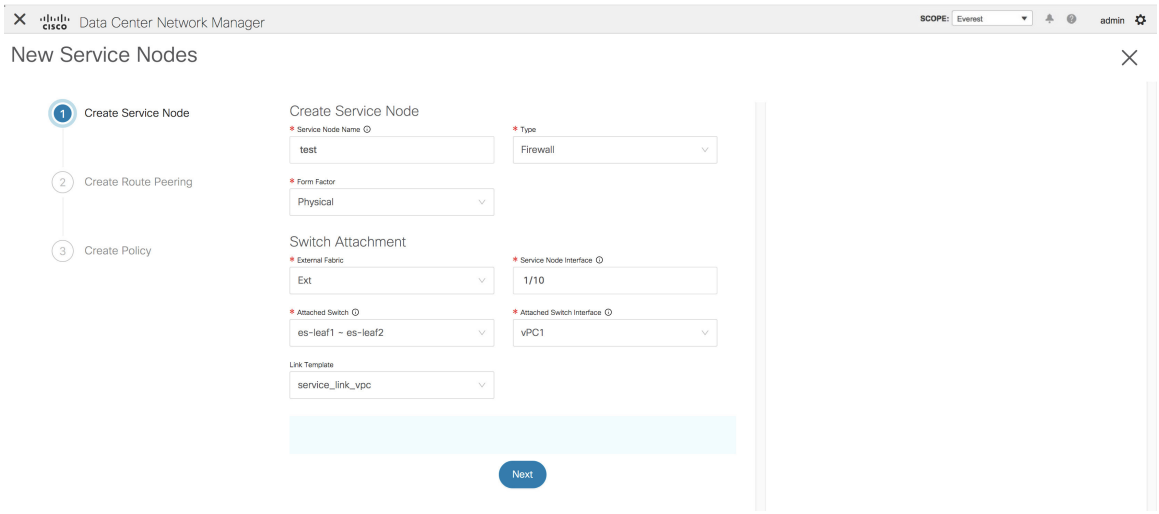

Figure 12: Example: Type - Virtual Network Function

Note From DCNM Release 11.5(1), one-arm Virtual Network Function is supported.

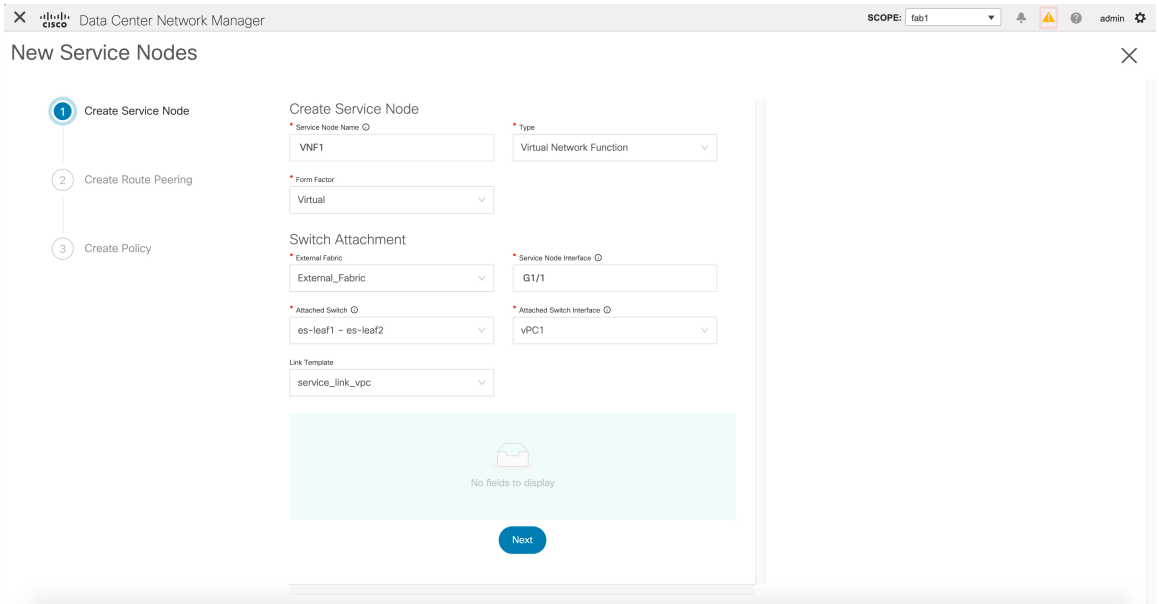

The fields in the **Create Service Node** window are as given below. It is mandatory to fill the fields marked with an asterisk. For more information on the fields in this window, hover over the **i** icon.

Create Service Node

Service Node Name - Enter a name for the service node. The name can have alphanumeric, underscore, or dash characters.

Type - Select Firewall or Load Balancer.

Form Factor - Select Physical or Virtual.

Switch Attachment

External Fabric - Specify the external fabric.

Service Node Interface - Specify the service node interface.

Attached Switch- Select a switch from the drop-down list.

Attached Switch Interface - Select the interface from the drop-down list. In case the vPC pair is selected from the **Attached Leaf Switch** drop-down list, the vPC channel will be shown in the **Attached Leaf Switch Interface** drop-down list. Otherwise, the port-channel and interfaces with trunk mode are shown in the **Attached Leaf Switch Interface** drop-down list.

Link Template - Select the service_link_trunk, service_link_port_channel_trunk, or the service_link_vpc template. For more information on template fields, refer Templates.

Now, click **Next**. A pop-up window is displayed stating that a new service node has been created successfully and the **Create Route Peering** window is displayed.

Create Route Peering

The fields that appear in the **Create Route Peering** window depend on the type of deployment chosen in the **Create Service Node** window. Depending on the type chosen (Firewall or Load Balancer), the types of deployments are Intra-Tenant Firewall, Inter-Tenant Firewall, One-Arm load balancer and Two-Arm load balancer.

Note Deletion of service network is not allowed on the **Control** > **Fabrics** > **Networks** window.

Example: Intra-Tenant Firewall Deployment

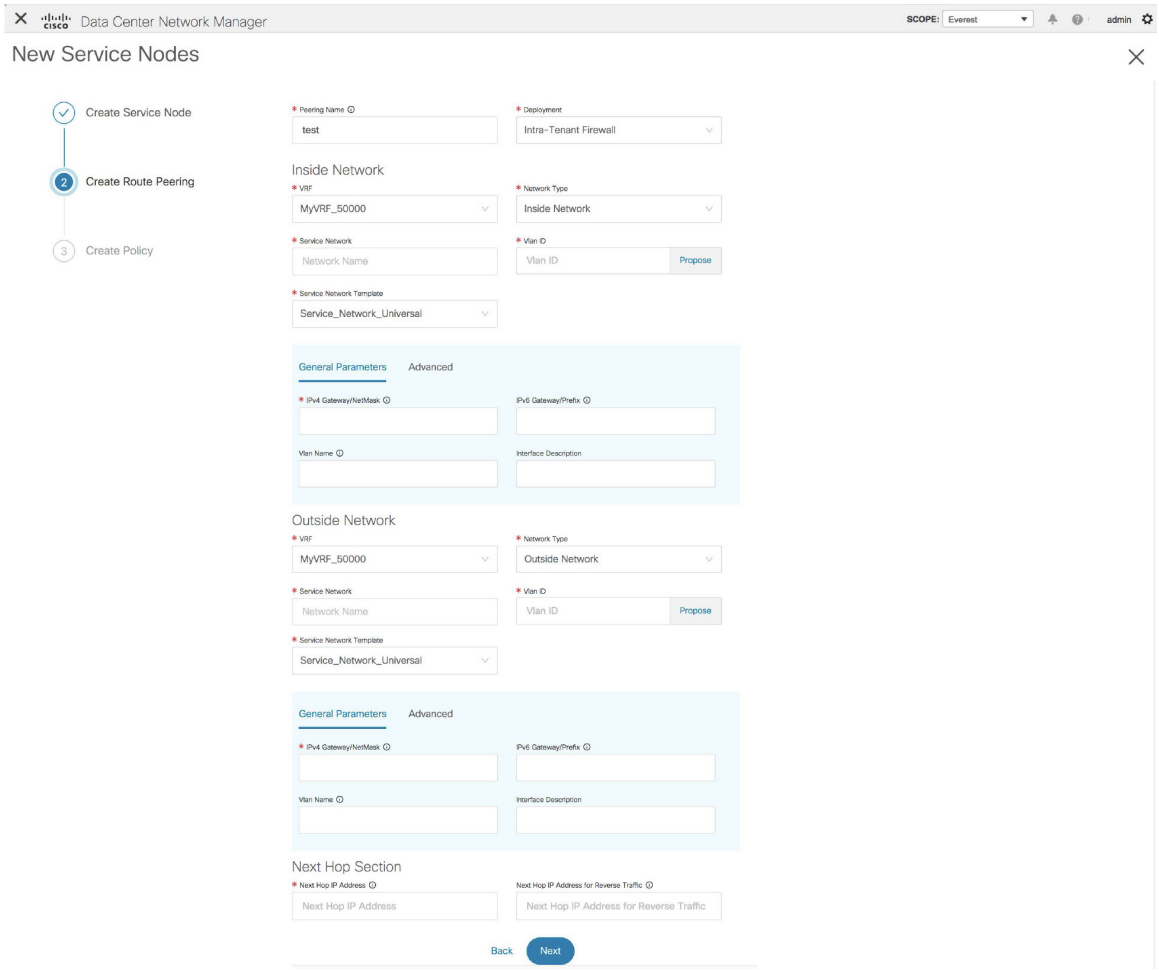

The fields in the **Create Route Peering** window for an Intra-Tenant Firewall deployment are as given below. It is mandatory to fill the fields marked with an asterisk. For more information on the fields in this window, hover over the **i** icon.

Peering Name - Specify a name for the peering. The name can have alphanumeric, underscore, or dash characters.

Deployment - Select Intra-Tenant Firewall.

Inside Network

VRF - Specify the VRF.

Network Type - Select Inside Network.

Service Network - Specify the name of the service network.

Vlan ID - Specify the VLAN ID. Valid IDs range from 2 to 3967. Click Propose to retrieve a value from the pre-defined service network VLAN range pool.

Service Network Template - Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates.

Outside Network

VRF - Specify the VRF.

Network Type - Select Outside Network.

Service Network - Specify the name of the service network.

Vlan ID - Specify the VLAN ID. Valid IDs range from 2 to 3967. Click Propose to retrieve a value from the pre-defined service network VLAN range pool.

Service Network Template - Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates.

Next Hop Section

Next Hop IP Address - Specify the next-hop IP address. This is the IP/VIP of the service node used for traffic redirection.

Next Hop IP Address for Reverse Traffic - Specify the next-hop IP address for reverse traffic. This is the IP/VIP of the service node used for traffic redirection.

Example: Inter-Tenant Firewall Deployment

Peering Option - Static Peering, Inside Network Peering Template - service_static_route, Outside Network Peering Template - service_static_route

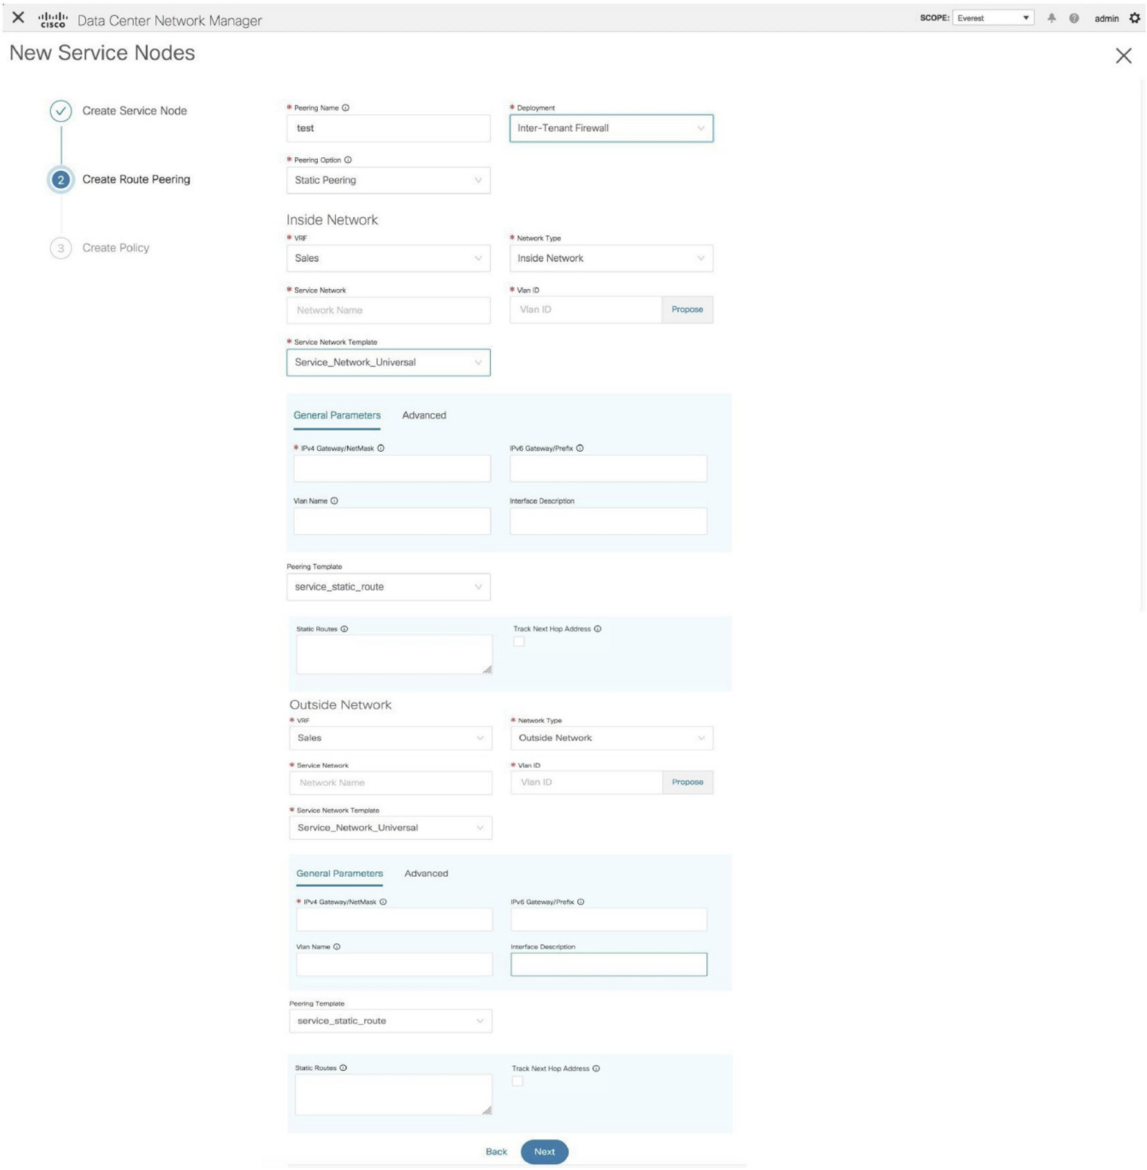

The fields in the **Create Route Peering** window for an Inter-Tenant Firewall deployment are as given below. It is mandatory to fill the fields marked with an asterisk.

Peering Name - Specify a name for the peering. The name can have alphanumeric, underscore, or dash characters.

Deployment - Select Inter-Tenant Firewall.

Peering Option - Select Static Peering or eBGP Dynamic Peering.

Inside Network

VRF - Select a VRF from the drop-down list..

Network Type - Select Inside Network.

Service Network - Select a service network name from the drop-down list.

Vlan ID - Specify the VLAN ID. Valid IDs range from 2 to 3967. Click **Propose** to retrieve a value from the pre-defined service network VLAN range pool.

Service Network Template - Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates.

Peering Template - Select service_static_route or service_ebgp_route from the drop-down list. For more information on the template fields, refer Templates.

Outside Network

VRF - Select a VRF from the drop-down list..

Network Type - Select Outside Network.

Service Network - Select a service network name from the drop-down list.

Vlan ID - Specify the VLAN ID. Valid IDs range from 2 to 3967. Click **Propose** to retrieve a value from the pre-defined service network VLAN range pool.

Service Network Template - Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates.

Peering Template - Select service_static_route or service_ebgp_route from the drop-down list. For more information on the template fields, refer Templates.

Example: One-Arm Mode Load Balancer

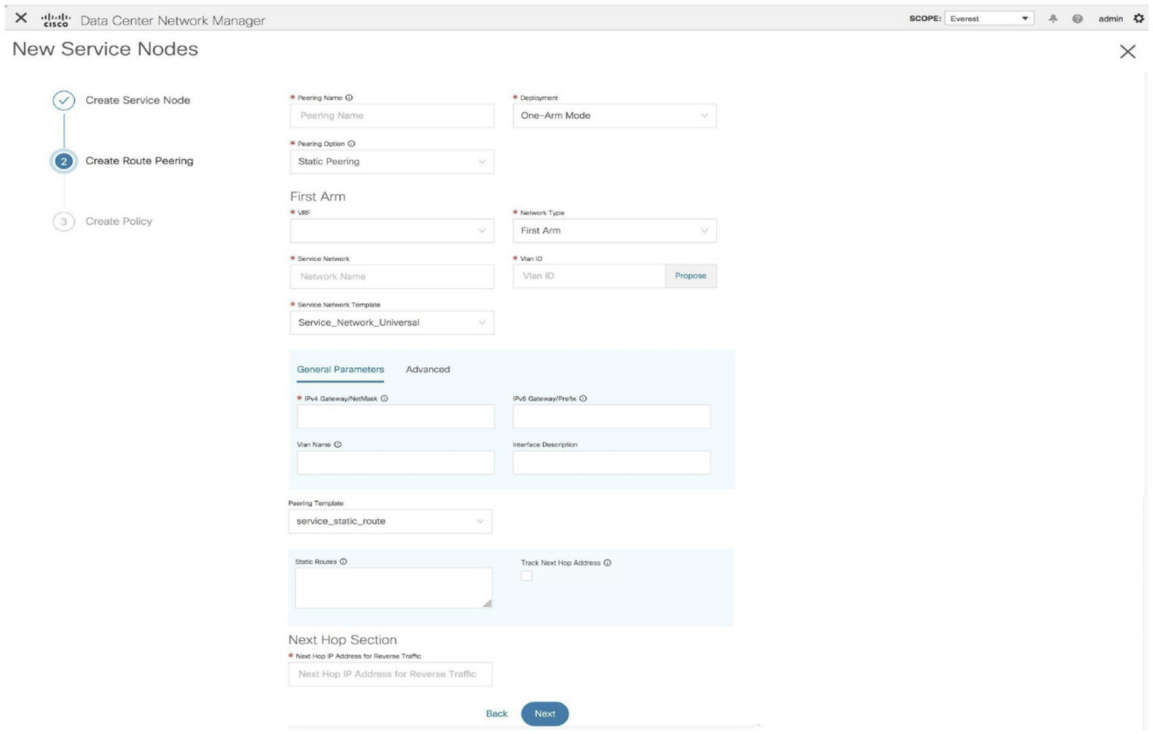

The fields in the **Create Route Peering** window for a One-Arm Mode load balancer deployment are as given below. It is mandatory to fill the fields marked with an asterisk.

Peering Name - Specify a name for the peering. The name can have alphanumeric, underscore, or dash characters.

Deployment - Select One-Arm Mode.

Peering Option - Select Static Peering or eBGP Dynamic Peering.

First Arm

VRF - Select a VRF from the drop-down list..

Network Type - Select First Arm.

Service Network - Select a service network name from the drop-down list.

Vlan ID - Specify the VLAN ID. Valid IDs range from 2 to 3967. Click **Propose** to retrieve a value from the pre-defined service network VLAN range pool.

Service Network Template - Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates.

Peering Template - Select service_static_route or service_ebgp_route from the drop-down list. For more information on the template fields, refer Templates.

Next Hop IP Address for Reverse Traffic - Specify the next-hop IP address for reverse traffic.

Example: Two-Arm Mode Load Balancer

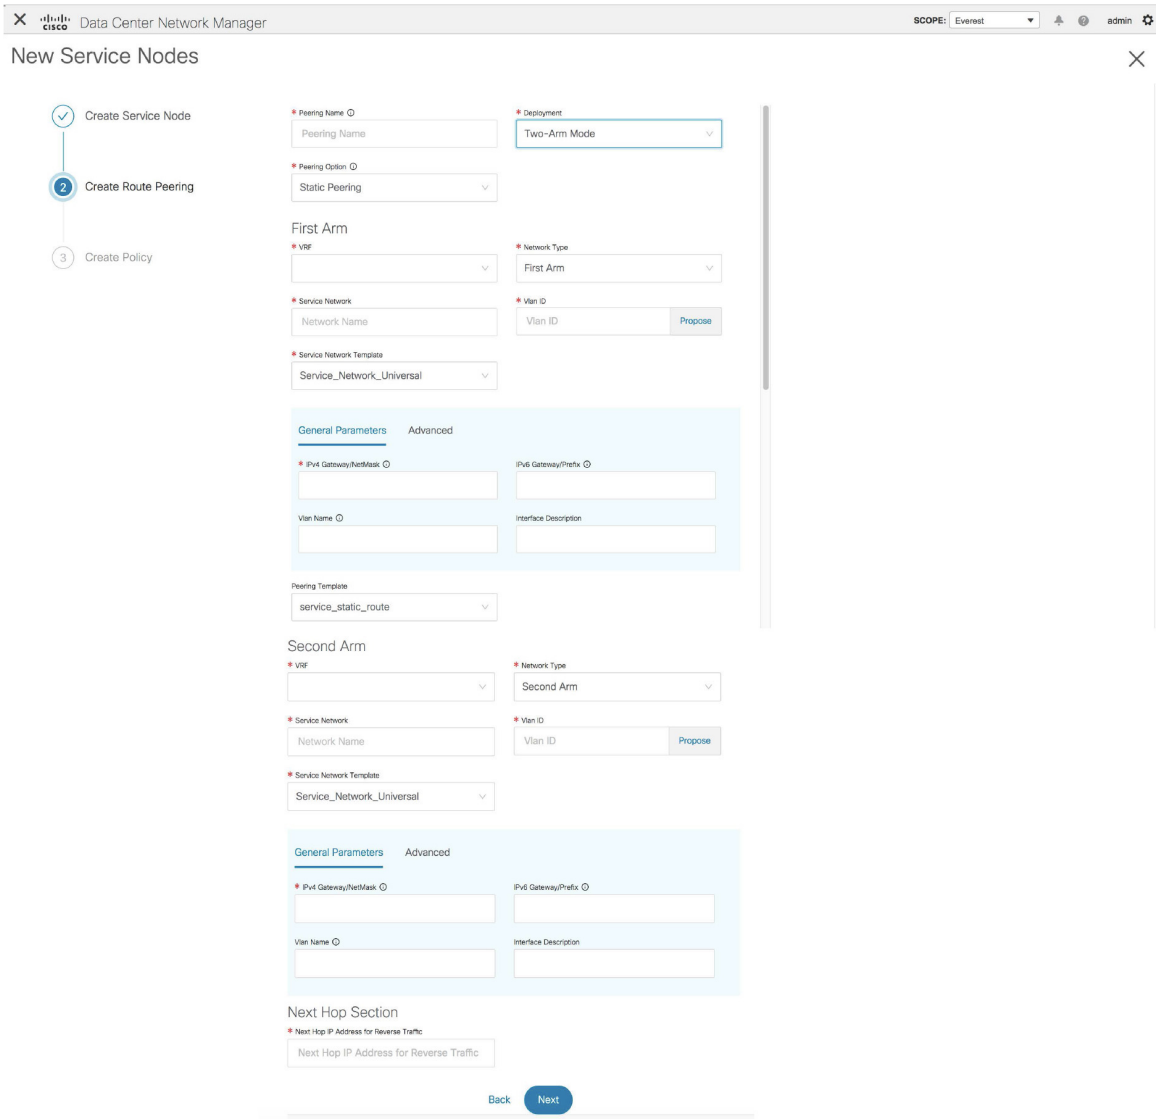

The fields in the Create Route Peering window for a Two-Arm Mode load balancer deployment are as given below. It is mandatory to fill the fields marked with an asterisk.

Peering Name - Specify a name for the peering. The name can have alphanumeric, underscore, or dash characters.

Deployment - Select Two-Arm Mode.

Peering Option - Select Static Peering or eBGP Dynamic Peering.

First Arm

VRF - Select a VRF from the drop-down list..

Network Type - Select First Arm.

Service Network - Select a service network name from the drop-down list.

Vlan ID - Specify the VLAN ID. Valid IDs range from 2 to 3967. Click **Propose** to retrieve a value from the pre-defined service network VLAN range pool.

Service Network Template - Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates.

Peering Template - Select service_static_route or service_ebgp_route from the drop-down list. For more information on the template fields, refer Templates.

Second Arm

VRF - Select a VRF from the drop-down list..

Network Type - Select Second Arm.

Service Network - Specify the name of the service network.

Vlan ID - Specify the VLAN ID. Valid IDs range from 2 to 3967. Click **Propose** to retrieve a value from the pre-defined service network VLAN range pool.

Service Network Template - Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates.

Next Hop Section

Next Hop IP Address for Reverse Traffic - Specify the next-hop IP address for reverse traffic.

Now, click **Next**. The **Create Policy** window is displayed.

Example: One-Arm Virtual Network Function

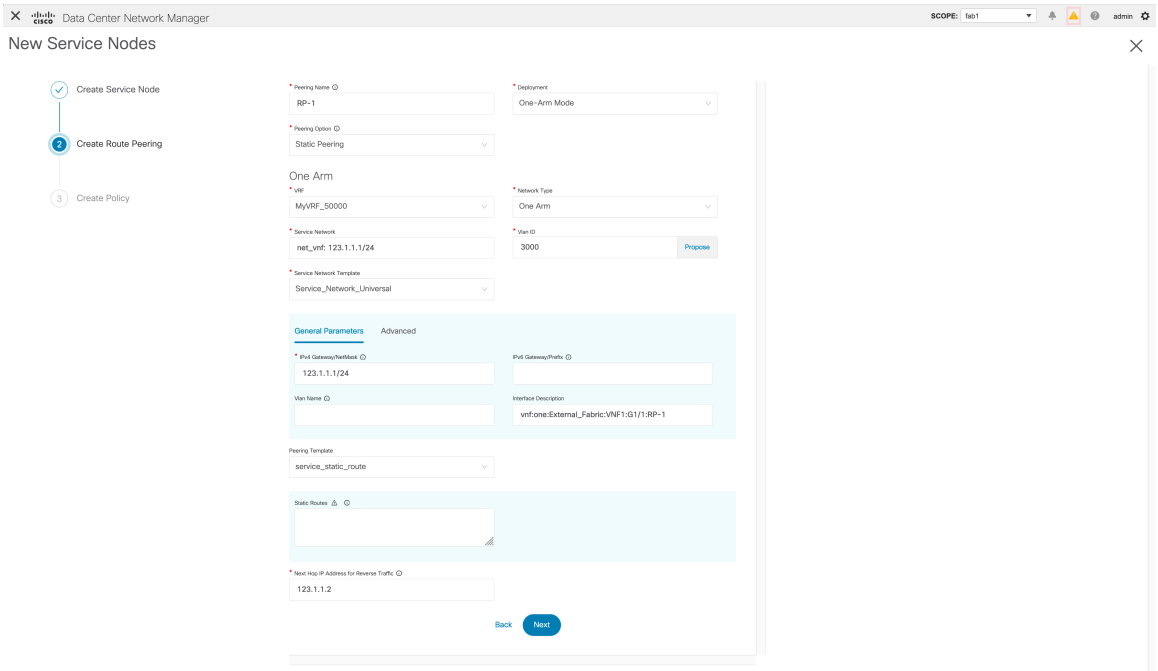

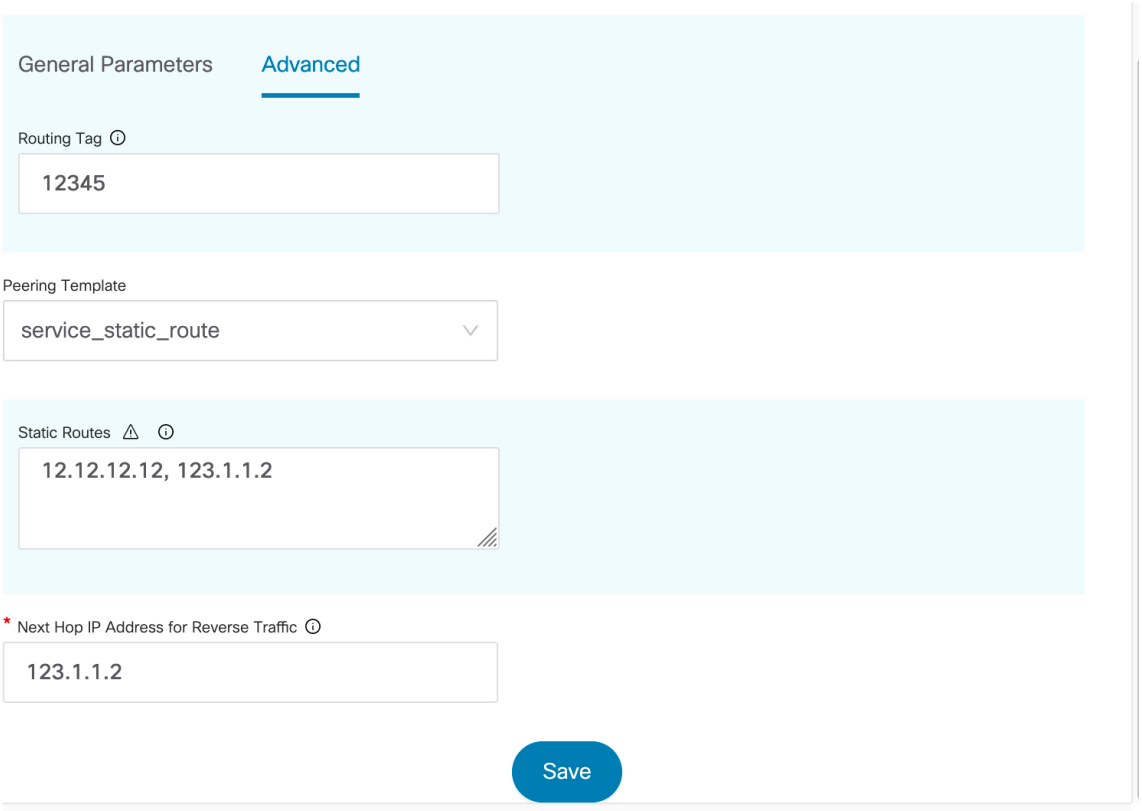

The fields in the Create Route Peering window for a One-Arm Mode Virtual Network Function deployment are as given below. It is mandatory to fill the fields marked with an asterisk.

Peering Name - Specify a name for the peering. The name can have alphanumeric, underscore, or dash characters.

Deployment - Select One-Arm Mode.

Peering Option - Select Static Peering or eBGP Dynamic Peering.

One Arm

VRF - Select a VRF from the drop-down list..

Network Type - Select One Arm.

Service Network - Select a service network name from the drop-down list.

Vlan ID - Specify the VLAN ID. Valid IDs range from 2 to 3967. Click Propose to retrieve a value from the pre-defined service network VLAN range pool.

Service Network Template - Select the Service_Network_Universal template from the drop-down list. For more information on the template fields, refer Templates.

IPv4 Gateway/Netmask - Specify the IPv4 gateway and netmask.

Peering Template - Select service_static_route or service_ebgp_route from the drop-down list. For more information on the template fields, refer Templates.

Next Hop IP Address for Reverse Traffic - Specify the next-hop IP address for reverse traffic.

Now, click Next. The Create Policy window is displayed.

Create Service Policy

The **Create Policy** window is displayed as given below.

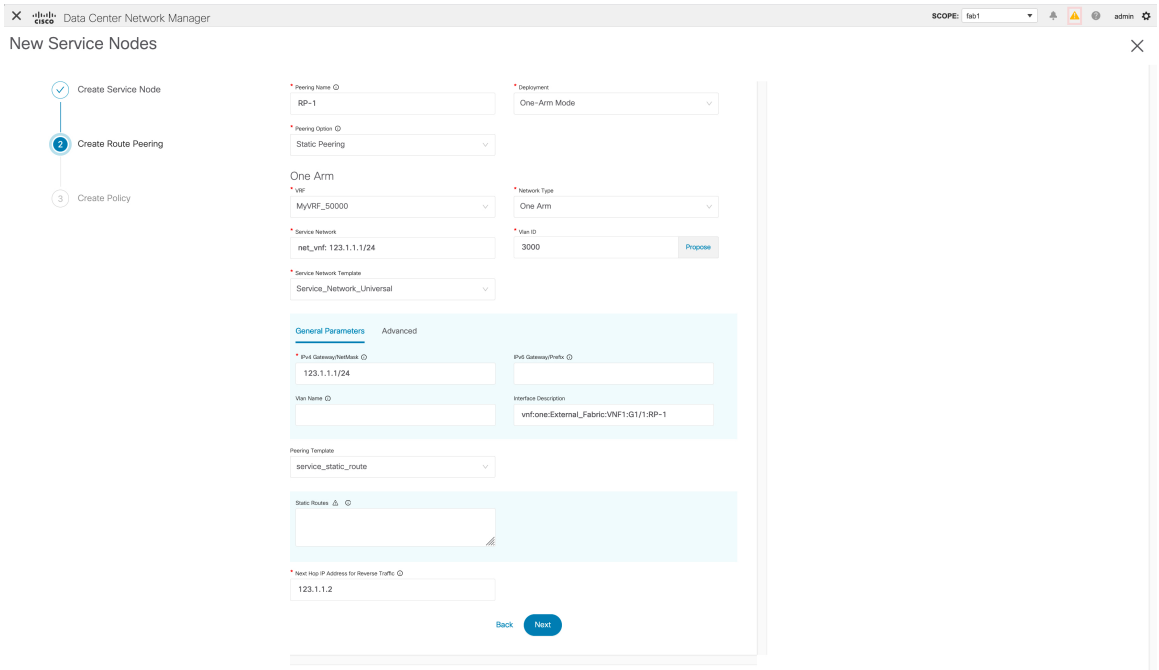

The fields in the **Create Policy** window are as given below. It is mandatory to fill the fields marked with an asterisk.

Policy Name - Specify a name for the policy.

Peering Name - Select a peering option from the drop-down list.

Source VRF Name - Select a source VRF from the drop-down list.

Destination VRF Name - Select a destination VRF from the drop-down list.

Source Network - Select an IP address from the drop-down list.

Destination Network - Select an IP address from the drop-down list.

Reverse Next Hop IP Address - The reverse next-hop IP address is displayed.

Policy Template Name - Select a template from the drop-down list. For more information on the template fields, refer Templates.

General Parameters

Protocol - Select a protocol from the drop-down list. The options are icmp, ip, tcp, and udp.

Source Port - Specify a source port number. In case the ip protocol is selected, this value is ignored.

Destination Port - Specify a destination port number. In case the ip protocol is selected, this value is ignored.

Starting from Cisco DCNM Release 11.4(1), the **Advanced** tab has been introduced. The options in this tab allow you to customize the matched traffic redirection. For example, you can specify matched traffic to be redirected using PBR, or for matched traffic to bypass a firewall and use routing table rules instead, or you can specify that any matched traffic has to be dropped. You can choose to override the route map match

sequence number for prioritization. You can also customize the ACL name, however ensure that the ACL name that you specify is unique and the same name is not used for another ACL. If you do not specify the route map match sequence number or ACL name, the sequence number will be auto-populated, as in Cisco DCNM Release 11.3(1), from the designated resource pool and the ACL name will be auto-generated based on 5-tuples. For more information on the fields in the **Advanced** tab, refer Templates.

Click **Create**. The service policy is created.

Note Deletion of any service network in Top-Down provisioning that is used by Services is not allowed. Deletion of any regular network that is used in a service policy is also not allowed.

Templates

Service Node Link Templates

service_link_trunk

General Parameters tab

MTU - Specifies the MTU for the interface. By default, this is set to jumbo.

SPEED - Specifies the speed of the interface. By default, this is set to Auto. You can change it to 100Mb, 1Gb, 10GB, 25Gb, 40Gb, or 100Gb, as required.

Trunk Allowed Vlans - Specify 'none',' all' or VLAN ranges. By default, none is specified.

Enable BPDU Guard - Specify an option from the drop-down list. The available options are true, false or no.

Enable Port Type Fast - Select the checkbox to enable spanning tree edge port behavior. By default, this is enabled.

Enable Interface - Uncheck the checkbox to disable the interface. By default, the interface is enabled.

Advanced tab

Source Interface Description - Enter a description for the source interface.

Destination Interface Description - Enter a description for the destination interface.

Source Interface Freeform Config - Enter any addition CLI for the source interface.

Destination Interface Freeform Config - Enter any addition CLI for the destination interface.

service_link_port_channel_trunk

Port Channel Mode - Select a port channel mode from the drop-down list. By default, active is specified.

Enable BPDU Guard - Specify an option from the drop-down list. The available options are true, false or no.

MTU - Specifies the MTU for the interface. By default, this is set to jumbo.

Trunk Allowed Vlans - Specify 'none',' all' or VLAN ranges. By default, none is specified.

Port Channel Description - Enter a description for the port channel.

Freeform Config - Specify the required freeform configuration CLIs.

Enable Port Type Fast - Select the checkbox to enable spanning tree edge port behavior. By default, this is enabled.

Enable Port Channel - Select the checkbox to enable the port channel. By default, this is enabled.

service_link_vpc

This template has no specifiable parameters.

Route Peering Service Network Template

Service_Network_Universal

General Parameters tab

IPv4 Gateway/Netmask - Specify the gateway IP address and mask of the service network.

IPv6 Gateway/Prefix - Specify the gateway IPv6 address and prefix of the service network.

Vlan Name - Specify a name for the VLAN.

Interface Description - Enter a description for the interface

Advanced tab

Routing Tag - Specify a routing tag. Valid values range from 0 to 4294967295.

Route Peering Templates

service_static_route

Enter the static routes in the **Static Routes** field. You can enter one static route per line.

service_ebgp_route

General Parameters tab

Neighbor IPv4 - Specify the IPv4 address of the neighbor.

Loopback IP - Specify the IP address of the loopback.

Advanced tab

Neighbor IPv6 - Specify the IPv6 address of the neighbor.

Loopback IPv6 - Specify the IPv6 address of the loopback.

Route-Map TAG - Specify route-map tag that is associated with the interface ID.

Interface Description - Enter a description for the interface.

Local ASN - Specify a local ASN to override the system ASN.

Advertise Host Routes - Select the checkbox to enable advertisement of /32 and /128 routes to edge routers.

Enable Interface - Uncheck the checkbox to disable the interface. By default, the interface is enabled.

Service Policy Template

service_pbr

General Parameters tab

Protocol - Select a protocol from the drop-down list. The options are icmp, ip, tcp, and udp.

Source port - Specify a source port number. In case the ip protocol is selected, this value is ignored.

Destination port - Specify a destination port number. In case the ip protocol is selected, this value is ignored.

Advanced tab

Route Map Action - Select an action from the drop-down list. The options are permit or deny. If you select **permit**, the matched traffic is redirected based on the next-hop option and the defined policy. If you select **deny**, the traffic is routed based on the routing table rules.

Next Hop Option - Specify an option for the next-hop. The options are **none**, **drop-on-fail**, and **drop**.If you select **none**, the matched traffic is redirected based on the defined PBR rules. If you select **drop-on-fail**, the matched traffic is dropped if the specified next hop is not reachable. If you select **drop**, the matched traffic is dropped.

ACL Name - Specify a name for the generated access control list (ACL). If not specified, this is auto-generated.

ACL Name for reversed traffic - Specify a name for the ACL that is generated for reversed traffic. If not specified, this is auto-generated.

Route map match number - Specify a route map match number. A valid value ranges from 1 to 65535. If not specified, a route map match sequence number will be retrieved from the predefined resource pool. This number is associated with the name of the ACL.

Route map match number for reversed traffic - Specify a route map match number for reversed traffic. A valid value ranges from 1 to 65535. If not specified, a route map match sequence number will be retrieved from the predefined resource pool. This number is associated with the name of the ACL that has been generated for reversed traffic.

You can also customize the templates based on specific requirements. For more information on templates, refer Template Library.

Adding a Route Peering

To add a route peering from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Click the **Add Route Peering** icon on the **Service Nodes** window.

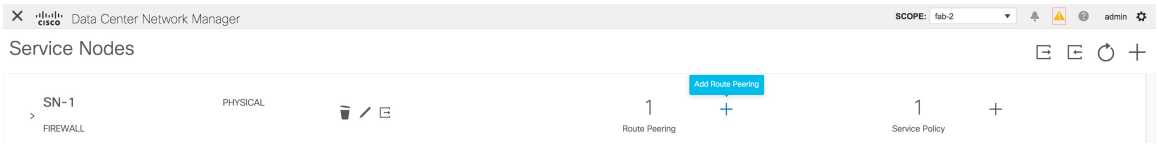

Step 2 The **Add Route Peering** window is displayed.

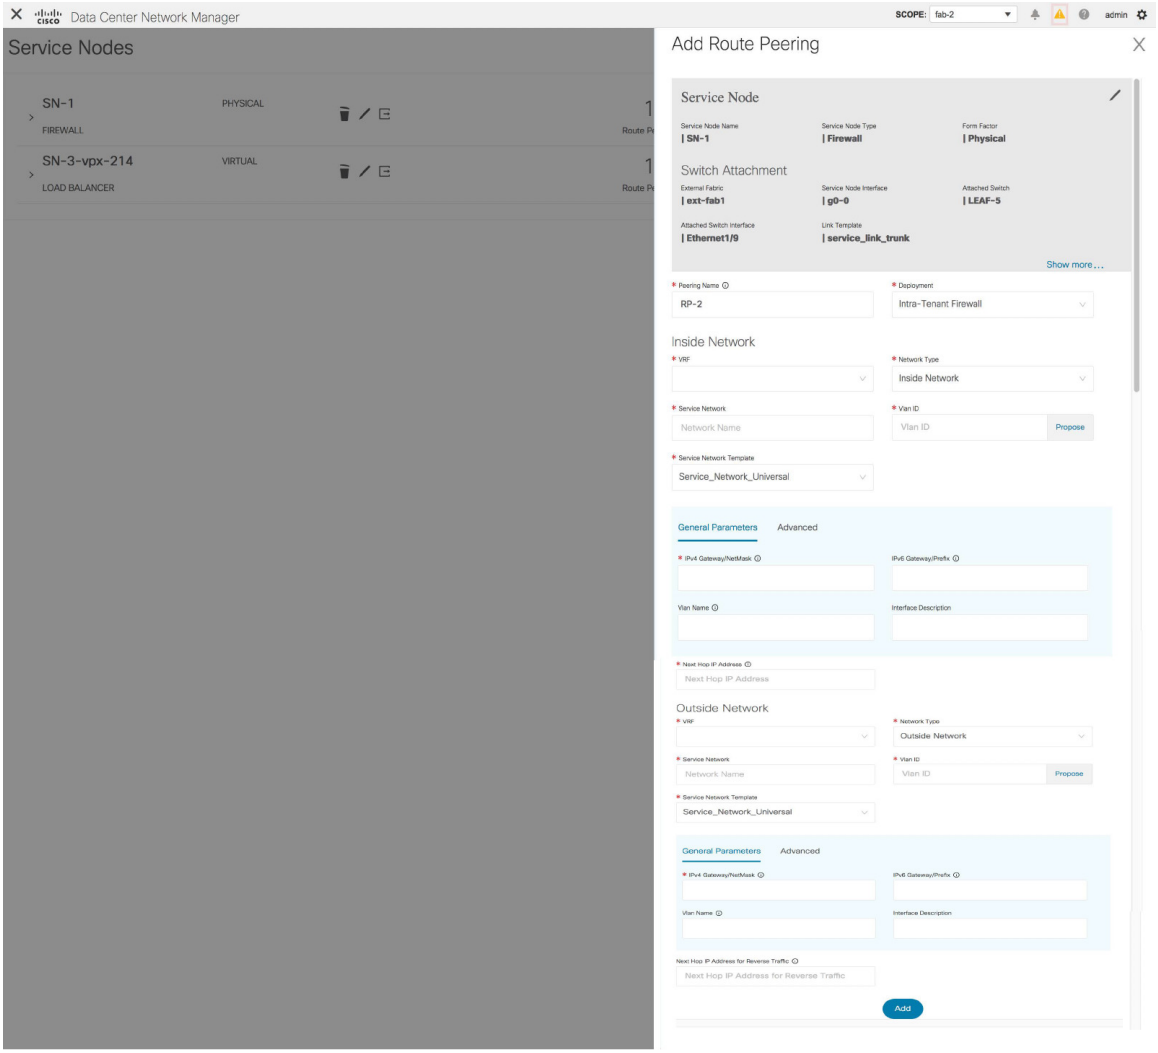

Specify the required parameters and click **Add**. For more information on specific fields, hover over the **i** icon.

Adding a Service Policy

To add a service policy from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Click the **Add Service Policy** icon on the **Service Nodes** window.

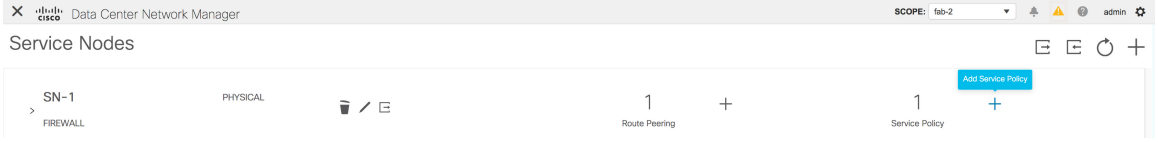

Step 2 The **Add Service Policy** window is displayed.

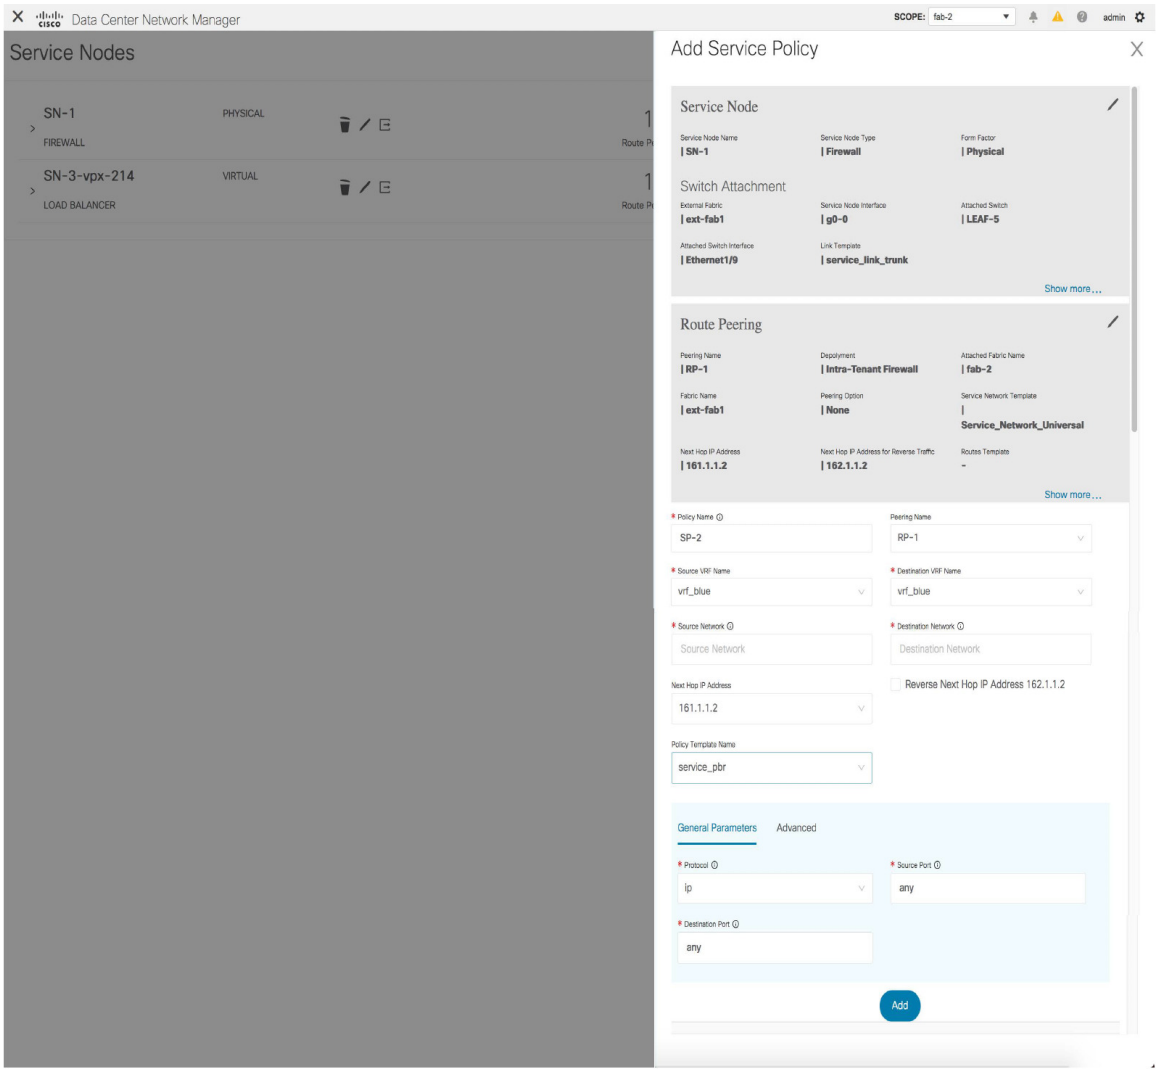

Specify the required parameters and click **Add**. For more information on specific fields, hover over the **i** icon.

Deleting a Service Node

To delete a service node from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Click the **Delete Service Node** icon on the **Service Nodes** window.

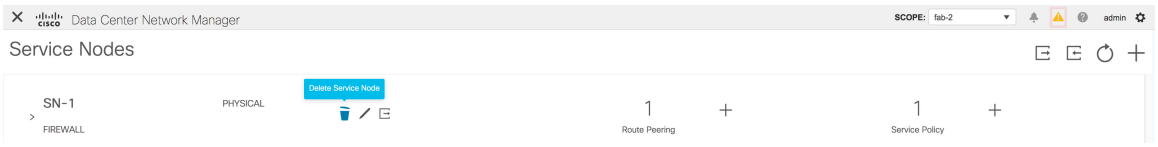

Step 2 A pop-up window comes up to confirm if the node has to be deleted. Click **Delete**.

Note Ensure that the service node that has to be deleted has no route peering or service policies associated with it. In case there are service policies or route peering associated with the service node, the deletion is blocked with a warning indicating that any route peering or service policies associated with the service node have to be removed before deleting the service node.

Editing a Service Node

To edit a service node from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Click the **Edit Service Node** icon on the **Service Nodes** window.

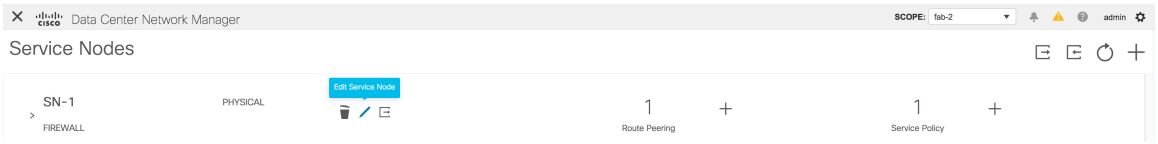

Step 2 The **Edit Service Node** window is displayed.

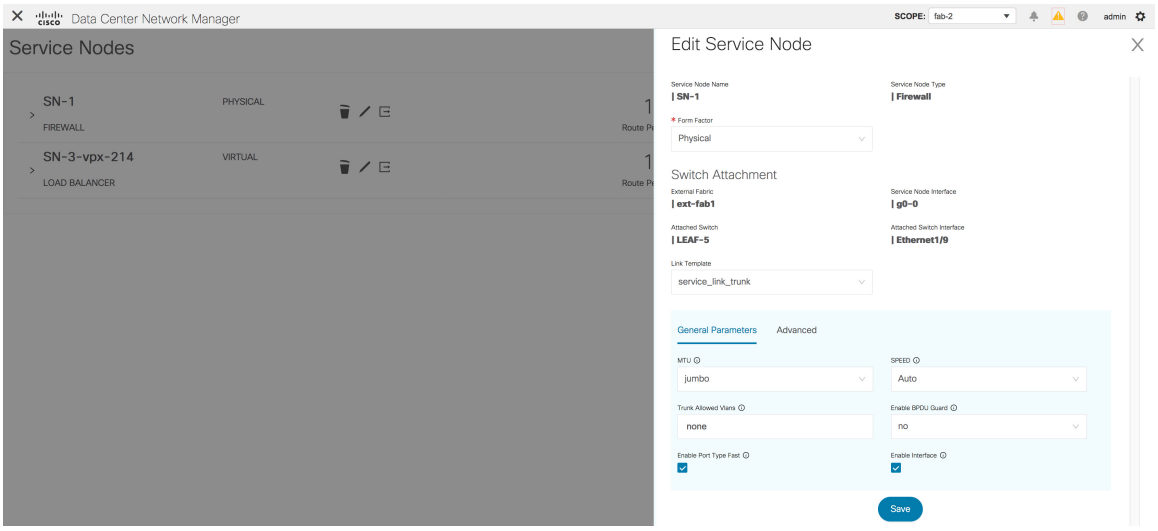

Make the required changes and click **Save**.

Refreshing the Service Policy and Route Peering List

To refresh the list of service policies or route peerings that is displayed in the **Service Nodes** window, click the **Refresh** icon that appears in the **Service Policy** tab or the **Route Peering** tab.

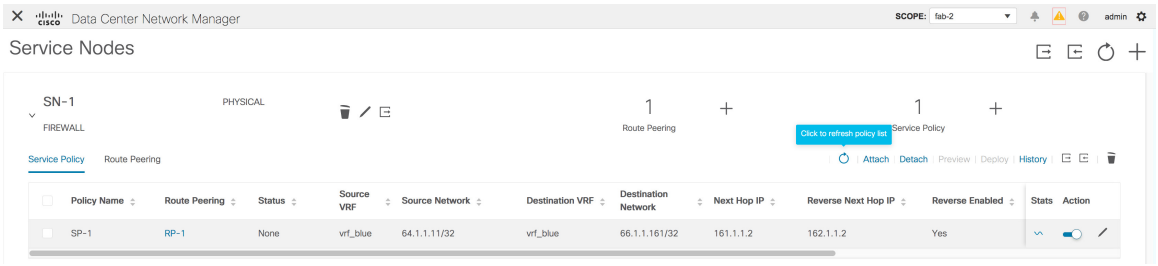

Refreshing a Specific Service Policy or Route Peering

From Cisco DCNM Release 11.5(1), click the **Refresh** icon that appears under the **Action** column to refresh a specific service policy or route peering.

Attaching a Service Policy or a Route Peering

To attach a specific service policy or route peering from a switch, select the checkbox next to the required service policy or route peering and click **Attach**.

Note From Cisco DCNM Release 11.5(1), bulk attachment, detachment, preview and deployment of route peering and service policies is supported and they are limited up to 10 route-peerings or 10 service policies only.

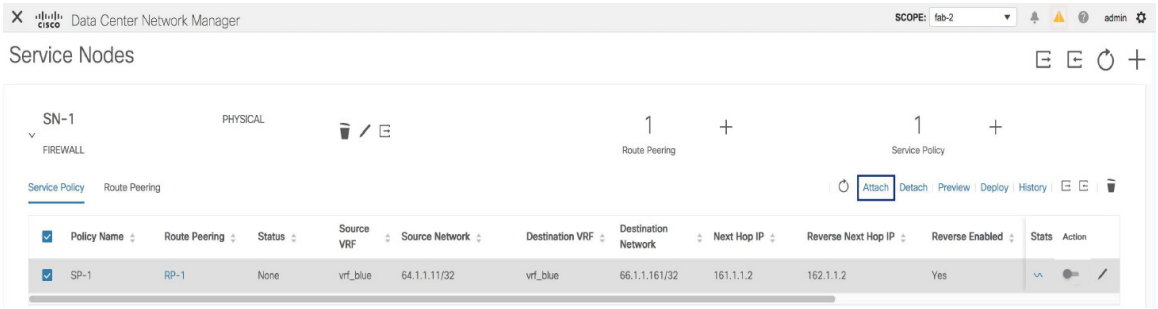

Detaching a Service Policy or a Route Peering

To detach a specific service policy or route peering from a switch, select the checkbox next to the required service policy or route peering and click **Detach**.

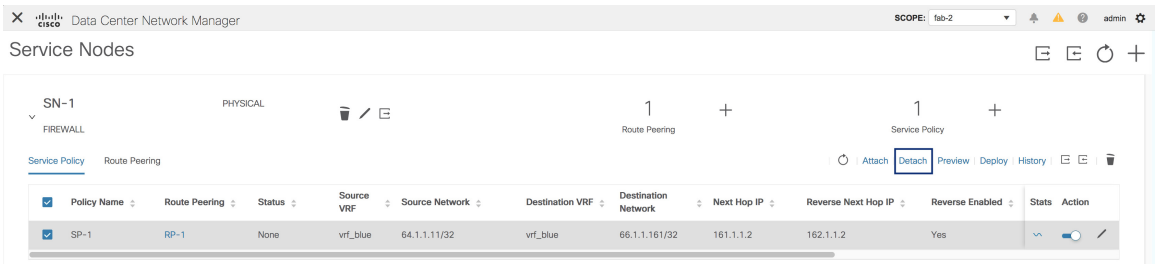

Preview a Service Policy or a Route Peering

To display the preview of a service policy or a route peering from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Select a service policy or route peering checkbox and click **Preview** on the **Service Nodes** window.

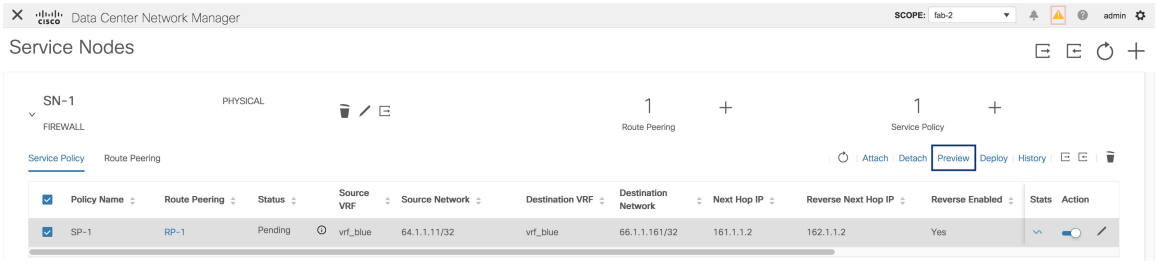

A **Preview Service Policy** or a **Preview Route Peering** window is displayed.

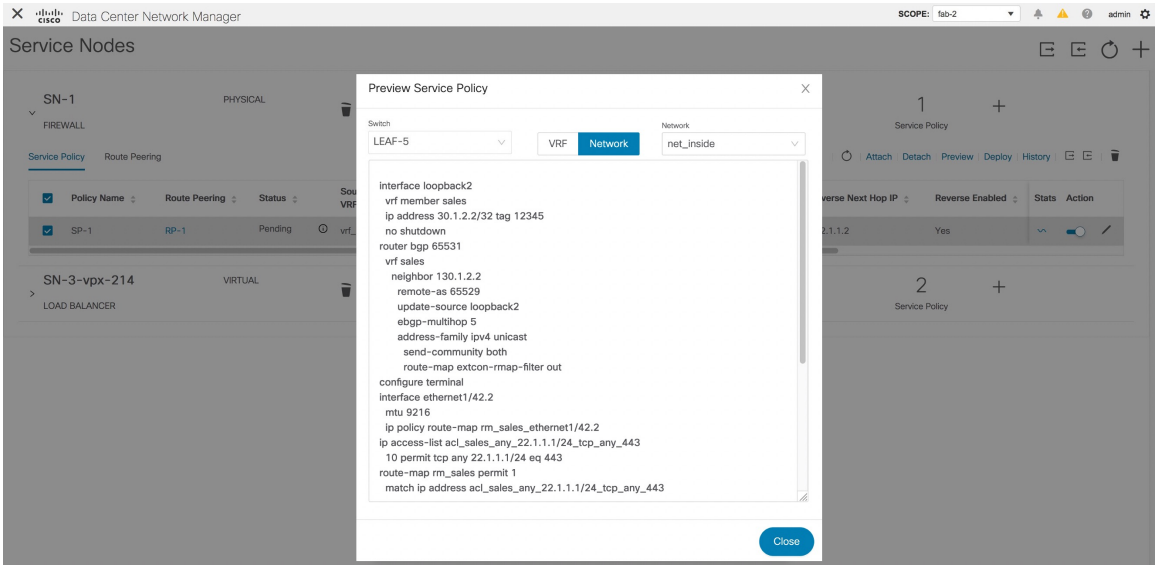

Step 2 Select a specific switch, network, or VRF from the respective drop-down lists to display the service policies or route peerings for specific switches, networks, and VRFs. Click Close to close the window.

Deploying a Service Policy or a Route Peering

To deploy a service policy or a route peering from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Select a service policy or route peering checkbox and click **Deploy** on the **Service Nodes** window.

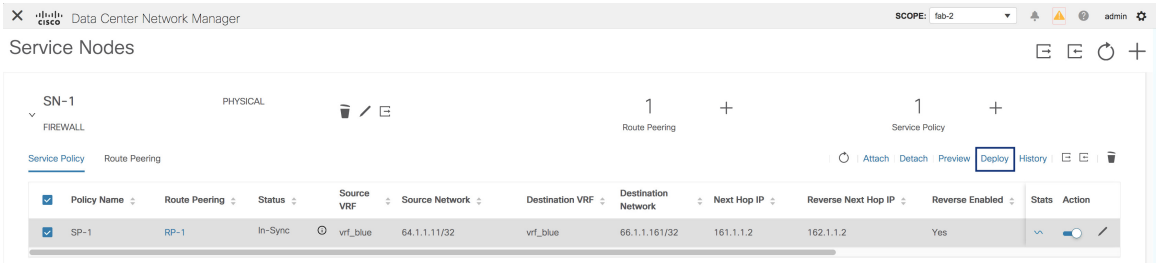

A pop-up window is displayed asking for confirmation to deploy.

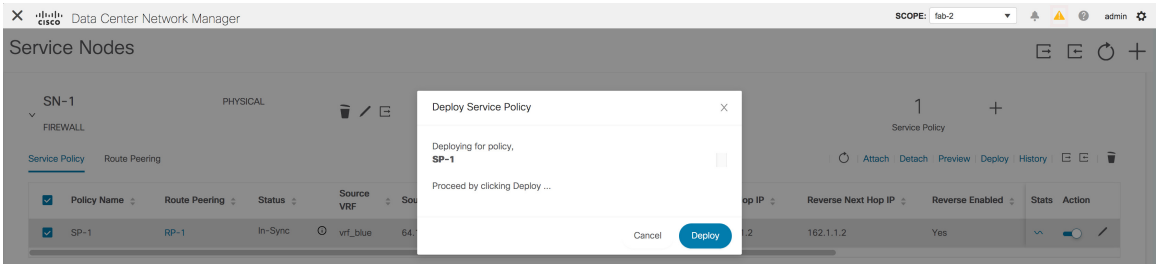

Step 2 Click **Deploy**.

Viewing Deployment History

To view deployment history of the switches and networks that are involved in the selected service policy or route peering, click **History** in the **Service Policy** tab or the **Route Peering** tab. The **Deployment History Service Policy** or the **Deployment History Route Peering** window is displayed.

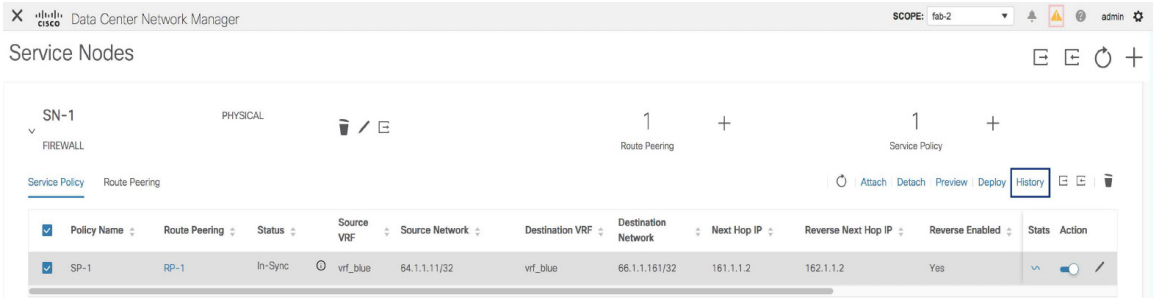

Information such as the name of the network, VRF, and switch, status, status details, and time of execution is displayed.

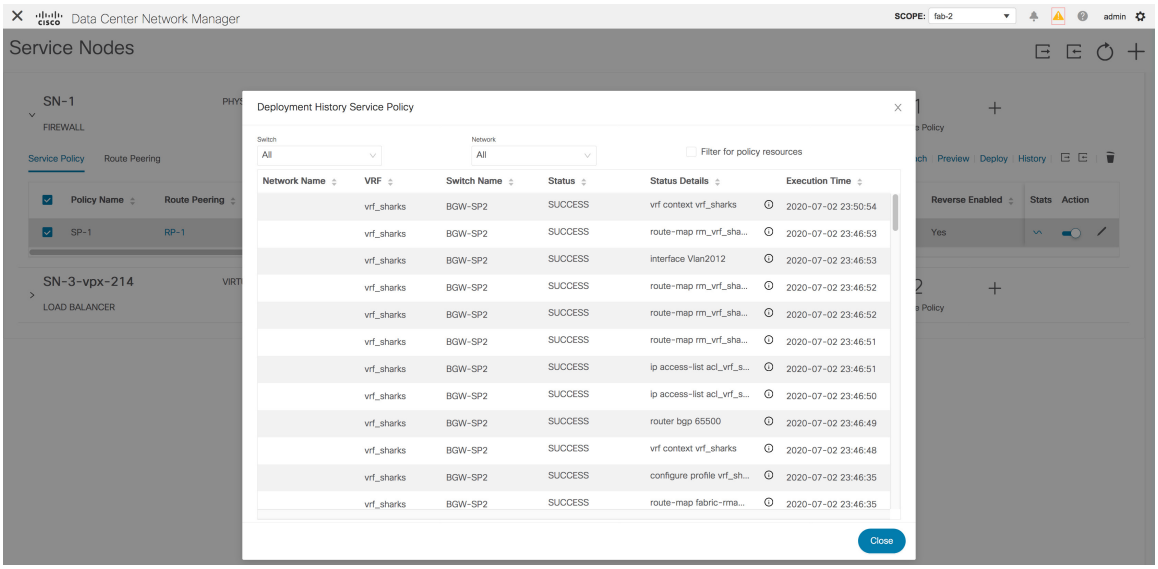

The first line in the list of CLIs is displayed in the **Status Details** column. This provides a peak into the deployed configuration. Hover over the **i** icon next to the **Status Details** field in each row to display more information.

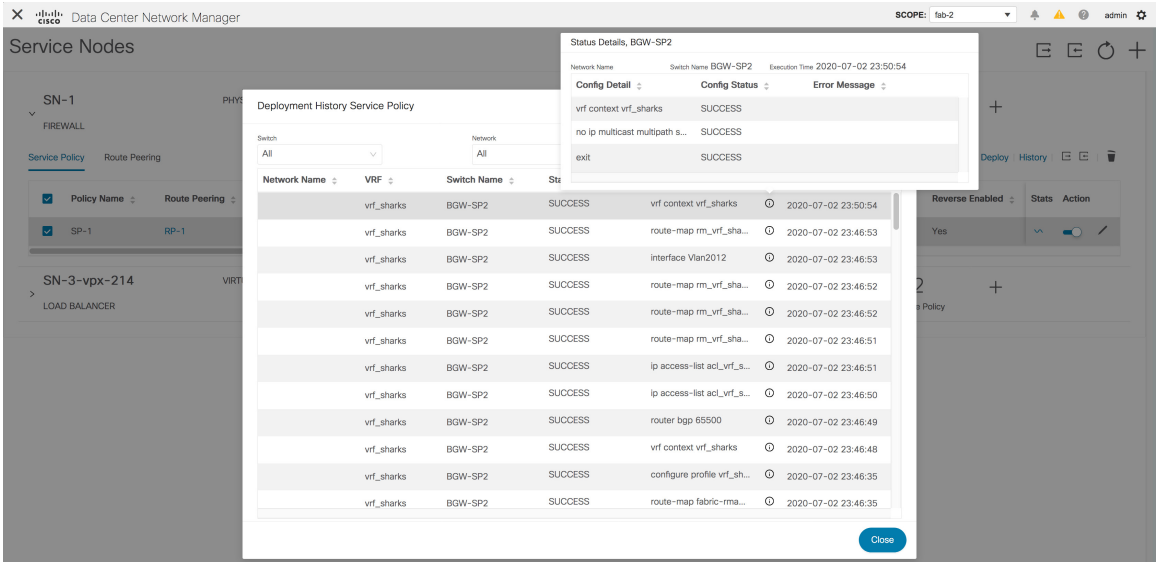

Select a switch from the **Switch** dropdown list to display information for the selected switch.

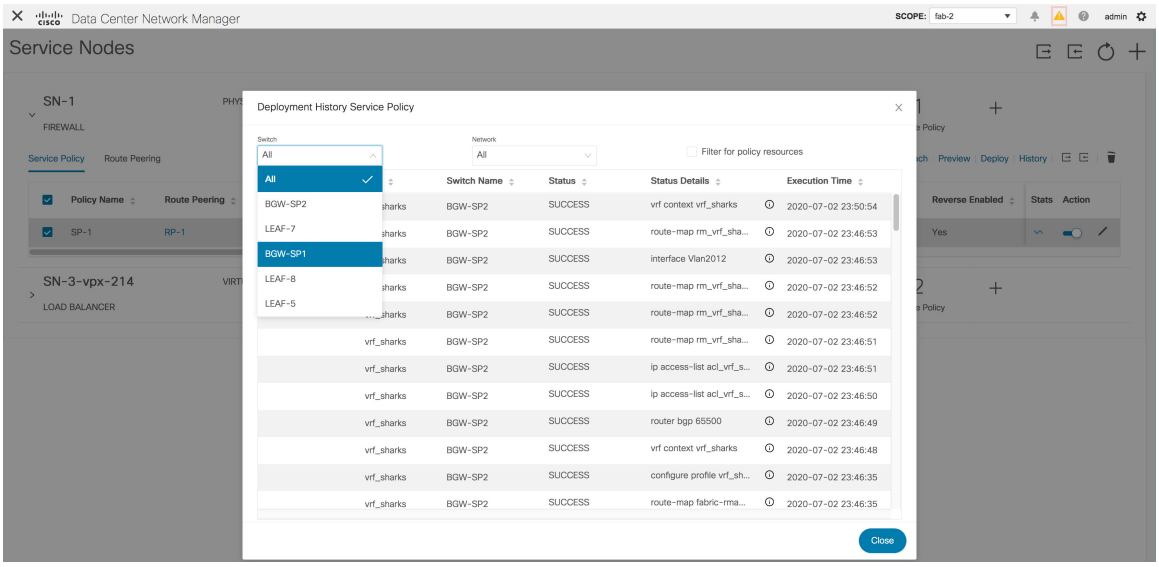

Select a network from the **Network** dropdown list to display information for the selected network.

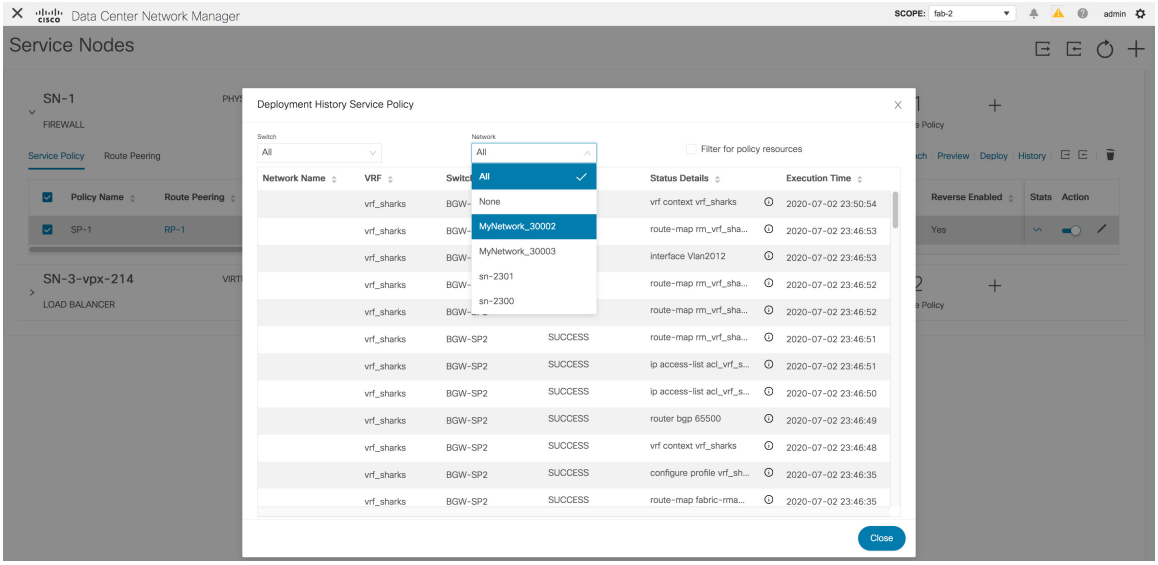

Select the **Filter for policy resources** checkbox to display only policy-related deployments such as ACLs, route maps and associated CLIs. This checkbox is available only in the **Deployment History Service Policy** window.

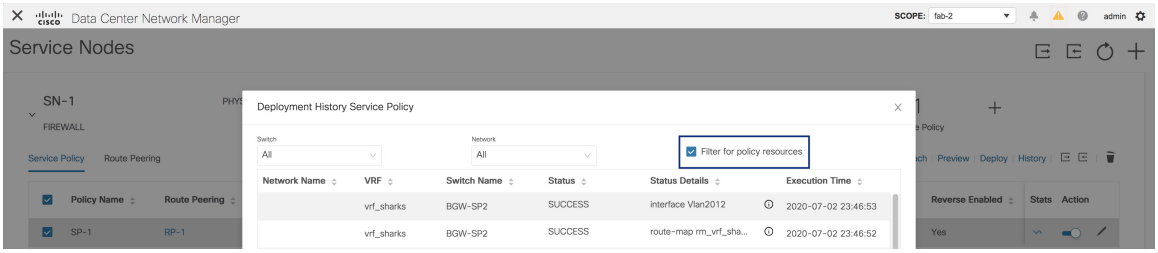

Exporting a Service Policy or a Route Peering Table

To export the service policy or route peering information as an Excel file, click the **Export** icon on the **Service Nodes** window. Click the **Export** icon on the **Service Policy** tab to export information about the service policies. Click the **Export** icon on the **Route Peering** tab to export information about the route peerings.

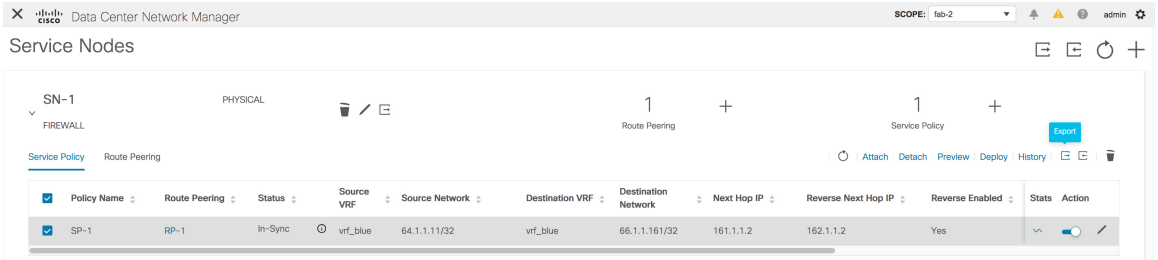

Importing a Service Policy or a Route Peering Table

To import service policy or route peering information as an Excel file, click the **Import** icon on the **Service Nodes** window. Click the **Import** icon on the **Service Policy** tab to export information about the service policies. Click the **Import** icon on the **Route Peering** tab to export information about the route peerings.

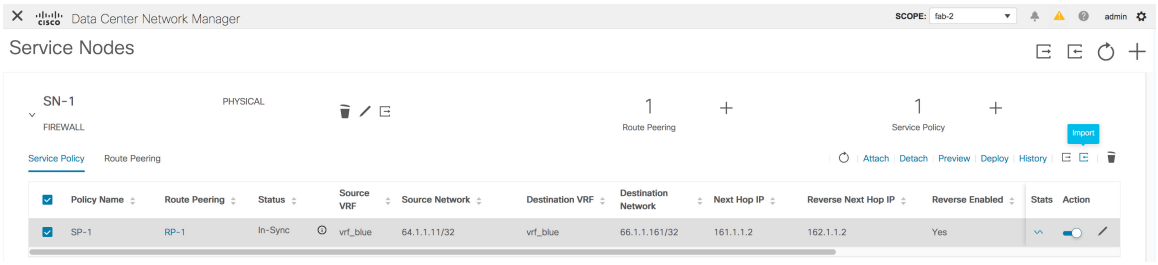

Deleting a Service Policy

To delete a service policy from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Select the service policy that has to be deleted by clicking the checkbox that is next to the name of the policy, and then click the **Delete** icon on the **Service Nodes** window.

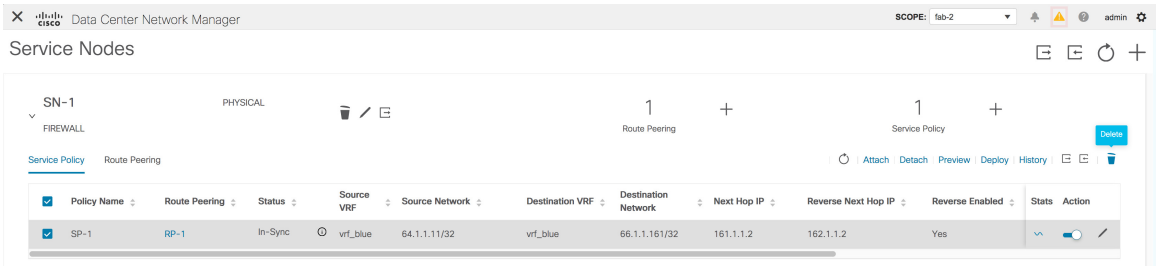

Step 2 A pop-up window is displayed asking for confirmation to delete. Click **Delete**. In case the service policy that has to be deleted is attached, the pop-up window indicates that the service policy has to be detached by using the toggle in the **Action** column, and deploying the changes (removing the policy) before it can be deleted.

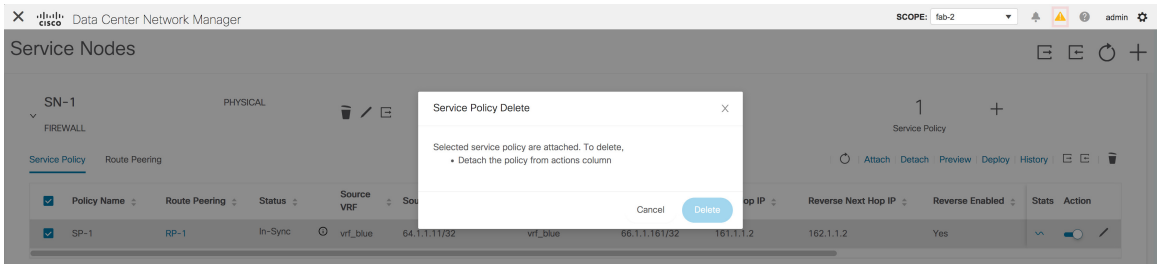

Deleting a Route Peering

To delete a route peering from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Select the route peering that has to be deleted by clicking the checkbox that is next to the name of the route peering, and then click the **Delete** icon on the **Service Nodes** window.

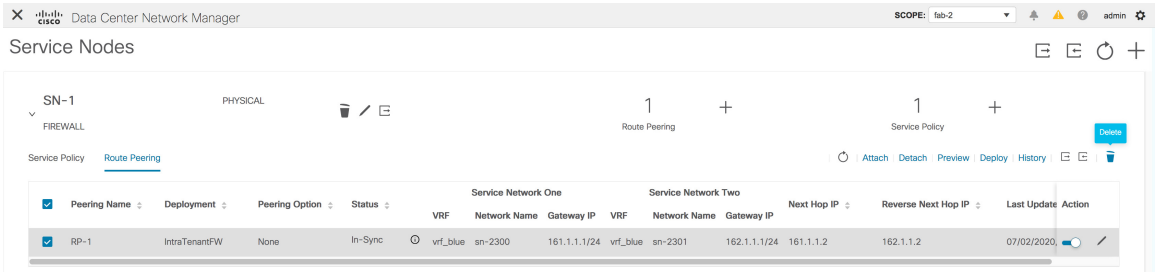

Step 2 A pop-up window is displayed asking for confirmation to delete. Click **Delete**. In case the route peering that has to be deleted is attached or if the service policy associated with the route peering is active, the pop-up window indicates that the peering has to be detached by using the toggle in the **Action** column, deploy the changes (remove the policy), and delete the service policy associated with the route peering before the route peering can be deleted.

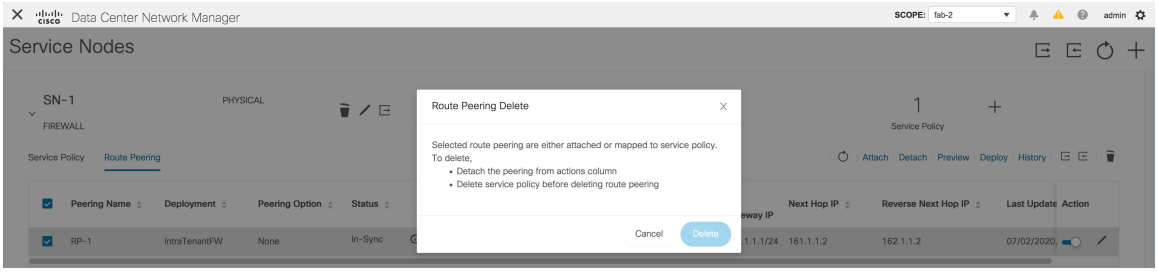

Viewing Service Policy Information

In the **Service Nodes** window, the **Service Policy** tab displays information about the configured service policies.

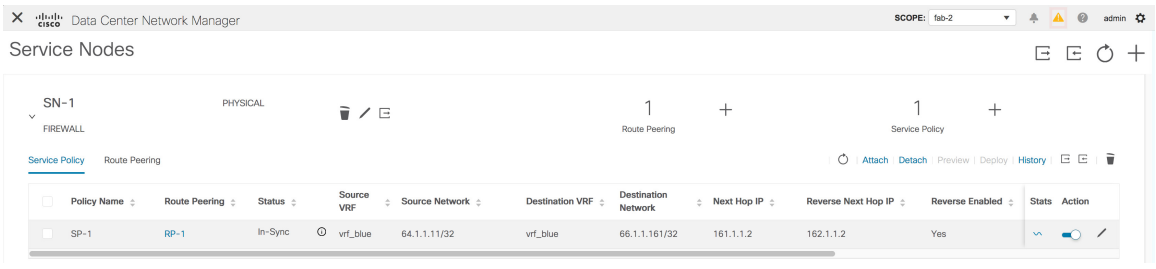

Table 27: Service Policy Table Field and Description

Field	Description
Policy Name	Displays the name of the policy.
Route Peering	Displays the route peering name given for the peering configuration. Click the specified peering name to display route peering information.
Status	Displays the status of the service policy.
Source VRF	Displays the Virtual Routing and Forwarding (VRF) source.
Source Network	Displays the source network.
Destination VRF	Displays the destination VRF.
Destination Network	Displays the destination network.
Next Hop IP	Displays the next-hop IP address.
Reverse Next Hop IP	Displays the reverse next-hop IP address.
Reverse Enabled	Displays if reverse next-hop is enabled or not.
Route Map Action	Displays the specified route map action.
Next Hop Option	Displays the specified next hop option.
Last Updated	Displays the time at which the service policy was last updated.
Stats	Click the graph line to display cumulative statistics for a policy in a specified time range. For more information, refer Stats.

Field	Description
Action	Use the toggle to enable/attach or disable/detach the service policy. When the service policy is attached or enabled, the corresponding policies are applied to the VRF (tenant), source, and destination networks. 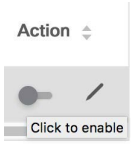 The toggle turns blue in color when the service policy is attached or enabled. Click the **Edit** icon to edit the service policy.

Stats

In the **Service Nodes** window, the **Service Policy** tab displays statistical information about the configured service policies. Select a time range for which the statistics should be displayed from the **Time Range** drop-down box. You can select the date from the calendar displayed on the window and the time by clicking **select time** at the bottom right corner of the window. You can also display statistics from the last 15 minutes, 1 hour, 6 hours, 1 day, and 1 week. Select the required time range and click **Apply**. Select a switch for which the statistics should be displayed from the **Switch** drop-down list. The statistics are then displayed for the selected switch in the specified time range. Starting from Cisco DCNM Release 11.4(1), you can click **Clear Stats** to reset the statistics for a specific policy on all involved switches. If multiple policies are sharing the same route map, then the statistics of other policies are also impacted.

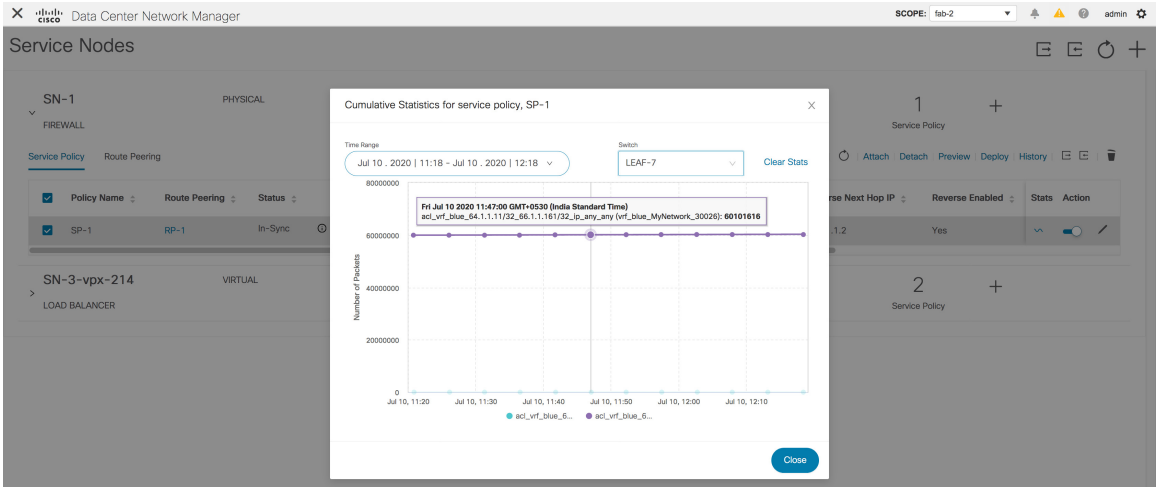

Viewing Route Peering Information

In the **Service Nodes** window, click **Route Peering**. The **Route Peering** tab displays route peering information.

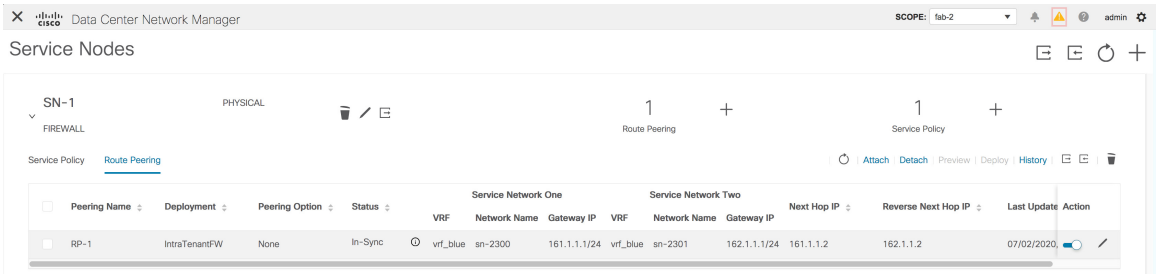

Table 28: Route Peering Table Field and Description

Field	Description
Peering Name	Displays the defined peering name.
Deployment	Displays the deployment - One-Arm mode or Two-Arm mode.
Peering Option	Displays the peering option - Static or eBGP Dynamic peering.
Status	Displays the status of the route peering.
Service Network VRF	Displays the service network VRF.
Service Network Name	Displays the name of the service network.
Service Network Gateway IP	Displays the gateway IP of the service network VRF.
Next Hop IP	Displays the next-hop IP address.
Reverse Next Hop IP	Displays the reverse next-hop IP address.

Field	Description
Last Updated	Displays the time at which the route peering was last updated.
Action	Use the toggle to enable/attach or disable/detach the route peering. When the route peering is enabled, the service networks defined in that route peering will be attached to the service leaf. 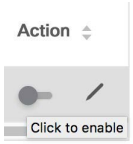 The toggle turns blue in color when the route peering is attached or enabled. Click the **Edit** icon to edit the route peering. 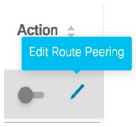

Service Node Backup and Restore

You can back up data at the service node level by clicking the **Export** icon to export data about the service nodes to an excel file. Data regarding all the service nodes, the respective route peerings and service policy, is exported.

You can also restore the service node level data by clicking the **Import** icon to import data about the service nodes from an excel file.

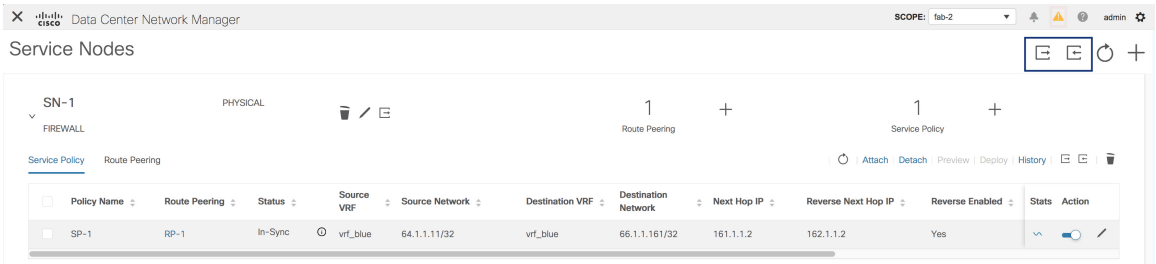

You can also export data for a specific service node by clicking the **Export** icon located next to the **Edit Service Node** icon.

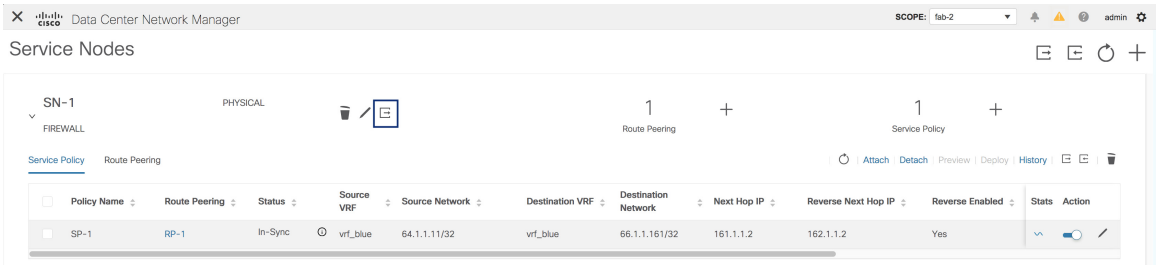

Fabric Backup and Restore

During easy fabric and parent MSD fabric backup, the service node connections, route peering and service policy configurations, such as composed ACL and route-map, are saved as part of the fabric, VRF and the tenant network intent. However, the definitions of service node, route peering and service policy are not saved. We recommend backing up the service data by clicking the **Export** icon at the service node level from the **Control > Services** window. While restoring easy fabrics and parent MSD fabrics, the service data can be restored by clicking the **Import** icon at the service node level from the **Control > Services** window. The service node connections, route peering and service policy configuration will be restored along with the associated fabric, VRF and the tenant network intent.

Brownfield Migration

During brownfield migration, the L4-L7 service configuration, such as ACLs and route-maps associated with networks and VRFs, are captured in the switch freeform policy linked to the tenant network and the VRF profile. No service node, route peering, or service policy is auto-generated as a result of brownfield migration. If you want to apply a new service policy to the same tenant network or VRF, remove the captured freeform configuration and configuration compliance will then generate the required CLIs that you can deploy later.

Audit History

From Cisco DCNM Release 11.5(1), click the Audit icon on the Service Nodes window to display the Audit History window.

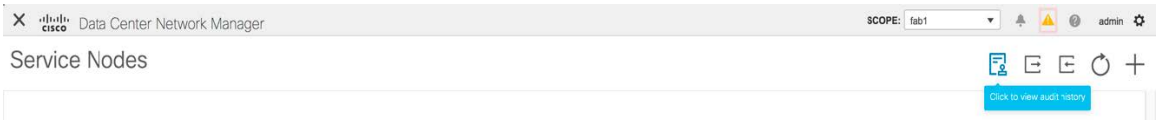

The Audit Logs table in the Audit History window displays information about all the actions that have been performed. Audit logs are generated when the following actions are performed:

- Creation of service nodes, route peering and service policies
- Deletion of service nodes, route peering and service policies
- Update of service nodes, route peering and service policies
- Attachment and detachment of route peering and service policies
- Deployment of route peering and service policies

This audit log is saved with the name of the user who has performed the action, the role of the user, the action taken, the entity on which the action was performed, details about the action, the status, and the time at which the action was performed.

To perform a search in each column, click the search icon in the required column and enter the search string.

To display more information about each row, click the + icon next to the user name.

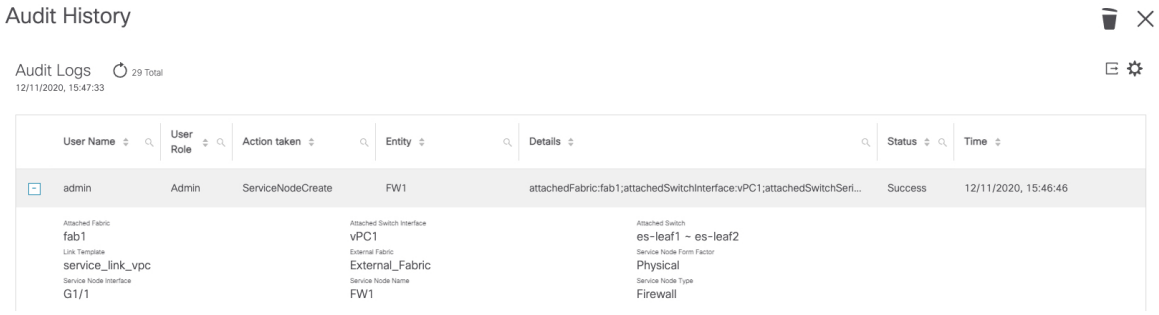

To export the data on this window to an Excel file, click the Export icon.

To selectively hide or show fields from the Audit Logs table, click the gear icon that is located next to the export icon to select the fields that have to be displayed in the Audit Logs table.

To delete older audit reports, click the icon, specify the maximum retained dates and confirm deletion. Note that only users with the admin role can delete audit log entries.

To display the latest audit log, click the Refresh icon that is located above the Audit Logs table.

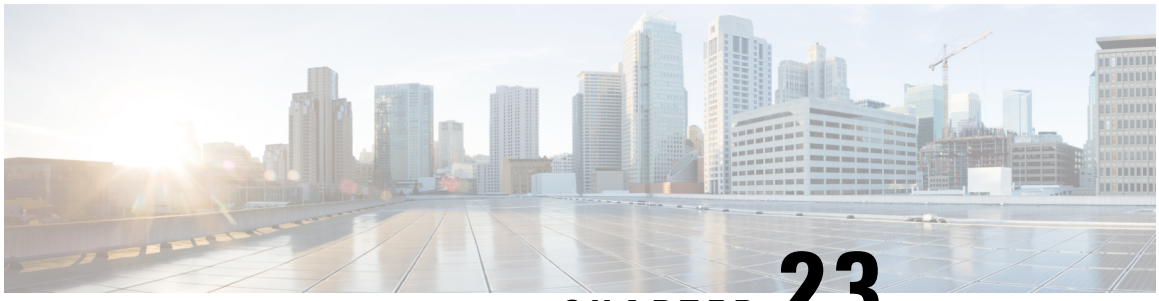

CHAPTER 23

L4-L7 Service Use Cases

- Use Case: Intra-tenant Firewall with Policy-based Routing, on page 903
- Use Case: Inter-tenant Firewall with eBGP Peering, on page 921
- Use Case: One-arm Load Balancer, on page 928

Use Case: Intra-tenant Firewall with Policy-based Routing

Refer the figure given below for topology details.

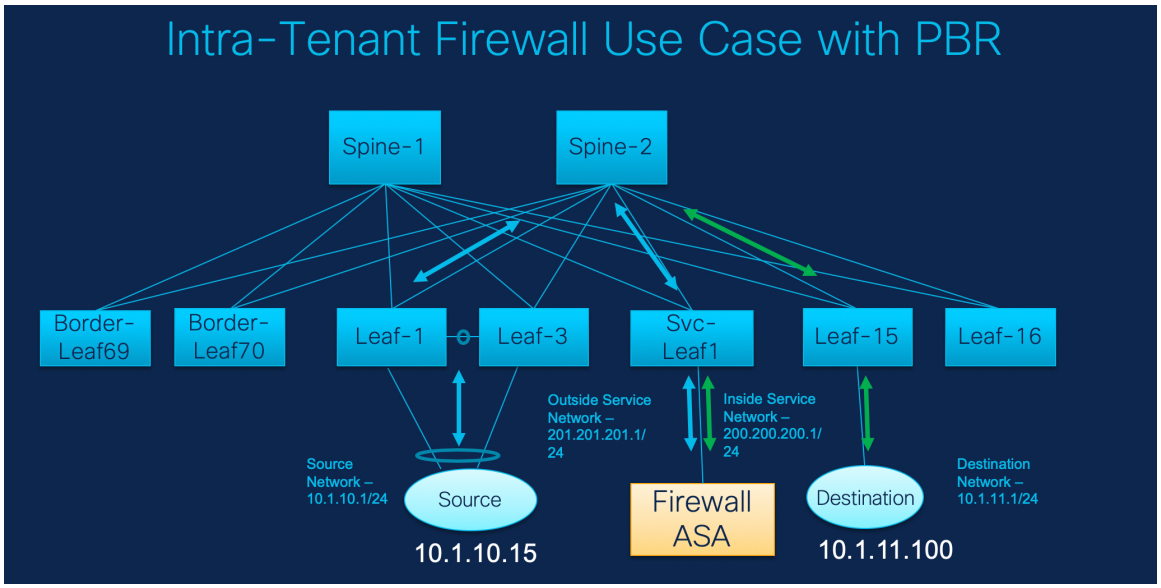

In this topology, Leaf1 and Leaf3 are a vPC pair and they are connected to **Source** (10.1.10.15) with the **Source Network** (10.1.10.1/24). The service leaf is connected to the virtual **Firewall ASA** and Leaf-15 is connected to **Destination** (10.1.11.100). In this use case, the source network refers to 'client' and the destination refers to 'server'.

Any traffic that is traversing from **Source** to **Destination** must go to the outside service network, and the firewall performs its function by allowing or denying traffic. This traffic is then routed to the inside service network and on to the Destination network. Since the topology is stateful, the traffic coming back from the destination to the source follows the same path.

Now, let us see how to perform service redirection in DCNM.

Note

- This use-case does not cover how to provision the **Site_A** VXLAN fabric. For information about this topic, refer to the Cisco DCNM LAN Fabric Configuration Guide.
- This use-case does not cover configurations on the service node (firewall or load balancer).

Select **Control** > **Fabrics** > **Services**.

This use-case consists of the following steps:

1. Create Service Node

Procedure

Step 1 From the **Scope** drop-down list, select **Site_A**.

Step 2 Click the **Add** icon in the **Service Nodes** window.

Step 3 Enter the node name and specify **Firewall** in the **Type** dropdown box. The **Service Node Name** has to be unique.

Step 4 From the **Form Factor** drop-down list, select **Virtual**.

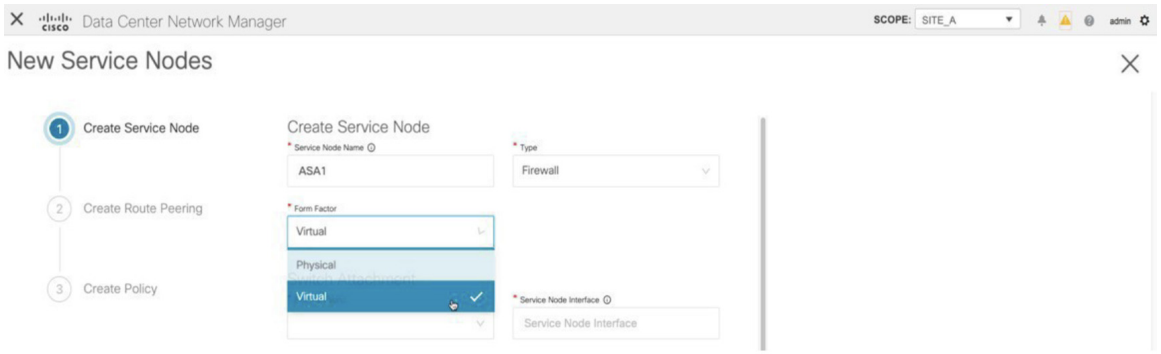

Step 5 In the **Switch Attachment** section, from the **External Fabric** drop-down list, select the external fabric in which the service node (for example, ASA firewall) is located. Note that the service nodes need to belong to the external fabric. This is a prerequisite before creating a service node.

Step 6 Enter the interface name of the service node that will be connected to the service leaf.

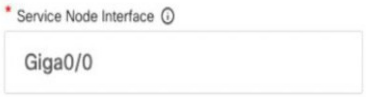

Step 7 Select the attached switch that is the service leaf, and the respective interface on the service leaf.

Step 8 Select the **service_link_trunk** template. DCNM supports trunk, port channel, and vPC link templates. The available link templates in the **Link Template** drop-down list are filtered based on the selected **Attached Switch Interface** type.

Step 9 Specify the **General Parameters** and **Advanced** parameters, if required. Some parameters are pre-filled with the default values.

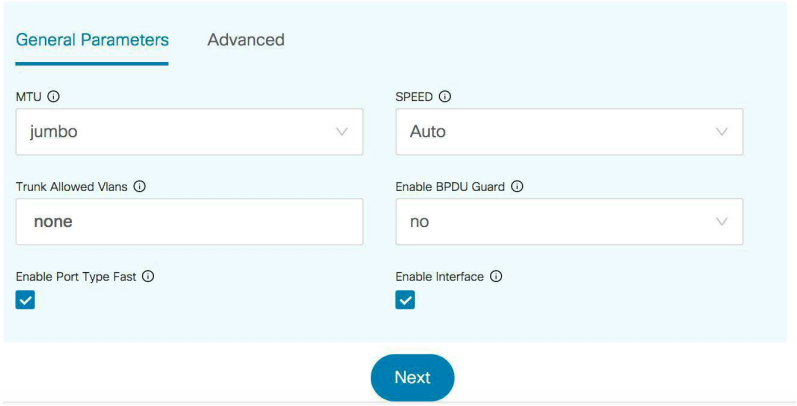

Step 10 Click **Next** to save the created service node.

2. Create Route Peering

Let us now configure the peering between a service leaf and a service node.

Procedure

Step 1 Enter the peering name and select **Intra-Tenant Firewall** from the **Deployment** drop-down list.

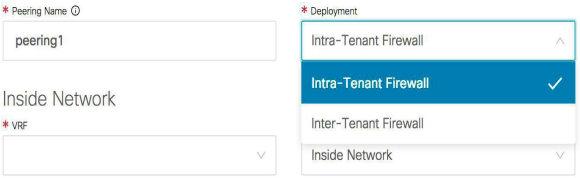

Step 2 Under **Inside Network**, from the **VRF** drop-down list, select a VRF that already exists and select **Inside Network** under **Network Type**.

Enter the name of the **Service Network** and specify the **Vlan ID**. You can also click **Propose** to allow DCNM to fetch the next available VLAN ID from the specified service network VLAN ID range in the fabric settings. The default **Service Network Template** is **Service_Network_Universal**.

Under the **General Parameters** tab, specify the gateway address for the service network. Specify the **Next Hop IP Address**. This next hop address has to be within the 'inside service network' subnet. Under the **Advanced** tab, the default **Routing Tag** value is 12345.

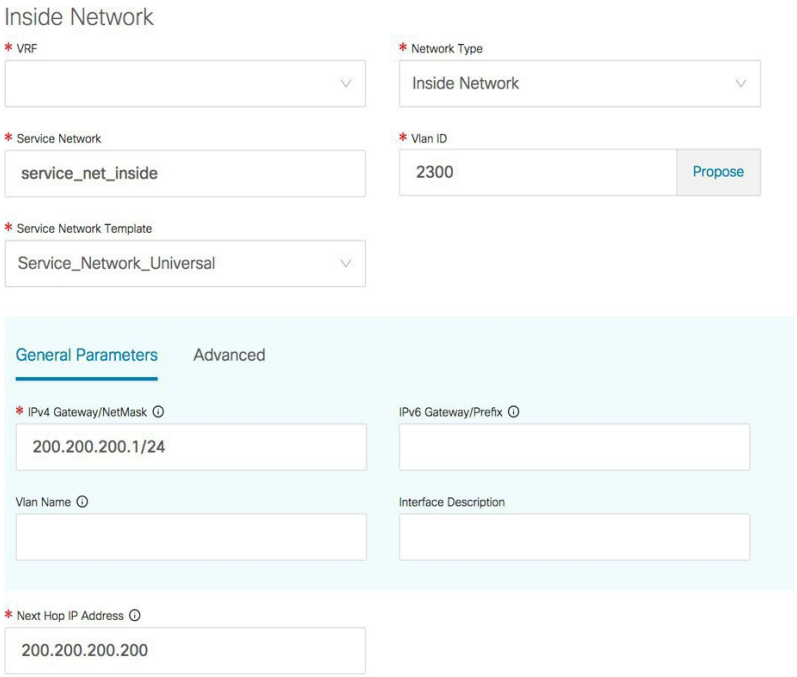

Step 3 Specify the required parameters under **Outside Network** and specify the **Next Hop IP Address for Reverse Traffic**. This next hop address for reverse traffic needs to be within the 'outside service network' subnet.

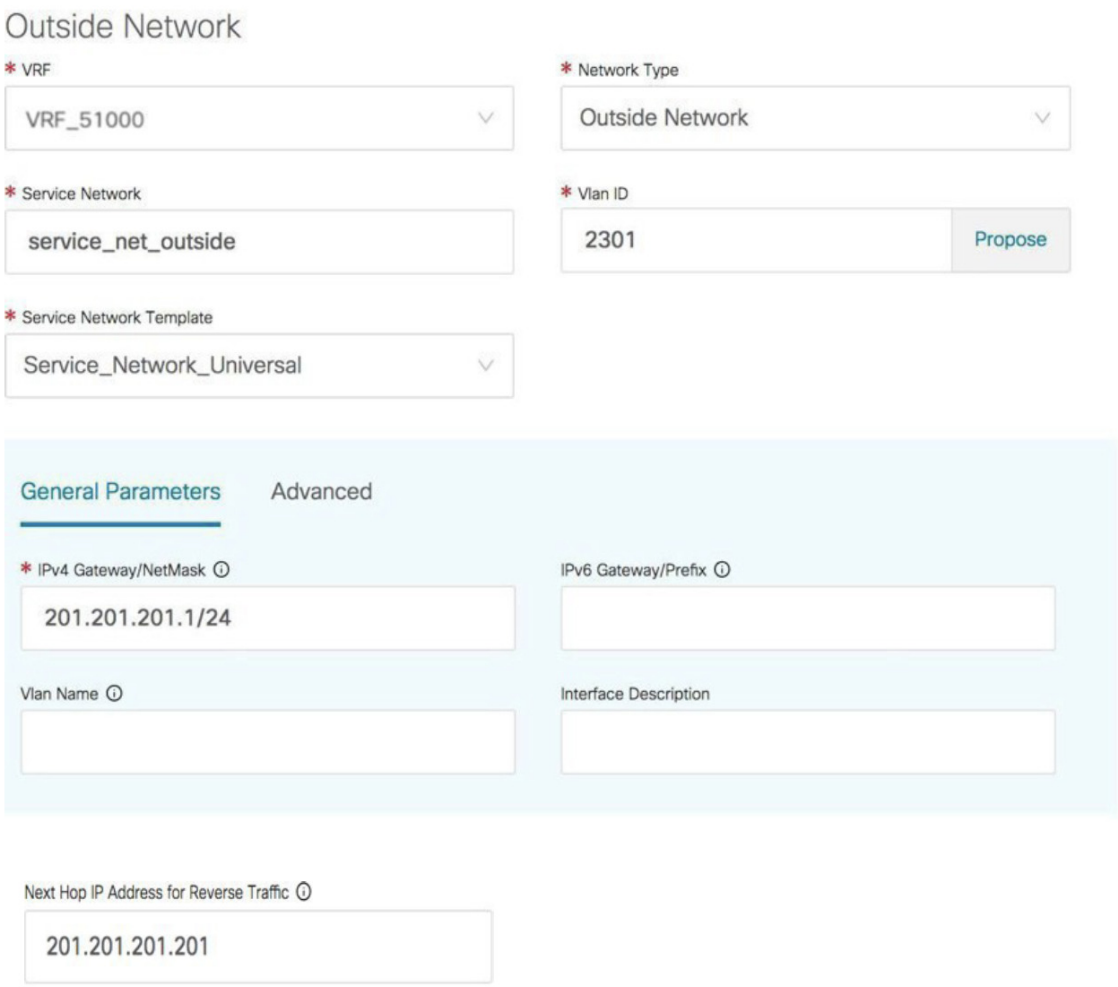

Step 4 Click **Next** to save the created route peering.

3. Create Service Policy

Procedure

Step 1 Specify a name for the policy and select the route peering from the **Peering Name** drop-down list.

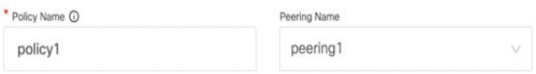

Step 2 Select the source and destination VRFs from the **Source VRF Name** and **Destination VRF Name** drop-down lists. The source and destination VRFs for an intra-tenant firewall deployment have to be the same.

Step 3 Select the source and destination networks from the **Source Network** and **Destination Network** drop-down lists, or specify the source or destination network that is within the network subnets defined in the **Control > Fabrics > Networks** window.

Step 4 The next hop and reverse next hop fields are populated based on the values entered while creating the route peering. Select the check box next to the **Reverse Next Hop IP Address** field to enable policy enforcement on reverse traffic.

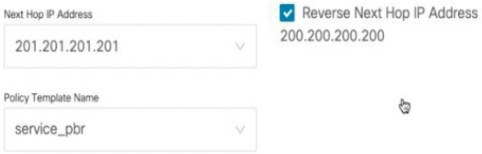

Step 5 Under the **General Parameters** tab in the policy template, select **ip** from the **Protocol** dropdown list, and specify **any** in the **Source Port** and the **Destination Port** fields.

Note For **ip** and **icmp** protocols, the **any** source and destination port is always used for ACL generation. You can also select a different protocol and specify the corresponding source and destination ports. DCNM will convert well-known port numbers to match the format required by the switch. For example, you can convert port 80 to ‘www’.

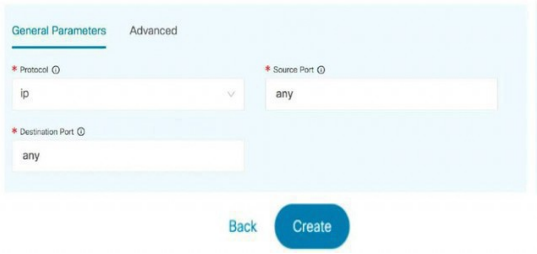

Step 6 Under the **Advanced** tab, by default, **permit** is selected for **Route Map Action** and **none** is selected for the **Next Hop Option**. You can change these values, and customize the ACL name and route map match sequence number, if required. For more information, refer Templates in the Layer 4-Layer 7 Service configuration guide.

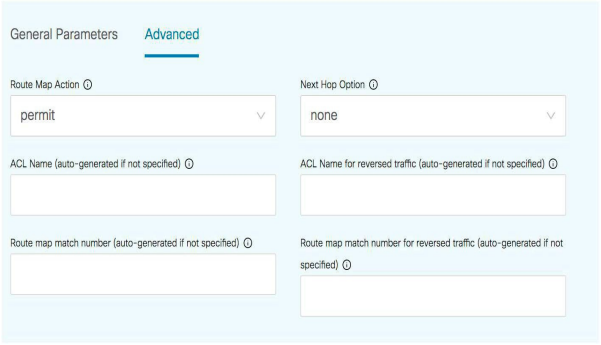

Step 7 Click **Create** to save the created service policy.

This completes the procedures that have to be performed to specify the flows for redirection.

4. Deploy Route Peering

Procedure

Step 1 In the **Service Nodes** window, select the required peering under the **Route Peering** tab.

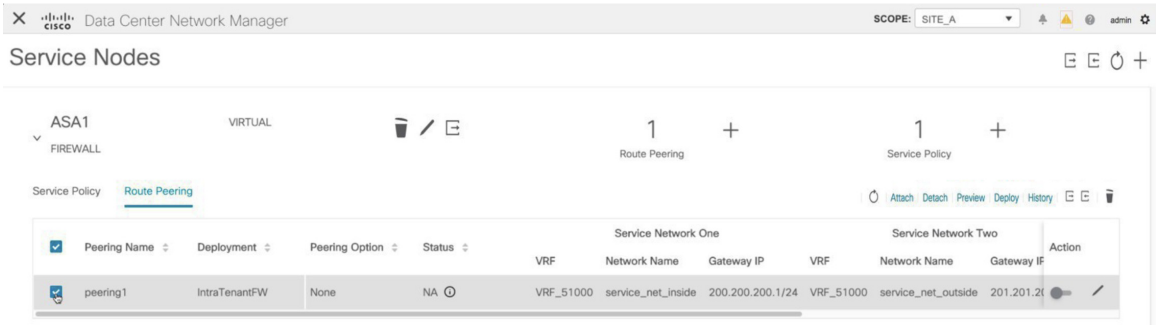

Step 2 Click the toggle button under **Action** to attach service networks to the service leafs.

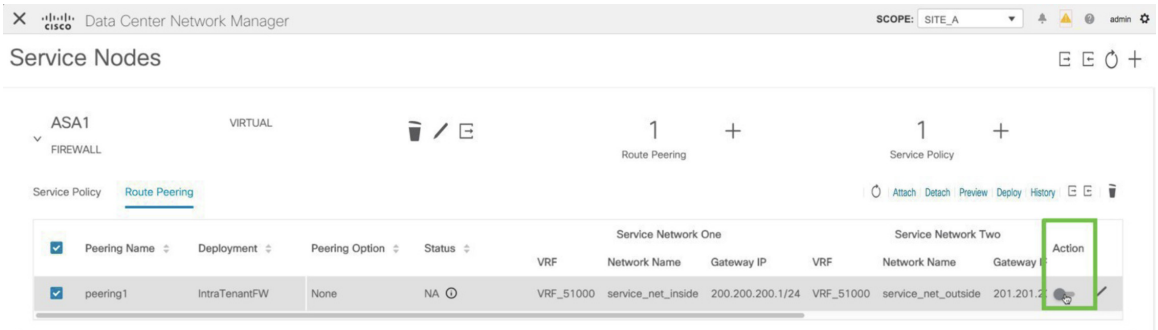

Step 3 Click **Preview** to view the configurations that will be pushed to the service leaf.

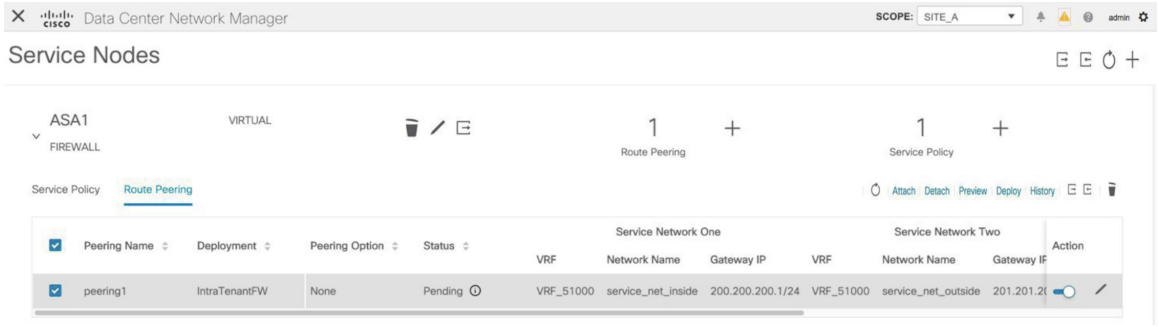

Previously, we had created inside and outside service networks. You can view these network configurations that will be pushed to the service leaf.

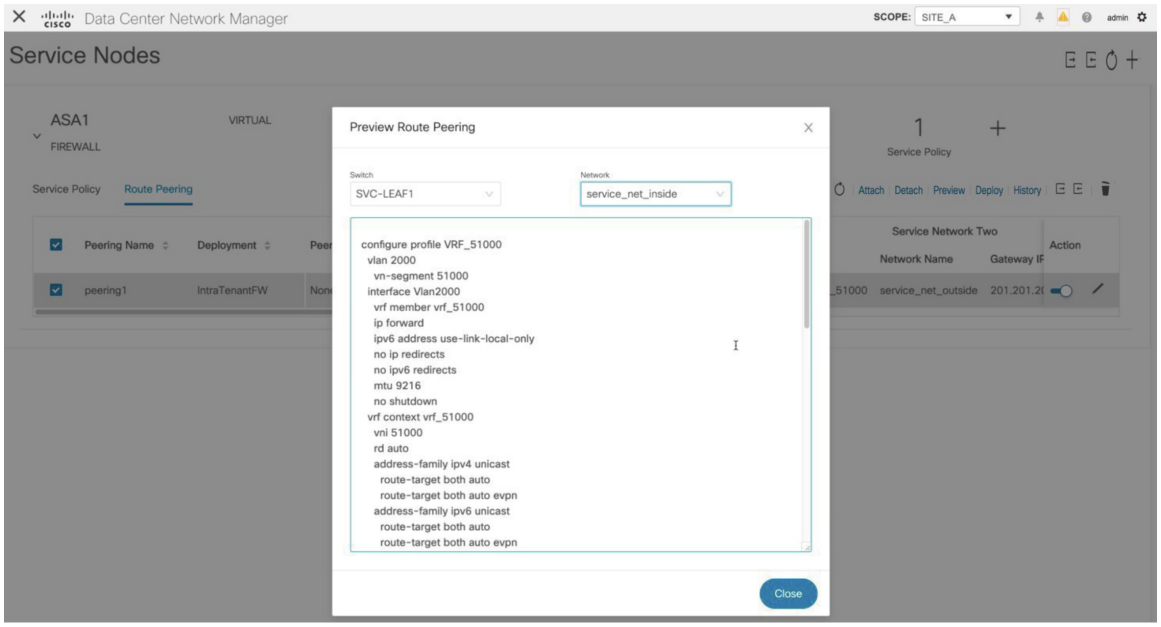

Step 4 Click **Close** to close the **Preview Route Peering** window.

Step 5 Click **Deploy** in the **Service Nodes** window to deploy the configuration to the attached switches (service leaf(s)) for route peering.

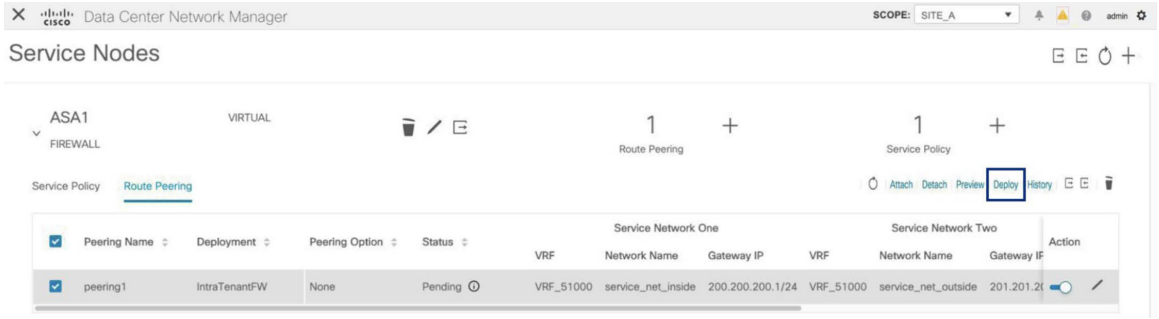

Click the **Deploy** button in the pop-up window to confirm deployment.

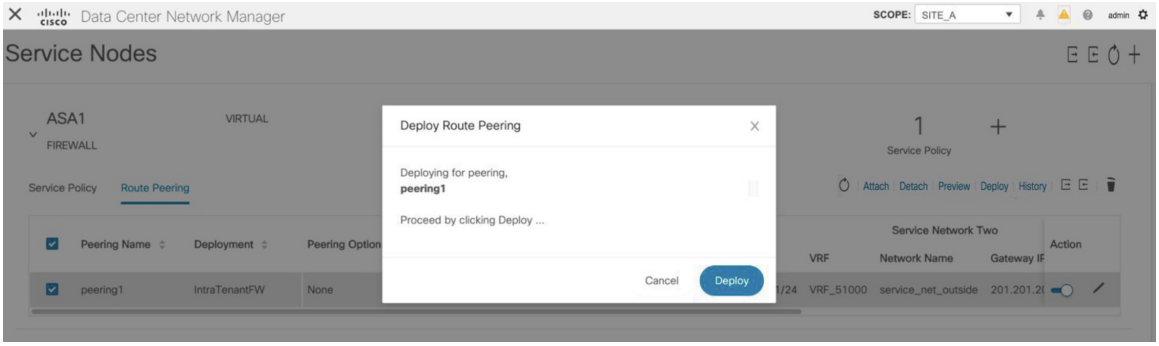

Step 6 Click the **Refresh** icon for the latest peering configuration attachment and deployment status.

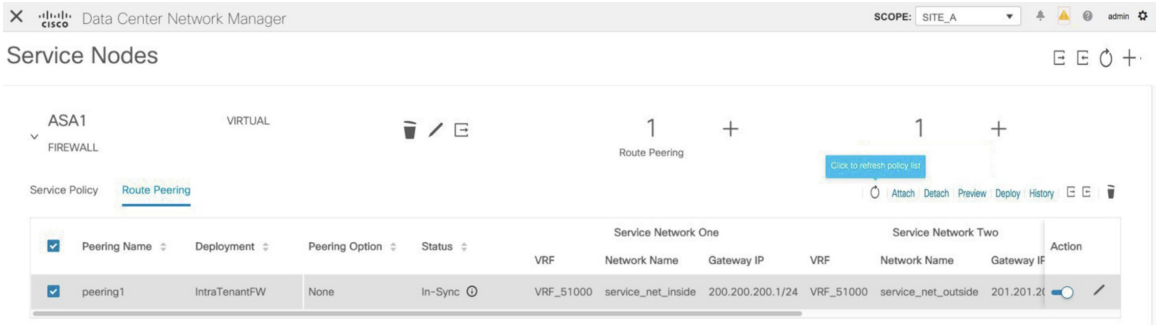

5. Deploy Service Policy

Perform the following procedure to deploy the service policy. This policy's corresponding configuration will be deployed to the switches that the source and destination network are attached to, and to the service leaf(s).

Procedure

Step 1 Select the checkbox next to the required policy under the **Service Policy** tab.

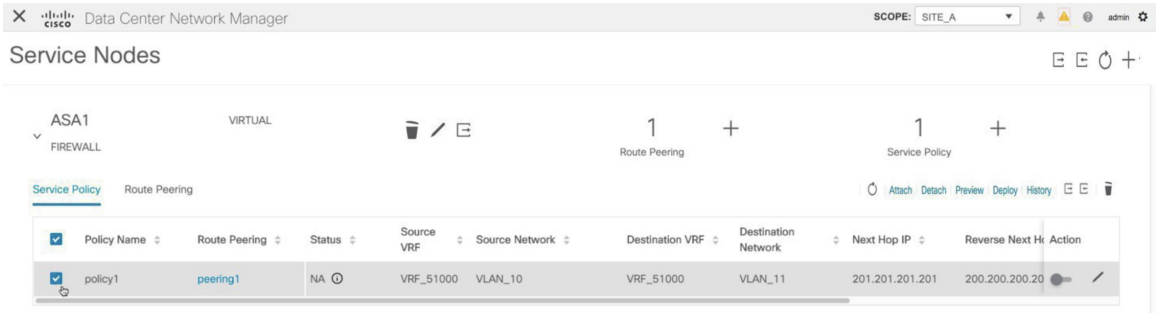

Step 2 Click the toggle button under **Action** to enable this policy.

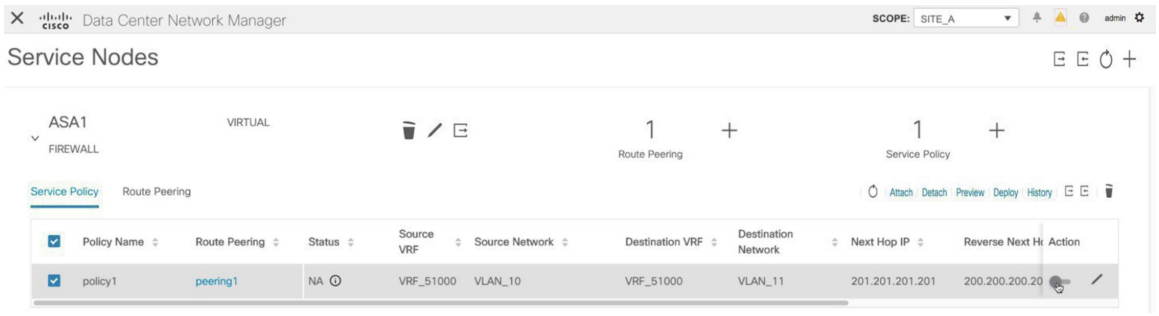

Step 3 Click **Preview** to view the configuration of the selected network.

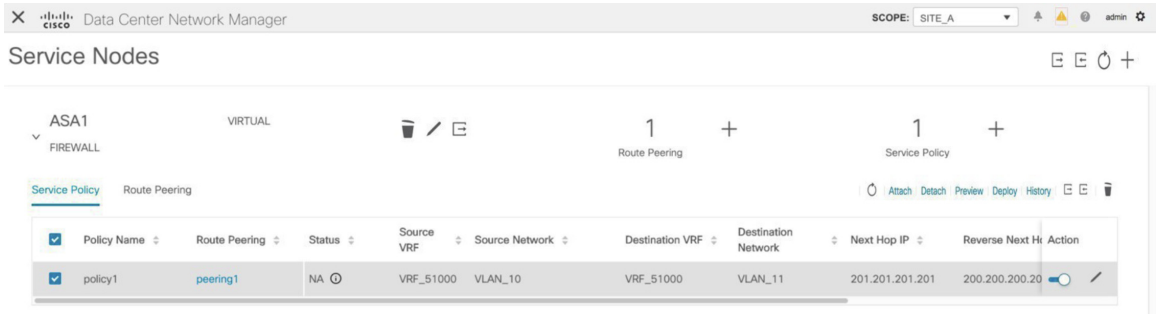

Step 4 Select a switch and a source, destination, or service network, from the drop-down lists to view the intended configuration of a specific source, destination, or service network, on the selected switch. In this window, you can see that there is an access list that will be created with a route map. This configuration will be pushed to the SVI.

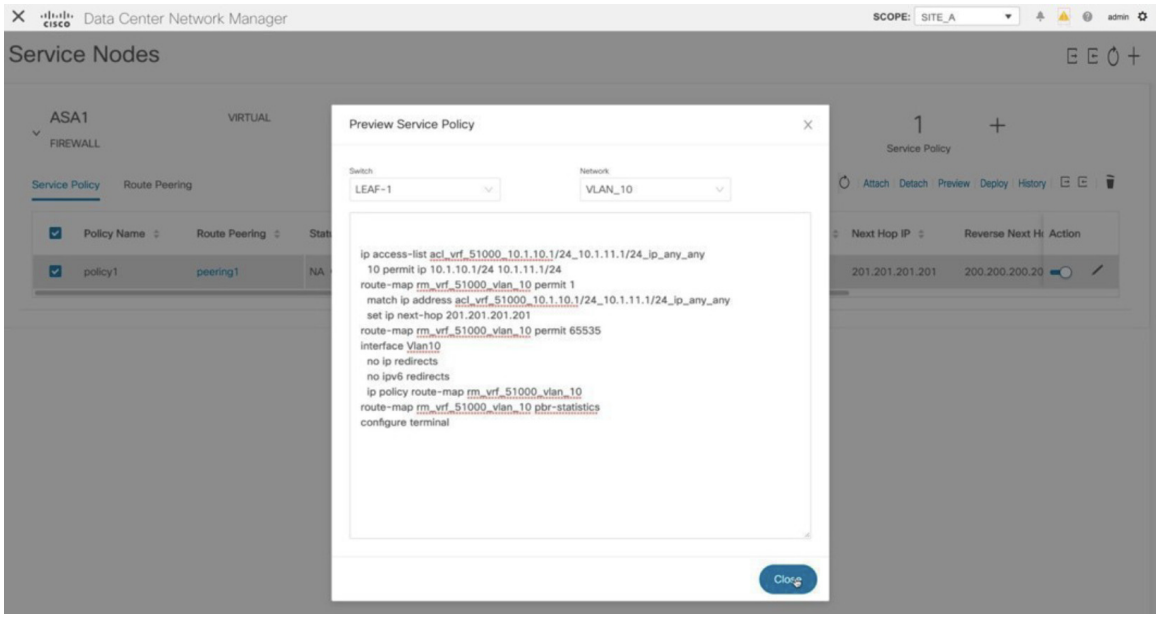

Click **Close** to close the Preview Service Policy window.

Step 5 Click **Deploy** in the **Service Nodes** window to deploy the configuration to the attached switches (service leaf(s)).

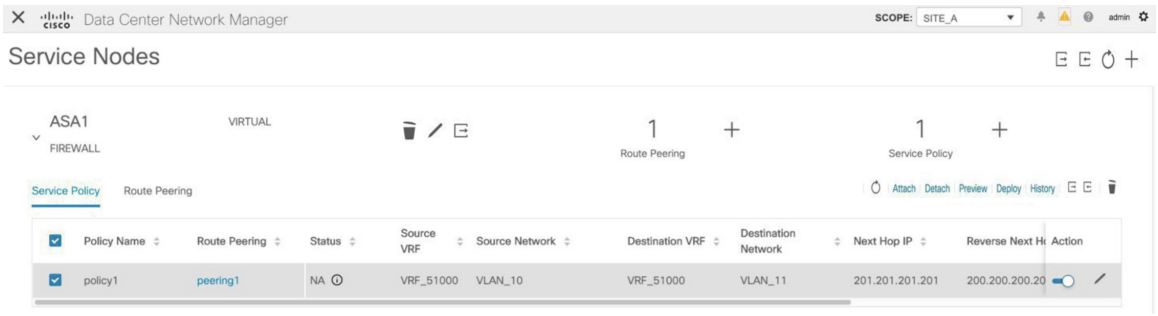

Click the **Deploy** button in the pop-up window to confirm deployment.

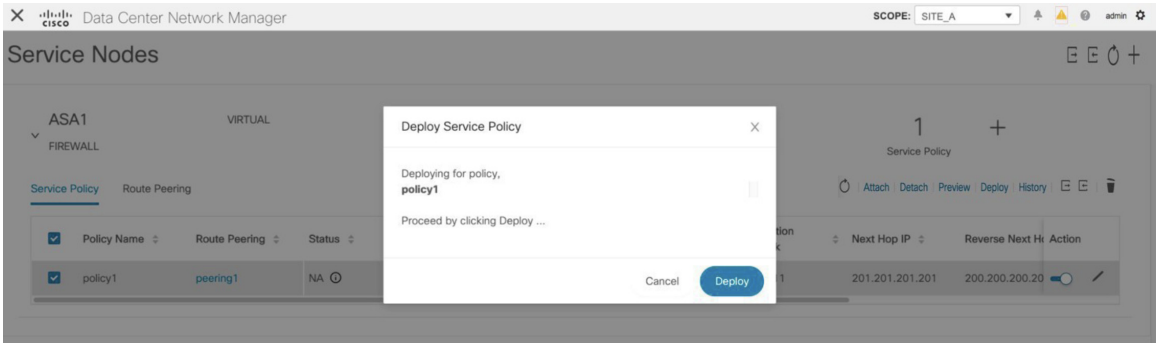

Step 6 Click the **Refresh** icon for the latest policy attachment and deployment status.

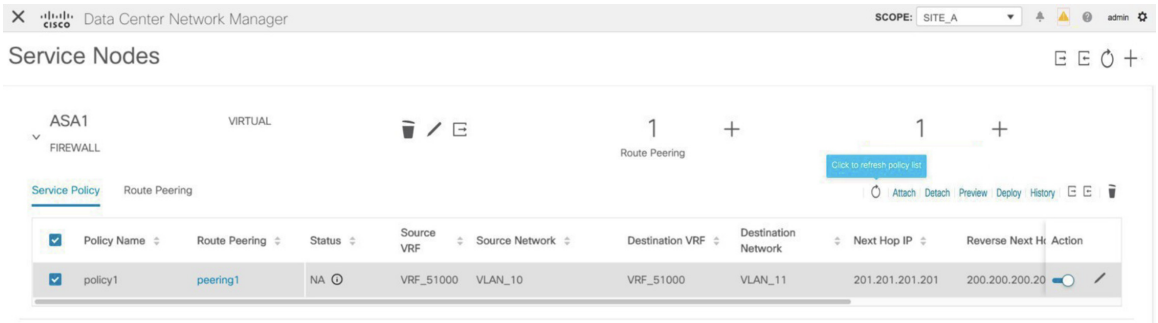

This policy will be pushed to the switches that the source and destination networks are attached to, as well as the service leaf(s). After pushing the policy, the status column shows **In-Sync**.

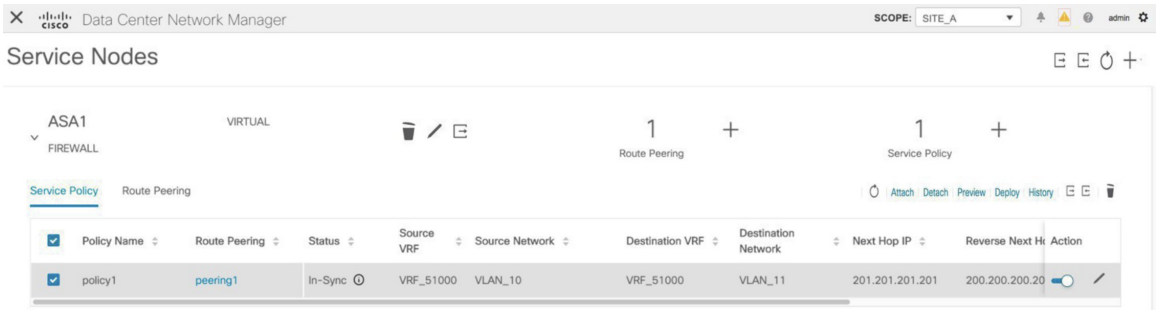

6. View Stats

Now that the respective redirection policies are deployed, ping traffic will be redirected to the firewall.

To visualize this scenario in DCNM, click the icon under the **Stats** column.

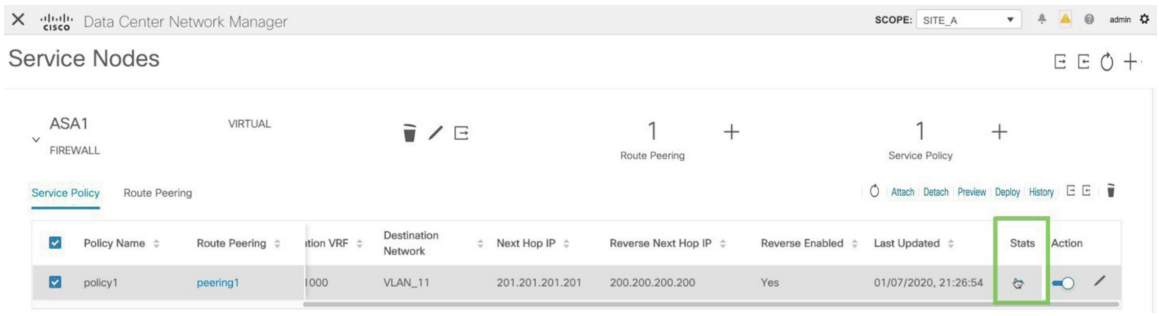

You can view the cumulative statistics for a policy in a specified time range.

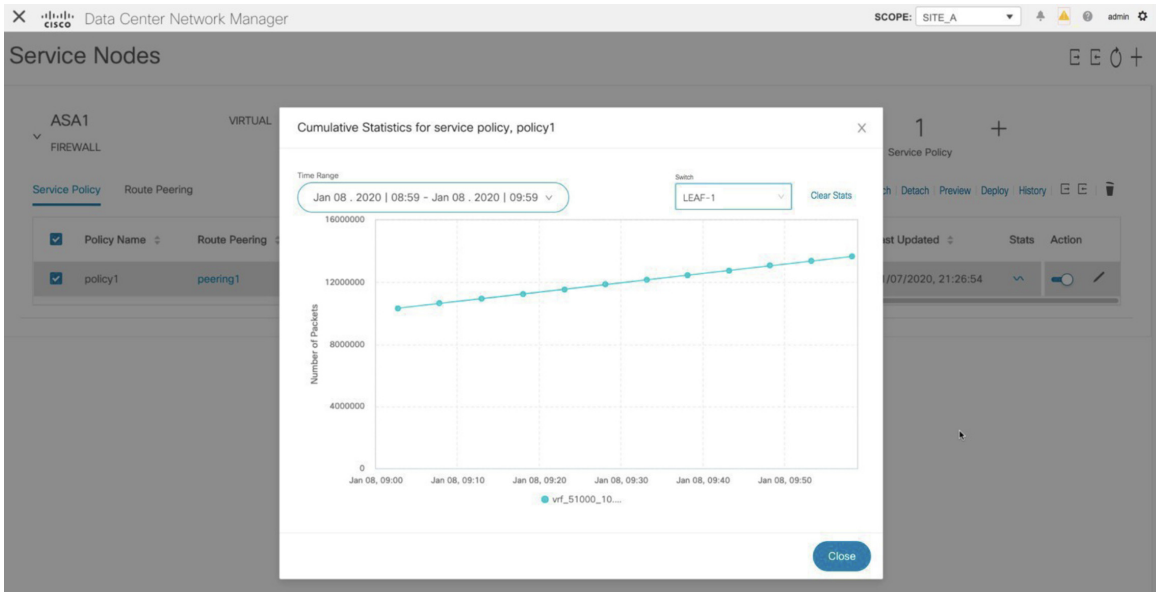

Statistics are displayed for forwarding traffic on the source switch, for reversed traffic on the destination switch, and for traffic in both directions on the service switch.

7. View Traffic Flow in Fabric Builder

The service node in the external fabric is attached to the service leaf, and this external fabric is shown as a cloud icon in the DCNM topology in the fabric builder.

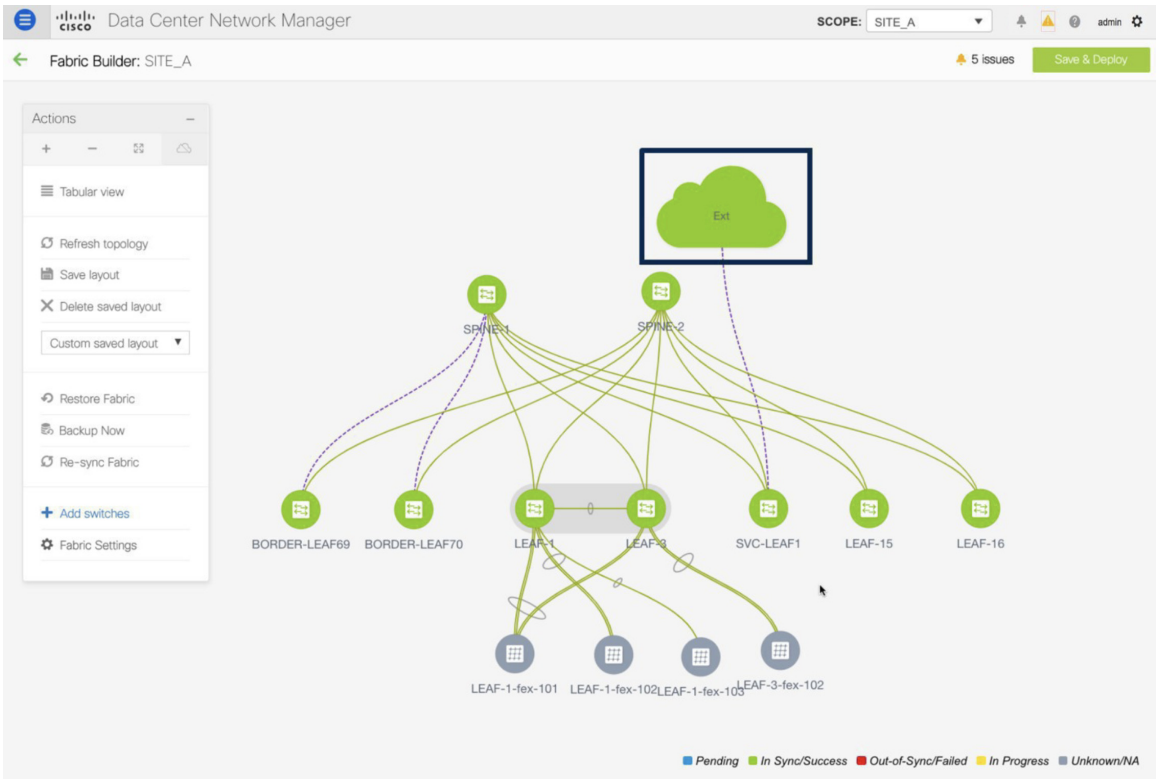

Procedure

Step 1 Click the service leaf and click **Show more flows**. You can see the flows that have been redirected.

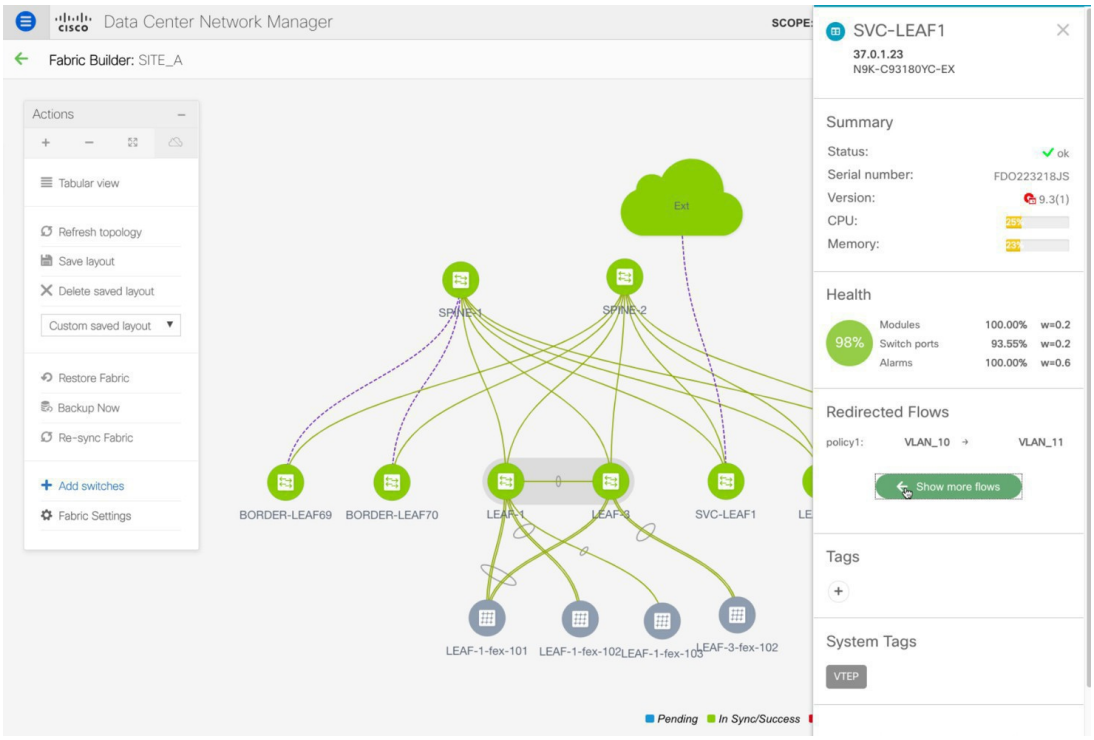

Step 2 Click **Details** in the **Service Flows** window to display attachment details.

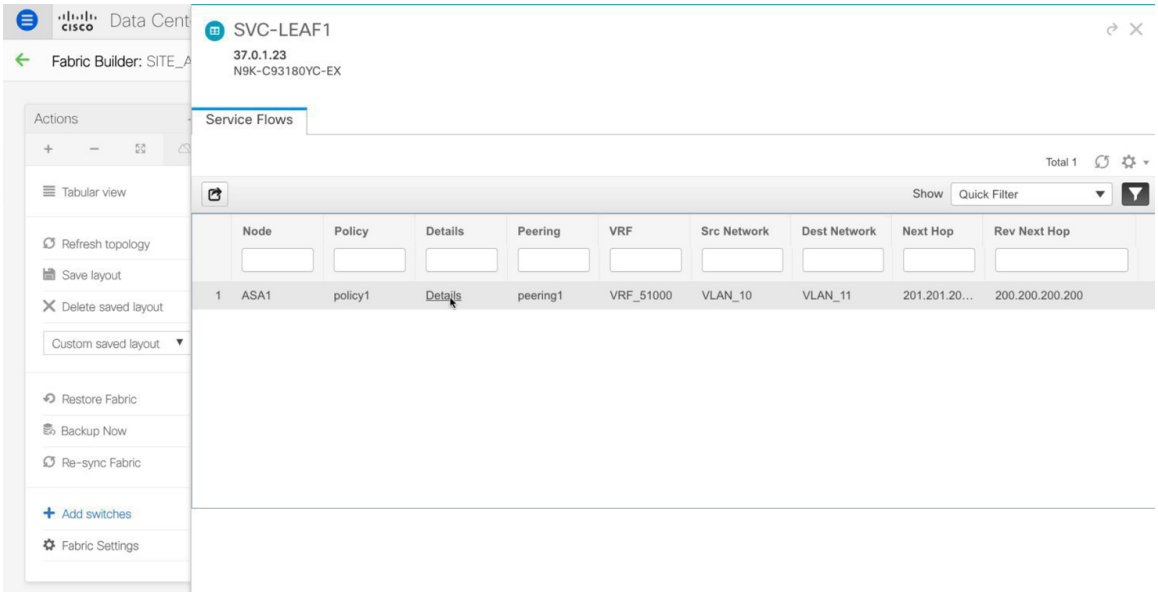

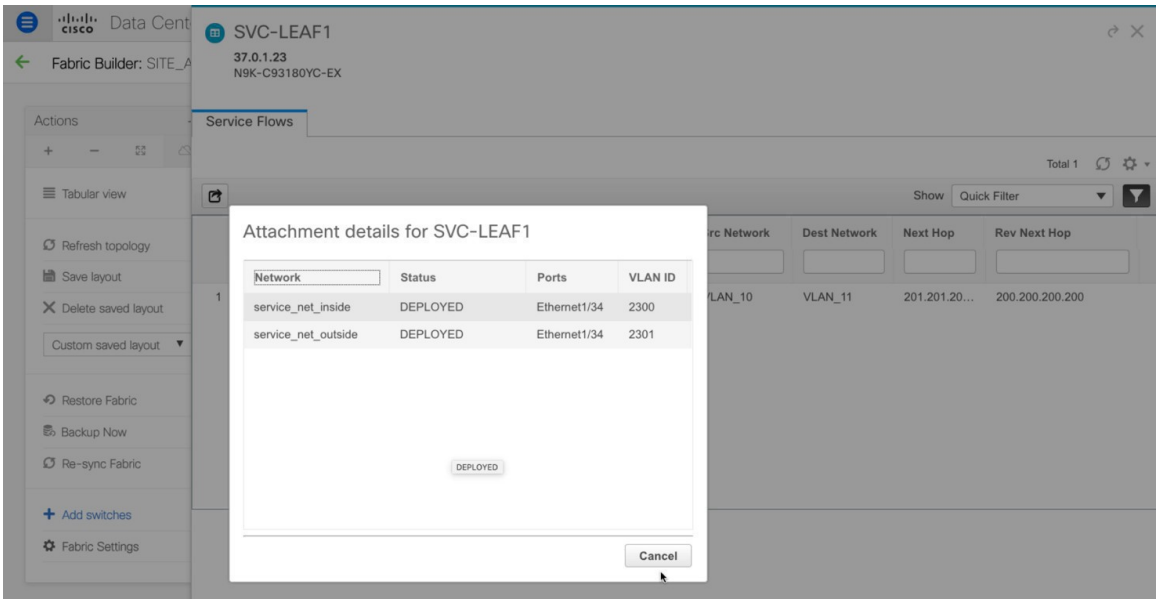

8. Visualize Redirected Flows to Destination in the Topology window

Procedure

Step 1 Click **Topology** and click on leafs to visualize the redirected flows to destination.

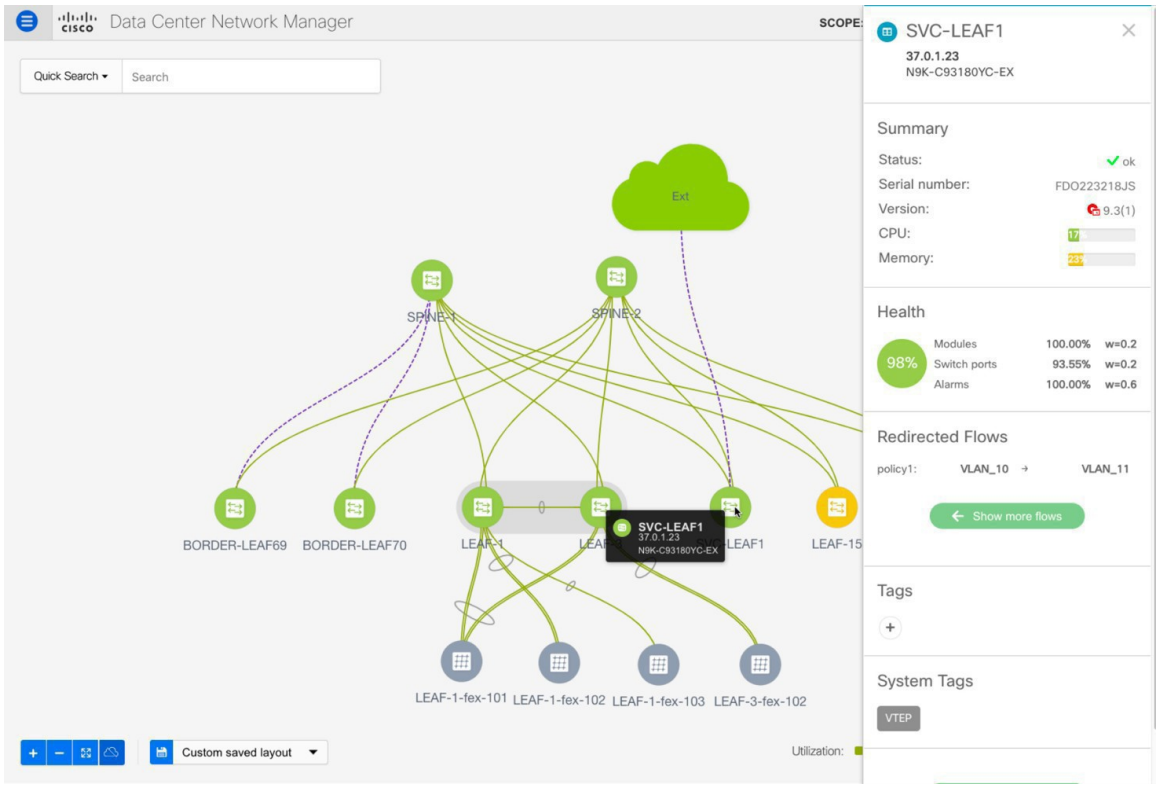

Step 2 Select **Redirected Flows** from the drop-down list.

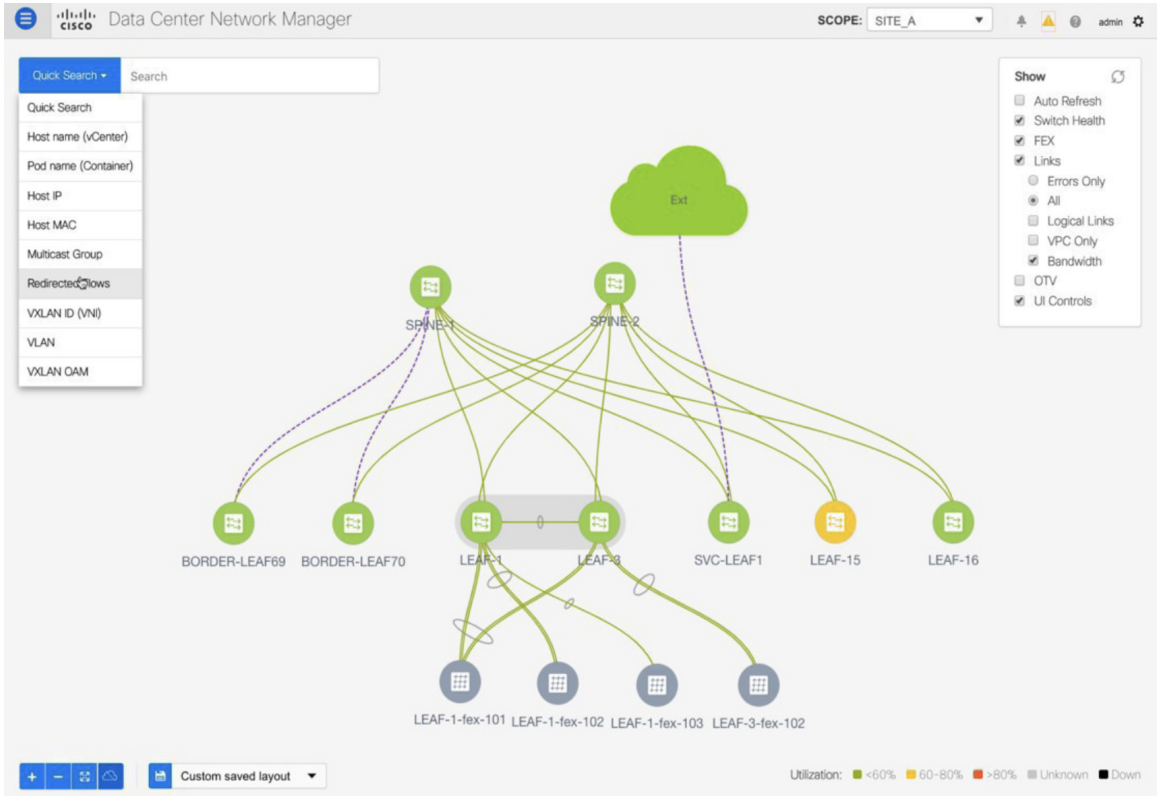

Step 3 Select a policy from the drop-down list or initiate a search by entering a policy name, source network and destination network in the search field. The search field is autopopulated based on your input.

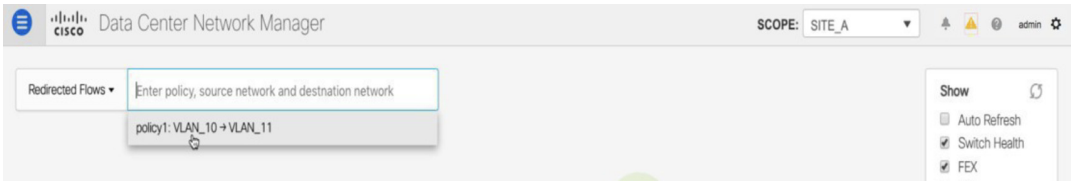

The switches, on which the source and destination network have been attached and the flows have been redirected, are highlighted.

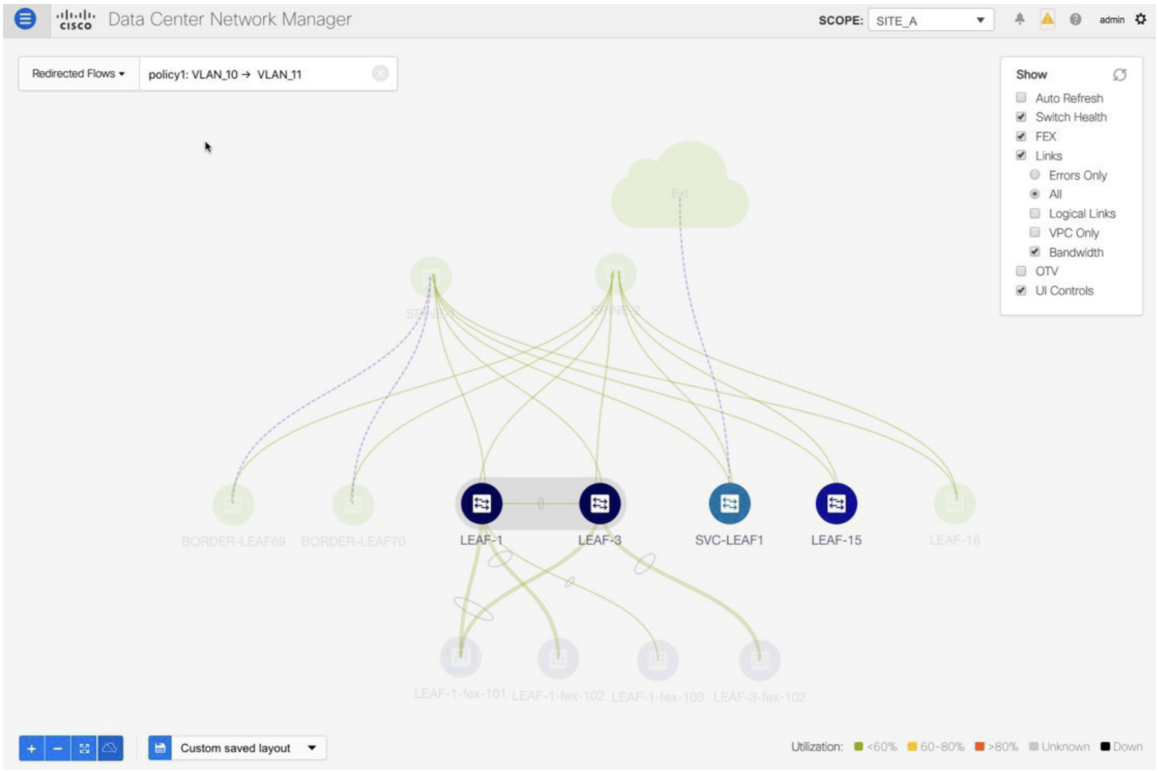

Step 4 The service node is shown as connected by a dotted line to the leaf switch on the topology window. Hover over the dotted line to get more information about the interface.

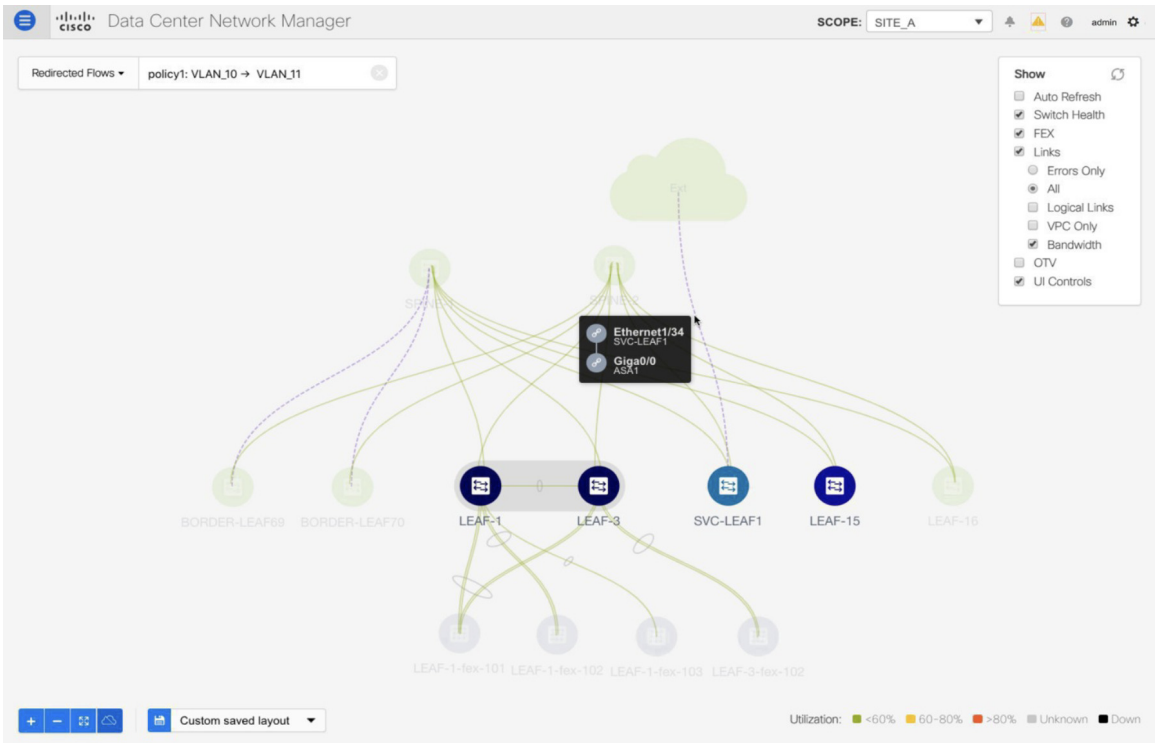

The traffic from **Source** traverses to the service leaf where the firewall is configured.

Based on firewall rules, traffic is allowed to reach the destination, Leaf 15.

Use Case: Inter-tenant Firewall with eBGP Peering

Refer the figure given below for topology details.

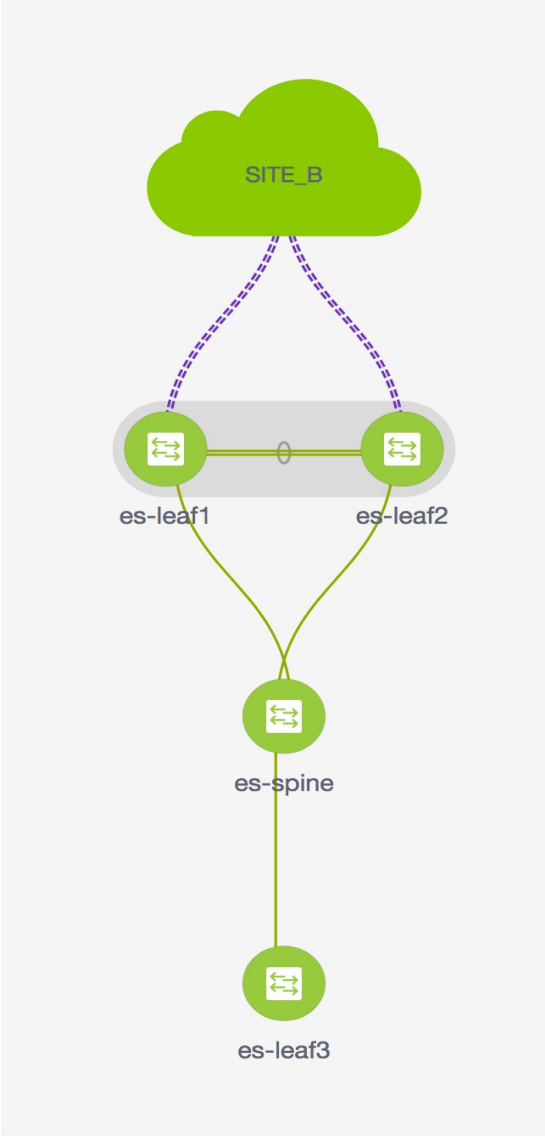

In this topology, es-leaf1 and es-leaf2 are vPC border leaf switches.

Now, let us see how to perform service redirection in DCNM.

Select **Control > Fabrics > Services**.

This use-case consists of the following steps:

Note

- As some steps are similar to the steps given in the Intra-tenant Firewall deployment use- case, reference links have been provided to the steps in that use-case.
- Service policies are not applicable on Inter-tenant firewall deployments.

1. Create Service Node

Procedure

Step 1 From the **Scope** drop-down list, select **Site_A**.

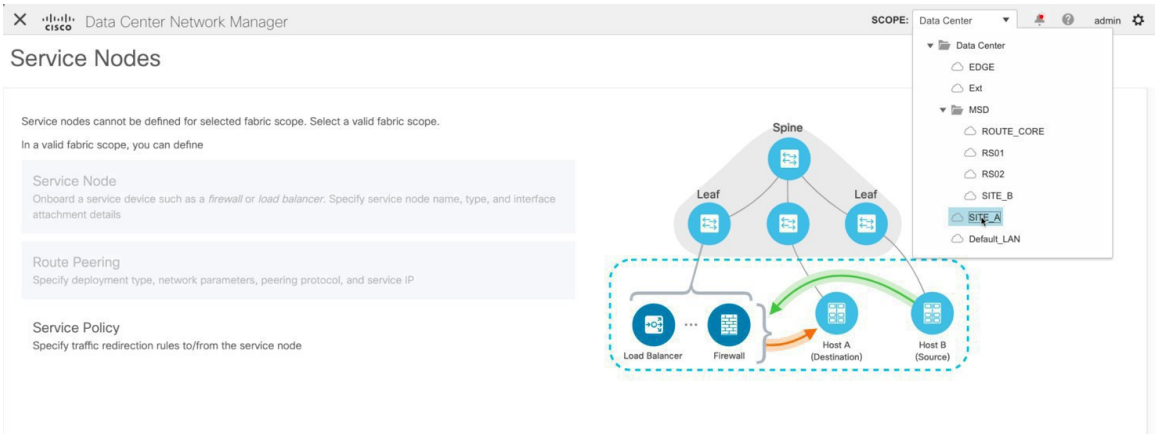

Step 2 Click the **Add** icon in the **Service Nodes** window.

Step 3 Enter the node name and specify **Firewall** in the **Type** dropdown box. The **Service Node Name** has to be unique.

Step 4 From the **Form Factor** drop-down list, select **Virtual**.

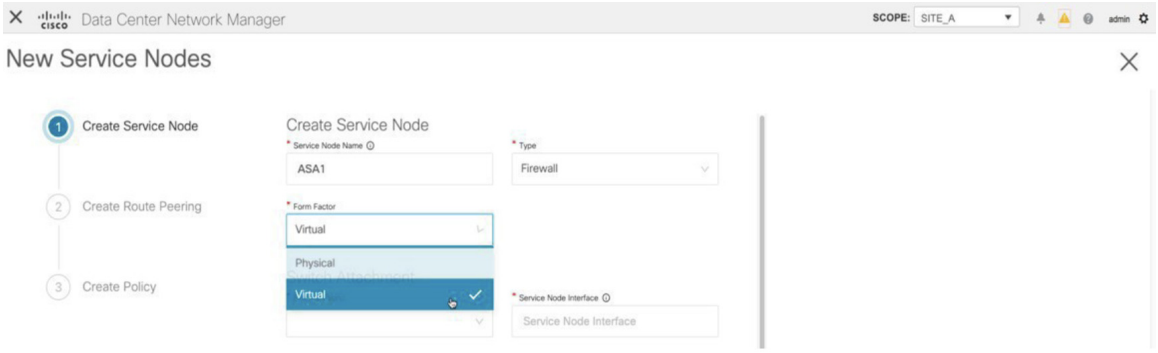

Step 5 In the **Switch Attachment** section, from the **External Fabric** drop-down list, select the external fabric in which the service node (for example, ASA firewall) is located. Note that the service nodes need to belong to the external fabric. This is a prerequisite before creating a service node.

Step 6 Enter the interface name of the service node that will be connected to the service leaf.

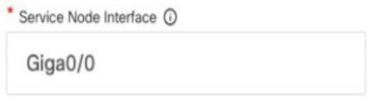

Step 7 Select the attached switch that is the service leaf, and the respective interface on the service leaf.

Step 8 Select the **service_link_trunk** template. DCNM supports trunk, port channel, and vPC link templates. The available link templates in the **Link Template** drop-down list are filtered based on the selected **Attached Switch Interface** type.

Step 9 Specify the **General Parameters** and **Advanced** parameters, if required. Some parameters are pre-filled with the default values.

Step 10 Click **Next** to save the created service node.

Note For more sample screenshots, refer 1. Create Service Node, on page 904 in the Intra-tenant firewall with policy-based routing use case.

2. Create Route Peering

Let us now configure the peering between a service leaf and a service node.

Procedure

Step 1 Enter the peering name and select **Inter-Tenant Firewall** from the **Deployment** drop-down list. From the **Peering Option** drop-down list, select **eBGP Dynamic Peering**.

Step 2 Under **Inside Network**, from the **VRF** drop-down list, select a VRF that already exists and select **Inside Network** under **Network Type**.

Enter the name of the **Service Network** and specify the **Vlan ID**. You can also click **Propose** to allow DCNM to fetch the next available VLAN ID from the specified service network VLAN ID range in the fabric settings. The default **Service Network Template** is **Service_Network_Universal**.

Under the **General Parameters** tab, specify the gateway address for the service network. Specify the **Next Hop IP Address**. This next hop address has to be within the 'inside service network' subnet. Under the **Advanced** tab, the default **Routing Tag** value is 12345.

Step 3 The default Peering Template for eBGP dynamic peering is **service_ebgp_route**.

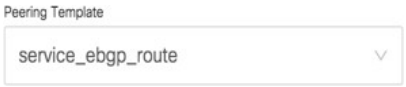

Under the **General Parameters** tab, specify the **Neighbor IPv4** address, **Loopback IP** address, and the **vPC Peer's Loopback IP** address. The border switches are a vPC pair.

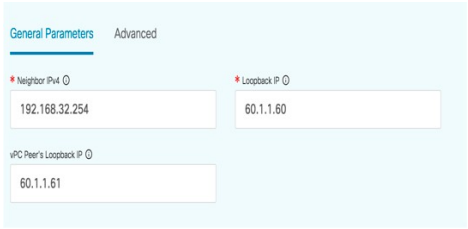

Step 4 Under the **Advanced** tab, specify the **Local ASN** and select the **Advertise Host Routes** checkbox. This local ASN value is used to override the system ASN on the switch and is required to avoid routing loops.

If the **Advertise Host Routes** checkbox is selected, the /32 and /128 routes are advertised. If this checkbox is not selected, the prefix routes will be advertised.

By default, the **Enable Interface** checkbox is selected.

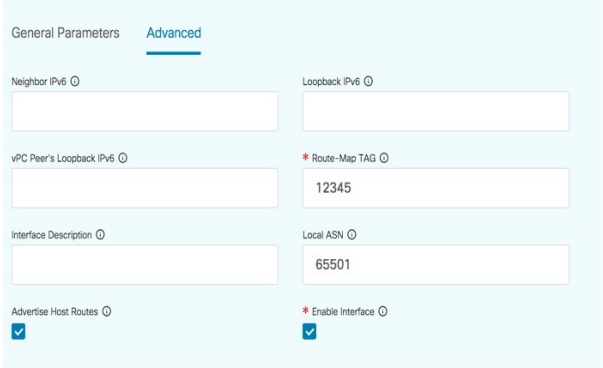

Step 5 Specify the required parameters under **Outside Network** and specify the **Next Hop IP Address for Reverse Traffic**. This next hop address for reverse traffic needs to be within the 'outside service network' subnet.

Step 6 The default Peering Template for eBGP dynamic peering is **service_ebgp_route**.

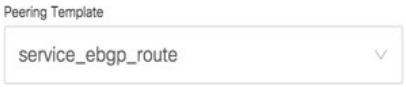

Under the **General Parameters** tab, **Neighbor IPv4** address, **Loopback IP** address, and the **vPC Peer's Loopback IP** address. The leaf switches are a vPC pair.

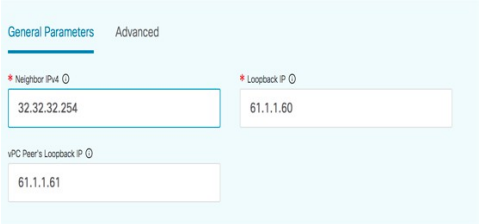

Step 7 Under the **Advanced** tab, specify the **Local ASN** and select the **Advertise Host Routes** checkbox. This local ASN value is used to override the system ASN on the switch and is required to avoid routing loops.

If the **Advertise Host Routes** checkbox is selected, the /32 and /128 routes are advertised. If this checkbox is not selected, the prefix routes will be advertised.

By default, the **Enable Interface** checkbox is selected.

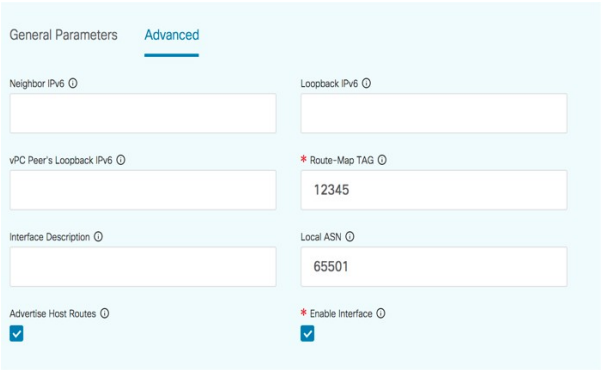

Step 8 Click **Next** to save the created route peering.

3. Deploy Route Peering

Refer 4. Deploy Route Peering, on page 910 of the Intra-Tenant Firewall deployment use-case. Note that **InterTenantFW** is displayed under **Deployment**.

The BGP configuration on the vPC border leaf for this use-case is given below.

```
router bgp 12345
 router-id 10.2.0.1
 address-family l2vpn evpn
  advertise-pip
 neighbor 10.2.0.4
  remote-as 12345
  update-source loopback0
  address-family l2vpn evpn
   send-community
    send-community extended
 vrf myvrf_50001
  address-family ipv4 unicast
   advertise l2vpn evpn
   redistribute direct route-map fabric-rmap-redist-subnet
```

```
   maximum-paths ibgp 2
  address-family ipv6 unicast
   advertise l2vpn evpn
   redistribute direct route-map fabric-rmap-redist-subnet
   maximum-paths ibgp 2
  neighbor 192.168.32.254
   remote-as 9876
  local-as 65501 no-prepend replace-as // Note: This configuration corresponds to the Local
 ASN template parameter value of the service_ebgp_route template of the inside network with
 VRF myvrf_50001. The no-prepend replace-as keyword is generated along with the local-as
command.
  update-source loopback2
  ebgp-multihop 5
  address-family ipv4 unicast
   send-community
   send-community extended
   route-map extcon-rmap-filter-allow-host out
vrf myvrf_50002
 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
 neighbor 32.32.32.254
  remote-as 9876
  local-as 65502 no-prepend replace-as // Note: This configuration corresponds to the Local
 ASN template parameter value of the service_ebgp_route template of the outside network
with VRF myvrf_50002. The no-prepend replace-as keyword is generated along with the local-as
 command.
  update-source loopback3
  ebgp-multihop 5
  address-family ipv4 unicast
   send-community
   send-community extended
   route-map extcon-rmap-filter-allow-host out
```

The loopback interface configuration on the vPC switch es-leaf1 for this use-case is given below. The loopback interfaces in the configuration correspond to the 'Loopback IP' parameter of the **service_ebgp_route** template. Two loopback interfaces are created automatically on each vPC switch for two separate VRF instances using the **Loopback IP** parameter values that are specified in the **service_ebgp_route** template.

```
interface loopback2
 vrf member myvrf_50001
 ip address 60.1.1.60/32 tag 12345
interface loopback3
 vrf member myvrf_50002
 ip address 61.1.1.60/32 tag 12345
```

The loopback interface config on vPC peer switch es-leaf2:

```
interface loopback2
 vrf member myvrf_50001
 ip address 60.1.1.61/32 tag 12345
interface loopback3
 vrf member myvrf_50002
 ip address 61.1.1.61/32 tag 12345
```

Use Case: One-arm Load Balancer

Refer the figure given below for topology details.

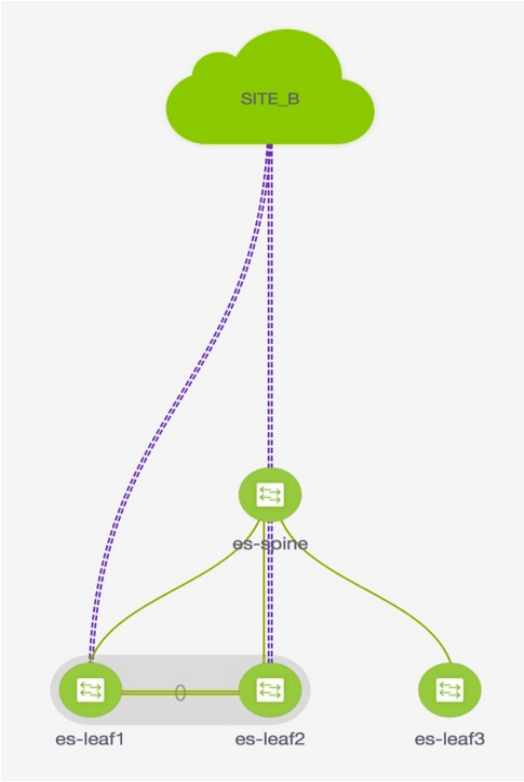

In this topology, es-leaf1 and es-leaf2 are vPC leafs.

Now, let us see how to perform service redirection in DCNM.

Select **Control > Fabrics > Services**.

This use-case consists of the following steps:

Note As some steps are similar to the steps given in the Intra-tenant Firewall deployment usecase, reference links have been provided to the steps in that use-case.

1. Create Service Node

Procedure

Step 1 From the **Scope** drop-down list, select **Site_A**.

Step 2 Click the **Add** icon in the **Service Nodes** window.

Step 3 Enter the node name and specify **Load Balancer** in the **Type** dropdown box. The **Service Node Name** has to be unique.

Step 4 From the **Form Factor** drop-down list, select **Virtual**.

Step 5 In the **Switch Attachment** section, from the **External Fabric** drop-down list, select the external fabric in which the service node (for example, ASA firewall) is located. Note that the service nodes need to belong to the external fabric. This is a prerequisite before creating a service node.

Step 6 Enter the interface name of the service node that will be connected to the service leaf.

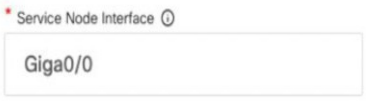

Step 7 Select the attached switch that is the service leaf, and the respective interface on the service leaf.

Step 8 Select the **service_link_trunk** template. DCNM supports trunk, port channel, and vPC link templates. The available link templates in the **Link Template** drop-down list are filtered based on the selected **Attached Switch Interface** type.

Step 9 Specify the **General Parameters** and **Advanced** parameters, if required. Some parameters are pre-filled with the default values.

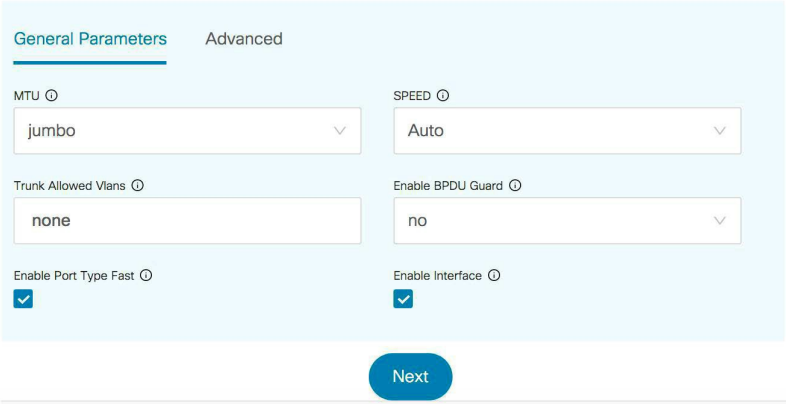

Step 10 Click **Next** to save the created service node.

Note For more sample screenshots, refer 1. Create Service Node, on page 904 in the Intra-tenant firewall with policy-based routing use case.

2. Create Route Peering

Let us now configure peering between a service leaf and a service node. In this use-case, we configure static route peering.

Procedure

Step 1 Enter the peering name and select **One-Arm Mode** from the **Deployment** drop-down list. Also, from the **Peering Option** dropdown list, select **Static Peering**.

Step 2 Under **First Arm**, specify the required values. From the **VRF** dropdown list, select a VRF that already exists and select **First Arm** under **Network Type**.

Step 3 Enter the name of the **Service Network** and specify the **Vlan ID**. You can also click Propose to allow DCNM to fetch the next available VLAN ID from the specified service network VLAN ID range in the fabric settings. The default **Service Network Template** is **Service_Network_Universal**.

Under the **General Parameters** tab, specify the gateway address for the service network. Specify the **Next Hop IP Address**. This next hop address has to be within the first arm's subnet. Under the **Advanced** tab, the default **Routing Tag** value is 12345.

Step 4 The default **Peering Template** is **service_static_route**. Add routes, as required, in the **Static Routes** field.

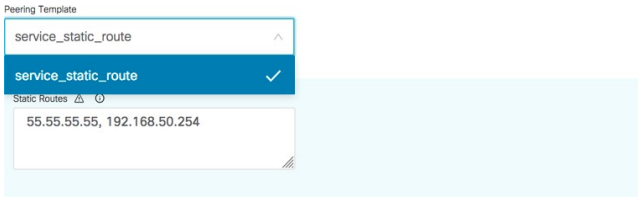

Step 5 Specify the **Next Hop IP Address** for Reverse Traffic.

Step 6 Click **Next** to save the created route peering.

3. Create Service Policy

Refer 3. Create Service Policy, on page 908 in the Intra-Tenant Firewall deployment use-case.

4. Deploy Route Peering

Refer 4. Deploy Route Peering, on page 910 in the Intra-tenant Firewall deployment use-case. Note that **OneArmADC** is displayed under **Deployment**.

5. Deploy Service Policy

Refer 5. Deploy Service Policy, on page 912 in the Intra-tenant Firewall deployment use-case. However, as there are two servers in this load balancer use-case, two service policies have to be defined with each server network.

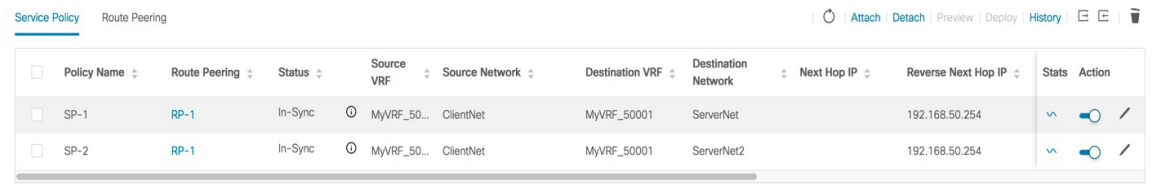

6. View Stats

Refer 6. View Stats, on page 915 in the Intra-Tenant Firewall deployment use-case.

7. View Traffic Flow in Fabric Builder

Refer 7. View Traffic Flow in Fabric Builder, on page 915 in the Intra-Tenant Firewall deployment use-case.

8. Visualize Redirected Flows to Destination in the Topology window

Refer 8. Visualize Redirected Flows to Destination in the Topology window, on page 918 in the Intra-Tenant Firewall deployment use-case.

The VRF configuration on the service leaf is as given below.

```
interface Vlan2000
 vrf member myvrf_50001
 ip policy route-map rm_myvrf_50001

interface Vlan2306
 vrf member myvrf_50001
vrf context myvrf_50001
vni 50001
 ip route 55.55.55.55/32 192.168.50.254 // Note: This is the static route
 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 12345
 vrf myvrf_50001
  address-family ipv4 unicast
   advertise l2vpn evpn
   redistribute direct route-map fabric-rmap-redist-subnet
   redistribute static route-map fabric-rmap-redist-static
   maximum-paths ibgp 2
  address-family ipv6 unicast
   advertise l2vpn evpn
   redistribute direct route-map fabric-rmap-redist-subnet
   redistribute static route-map fabric-rmap-redist-static
   maximum-paths ibgp 2
```

PART VI

Public Cloud Connectivity

- Connecting Cisco Data Center and a Public Cloud, on page 935

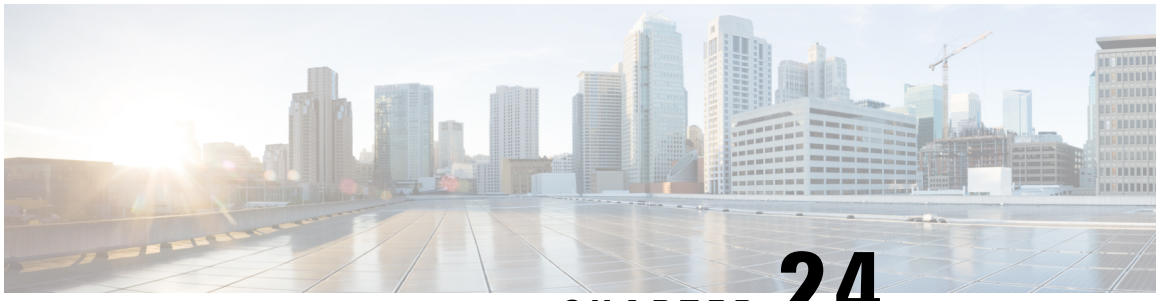

CHAPTER 24

Connecting Cisco Data Center and a Public Cloud

• Connecting Cisco Data Center and a Public Cloud, on page 935

Connecting Cisco Data Center and a Public Cloud

This section explains the 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 Cisco CSR 1000v. This functionality is a preview feature in Cisco DCNM Release 11.2(1). After an inline upgrade to Cisco DCNM Release 11.3(1), this feature is enabled by default.

Topology Overview

Figure 13: Topology Overview

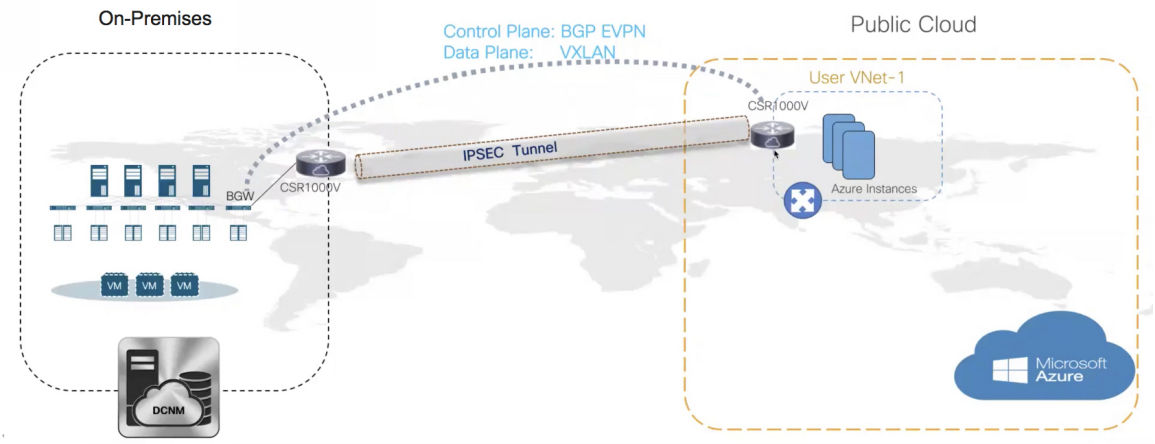

The on-premise data center has the required switches. One of these switches is a border gateway (BGW) that interfaces with an core router for WAN connectivity to the public cloud. The Cisco CSR 1000v is the core router in this use case. You can import this core router into an external fabric in Cisco DCNM. The following figure depicts the sample topology that is employed.

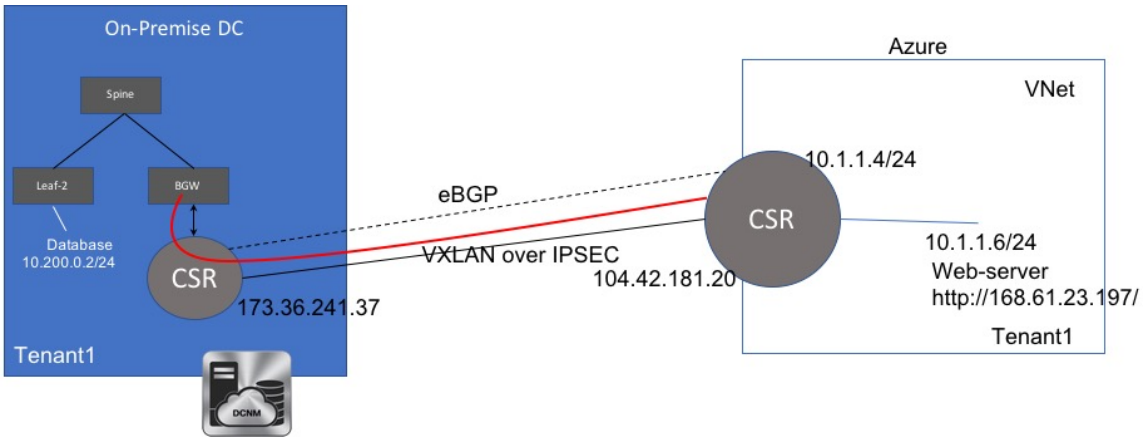

In this example, we list the tasks that are required to provide a layer-3 connectivity between a VM behind standalone leaf and a VM in the Microsoft Azure cloud in a specific user VNET.

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:

- 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 955 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. Set the polling time.
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.

Setting the Polling Time

Cisco DCNM queries the on-premises core router and updates the state of the routing table depending on the polling time you set. To set the polling time 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 polling time in the **private_public_cloud_connectivity.stats.polling_time** field.

The value is in milliseconds.

```
# Private and public cloud connectivity
#
                                    preview_features.enable true

        private_public_cloud_connectivity.stats.polling_time 300000


#_____________________________________________________________________
```

Step 4 Click **Apply Changes**.

Step 5 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 **IOS XE** from the drop-down list, and click the **CSR** radio button.
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 the role of the router to **Core Router**. Right-click the router, choose **Set role > Core Router**.

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 **IOS XE** from the drop-down list, and click the **CSR** radio button.
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

Set the role of the router to **Core Router**. Right-click the router, choose **Set role > Core Router**.

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.

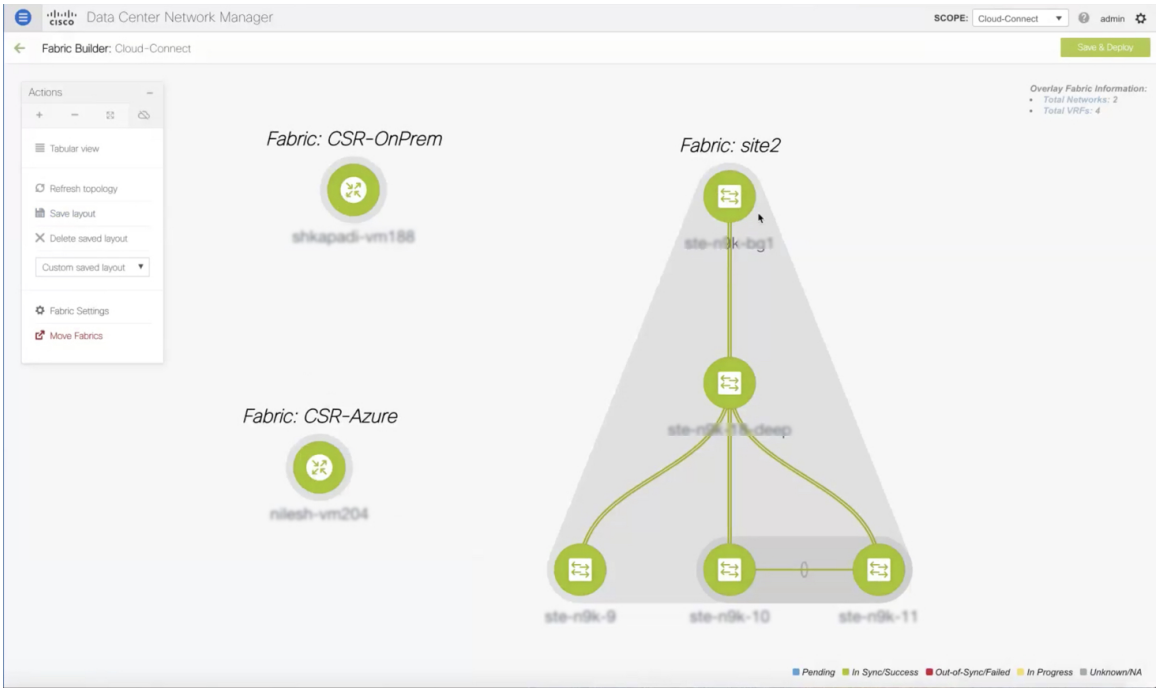

What to do next

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 **Inter-Fabric** link type from the drop-down list.
Link Sub-Type	Choose the **MULTISITE_UNDERLAY** link sub-type from the drop-down list.
Link Template	Choose the **csr_ext_multisite_underlay_setup** link template from the drop-down list. **Note** This template is available only after you enable the preview functionality and restart the DCNM.
Source Fabric	Choose **site2** as the source fabric from the drop-down list.
Destination Fabric	Choose **CSR-OnPrem** 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.

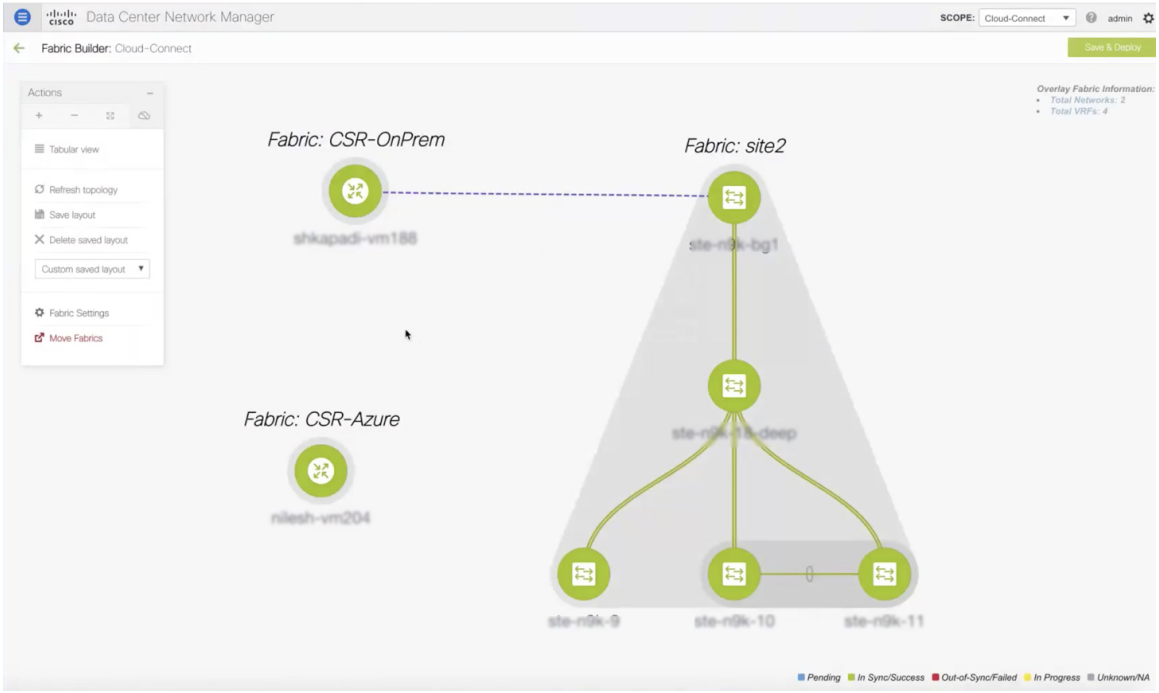

What to do next

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 **Inter-Fabric** link type from the drop-down list.
Link Sub-Type	Choose the **BGP_OVER_IPSEC** link sub-type from the drop-down list.
Link Template	Choose the **csr_link_template** link template from the drop-down list.
Source Fabric	Choose **CSR-OnPrem** as the source fabric from the drop-down list.
Destination Fabric	Choose **CSR-Azure** 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.

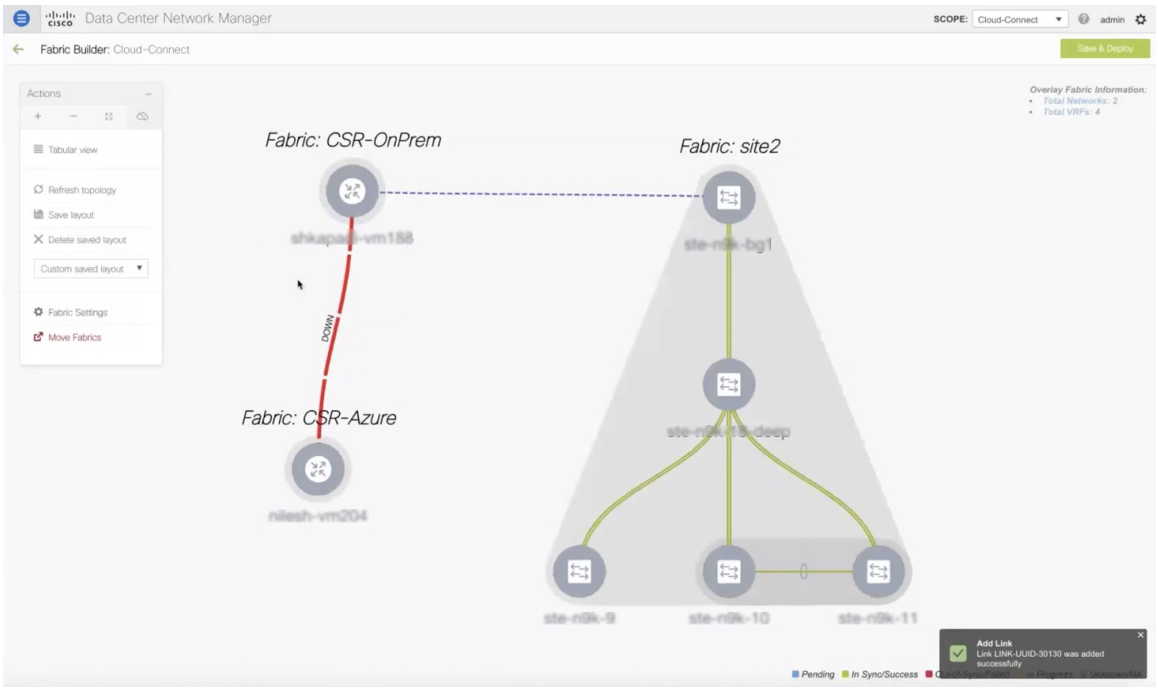

What to do next

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 **Inter-Fabric** link type from the drop-down list.
Link Sub-Type	Choose the **MULTISITE_OVERLAY** link sub-type from the drop-down list.

Field	Description
Link Template	Choose the **csr_ext_evpn_multisite_overlay_setup** link template from the drop-down list.
Source Fabric	Choose **site2** as the source fabric from the drop-down list.
Destination Fabric	Choose **CSR-Azure** 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. **Note** 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 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 subnet.
NEIGHBOR_IP	Enter the IPV4 address of the destination interface.

Step 5 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.

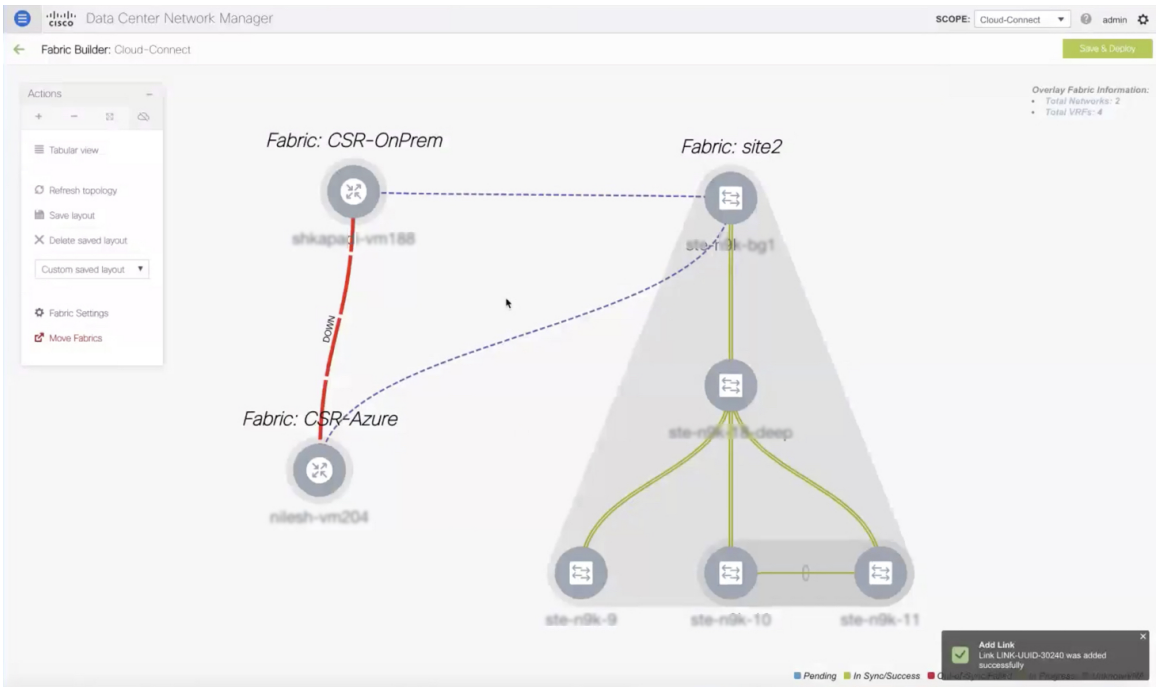

What to do next

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.

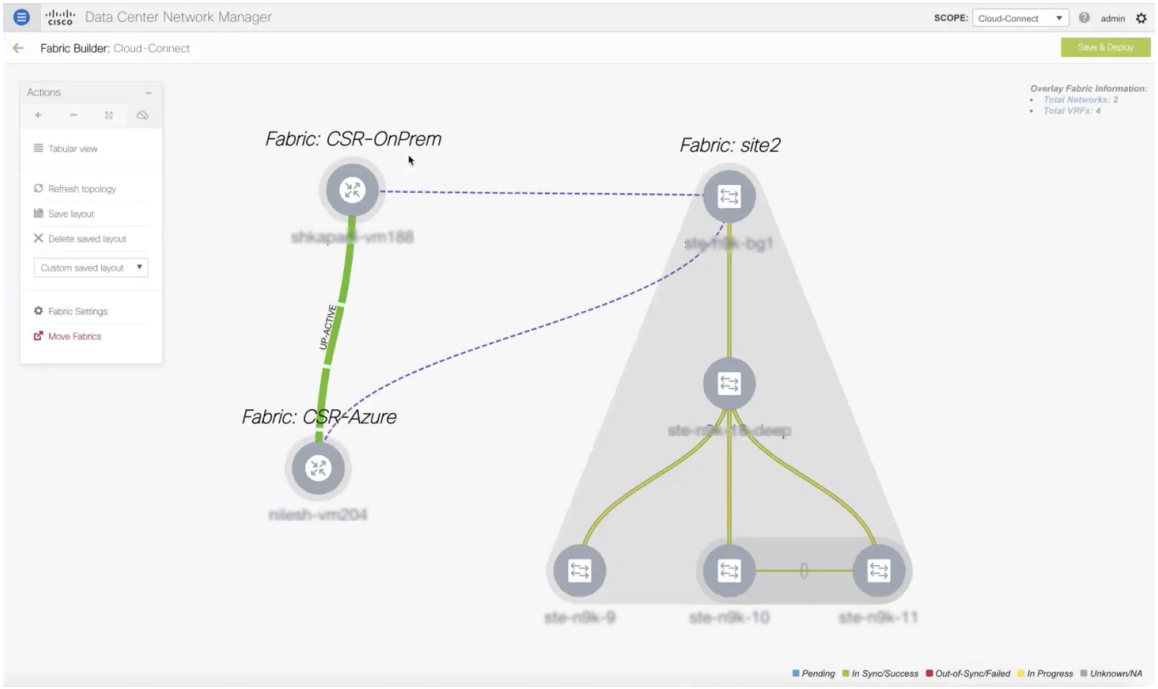

What to do next

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.

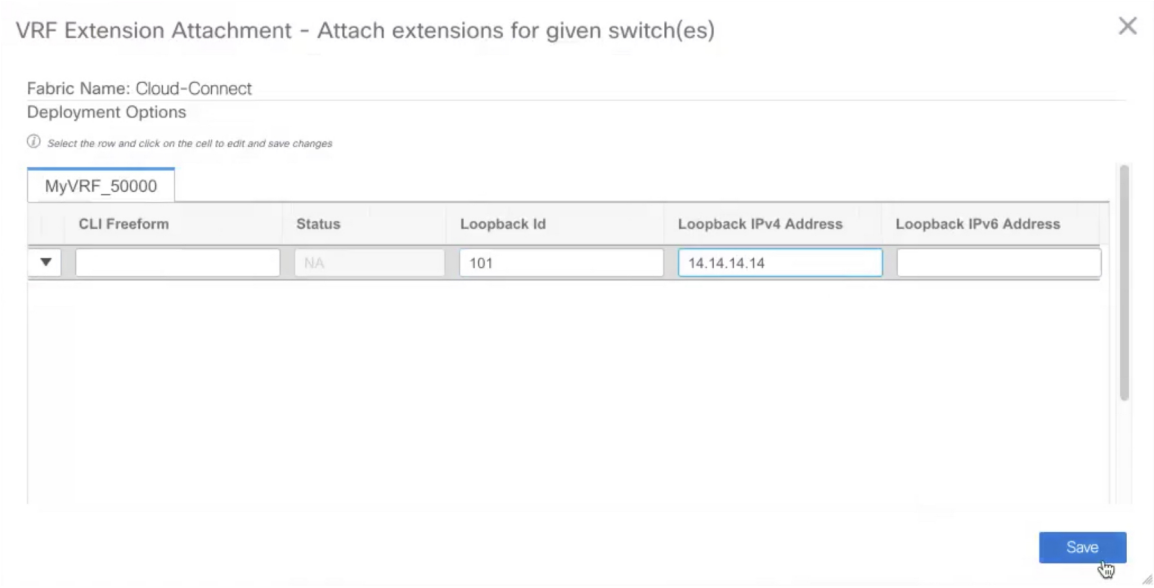

Step 7 Click **Save**.

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.

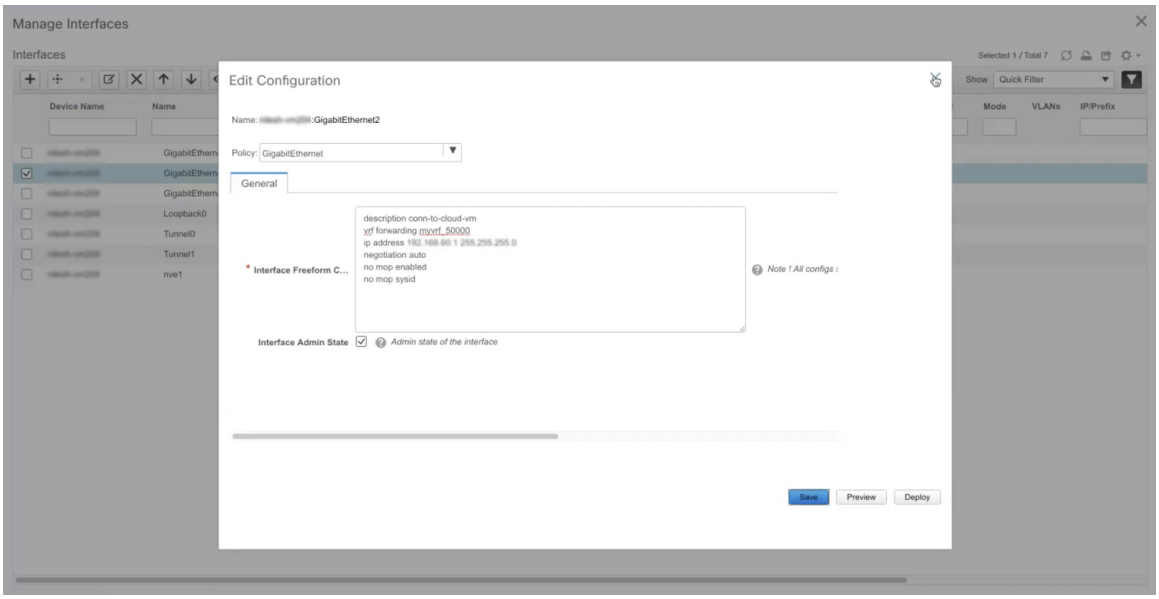

The fabric topology window appears.

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.

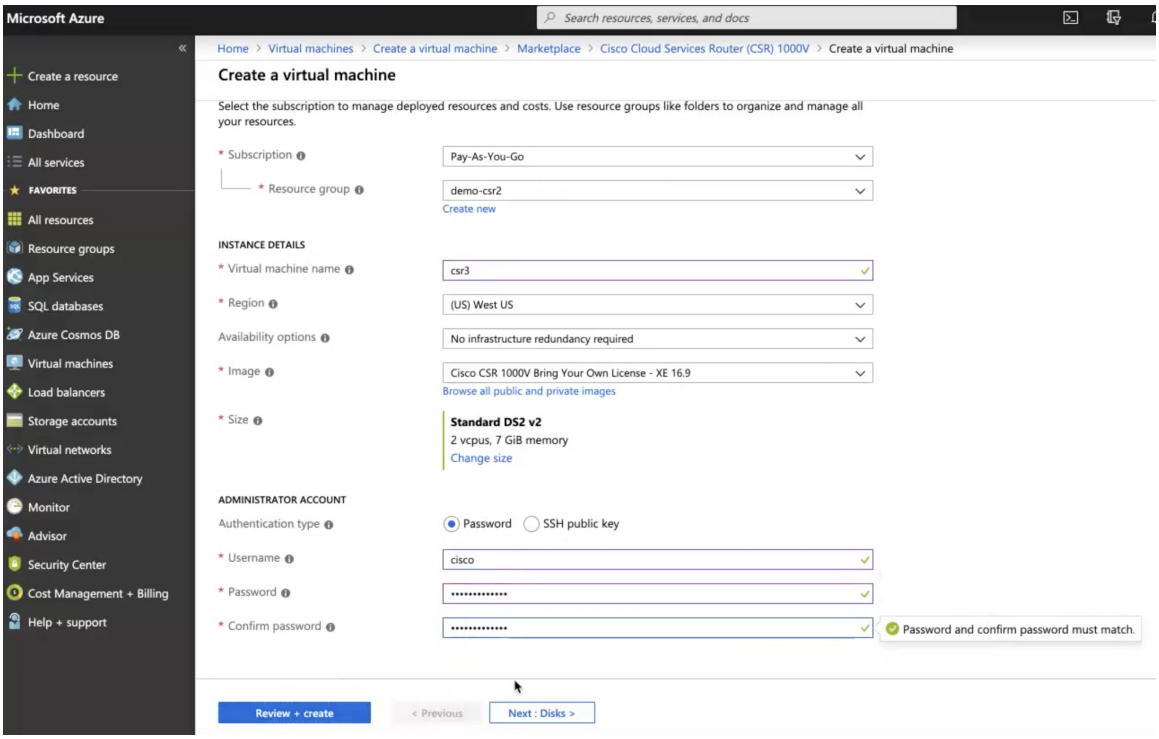

Step 11 Click **Next : Disks >**.

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.

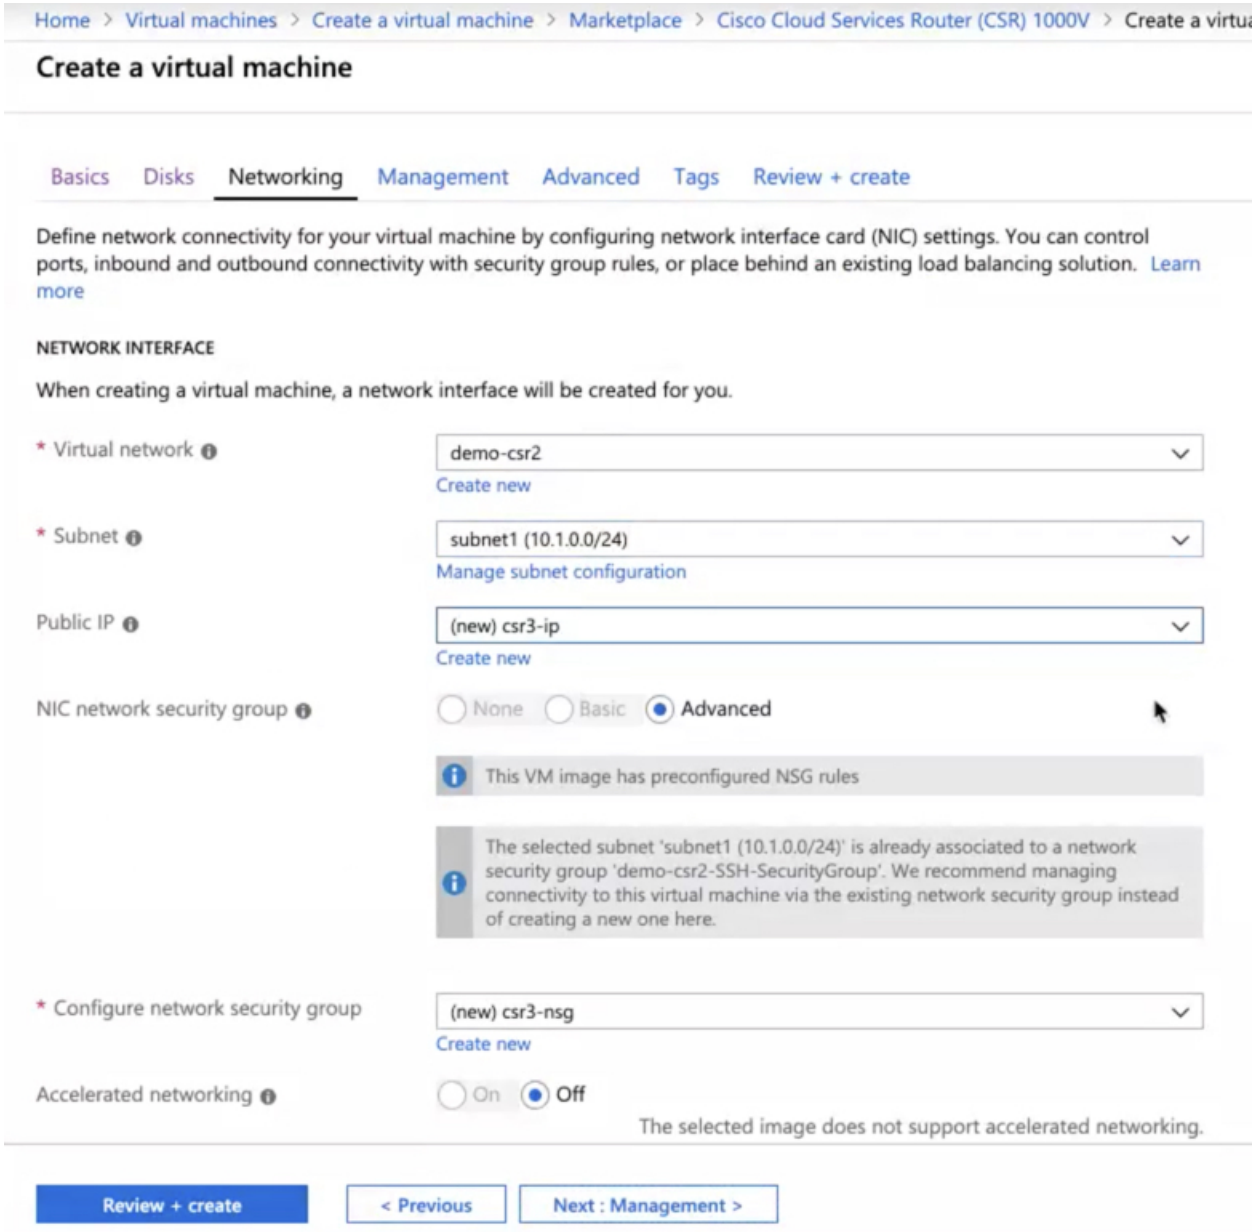

Step 16 Use the default values in other fields.

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.

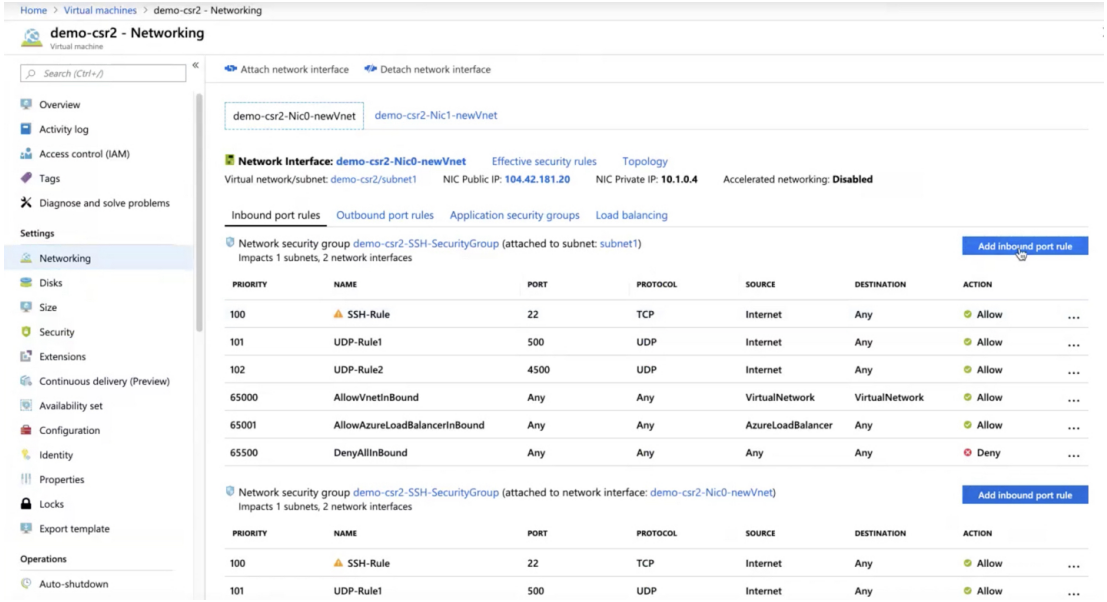

- Create routes in the **Routes** setting of the VM to create traffic routes between the on-prem data center and Microsoft Azure. You can use the default route to redirect traffic from the VNet to Cisco CSR 1000v.

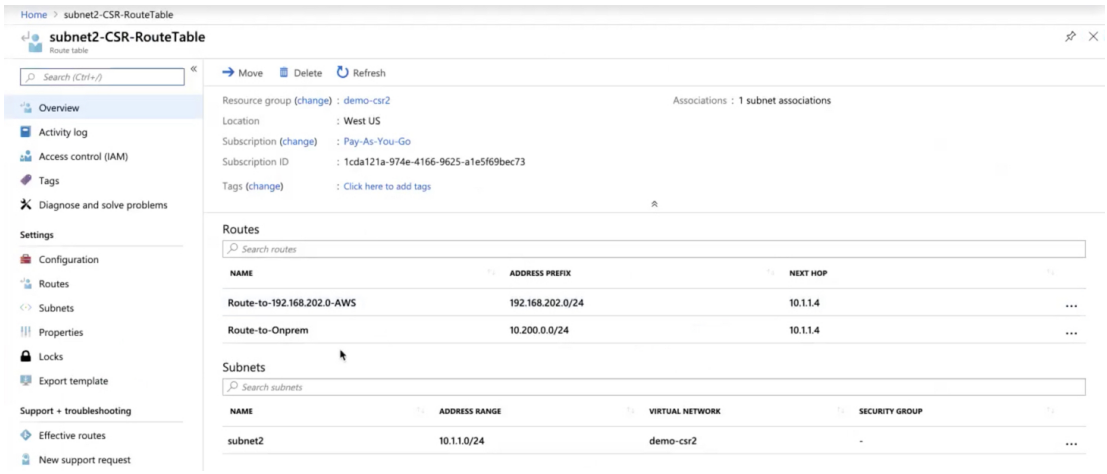

See *Cisco CSR 1000v Deployment Guide for Microsoft Azure* for more information.

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 959 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.

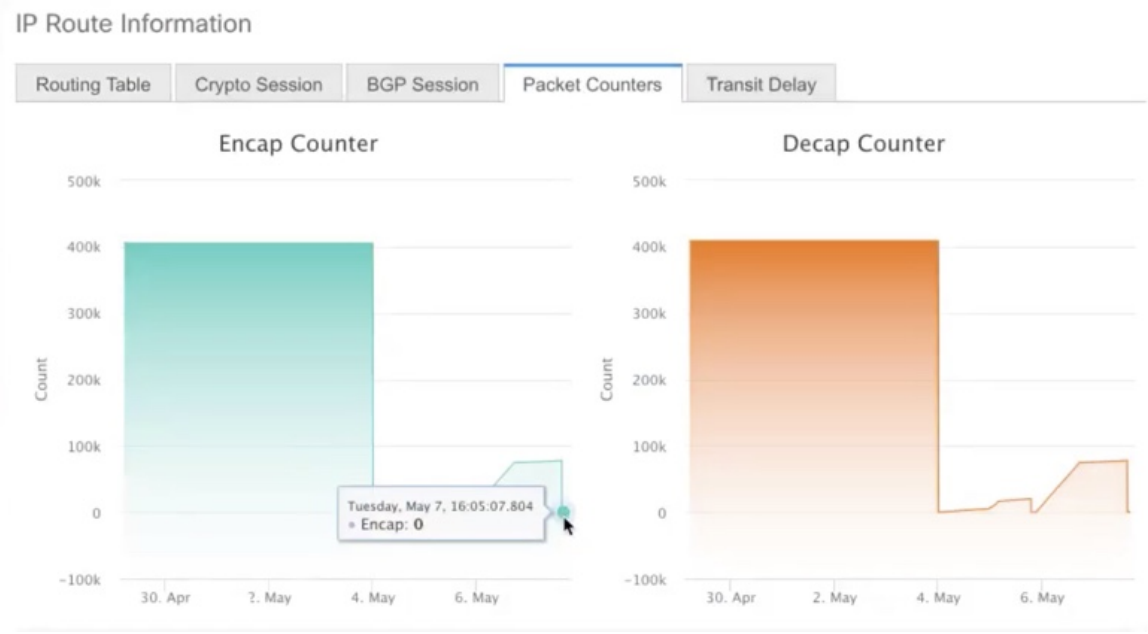

PART VII

Easy Provisioning of MSDC Deployments

- Managing BGP-Based Routed Fabrics, on page 963

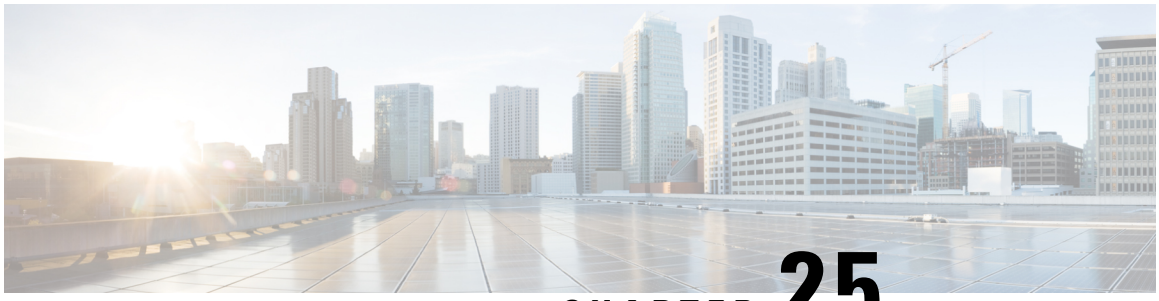

CHAPTER 25

Managing BGP-Based Routed Fabrics

This chapter describes how to configure a typical spine-leaf based routed fabric with eBGP as the routing protocol of choice. This is the preferred deployment choice for Massively Scalable Data Center (MSDC) networks. Both Single-AS and Multi-AS options are supported. A routed fabric has no layer-2 stretch or subnet stretch across leafs. In other words, networks are localized to a pair of leafs or a rack, with leafs hosting the default gateway for the directly attached server workloads. Subnet advertisement across racks are communicated over eBGP via the spine thereby providing any-to-any reachability within the routed fabric.

- **Creating an eBGP-based Fabric**, on page 963
- Adding Switches to a Fabric, on page 973
- Deploying Fabric Underlay eBGP Policies, on page 987
- Deploying Networks in eBGP-based Fabrics, on page 988

Creating an eBGP-based 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.

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.

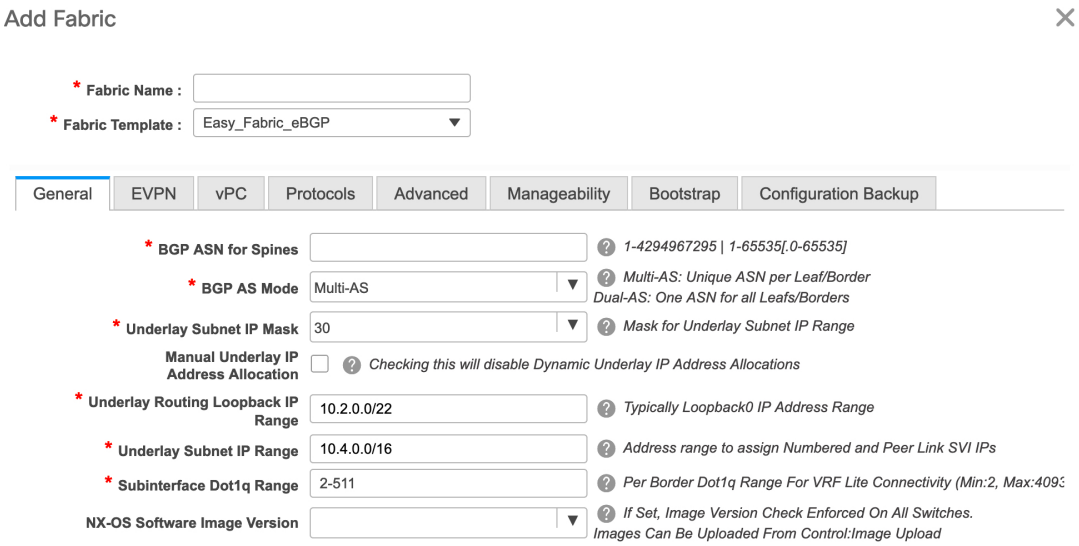

3. The **General** tab is displayed by default. The fields in this tab are:

   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.

   Manual Underlay IP Address Allocation – Select this check box to disable Dynamic Underlay IP Address Allocations.

   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**. The Enable EVPN VXLAN Overlay option must be explicitly disabled. Note that this checkbox is enabled by default. This option should be enabled only for use-cases where customers want to build an eBGP-underlay/overlay based VXLAN EVPN fabric.

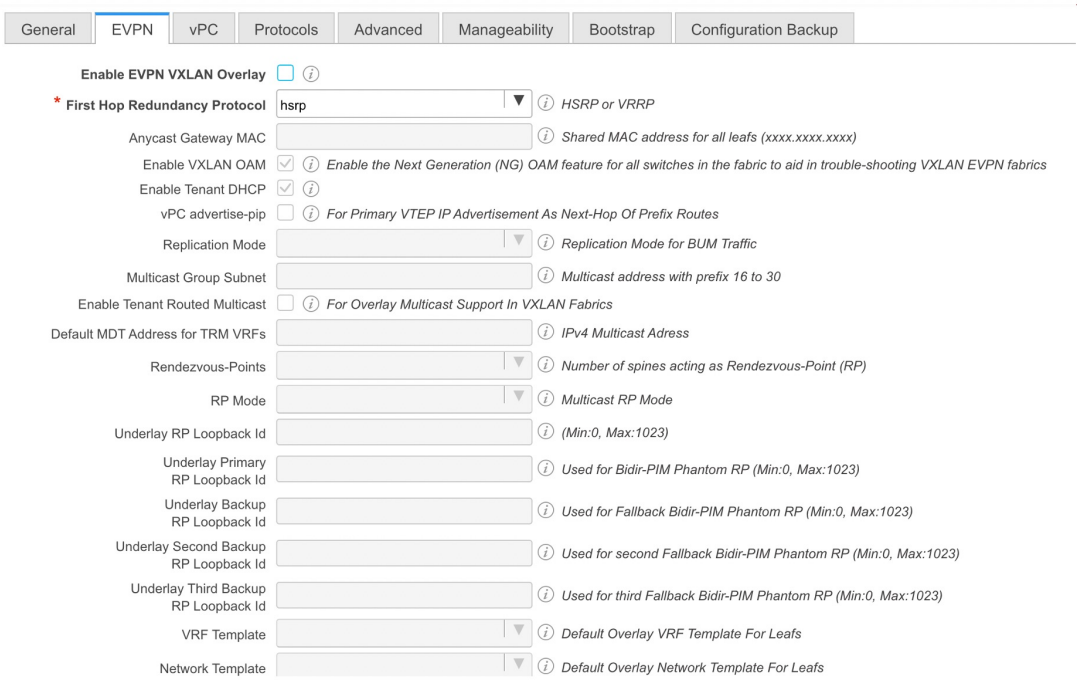

Routed Fabric: In a Routed Fabric, once the IP reachability between the spine—leaf network has been established, you can easily create and deploy networks on the leafs using either HSRP or VRRP as the First-Hop Routing Protocol (FHRP) of choice. For more information, see Overview of Networks in a Routed Fabric, on page 988.

When you create an eBGP Routed 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.

Note that **Routed_Network_Universal Template** is only applicable to a Routed Fabric.

First Hop Redundancy Protocol: Specifies the FHRP protocol. Choose either **hsrp** or **vrrp**. This field is only applicable to a Routed Fabric.

Note

- After a network has been created, you cannot change this fabric setting. You should delete all netw and then change the FHRP setting.
- The rest of the fields in the EVPN tab section are only applicable if you enable the EVPN VXLA Overlay.

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

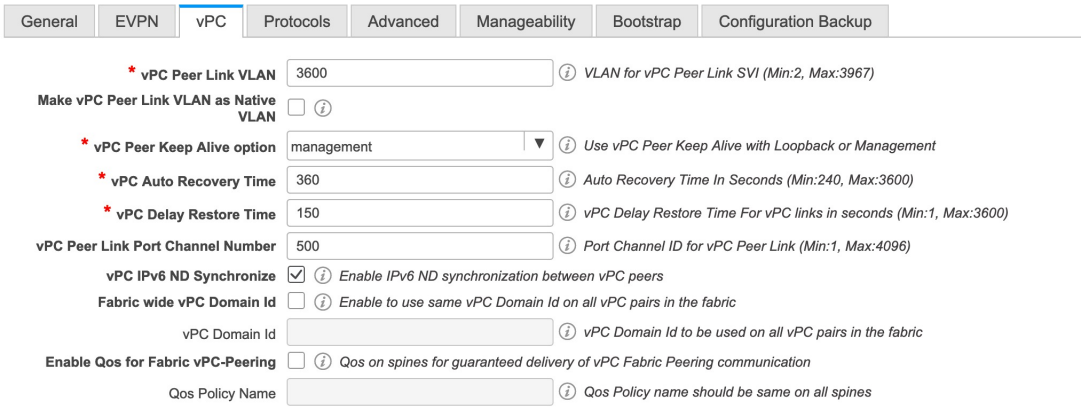

vPC Peer Link VLAN: VLAN used for the vPC peer link SVI.

Make vPC Peer Link VLAN as Native VLAN - Enables vPC peer link VLAN as Native VLAN.

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.

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

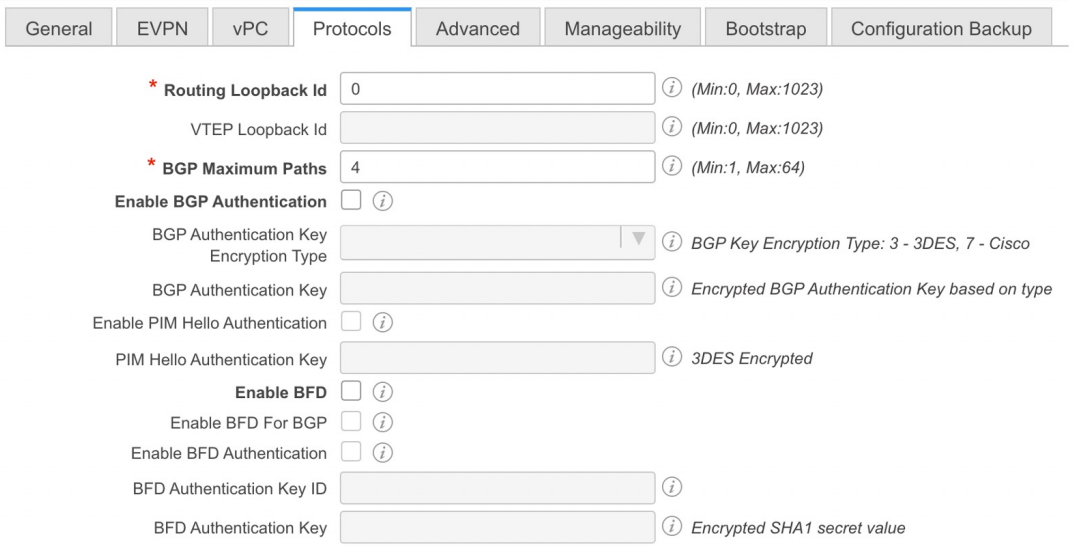

Routing Loopback Id - The loopback interface ID is populated as 0 by default. It is used as the BGP router ID.

BGP Maximum Paths - Specifies the BGP maximum paths.

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 Authentication Key field. Refer the Retrieving the Authentication Key section for details.

Enable BFD: Select the check box to enable **feature bfd** on all switches in the fabric. This feature is valid only on IPv4 underlay and the scope is within a fabric.

From Cisco DCNM Release 11.3(1), BFD within a fabric is supported natively. The BFD feature is disabled by default in the Fabric Settings. If enabled, BFD is enabled for the underlay protocols with the default settings. Any custom required BFD configurations must be deployed via the per switch freeform or per interface freeform policies.

The following config is pushed after you select the **Enable BFD** check box:

```
feature bfd
```

Note After you upgrade from DCNM Release 11.2(1) with BFD enabled to DCNM Release 11.3(1), the follo configs are pushed on all P2P fabric interfaces:

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

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 BGP: Select the check box to enable BFD for the BGP neighbor. This option is disabled by default.

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.

BFD Authentication Key ID: Specifies the BFD authentication key ID for the interface authentication.

BFD Authentication Key: Specifies the BFD authentication key.

For information about how to retrieve the BFD authentication parameters, see *Retrieving the Encrypted BFD Authentication Key, in Cisco DCNM LAN Fabric Configuration Guide*.

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

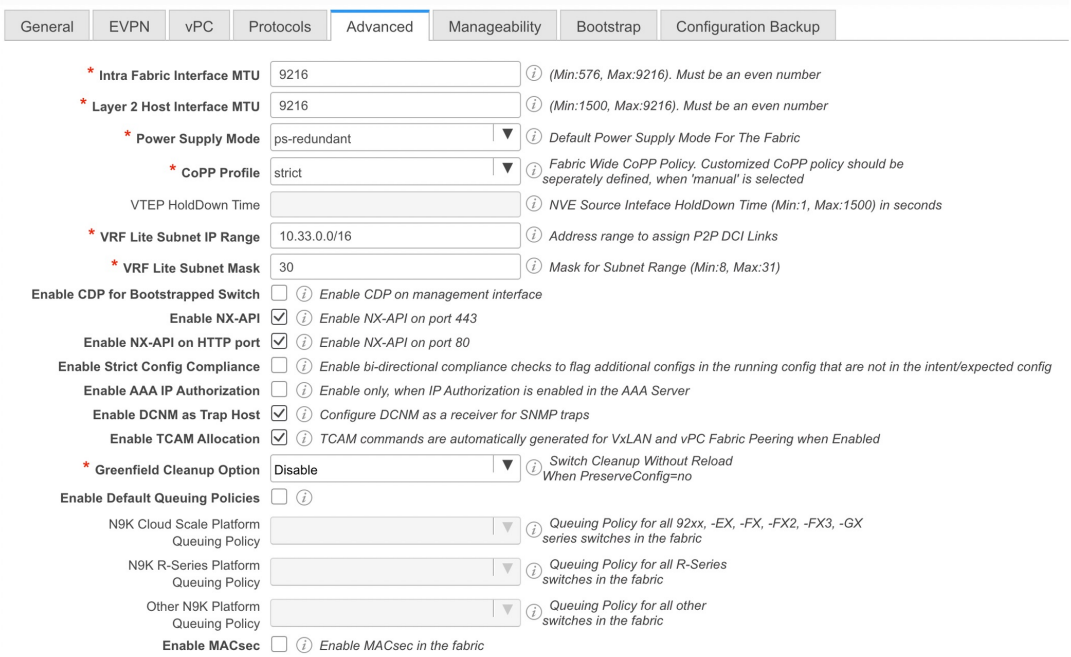

Intra Fabric Interface MTU - Specifies the MTU for the intra fabric interface. This value should be an even number.

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

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

Enable CDP for Bootstrapped Switch - Select the check box to enable CDP for bootstrapped switch.

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

Enable NX-API on HTTP - Specifies enabling of NX-API on HTTP. Enable this check box and the **Enable NX-API** check box to use HTTP. This check box is checked by default. If you uncheck this check box, the applications that use NX-API and supported by Cisco DCNM, such as Endpoint Locator (EPL), Layer 4-Layer 7 services (L4-L7 services), VXLAN OAM, and so on, start using the HTTPS instead of HTTP.

Note If you check the **Enable NX-API** check box and the **Enable NX-API on HTTP** check box, applications u HTTP.

Enable Strict Config Compliance - Enable the Strict Config Compliance feature by selecting this check box.

For Strict Configuration Compliance, see *Enhanced Monitoring and Monitoring Fabrics Guide*.

Note If Strict Config Compliance is enabled in a fabric, you cannot deploy Network Insights for Resources Cisco DCNM.

Enable AAA IP Authorization - Enables AAA IP authorization, when IP Authorization is enabled in the AAA Server.

Enable DCNM as Trap Host - Select this check box to enable DCNM as a trap host.

Enable TCAM Allocation: TCAM commands are automatically generated for VXLAN and vPC Fabric Peering when enabled.

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 Default Queuing Policies: Check this check box to apply QoS policies on all the switches in this fabric. To remove the QoS policies that you applied on all the switches, uncheck this check box, update all the configurations to remove the references to the policies, and save and deploy. From Cisco DCNM Release 11.3(1), pre-defined QoS configurations are included that can be used for various Cisco Nexus 9000 Series Switches. When you check this check box, the appropriate QoS configurations are pushed to the switches in the fabric. The system queuing is updated when configurations are deployed to the switches. You can perform the interface marking with defined queuing policies, if required, by adding the required configuration to the per interface freeform block.

Review the actual queuing policies by opening the policy file in the template editor. From Cisco DCNM Web UI, choose **Control > Template Library**. Search for the queuing policies by the policy file name, for example, **queuing_policy_default_8q_cloudscale**. Choose the file and click the **Modify/View template** icon to edit the policy.

See the *Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide* for platform specific details.

N9K Cloud Scale Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all Cisco Nexus 9200 Series Switches and the Cisco Nexus 9000 Series Switches that ends with EX, FX, and FX2 in the fabric. The valid values are **queuing_policy_default_4q_cloudscale** and **queuing_policy_default_8q_cloudscale**. Use the **queuing_policy_default_4q_cloudscale** policy for FEXes. You can change from the **queuing_policy_default_4q_cloudscale** policy to the **queuing_policy_default_8q_cloudscale** policy only when FEXes are offline.

N9K R-Series Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all Cisco Nexus switches that ends with R in the fabric. The valid value is **queuing_policy_default_r_series**.

Other N9K Platform Queuing Policy: Choose the queuing policy from the drop-down list to be applied to all other switches in the fabric other than the switches mentioned in the above two options. The valid value is **queuing_policy_default_other**.

Enable MACsec: Enables MACsec for the fabric. For more information, see MACsec Support in Easy Fabric and eBGP Fabric, on page 194.

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.

Intra-fabric Links Additional Config - Add CLIs that should be added to the intra-fabric links.

8. Click the **Manageability** tab.

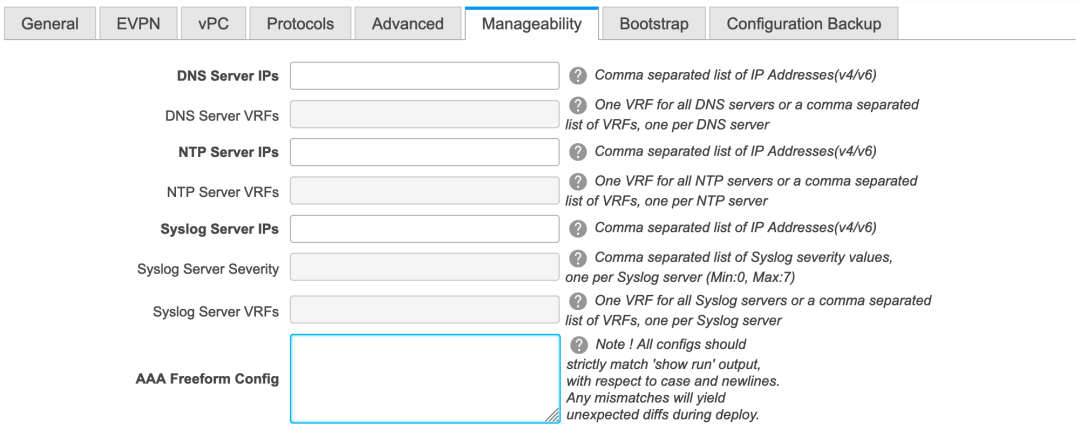

The fields in this tab are:

DNS Server IPs - Specifies the comma separated list of IP addresses (v4/v6) of the DNS servers.

DNS Server VRFs - Specifies one VRF for all DNS servers or a comma separated list of VRFs, one per DNS server.

NTP Server IPs - Specifies comma separated list of IP addresses (v4/v6) of the NTP server.

NTP Server VRFs - Specifies one VRF for all NTP servers or a comma separated list of VRFs, one per NTP server.

Syslog Server IPs – Specifies the comma separated list of IP addresses (v4/v6) IP address of the syslog servers, if used.

Syslog Server Severity – Specifies the comma separated list of syslog severity values, one per syslog server. The minimum value is 0 and the maximum value is 7. To specify a higher severity, enter a higher number.

Syslog Server VRFs – Specifies one VRF for all syslog servers or a comma separated list of VRFs, one per syslog server.

AAA Freeform Config – Specifies the AAA freeform configs.

If AAA configs are specified in the fabric settings, **switch_freeform** PTI with source as **UNDERLAY_AAA** and description as "**AAA Configurations**" will be created.

9. Click the **Bootstrap** tab.

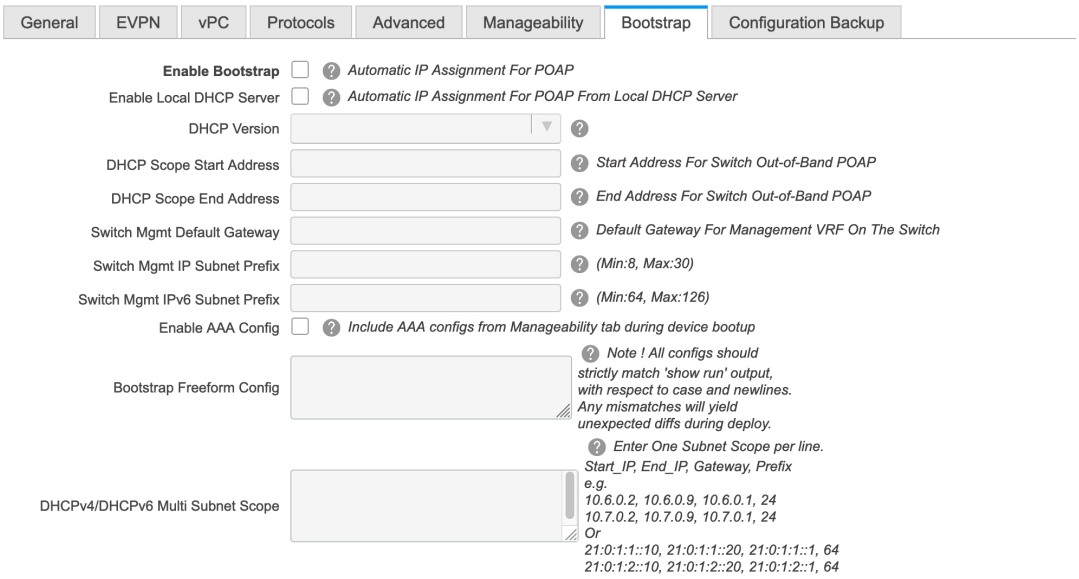

Enable Bootstrap - Select this check box to enable the bootstrap feature.

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

- External DHCP Server: Enter information about the external DHCP server in the **Switch Mgmt Default Gateway** and **Switch Mgmt IP Subnet Prefix** fields.
- Local DHCP Server: Enable the **Local DHCP Server** check box 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 Version – Select DHCPv4 or DHCPv6 from this drop-down list. When you select DHCPv4, the **Switch Mgmt IPv6 Subnet Prefix** field is disabled. If you select DHCPv6, the **Switch Mgmt IP Subnet Prefix** is disabled.

Note Cisco DCNM IPv6 POAP is not supported with Cisco Nexus 7000 Series Switches. Cisco Nexus 9000 3000 Series Switches support IPv6 POAP only when switches are either L2 adjacent (eth1 or out-of-ba subnet must be a /64) or they are L3 adjacent residing in some IPv6 /64 subnet. Subnet prefixes other t /64 are not supported.

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 Mgmt Default Gateway: Specifies the default gateway for the management VRF on the switch.

Switch Mgmt IP 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.

Switch Mgmt IPv6 Subnet Prefix - Specifies the IPv6 prefix for the Mgmt0 interface on the switch. The prefix should be between 112 and 126. This field is editable if you enable IPv6 for DHCP.

Enable AAA Config – Select this check box to include AAA configs from the Manageability tab during device bootup.

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*.

DHCPv4/DHCPv6 Multi Subnet Scope - Specifies the field to enter one subnet scope per line. This field is editable after you check the **Enable Local DHCP Server** check box.

The format of the scope should be defined as:

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

For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24

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

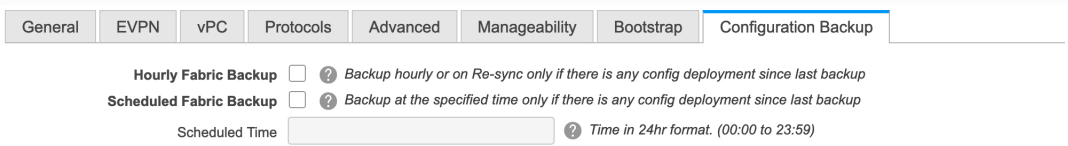

Hourly Fabric Backup: Select the check box to enable an hourly backup of fabric configurations and the intent.

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.

The backup process is initiated after you click **Save**.

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

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

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

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

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

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

Salient Points

- Deploy the leaf underlay policies on all leaf switches at once, since they have a common AS number.
- Brownfield migration is not supported for eBGP fabrics.
- 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.
- The supported roles are leaf, spine, and border leaf.
- On the border device, VRF-Lite is supported with manual mode.
- You must apply policies on the leaf and spine switches for a functional 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 and interfaces. For more information, see Pre-provisioning a Device , on page 86 and Pre-provisioning an Ethernet Interface, on page 90.

Note When DCNM discovers a switch with the hostname containing the period character (.), it is treated as a domain-name and truncated. Only the text prior to the period character (.) is considered as a hostname. For example:

- If hostname is **leaf.it.vxlan.bgp.org1-XYZ**, DCNM shows only **leaf**
- If hostname is **leaf-itvxlan.bgp.org1-XYZ**, DCNM shows only **leafit-vxlan**

Discovering Existing Switches

1. After clicking on **Add Switches**, use the **Discover Existing Switches** tab to add one or more existing switches into the fabric. In this case, a switch with known credentials and a pre-provisioned IP address, is added to the fabric. The IP address (Seed IP), administrator username, and password (**Username** and **Password** fields) of the switch are provided as the input by a user. The **Preserve Config** knob is set to **yes** by default. This is the option that a user would select for a brownfield import of a device into the fabric. For a greenfield import where the device configuration will be cleaned up as part of the import process, the user should set the **Preserve Config** knob to **no**.

Note Easy_Fabric_eBGP does not support brownfield import of a device into the fabric.

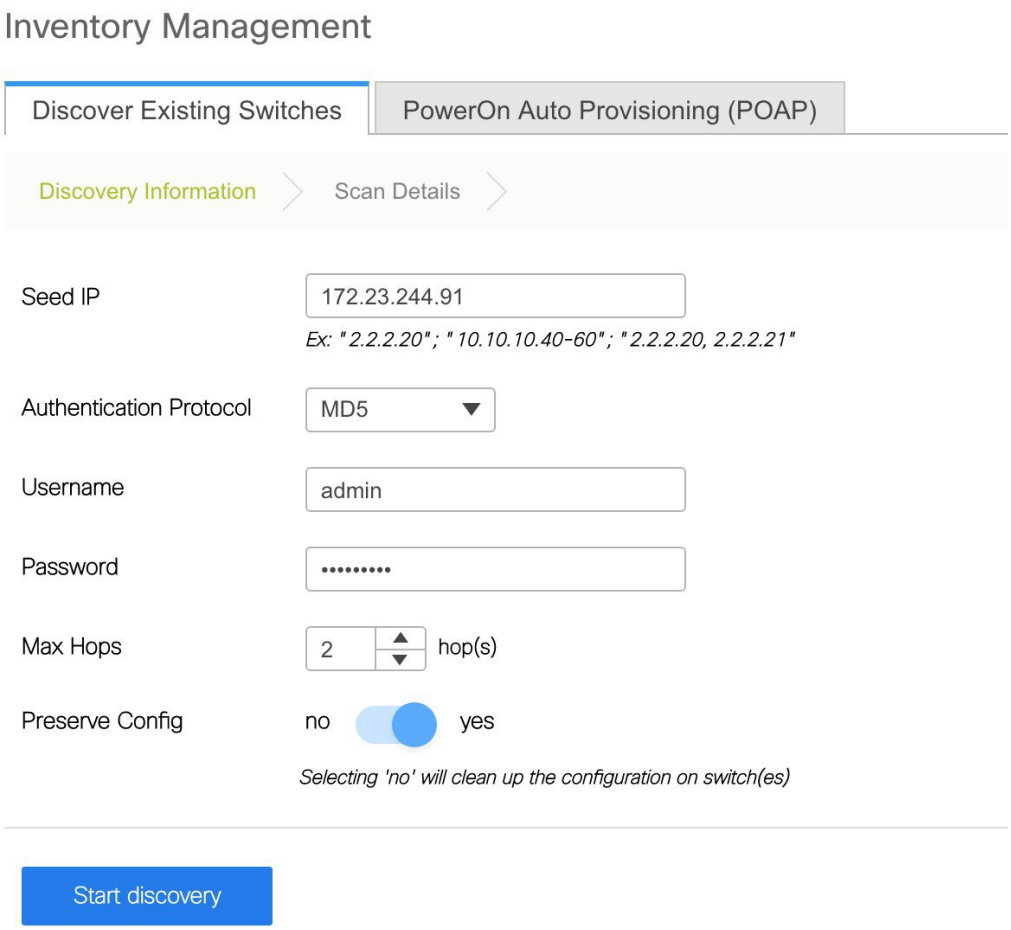

2. Click **Start discovery**. The **Scan Details** window comes up shortly. Since the **Max Hops** field was populated with 2 (by default), the switch with the specified IP address (leaf-91) and switches two hops from that switch, are populated in the **Scan Details**result.

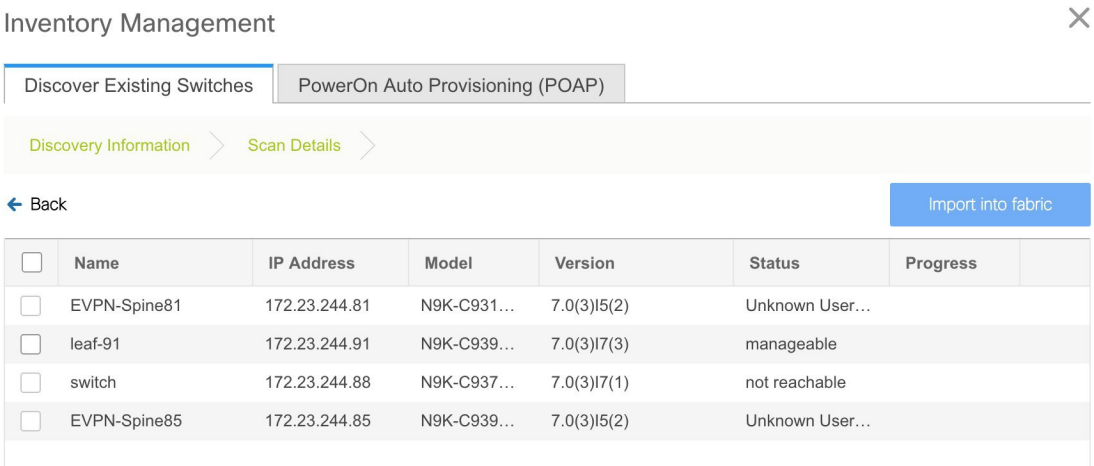

3. If the DCNM was able to perform a successful shallow discovery to a switch, the status will show up as **Manageable**. Select the check box next to the appropriate switch(es) and click **Import into fabric**.

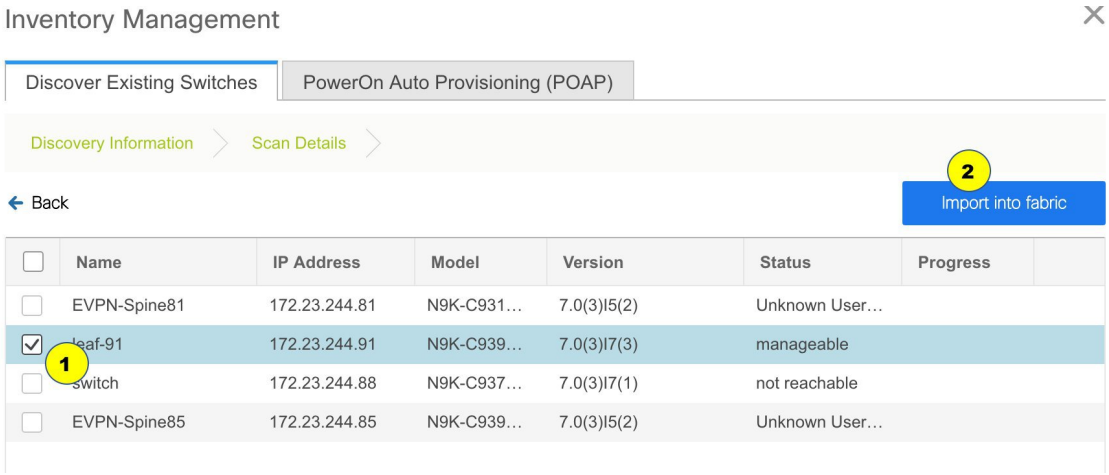

Though this example describes the discovery of one switch, multiple switches can be discovered at once.

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 a 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.

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.

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 devices, assign an appropriate role to each device. For this purpose, tight click the device, and use the **Set role** option to set the appropriate role. Alternatively, the tabular view may be employed to assign the same role to multiple devices at one go.

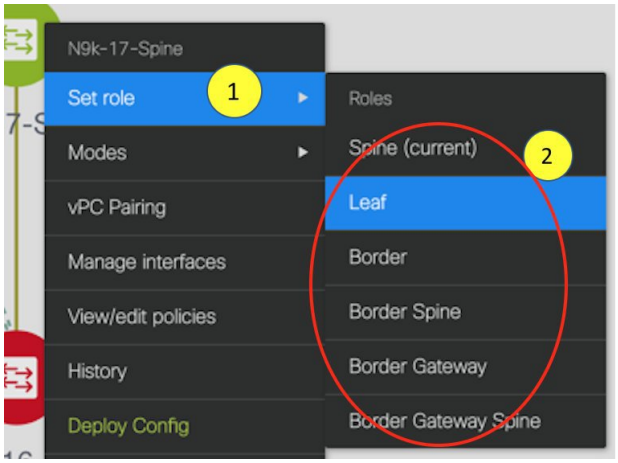

If you choose the Hierarchical layout for display (in the Actions panel), the topology automatically gets aligned as per role assignment, with the leaf devices at the bottom, the spine devices connected on top of them, and the border devices at the top.

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

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

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

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

The template and interface configurations form the underlay network configuration on the switches. Also, freeform CLIs that were entered as part of fabric settings (leaf and spine switch freeform configurations

entered in the Advanced tab) are deployed. For more details on freeform configurations, refer Enabling Freeform Configurations on Fabric Switches .

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

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

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

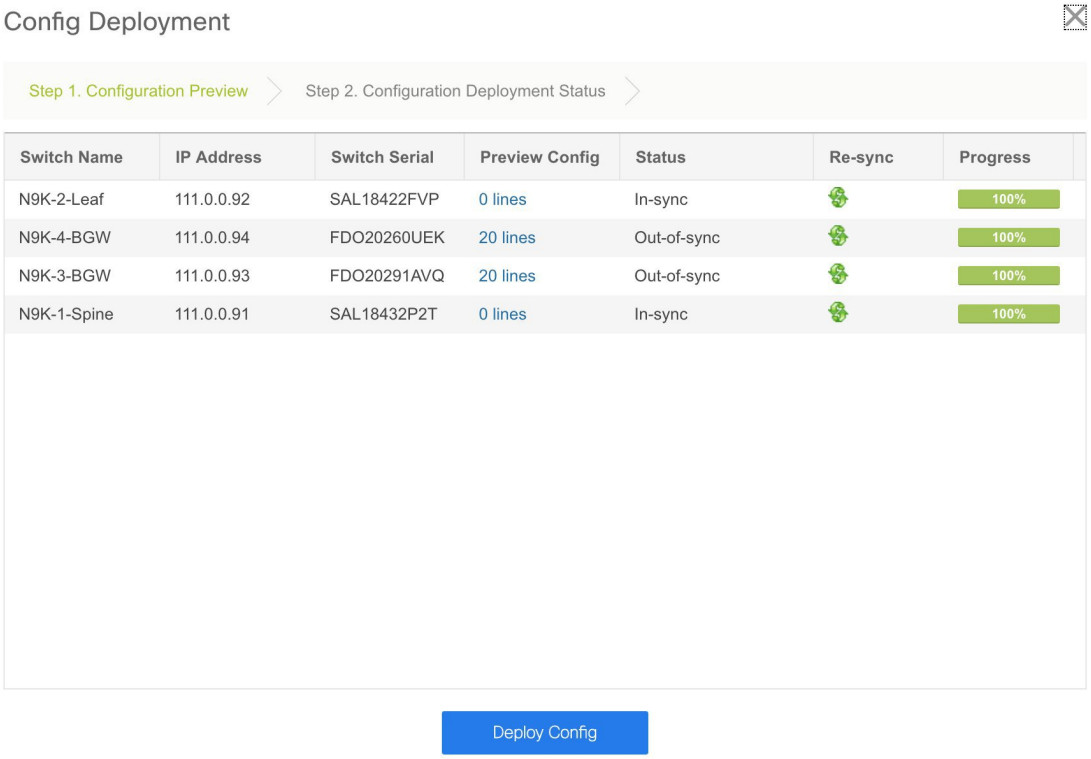

If the status is out-of-sync, it suggests that there is inconsistency between the DCNM and configuration on the device.

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

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

The **PendingConfig** tab displays the pending configurations for successful deployment.

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

In DCNM 11, 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 are 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. The pending state indicates that there is a pending deployment or pending recomputation. You can click on the switch and review the pending deployments using **Preview** or **Deploy Config** options, or click **Save & Deploy** to recompute the state of the switch.

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

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.

From Cisco NX-OS Release 11.4(1), if you uncheck the **FEX** check box in the **Topology** window, FEX devices are hidden in the **Fabric Builder** topology window as well. To view FEX in **Fabric Builder**, you need to check this check box. This option is applicable for all fabrics and it is saved per session or until you log out of DCNM. If you log out and log in to DCNM, the FEX option is reset to default, that is, enabled by default. For more information, see Show Panel, on page 24.

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

Discovering New Switches

1. When a new Cisco NX-OS device is powered on, typically that device has no startup configuration or any configuration state for that matter. Consequently, it powers on with NX-OS and post initialization, goes into a POAP loop. The device starts sending out DHCP requests on all the interfaces that are up including the mgmt0 interface.
2. As long as there is IP reachability between the device and the DCNM, the DHCP request from the device, will be forwarded to the DCNM. For easy day-0 device bring-up, the bootstrap options should be enabled in the **Fabric Settings** as mentioned earlier.

3. With bootstrap enabled for the fabric, the DHCP request coming from the device will be serviced by the DCNM. The temporary IP address allocated to the device by the DCNM will be employed to learn basic information about the switch including the device model, device NX-OS version, etc.

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

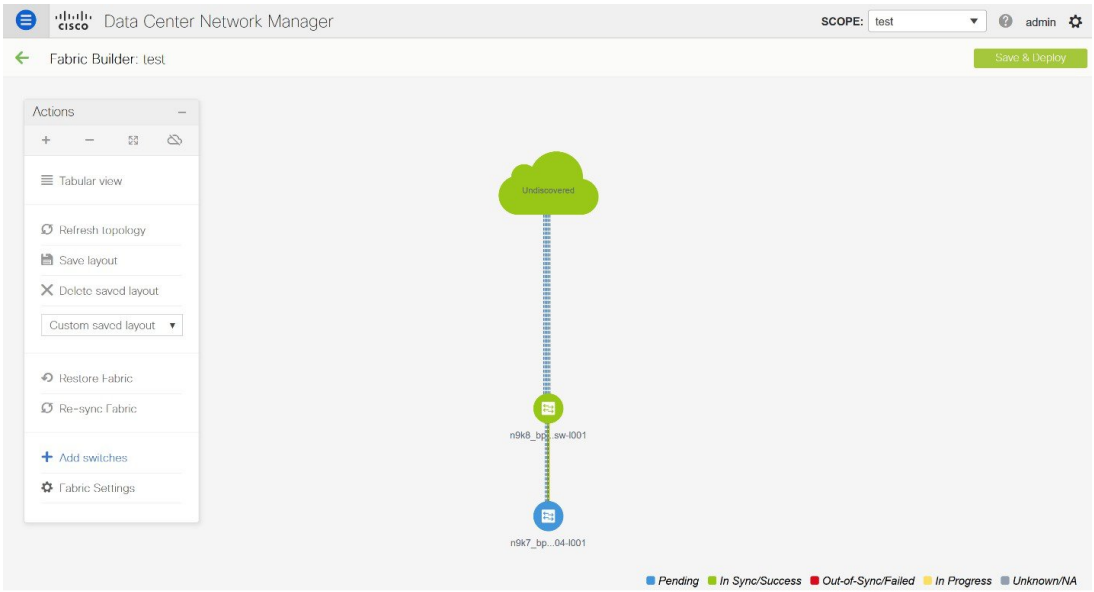

Go to the fabric topology window and click the **Add switches** option from the **Actions** panel. The Inventory Management window comes up.

5. Click the **POAP** tab.

As mentioned earlier, DCNM retrieves the serial number, model number, and version from the device 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

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

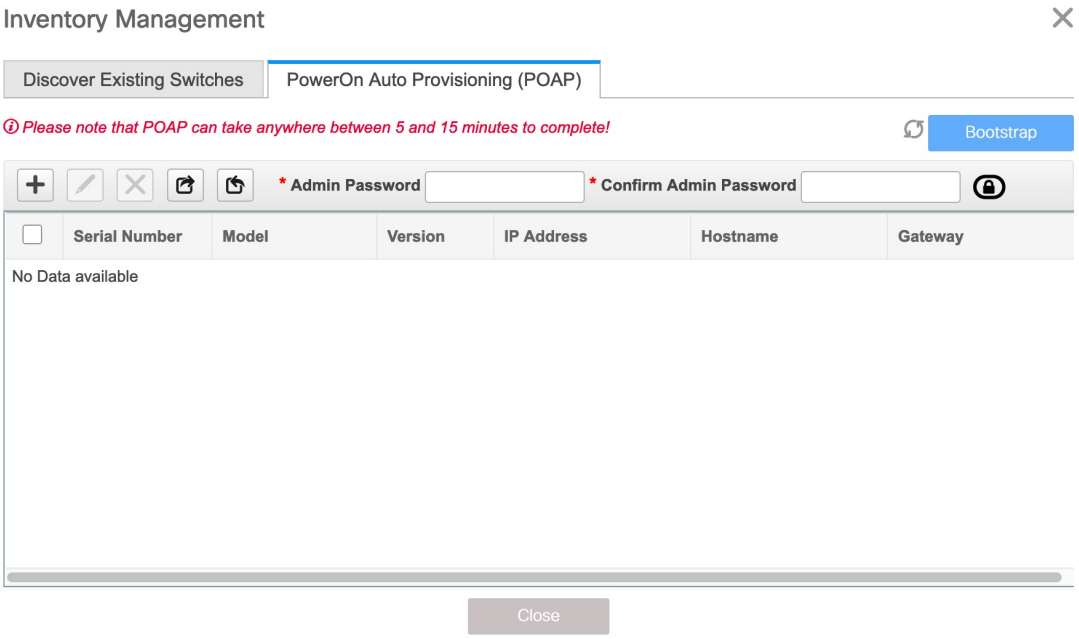

Select the checkbox next to the switch and enter the switch credentials: IP address and host name.

Based on the IP address of your device, you can either add the IPv4 or IPv6 address in the **IP Address** field.

Beginning with Release 11.2(1), you can provision devices in advance. To pre-provision devices, refer to Pre-provisioning a Device , on page 86.

6. 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 authentic that is, RADIUS or TACACS credentials for discovery only.

7. (Optional) Use discovery credentials for discovering switches.

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

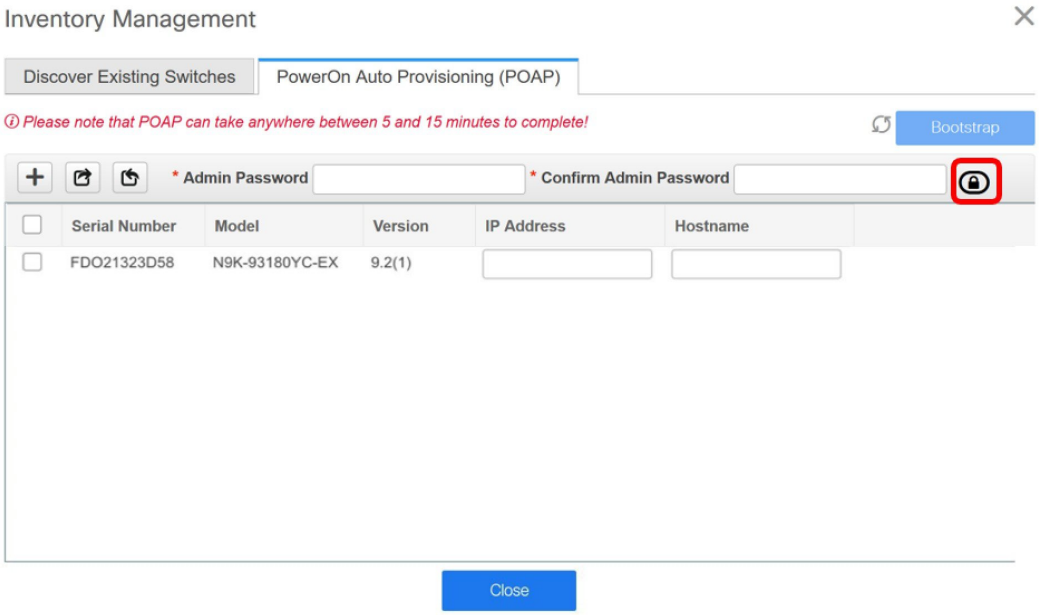

b. In the **Discovery Credentials** window, enter the discovery credentials such as discovery username and password.

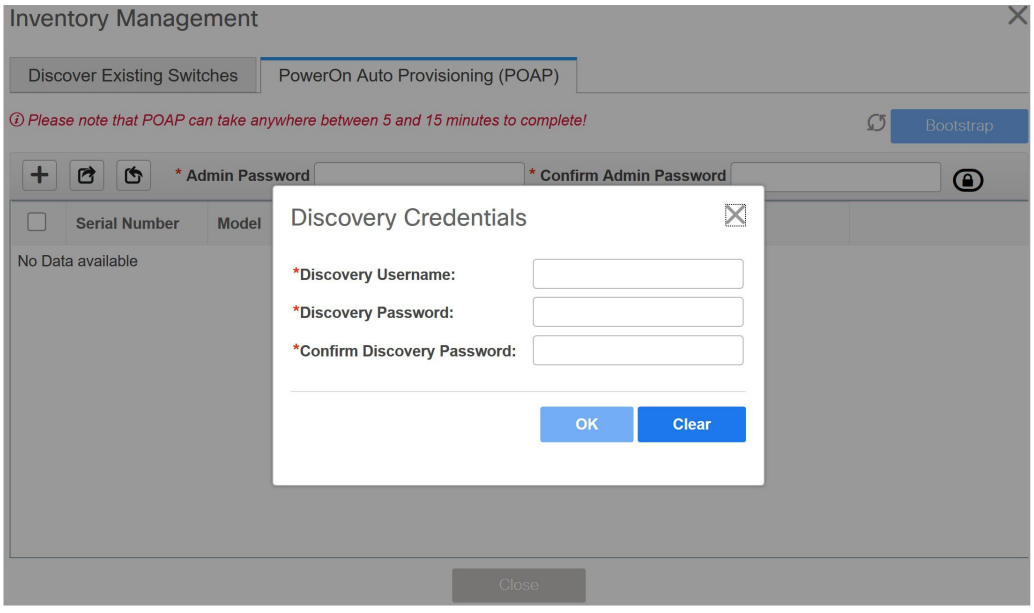

Click **OK** to save the discovery credentials.

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

8. 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.

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

10. After the added switch completes POAP, the fabric builder topology page is refreshed with the added switch thereby depicting its discovered physical connections. Set the appropriate role for the switch followed by a Save & Deploy operation at the fabric level. The Fabric Settings, switch role, the topology etc. are evaluated by the Fabric Builder and the appropriate intended configuration for the switch is generated as part of the Save operation. The pending configuration will provide a list of the configurations that need to be deployed to the new switch in order to bring it IN-SYNC with the intent.

Note For any changes on the fabric that results in the Out-of-Sync, then you must deploy the changes. The pr 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 n update the AAA server password on each switch. Else, switch discovery fails.

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

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

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

14. 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.

15. 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. 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 pairing/un-pairing 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:

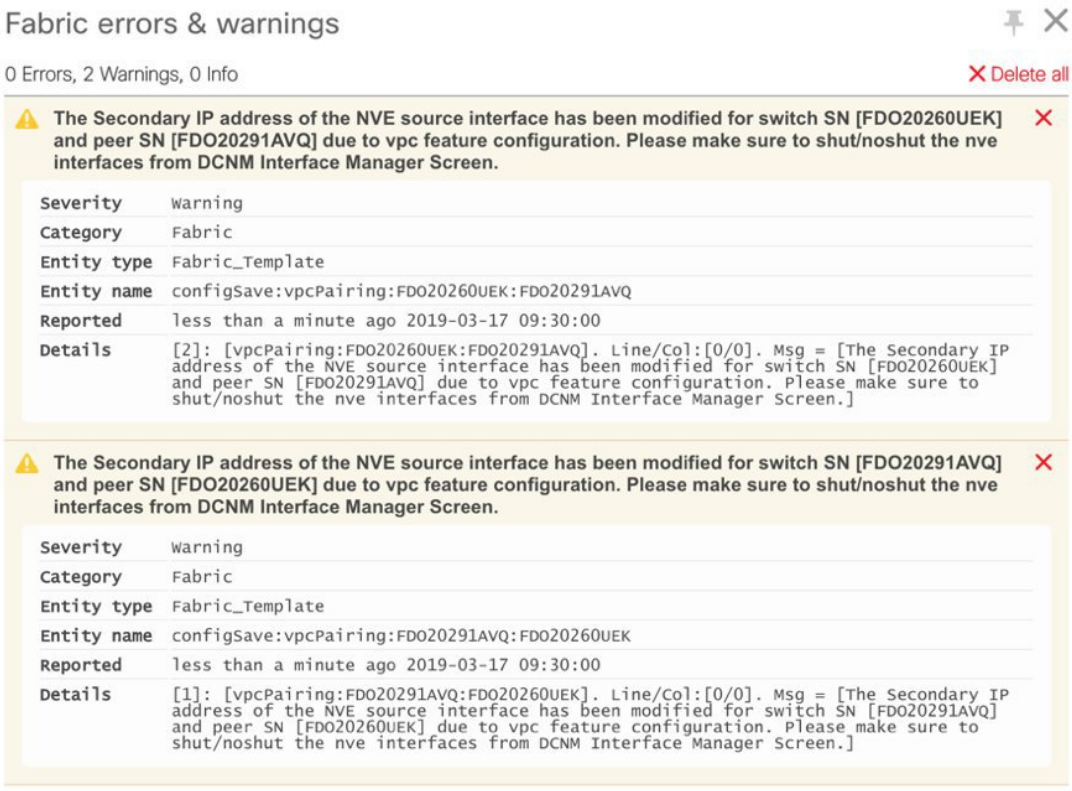

To resolve, go to the Control > Interfaces screen and deploy the Shutdown operation on the nve interface followed by a No Shutdown configuration. This is depicted in the figure below where the up arrow corresponds to a No Shutdown operation while a down arrow corresponds to a Shutdown operation.

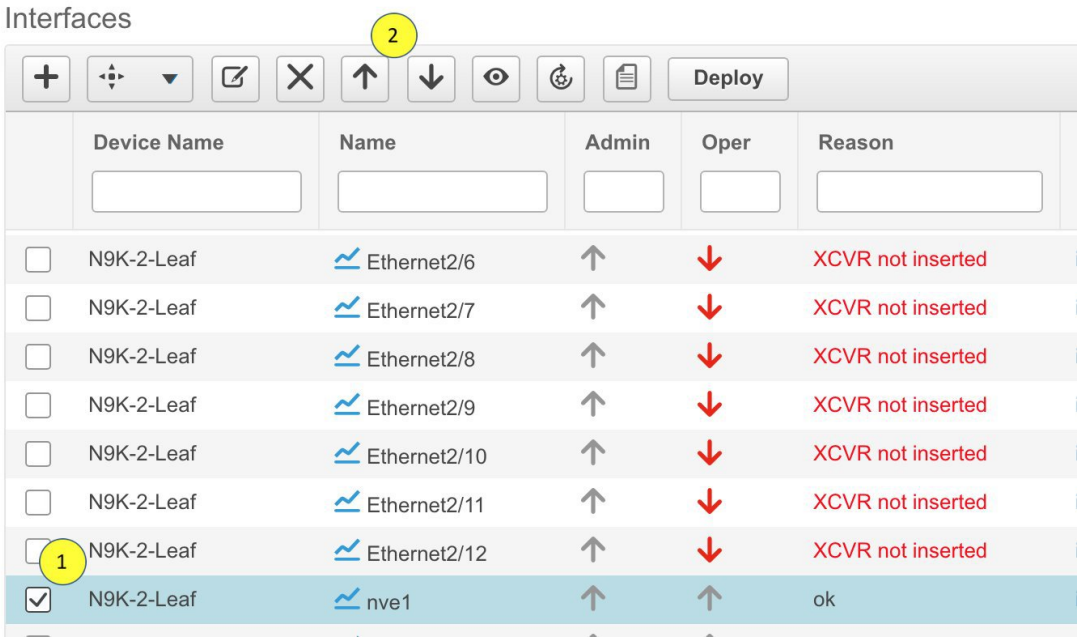

You can right-click the switch to view various options:

- **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 205.

- **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 and policy change history.

  The **Policy Change History** tab lists the history of policies along with the users who made the changes like add, update, or delete.

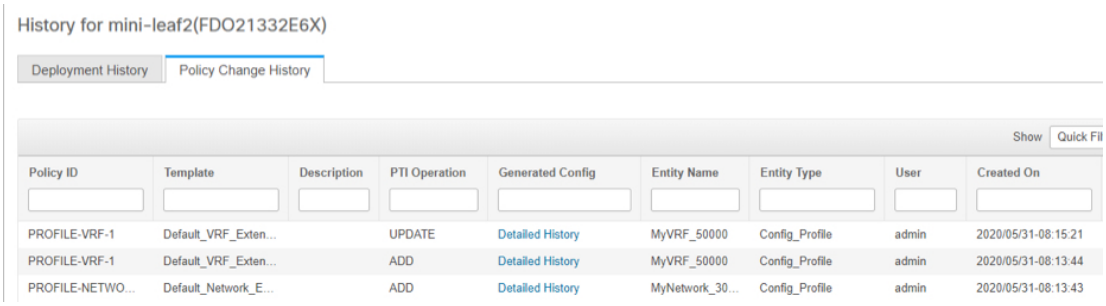

Under the **Policy Change History** tab, for a policy, click **Detailed History** under the **Generated Config** column to view the generated config before and after.

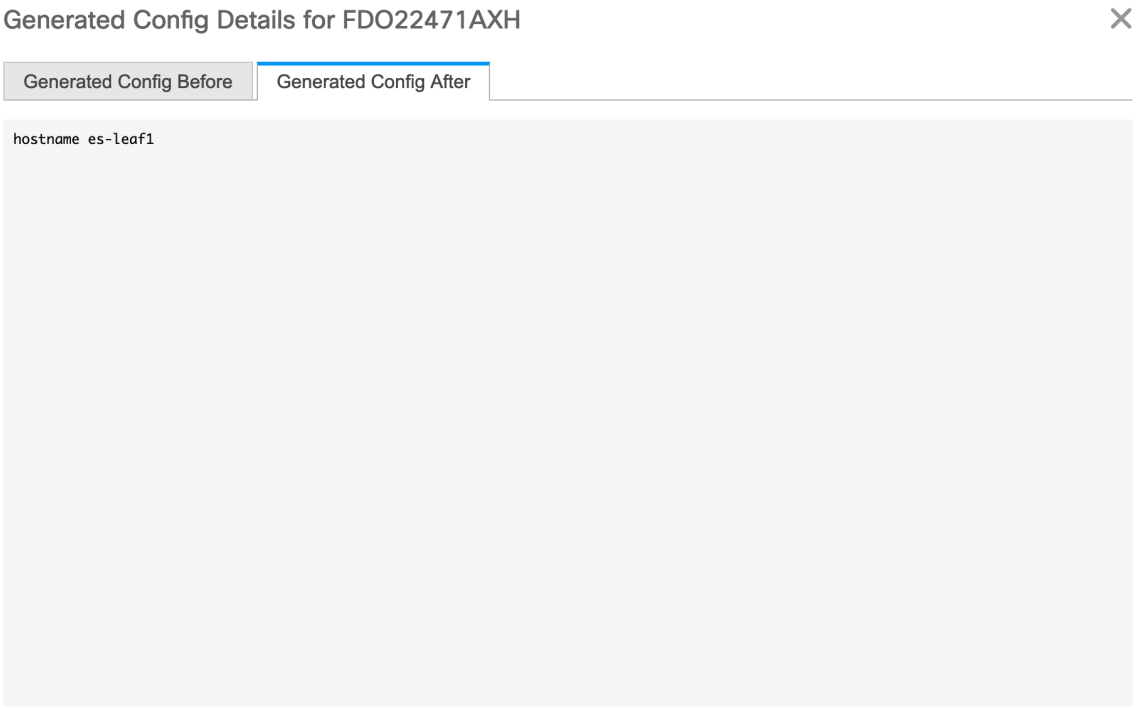

The following table provides the summary of generated config before and after for Policy Template Instances (PTIs).

PTI Operations	Generated Config Before	Generated Config After
Add	Empty	Contains the config
Update	Contains config before changes	Contains config after changes
Mark-Delete	Contains the config to be removed.	Contains the config to be removed with colour change.
Delete	Contains the config	Empty

Note When a policy or profile template is applied, an instance is created for each application of the template, which is known as Policy Template Instance or PTI.

- **Preview Config** - View the pending configuration and the side-by-side comparison of the running and expected configuration.
- **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 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 networks and deploy them on the switches. [Refer Creating and Deploying Networks and VRFs].

Deploying Fabric Underlay eBGP Policies

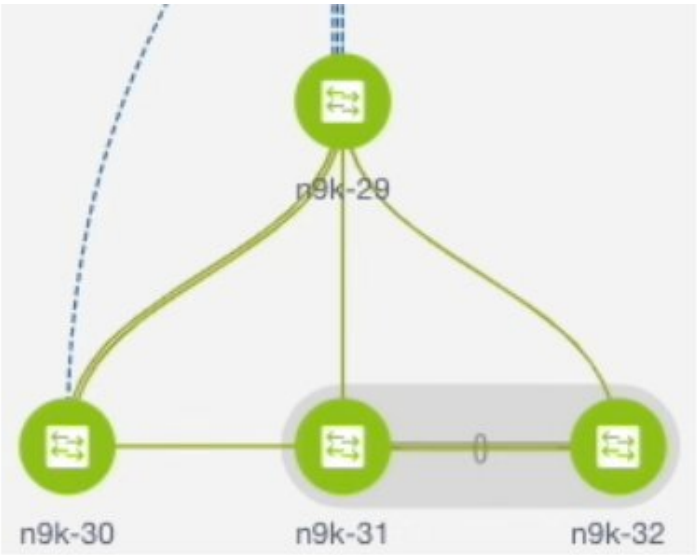

The topology shows a Routed fabric enabled with eBGP as the routing protocol for distributing reachability information. In DCNM, a fabric with the **Easy_Fabric_eBGP** template is created. One spine switch (n9k-29) and three leaf switches (n9k-30, and vPC switch pair n9k-31 and n9k-32) are imported to 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 mod 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), a 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 Networks in eBGP-based Fabrics

Overview of Networks in a Routed Fabric

From Cisco DCNM Release 11.3(1), you can create a top-down network configuration for a routed fabric using DCNM. A routed fabric is run in one VRF, which is the default VRF. Note that creating VRFs manually is disabled for a routed fabric. Since the fabric is an IPv4 fabric, IPv6 address within the network is not supported. In a routed fabric, a network can only be attached to one device or a pair of vPC devices, unless it is a Layer 2 only network.

Note A routed fabric network configuration will not be put under a config-profile.

When the eBGP fabric is configured as Routed Fabric (EVPN is disabled), at the fabric level, you can select the first hop redundancy protocol (FHRP) for host traffic to be either HSRP or VRRP. HSRP is the default value.

For a vPC pair, DCNM generates network level HSRP or VRRP configuration based on the fabric setting. If HSRP is chosen, each network is configured with one HSRP group, and the HSRP VIP address. By default, all the networks will share the same HSRP group number allocated by DCNM, while you can overwrite it per network. VRRP support is similar to HSRP.

Guidelines

- HSRP authentication or VRRP authentication is not supported. If you want to use authentication, you can enter the applicable commands in the network freeform config.
- vPC peer gateway can be used to minimize peer link usage in the case that some third-party devices ignore the HSRP virtual-MAC and use the ARP packet source MAC for ARP learning. In Routed fabric mode, DCNM generates vPC peer gateway command for VPC devices.
- For an eBGP fabric, changing between routed fabric type and EVPN fabric type, or HSRP and VRRP, is not allowed with the presence of networks and VRFs. You need to undeploy and delete these networks and VRFs before changing the fabric type or FHRP. For more information, see *Undeploying Networks for the Standalone Fabric* and *Undeploying VRFs for the Standalone Fabric*.
- After the upgrade from DCNM Release 11.2(1) to 11.3(1), if the fabric was running in Routed Fabric mode previously, the default fabric values such as FHRP protocol and network VLAN range are internally set for a Routed Fabric. You need to edit the fabric settings if you want to configure different values. Before deploying a network configuration, you need to update the FHRP protocol fabric setting and click **Save & Deploy**.

Creating and Deploying a Network in a Routed Fabric

This procedure shows how to create and deploy a network in a routed fabric.

Before you begin

Create a routed fabric and deploy the necessary leaf and spine policies.

Procedure

Step 1 Navigate to **Control > Networks**.

Step 2 From the **SCOPE** drop-down list, choose a routed fabric.

Step 3 Click the **Add** button in the **Networks** window to create a network.

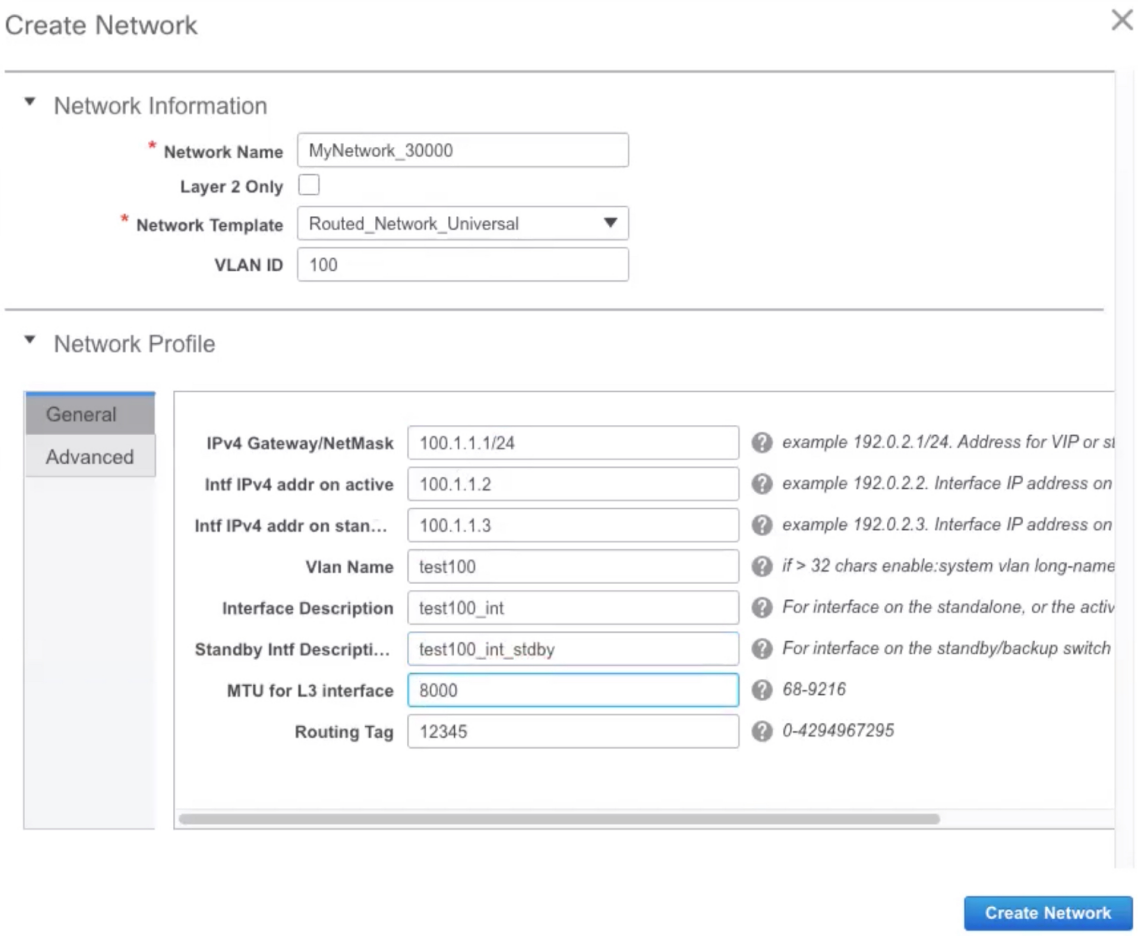

Network Name: Specifies the name of the network. The network name should not contain any white spaces or special characters except underscore (_) and hyphen (-).

Layer 2 Only: Optional. Specifies whether the network is a Layer 2 only network. FHRP configuration is not generated in a Layer 2 only network.

Note When an L3 Network template is attached to a standalone device, no FHRP configuration is generated.

Network Template: Select the **Routed_Network_Universal** template.

VLAN ID: Optional. 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 gateway address with subnet.

Intf IPv4 addr on active: Specifies the IPv4 interface address on an active device in a vPC pair. This field is applicable only when you are creating and deploying a network for a vPC pair of devices.

Intf IPv4 addr on standby: Specifies the IPv4 interface address on a standby/backup device in a vPC pair. This field is applicable only when you are creating and deploying a network for a vPC pair of devices.

Note The IPv4 gateway address and interface addresses should be in the same subnet.

The following fields under the **General** tab are optional:

Vlan Name: Specifies the VLAN name.

Interface Description: Specifies the description for the interface.

Standby Intf Description: Specifies the description for the standby interface in a vPC pair.

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

Routing Tag: Specifies the routing tag that is associated with each gateway IP address prefix.

Advanced tab: This tab is applicable only when you are creating and deploying a network for a vPC pair of devices.

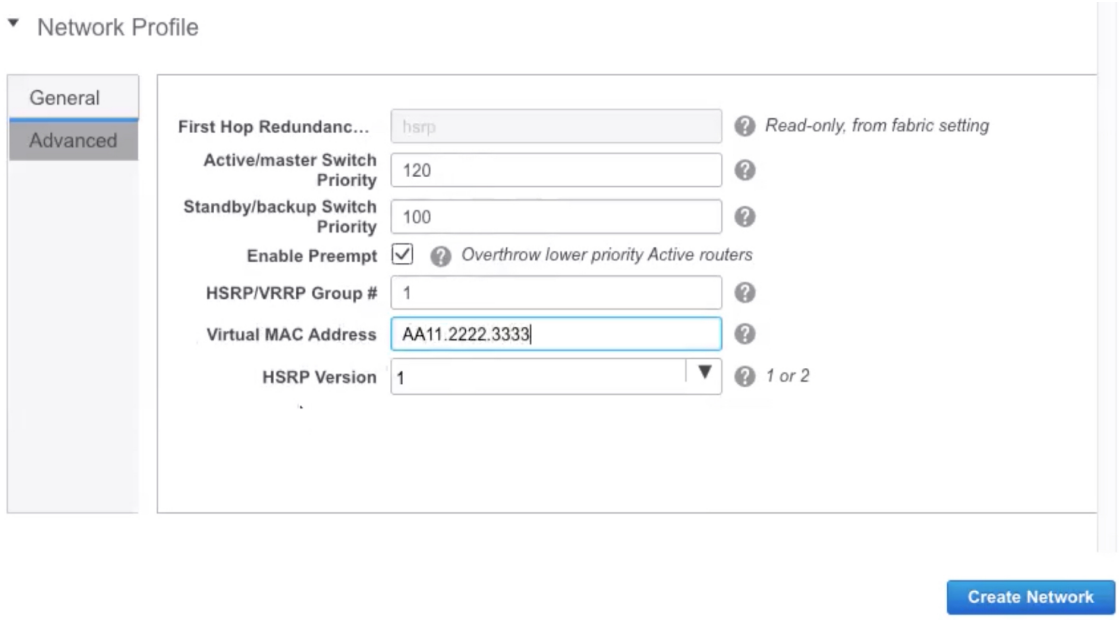

First Hop Redundancy Protocol: A read-only field that specifies FHRP selected in the fabric settings.

Active/master Switch Priority:Specifies the priority of the active or master device.

Standby/backup Switch Priority: Specifies the priority of the standby or backup device. The default value is 100. Note that this default value is not displayed when you preview the network configuration before deployment.

Enable Preempt: Specifies whether the standby/backup device can preempt an active device.

HSRP/VRRP Group #: Specifies the HSRP or VRRP group number. By default, HSRP group number is 1.

Virtual MAC Address: Optional. Specifies the virtual MAC address. By default, VMAC is internally generated based on the HSRP group number (0000.0c9f.f000 + group number). The virtual MAC address is only applicable when **hsrp** is selected in the fabric settings.

HSRP Version: Specifies the HSRP version. The default value is 1. The **HSRP version** field is only applicable for HSRP.

Step 4 Click **Create Network**.

Step 5 In the **Networks** window, select the check box next to a network and click **Continue**.

Note A non Layer 2 network can be only applied to a vPC pair of devices or a single device. For example, if you have deployed a network on a single device, you cannot deploy the same network on another device or a vPC pair of devices.

Step 6 Select a device or a vPC pair to deploy a network.

Note In a routed fabric, when you try to attach a network on a vPC pair without active or standby IP addresses, an error is displayed saying that the IP address fields are not filled. After you add the IP addresses and save the network, the network state changes to **PENDING** without the need to attach the network again.

Step 7 In the **Network Attachment** window, for a vPC pair, assign the active state for a device.

Enter **true** under the **isActive** column for an active device and **false** for a standby device.

Click **Save**.

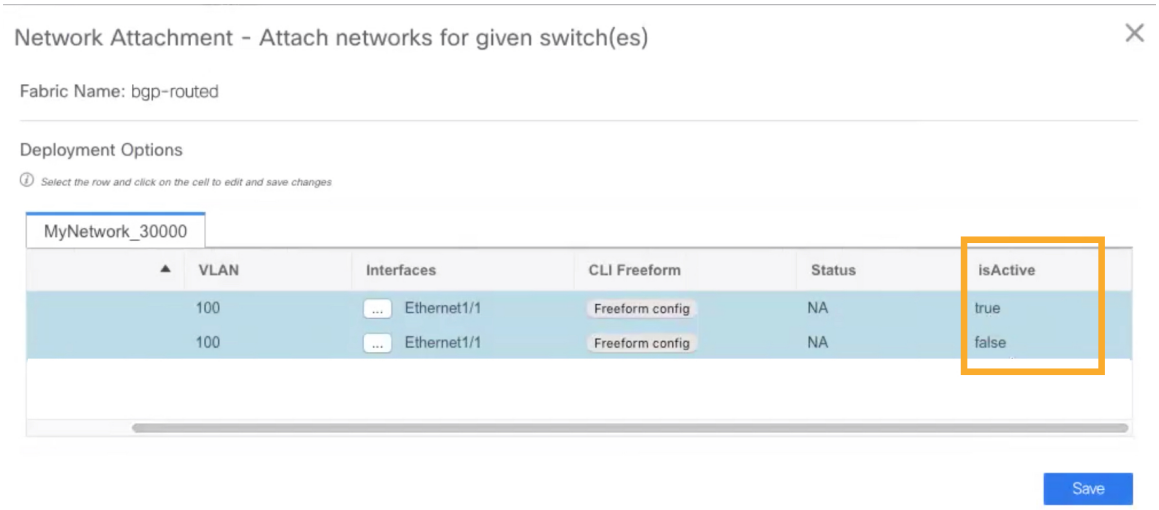

Note In a routed fabric, when you edit a deployed network and save without making any changes, the status of the network changes to **Pending**. Similarly, if a **Network Attachment** window is opened for a deployed network, and saved without any changes, the status of the network changes to **Pending**. In these cases, click the **Preview** icon to preview the config. This action changes the network status back to **Deployed**.

Step 8 (Optional) Click the **Preview** icon to preview the configs that will be deployed on devices.

The **Preview Configuration** window is displayed.

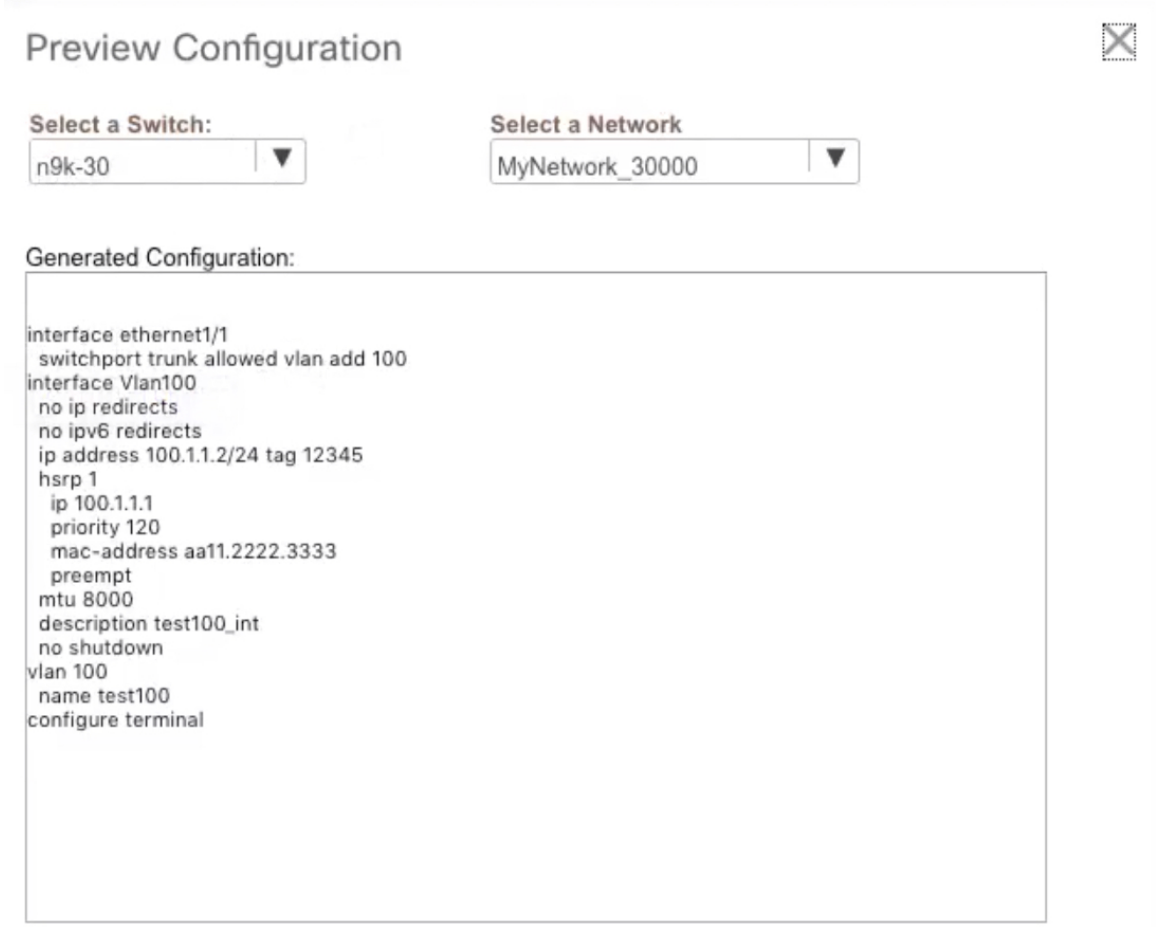

Step 9 Click the **Deploy** button in the **Network / VRF Deployment** window.

You can also deploy the network by navigating to the **Fabric Builder** window and clicking the **Deploy** button.

Creating Inter-Fabric Links Between a Routed Fabric and an External Fabric

From DCNM Release 11.3(1), you can use an inter-fabric link to connect a route fabric to an edge router. This link configures an IP address on the physical interface and establish eBGP peering with the edge router on default vrf. The BGP configuration includes advertising default route to leaf switches.

Note The **Fabric Monitor Mode** check box in the external fabric settings can be unchecked. Unchecking the **Fabric Monitor Mode** check box enables DCNM to deploy configurations to the external fabric. For more information, see Creating an External Fabric.

Procedure

Step 1 Navigate to **Control** > **Fabric Builder**.

Step 2 Click a routed a fabric in the **Fabric Builder** window.

Step 3 Click **Tabular view** in the **Actions** panel that is displayed at the left part of the window.

Step 4 Click the **Links** tab.

Step 5 Click the **Add** icon to add a link.

The **Link Management – Add Link** window is displayed.

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

Link Sub-Type – This field populates the IFC type. Choose **ROUTED_FABRIC** from the drop-down list.

Link Template: The link template is populated. The templates are autopopulated with corresponding pre-packaged default templates that are based on your selection. For a routed fabric, the **ext_routed_fabric** template is populated.

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 or port channel interface that connects to the destination device. Only device with the border role can be chosen.

Destination Device and **Destination Interface**—Choose the destination device and Ethernet or port channel 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.

BGP Local ASN: In this field, the AS number of the leaf is autopopulated if you have created and applied the **leaf_bgp_asn** policy.

IP Address/Mask: Fill up this field with the IP address of the source interface that connects to the destination device.

BGP Neighbor IP: Fill up this field with the IP address of the destination interface.

BGP Neighbor ASN: In this field, the AS number of the destination device is autopopulated.

BGP Maximum Paths: Specifies the maximum supported BGP paths.

The **Advanced** tab contains the following optional fields:

Source Interface Description and **Destination Interface Description** – Describe the links for later use. After **Save & Deploy**, this description 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 to *Enabling Freeform Configurations on Fabric Switches*.

Step 6 Click **Save** to finish adding a link.

Step 7 Click the **Back** icon to navigate back to the Fabric Builder window.

Step 8 Right-click the device which is connecting to the edge router in the external fabric, and select **Deploy Config**.

Step 9 In the **Config Deployment** window, click **Deploy Config**.

Step 10 Navigate to the external fabric in the **Fabric Builder** window, and click **Tabular view** in the **Actions** panel. Click the **Links** tab to see all the links for the external fabric.

You can see the inter-fabric link that has been created.

Note The inter-fabric link is created if the External fabric is not in the monitor mode.

Step 11 Click the **Back** icon twice to navigate back to the **Fabric Builder** window.

Step 12 Click the external fabric connecting to the routed fabric.

Step 13 Right-click the device which is connecting to the routed fabric, and select **Deploy Config**.

Step 14 In the **Config Deployment** window, click **Deploy Config**.

PART VIII

Templates Usage

- Template Usage in Cisco DCNM LAN Fabric Deployment, on page 999
- Guidelines for Programmable Reports, on page 1015
- Cisco DCNM Programmable Report APIs, on page 1021

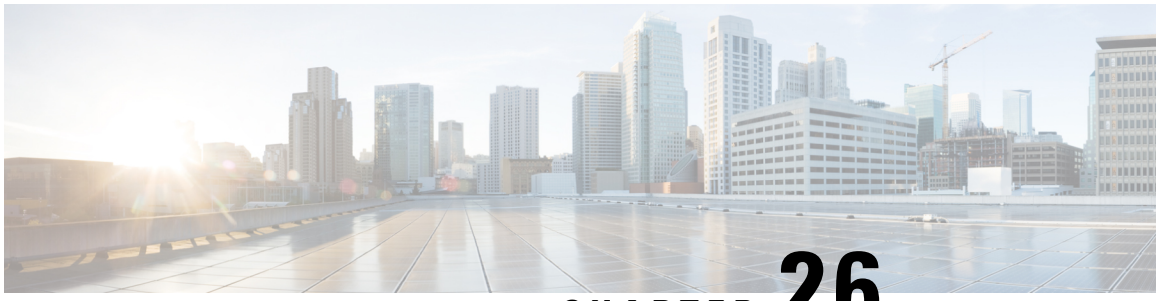

CHAPTER 26

Template Usage in Cisco DCNM LAN Fabric Deployment

templateType	Specifies the type of Template used.	• CLI • POLICY • SHOW • PROFILE • ABSTRACT

- Policy Template, on page 999
- Fabric Template, on page 1003
- Profile Template, on page 1003
- Viewing, Editing, and Adding Policies, on page 1004
- Deploying New Configurations, on page 1008
- switch_freeform Template Usage, on page 1009
- Changing the Contents of a Template in Use, on page 1012

Policy Template

For the policy template, there are two template content types: CLI and PYTHON. With CLI content type, the policy templates are 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:

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:

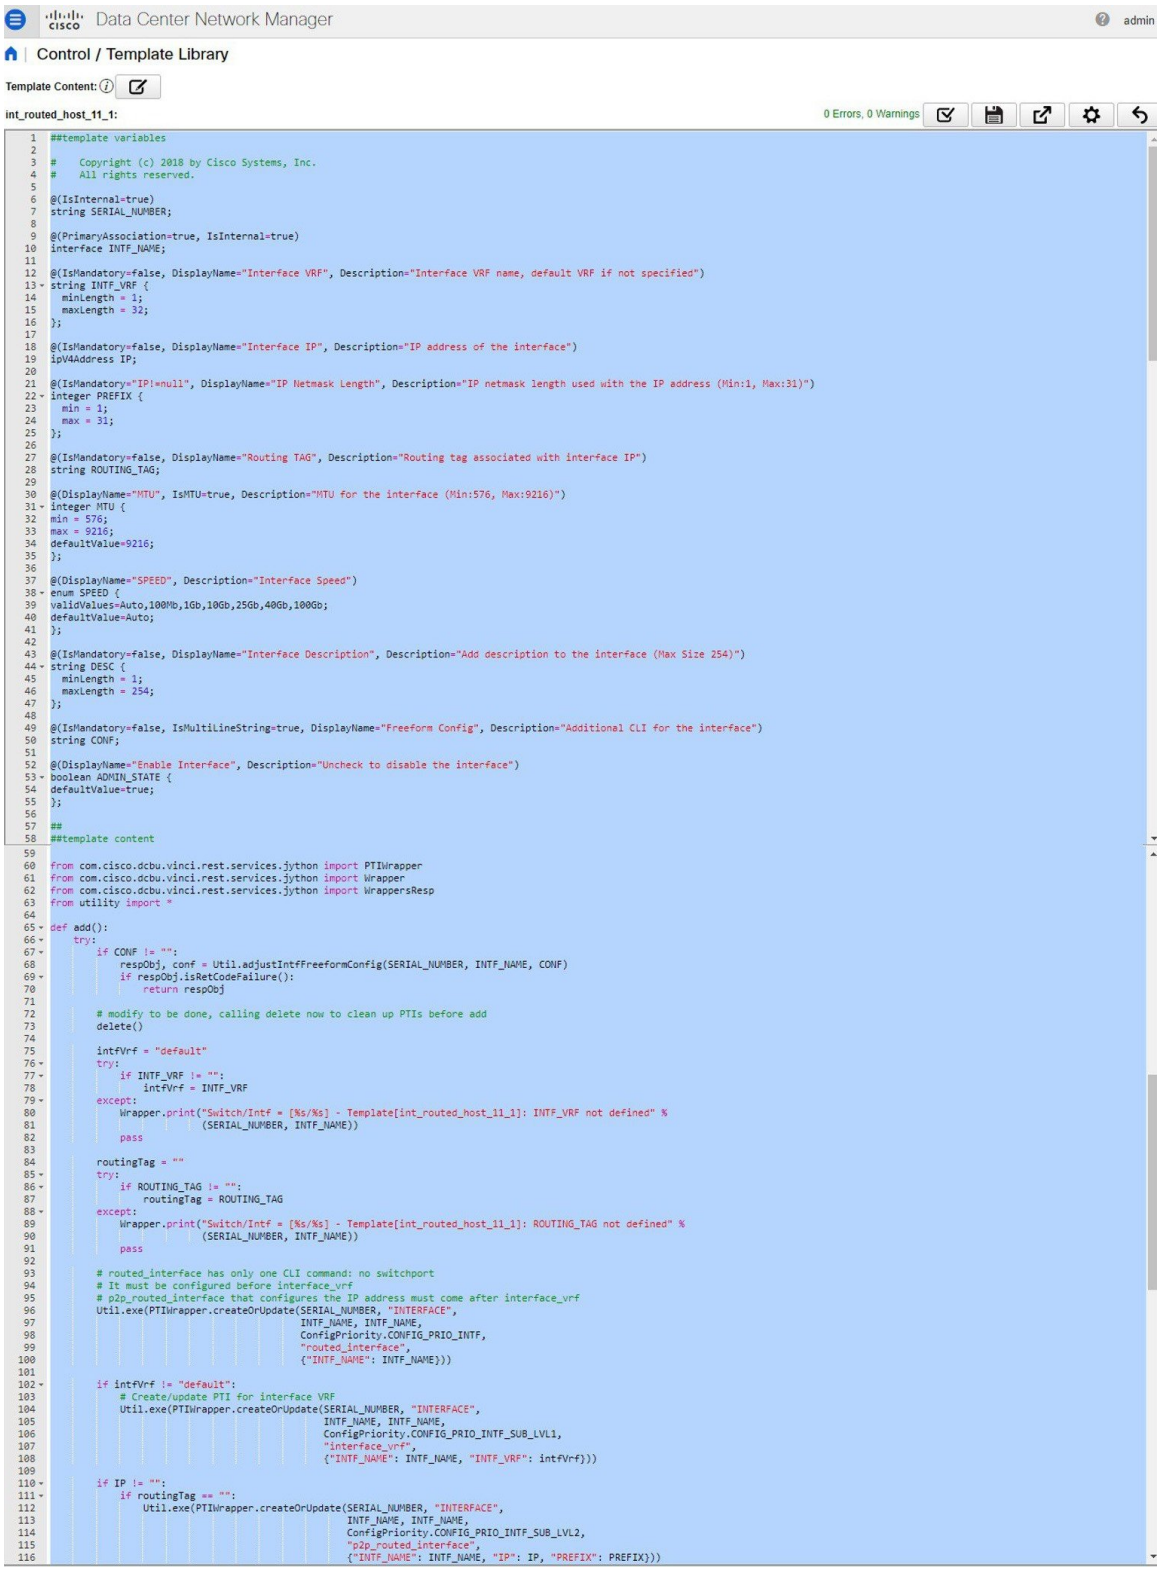


```
##template variables

#    Copyright (c) 2018 by Cisco Systems, Inc.
#    All rights reserved.

@(IsInternal=true)
string SERIAL_NUMBER;

@(PrimaryAssociation=true, IsInternal=true)
interface INTF_NAME;

@(IsMandatory=false, DisplayName="Interface VRF", Description="Interface VRF name, default VRF if not specified")
string INTF_VRF {
  minLength = 1;
  maxLength = 32;
};

@(IsMandatory=false, DisplayName="Interface IP", Description="IP address of the interface")
ipV4Address IP;

@(IsMandatory="IP!=null", DisplayName="IP Netmask Length", Description="IP netmask length used with the IP address (Min:1, Max:31)")
integer PREFIX {
  min = 1;
  max = 31;
};

@(IsMandatory=false, DisplayName="Routing TAG", Description="Routing tag associated with interface IP")
string ROUTING_TAG;

@(DisplayName="MTU", IsMTU=true, Description="MTU for the interface (Min:576, Max:9216)")
integer MTU {
min = 576;
max = 9216;
defaultValue=9216;
};

@(DisplayName="SPEED", Description="Interface Speed")
enum SPEED {
validValues=Auto,100Mb,1Gb,10Gb,25Gb,40Gb,100Gb;
defaultValue=Auto;
};

@(IsMandatory=false, DisplayName="Interface Description", Description="Add description to the interface (Max Size 254)")
string DESC {
  minLength = 1;
  maxLength = 254;
};

@(IsMandatory=false, IsMultiLineString=true, DisplayName="Freeform Config", Description="Additional CLI for the interface")
string CONF;

@(DisplayName="Enable Interface", Description="Uncheck to disable the interface")
boolean ADMIN_STATE {
defaultValue=true;
};

##
##template content

from com.cisco.dcbu.vinci.rest.services.jython import PTIWrapper
from com.cisco.dcbu.vinci.rest.services.jython import Wrapper
from com.cisco.dcbu.vinci.rest.services.jython import WrappersResp
from utility import *

def add():
    try:
        if CONF != "":
            respObj, conf = Util.adjustIntfFreeformConfig(SERIAL_NUMBER, INTF_NAME, CONF)
            if respObj.isRetCodeFailure():
                return respObj

        # modify to be done, calling delete now to clean up PTIs before add
        delete()

        intfVrf = "default"
        try:
            if INTF_VRF != "":
                intfVrf = INTF_VRF
        except:
            Wrapper.print("Switch/Intf = [%s/%s] - Template[int_routed_host_11_1]: INTF_VRF not defined" %
                          (SERIAL_NUMBER, INTF_NAME))
            pass

        routingTag = ""
        try:
            if ROUTING_TAG != "":
                routingTag = ROUTING_TAG
        except:
            Wrapper.print("Switch/Intf = [%s/%s] - Template[int_routed_host_11_1]: ROUTING_TAG not defined" %
                          (SERIAL_NUMBER, INTF_NAME))
            pass

        # routed_interface has only one CLI command: no switchport
        # It must be configured before interface_vrf
        # p2p_routed_interface that configures the IP address must come after interface_vrf
        Util.exe(PTIWrapper.createOrUpdate(SERIAL_NUMBER, "INTERFACE",
                                           INTF_NAME, INTF_NAME,
                                           ConfigPriority.CONFIG_PRIO_INTF,
                                           "routed_interface",
                                           {"INTF_NAME": INTF_NAME}))

        if intfVrf != "default":
            # Create/update PTI for interface VRF
            Util.exe(PTIWrapper.createOrUpdate(SERIAL_NUMBER, "INTERFACE",
                                               INTF_NAME, INTF_NAME,
                                               ConfigPriority.CONFIG_PRIO_INTF_SUB_LVL1,
                                               "interface_vrf",
                                               {"INTF_NAME": INTF_NAME, "INTF_VRF": intfVrf}))

        if IP != "":
            if routingTag == "":
                Util.exe(PTIWrapper.createOrUpdate(SERIAL_NUMBER, "INTERFACE",
                                                   INTF_NAME, INTF_NAME,
                                                   ConfigPriority.CONFIG_PRIO_INTF_SUB_LVL2,
                                                   "p2p_routed_interface",
                                                   {"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.

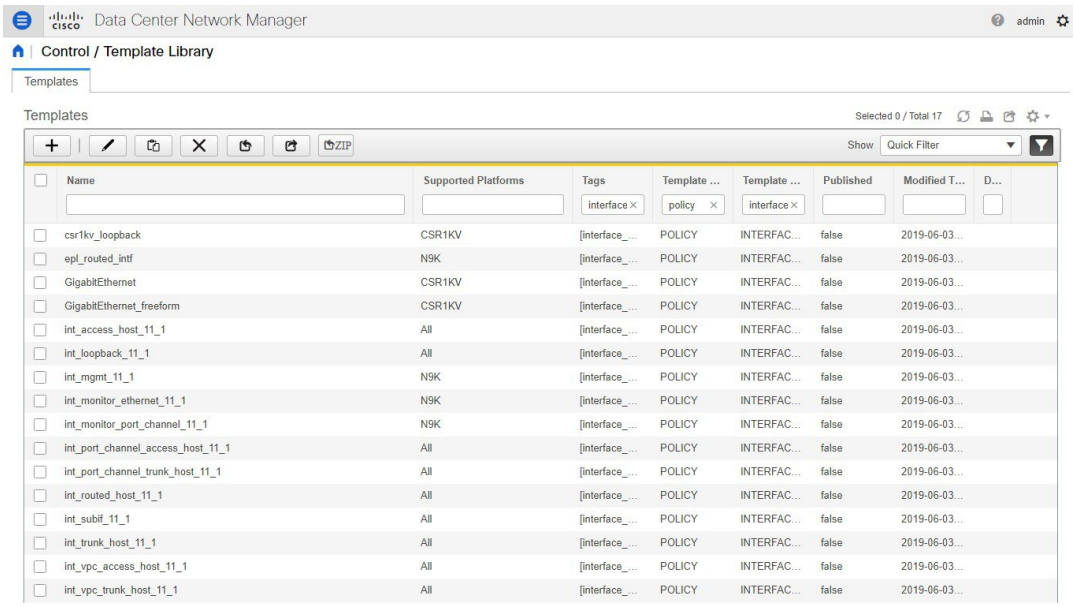

In the View/Edit Policies window on the Fabric Builder, you will only see device policy templates.

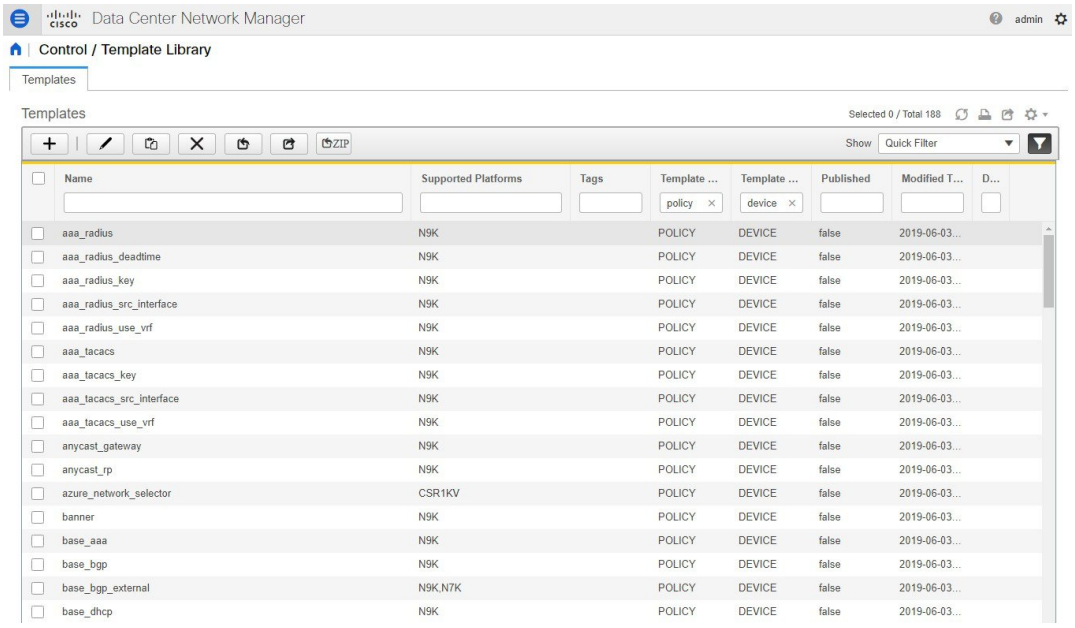

You can make a copy of any of these templates and customize them as per their needs. That is the typical use-case for customization. **Do not** modify existing policies but make a copy, and then customize as per the requirements. Otherwise, after a DCNM upgrade, the changes may be lost.

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).

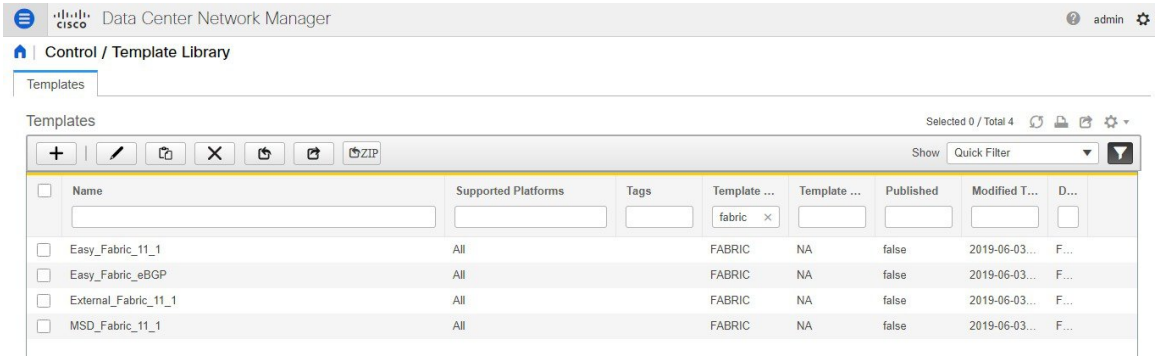

Profile Template

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*’)

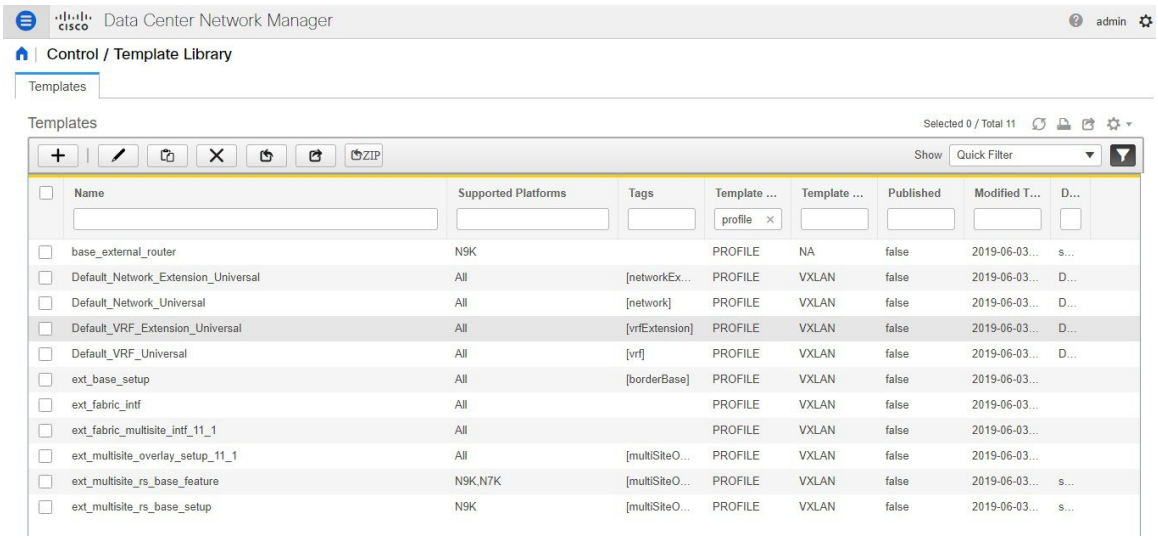

For more information about how to apply overlay configuration via the Networks & VRFs workflow in DCNM, see *Creating and Deploying Networks and VRFs* section.

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* section.

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.

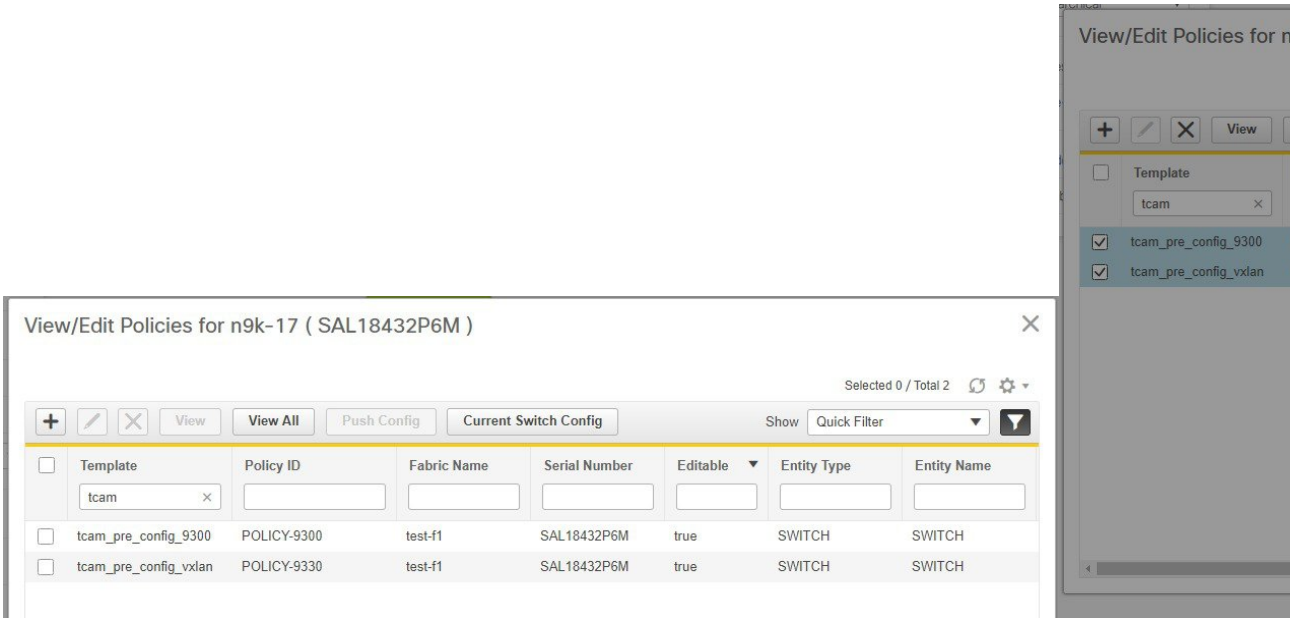

Example: Viewing Policies for an Interface

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.

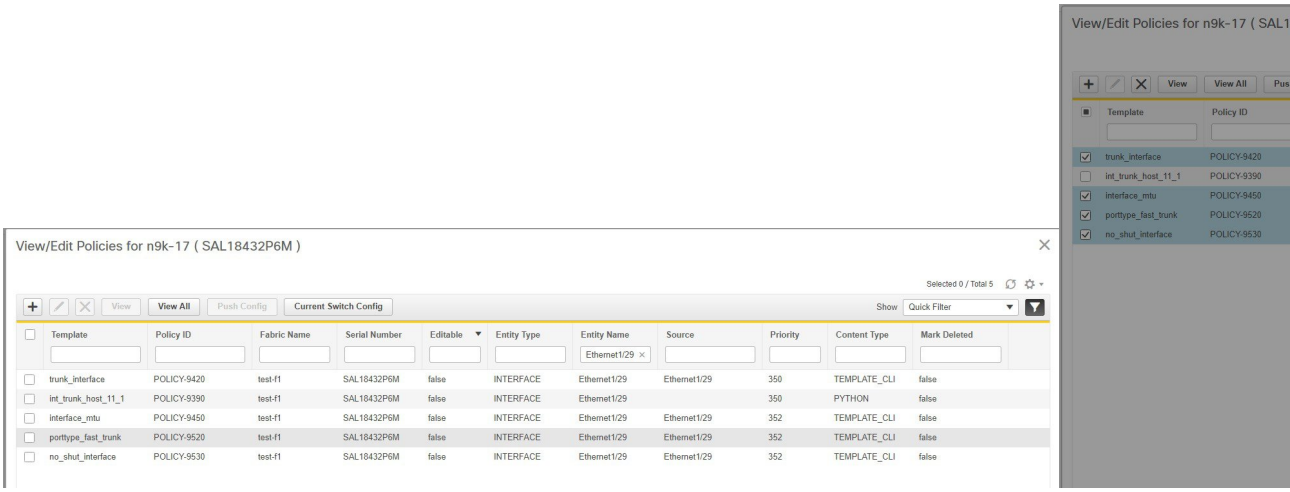

Note

- 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.

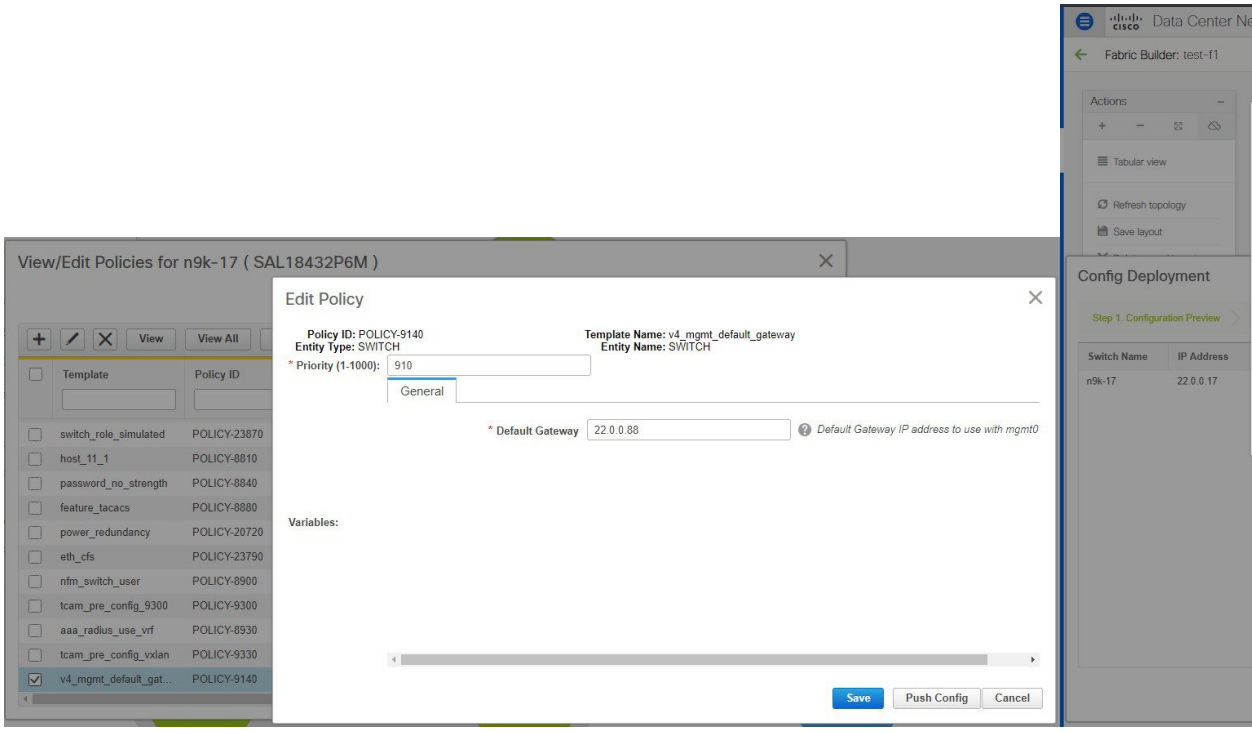

Adding Policies

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.

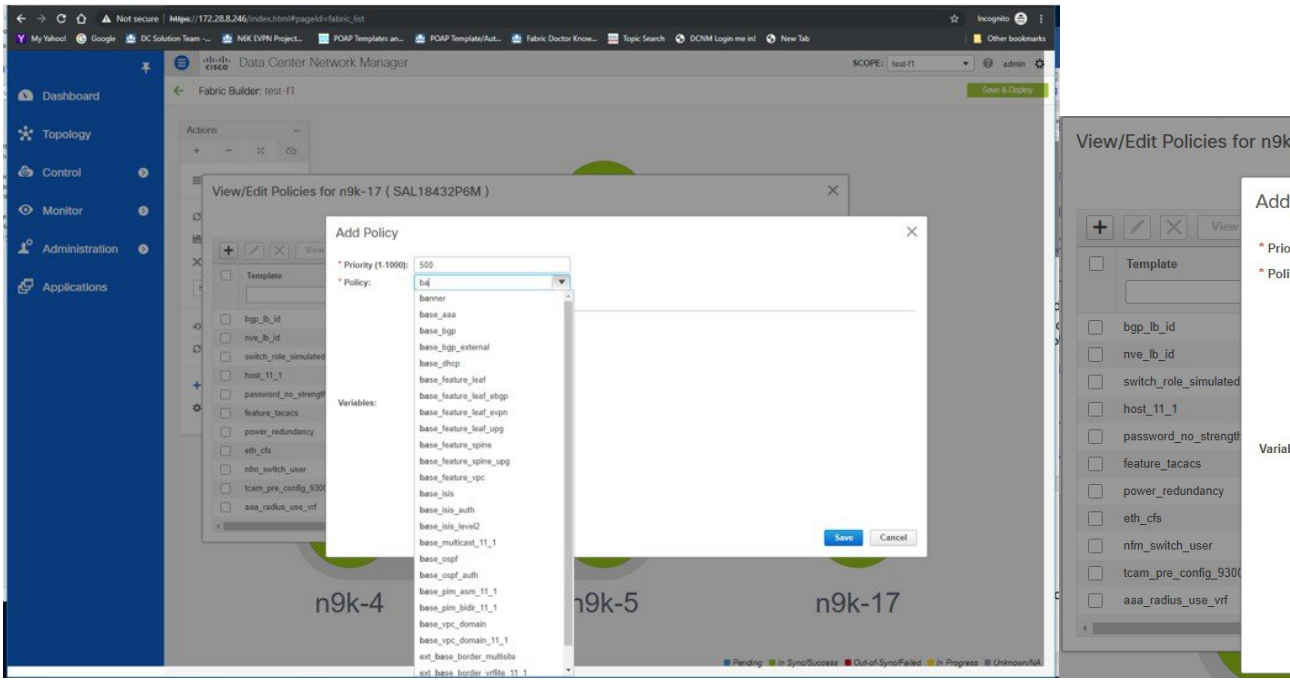

Deploying New Configurations

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.

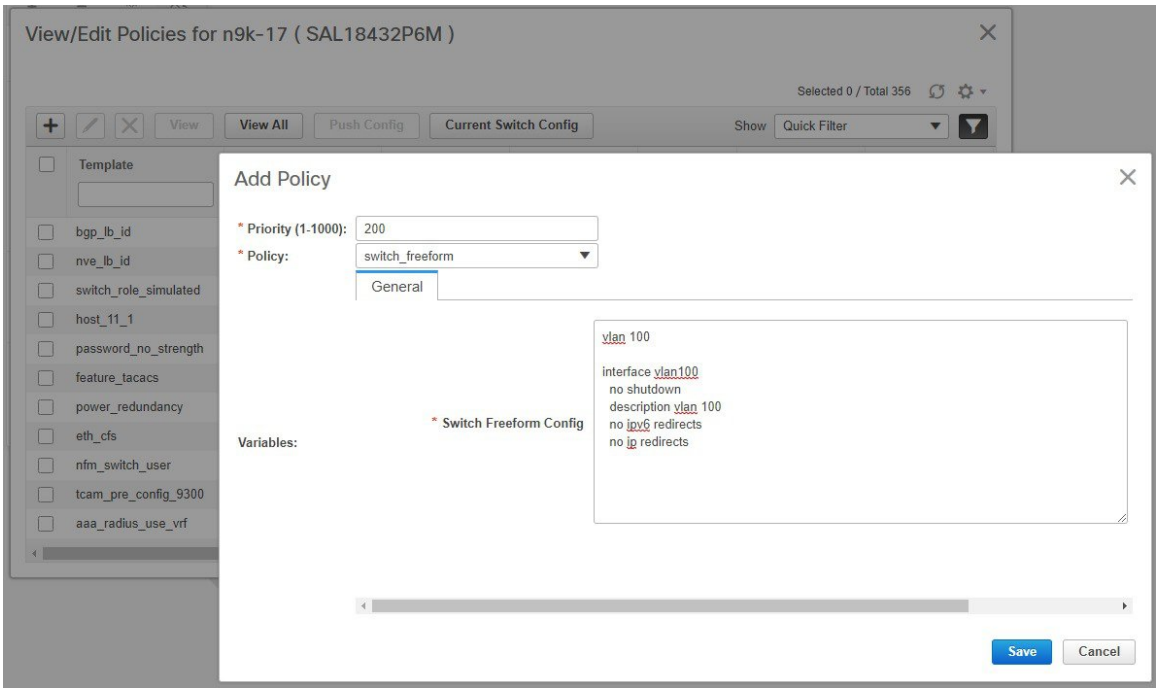

Step 2 View the intent config of the **switch_freeform** policy.

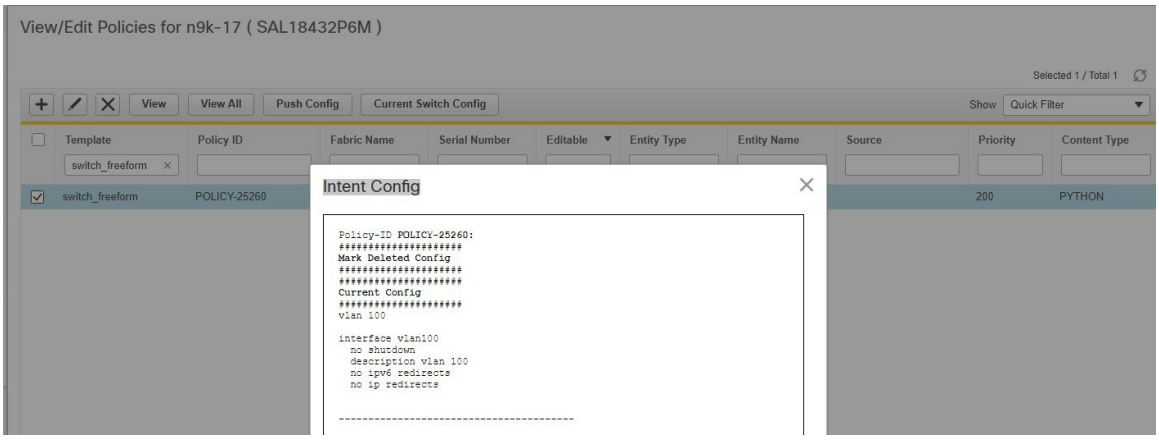

Step 3 Deploy the switch_freeform policy from Fabric Builder.

Step 4 Edit the switch_freeform policy from the View/Edit Policies window.

Change the config.

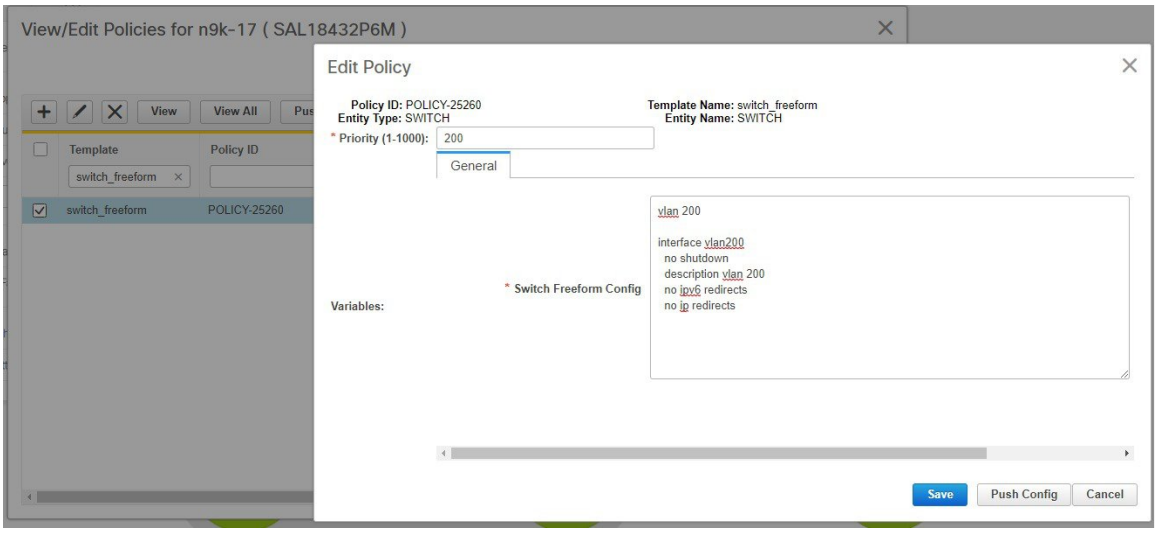

Step 5 Save the change.

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.

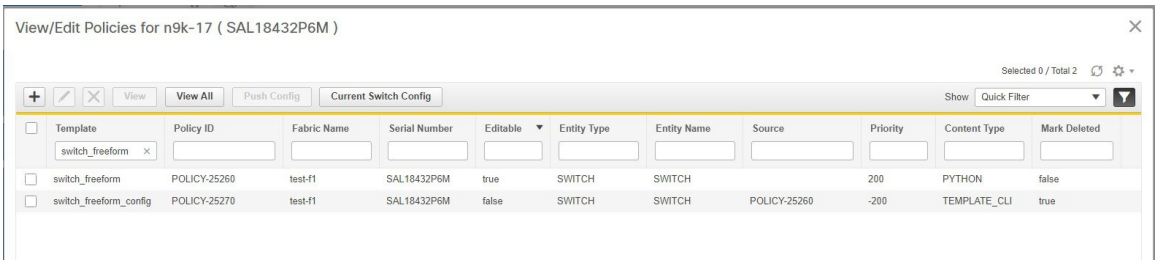

Step 6 View the intent config of the **mark deleted** internal policy.

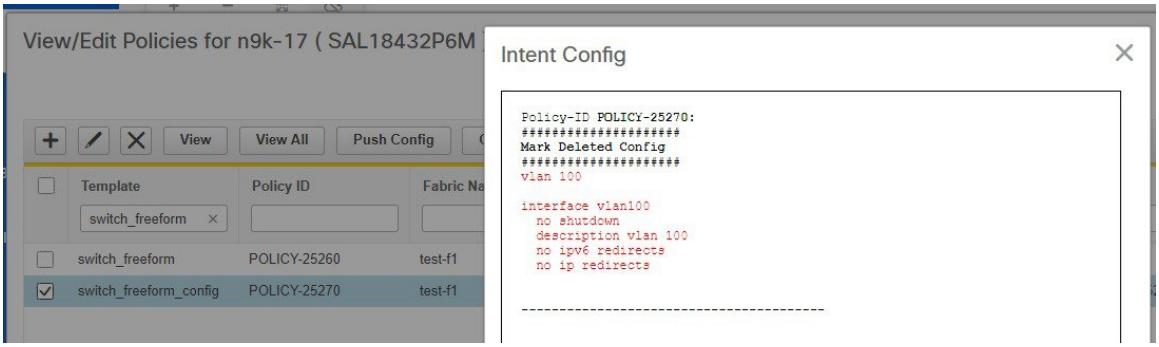

Step 7 View the intent config of the **changed** switch_freeform policy before deployment.

Note that both the **mark-deleted** and **current configs** are shown.

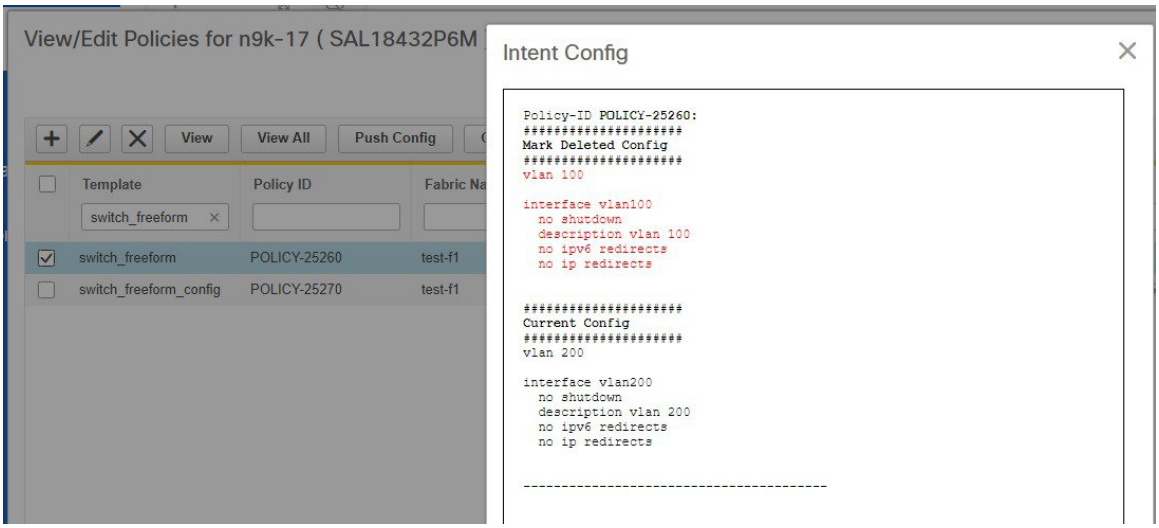

Step 8 Deploy the changed config from Fabric Builder.

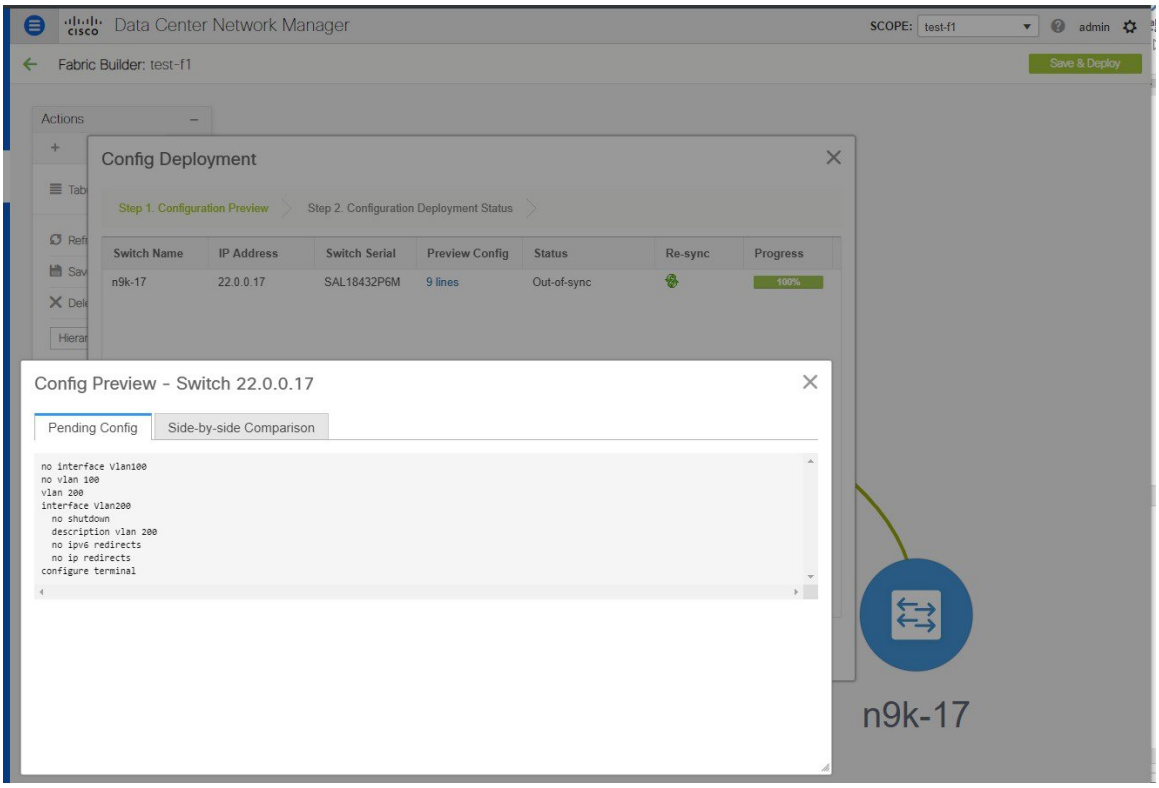

Changing the Contents of a Template in Use

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.

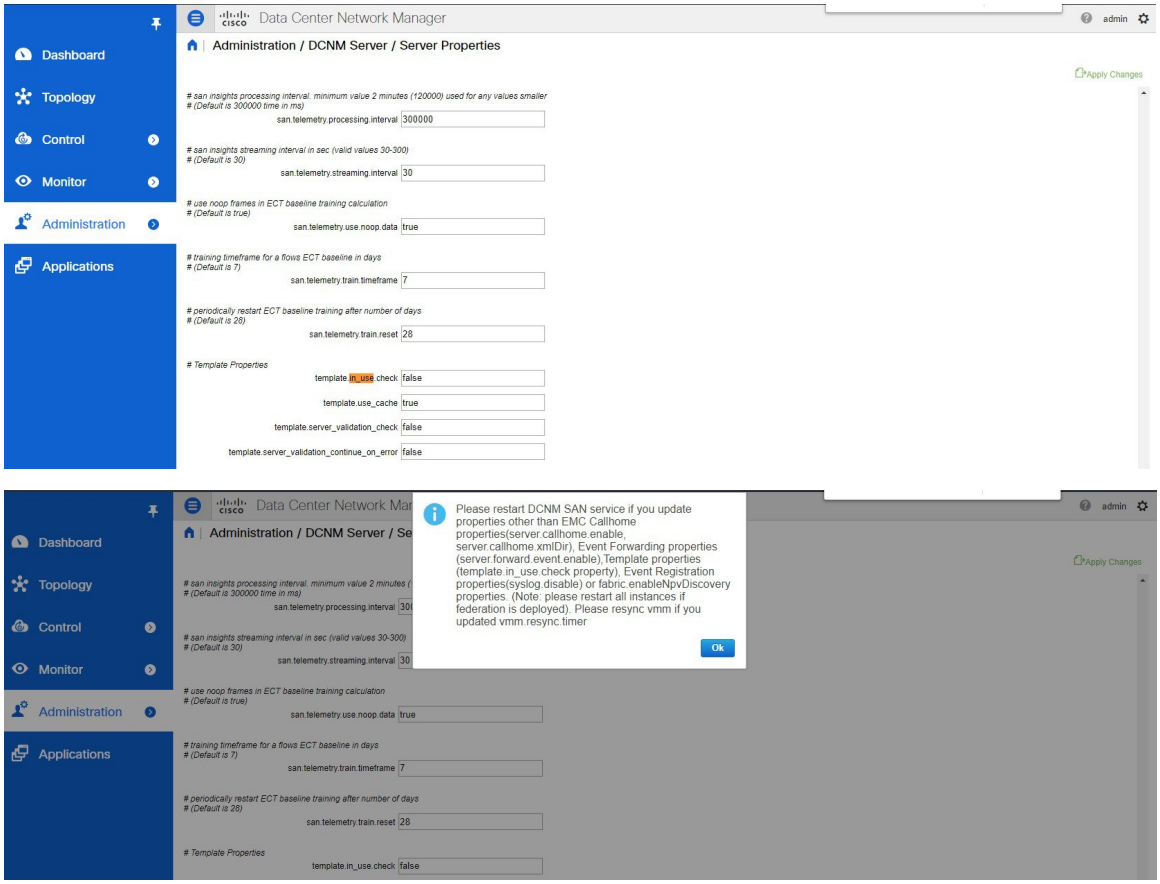

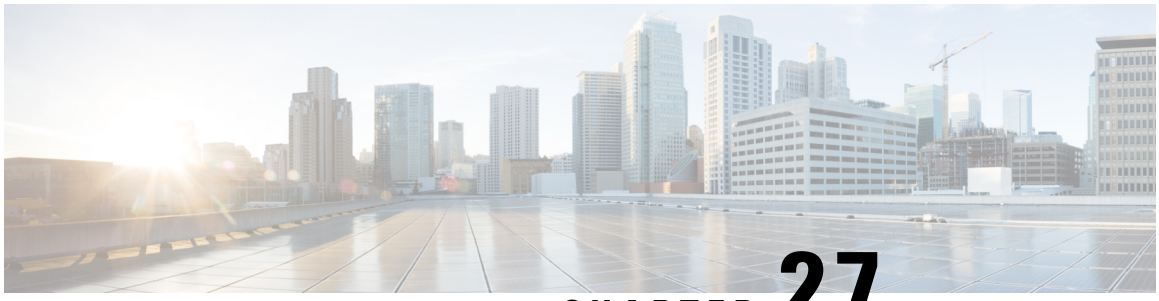

CHAPTER 27

Guidelines for Programmable Reports

- Prerequisites, on page 1015
- CLI Output Processing, on page 1016
- Report Template, on page 1017
- Template Content, on page 1018

Prerequisites

Planning

1. Determine if the report is meant to be run for a whole fabric or device(s).
2. Find out which **show** command(s) should be run on switch(es) to collect the required data.
   - Find out if the CLI output supports xml, json, or neither.
   - If neither, plain CLI output is returned by the switch.
   - Determine if the CLI response, including the command executed, needs to be stored in elasticsearch. You should be cautious as storing responses may increase storage tremendously.
3. Determine if you need to pre-validate the report creation input such as Recurrence, Period, etc. For example, does the report job support a periodic report, and how frequent should the job run?

Report Presentation

1. If you want **Summary**, choose the format to display the data:
   - Key-value pairs
   - Table
   - Chart (column, pie, line)
2. For **Section** (Detailed View), choose which format to display the data:
   - Key-value pairs
   - Array of JSON objects

- Chart (column, pie, line)

3. For **Formatters**, the following are applicable:
   - Add additional formatting to values displayed in UI
   - Markers supported: ERROR, SUCCESS, WARNING, INFO

Data Comparison between Reports

1. Determine if the report needs to compare data between the current report and an older report.
2. If yes, report infra APIs allow you to retrieve previous report(s) like:
   - One or more previously generated reports
   - The oldest report in the report job
   - Summary of a particular report
   - A particular section of a reportCLI Output Processing

CLI Output Processing

XML Format

If the CLI output returns data in XML format, you can use the report infra provided XML utilities to read the XML data out:

From reportlib.preport import *:

- getxmltree(xml_string, tag)
- getxmlrows(xml_tree, tag_xpath)
- getnodevalue(xml_tree, node_xpath)
- has_tag(xml_tree, tag)

For an example, see the report template **switch_inventory**.

JSON Format

If the CLI output returns data in JSON format, import Python's json module and use the **json.loads()** method to parse a JSON string.

import json

json_string = <CLI response>

json_obj = json.loads(json_string)

For an example, see the report template **fabric_nve_vni_counter**.

Plain CLI Output

If the CLI output returns data in the same format as seen on the CLI UI, you need to write your own parsing method to read the data out in the CLI response.

Logger

Logger allows you to log messages from a report template. Logged messages are written to /usr/local/cisco/dcm/fm/logs/preport_jython.log.

Report Template

Template Properties

Specify the following mandatory template properties:

```
name = <template-name>;
tags = fabric or device;
userDefined = true or false;
templateType = REPORT;
templateSubType = GENERIC;
contentType = PYTHON;
```

Note

- Set **tags = fabric** if report is run for a fabric; **tags = device** if report is run for a device
- Set **userDefined = true** if template is created by a customer; **userDefined = false** if template is created by a DCNM developer.

Template Variables

```
Specify the following template variables:
##template variables
@(IsInternal=true)
string fabric_name or serial_number;
string user_input;
```

Note

- Configure variable *fabric_name* if **tags = fabric**
- Configure variable *serial_number* if **tags = device**
- User variables are optional. All data types and annotations supported by DCNM template infra can be used.

Template Content

Imported Libraries

The following 2 python libraries are required. Note that **reportlib.preport** contains all reporting infrastructure APIs.

```
##template content
from com.cisco.dcbu.vinci.rest.services.jython import WrappersResp
from reportlib.preport import *
```

Template Functions

generateReport()

generateReport() is the entry function and is invoked while generating a report. All the report implementation logic should be provided here. This function takes a context object. The 'context' parameter is created by the report infrastructure when a report job is created.

```
def generateReport(context):

    report = Report("Report title")   ## Create a report object
     ## Gather data and fill in content for the report

    respObj = WrappersResp.getRespObj()
    respObj.setSuccessRetCode()
    respObj.setValue(report)
    return respObj
```

Note

- This function must return a WrappersResp object.
- If there is no error in generating the report, the report object created within this function must be set in **WrappersResp.setValue()** before the **WrappersResp** object is returned.

Run CLI and Process CLI Response

Below is a sample code on how to send **show** command(s) to one or multiple devices, and on how to process the responses from the device(s).

```
show_cmd1 = 'show xxx'
show_cmd2 = 'show yyy'
device_list = [device1,device2]
## run the command(s) on each device in the device_list


cli_responses = show(device_list, show_cmd1, show_cmd2)
## run the command(s) on each device in the device_list and store the CLI response(s)

cli_responses = show_and_store(device_list, show_cmd1, show_cmd2)
```

For resp in cli_responses:

```
command = resp['command'].strip()

    if show_cmd1 in command:
        cmd1_response = resp['response'].strip()
        ## process show_cmd1 response
```

```
    elif show_cmd2 in command:
        cmd2_response = resp['response'].strip()
        ## process show_cmd1 response
```

validate()

The **validate()** function is an optional function and is used to perform pre-validation of report creation input such as Recurrence, Period, etc. This function, if defined, is called while creating a report job. The report job is only created if this function returns a WrappersResp with SuccessRetCode. If validation fails, a WrappersResp with FailureRetCode with errors should be returned.

```
def validate(context):
    respObj = WrappersResp.getRespObj()
    ## Validation content

    if validation_failed:
        respObj.addErrorReport(…)
        respObj.setFailureRetCode()
    else:
        respObj.setSuccessRetCode()
    return respObj
```

report.add_summary

Each report can have one summary and the content is a python dictionary.

```
summary = report.add_summary()
summary[key] = value
summary.add_message(msg)
## Present the summary in a table format

table = summary.add_table(title,_id)   ## _id must be a unique id for the table
table.append(value,_id)   ## adding rows to table
## Present the summary in a chart format

chart = summary.add_chart(ChartType,_id)
## ChartTypes: ChartTypes.COLUMN_CHART, ChartTypes.PIE_CHART, ChartTypes.LINE_CHART
```

report.add_section

Section is a logical grouping of report content. Section details are displayed in **View Details**.

```
section = report.add_section(title,_id)   ## _id must be a unique id for the section
section[key] = value
section.append(key, json_obj, _id)   ## adding rows of json objects to section
## Present the section details in a chart format

chart = section.add_chart(ChartType,_id)
## ChartTypes: ChartTypes.COLUMN_CHART, ChartTypes.PIE_CHART, ChartTypes.LINE_CHART
```

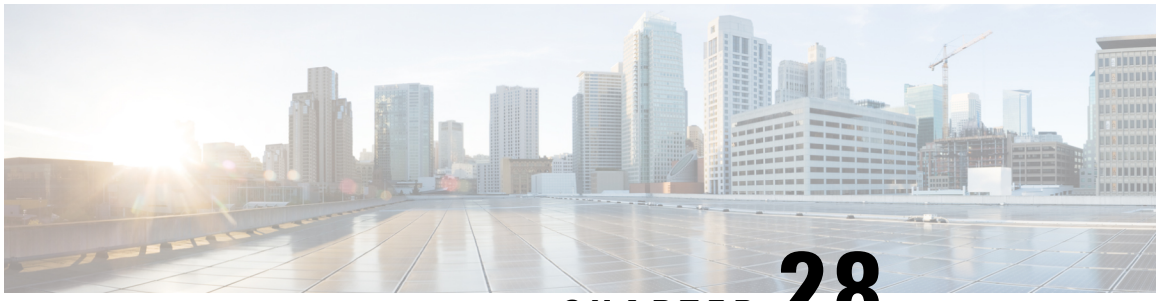

CHAPTER 28

Cisco DCNM Programmable Report APIs

- Template, on page 1021
- Template Functions, on page 1022
- Report Layout, on page 1023
- Report Python Library, on page 1025

Template

In Cisco DCNM Release 11.4(1), the new template type “REPORT” is added with following two subtypes: UPGRADE and GENERIC. The template type is python and requires providing an implementation of method “generateReport”.

UPGRADE

UPGRADE templates are used for ISSU pre and post ISSU. These templates will be listed in ISSU wizard.

GENERIC

The GENERIC template can be used for any generic reporting purpose. For example, collecting inventory report.

Template Structure

The following image shows an example template structure:

```
##template variables


@(IsInternal=true)
String serial_number/fabric_name;

String user_input;

##
##template content


from com.cisco.dcbu.vinci.rest.services.jython import WrappersResp
from reportlib.preport import *

def validate(context):

    respObj = WrappersResp.getRespObj()
    respObj.setSuccessRetCode()
    return respObj


def generateReport(context):

    report = Report("Report title")

    ##Report content

    respObj = WrappersResp.getRespObj()
    respObj.setSuccessRetCode()
    respObj.setValue(report)
    return respObj
##
```

serial_number or fabric_name based on the scope selected while scheduling the report. In case of UPGRADE report, always serial number will be injected

Template variable section. **All data types, annotations supported in DCNM template can be used here. User should provide these inputs while creating the report

Import necessary python lib.

Report can have optional validation method, This method will be invoked while creating the report job. Job will be created only if this method return success. This method is invoked only once and serial_number or fabric_name will not be available inside this method. More information check the API guide

report must provide implementation of generateReport(context) method.

generateReport method should return an object of type WrapperResp. Report object created above must be stored in wrappersResp using wrappersResp.setValue() API

Template Functions

generateReport Method

The generateReport method is invoked while generating the report. All the report implementation logic should be provided. This method accepts context object. As mentioned above, this method should return WrappersResp object.

Validation Method

The Validation method is optional. If the template defines this method, report application calls this method to perform pre validation while creating the job. This method is called only when the job is created and invoked only once irrespective of device or fabrics selected.

If the validation is passed, this method should return WrappersResp with SuccessRetCode, and for failure FailureRetCode with error in error list.

For example:

Validation failed

```
def validate (context):
    respObj = WrappersResp.getRespObj()
```

```
    ## Validation logic here

    respObj.setFailureRetcode()
    respObj.addErrorReport(template_name,error)
    return respObj
```

Validation success

```
def validate (context):
    respObj = WrappersResp.getRespObj()

    ## Validation logic here

    respObj.setSuccessRetcode()
    return respObj
```

You can perform validation based on content of context parameter.

Context Parameter

Context parameter consists of following attributes:

1. User name: Name of the user who created the job
2. User role: Role of the user who created the job
3. Job Id
4. Recurrence: NOW, ONCE, DAILY, WEEKLY, MONTHLY, ONDEMAND, or PERIODIC
5. Period: If the recurrence is periodic, then period will have frequency selected. For example, 10 MINUTES.

To read these values from context, see the APIs mentioned in *Get Job Context Information*.

Report Layout

A report contains the following components:

1. Summary
   a. Key and values
   b. Messages – Inferences
2. Details/Sections
   a. Key and values
   b. A JSON document – Cards
   c. Array of JSON Documents – Tables
3. Command log

Summary View

This view shows summary for each entity included in the report.

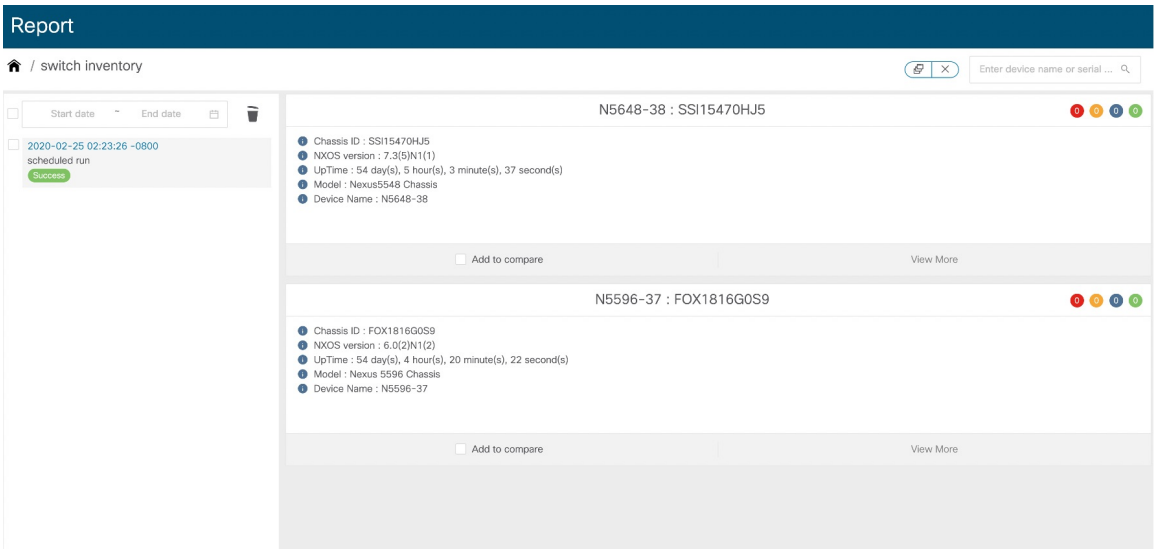

Detail View

The Detail view displays complete report JSON data along with summary. Report detail is logically grouped into sections. Each section is displayed separately with a collapsible widget.

Both summary and detail views provide counts of number of errors, warnings, info, success messages generated in the report.

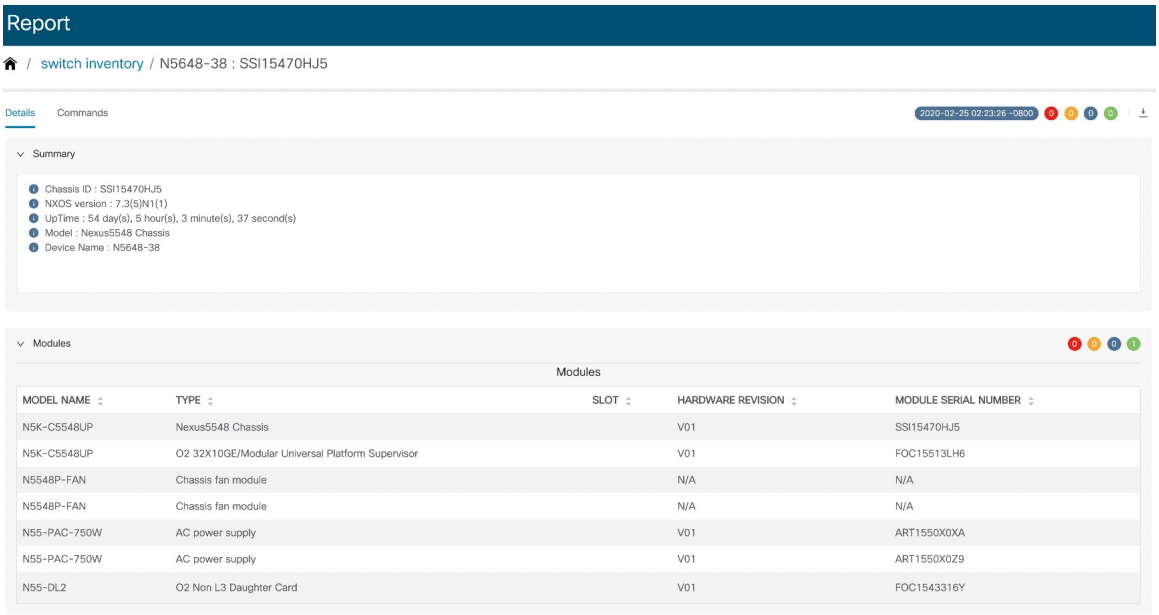

Command Log

Command log contains all commands executed in the report, based on the API used to execute the commands.

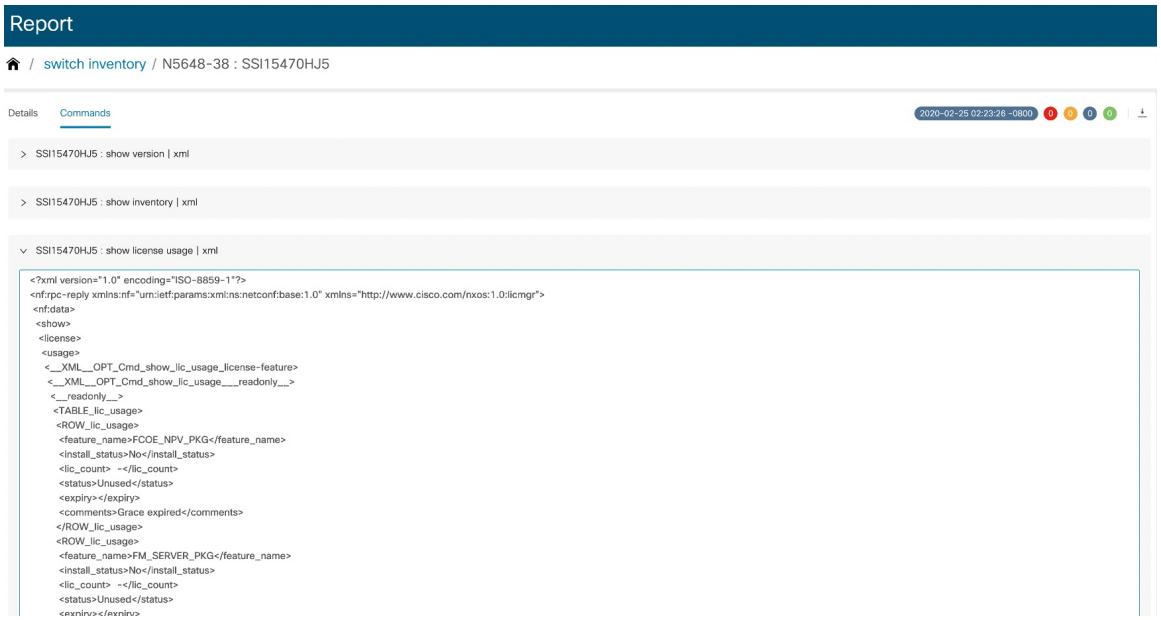

Report Python Library

Reporting infrastructure provides an easy to use and light weight python library to generate the report JSON model. To use this API, you should add following import statement in the template:

```
from reportlib.preport import Report
```

Report APIs

Create Report Object

Every report should create a “Report” object as the first step.

```
report = Report (“Report title”)
```

Add Summary

Every report can have one summary and it’s a python dictionary. Summary can be added as follows:

```
summary = report.add_summary()
```

Adding Content to Summary

Key and values

```
summary [‘NXOS Version’] = ‘8.1(0)’
```

Messages – Inferences

```
summary.add_message (“Simple message”)
```

Note In DCNM 11.4(1), DCNM does not support JSON object as value in summary. Following example is not supported.

```
summary[“info”] = {“key”:”value”,”key-2”:”value-2”}
```

Tables in Summary

```
table = summary.add_table(title,_id)
```

- title: Table title
- _id : Unique identifier for the table

Adding Rows to Table

```
table.append(value, _id)
```

- value: A JSON object. Nested json not supported.
- _id : Unique identifier for the table

For example:

```
table.append({'column1': 'value1','column2’:’value2’},” FOX1816G0S9”)
```

Add Section

Section is a logical grouping of report contents. It’s up to the user to create these sections and add information to be displayed.

Section can be added as shown:

```
section = report.add_section (“Section title”,_id)
```

- _id : Unique identifier for the table
- section : It is a dictionary

Adding Content to a Section

Key and values

You can add simple key and value pair to section as shown below:

```
section[‘key’] = ‘value’
```

A JSON document – Cards

A single JSON document can be added as same as any key value pair. Nested JSON is not supported in 11.4(1)

```
section[‘key’] = {‘key’:’value’,’key-2’:’value’}
```

The JSON document is displayed in a card widget as shown:

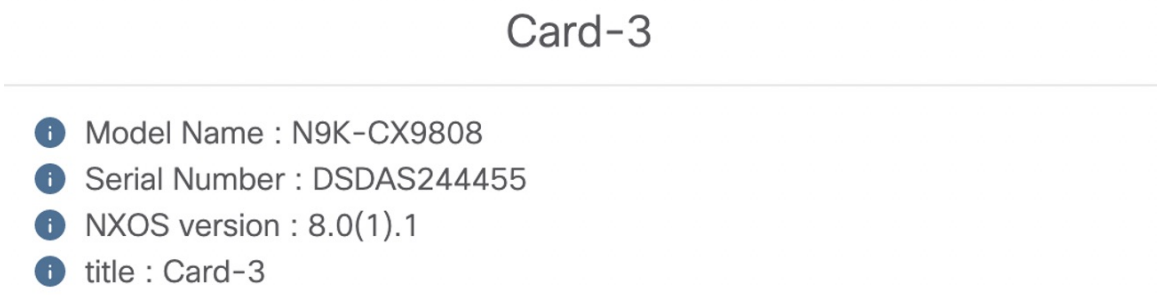

Array of JSON Documents – Tables

The **section.append** API allows user to create a table and add rows to it with following restriction:

1. All JSON document should have same set of keys
2. Nested JSON is not supported

```
section.append(key, dictionary, _id)
```

_id: Unique identifier which uniquely identifies a row in a table. Duplicate _id resultx in Unique id violation error.

For example:

```
section.append('Switch Details', {'name': 'N5K'},'DSDAS244455')
section.append('Switch Details', {'name': 'N6K'}, 'CSDAS244456')
section.append('Switch Details', {'name': 'N7K'}, 'ASDAS244457')
```

Formatters

Formatter allows you to add additional formatting to values displayed in UI.

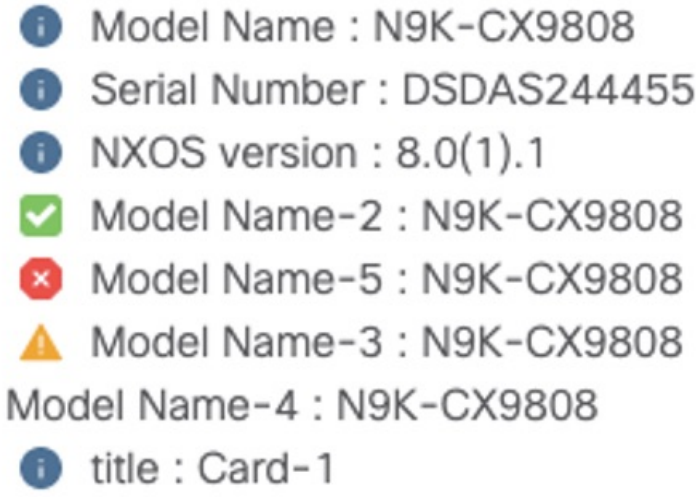

As shown, you can mark values as:

1. ERROR
2. SUCCESS
3. WARNING
4. INFO

When you add these markers to report, corresponding counts error, warning, success, info are automatically updated to be displayed in the UI.

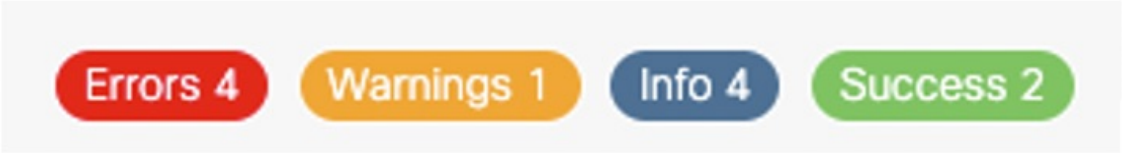


```
Formatter.add_marker(value,marker)
```

- value: Value to add marker.
- marker: Marker.ERROR,Marker.SUCCESS,Marker.WARNING,Marker.INFO

For example:

```
Formatter.add_marker ("NXOS version",Marker.INFO)
```

Chart

Report supports adding chart in both summary and section.

Adding Chart to Summary

```
report = Report("title")
summary = report.add_summary()
summary.add_chart(ChartType, _id)
```

- **ChartType**: ChartTypes.COLUMN_CHART, ChartTypes.PIE_CHART, ChartTypes.LINE_CHART.
- **_id**: Unique ID for the chart.

Adding Chart to Section

```
report = Report("title")
section = report.add_section("Section title",_id)
section.add_chart(ChartType, _id)
```

- **ChartType**: ChartTypes.COLUMN_CHART, ChartTypes.PIE_CHART, ChartTypes.LINE_CHART
- **_id**: Unique ID for the chart

Pie Chart

Set and subtitle title

```
pie_chart.set_title("Chart title")
pie_chart.set_subtitle("Sub title")
```

Add value

```
pie_chart.add_value("key",value)
```

- key: String key
- value: Numeric value

Column Chart

Set and subtitle title

```
column_chart.set_title("Chart title")
column_chart.set_subtitle("Sub title")
```

Set X-Axis and Y-Axis title

```
column_chart.set_xAxis_title("X-Axis title")
column_chart.set_yAxis_title("y-Axis title")
```

Add Value

```
bar_chart.add_value("key",value,category)
```

- key: String key
- value: Numeric value
- category: Bar chart divides the data into logical group called "category". A given key should have value in each category.

For example, device count is a key and Fabric Names are categories. Chart should have Device count for each fabric that is, each category.

Line Chart

Set and Subtitle Title

```
line_chart.set_title("Chart title")
line_chart.set_subtitle("Sub title")
```

Set X-Axis and Y-Axis title

```
line_chart.set_xAxis_title("X-Axis title")
line_chart.set_yAxis_title("y-Axis title")
```

Add Value

```
line_chart.add_value("key",value,category)
```

- key: String key
- value: Numeric value
- category: Line chart divides the data into logical group called "category". A given key should have value in each category.

For example, device count is a key and Fabric Names are categories. Chart should have Device count for each fabric, that is, each category.

Run CLIs on Device

Show Command

```
from reportlib.preport import show
cli_responses = show (serial_number ,*commands)
```

- serial_number: Serial number of the device to run commands. In case of VDC serial number should be serial_number:vdc_name. You can pass list of serial number to execute the same set of commands on multiple devices.
- *commands: Commands to run on device. It's a var args. You can specify multiple commands.

Examples:

- Executing command on single switch:

  ```
  cli_responses = show("FOX1816G0S9",'show version | xml', 'show inventory | xml', 'show
   license usage | xml')
  ```

- Executing command on multiple switches:

  ```
  cli_responses = show( ["FOX1816G0S9","SSI15470HJ5"],'show version | xml', 'show inventory
   | xml', 'show license usage | xml')
  ```

Show Commands and Store Response

```
from reportlib.preport import show_and_store
cli_responses = show_and_store(report,serial_number,*commands)
```

report: Report Object created.

serial_number: Serial number of the device to run commands. In case of VDC, serial number should be serial_number:vdc_name. You can pass a list of serial number to execute the same set of commands on multiple devices

*commands: Commands to run on device. It's a var args. You can specify multiple commands.

Examples:

- Executing command on single switch:

  ```
  cli_responses = show_and_store(report, "FOX1816G0S9", 'show version | xml', 'show
  inventory | xml', 'show license usage | xml')
  ```

- Executing command on multiple switches:

  ```
  cli_responses = show_and_store(report, ["FOX1816G0S9","SSI15470HJ5"], 'show version |
  xml', 'show inventory | xml', 'show license usage | xml')
  ```

Caution: This API stores the response from the device in elasticsearch along with report. User should be cautious while using this API, since storing all response may increase storage drastically.

Return Value

The Return Value API will return list of responses, and each response is a dictionary with following structure:

```
{
 'status': 'success' | 'failed,
 'response':<response from device>,
 'command':<cli command>,
 'serial_number': <device serial number>
}
```

In case of multiple switches, the response still be a list of responses with entries for each switch.

```
[
{
 'status': 'success',
 'response':<response from device>,
 'command':'show version',
 'serial_number': 'FOX1816G0S9'
},
{
 'status': 'success',
 'response':<response from device>,
 'command':'show version',
 'serial_number': 'SSI15470HJ5
}
]
```

Get Job Context Information

Get Recurrence Selected While Scheduling the Job from APP

```
get_recurrence(context)
```

This API returns the recurrence selected while creating the job. Returns value can be NOW,ONCE,DAILY,WEEKLY,MONTHLY,ONDEMAND, and PERIODIC.

get_period

If job is scheduled as Periodic, then period information can be accessed using the API:

```
period = get_period(context)
period.get_period() will return the period
period.get_time_unit() will return time Unit (HOURS, MINUTES)
```

Analyze with Historical Reports

Get Previously Generated Reports

The "get_previous_reports()" method allows to get reports generated in the past. This can be used to perform analysis based on current data and historical data. This API will return the report in descending order of created time.

```
List of reports = get_previous_reports (context,entity,count)
```

This API returns a list of reports.

context: The object received as input from generateReport(context) method.

entity: serial_number or fabric name.

count: Number of reports to fetch.

Get Oldest Report

```
oldest_report = get_oldest_report(context,entity)
```

context: The object received as input from generateReport(context) method

entity: serial_number or fabric name

Both the above APIs return Report object with the following API to retrieve information:

1. **Get summary** : report.get_summary()
2. **Get section** : report.get_section(_id)

```
report.get_section(_id)
```

_id: Unique Identifier for the section.

XML Utilities

Get XML Tree

```
from reportlib.preport import getxmltree
xml_element_tree = getxmltree(xml_string,tag)
```

This API returns the XML tree under the given tag.

xml_string: XML response from device.

tag: XML tag. Complete XML under this tag will be returned as ElementTree.

xml_element_tree: This API returns xml.etree.ElementTree object.

Get XML Rows

If the CLI response has rows, you can get the array of rows by using the getxmlrows API.

```
from reportlib.preport import getxmlrows
rows = getxmlrows(xml_tree,tag_xpath)
```

xml_tree: xml.etree.ElementTree object

tag_xpath: xpath of the XML record. For more info, see https://docs.python.org/2/library/xml.etree.elementtree.html#xpath-support.

rows: Array of rows.

Get Node Value

XML node value can read using the **getnodevalue** API. This API should be used get the node value of primitive type.

```
from reportlib.preport import getnodevalue
value = getnodevalue(xml_tree,node_xpath)
```

Check Whether Node Exists

```
from reportlib.preport import has_tag
has_tag(xml_tree,tag)
```

This API returns true or false based on whether the given tag is present in XML tree.

WrapperResp

Every report should return an object of the type WrapperResp.

WrapperResp can be instantiated as:

```
respObj = WrappersResp.getRespObj()
```

The return code in WrapperResp indicates whether the report ran successfully or not.

1. If all commands are run and required information is extracted, then report returns success **respObj.setSuccessRetCode()**.
2. In case of any exception like commands failure, then report returns failure **respObj.setFailureRetCode()**.
3. In case of an error, you can add the reason for error as **respObj.addErrorReport(template_name,error_message)**.

The report object created in the Report section should be set to value of WrappersResp as shown:

```
respObj.setValue(report)
```

Logger

Logger allows you to log messages from report template. All information logged using the logger is logged to: "`/usr/local/cisco/dcm/fm/logs/preport_jython.log`".

```
Logger.info("message")
Logger.debug("message")
Logger.error("message")
Logger.trace("message")
Logger.warn("message")
```