Easy Provisioning in Cisco DCNM LAN Fabric Deployments, Release 11.3(1)

First Published: 2020-05-29

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

CONTENTS

CHAPTER 1 **Configuring a Fabric with eBGP Underlay** 1

Creating a New Fabric for EBGP-Based Underlay 1

Applying Policies On A Fabric With An eBGP Underlay 15

Deploying Fabric Underlay Policies 15

Deploying Fabric Overlay Policies 16

Deploying Spine Switch Overlay Policies 16

Deploying Leaf Switch Overlay Policies 17

Dual-AS Fabric Deployment 18

Overview of Networks in a Routed Fabric 19

Creating and Deploying a Network in a Routed Fabric 20

Creating Inter-Fabric Links Between a Routed Fabric and an External Fabric 23

CHAPTER 2 **Configuring a VXLANv6 Fabric** 27

Overview 27

Creating a VXLAN Fabric with IPv6 Underlay 28

CHAPTER 3 **Configuring ToR Switches and Deploying Networks** 33

Overview 33

Supported Topologies for ToR Switches 33

Configuring ToR Switches 39

Deploying Networks on ToR Switches 44

CHAPTER 4 **Connecting Cisco Data Center and a Public Cloud** 47

Connecting Cisco Data Center and a Public Cloud 47

Topology Overview 48

Guidelines and Limitations 49

Prerequisites 49
Task Summary 49
Setting the Polling Time 50
Setting Up the On-premise External Fabric with CSR 1000v 51
Creating an External Fabric 51
Discovering the On-Premises Core Router 51
Setting Up the VXLAN EVPN Fabric 52
Creating a VXLAN EVPN Fabric 52
Assigning the BGW Role 53
Setting Up the External Fabric with CSR in Azure 53
Creating an External Fabric 53
Discovering the Core Router 54
Setting Up the MSD Fabric for Connectivity 55
Creating an MSD Fabric 55
Moving Other Fabrics into the MSD Fabric 56
Setting Up Connections 57
Connecting the On-Premises BGW and the On-Premises Core Router 57
Connecting the On-prem Core Router and the Public-cloud Core Router with IPsec Tunnel 58
Connecting the On-prem BGW and the Public-cloud Core Router using EVPN Peering 60
Saving and Deploying Configurations 62
Extending VRFs 63
Deploying and Extending the VRF On-prem Core Router 63
Creating and Deploying VRF on Public Cloud 64
Configuring Default Gateway for the VM 65
Verifying the Connectivity 66
Deploying Cisco CSR 1000v on Microsoft Azure 67
Viewing Links and Core Routers Details 71
Resetting Packet Counter Using API 71

CHAPTER 5 **Managing a Brownfield VXLAN BGP EVPN Fabric 73**
Overview 73
Prerequisites 74
Guidelines and Limitations 74
Fabric Topology Overview 76

DCNM Brownfield Deployment Tasks 77
Verifying the Existing VXLAN BGP EVPN Fabric 77
Creating a VXLAN BGP EVPN Fabric 80
Adding Switches and Transitioning VXLAN Fabric Management to DCNM 94
Verifying the Import of the VXLAN BGP EVPN Fabric 107
Verifying VXLANs and Commands on Switches 107
Verifying Resources 111
Verifying Networks 112
Configuration Profiles Support for Brownfield Migration 115
Migrating a Bottom-Up VXLAN Fabric to DCNM 115
Resolving Config Compliance Error on Switches with Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) Images 124
Modifying VLAN Names in a Switch with Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) Images 128
Changing a Brownfield Imported BIDIR Configuration 130
Manually Adding PIM-BIDIR Configuration for Leaf or Spine Post Brownfield Migration 131
Migrating an MSD Fabric with Border Gateway Switches 131

CHAPTER 6 **Layer 4-Layer 7 Service 135**
Layer 4-Layer 7 Service 135
Guidelines and Limitations for Layer 4-Layer 7 Service 136
Configuring Fabric Settings for Layer 4-Layer 7 Service 137
Configuring Layer 4-Layer 7 Service 139
Create Service Node 139
Create Route Peering 142
Create Service Policy 149
Templates 151
Adding a Route Peering 152
Adding a Service Policy 154
Deleting a Service Node 155
Editing a Service Node 155
Preview a Service Policy or a Route Peering 156
Deploying a Service Policy or a Route Peering 157
Exporting a Service Policy or a Route Peering Table 158

Importing a Service Policy or a Route Peering Table 158
Deleting a Service Policy 158
Deleting a Route Peering 159
Viewing Service Policy Information 160
Viewing Route Peering Information 162

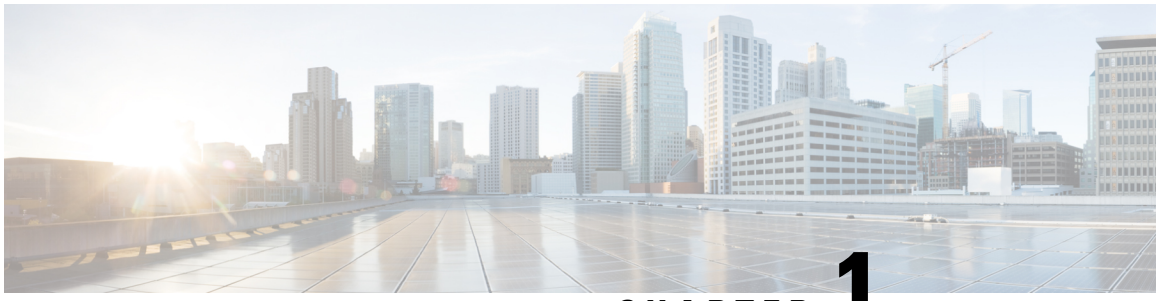

CHAPTER 1

Configuring a Fabric with eBGP Underlay

This chapter describes how to configure a fabric with eBGP-based underlay.

- Creating a New Fabric for EBGP-Based Underlay, on page 1
- Overview of Networks in a Routed Fabric, on page 19

Creating a New Fabric for EBGP-Based Underlay

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

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

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

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

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

   The fields are explained:

   Fabric Name - Enter the name of the fabric.

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

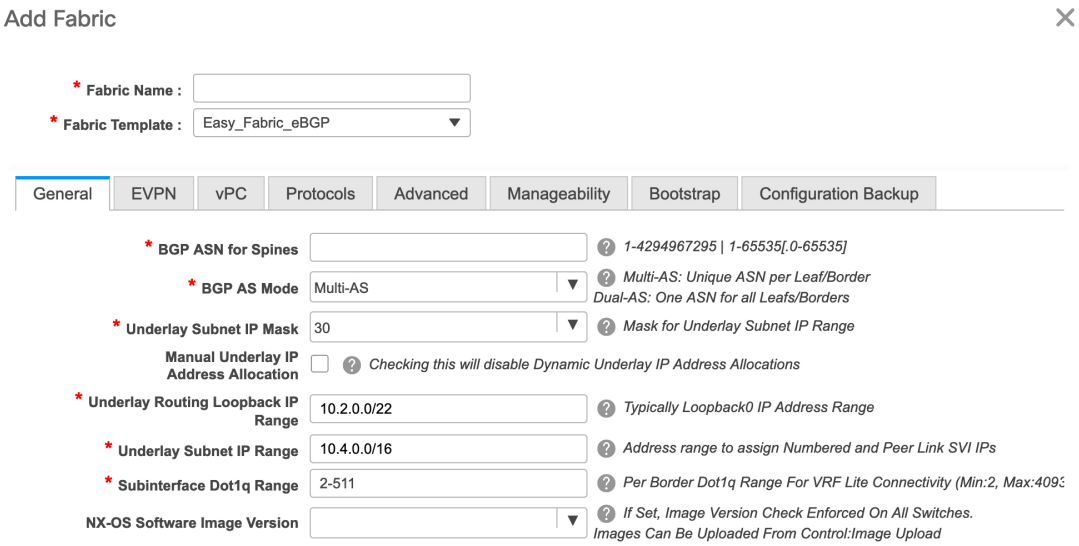

3. The **General** tab is displayed by default. The fields in this tab are:

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

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

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

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

   The fabric is identified by the spine switch AS number.

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

   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:

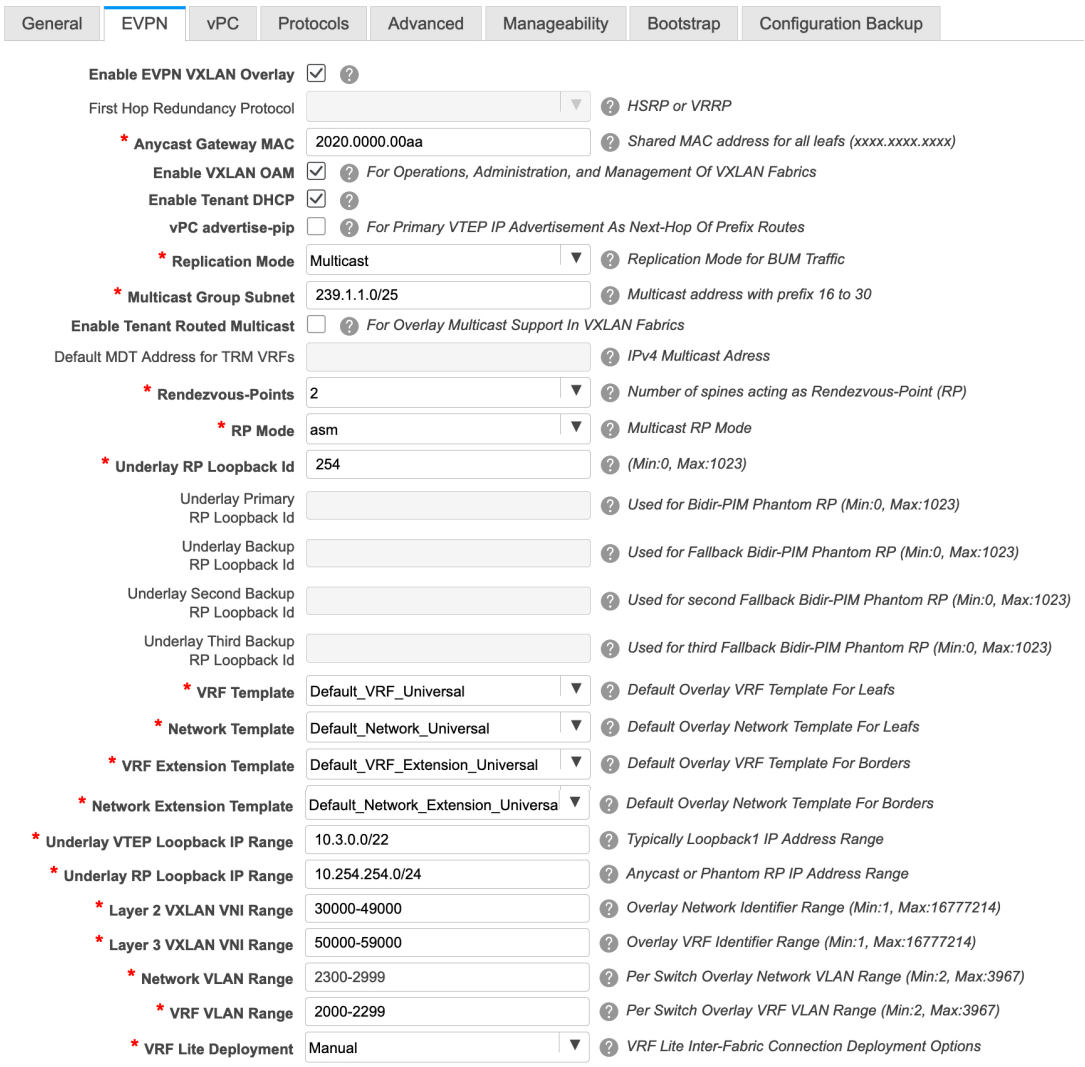

Enable EVPN VXLAN Overlay: Enables the VXLAN overlay provisioning for the fabric.

You can convert a routed fabric to a VXLAN enabled fabric by selecting this option. When the fabric is VXLAN enabled, you can create and deploy overlay networks or VRFs. The procedure for creating and deploying networks or VRFs is the same as in Easy_Fabric_11_1.

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

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

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

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.

First Hop Redundancy Protocol: Specifies the FHRP protocol. Choose either **hsrp** or **vrrp**.

Note After a network has been created, you cannot change this fabric setting. You should delete all networks, and then change the FHRP setting.

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

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

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

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

Enable Tenant DHCP: Enables tenant DHCP support.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

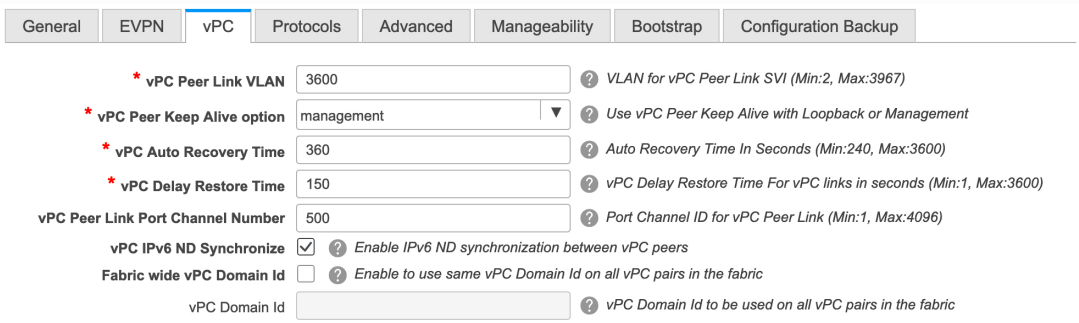

vPC Peer Link VLAN: VLAN used for the vPC peer link SVI.

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

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

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

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

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

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.

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

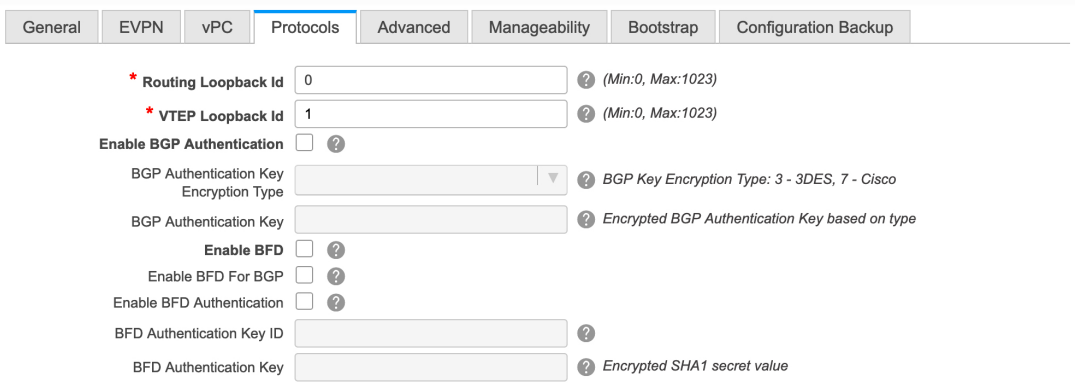

Routing Loopback Id - The loopback interface ID is populated as 0 by default. It is used as the BGP router ID.

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

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

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

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

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

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 following 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:

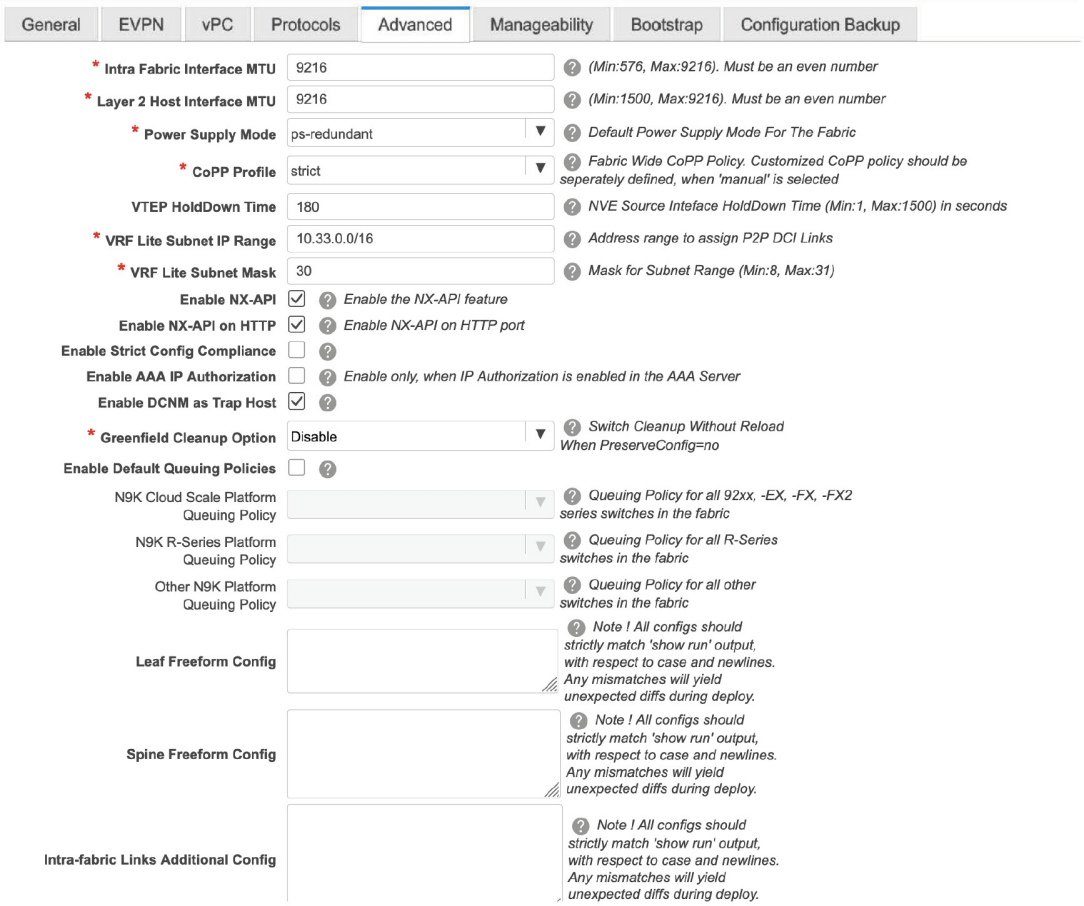

Intra Fabric Interface MTU - Specifies the MTU for the intra fabric interface. This value should be an even number.

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 NX-API - Specifies enabling of NX-API.

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

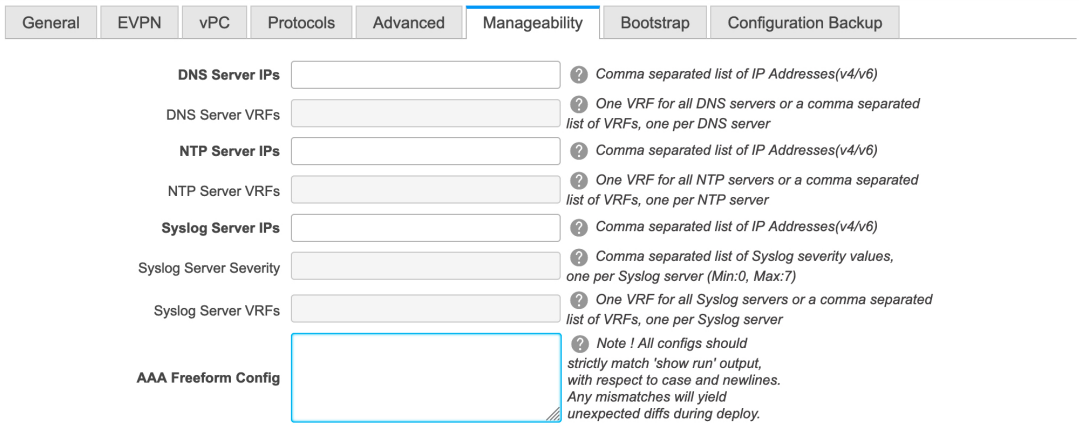

The fields in this tab are:

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.

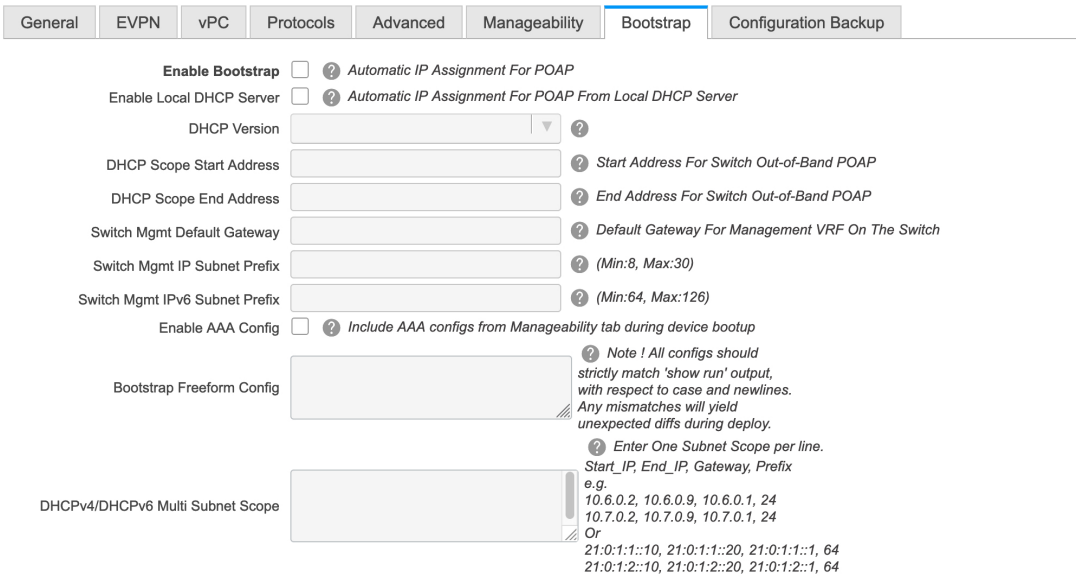

Enable Bootstrap - Select this check box to enable the bootstrap feature.

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 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 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:

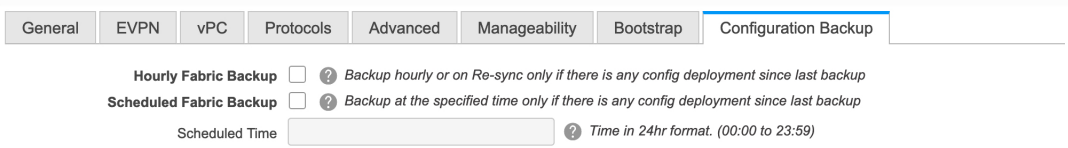

Hourly Fabric Backup: Select the check box to enable an hourly backup of fabric configurations and the intent.

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.

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

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

The **Actions** panel at the left part of the screen allows you to perform various functions. One of them is the **Add switches** option to add switches to the fabric. After you create a fabric, you should add fabric devices. The 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 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 *Restoring Fabrics* 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. 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 part of the screen. The current fabric is highlighted. An MSD and its member fabrics are distinctly displayed, wherein the member fabrics are indented, under the MSD fabric.

VXLAN Fabric With eBGP Underlay – Pointers

- 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.
- When you convert a non-VXLAN (or routed fabric) to a VXLAN enabled fabric, you can create and deploy overlay networks and VRFs.

Applying Policies On A Fabric With An eBGP Underlay

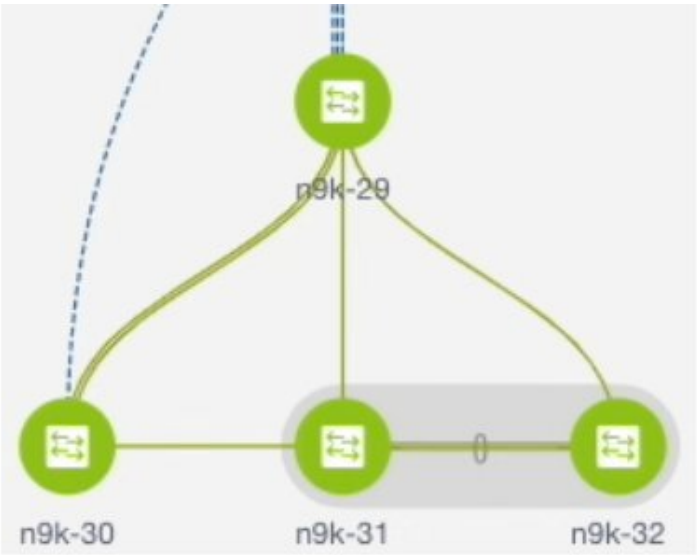

The topology shows a VXLAN fabric enabled with eBGP for the underlay. In DCNM, a fabric with the Easy_Fabric_eBGP template is created. One spine switch (n9k-29) and three leaf switches (n9k-30, and vPC switch pair n9k-31 and n9k-32) are imported to it.

This topic covers the following:

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

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

Deploying Fabric Underlay Policies

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

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

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

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

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

7. Close the screen.
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 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.

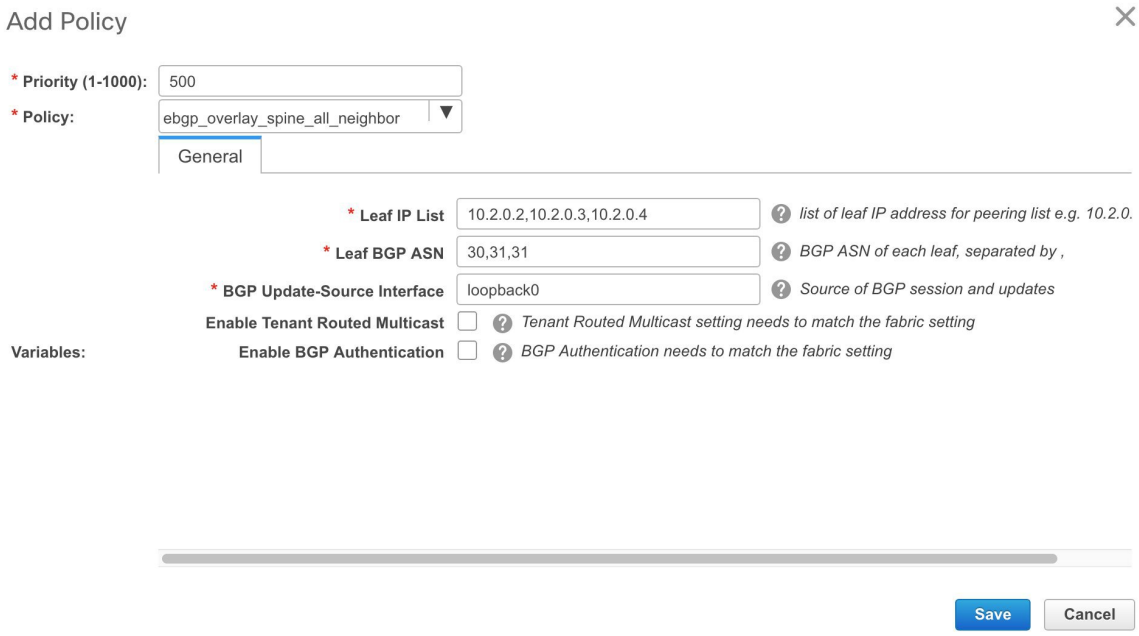

The fields on the screen are:

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 for this field.

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.

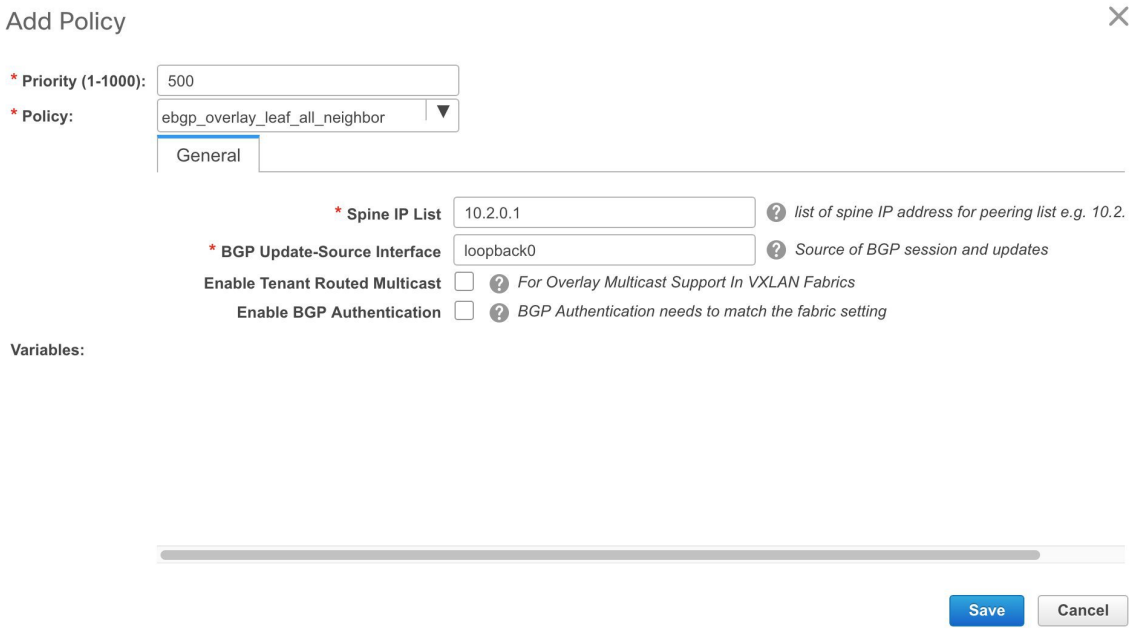

The fields on the screen are:

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 for this field.

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.

Dual-AS Fabric Deployment

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

- Deploy the spine overlay policy as explained in the Multi-AS fabric section.
- Deploy the leaf overlay and underlay policies on all leaf switches at once, since they have a common AS number.

Additional Pointers

- Brownfield migration is not supported for eBGP fabric.

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

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.

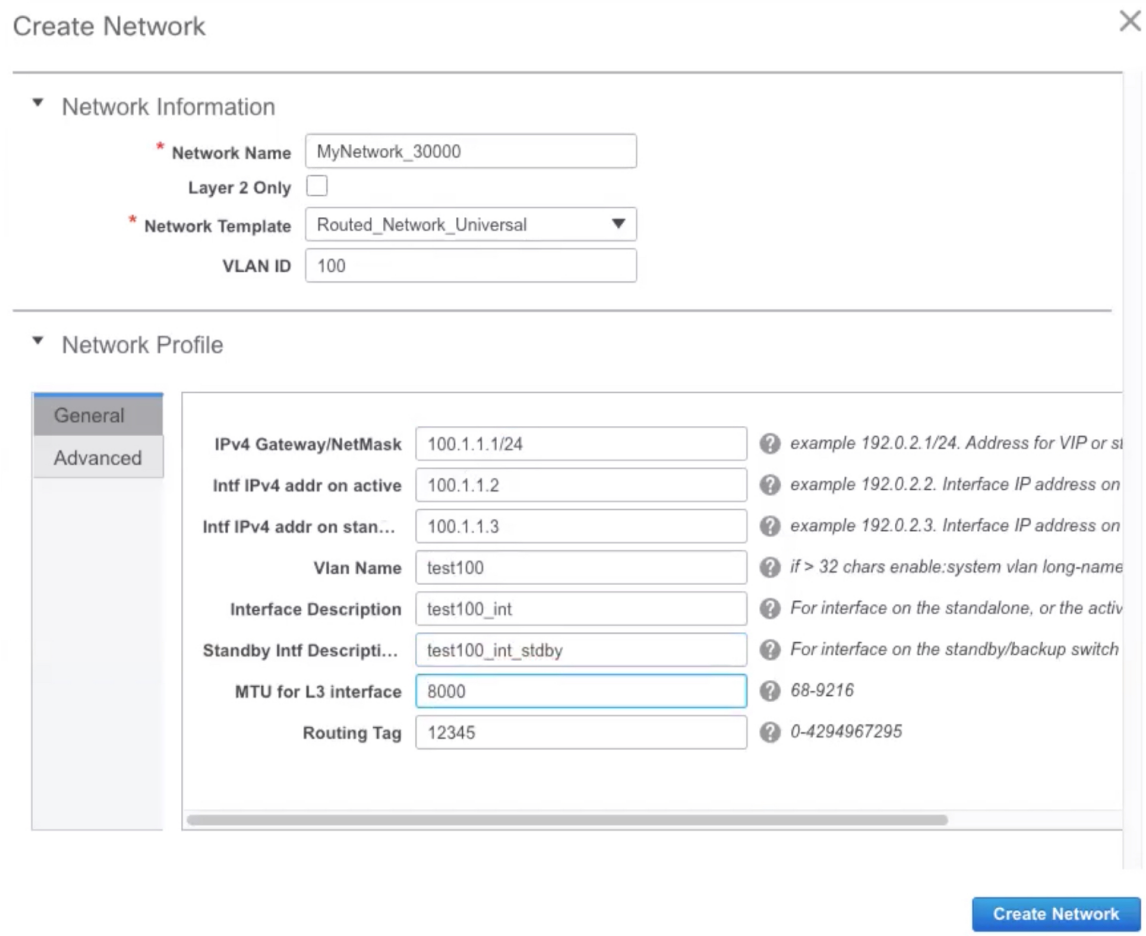

Network Name: Specifies the name of the network. The network name should not contain any white spaces or special characters except underscore (_) and hyphen (-).

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/master 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.

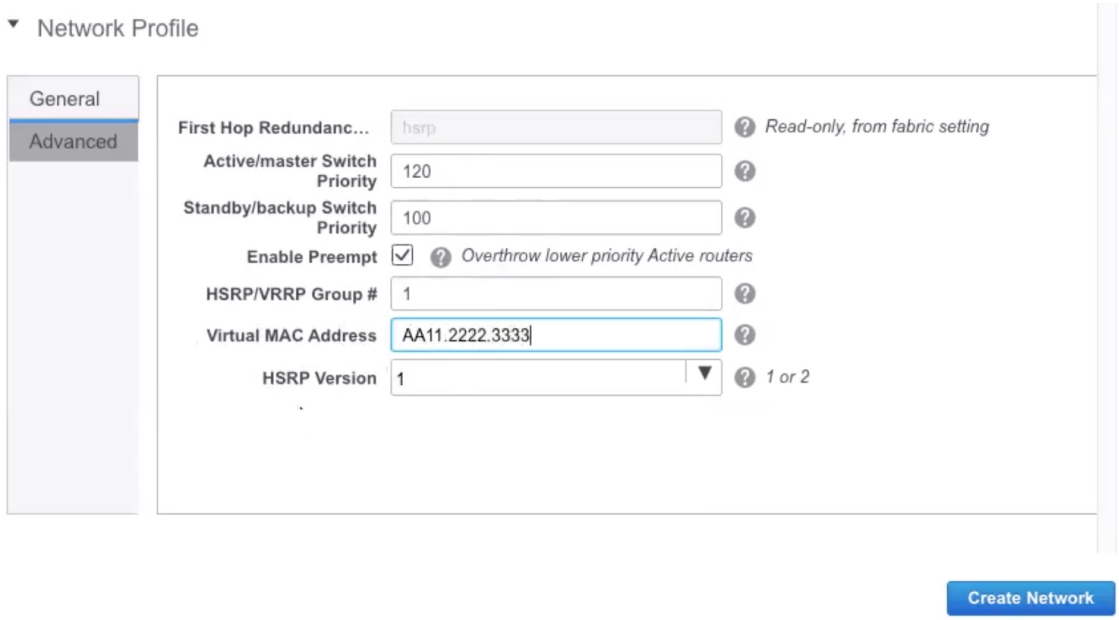

First Hop Redundancy Protocol: A read-only field that specifies FHRP selected in the fabric settings.

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 a active or master 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**.

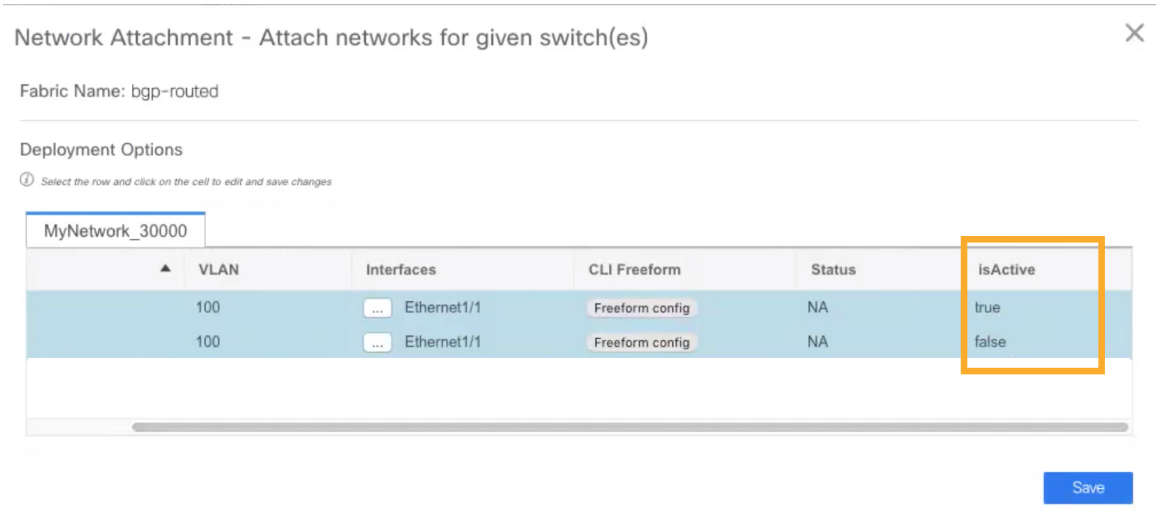

Note In a routed fabric, when you edit a deployed network and save without making any changes, the status of the network changes to **Pending**. Similarly, if a **Network Attachment** window is opened for a deployed network, and saved without any changes, the status of the network changes to **Pending**. 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.

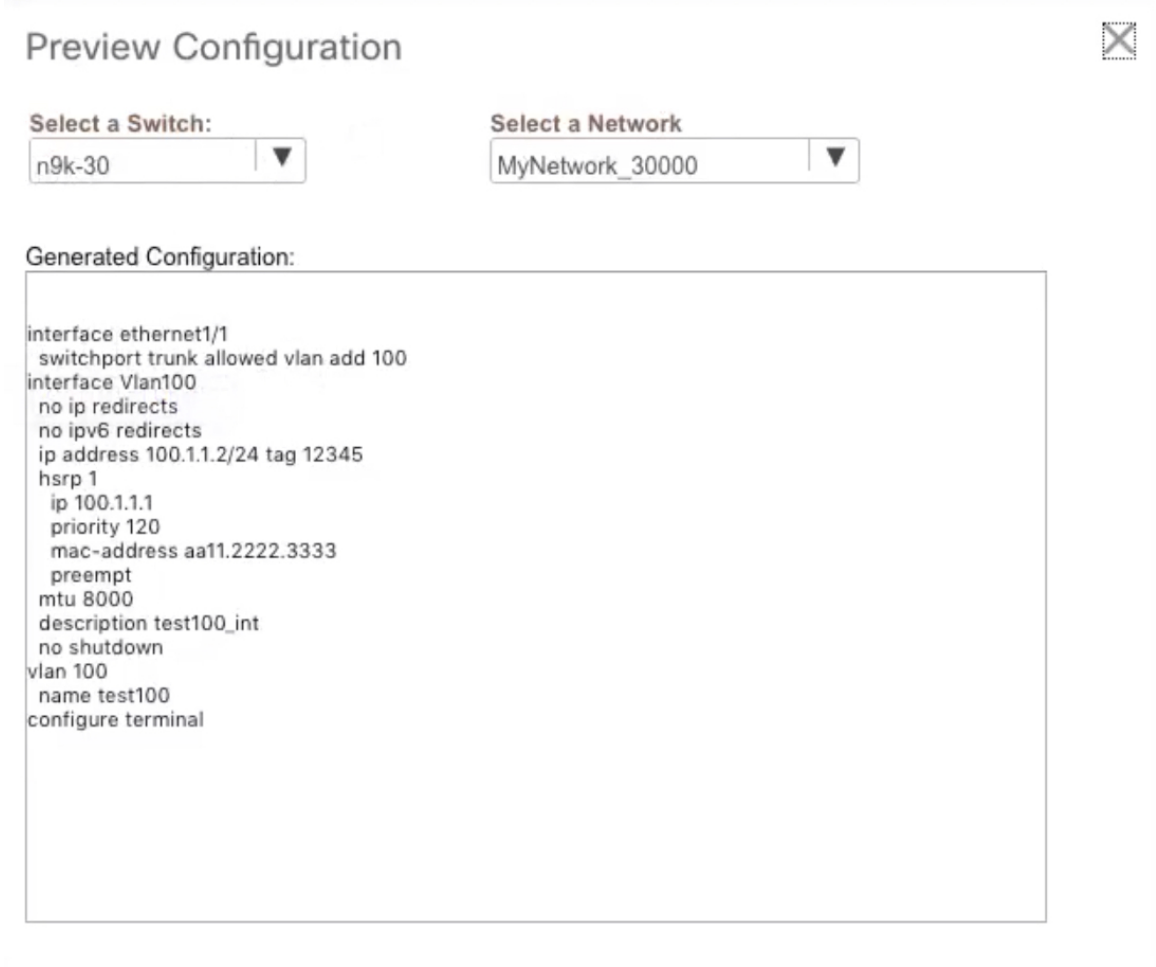

Step 9 Click the **Deploy** button in the **Network / VRF Deployment** window.

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.

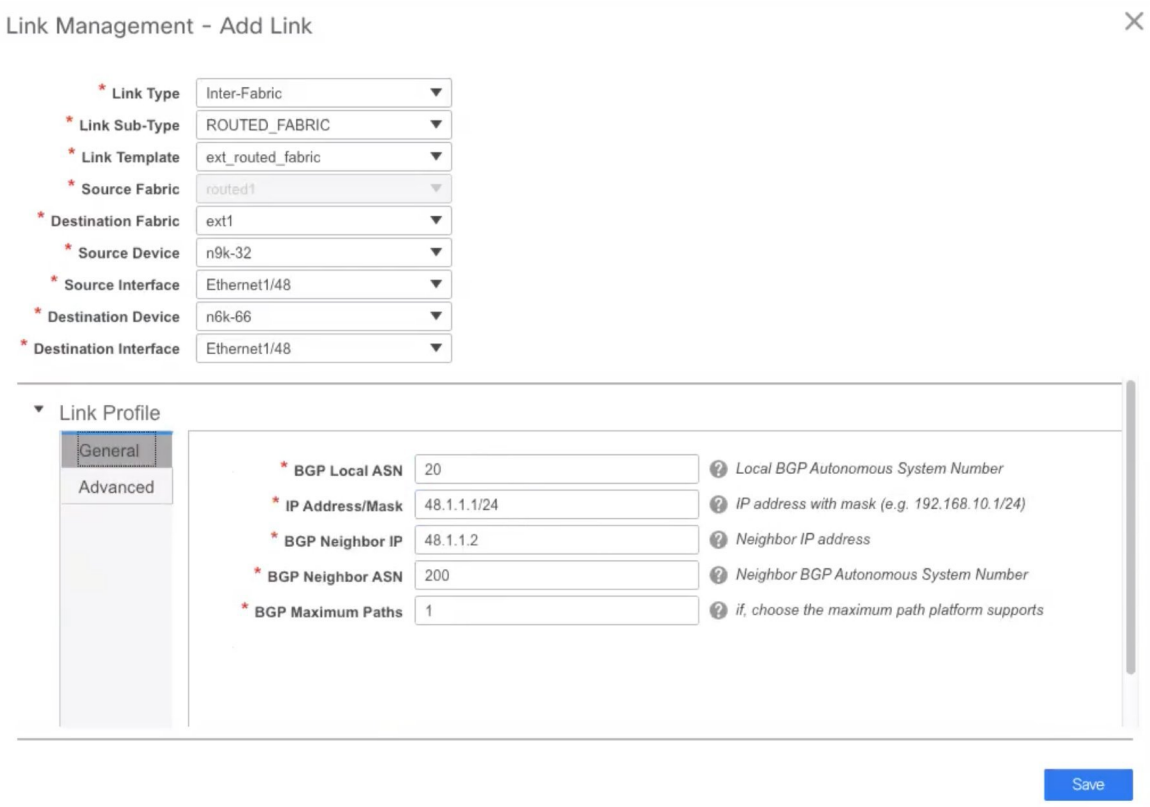

Link Type – Choose **Inter-Fabric** to create an inter-fabric connection between two fabrics, via their border switches or edge routers.

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.

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**.

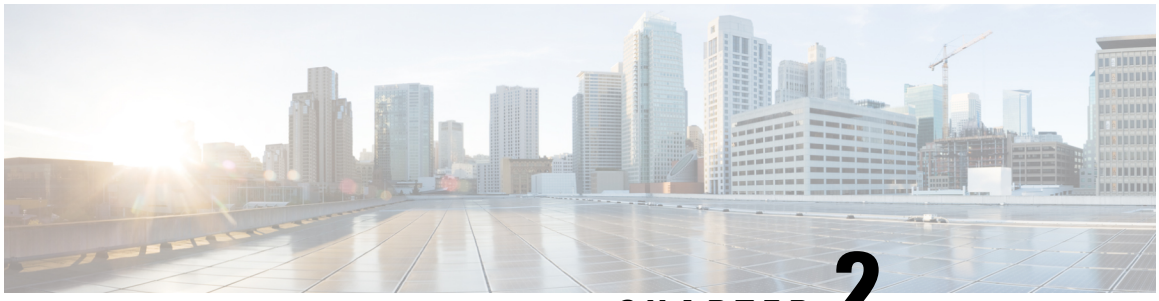

CHAPTER 2

Configuring a VXLANv6 Fabric

This chapter describes how to configure a VXLAN fabric with IPv6 underlay.

- Overview, on page 27
- Creating a VXLAN Fabric with IPv6 Underlay, on page 28

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.

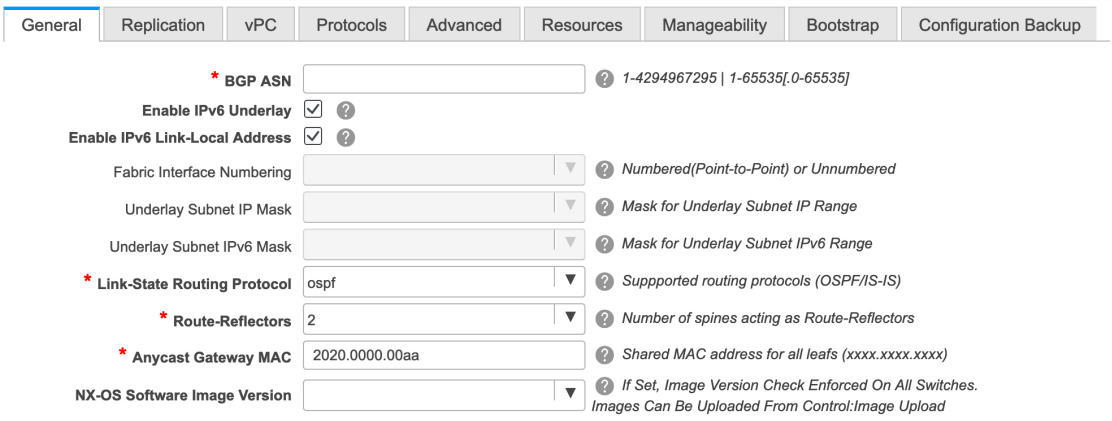

BGP ASN: Enter the BGP AS number for the fabric. You can enter either the 2 byte BGP ASN or 4 byte BGP ASN.

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.

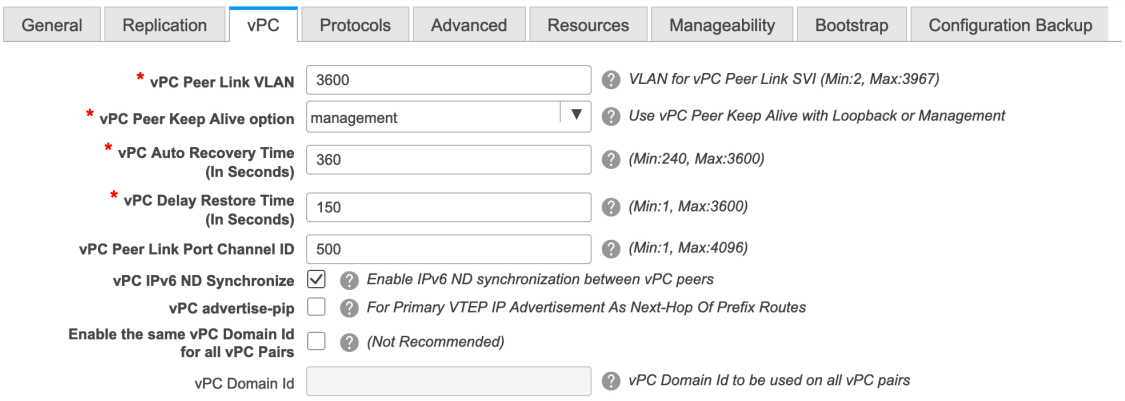

vPC Peer Keep Alive option – Choose the management or loopback option. If you want to use IP addresses assigned to the management port and the management VRF, choose management. If you use IP addresses assigned to loopback interfaces (and a non-management VRF), choose loopback. Both the options are supported for IPv6 underlay.

Step 6 Click the **Protocols** tab.

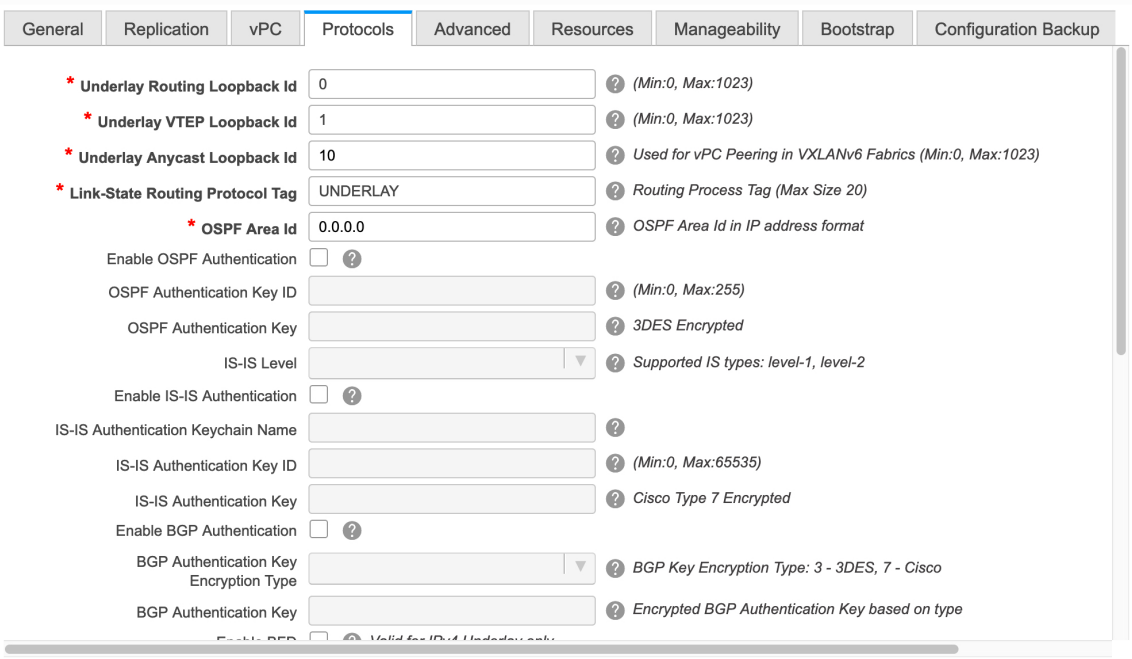

Underlay Anycast Loopback Id: Specifies the underlay anycast loopback ID for IPv6 underlay. Since an IPv6 address cannot be configured as secondary, an additional loopback interface is allocated on each vPC device. Its IPv6 address will be used as the VIP.

Step 7 Click the **Resources** tab.

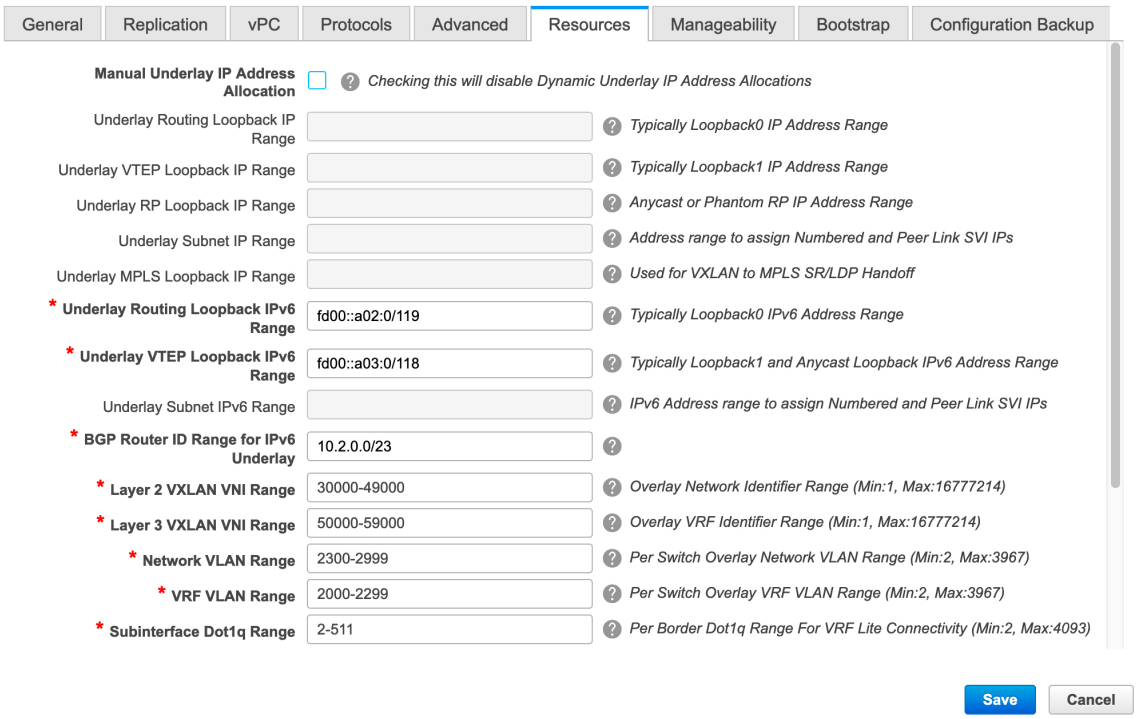

Manual Underlay IP Address Allocation: Select this check box to manually allocate underlay IP addresses. The dynamic underlay IP addresses fields are disabled.

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.

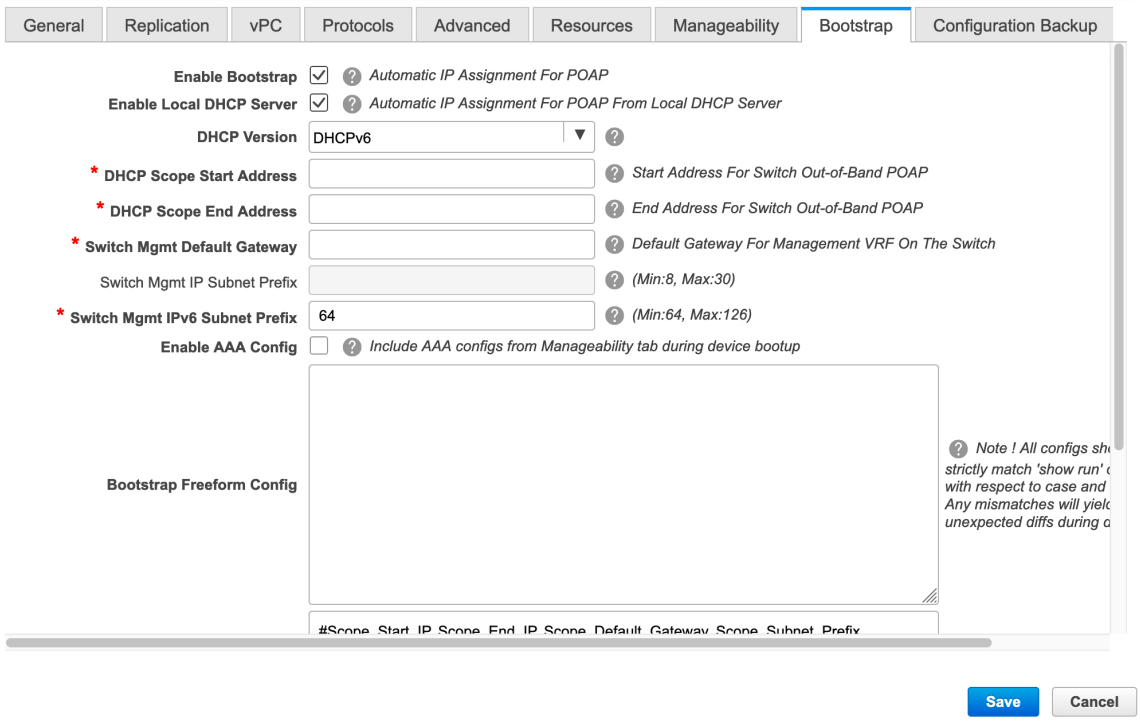

DHCP Version – Select DHCPv4 or DHCPv6 from this drop-down list. When you select DHCPv4, the **Switch Mgmt IPv6 Subnet Prefix** field is disabled. If you select DHCPv6, the **Switch Mgmt IP Subnet Prefix** is disabled.

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

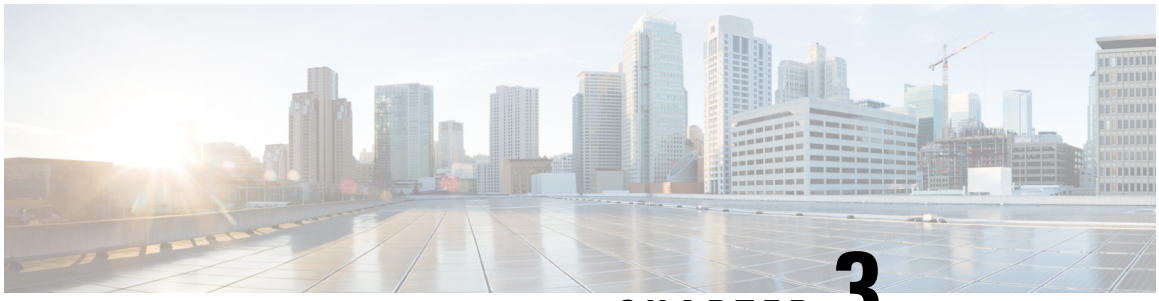

CHAPTER 3

Configuring ToR Switches and Deploying Networks

This chapter describes how to configure the Top-of-Rack (ToR) switches and deploy networks in DCNM.

- Overview, on page 33
- Supported Topologies for ToR Switches, on page 33
- Configuring ToR Switches, on page 39
- Deploying Networks on ToR Switches, on page 44

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.

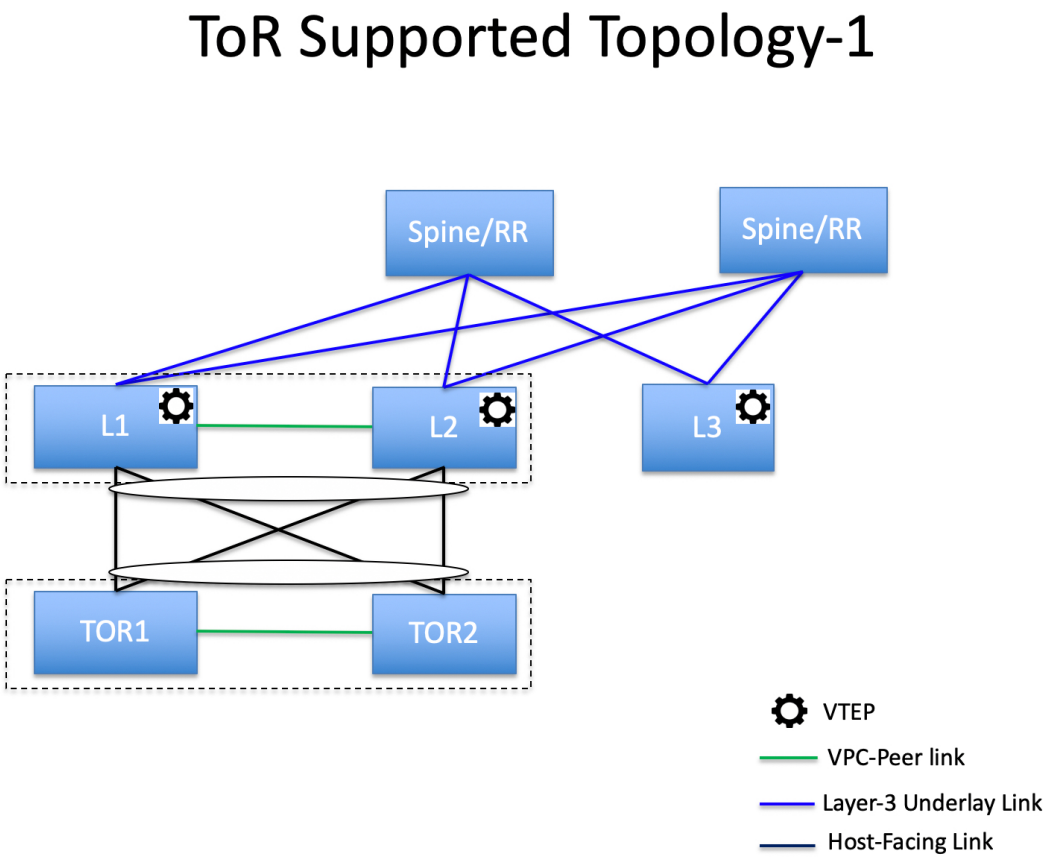

- ToR switches with port channels connected to both the leaf switches. The L1 and L2 switches are connected as a vPC pair.

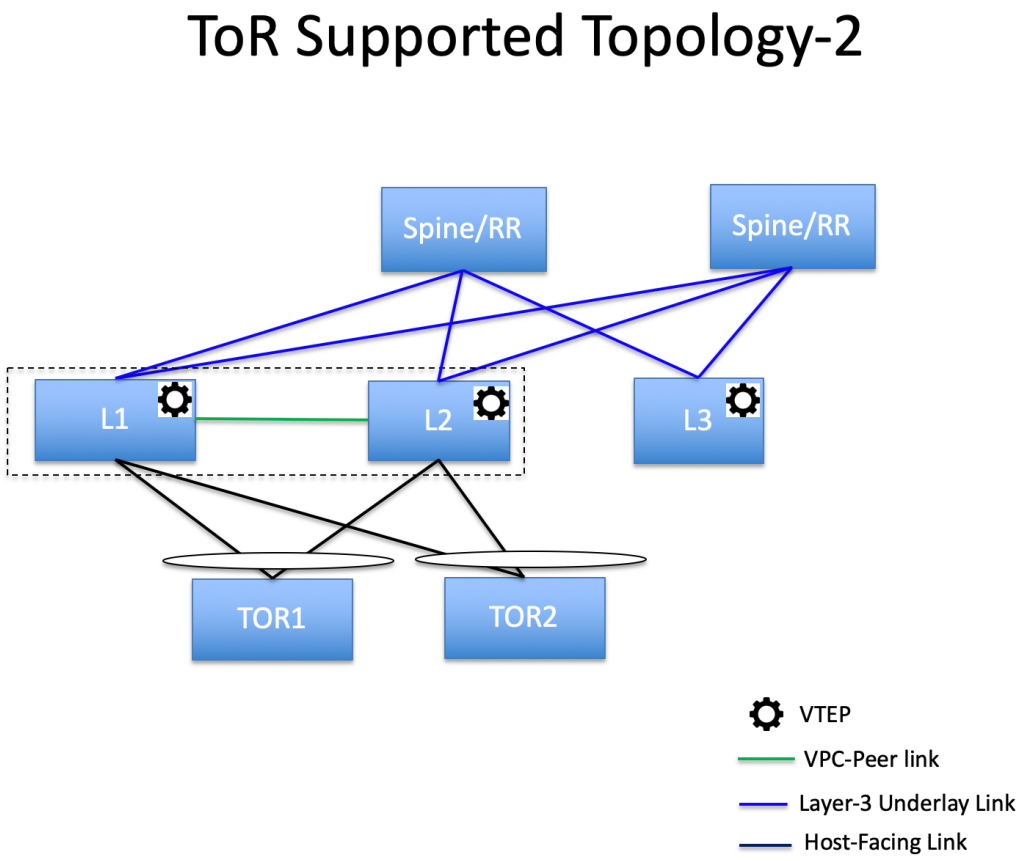

- ToR switches with port channels directly connected to the leaf switches. The L1 and L2 switches are connected as a vPC pair.

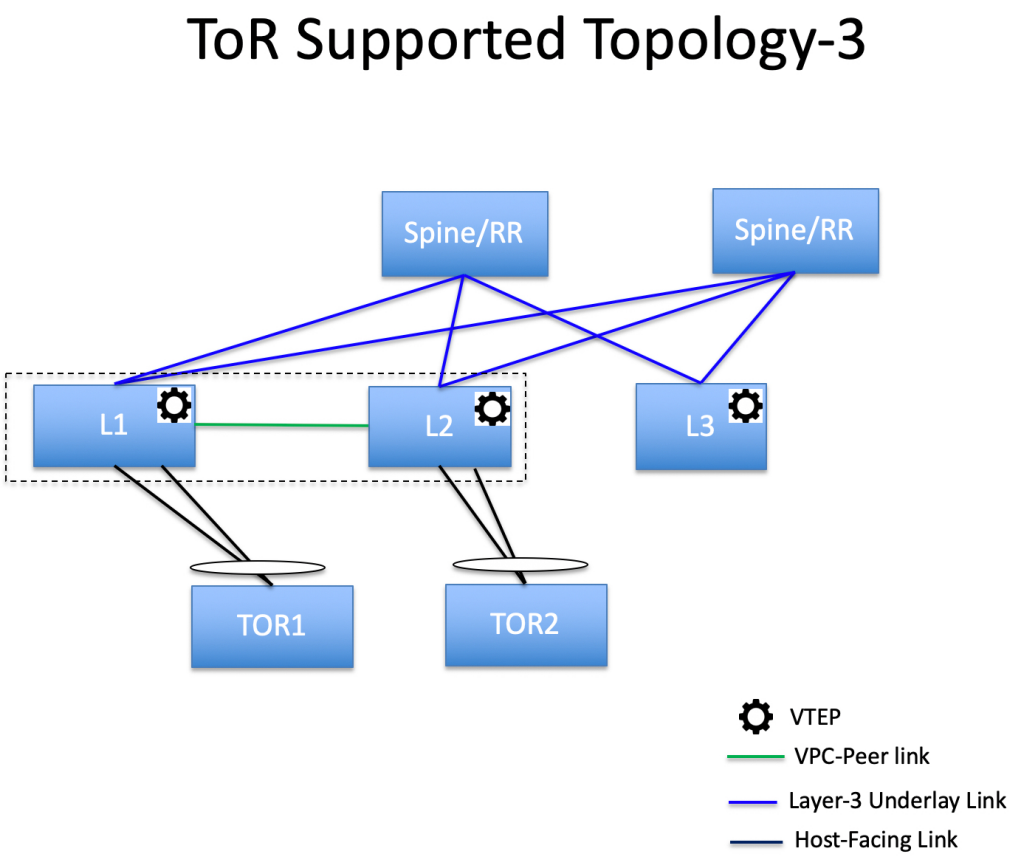

- ToR switches with port channels directly connected to the leaf switches. vPC pairs are not configured for the leaf or ToR switches.

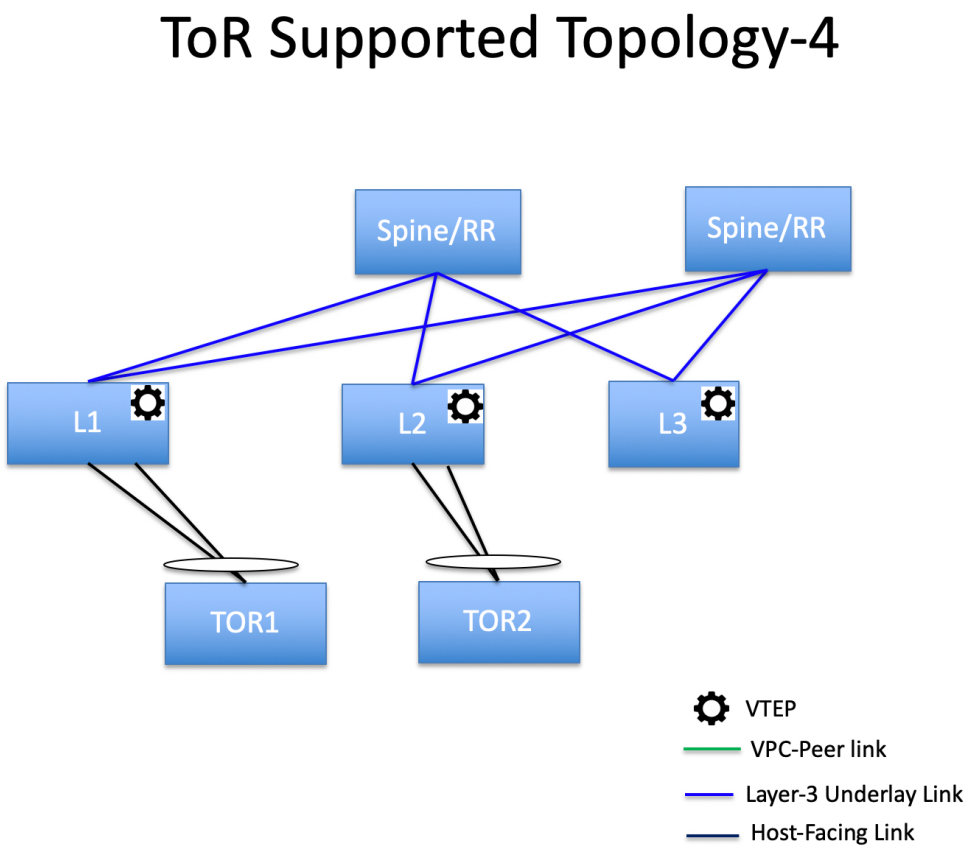

- ToR switches directly connected to the leaf switches. vPC pairs are not configured for the leaf or ToR switches.

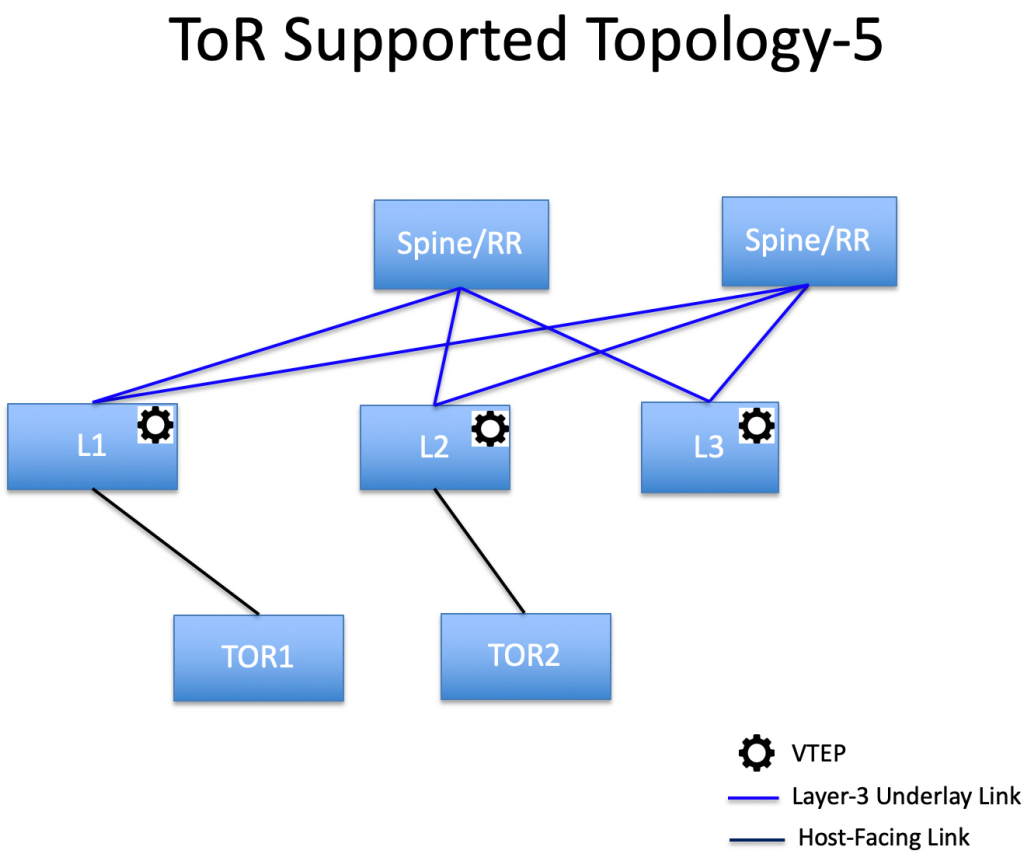

The following topology with ToR switches is not supported in DCNM:

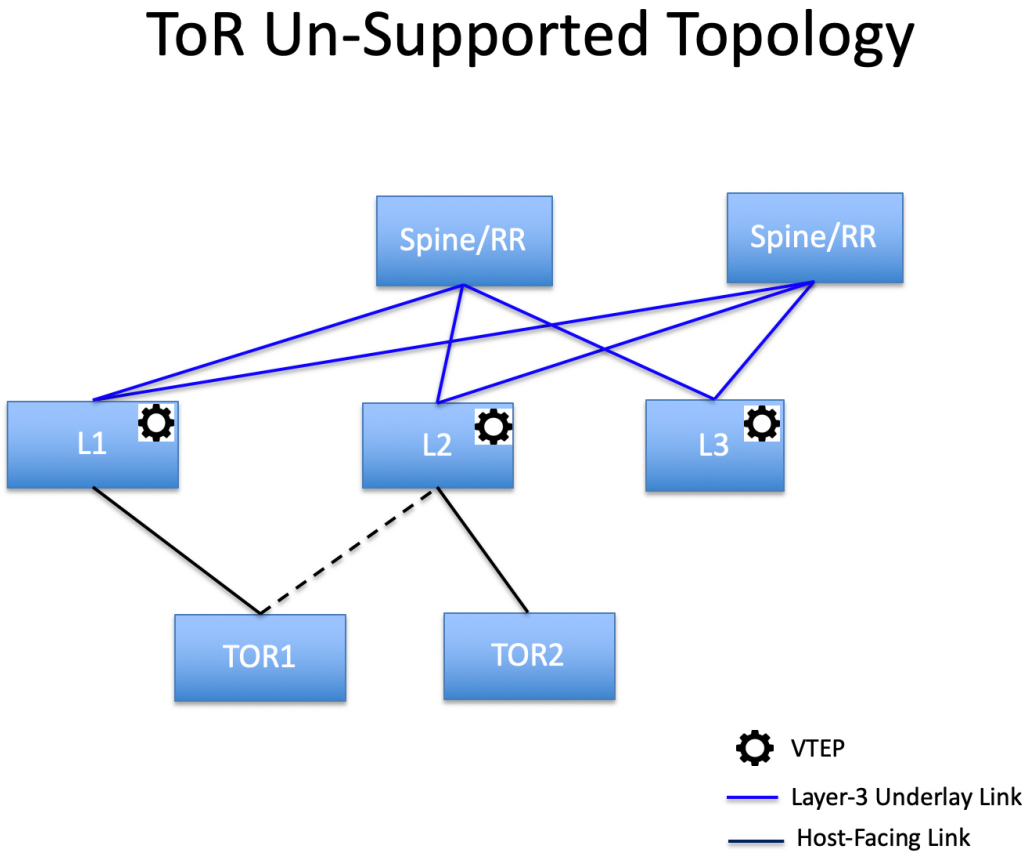

Configuring ToR Switches

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.

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.

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.

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.

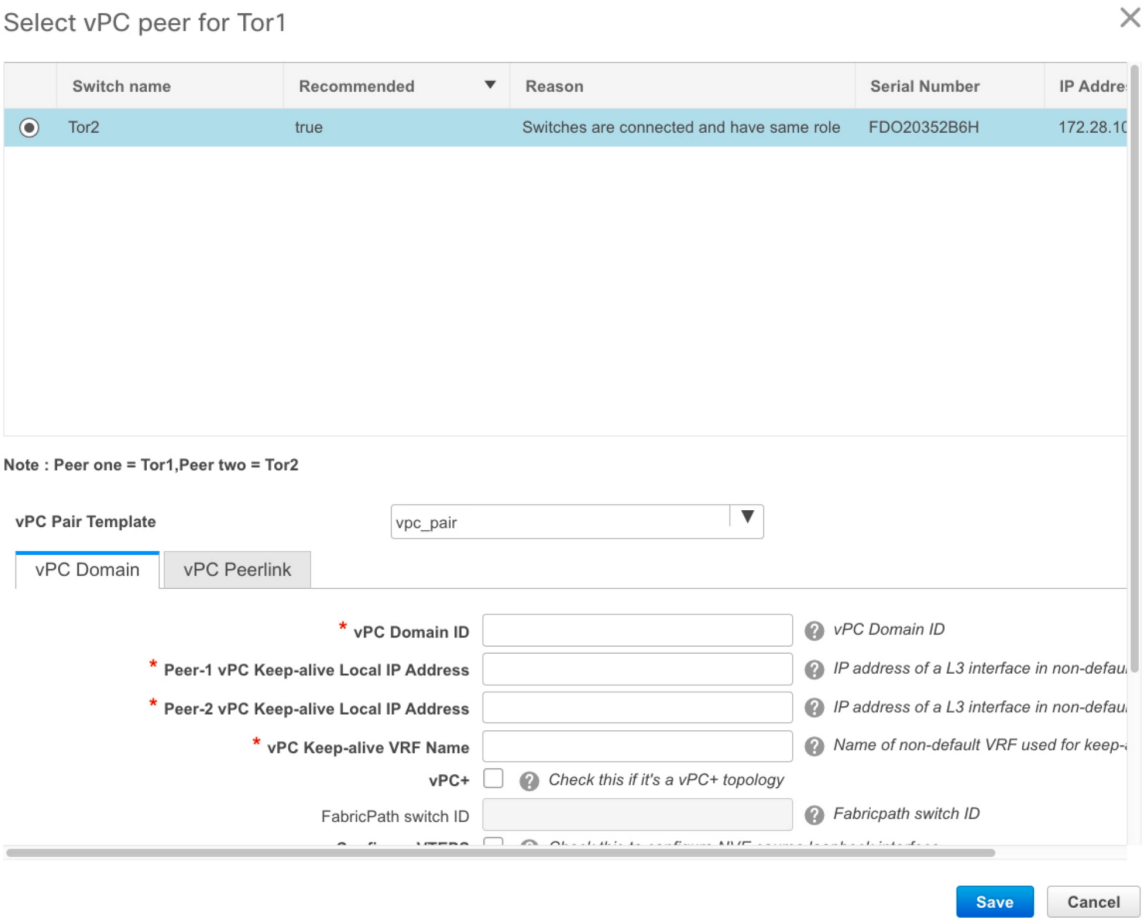

Step 4 Click **Save & Deploy**, and then click **Deploy Config**.

Step 5 After the progress bar shows 100% in the **Config Deployment** window, click **Close**.

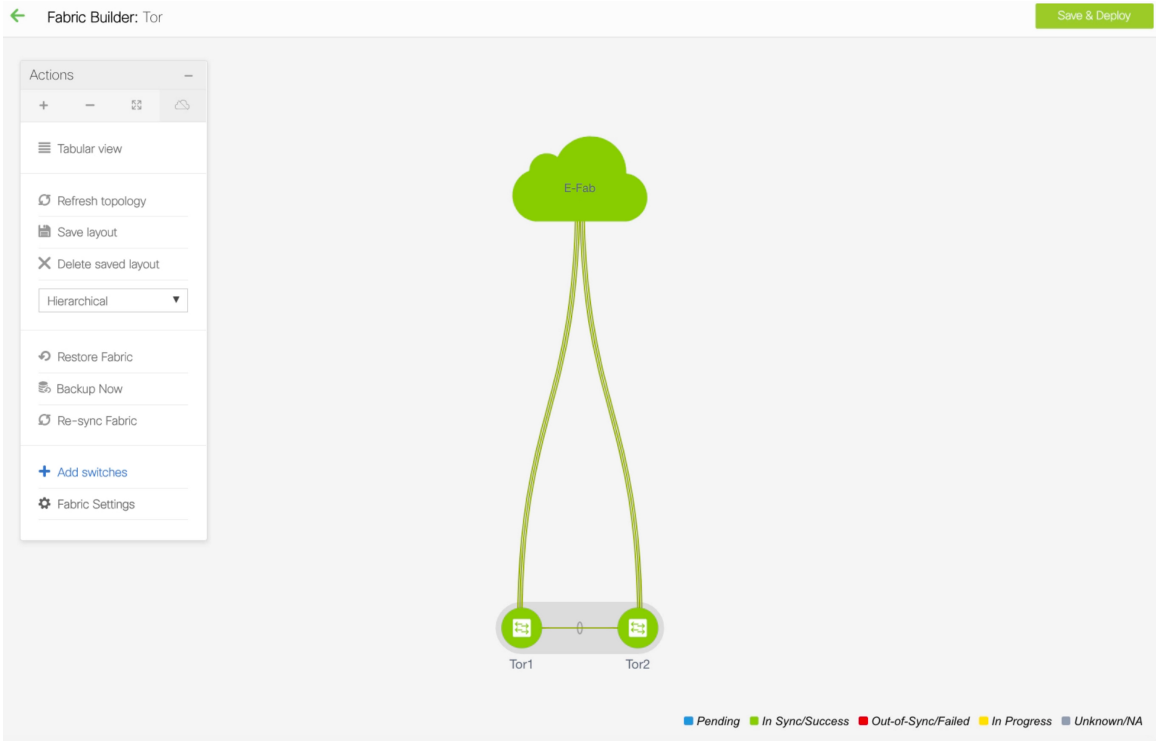

Step 6 Create an MSD fabric.

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.

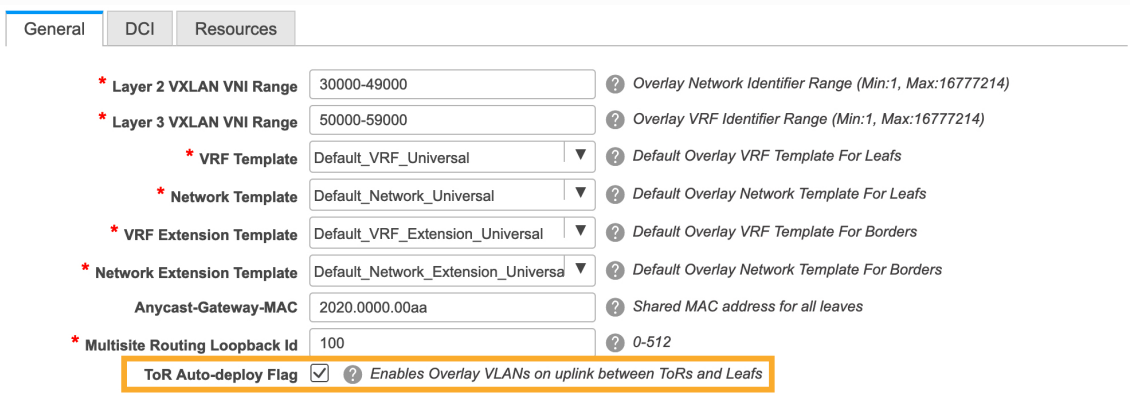

Step 7 Click **Move Fabric** in the **Action** panel. In the **Move Fabric** window, select the Easy Fabric and click **Add**.

Similarly, move the external fabric that contains the ToR switches to the MSD fabric.

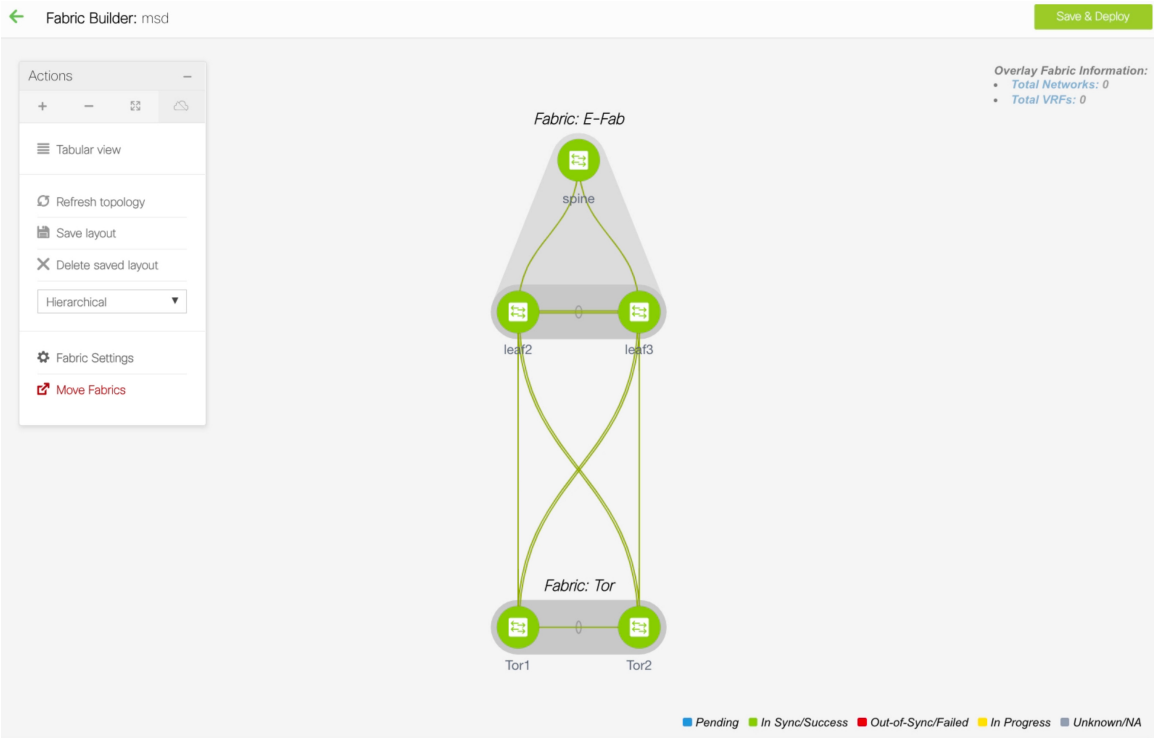

Step 8 Click the **Back** icon and click the Easy fabric containing the leaf switches.

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**.

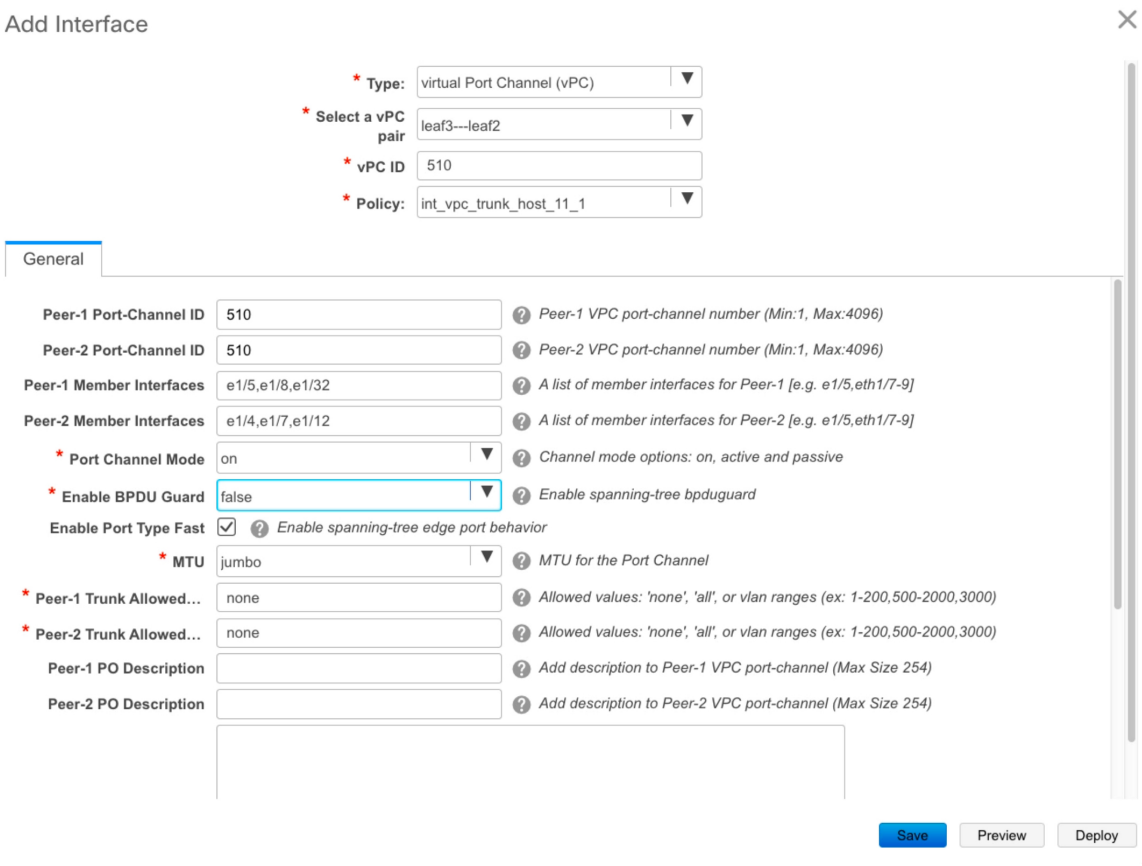

For more information about the fields in this window, see Adding Interfaces.

After saving all the information, click **Deploy**.

Similarly, follow the Steps 9 and 10 to create a vPC in the ToR switch as well.

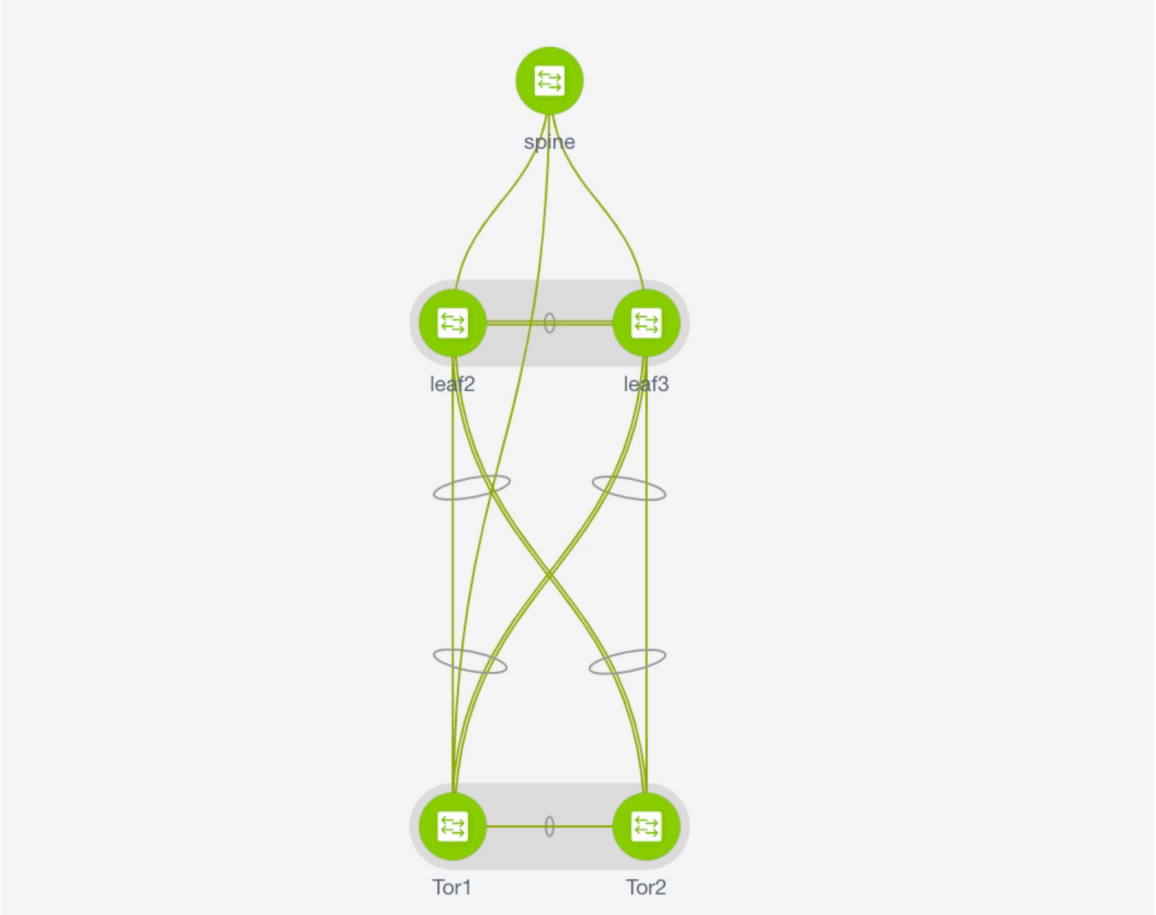

Deploying Networks on ToR Switches

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.

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.

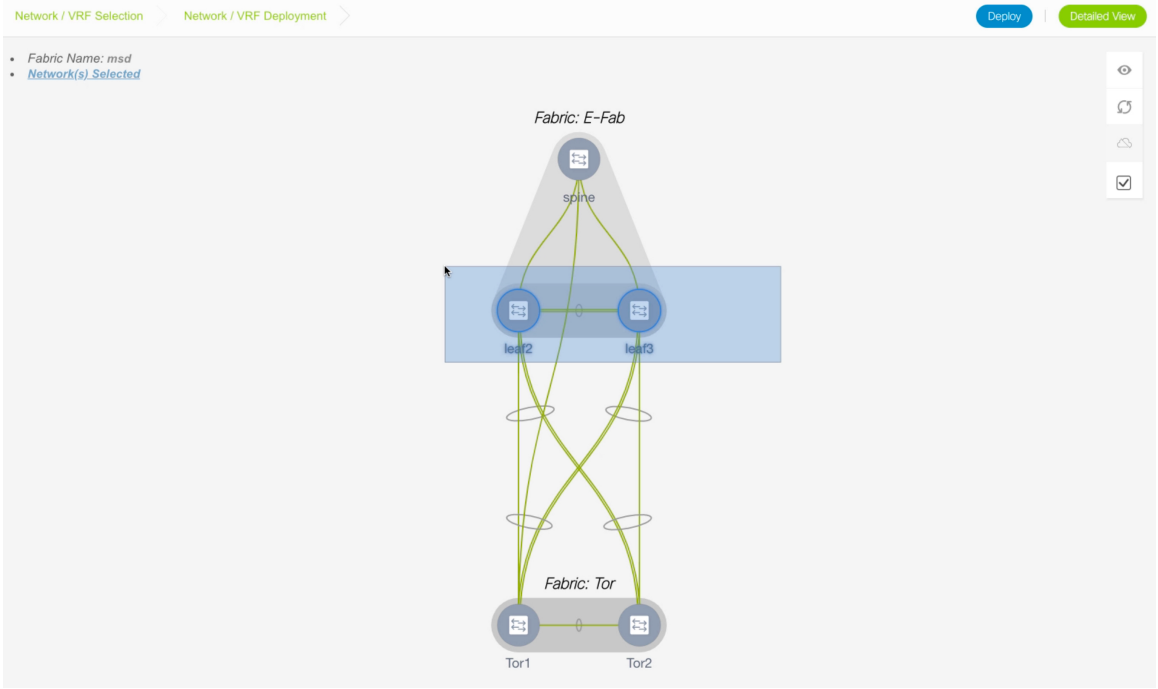

Step 5 In the **Network Attachment** window, click **...** in the **Interfaces** column.

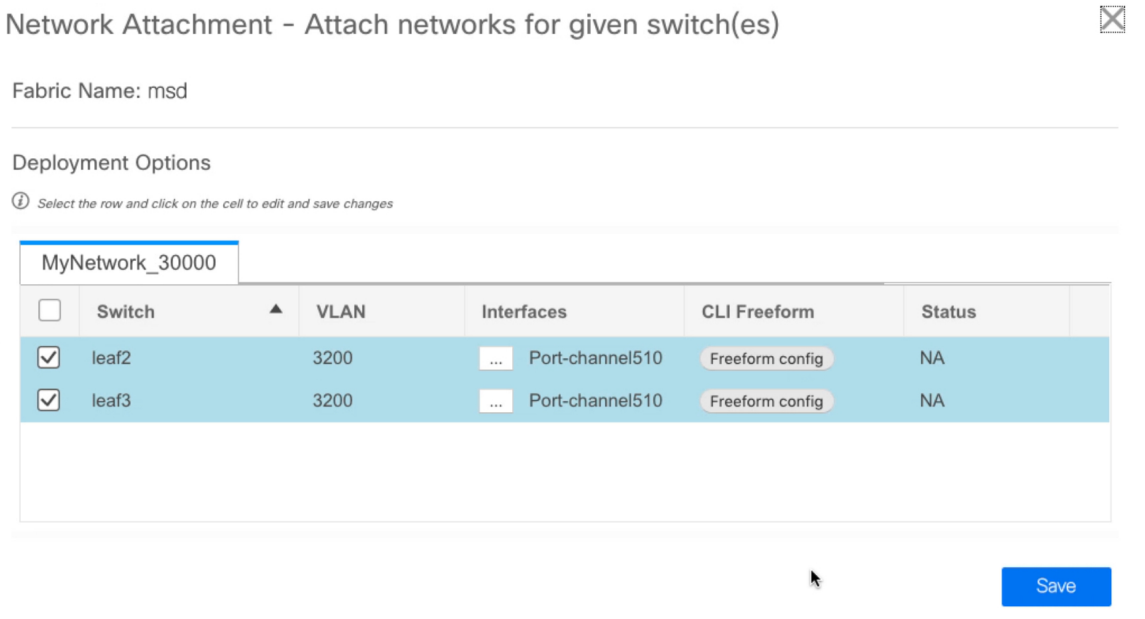

The **Interfaces** window lists interfaces or port channels. You can select interfaces/port channels to associate them with the selected network. These port channels connect the leaf switches to the ToR switches. The networks will be deployed on these port channels.

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.

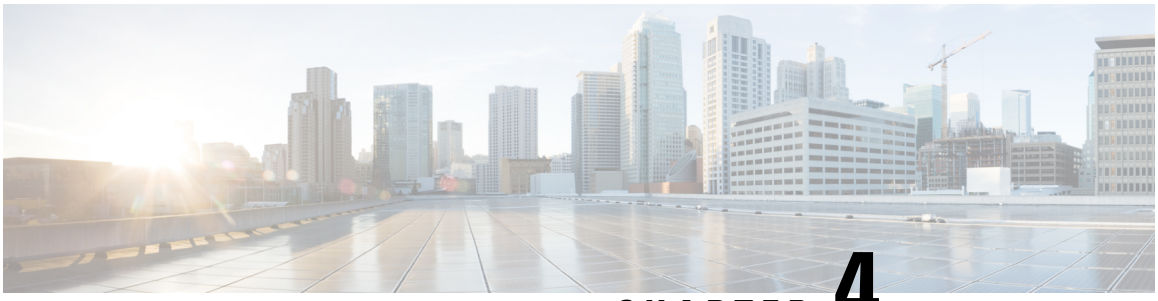

CHAPTER 4

Connecting Cisco Data Center and a Public Cloud

• Connecting Cisco Data Center and a Public Cloud, on page 47

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 1: Topology Overview

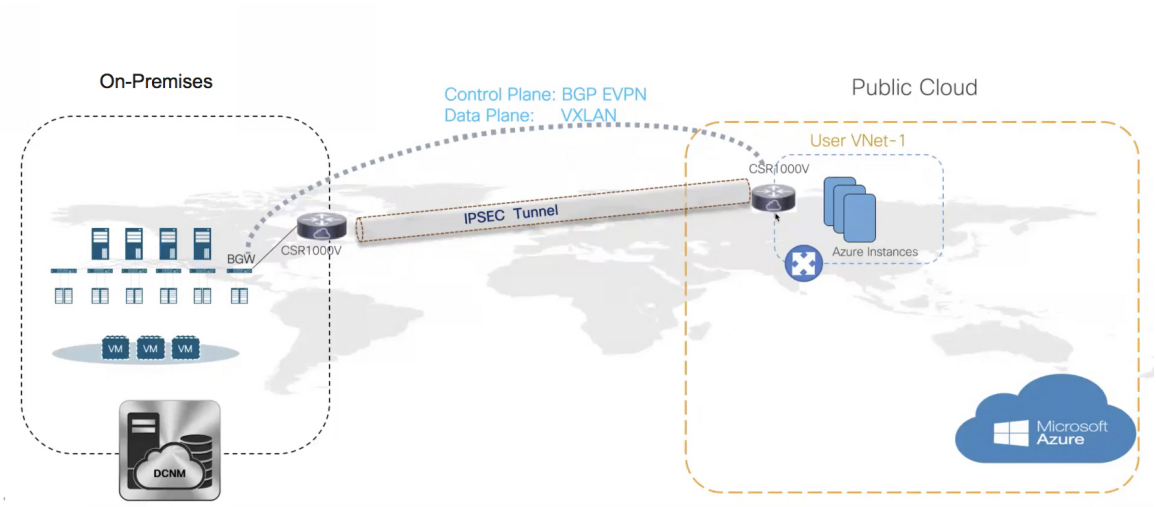

The on-premise data center has the required switches. One of these switches is a border gateway (BGW) that interfaces with an core router for WAN connectivity to the public cloud. The Cisco CSR 1000v is the core router in this use case. You can import this core router into an external fabric in Cisco DCNM. The following figure depicts the sample topology that is employed.

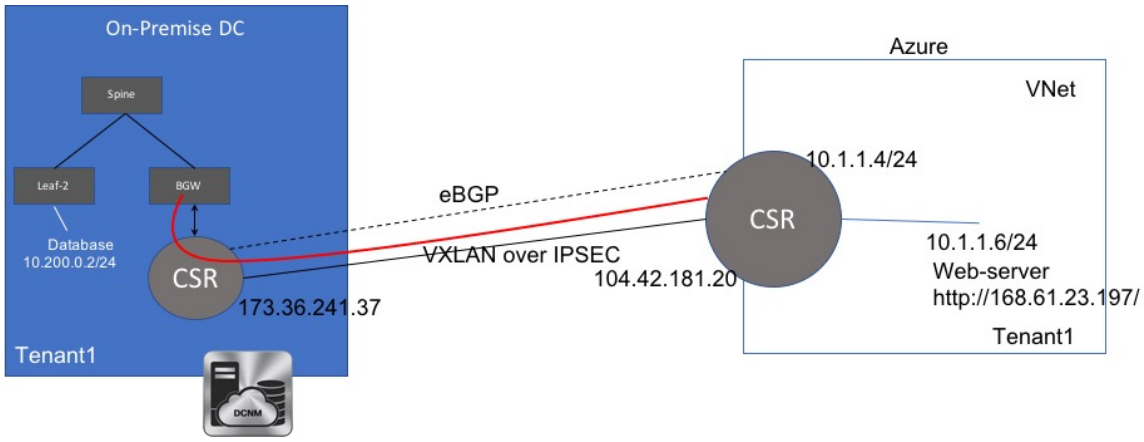

In this example, we list the tasks that are required to provide a layer-3 connectivity between a VM behind standalone leaf and a VM in the Microsoft Azure cloud in a specific user VNET.

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 67 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.

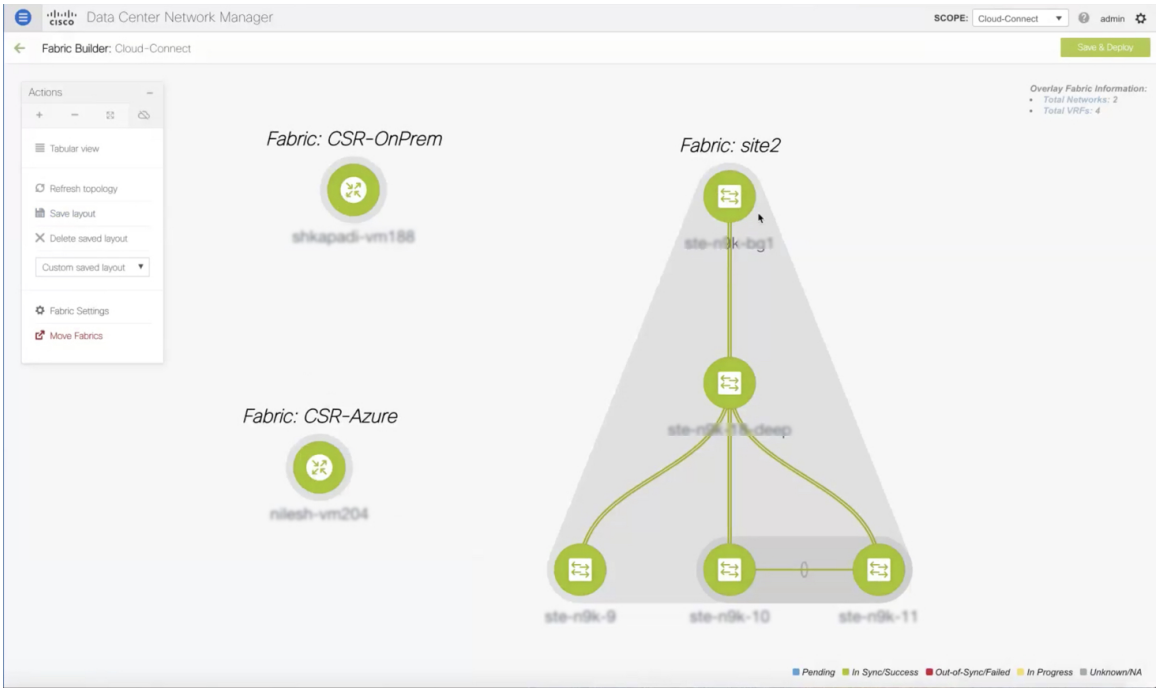

What to do next

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.

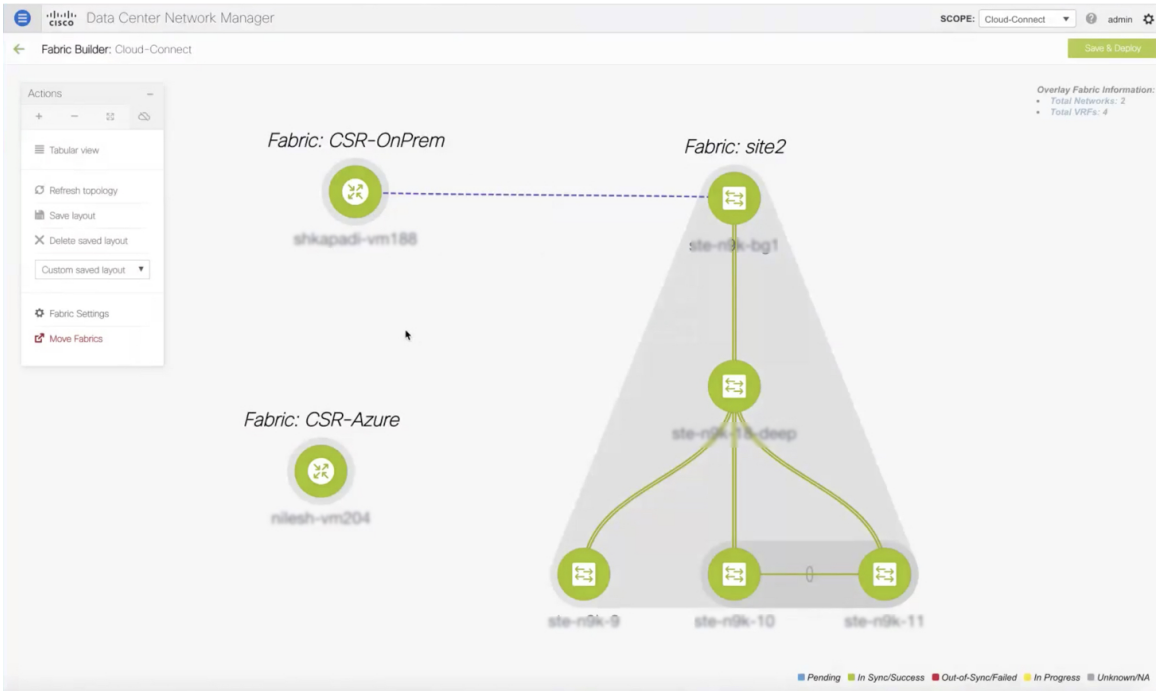

What to do next

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.

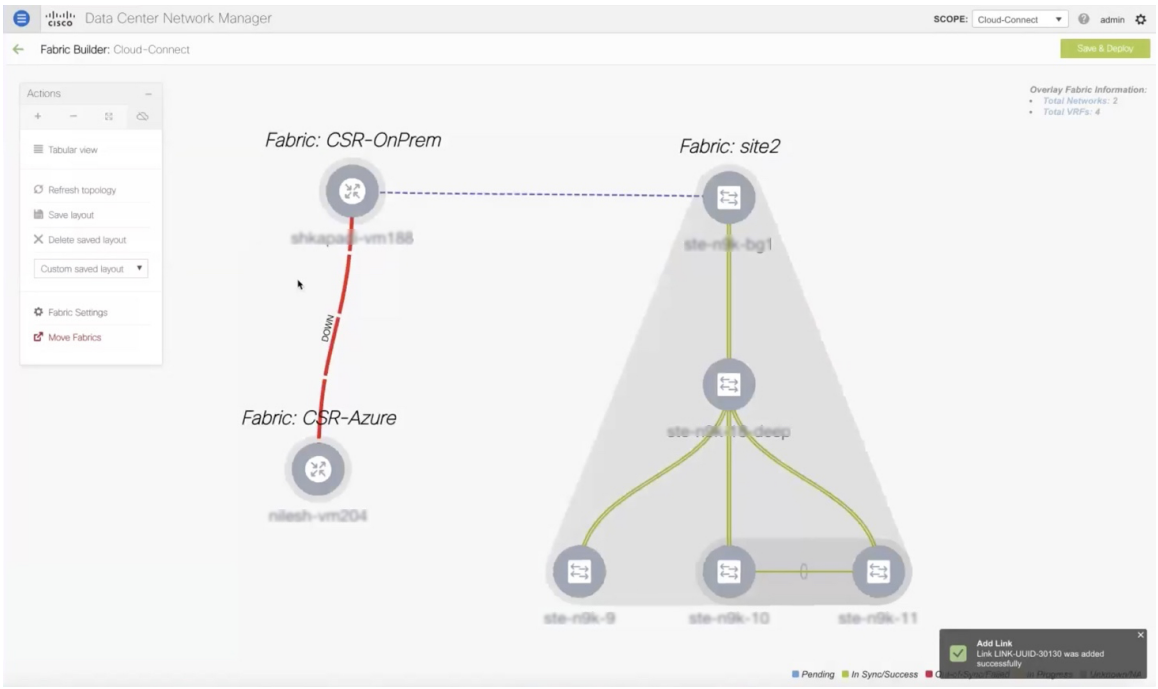

What to do next

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.
Link Template	Choose the **csr_ext_evpn_multisite_overlay_setup** link template from the drop-down list.

Field	Description
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.

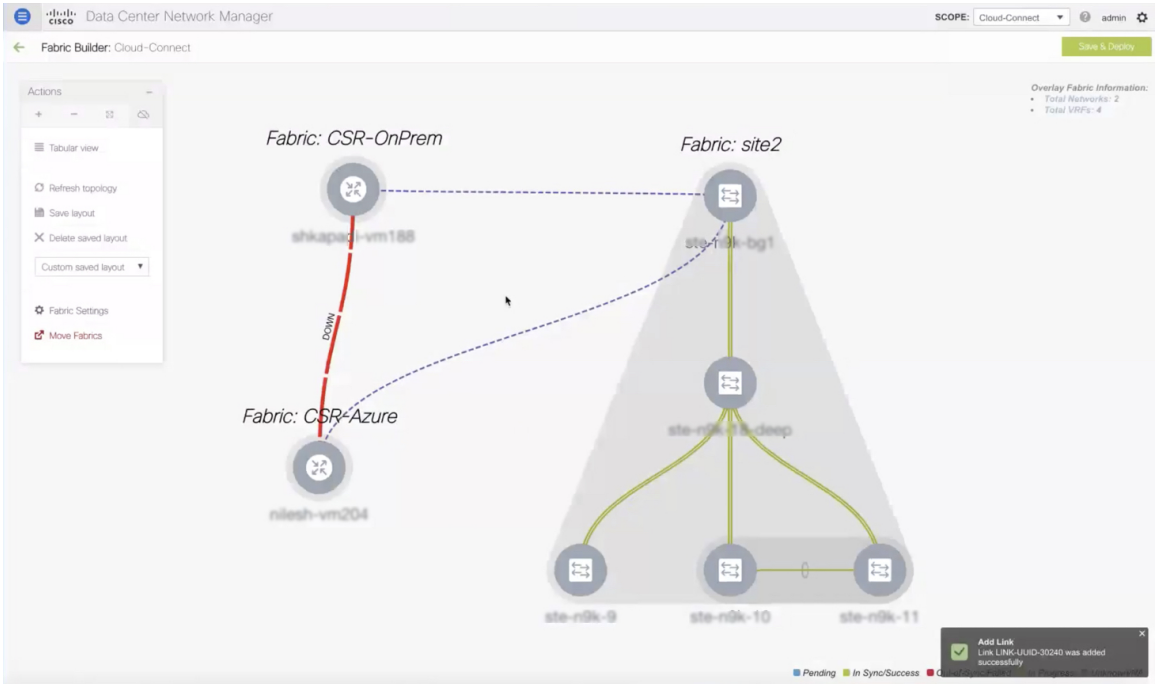

What to do next

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.

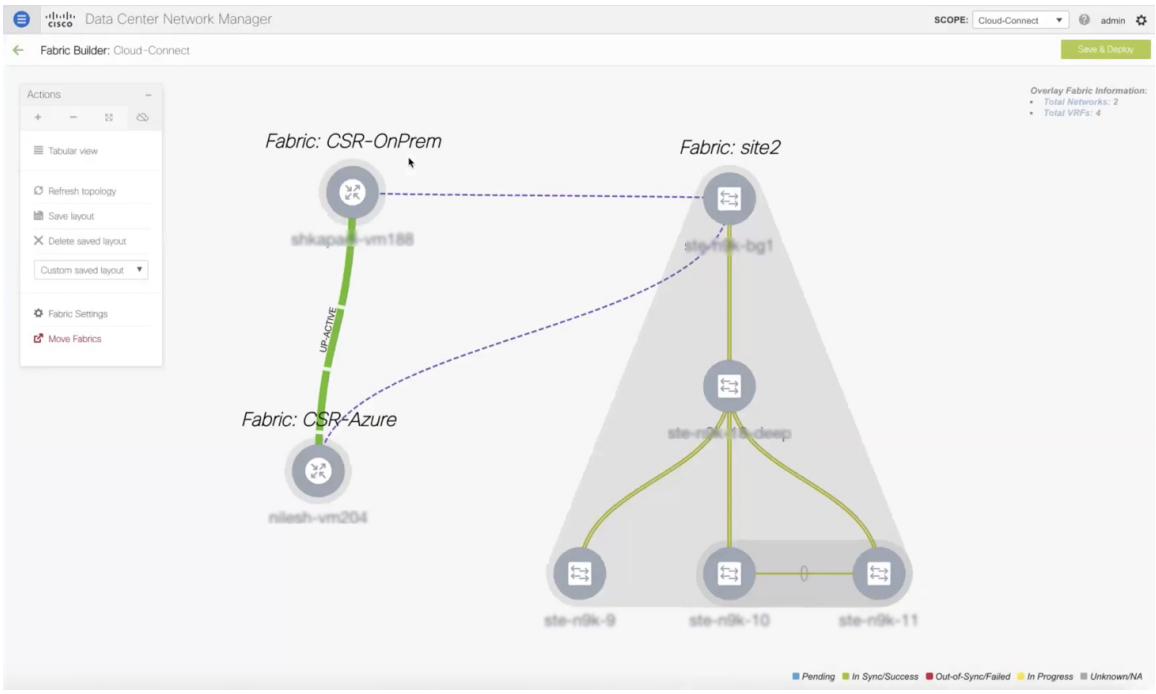

What to do next

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.

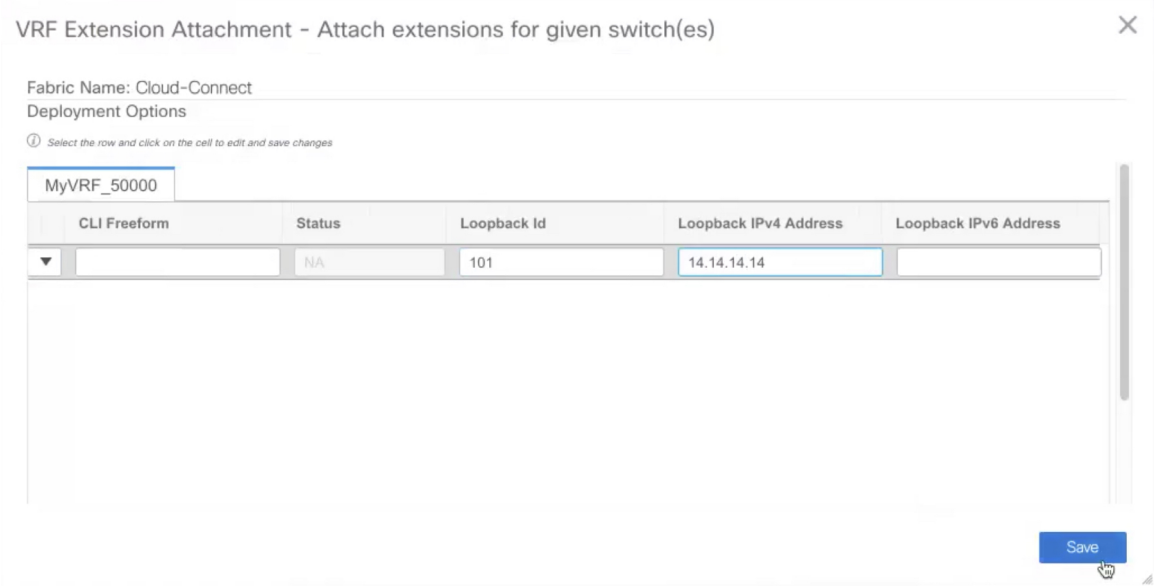

Step 7 Click **Save**.

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.

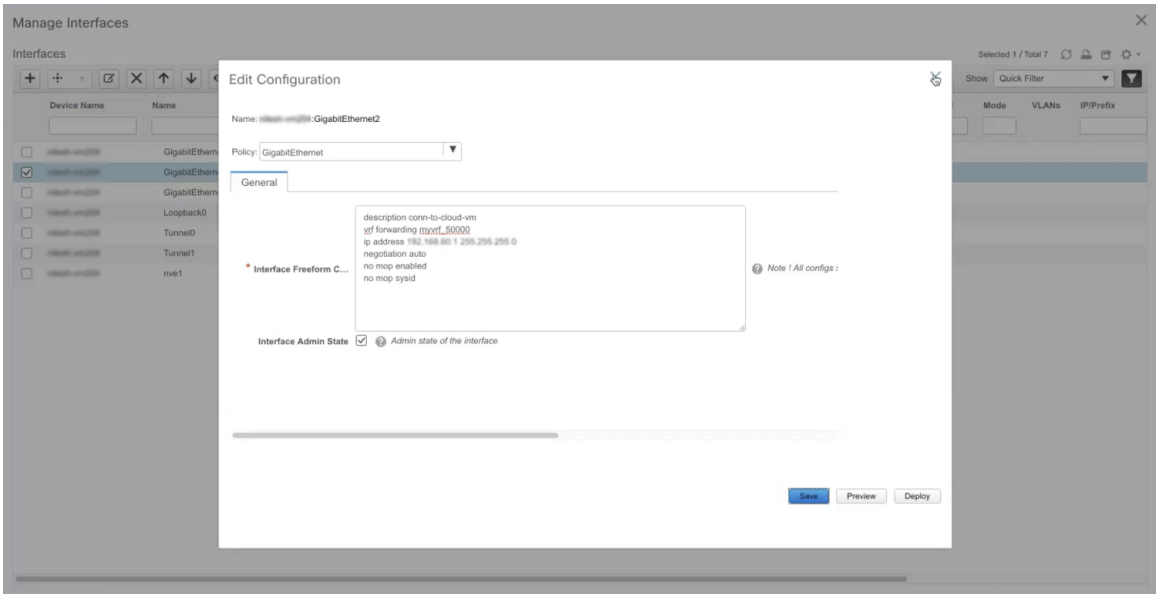

The fabric topology window appears.

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.

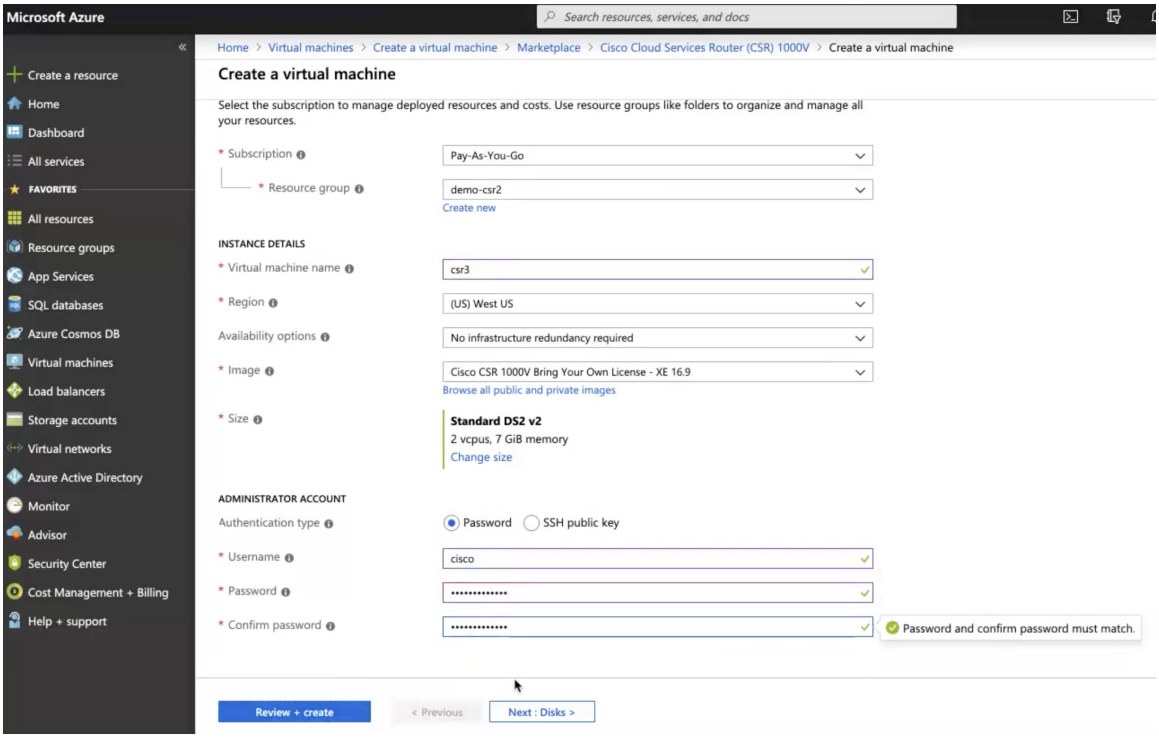

Step 11 Click **Next : Disks >**.

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.

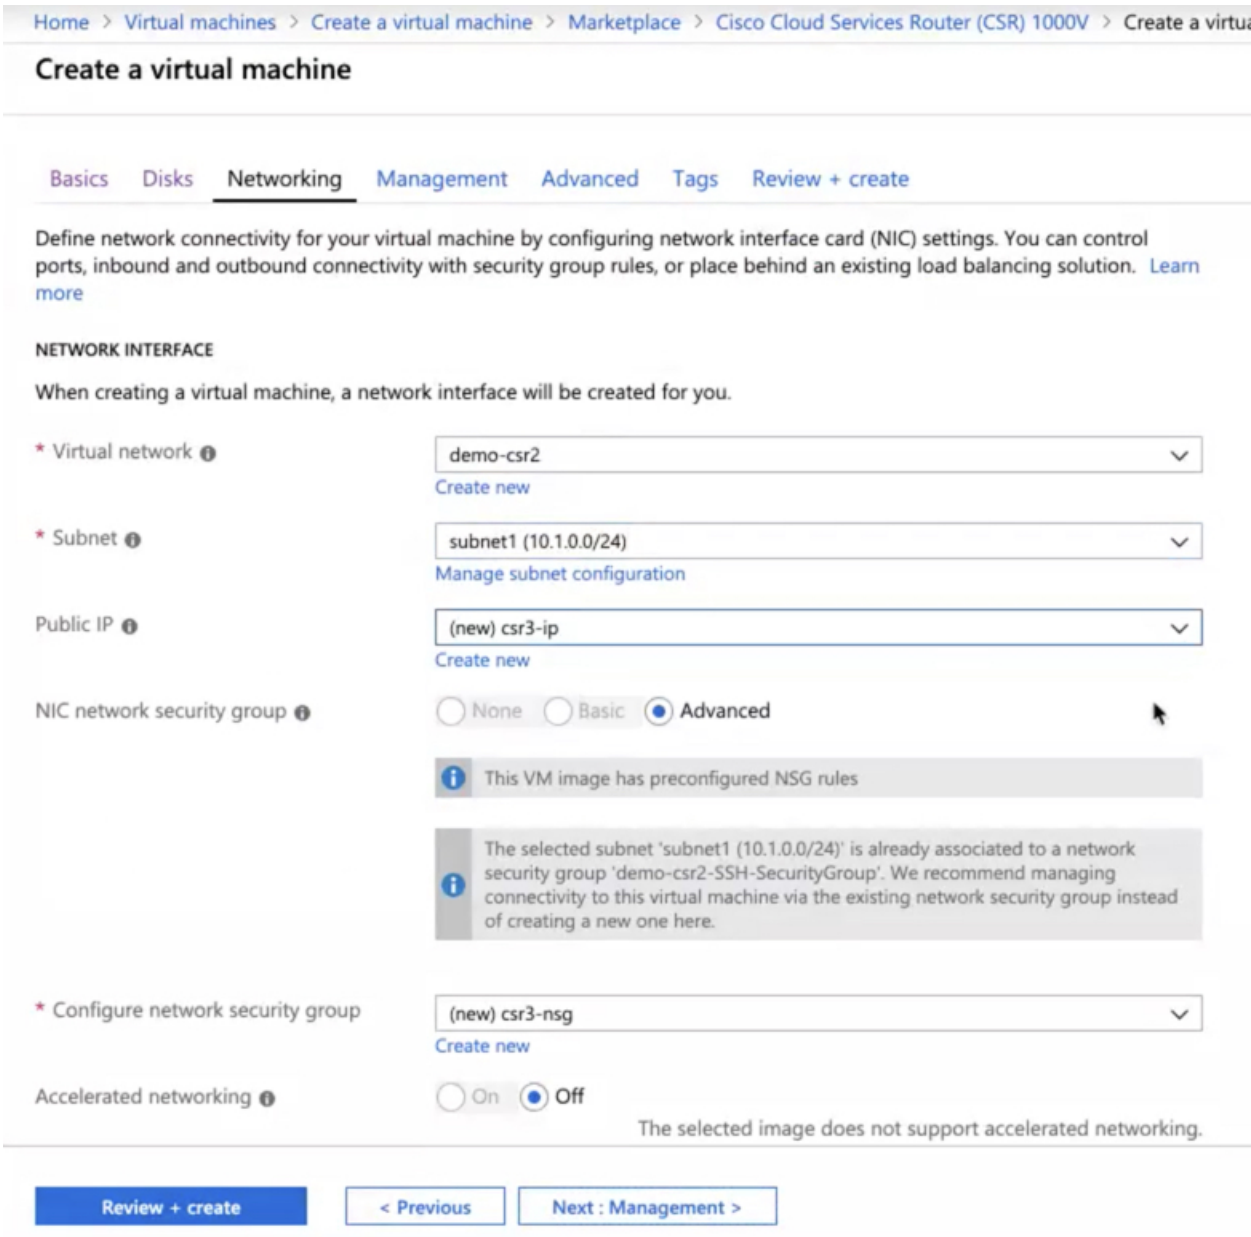

Step 16 Use the default values in other fields.

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.

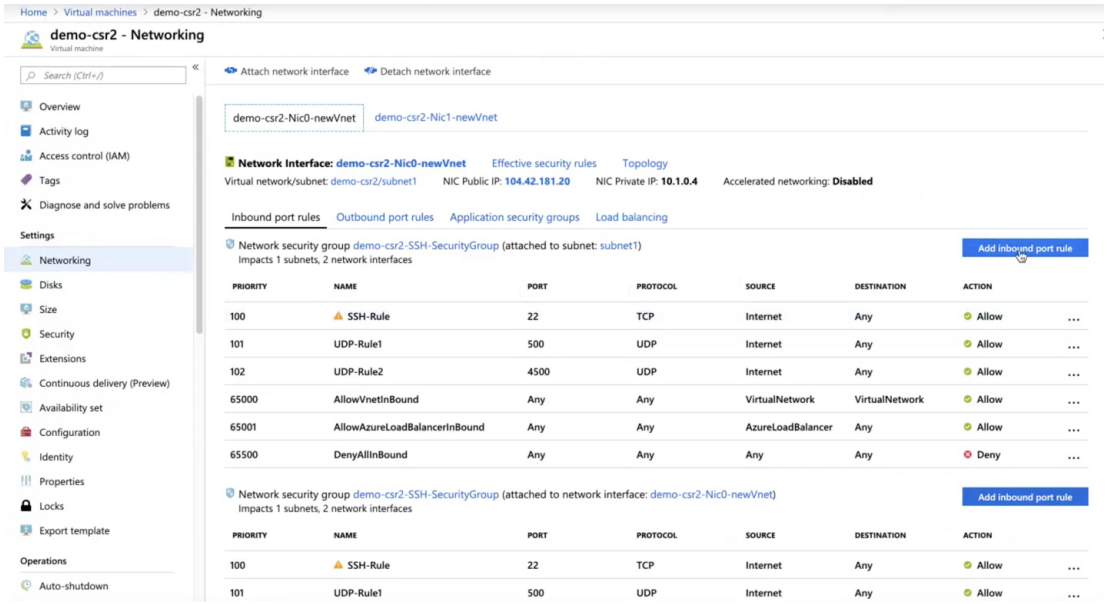

- Create routes in the **Routes** setting of the VM to create traffic routes between the on-prem data center and Microsoft Azure. You can use the default route to redirect traffic from the VNet to Cisco CSR 1000v.

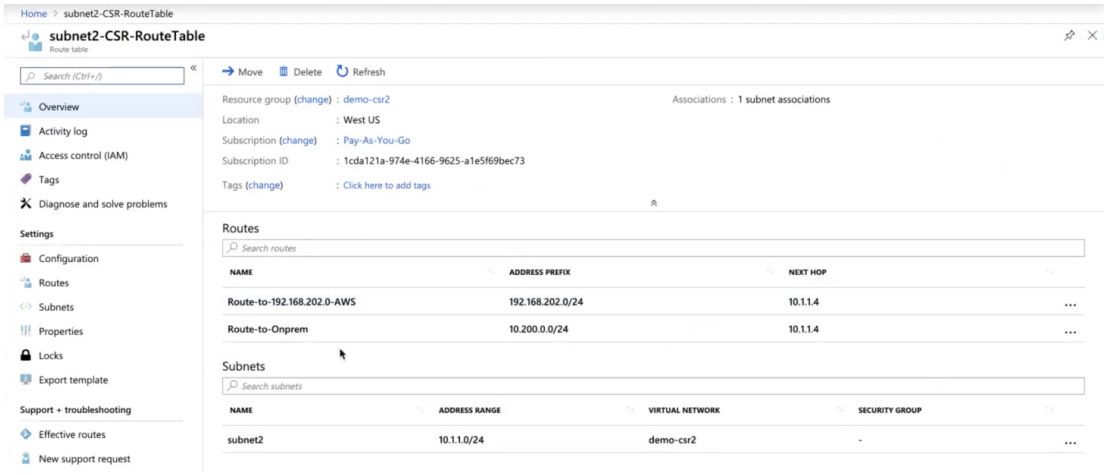

See *Cisco CSR 1000v Deployment Guide for Microsoft Azure* for more information.

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 71 section more information.

Resetting Packet Counter Using API

To reset the packet counter, perform the following steps:

Procedure

Step 1 Log into Cisco DCNM.

Step 2 Navigate to the https://*DCNM-IP*/api-docs URL.

Step 3 Expand the GET **/cloud-extension/status/{ipAddress}** API under cloud extension.

Step 4 Enter the IP address of the on-prem core router.

Step 5 Set the **fetchLatestFromSwitch** value to **true**.

Step 6 Click **Try it out**.

The packet counter is cleared and the count drops to zero.

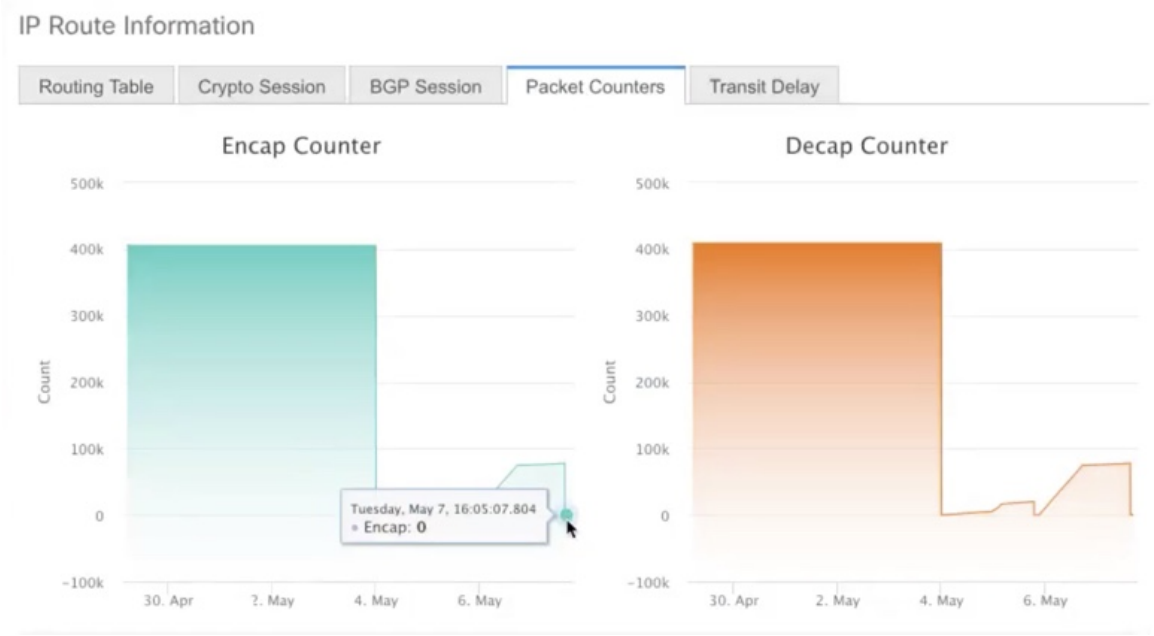
IP Route Information
Routing Table
Crypto Session
BGP Session
Packet Counters
Transit Delay
Encap Counter
Decap Counter
500k
400k
300k
200k
100k
0
-100k
Count
30. Apr
?. May
4. May
6. May
2. May
Tuesday, May 7, 16:05:07.804
Encap: 0

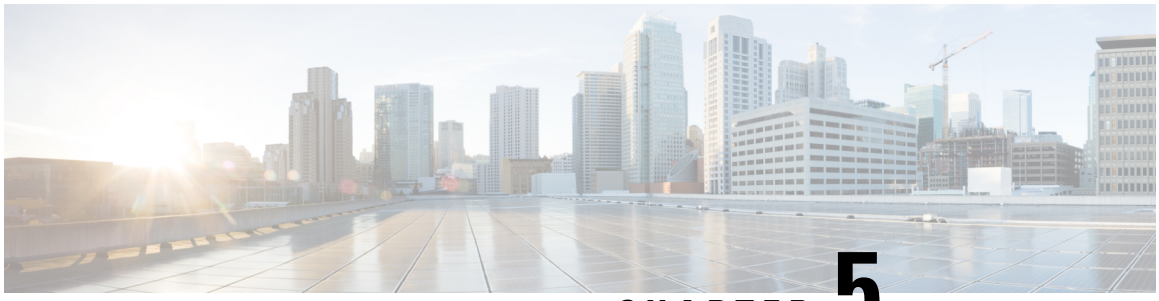

CHAPTER 5

Managing a Brownfield VXLAN BGP EVPN Fabric

This chapter explains how to migrate a Brownfield fabric into Cisco DCNM.

- Overview, on page 73
- Prerequisites, on page 74
- Guidelines and Limitations, on page 74
- Fabric Topology Overview, on page 76
- DCNM Brownfield Deployment Tasks, on page 77
- Verifying the Existing VXLAN BGP EVPN Fabric, on page 77
- Creating a VXLAN BGP EVPN Fabric, on page 80
- Adding Switches and Transitioning VXLAN Fabric Management to DCNM, on page 94
- Verifying the Import of the VXLAN BGP EVPN Fabric, on page 107
- Configuration Profiles Support for Brownfield Migration, on page 115
- Migrating a Bottom-Up VXLAN Fabric to DCNM, on page 115
- 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 124
- 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 128
- Changing a Brownfield Imported BIDIR Configuration, on page 130
- Manually Adding PIM-BIDIR Configuration for Leaf or Spine Post Brownfield Migration , on page 131
- Migrating an MSD Fabric with Border Gateway Switches , on page 131

Overview

This use case shows how to migrate an existing VXLAN BGP EVPN fabric to Cisco DCNM. The transition involves migrating existing network configurations to DCNM.

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

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

Prerequisites

- DCNM-supported NX-OS software versions. For details, refer *Cisco DCNM Release Notes, Release 11.3(1)*.
- Underlay routing protocol is OSPF or IS-IS.
- The supported underlay is based on the DCNM 10.2(1) POAP template's best practices for the VXLAN fabric (dcnm_ip_vxlan_fabric_templates.10.2.1.ST.1.zip) available on Cisco.com.
- The following fabric-wide loopback interface IDs must not overlap:
  - Routing loopback interface for IGP/BGP.
  - VTEP loopback ID
  - Underlay rendezvous point loopback ID if ASM is used for multicast replication.
- BGP configuration uses the ‘router-id’, which is the IP address of the routing loopback interface.
- If the iBGP peer template is configured, then it must be configured on the leaf switches and route reflectors. The template name that needs to be used between leaf and route reflector should be identical.
- The BGP route reflector and multicast rendezvous point (if applicable) functions are implemented on spine switches. Leaf switches do not support the functions.
- Familiarity with VXLAN BGP EVPN fabric concepts and functioning of the fabric from the DCNM perspective.
- Fabric switch nodes are operationally stable and functional and all fabric links are up.
- vPC switches and the peer links are up before the migration. Ensure that no configuration updates are in progress or changes pending.
- Create an inventory list of the switches in the fabric with their IP addresses and credentials. DCNM uses this information to connect to the switches.
- Shut down any other controller software you are using presently so that no further configuration changes are made to the VXLAN fabric. Alternatively, disconnect the network interfaces from the controller software (if any) so that no changes are allowed on the switches.
- The switch overlay configurations must have the mandatory configurations defined in the shipping DCNM Universal Overlay profiles. Additional network or VRF overlay related configurations found on the switches are preserved in the freeform configuration associated with the network or VRF DCNM entries.
- All the overlay network and VRF profile parameters such as VLAN name and route map name should be consistent across all devices in the fabric for the brownfield migration to be successful.

Guidelines and Limitations

- Fabric interfaces can be numbered or unnumbered.
- Various other interface types are supported.

- The following features are unsupported.
  - eBGP underlay
  - Layer 3 port channel
- 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.

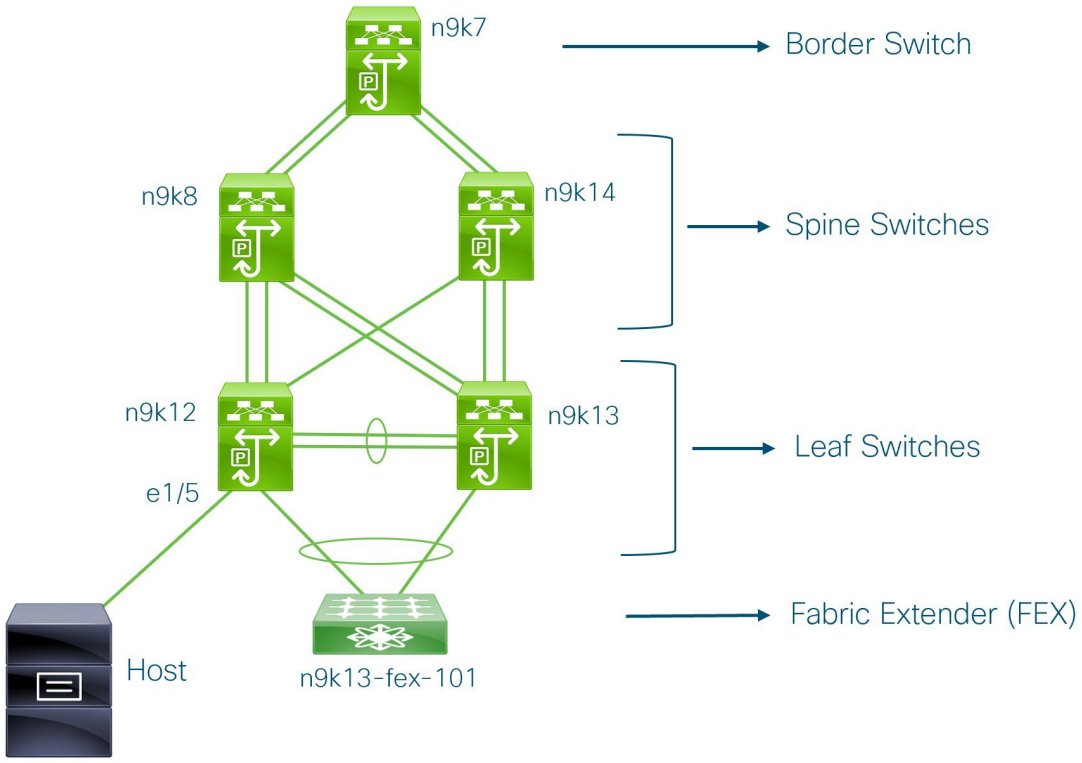

You can see that there is a border switch, two spine switches, two leaf switches, and a Fabric Extender or FEX.

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 77
2. Creating a VXLAN BGP EVPN Fabric, on page 80
3. Adding Switches and Transitioning VXLAN Fabric Management to DCNM, on page 94
4. Verifying the Import of the VXLAN BGP EVPN Fabric, on page 107

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:

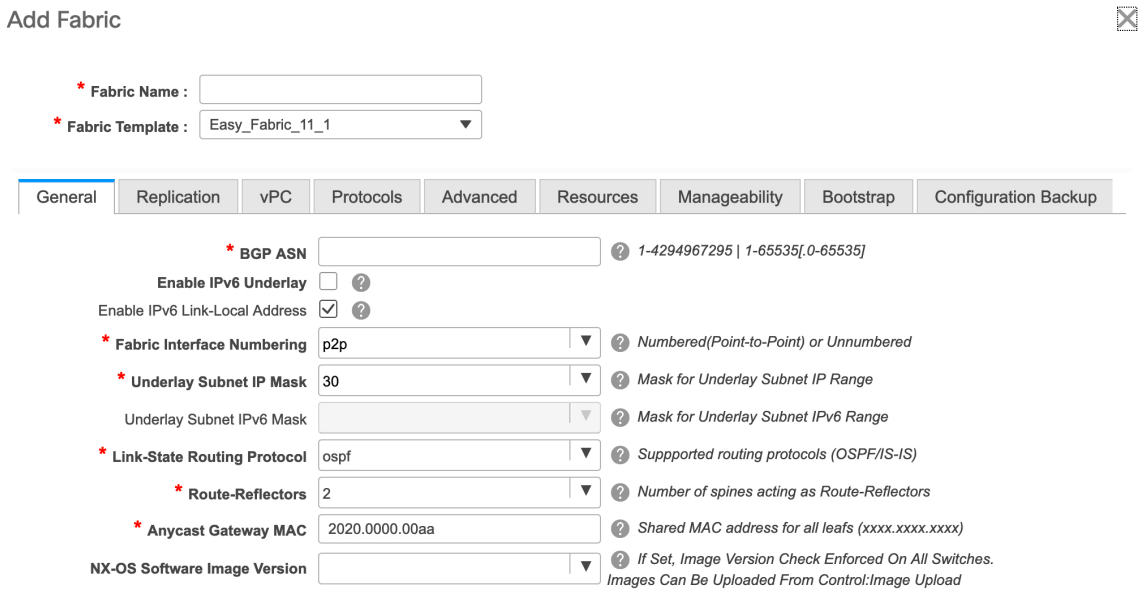

BGP ASN: Enter the BGP AS number the fabric is associated with.

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.

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

Route-Reflectors – The Route Reflector count is only applicable post-migration. The existing route reflector configuration is honored when importing into the DCNM setup.

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

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

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

Decreasing the count

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

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

b. Identify the spine switches designated as route reflectors.

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

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

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

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

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

Anycast Gateway MAC: Enter the Anycast gateway MAC address of the existing fabric.

NX-OS Software Image Version: Leave this field blank. You can update this post-transition, as desired.

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

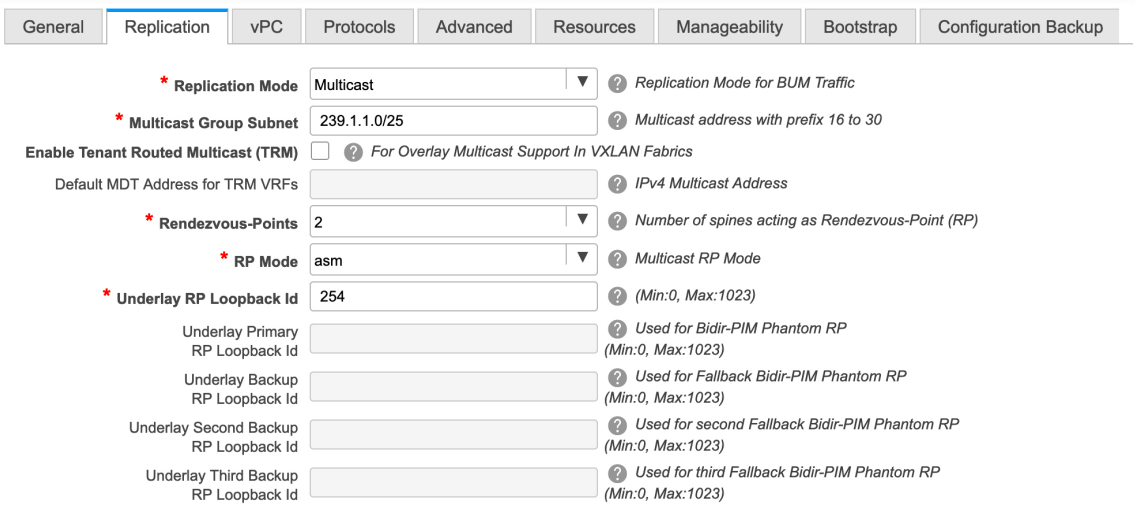

Replication Mode: The mode of replication that is used in the existing fabric, Ingress Replication, or Multicast.

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

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

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

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

If you enable TRM, the Multicast address for TRM must be entered. All the TRM specific tenant configuration is captured in the switch freeform policy linked to the tenant network and VRF profile.

Note that the TRM feature is unsupported on switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images.

Default MDT address for TRM VRFs – Enter the default multicast distribution tree (MDT) IPv4 address for TRM VRFs.

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

RP mode – Select **asm** (Any-Source Multicast) or **bidir** (Bidirectional PIM) mode.

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

The **asm** RP mode supports up to 4 RPs.

The **bidir** mode supports up to 2 RPs. An error message is displayed if the BIDIR configuration indicates that more than 2 RPs are used.

After brownfield migration, only 2 RPs are supported in the migrated fabric. An error message is displayed when you click **Save & Deploy** after changing the RP count to 4.

If an RP is down or deleted from the fabric, this RP cannot be replaced by another spine as Easy Fabric does not remember the configuration of a removed switch. Easy Fabric uses a specific scheme to generate RP configuration for Bidir. Therefore, the generated Bidir configuration will not work with the brownfield imported configuration. After brownfield migration, if you change the RP count or add new spine or leaf switches, you should manually configure the PIM-Bidir feature. If a manual configuration is required, a warning message is displayed after you click **Save & Deploy**. For more information, see *Manually Adding PIM-BIDIR Configuration for Leaf or Spine Post Brownfield Migration*.

You can also modify a brownfield imported bidir configuration to use the configuration generated by **Fabric Builder**. For more information, see *Changing a Brownfield Imported BIDIR Configuration*.

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

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

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

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

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

Underlay Second Backup RP Loopback Id – The second fallback loopback ID for Phantom RP, for multicast protocol peering purposes in the fabric underlay.

Underlay Third Backup RP Loopback Id – The third fallback loopback ID for Phantom RP, for multicast protocol peering purposes in the fabric underlay.

Step 5 Click the **vPC** tab. Most of the fields are auto generated.

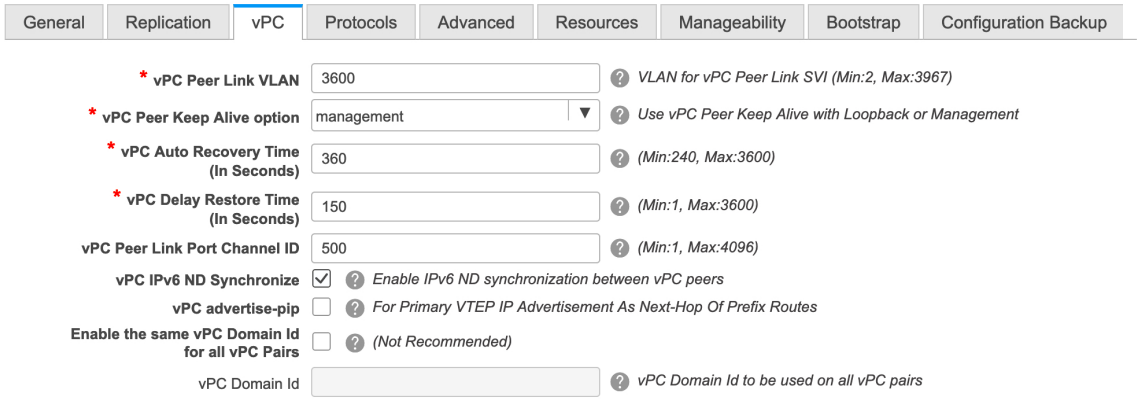

vPC Peer Link VLAN - Enter the VLAN ID used for the vPC peer link SVI in the existing fabric.

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

If you only use IPv6 addresses on the management interface, you must use the loopback option.

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

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

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

vPC Peer Link Port Channel 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.

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

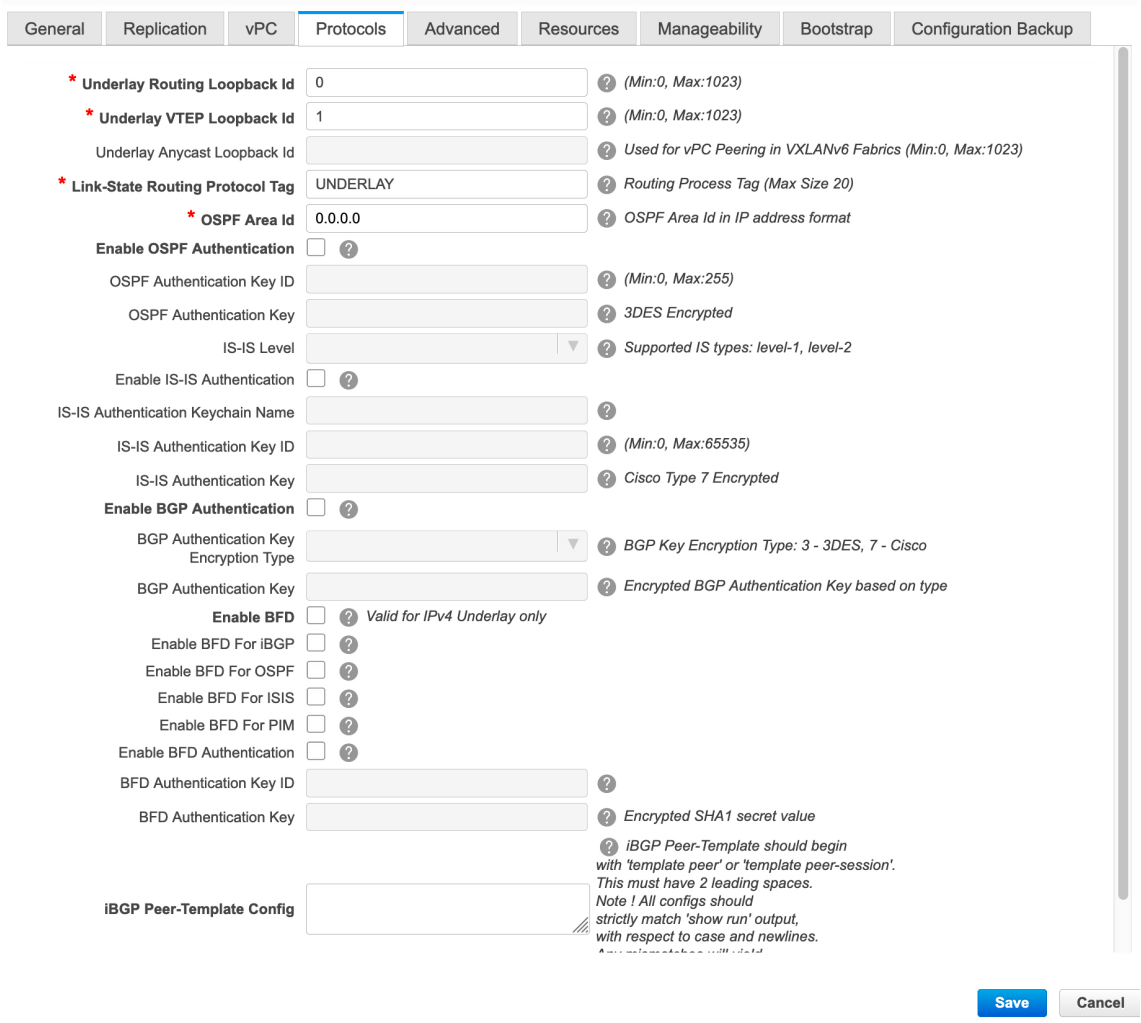

Underlay Routing Loopback Id - The loopback interface ID is populated as 0 since loopback0 is usually used for fabric underlay IGP peering purposes. This must match the existing configuration on the switches. This must be the same across all the switches.

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

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

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

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

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

OSPF Authentication Key ID – Enter the OSPF authentication key ID.

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

Note Plain text passwords are not supported. Login to the switch, retrieve the OSPF authentication details.

You can obtain the OSPF authentication details by using the **show run ospf** command on your switch.

```
# show run ospf | grep message-digest-key
ip ospf message-digest-key 127 md5 3 c7c83ec78f38f32f3d477519630faf7b
```

In this example, the OSPF authentication key ID is **127** and the authentication key is **c7c83ec78f38f32f3d477519630faf7b**.

For information about how to configure a new key and retrieve it, see *Retrieving the Authentication Key*.

IS-IS Level - Select the IS-IS level from this drop-down list.

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

IS-IS Authentication Keychain Name - Enter the keychain name.

IS-IS Authentication Key ID - Enter the IS-IS authentication key ID.

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

Note Plain text passwords are not supported. Login to the switch, retrieve the IS-IS authentication details.

You can obtain the IS-IS authentication details by using the **show run | section "key chain"** command on your switch.

```
# show run | section "key chain"
key chain CiscoIsisAuth
  key 127
      key-string 7 075e731f
```

In this example, the keychain name is **CiscoIsisAuth**, the key ID is **127**, and the type 7 authentication key is **075e731f**.

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

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

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

Note Plain text passwords are not supported. Login to the switch, retrieve the BGP authentication details.

You can obtain the BGP authentication details by using the **show run bgp** command on your switch.

```
# show run bgp
neighbor 10.2.0.2
remote-as 65000
password 3 sd8478fswerdfw3434fsw4f4w34sdsd8478fswerdfw3434fsw4f4w3
```

In this example, the BGP authentication key is displayed after the encryption type **3**.

Enable BFD feature – Select the check box to enable the BFD feature.

The BFD feature is disabled by default.

Make sure that the BFD feature setting matches with the switch configuration. If the switch configuration contains **feature bfd** but the BFD feature is not enabled in the fabric settings, config compliance generates diff to remove the BFD feature after brownfield migration. That is, **no feature bfd** is generated after migration.

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.
- 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.

Step 7 Click the **Advanced** tab. Most of the fields are auto generated.

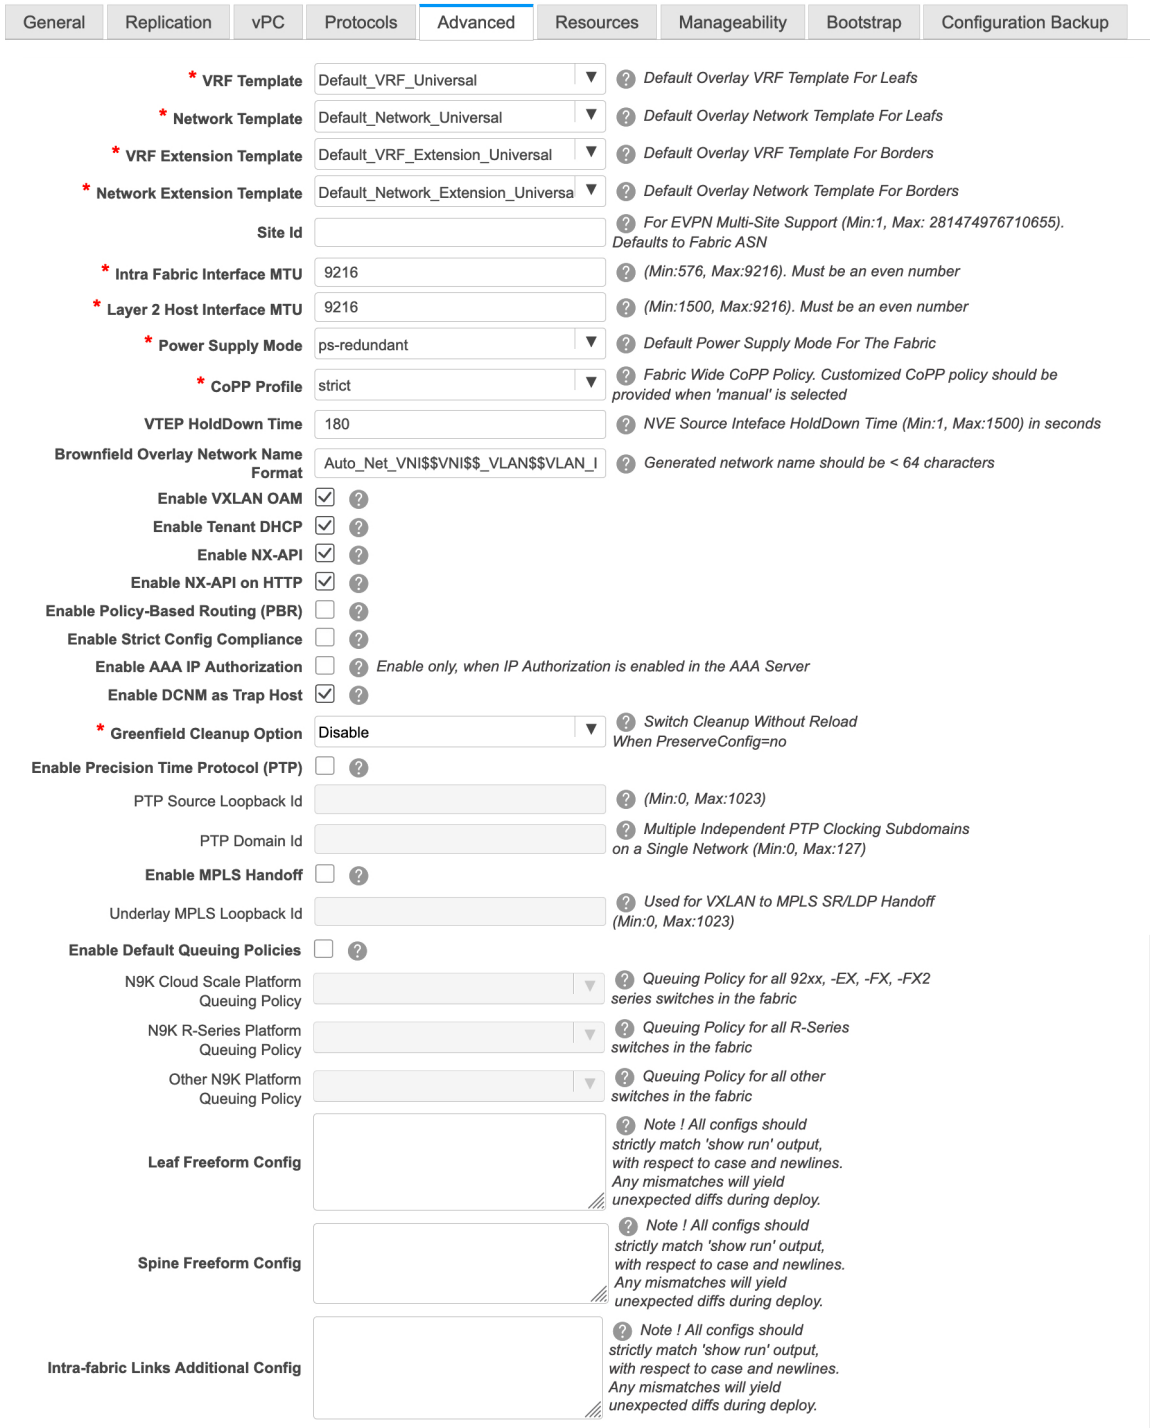

VRF Template and **VRF Extension Template**: Specifies the VRF template for creating VRFs, and the VRF extension template for enabling VRF extension to other fabrics.

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

Intra-fabric Links Additional Config - You can enter this field after fabric transitioning is complete, as needed.

Step 8 Click the **Resources** tab.

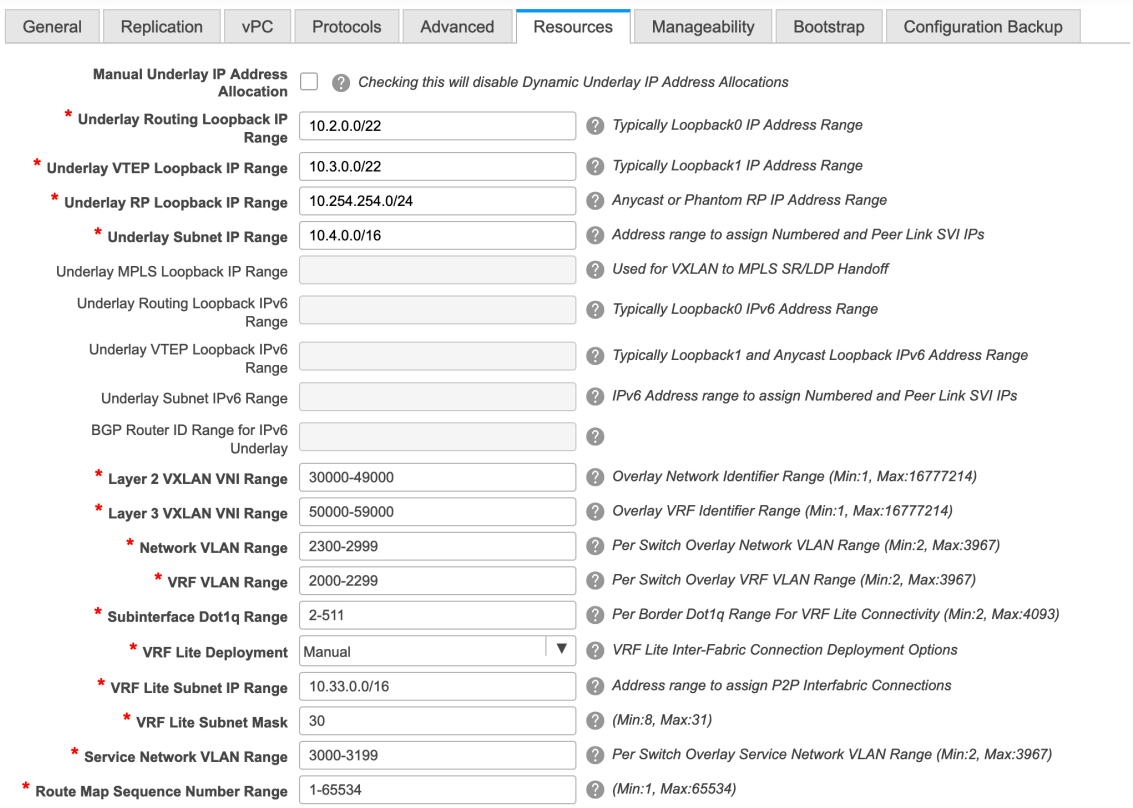

Manual Underlay IP Address Allocation – *Do not* select this check box if you are transitioning your VXLAN fabric management to DCNM.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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.

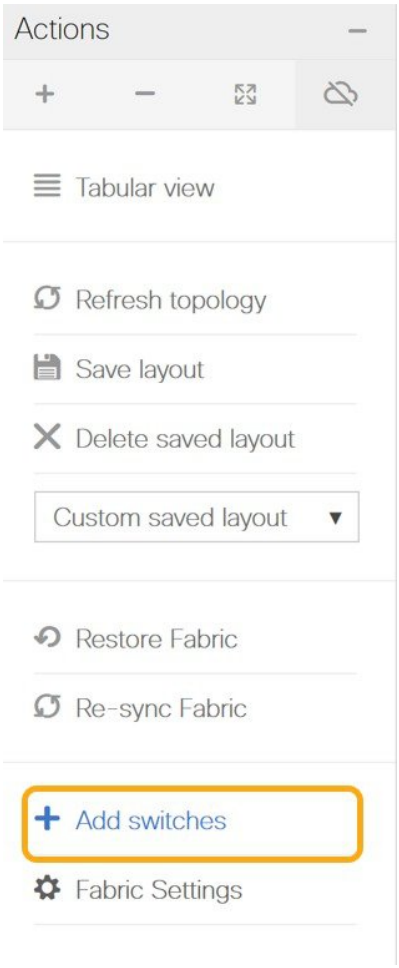

Step 2 Under the **Discover Existing Switches** tab, enter the IP address of the switch in the **Seed IP** field. Enter the username and password of the switches that you want to discover.

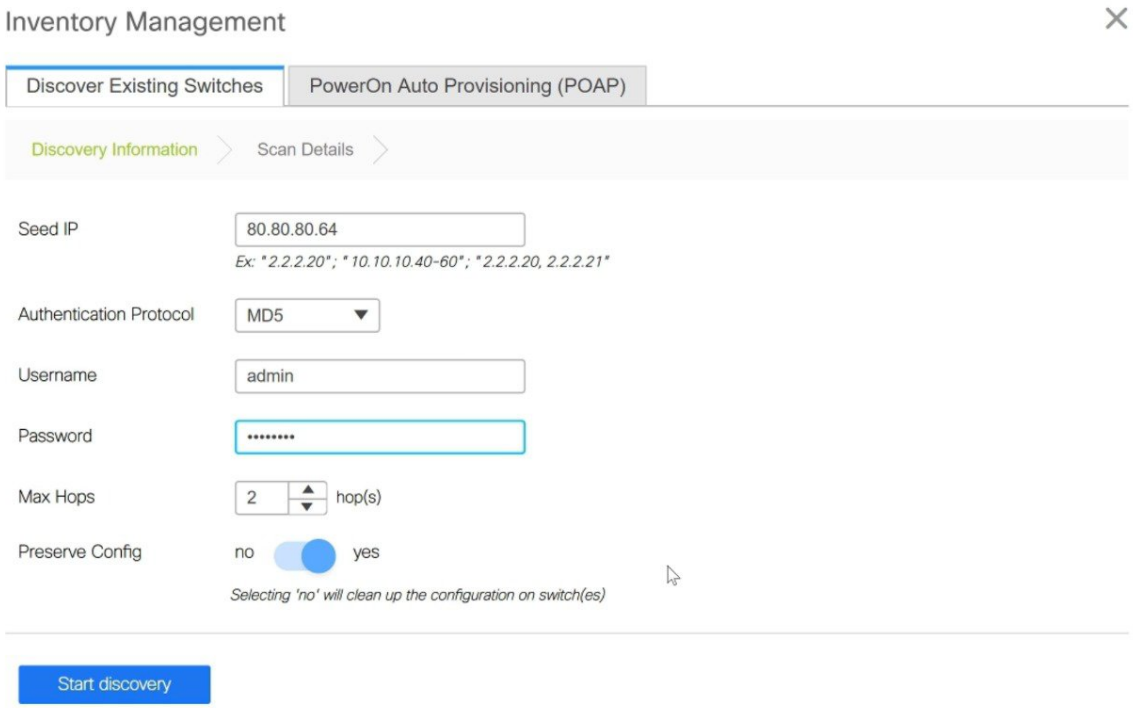

By default, the value in the **Max Hops** field is **2**. The switch with the specified IP address and the switches that are 2 hops from it will be populated after the discovery is complete.

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**.

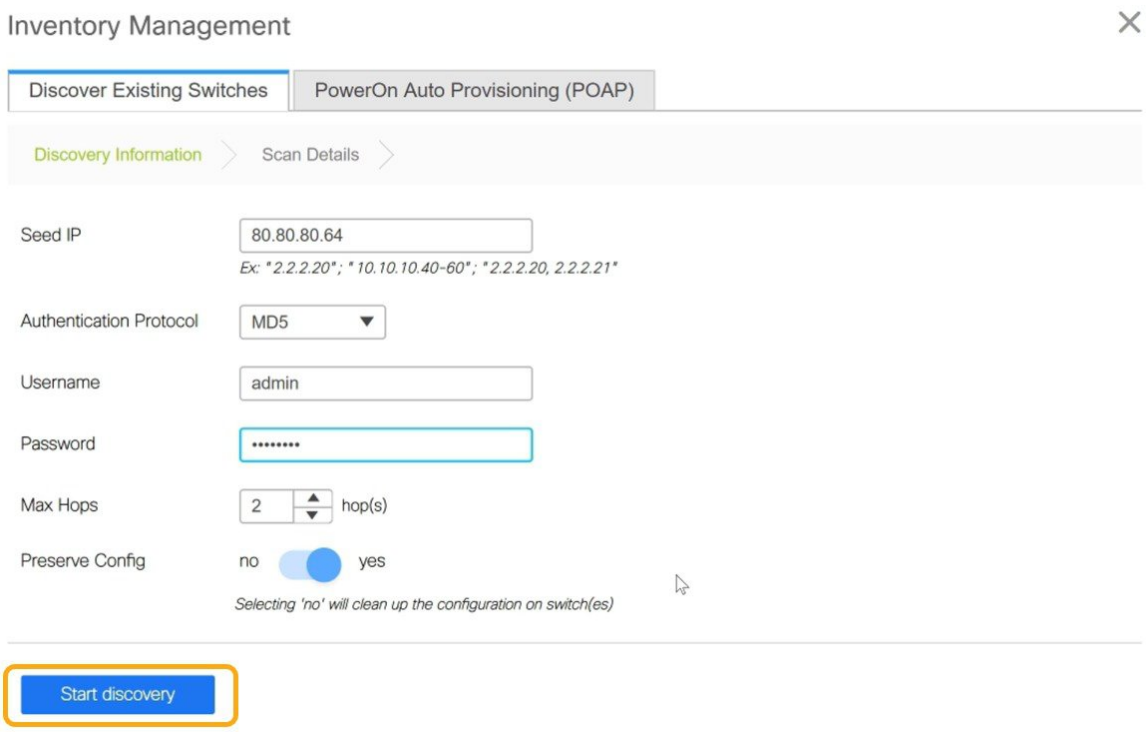

The switch with the specified IP address and switches up to two hops away (depending on the setting of Max Hops) from it are populated in the Scan Details section.

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**.

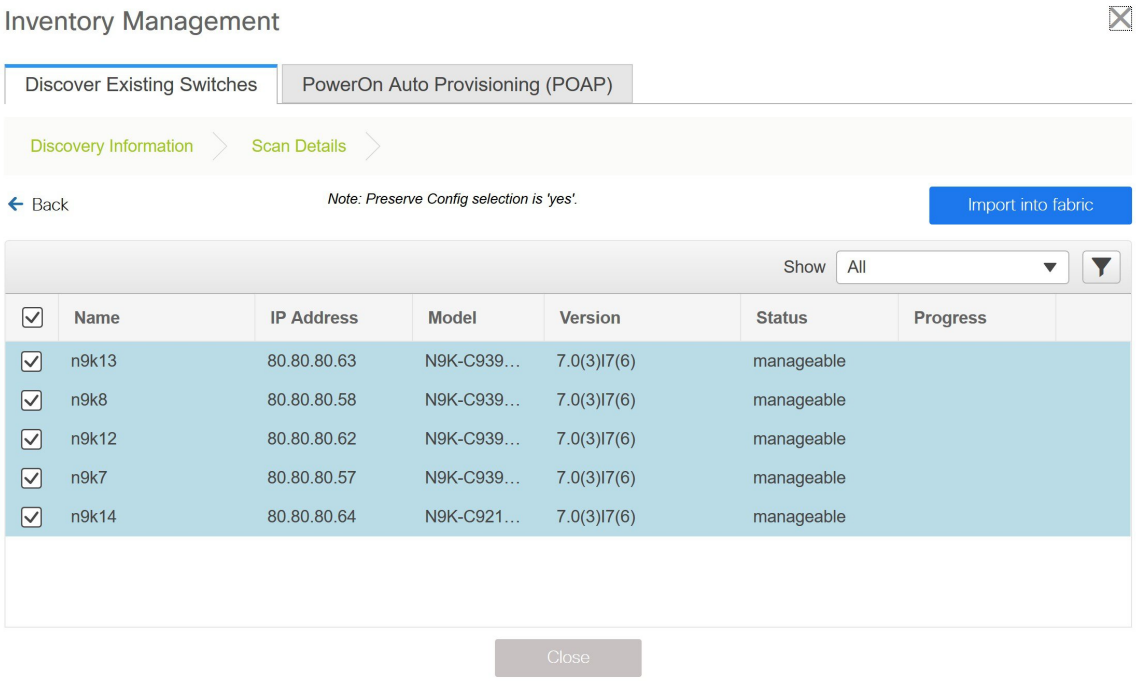

Step 5 Click **Import into fabric**.

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**.

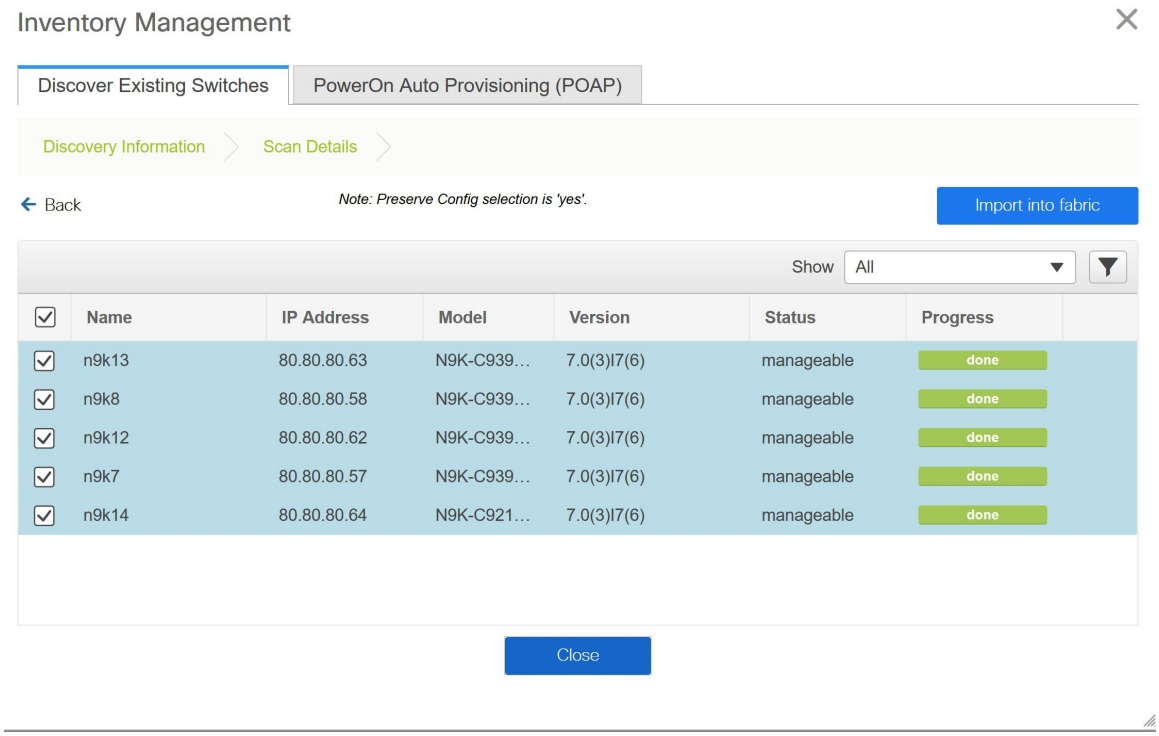

After closing the window, the fabric topology window comes up again. The switch is in Migration Mode now, and the Migration mode label is displayed on the switch icons.

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.

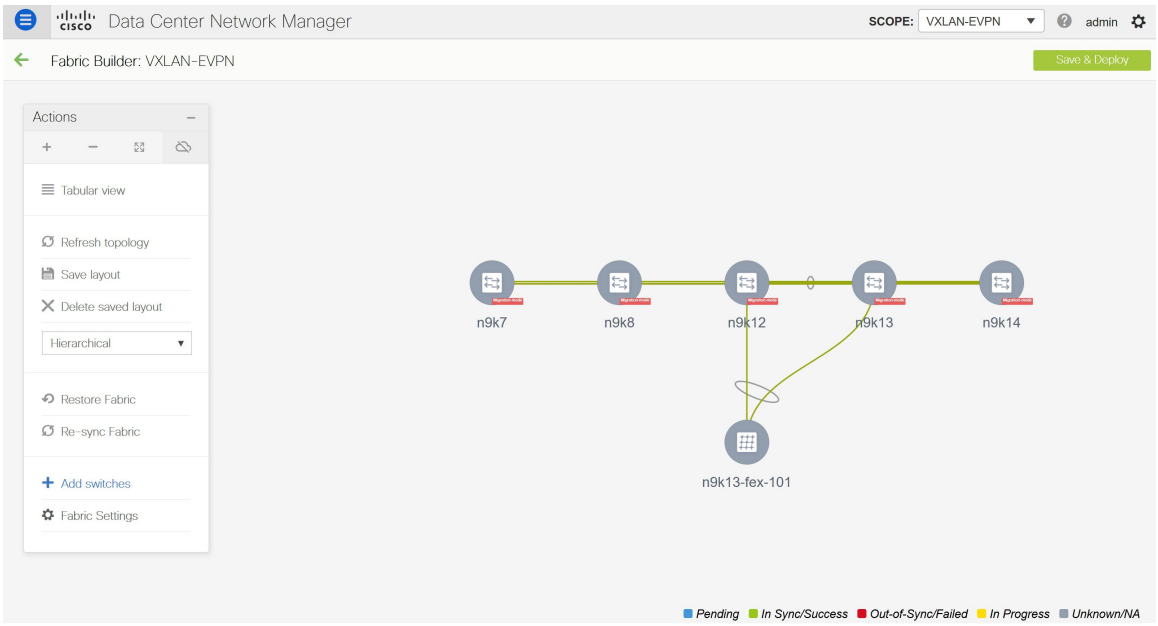

The switch discovery process might fail for a few switches, and the Discovery Error message is displayed. However, such switches are still displayed in the fabric topology. You should remove such switches from the fabric (Right-click the switch icon and click **Discovery** > **Remove** from fabric), 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.

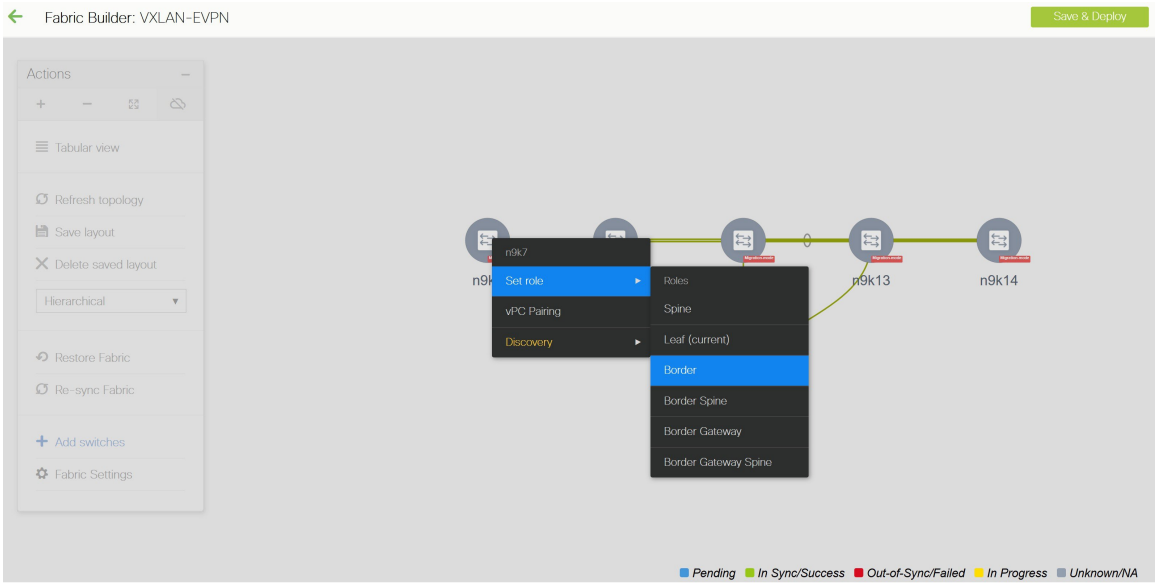

Similarly, set the **Spine** role for the **n9k-14** and **n9k-8** spine switches.

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.

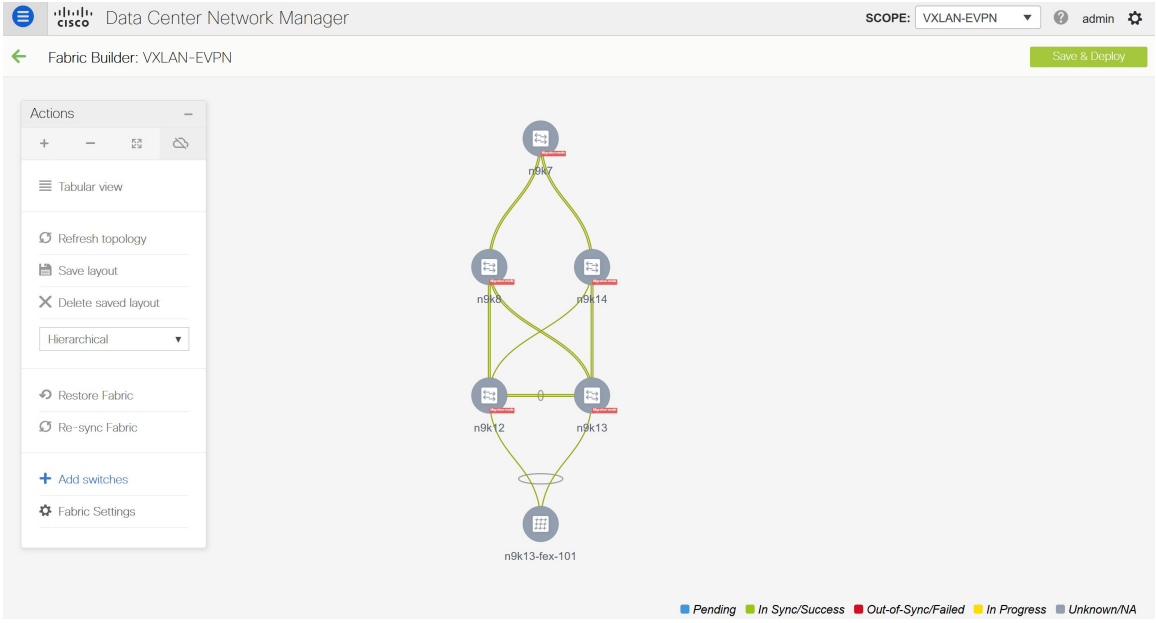

vPC Pairing - The vPC pairing must be done for switches where the Layer 3 vPC peer-keep alive is used. The vPC configuration is automatically picked up from the switches when the vPC peer keep alive is established through the management option. This pairing reflects in the GUI only after the migration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

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

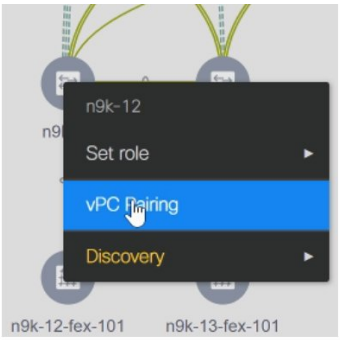

The Select vPC peer screen comes up. It lists potential vPC peer switches.

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

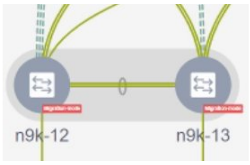

Note Check if you have added all switches from the current fabric. If you have missed adding switches, add them now. Once you are certain that you have imported all existing switches, move to the next step, the Save and Deploy option.

Step 9 Click **Save & Deploy**.

When you click **Save & Deploy**, DCNM obtains switch configurations and populates the state of every switch from the current running config to the current expected config, which is the intended state maintained in DCNM.

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

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

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

Note
- The brownfield migration requires best practices to be followed on the existing fabric such as maintain consistency of the overlay configurations. For more information, see Guidelines and Limitations.
- Any errors or inconsistencies that are found during the migration is reported in fabric errors. The switches continue to remain in the Migration mode. You should fix these errors and complete the migration again by clicking **Save & Deploy** until no errors are reported.

Step 10 After the configurations are generated, review them by clicking the links in the **Preview Config** column.

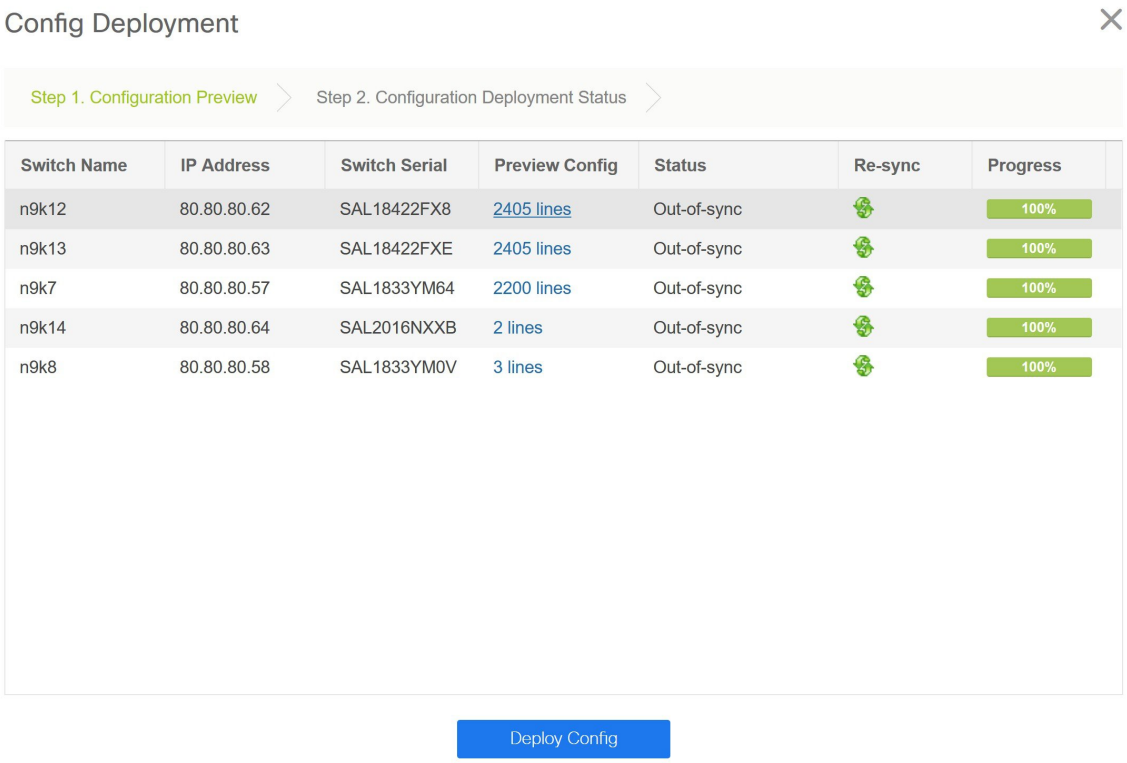

We strongly recommend that you preview the configuration before proceeding to deploy it on the switches. Click the Preview Config column entry. The Config Preview screen comes up. It lists the pending configurations on the switch.

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.

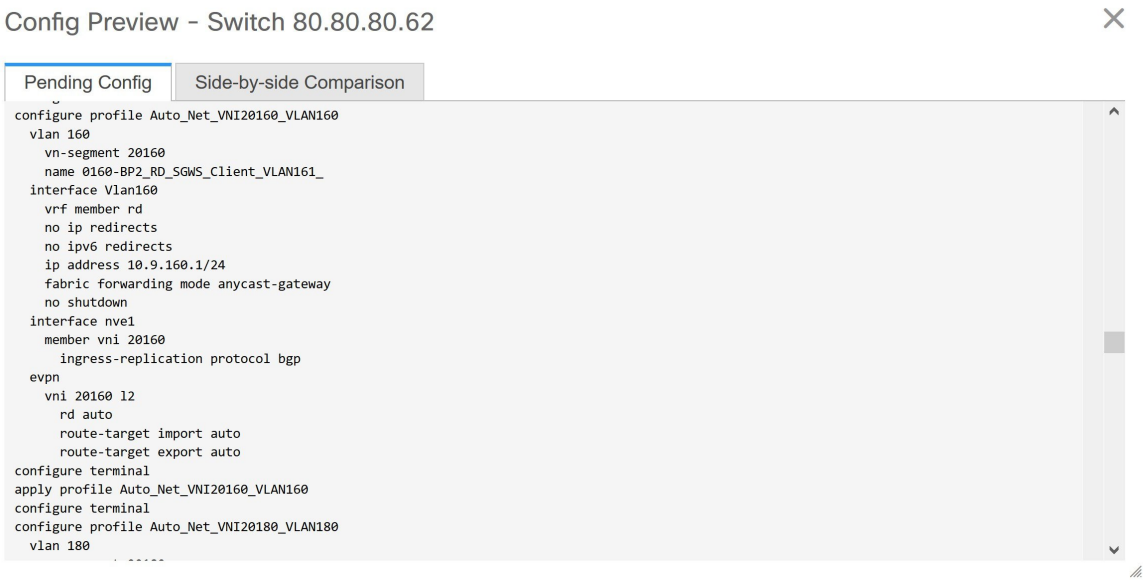

As part of the import process, after the configuration profiles are applied, the original CLI based configuration references will be removed from the switches. These are the 'no' CLIs that will be seen towards the end of the diffs. The VXLAN configurations on the switches will be persisted in the configuration 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.

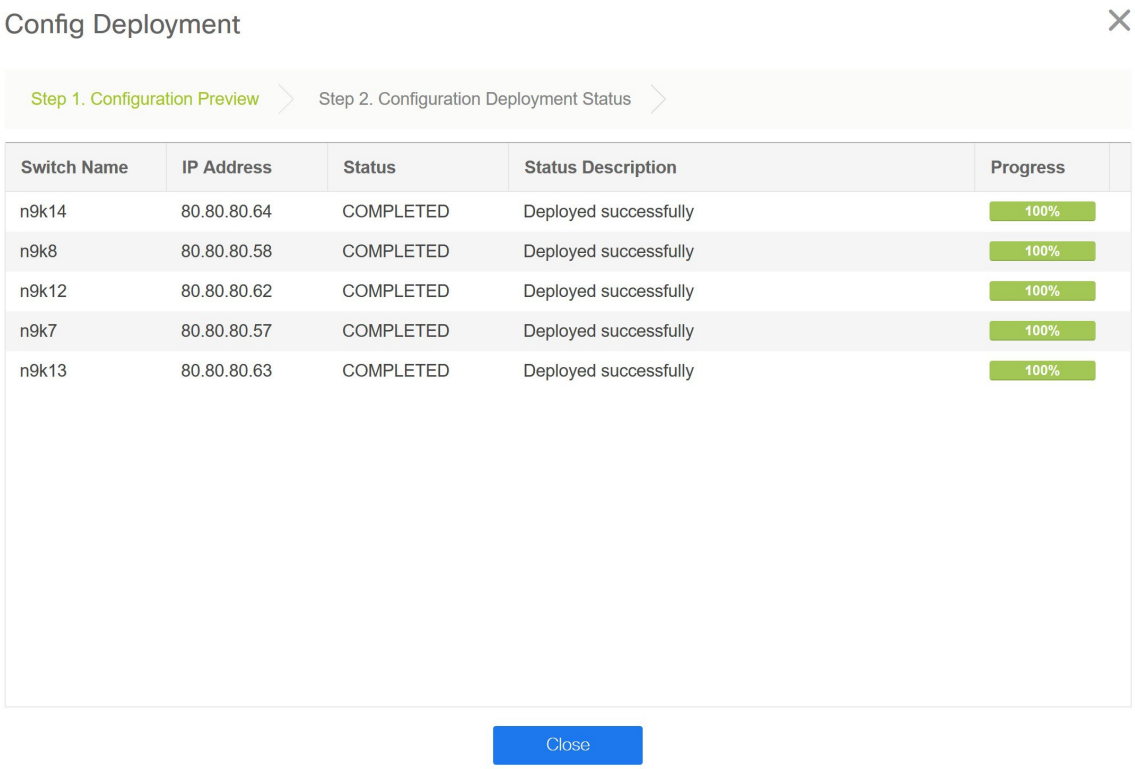

If the **Status** column displays **FAILED**, investigate the reason for failure to address the issue.

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.

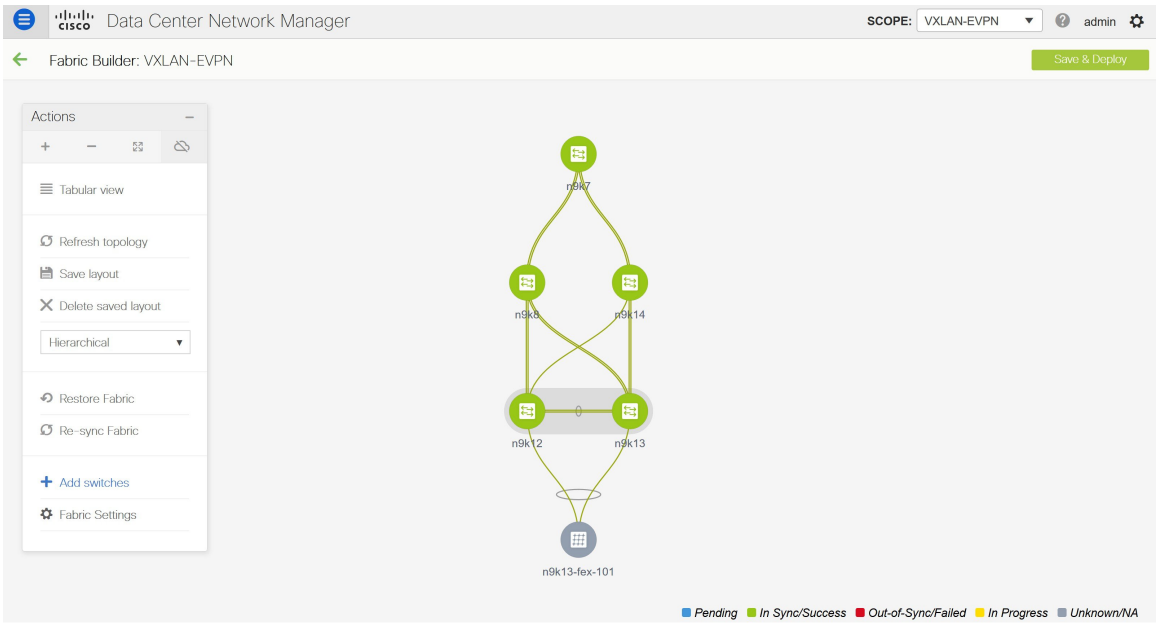

Post-transitioning of VXLAN fabric management to DCNM - This completes the transitioning process of VXLAN fabric management to DCNM. Now, you can add new switches and provision overlay networks for your fabric. For details, refer the respective section in the Fabrics topic in the configuration guide.

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.

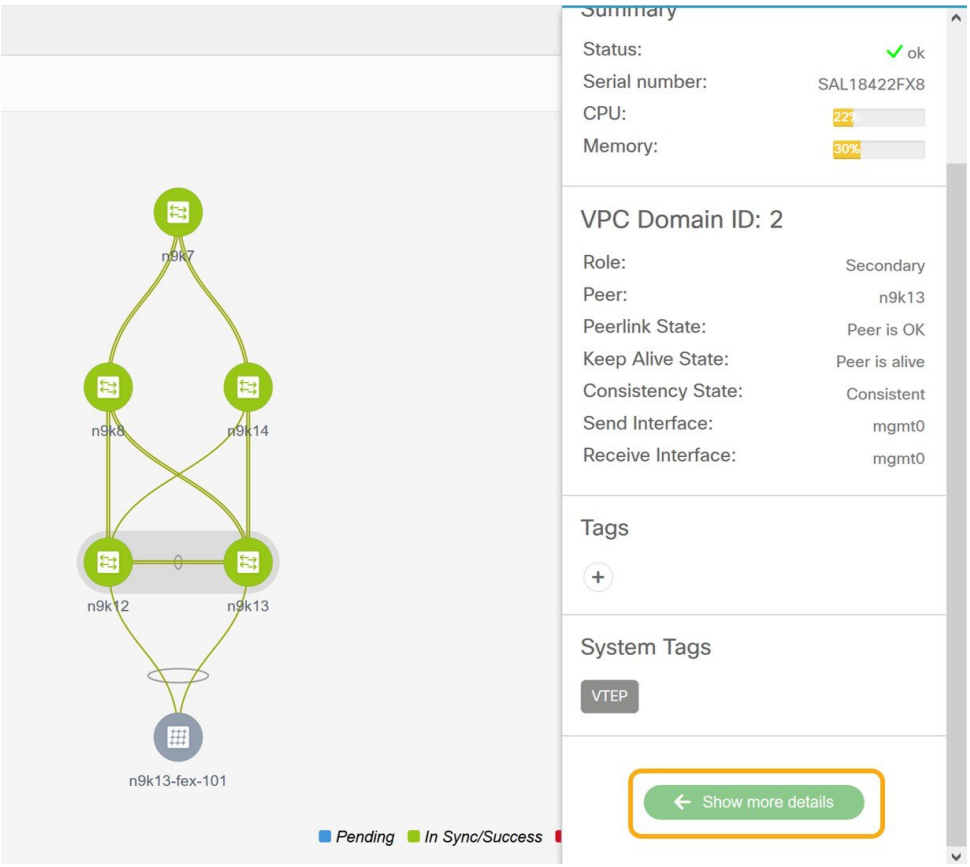

Step 2 Click the **VXLAN** tab.

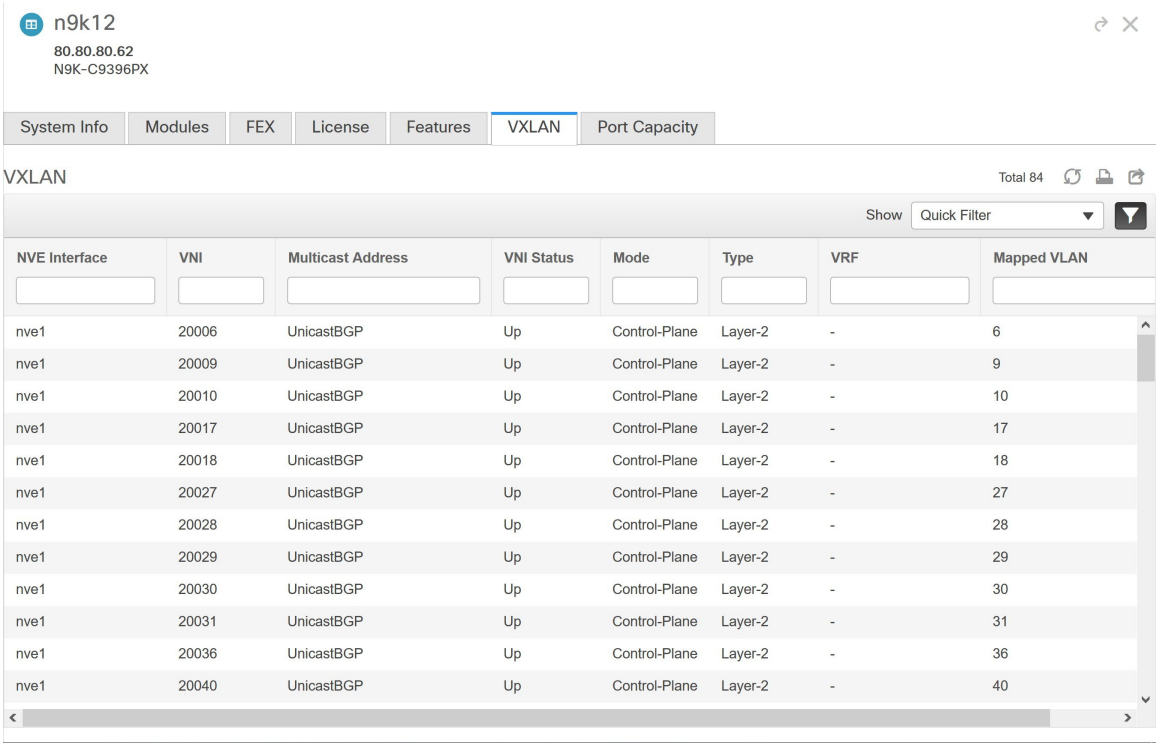

You can see that all the VXLANs have been migrated successfully.

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.

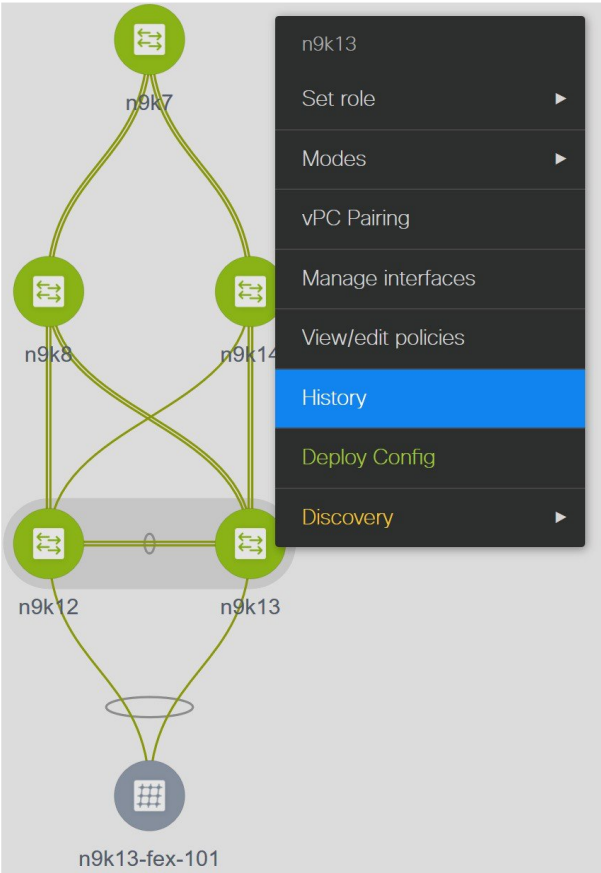

Step 4 Click the **Success** hyperlink under the **Status** column to view the commands pushed by DCNM.

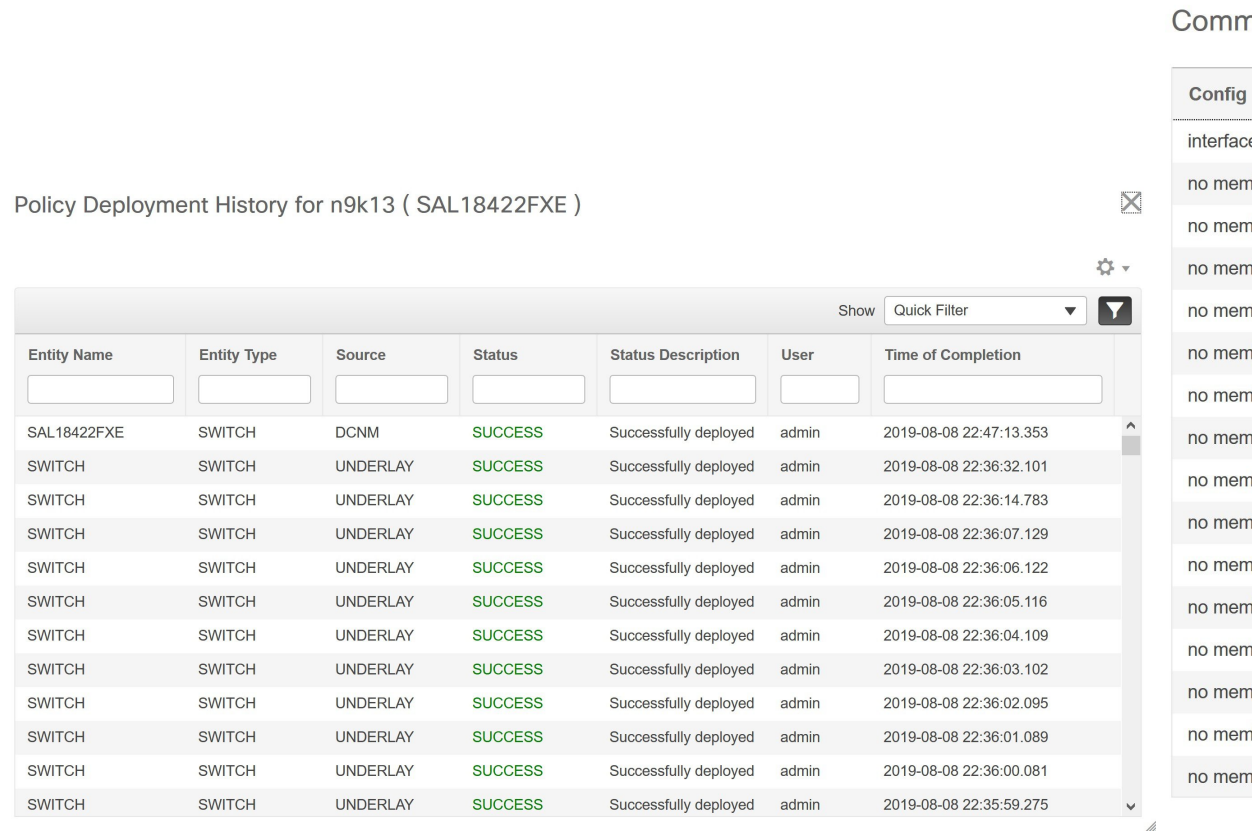

Verifying Resources

DCNM has a resource manager that tracks all the resources used in a fabric. Navigate to **Control > Management > Resources** in the left menu.

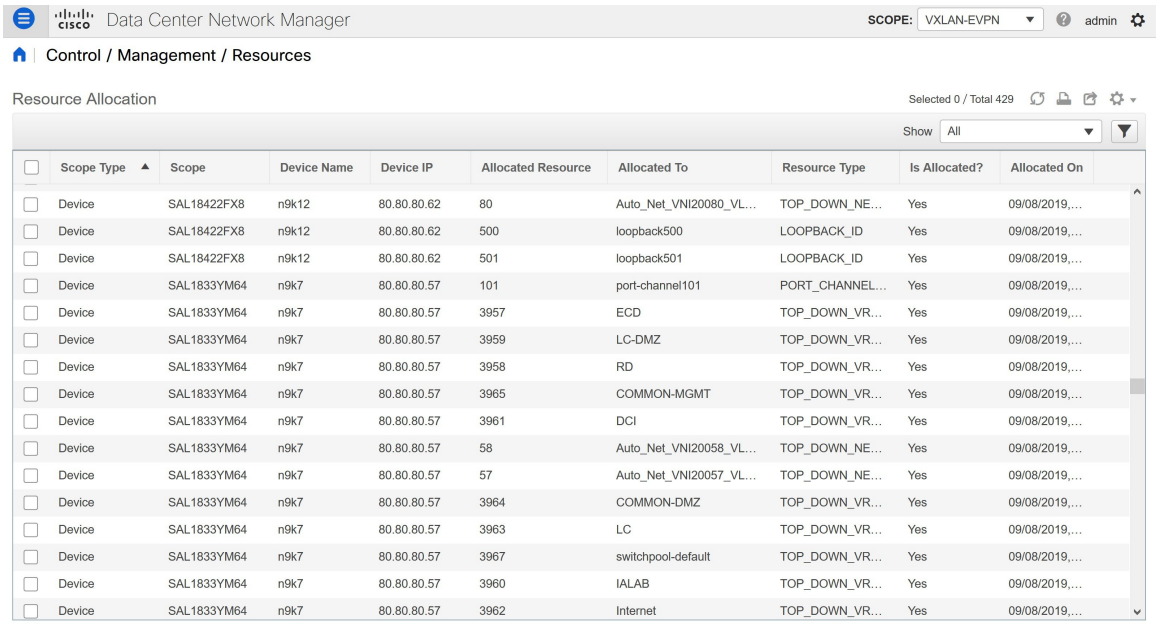

The resources that are being utilized by the VXLAN EVPN fabric such as VLAN IDs, port channel IDs, point to point IP addresses, and loopback IDs are displayed in this window.

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.

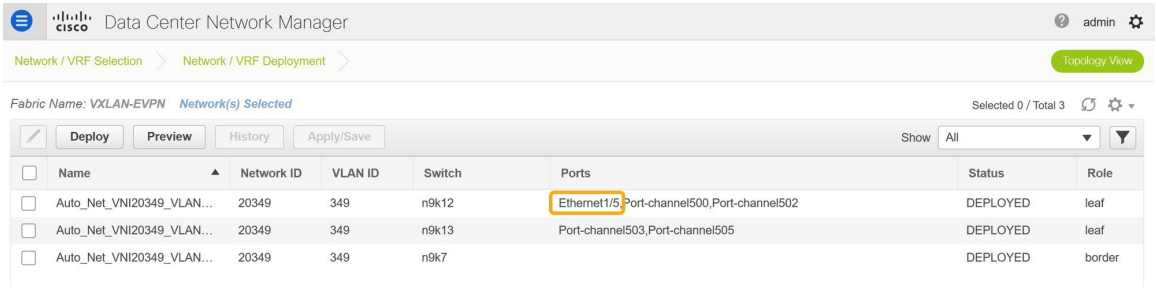

Let us verify the overlay network associated with this interface.

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**.

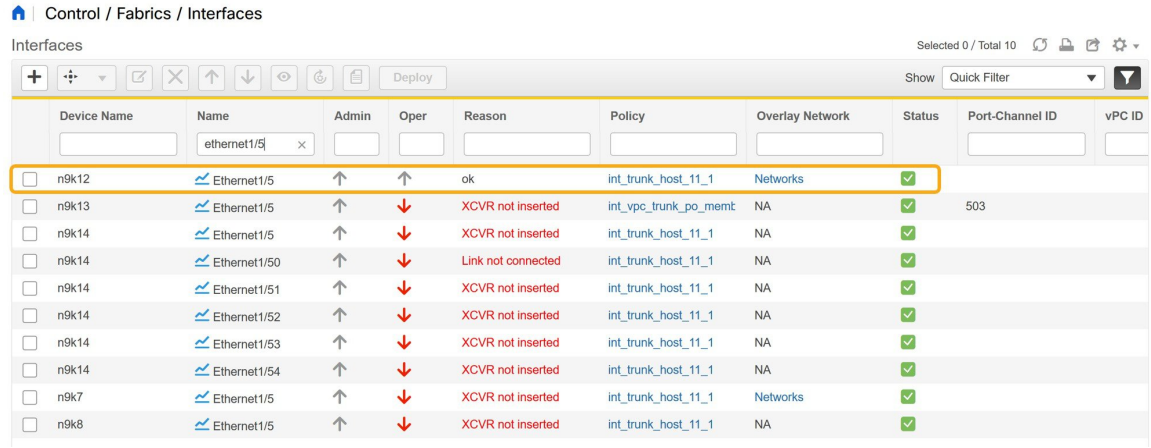

This interface is attached to the host through the **n9k-12 switch**.

Step 7 In the **Overlay Networks** column, click **Networks** that corresponds to the n9k-12 switch and the Ethernet 1/5 interface.

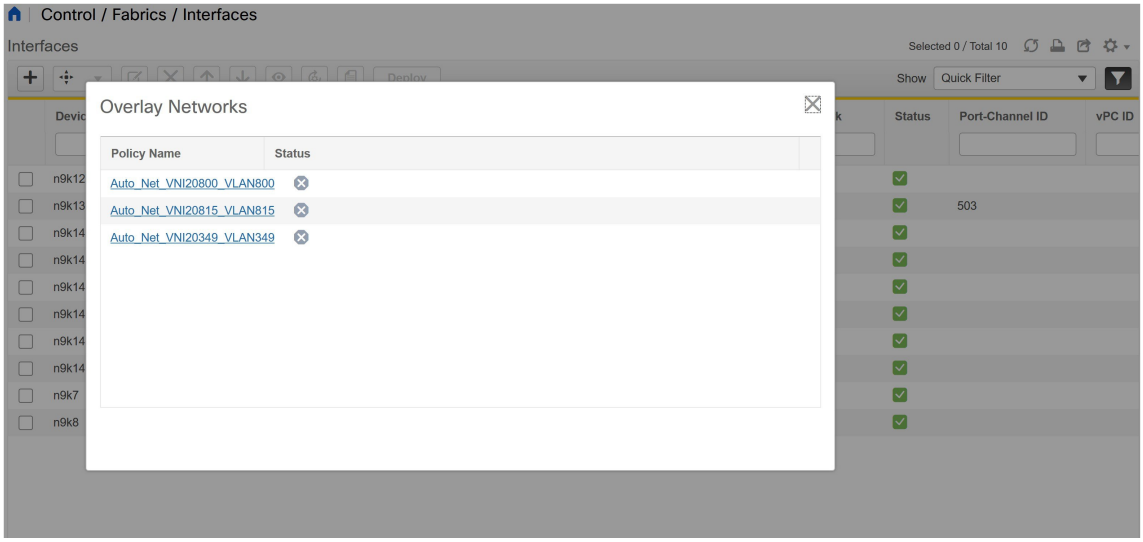

These are the networks that are attached to the **Ethernet 1/5** interface.

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.

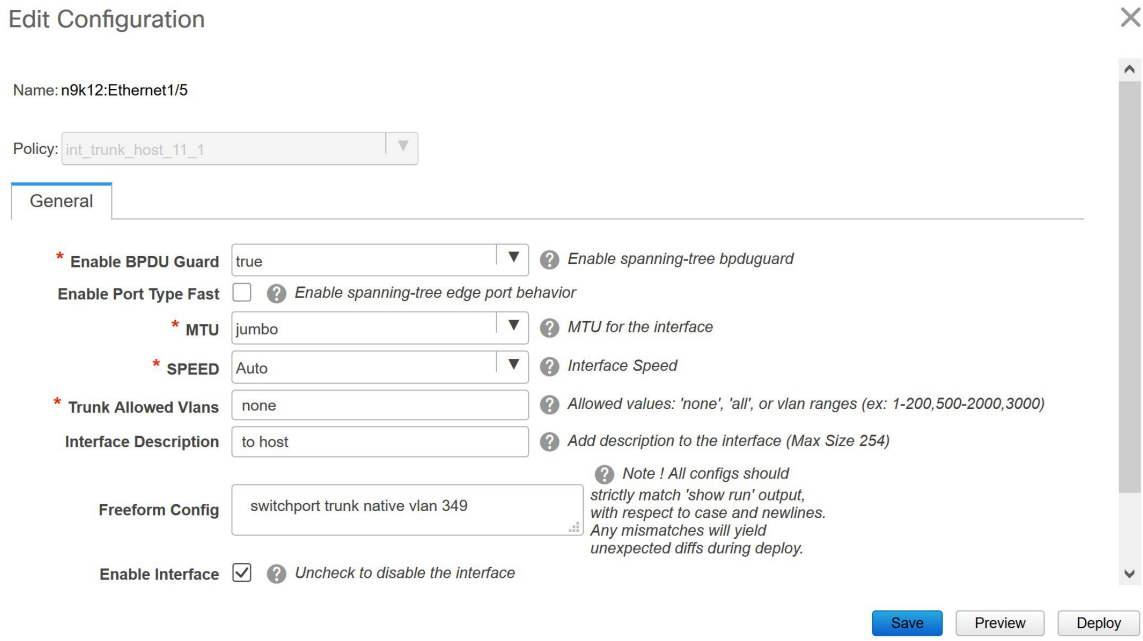

You can see that all the settings for this interface have been successfully imported, including the BPDU guard settings and the interface description.

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.

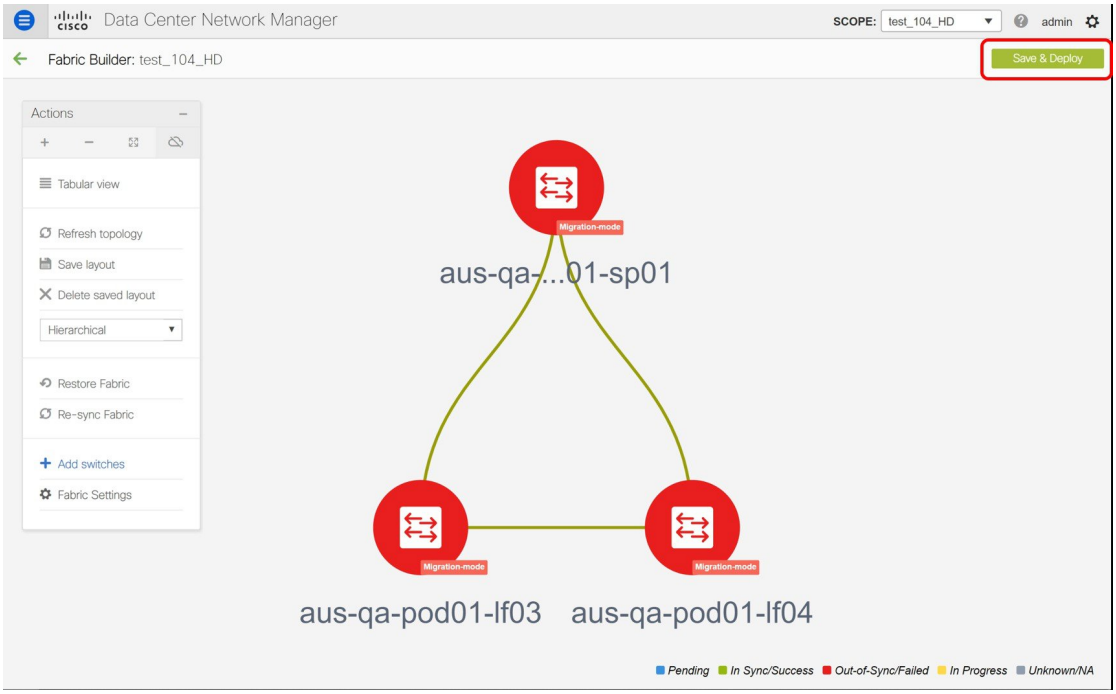

If the added switches contain bottom-up configurations, an error is displayed saying – Reloading switch to clean up bottom up config. Please wait for switch to come online and try **Save & Deploy**.

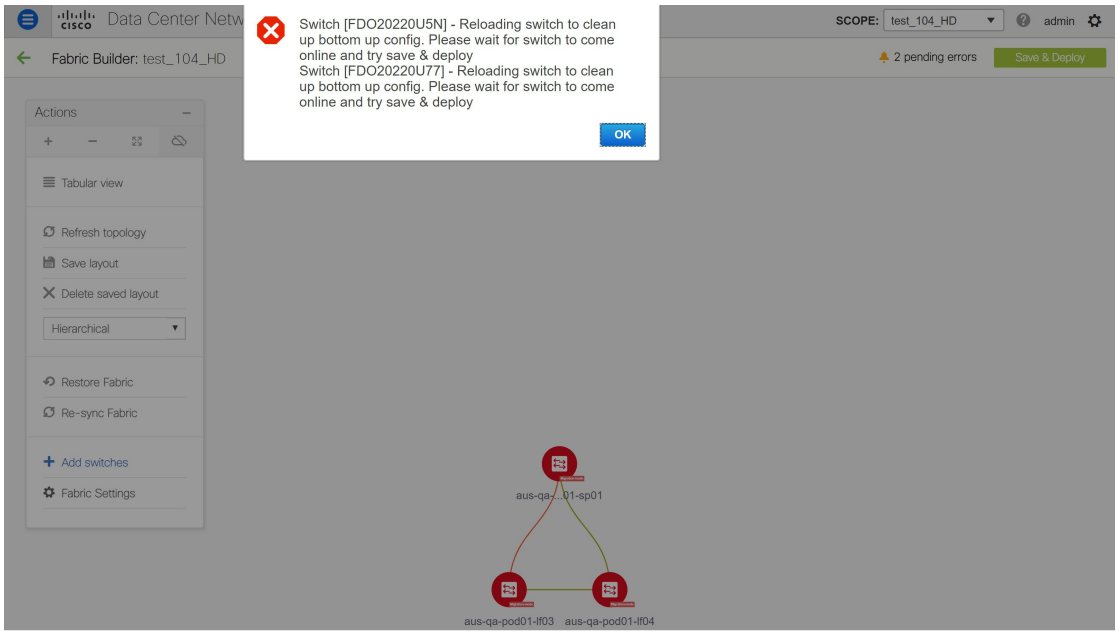

4. Wait for the switches to complete the reload operation. Click **Tabular view** under the **Actions** menu to view the status of the switches.

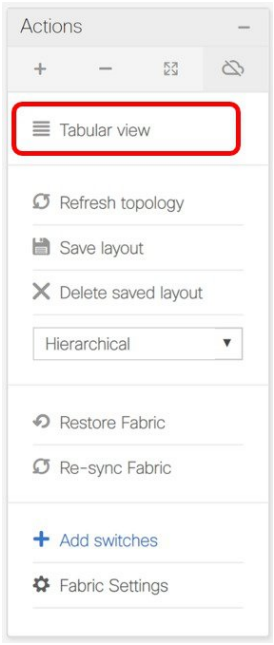

5. (Optional) Rediscovery of the reloaded switches occurs every 5 minutes. If you want to manually rediscover switches, select the switches and click the **Rediscover switch** icon.

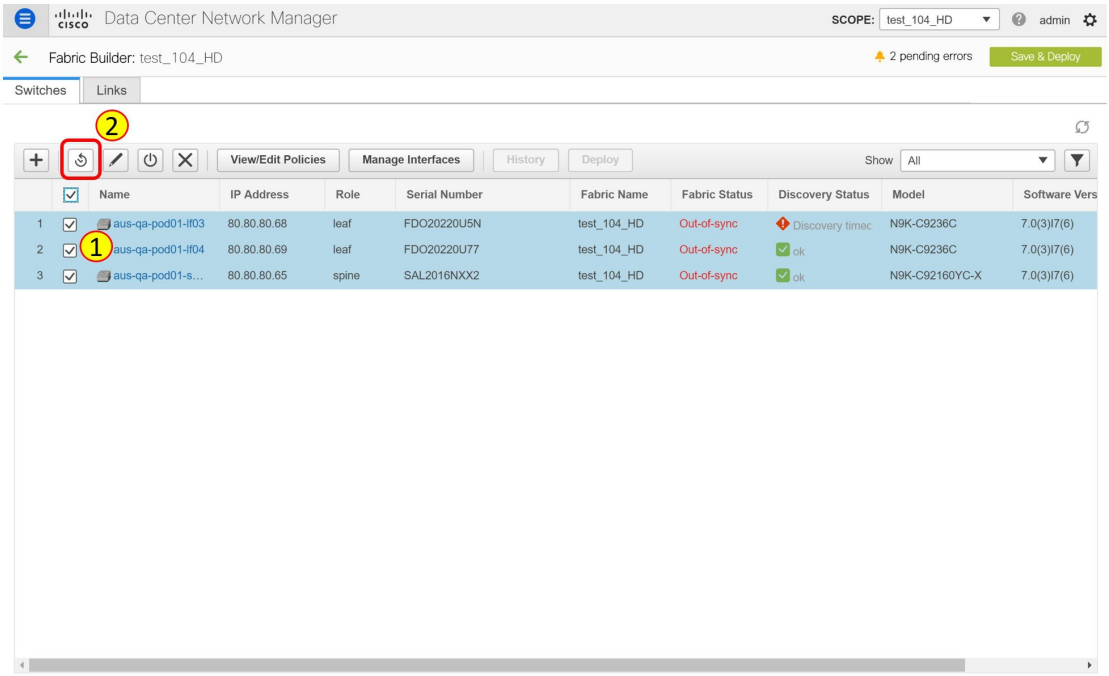

Note Click the **Refresh** icon to refresh the **Fabric Builder** window and see the updated discovery status of switches.

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 in other columns. For example, if the switch was in **RUNNING** tracker status before it becomes unreachable, the value under the **Tracker Status** column for this switch will still be **RUNNING** despite the switch being in **Unreachable** discovery status.

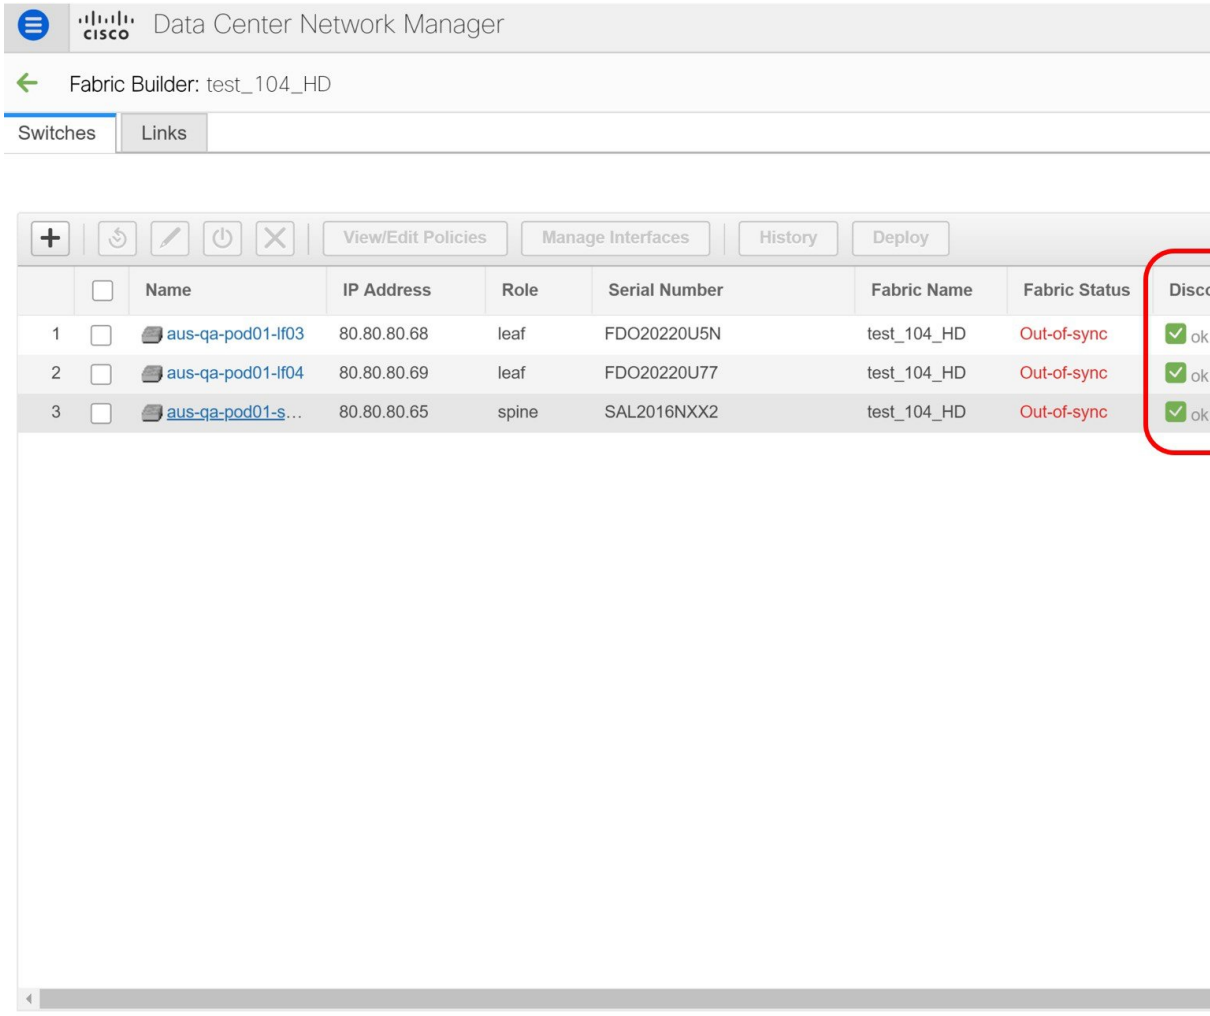

7. Click **Save & Deploy** again to sync configurations between the switches and DCNM.

   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.

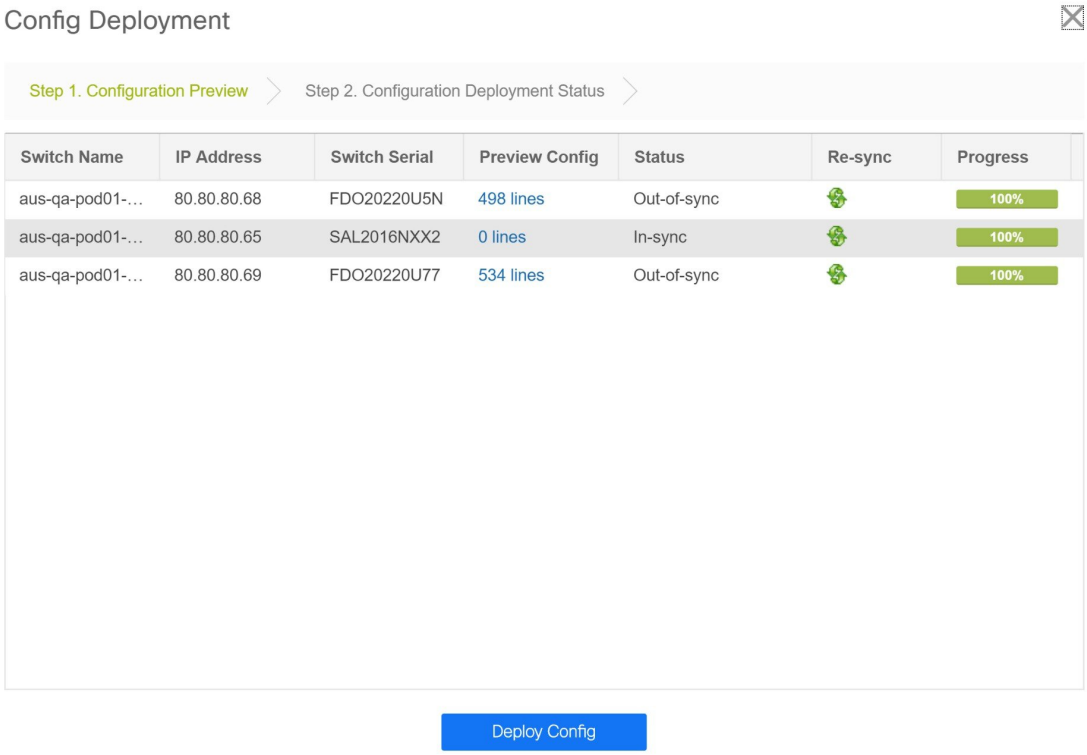

The **Preview Config** column is updated with entries denoting a specific number of lines.

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.

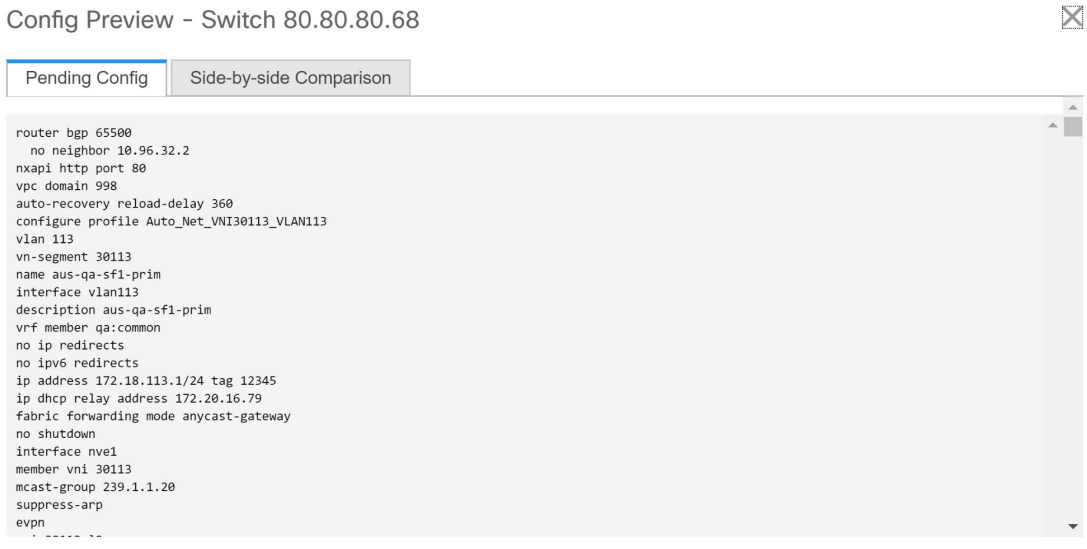

Close the **Config Preview** window.

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.

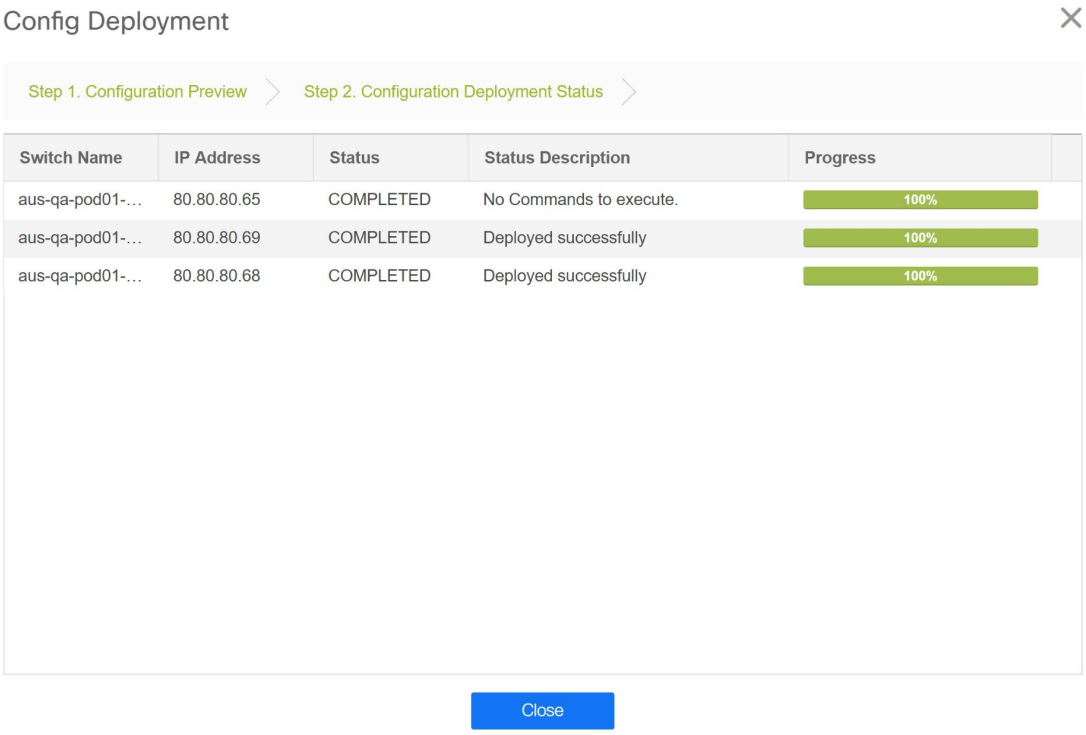

The progress bar shows 100% for each switch. After correct provisioning and successful configuration compliance, close the **Config Deployment** window.

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.

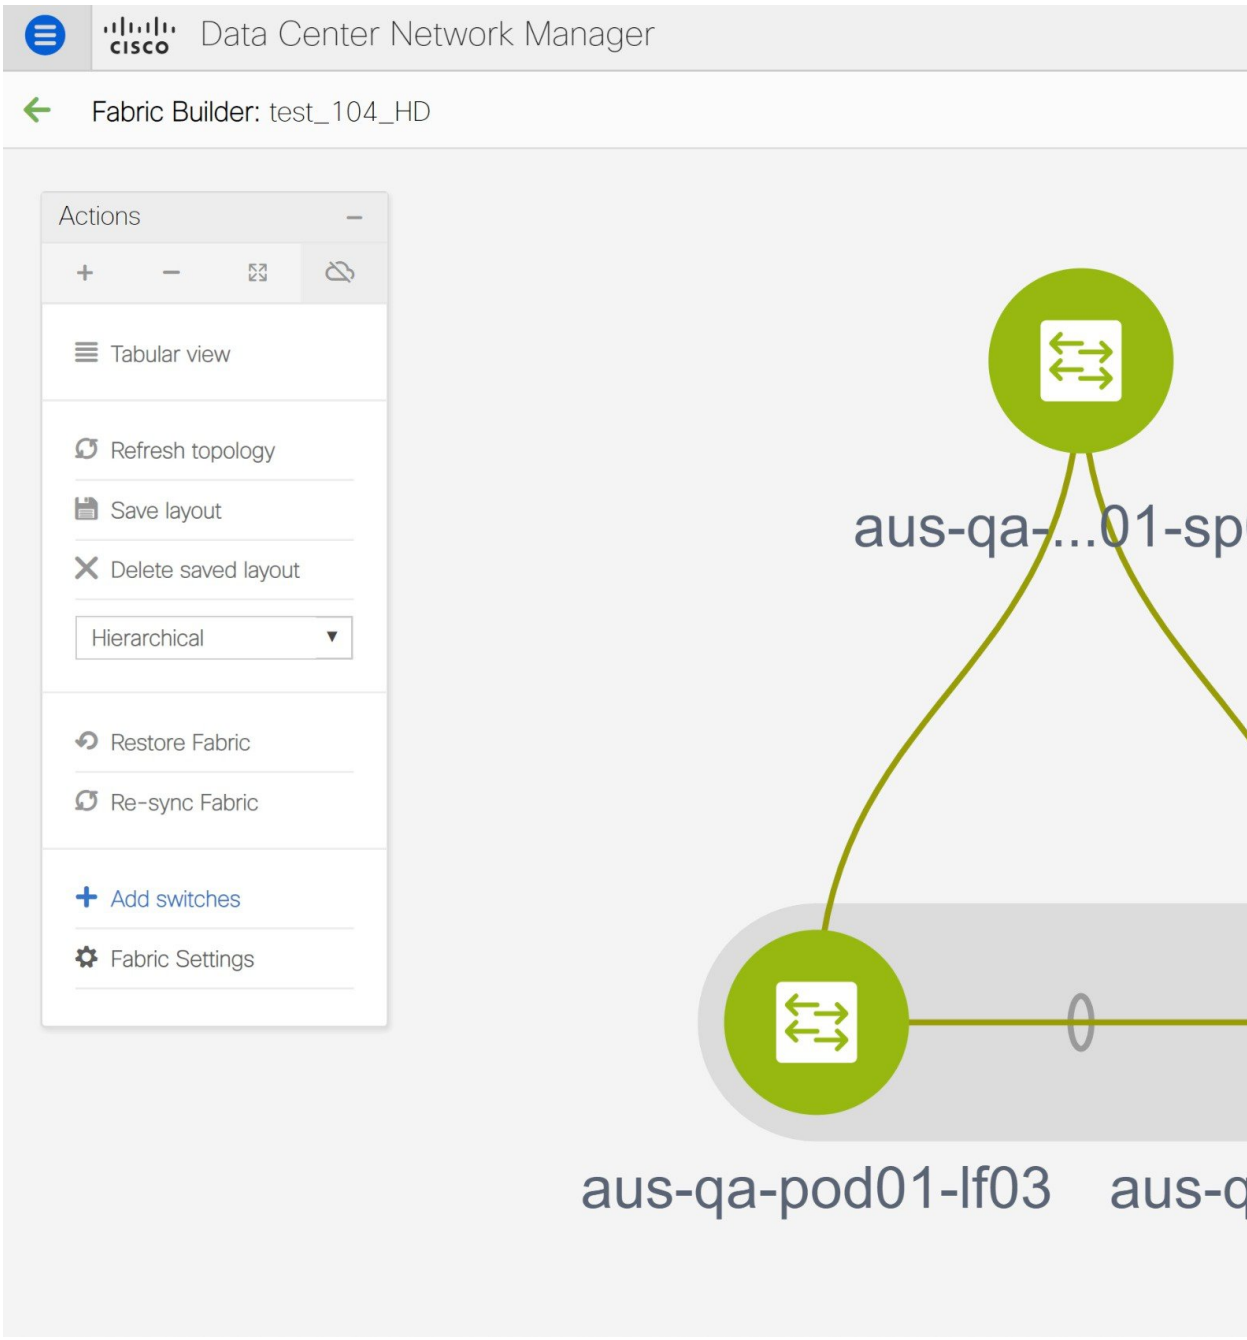

This completes the migration process of bottom-up VXLAN fabric to DCNM.

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**.

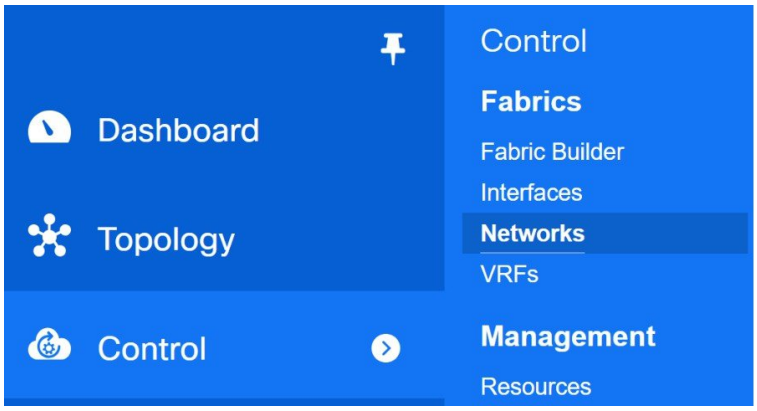

2. Select the fabric from the **SCOPE** drop-down list in the **Networks** window.

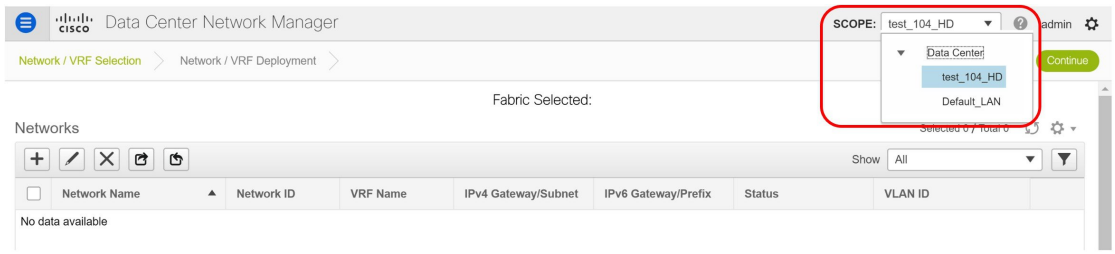

3. Check the networks that are migrated from the bottom-up VXLAN fabric and their deployment status.

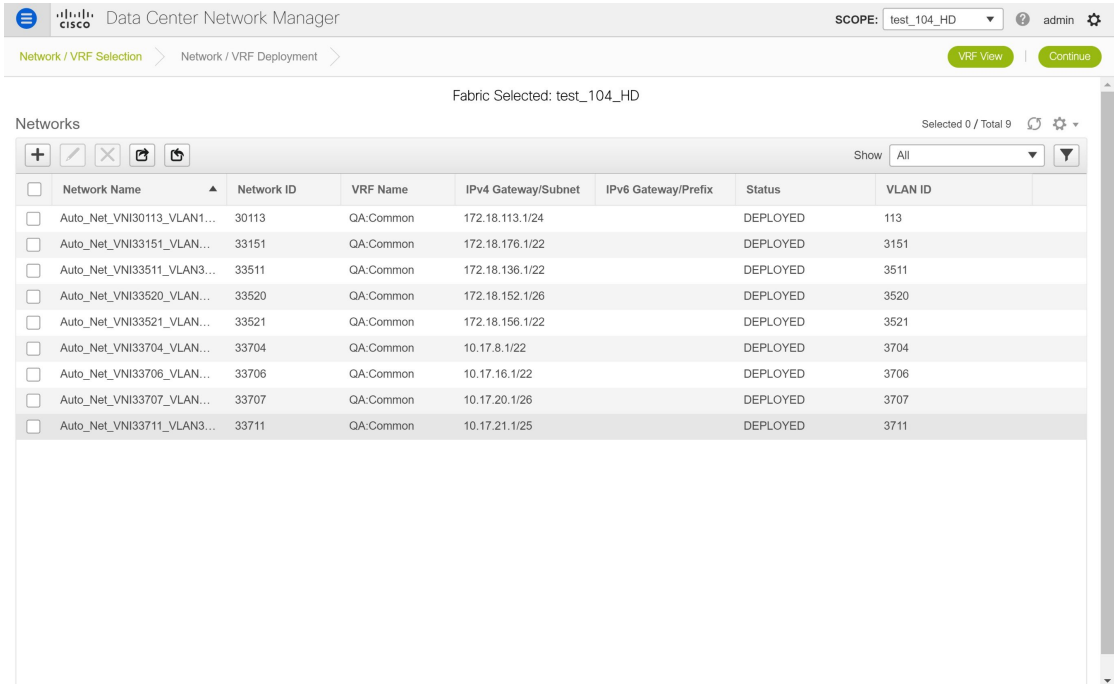

Resolving Config Compliance Error on Switches with Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) Images

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.

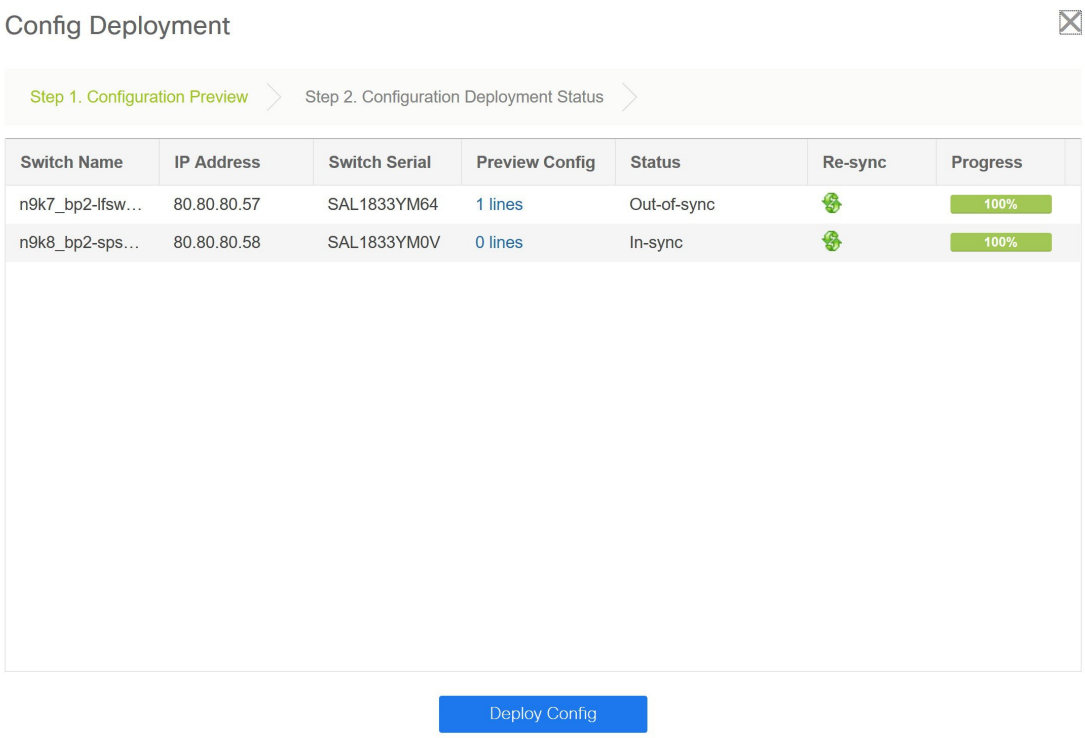

4. (Optional) Click the entry showing **1 lines** under the **Preview Config** column.

   You can see the TCAM command under the **Pending Config** tab in the **Config Preview** window.

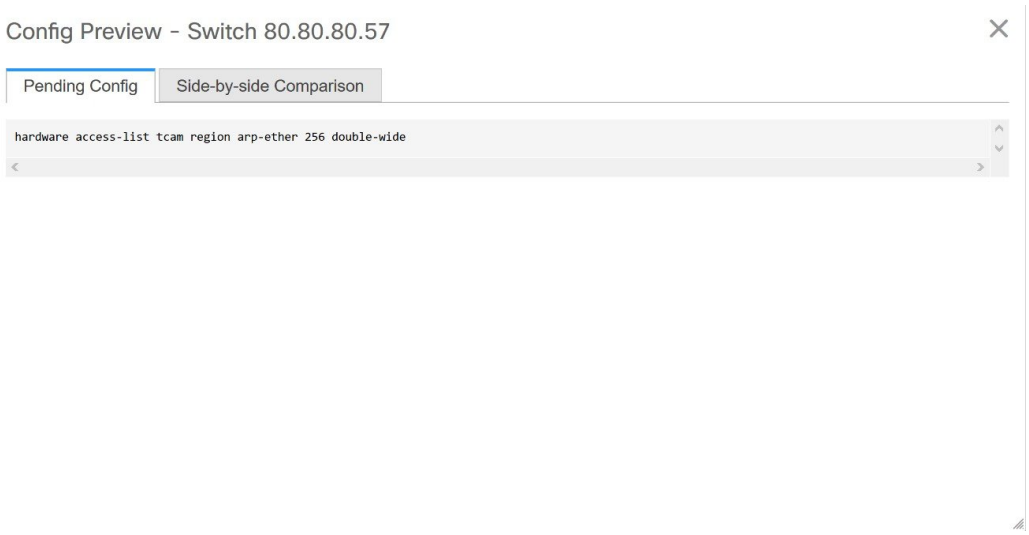

Close the **Config Preview** window.

5. Right-click a switch and click **View/Edit Policies**.

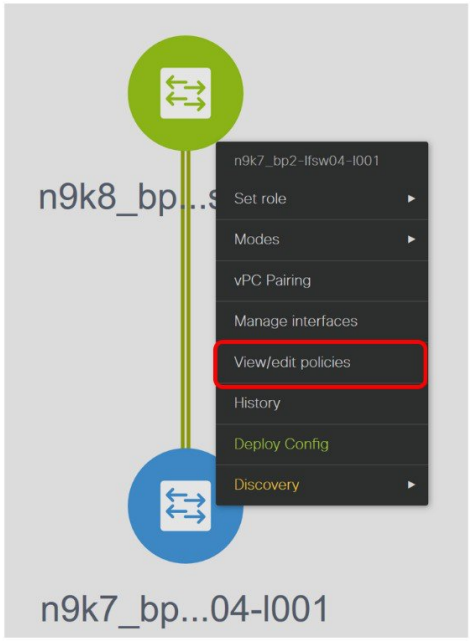

6. Search for the **tcam_pre_config_vxlan** policy in the **Template** search field.
7. Select the **tcam_pre_config_vxlan** policy and click the **Delete** icon to delete the policy.

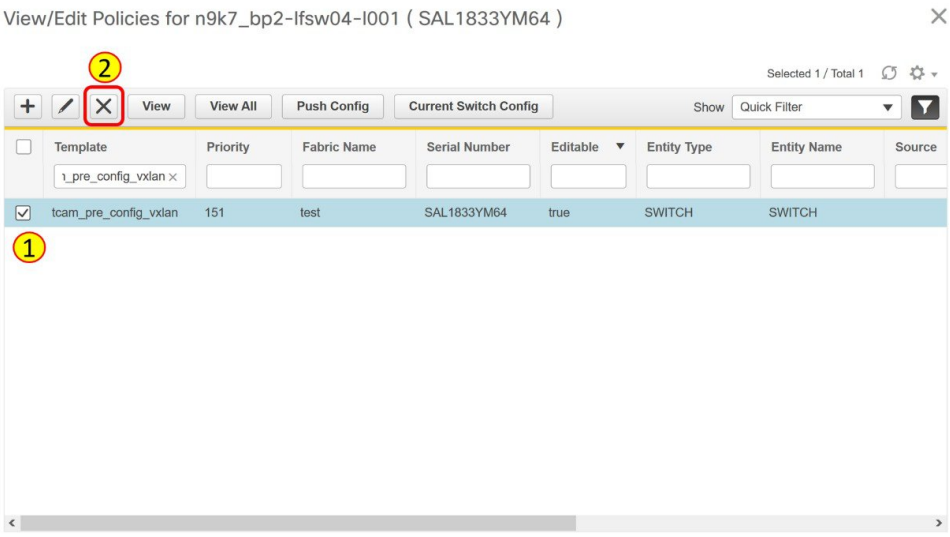

Close the **View/Edit Policies** window.

8. (Optional) Click **Save & Deploy** to verify whether there are any pending configs.

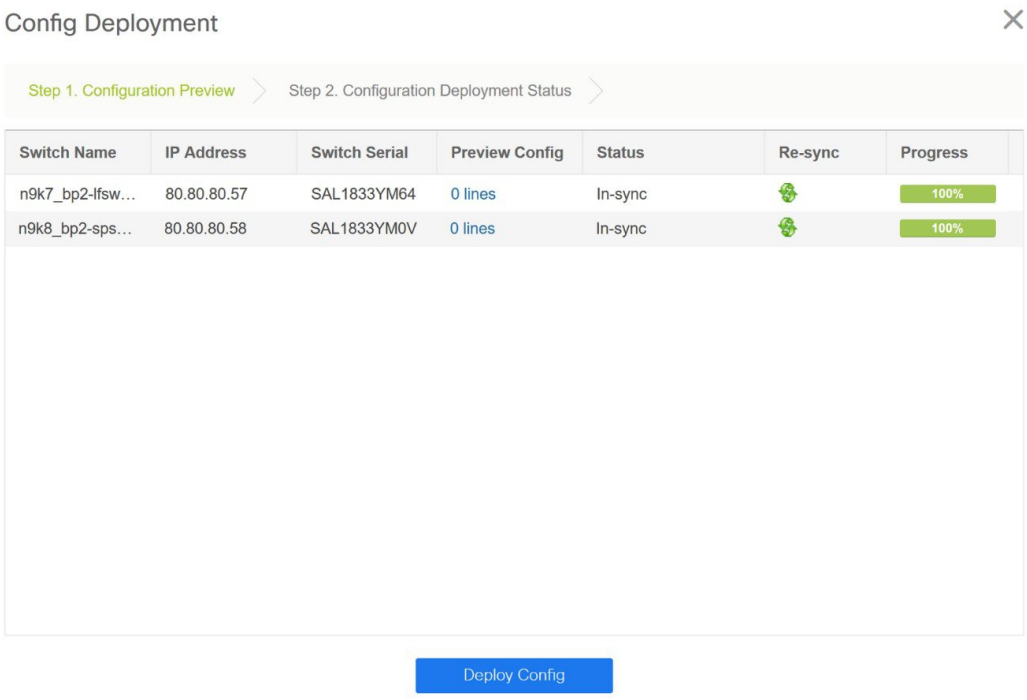

Resolving Config Compliance Error on Switches for RMA, and Write Erase and Reload Operations

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.

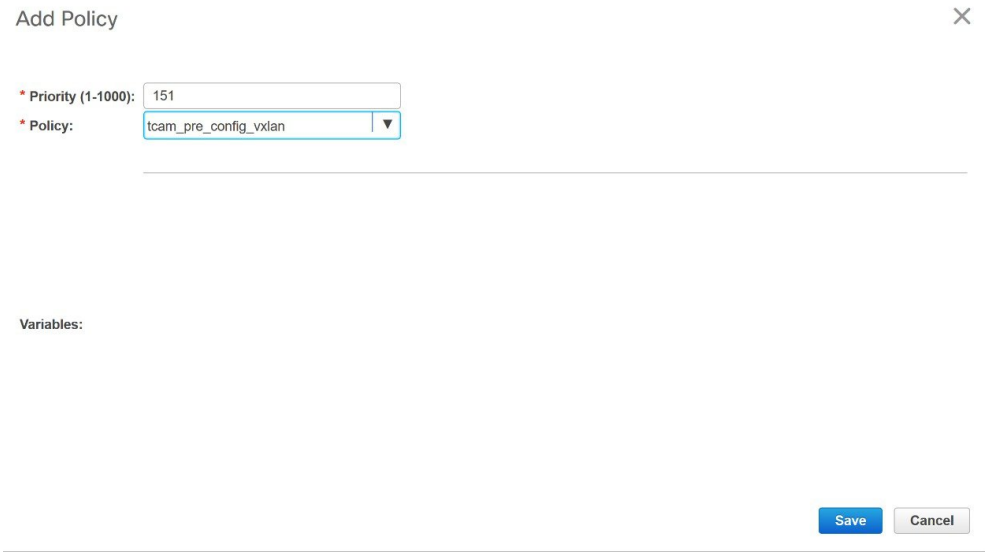

6. Click **Save**.
7. Complete the RMA or Write Erase and Reload operation.

   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.

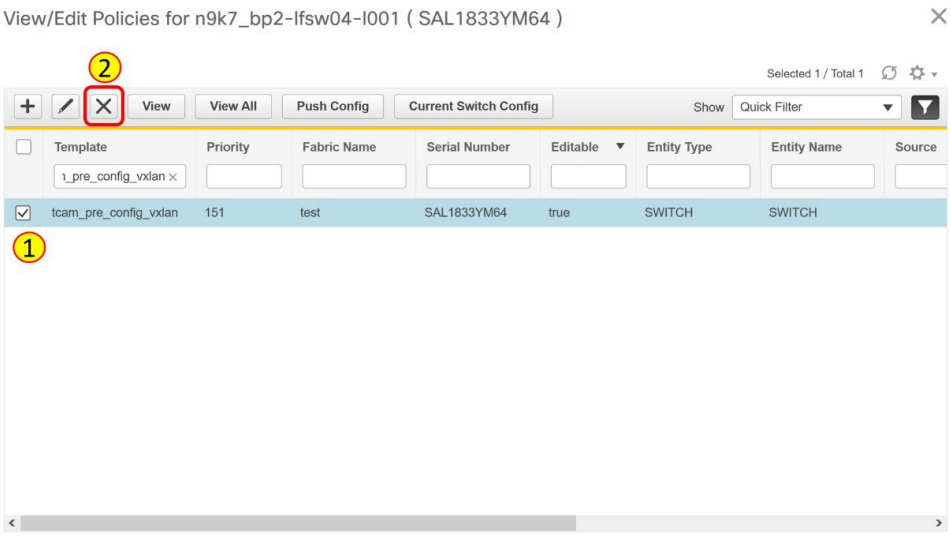

Close the **View/Edit Policies** window.

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.

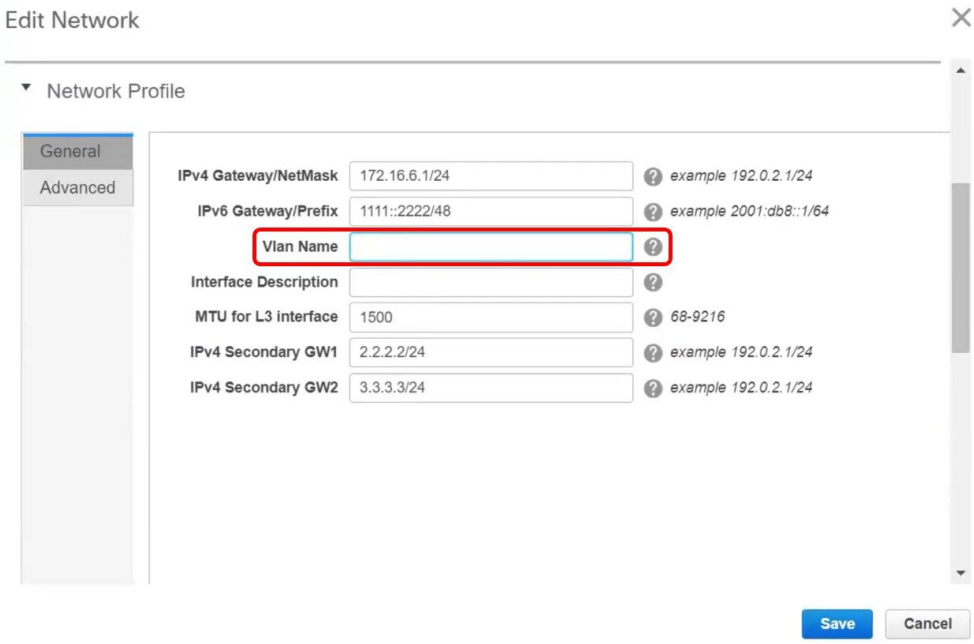

In the **Edit Network** window, the **Vlan Name** field is empty because DCNM has not captured this info in the overlay profile. Instead, the VLAN name is captured in the freeform config associated with the overlay network or VRF.

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.

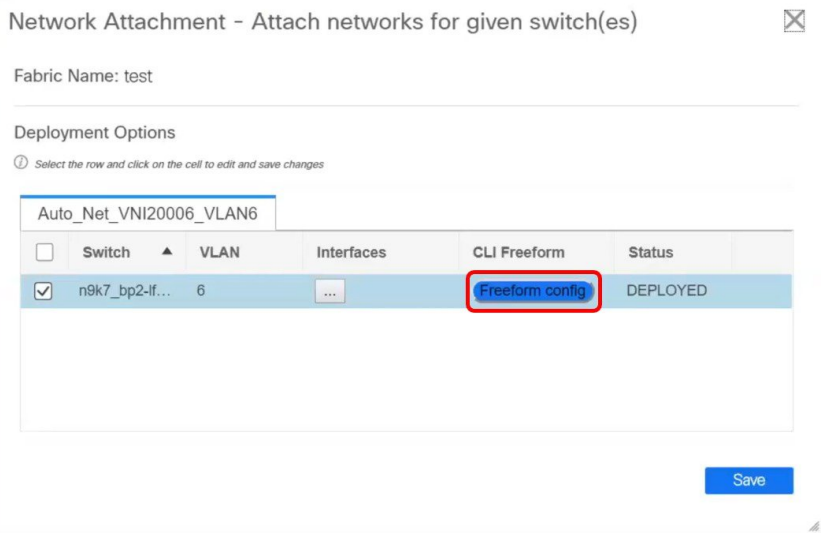

Step 7 Verify the VLAN name in the **Free Form Config** window.

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
.
.
.
```

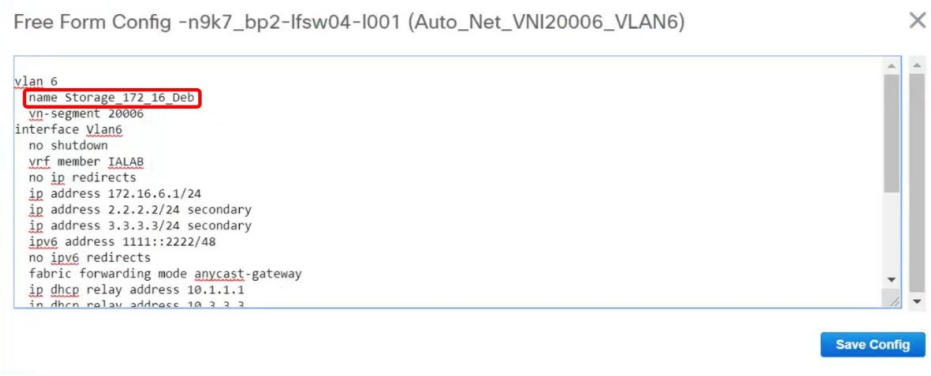

Step 9 Click **Save** in the **Network Attachment** window.

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:

- 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**.

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

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

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

1. Import the switches into all the required fabrics and set roles accordingly.
2. Click **Save & Deploy** on each of the fabrics, but do not click **Deploy Config**.
3. Create an MSD. The **Multi-Site Underlay IFC Auto Deployment Flag** can be enabled or disabled. For more information, see *Creating an MSD Fabric in Cisco DCNM LAN Fabric Configuration Guide*.
4. Ensure that the fabric settings for MSD are correct including settings such as profile selection, the multisite loopback ID, and anycast GW MAC.
5. Move the member fabrics into the MSD. For more information, see *Moving the Member1 Fabric Under MSD-Parent-Fabric in Cisco DCNM LAN Fabric Configuration Guide*.

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

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

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

   If **Multi-Site Overlay IFC Deployment Method** is set to **Direct_To_BGWS**, then overlay IFCs are created after brownfield migration, and associated with appropriate **MULTISITE_OVERLAY** policy.

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

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

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

7. Edit the 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.

8. Navigate to each member fabric, click **Save & Deploy**, and then click **Deploy Config**.
9. To verify, you can select networks or VRFs and corresponding BGWs, and see the expected configurations. You can now manage all the networks or VRFs for BGWs by using the regular top-down workflow.

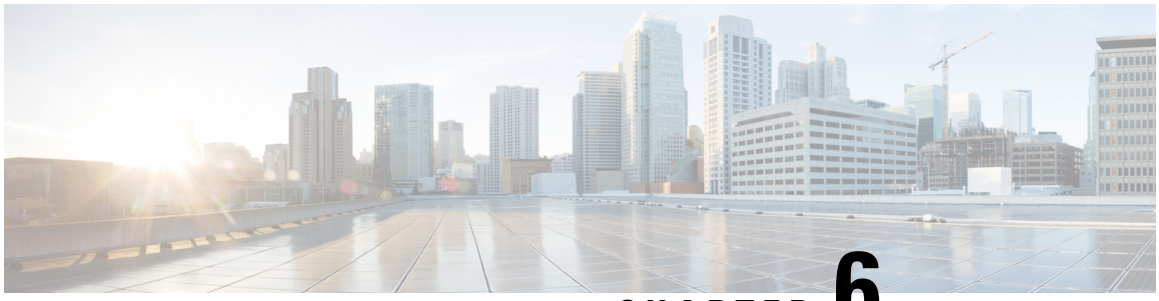

CHAPTER 6

Layer 4-Layer 7 Service

• Layer 4-Layer 7 Service, on page 135

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.

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 switch. 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. 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

You can only define the service policy between the created networks. The L4-L7 service does not create any VRF or network other than the service networks that are defined during route peering. 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. Note that the source or destination network can also be defined by using the **any** CLI keyword. This means that any IP address traffic is supported. 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.

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.
- This feature is supported only on VXLAN BGP EVPN fabrics.
- 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, or border super spine 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.
- Load sharing is not supported.
- This feature does not support Multi-Site Domains (MSD).
- This feature creates, updates, and deletes the service network, as required. Service networks cannot be created or deleted from the **Control > Fabrics > Networks** window.

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.

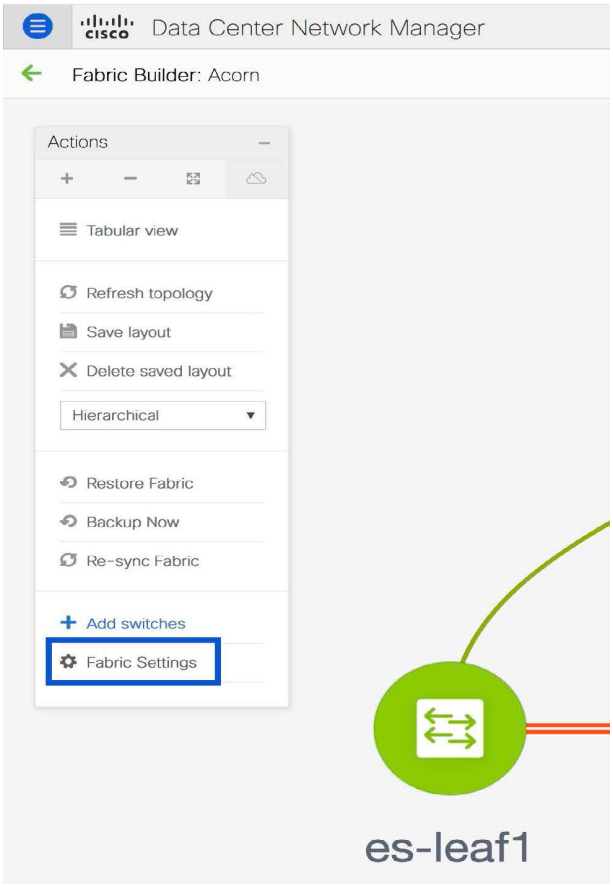

The **Edit Fabric** window is displayed. Click **Advanced**. Select the **Enable Policy-Based Routing (PBR)** checkbox to enable routing of packets based on the specified policy.

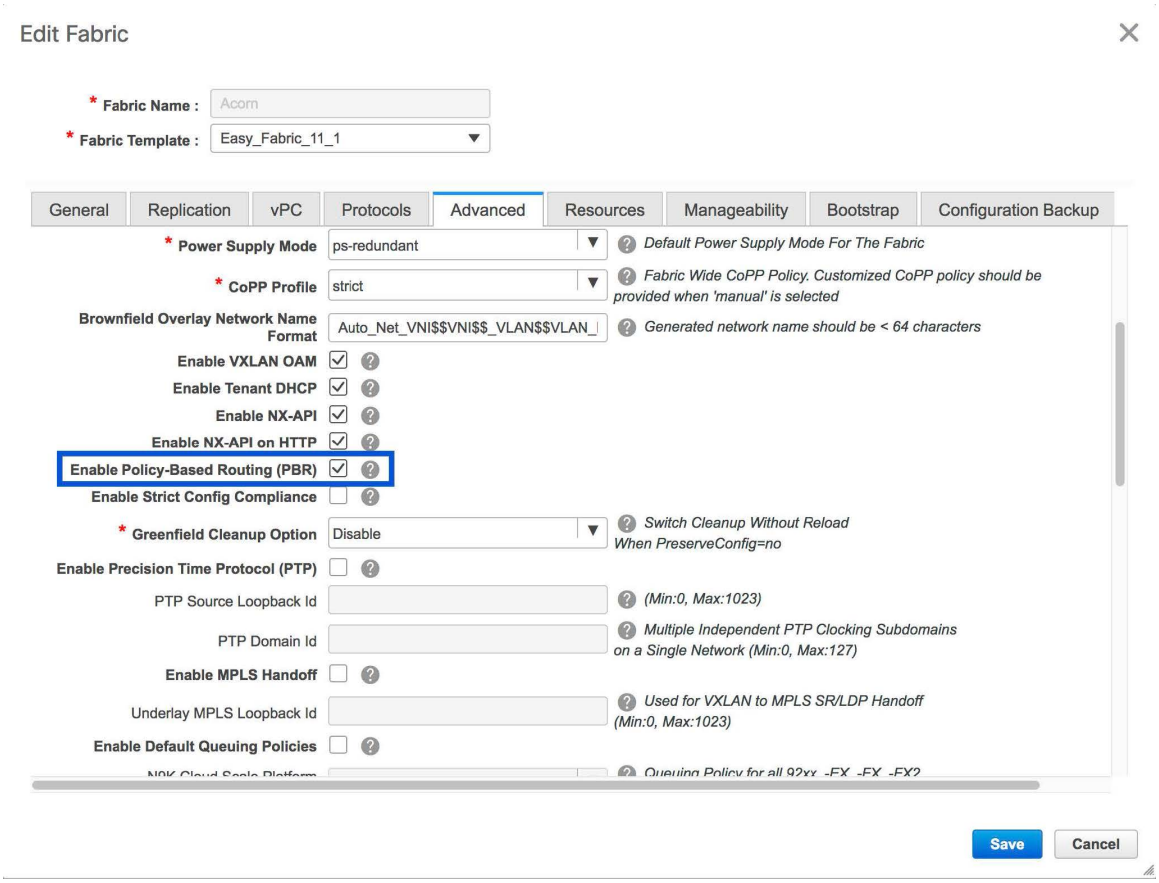

Now, click **Resources**. Specify a VLAN range in the **Service Network VLAN Range** field. This is a per switch overlay service network VLAN range. The minimum allowed value is 2 and the maximum allowed value is 3967. Also, specify a value for the **Route Map Sequence Number Range** field. The minimum allowed value is 1 and the maximum allowed value is 65534. Click **Save and Deploy** to deploy the updated configuration.

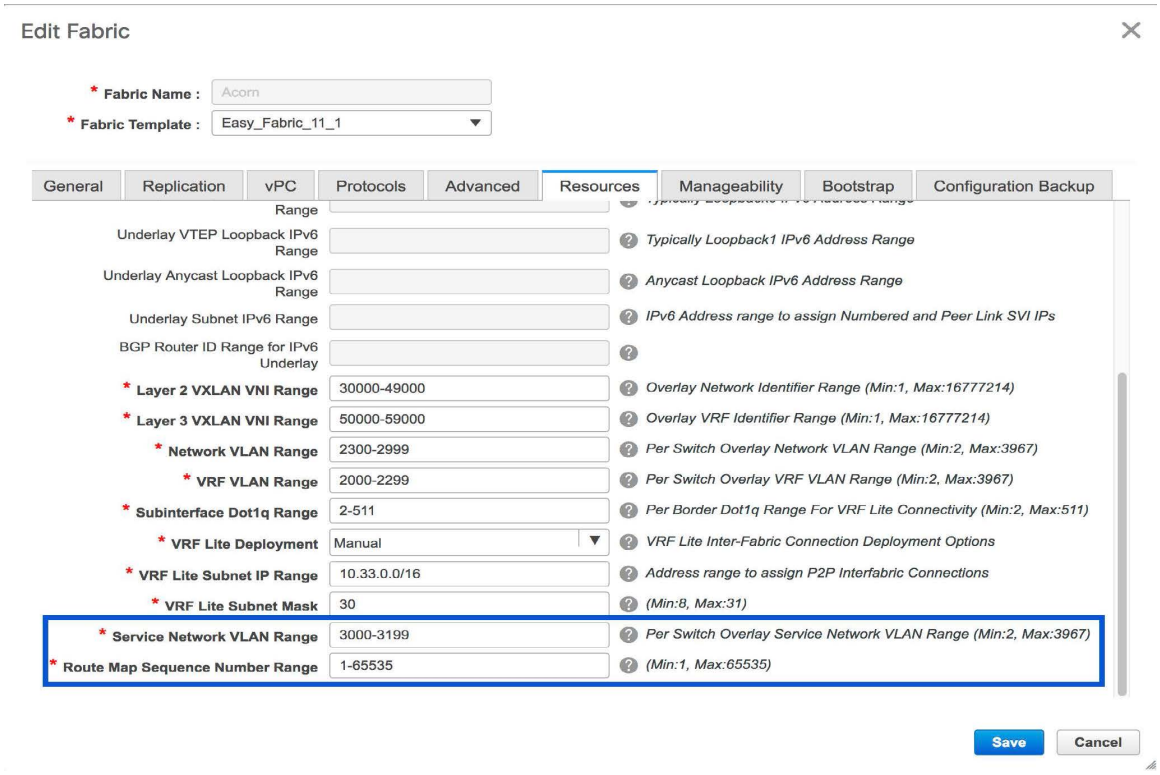

Configuring Layer 4-Layer 7 Service

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.

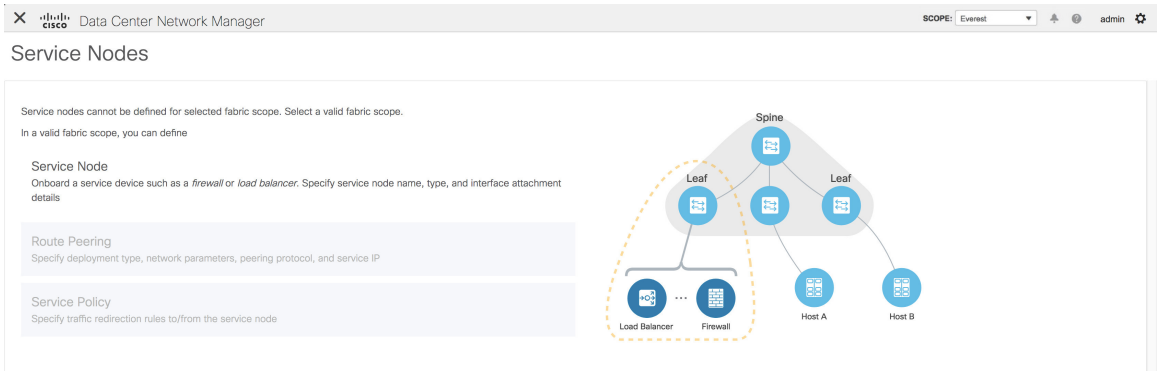

The L4-L7 service configuration procedure consists of the following steps:

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.

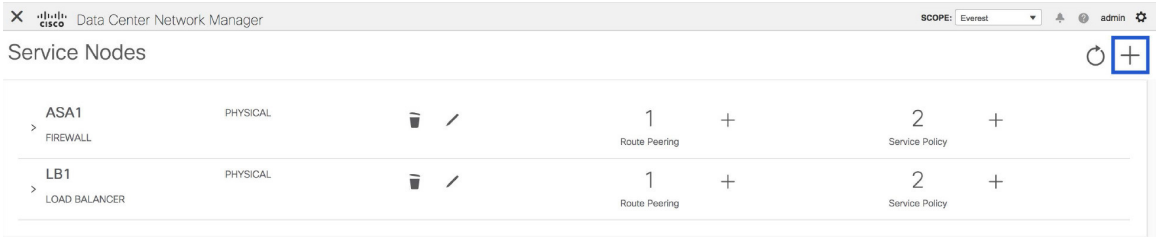

The **New Service Nodes** window has three steps, **Create Service Node, Create Route Peering** and **Create Service Policy**.

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 2: Example: Link Template - service_link_trunk

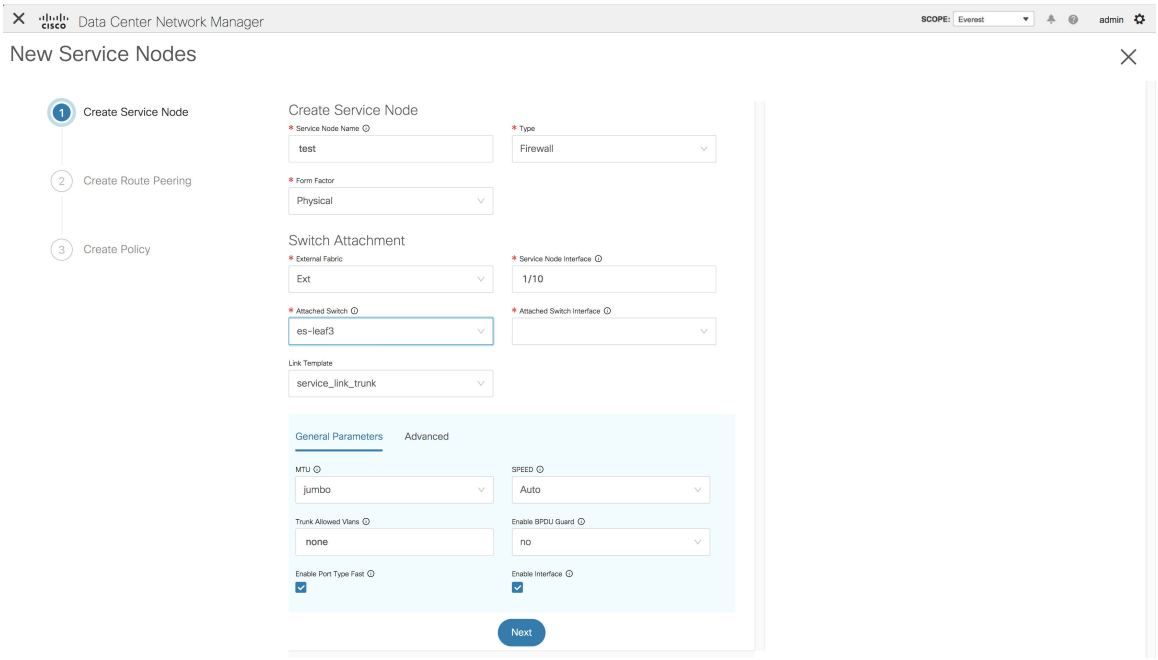

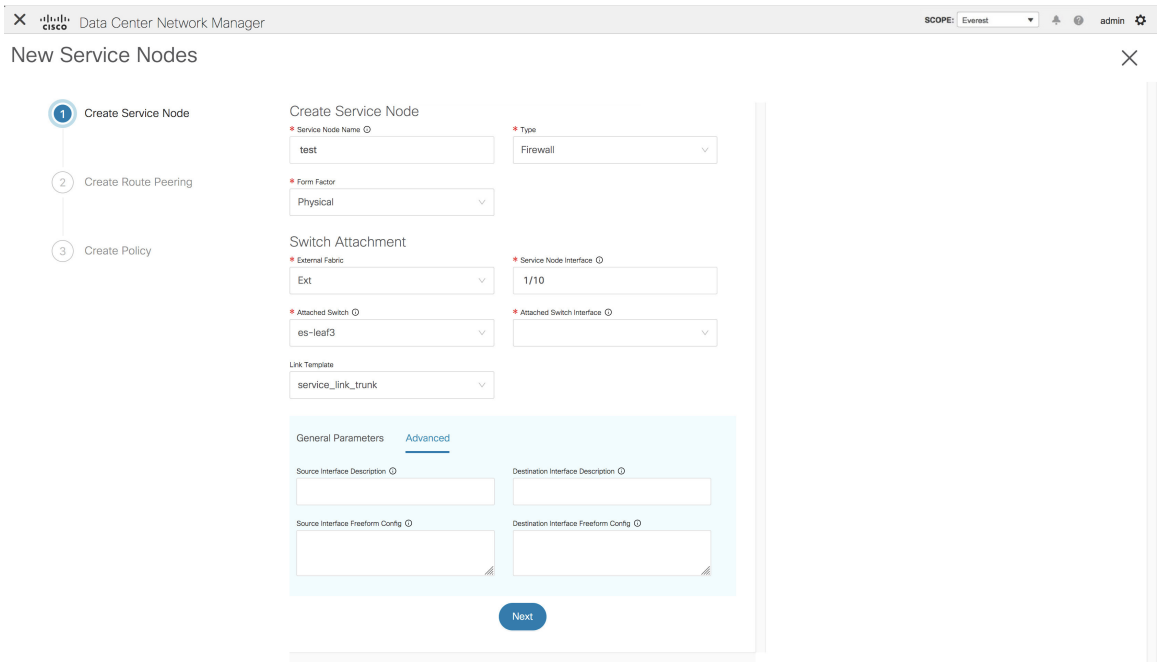

Figure 3: Example: Link Template - service_link_port_channel_trunk

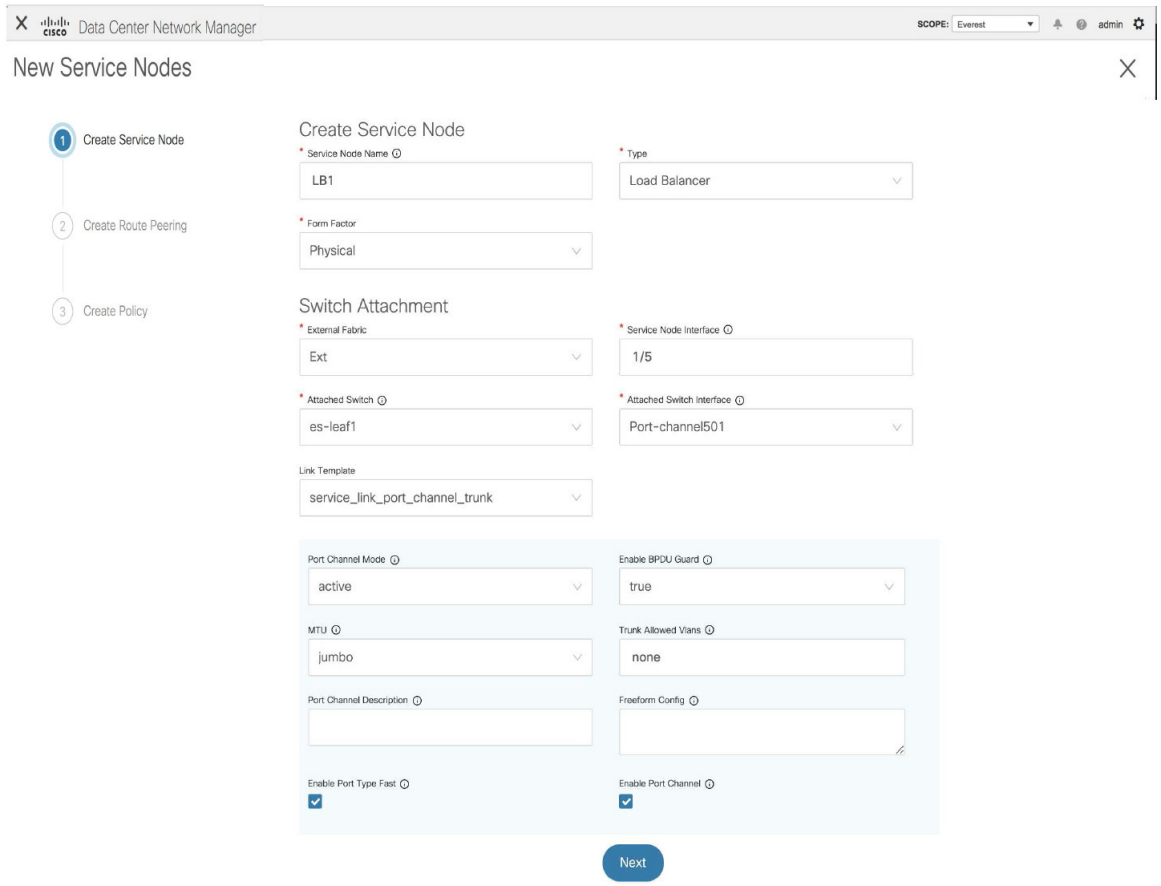

Figure 4: Example: Link Template - service_link_vpc

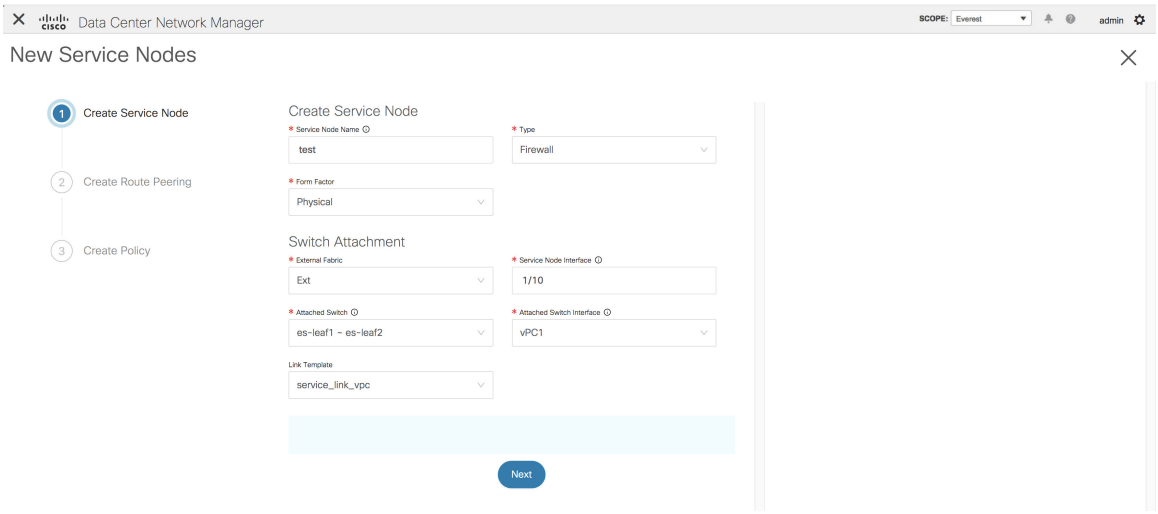

The fields in the **Create Service Node** window are as given below. It is mandatory to fill the fields marked with an asterisk. For more information on the fields in this window, hover over the **i** icon.

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 supported in Top-down provisioning.

Example: Intra-Tenant Firewall Deployment

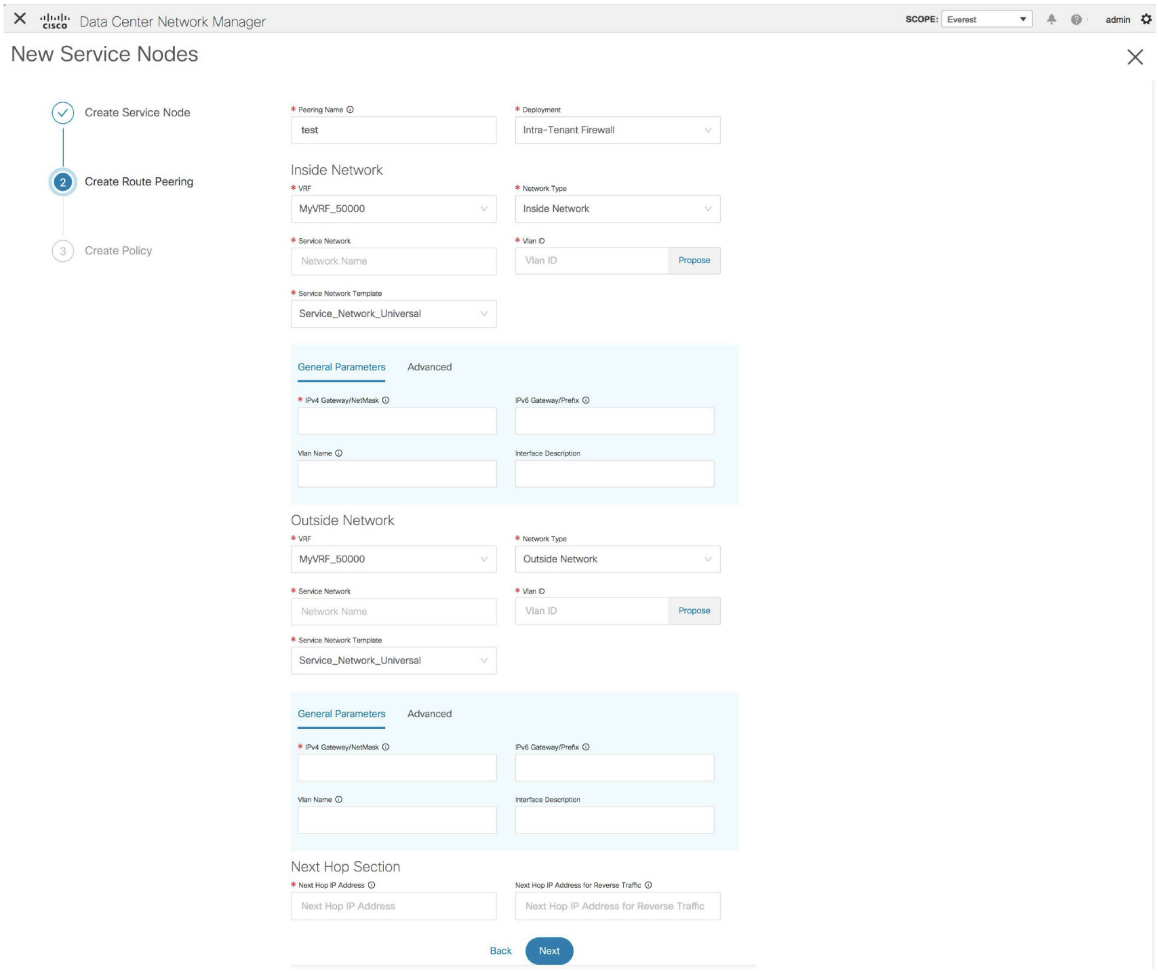

The fields in the **Create Route Peering** window for an Intra-Tenant Firewall deployment are as given below. It is mandatory to fill the fields marked with an asterisk. For more information on the fields in this window, hover over the **i** icon.

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

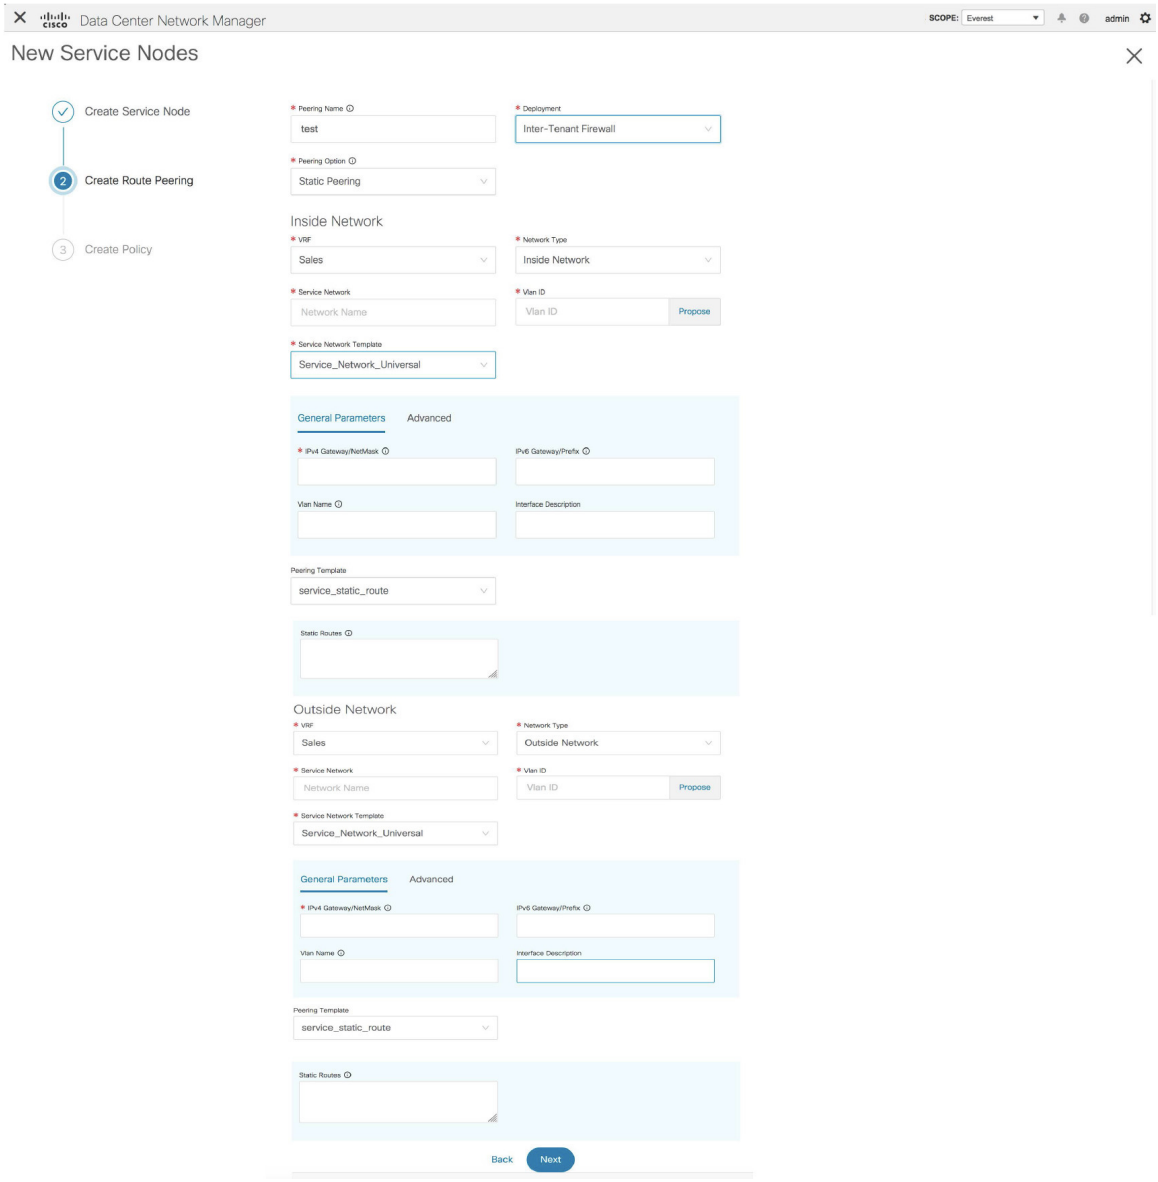

The fields in the **Create Route Peering** window for an Inter-Tenant Firewall deployment are as given below. It is mandatory to fill the fields marked with an asterisk.

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.

Note Inter-tenant firewall deployment with eBGP dynamic peering option is not supported.

Example: One-Arm Mode Load Balancer

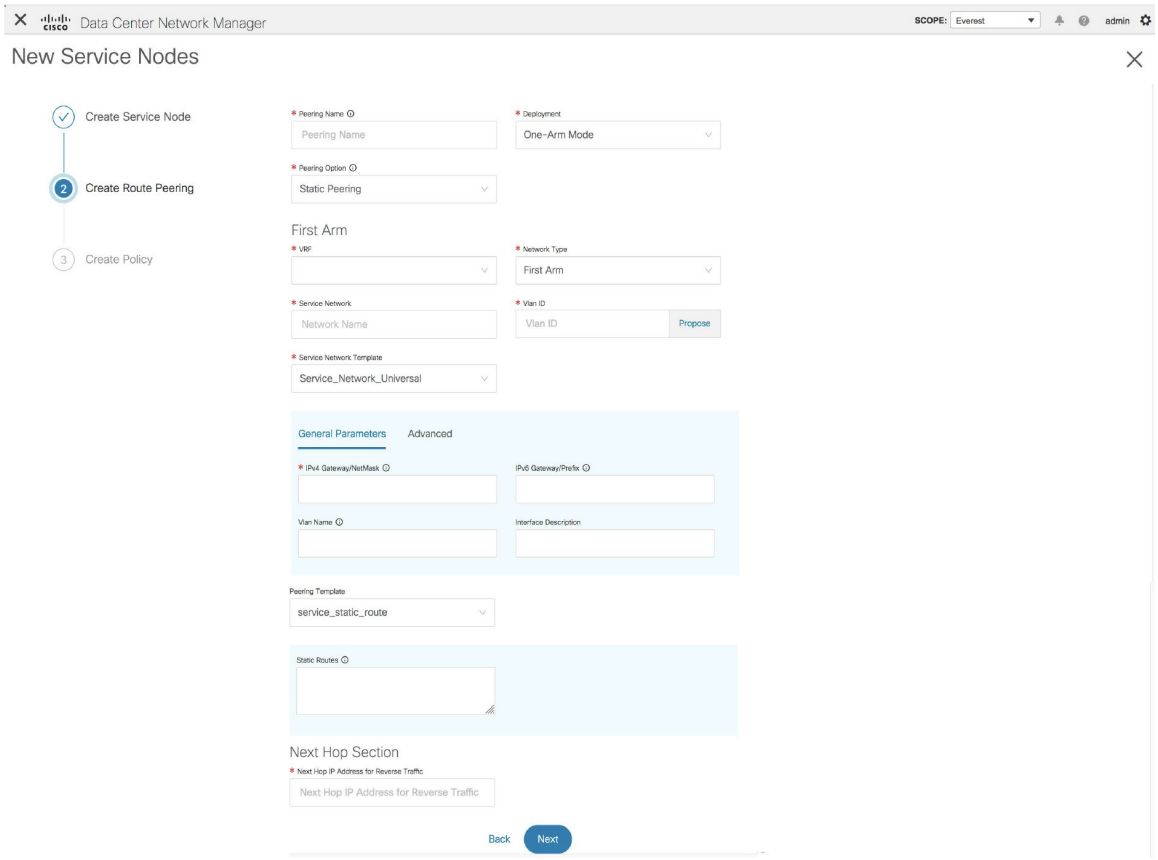

The fields in the **Create Route Peering** window for a One-Arm Mode load balancer deployment are as given below. It is mandatory to fill the fields marked with an asterisk.

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

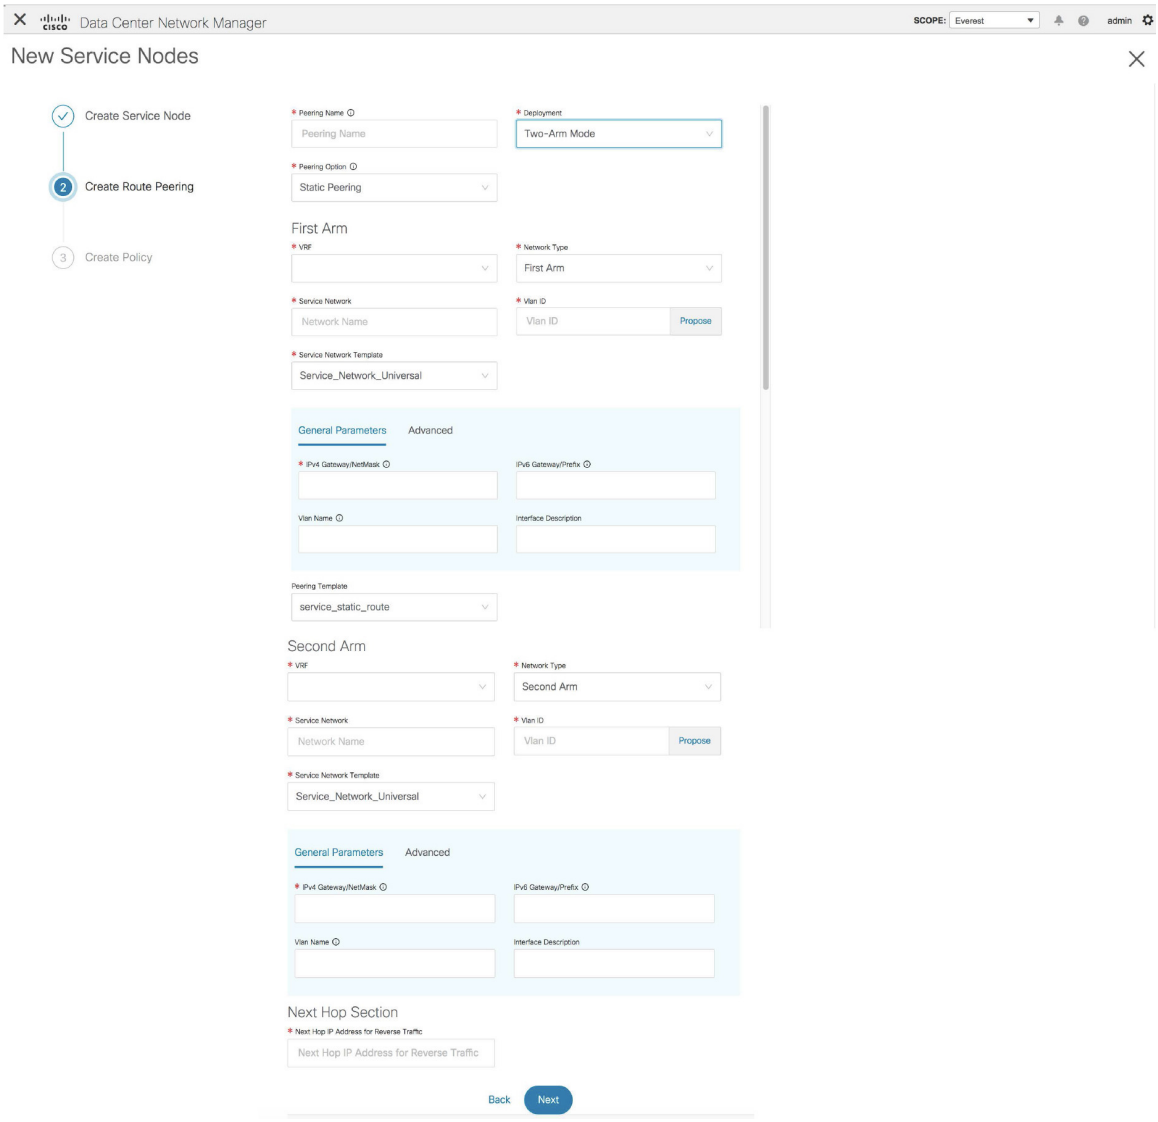

The fields in the Create Route Peering window for a Two-Arm Mode load balancer deployment are as given below. It is mandatory to fill the fields marked with an asterisk.

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.

Create Service Policy

The **Create Policy** window is displayed as given below.

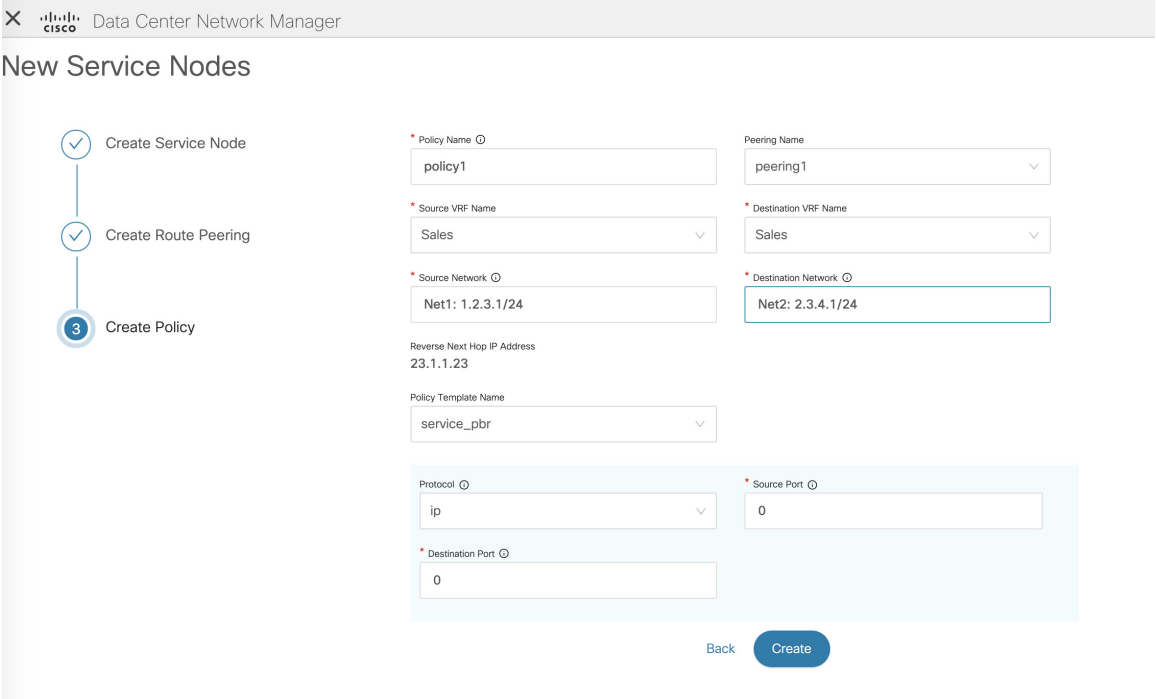

The fields in the **Create Policy** window are as given below. It is mandatory to fill the fields marked with an asterisk.

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.

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.

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.

Enable Interface - Uncheck the checkbox to disable the interface. By default, the interface is enabled.

Service Policy Template

service_pbr

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.

You can also customize the templates based on specific requirements. For more information on templates, refer *Template Library section in Cisco LAN Fabric Configuration Guide*.

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.

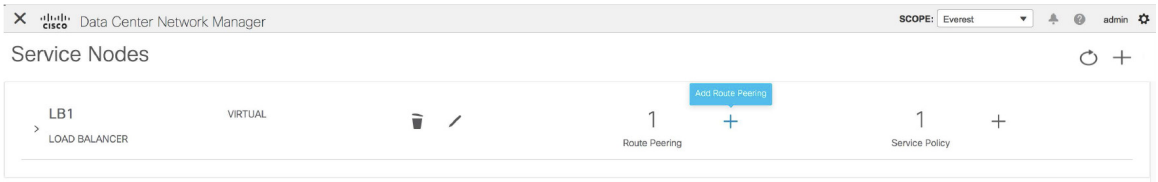

Step 2 The **Add Route Peering** window is displayed.

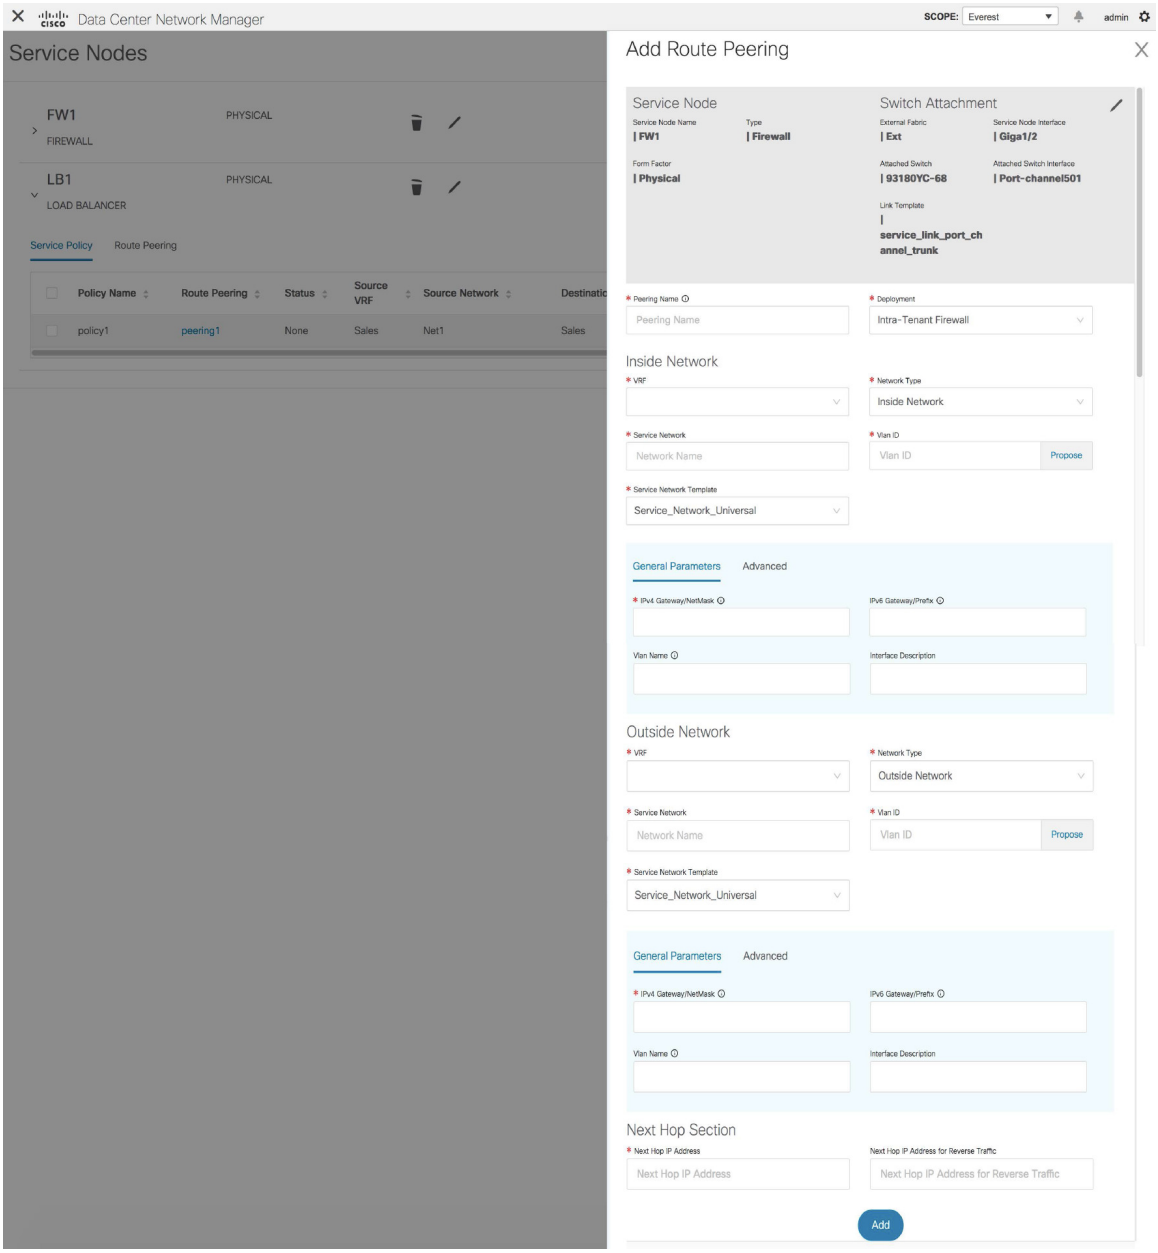

Specify the required parameters and click **Add**. For more information on specific fields, hover over the **i** icon.

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.

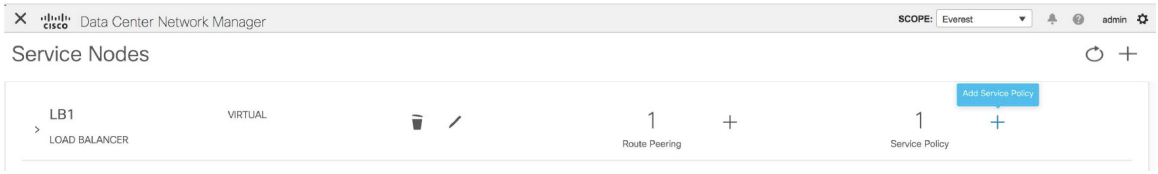

Step 2 The **Add Service Policy** window is displayed.

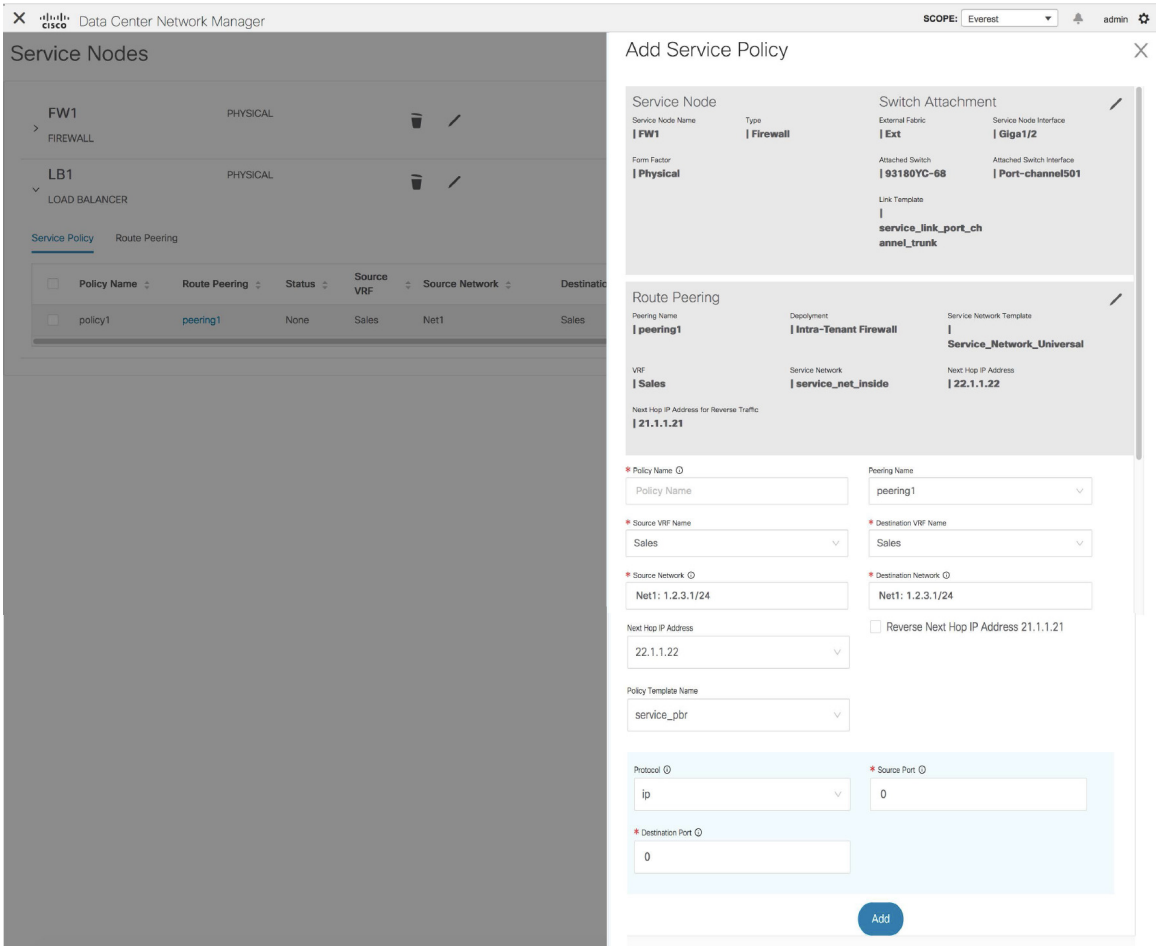

Specify the required parameters and click **Add**. For more information on specific fields, hover over the **i** icon.

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.

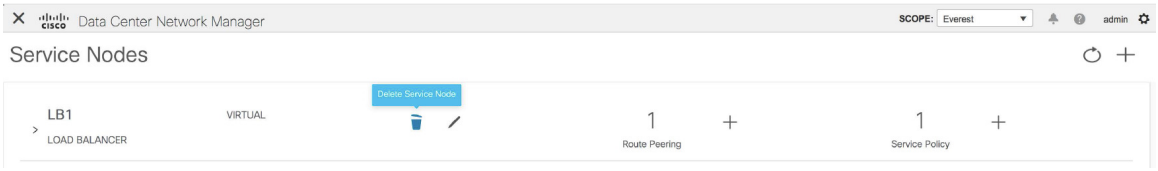

Step 2 A pop-up window comes up to confirm if the node has to be deleted. Click **Delete**.

Note Ensure that the service node that has to be deleted has no pairings or policies associated with it. In case there are pairings or policies associated with the service node, the deletion is blocked with a warning indicating that any pairings or 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.

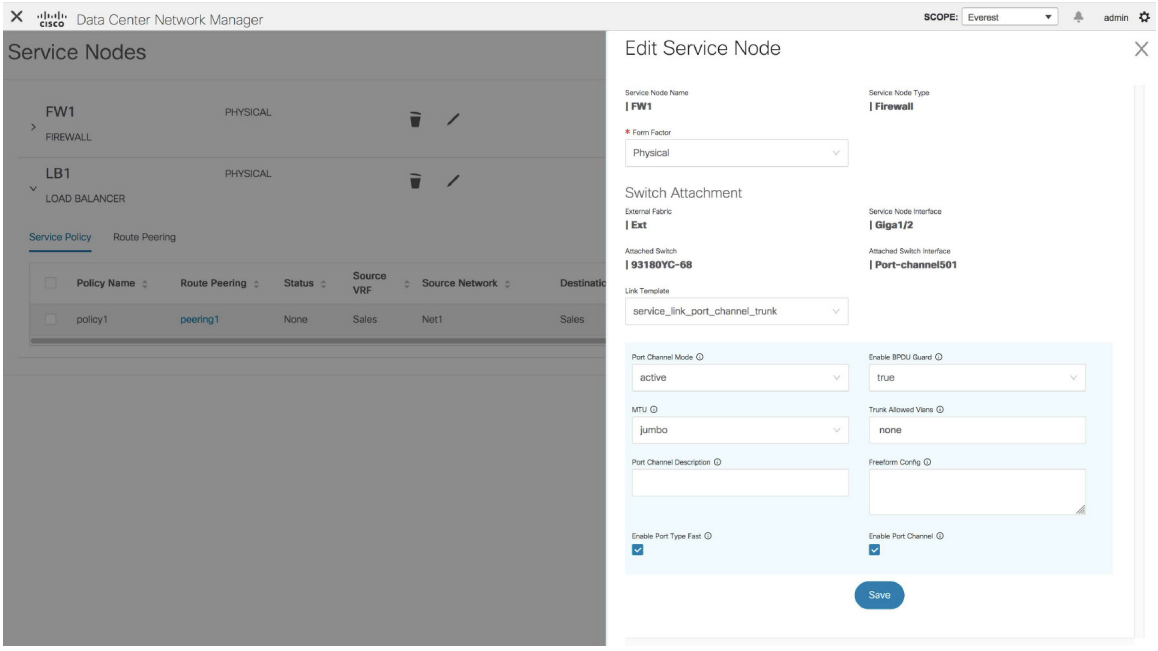

Make the required changes and click **Save**.

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.

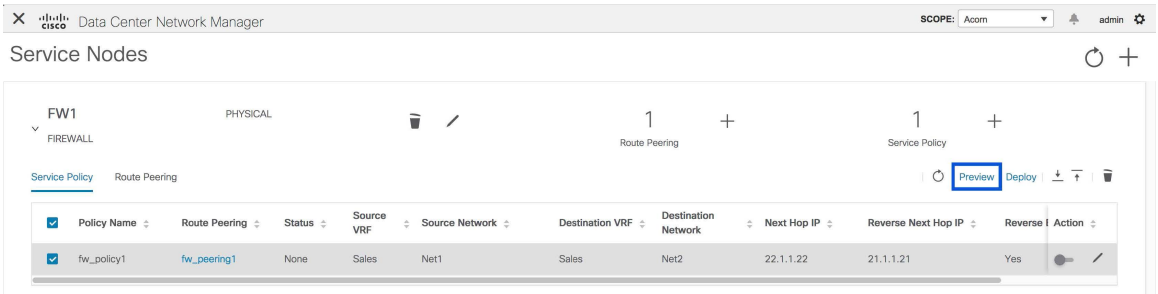

A **Preview Service Policy** or a **Preview Route Peering** window is displayed.

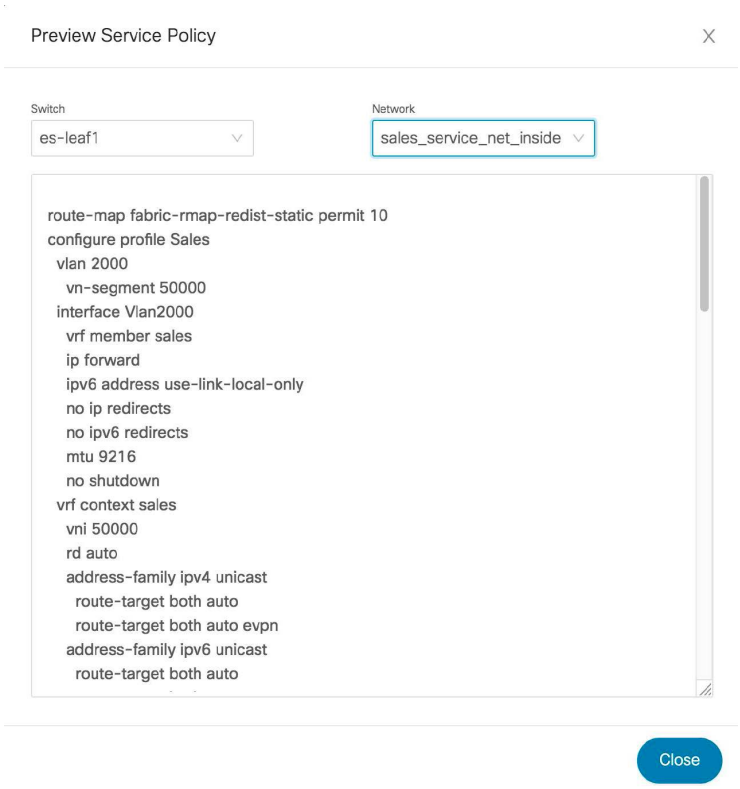

Step 2 Select a specific switch or network from the respective drop-down lists to display the service policies or route peerings for specific switches and networks. Click **Close** to close the window.

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.

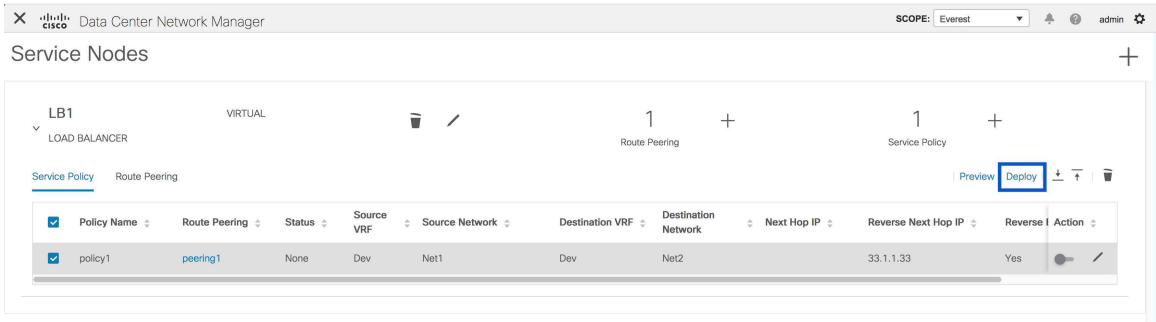

A pop-up window is displayed asking for confirmation to deploy.

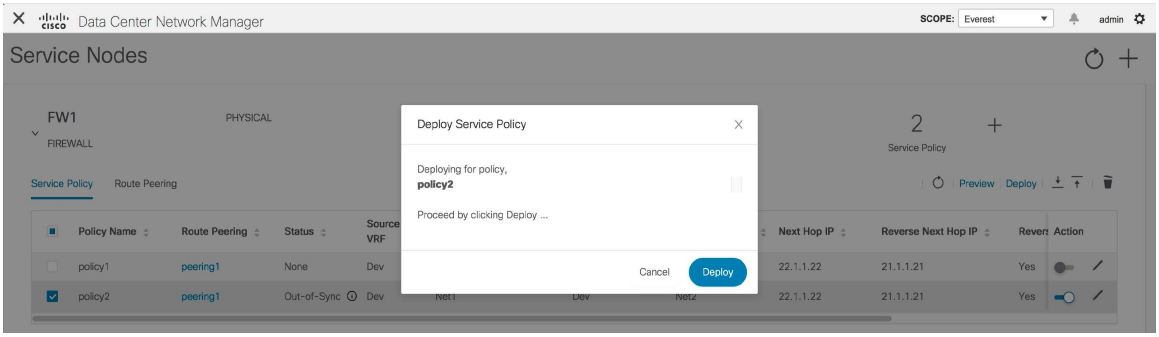

Step 2 Click **Deploy**.

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.

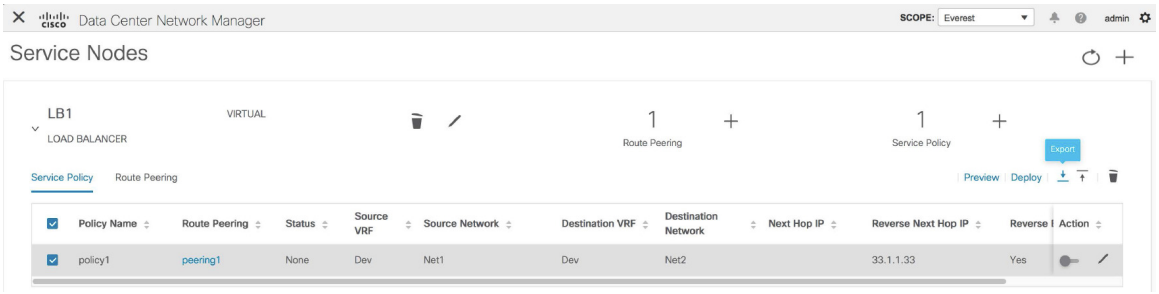

Importing a Service Policy or a Route Peering Table

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.

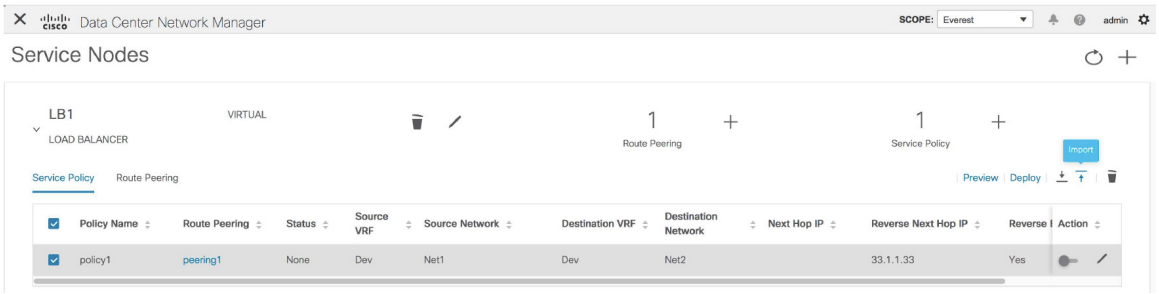

Deleting a Service Policy

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.

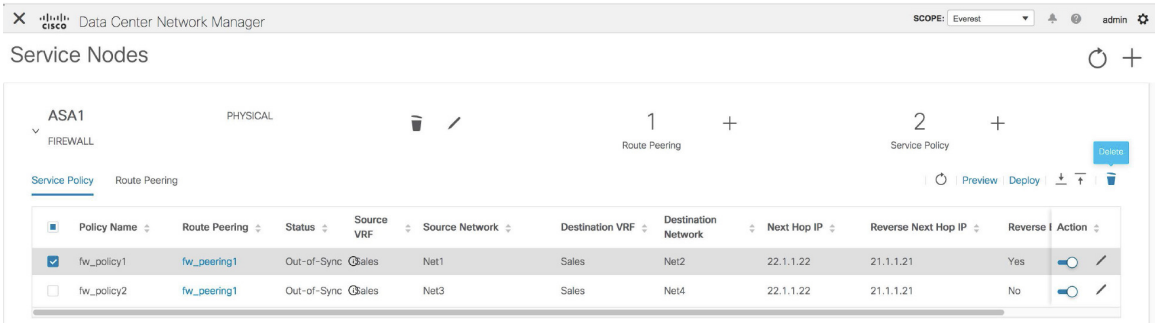

Step 2 A pop-up window is displayed asking for confirmation to delete. Click **Delete**.In case the service policy that has to be deleted is attached, the pop-up window indicates that the service policy has to be detached by using the toggle in the **Action** column, and deploying the changes (removing the policy) before it can be deleted.

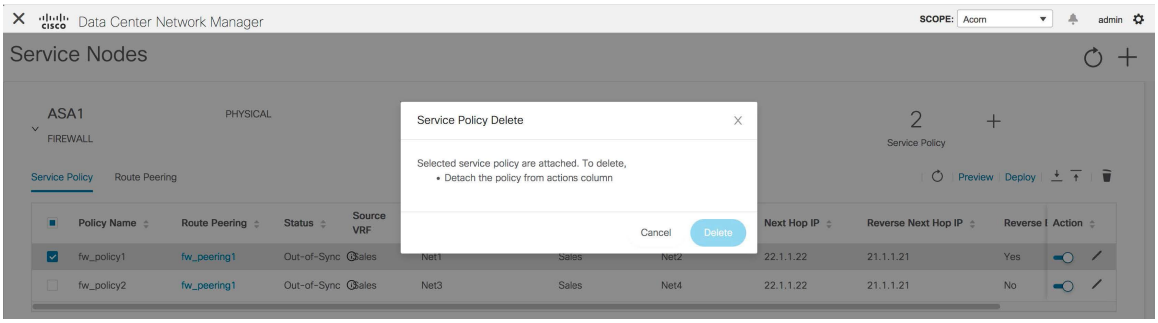

Deleting a Route Peering

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.

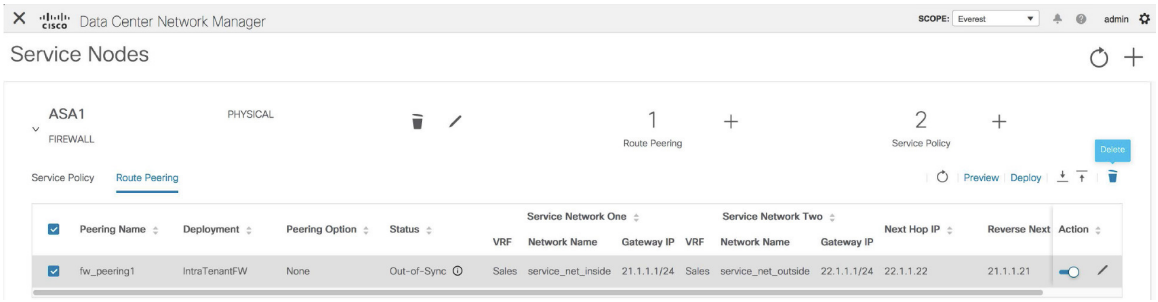

Step 2 A pop-up window is displayed asking for confirmation to delete. Click **Delete**. In case the route peering that has to be deleted is attached or if the service policy associated with the route peering is active, the pop-up window indicates that the peering has to be detached by using the toggle 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.

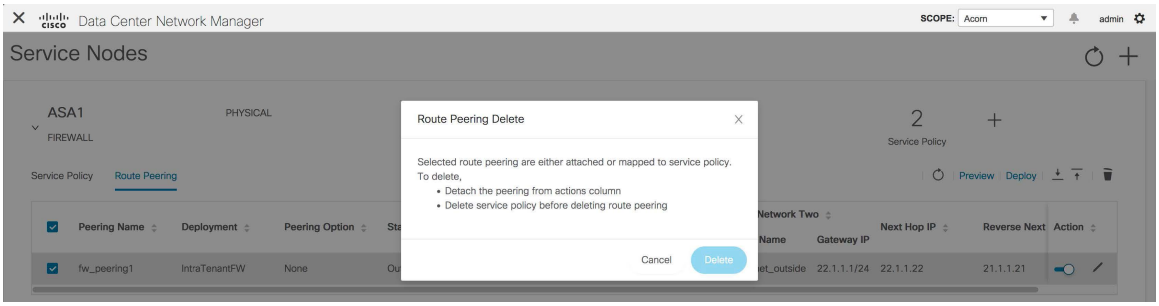

Viewing Service Policy Information

In the **Service Nodes** window, the **Service Policy** tab displays information about the configured service policies.

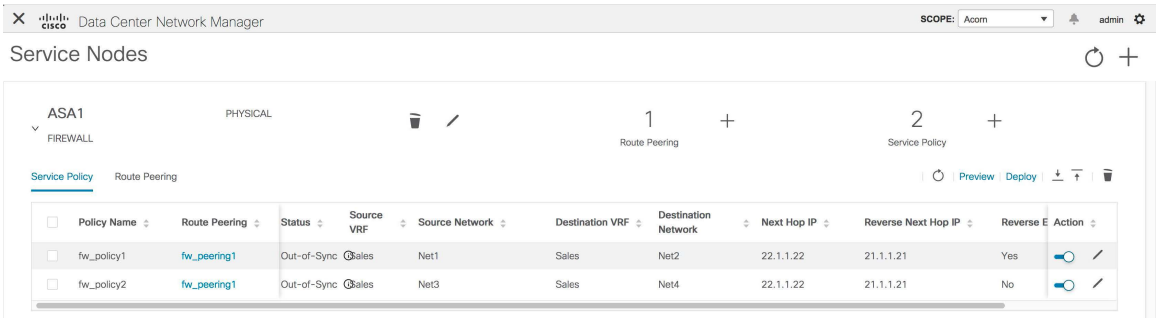

Table 1: Service Policy Table Field and Description

Field	Description
Policy Name	Displays the name of the policy.

Field	Description
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.
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.
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. 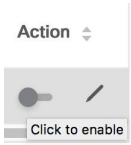 The toggle turns blue in color when the service policy is attached or enabled. 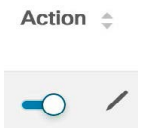 Click the **Edit** icon to edit the service policy. 

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, 1 week, and 1 month. 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.

Viewing Route Peering Information

In the **Service Nodes** window, click **Route Peering**. The **Route Peering** tab displays route peering information.

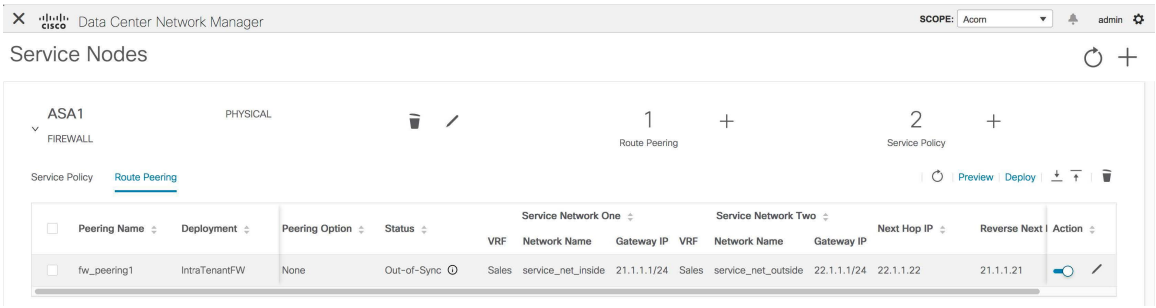

Table 2: Route Peering Table Field and Description

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.
Last Updated	Displays the time at which the route peering was last updated.

Field	Description
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. 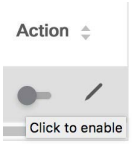 The toggle turns blue in color when the route peering is attached or enabled.  Click the **Edit** icon to edit the route peering.