



Managing a Greenfield VXLAN BGP EVPN Fabric

This chapter describes how to manage a greenfield VXLAN BGP EVPN fabric.

- [VXLAN BGP EVPN Fabrics Provisioning, on page 1](#)
- [Creating a New VXLAN BGP EVPN Fabric, on page 4](#)
- [Adding Switches to a Fabric, on page 24](#)
- [VXLAN EVPN Deployment with eBGP EVPN, on page 37](#)

VXLAN BGP EVPN Fabrics Provisioning

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

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

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

The DCNM GUI functions for creating and deploying VXLAN BGP EVPN fabrics are as follows:

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

Create, edit, and delete a fabric:

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

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

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

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

Underlay provisioning:

- Create, deploy, view, edit and delete a port-channel, vPC switch pair, Straight Through FEX (ST-FEX), Active-Active FEX (AA-FEX), loopback, subinterface, etc.
- Create breakout and unbreakout ports.
- Shut down and bring up interfaces.
- Rediscover ports and view interface configuration history.

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

Overlay network provisioning.

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

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

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

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

Guidelines for VXLAN BGP EVPN Fabrics Provisioning

- For any switch to be successfully imported into DCNM, the user specified for discovery/import, should have the following permissions:
 - SSH access to the switch
 - Ability to perform SNMPv3 queries
 - Ability to run the **show** commands including show run, show interfaces, etc.
- The switch discovery user need not have the ability to make any configuration changes on the switches. It is primarily used for read access.

- When an invalid command is deployed by DCNM to a device, for example, a command with an invalid key chain due to an invalid entry in the fabric settings, an error is generated displaying this issue. This error is not cleared after correcting the invalid fabric entry. You need to manually cleanup or delete the invalid commands to clear the error.

Note that the fabric errors related to the command execution are automatically cleared only when the same failed command succeeds in the subsequent deployment.

- LAN credentials are required to be set of any user that needs to be perform any write access to the device. LAN credentials need to be set on the DCNM, on a per user per device basis. When a user imports a device into the Easy Fabric, and LAN credentials are not set for that device, DCNM moves this device to a migration mode. Once the user sets the appropriate LAN credentials for that device, a subsequent Save & Deploy will retrigger the device import process.
- The **Save & Deploy** button triggers the intent regeneration for the entire fabric as well as a configuration compliance check for all the switches within the fabric. This button is required but not limited to the following cases:
 - A switch or a link is added, or any change in the topology
 - A change in the fabric settings that must be shared across the fabric
 - A switch is removed or deleted
 - A new vPC pairing or unpairing is done
 - A change in the role for a device

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

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

- Persistent configuration diff is seen for the command line: **system nve infra-vlan int force**. The persistent diff occurs if you have deployed this command via the freeform configuration to the switch. Although the switch requires the **force** keyword during deployment, the running configuration that is obtained from the switch in DCNM does not display the **force** keyword. Therefore, the **system nve infra-vlan int force** command always shows up as a diff.

The intent in DCNM contains the line:

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

The running config contains the line:

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

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

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

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

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

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

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

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

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

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

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

Creating a New VXLAN BGP EVPN Fabric

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

This procedure contains descriptions for the IPv4 underlay. For information about IPv6 underlay, see [IPv6 Underlay Support for Easy Fabric](#).

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

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

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

2. Click **Create Fabric**, the **Add Fabric** screen appears.

The fields are explained:

Fabric Name - Enter the name of the fabric.

Fabric Template - From the drop-down menu, choose the **Easy_Fabric_11_1** fabric template. The fabric settings for creating a standalone fabric appear.

Add Fabric ✕

* Fabric Name :

* Fabric Template : Easy_Fabric_11_1

General	Replication	vPC	Protocols	Advanced	Resources	Manageability	Bootstrap	Configuration Backup
* BGP ASN <input type="text"/> <small>1-4294967295 1-65535[0-65535]</small>								
Enable IPv6 Underlay <input type="checkbox"/>								
Enable IPv6 Link-Local Address <input checked="" type="checkbox"/>								
* Fabric Interface Numbering <input type="text" value="p2p"/> <small>Numbered(Point-to-Point) or Unnumbered</small>								
* Underlay Subnet IP Mask <input type="text" value="30"/> <small>Mask for Underlay Subnet IP Range</small>								
Underlay Subnet IPv6 Mask <input type="text"/>								
* Link-State Routing Protocol <input type="text" value="ospf"/> <small>Supported routing protocols (OSPF/IS-IS)</small>								
* Route-Reflectors <input type="text" value="2"/> <small>Number of spines acting as Route-Reflectors</small>								
* Anycast Gateway MAC <input type="text" value="2020.0000.00aa"/> <small>Shared MAC address for all leafs (xxxx.xxxx.xxxx)</small>								
NX-OS Software Image Version <input type="text"/> <small>If Set, Image Version Check Enforced On All Switches. Images Can Be Uploaded From Control:Image Upload</small>								

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.

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

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

Enable IPv6 Underlay: Enable the IPv6 underlay feature. For information, see [IPv6 Underlay Support for Easy Fabric](#).

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

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

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

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

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

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

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

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

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

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

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

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

- d. Click **Save & Deploy** in the fabric topology window.

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

Anycast Gateway MAC : Specifies the anycast gateway MAC address.

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

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

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

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

The screenshot shows the configuration page for the Replication tab. The tabs at the top are: General, Replication (selected), vPC, Protocols, Advanced, Resources, Manageability, Bootstrap, and Configuration Backup. The configuration fields are as follows:

- * Replication Mode**: Multicast (dropdown menu). Help: Replication Mode for BUM Traffic
- * Multicast Group Subnet**: 239.1.1.0/25. Help: Multicast address with prefix 16 to 30
- Enable Tenant Routed Multicast (TRM)**: . Help: For Overlay Multicast Support In VXLAN Fabrics
- Default MDT Address for TRM VRFs**: (empty text field). Help: IPv4 Multicast Address
- * Rendezvous-Points**: 2 (dropdown menu). Help: Number of spines acting as Rendezvous-Point (RP)
- * RP Mode**: asm (dropdown menu). Help: Multicast RP Mode
- * Underlay RP Loopback Id**: 254. Help: (Min:0, Max:1023)
- Underlay Primary RP Loopback Id**: (empty text field). Help: Used for Bidir-PIM Phantom RP (Min:0, Max:1023)
- Underlay Backup RP Loopback Id**: (empty text field). Help: Used for Fallback Bidir-PIM Phantom RP (Min:0, Max:1023)
- Underlay Second Backup RP Loopback Id**: (empty text field). Help: Used for second Fallback Bidir-PIM Phantom RP (Min:0, Max:1023)
- Underlay Third Backup RP Loopback Id**: (empty text field). Help: Used for third Fallback Bidir-PIM Phantom RP (Min:0, Max:1023)

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

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

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

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

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

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

For more information, see [Overview of Tenant Routed Multicast](#).

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

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

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



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

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

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

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

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

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

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

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

General	Replication	vPC	Protocols	Advanced	Resources	Manageability	Bootstrap	Configuration Backup
		* vPC Peer Link VLAN	3600	VLAN for vPC Peer Link SVI (Min:2, Max:3967)				
		Make vPC Peer Link VLAN as Native VLAN	<input type="checkbox"/>					
		* vPC Peer Keep Alive option	management	Use vPC Peer Keep Alive with Loopback or Management				
		* vPC Auto Recovery Time (In Seconds)	360	(Min:240, Max:3600)				
		* vPC Delay Restore Time (In Seconds)	150	(Min:1, Max:3600)				
		vPC Peer Link Port Channel ID	500	(Min:1, Max:4096)				
		vPC IPv6 ND Synchronize	<input checked="" type="checkbox"/>	Enable IPv6 ND synchronization between vPC peers				
		vPC advertise-pip	<input type="checkbox"/>	For Primary VTEP IP Advertisement As Next-Hop Of Prefix Routes				
		Enable the same vPC Domain Id for all vPC Pairs	<input type="checkbox"/>	(Not Recommended)				
		vPC Domain Id		vPC Domain Id to be used on all vPC pairs				
		vPC Domain Id Range	1-1000	vPC Domain id range to use for new pairings				
		Enable Qos for Fabric vPC-Peering	<input type="checkbox"/>	Qos on spines for guaranteed delivery of vPC Fabric Peering communication				
		Qos Policy Name		Qos Policy name should be same on all spines				

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

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

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

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

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

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

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

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

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

You can enable the advertise PIP feature on a specific vPC as well. For more information, see [Advertising PIP on vPC](#).

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

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

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

Enable Qos for Fabric vPC-Peering - Enable QoS on spines for guaranteed delivery of vPC Fabric Peering communication. For more information, see [QoS for Fabric vPC-Peering](#).



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

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

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

Add Fabric ✕

* Fabric Name :

* Fabric Template : Easy_Fabric_11_1

© Fabric Template for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches.

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

Enable BFD For PIM ⓘ

Enable BFD Authentication ⓘ Valid for P2P Interfaces only

BFD Authentication Key ID ⓘ

BFD Authentication Key ⓘ Encrypted SHA1 secret value

IBGP Peer-Template Config

Leaf/Border/Border Gateway IBGP Peer-Template Config

Specifies the config used for RR and spines with border or border gateway role. This field should begin with 'template peer' or 'template peer-session'. This must have 2 leading spaces. Note! All configs should strictly match 'show run' output, with respect to case and newlines. Any mismatches will yield unexpected diffs during deploy.

Specifies the config used for leaf, border or border gateway. If this field is empty, the peer template defined in IBGP Peer-Template Config is used on all BGP enabled devices (RRs, leafs, border or border gateway roles). This field should begin with 'template peer' or 'template peer-session'. This must have 2 leading spaces. Note! All configs should strictly match 'show run' output, with respect to case and newlines. Any mismatches will yield unexpected diffs during deploy.

Save Cancel

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

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

Underlay Routing Protocol Tag - The tag defining the type of network.

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



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

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

OSPF Authentication Key ID - The Key ID is populated.

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



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

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

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

IS-IS Authentication Keychain Name - Enter the Keychain name, such as CiscoisAuth.

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

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



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

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



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

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

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



Note Plain text passwords are not supported. Log in 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 PIM Hello Authentication - Enables the PIM hello authentication.

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

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

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

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

```
feature bfd
```



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

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

For information about BFD feature compatibility, refer your respective platform documentation and for information about the supported software images, see *Compatibility Matrix for Cisco DCNM*.

Enable BFD for iBGP: Select the check box to enable BFD for the iBGP neighbor. This option is disabled by default.

Enable BFD for OSPF: Select the check box to enable BFD for the OSPF underlay instance. This option is disabled by default, and it is grayed out if the link state protocol is ISIS.

Enable BFD for ISIS: Select the check box to enable BFD for the ISIS underlay instance. This option is disabled by default, and it is grayed out if the link state protocol is OSPF.

Enable BFD for PIM: Select the check box to enable BFD for PIM. This option is disabled by default, and it is be grayed out if the replication mode is Ingress.

Here are the examples of the BFD global policies:

```
router ospf <ospf tag>
  bfd

router isis <isis tag>
  address-family ipv4 unicast
  bfd

ip pim bfd

router bgp <bgp asn>
  neighbor <neighbor ip>
  bfd
```

Enable BFD Authentication: Select the check box to enable BFD authentication. If you enable this field, the **BFD Authentication Key ID** and **BFD Authentication Key** fields are editable.



Note BFD Authentication is not supported when the **Fabric Interface Numbering** field under the **General** tab is set to **unnumbered**. The BFD authentication fields will be grayed out automatically. BFD authentication is valid for only for P2P interfaces.

BFD Authentication Key ID: Specifies the BFD authentication key ID for the interface authentication. The default value is 100.

BFD Authentication Key: Specifies the BFD authentication key.

For information about how to retrieve the BFD authentication parameters, see [Retrieving the Encrypted BFD Authentication Key](#).

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

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

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

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

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

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

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

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

General	Replication	vPC	Protocols	Advanced	Resources	Manageability	Bootstrap	Configuration Backup
				<p>* VRF Template <input type="text" value="Default_VRF_Universal"/> ? <i>Default Overlay VRF Template For Leafs</i></p> <p>* Network Template <input type="text" value="Default_Network_Universal"/> ? <i>Default Overlay Network Template For Leafs</i></p> <p>* VRF Extension Template <input type="text" value="Default_VRF_Extension_Universal"/> ? <i>Default Overlay VRF Template For Borders</i></p> <p>* Network Extension Template <input type="text" value="Default_Network_Extension_Universa"/> ? <i>Default Overlay Network Template For Borders</i></p> <p>Site Id <input type="text"/> ? <i>For EVPN Multi-Site Support (Min:1, Max: 281474976710655). Defaults to Fabric ASN</i></p> <p>* Intra Fabric Interface MTU <input type="text" value="9216"/> ? <i>(Min:576, Max:9216). Must be an even number</i></p> <p>* Layer 2 Host Interface MTU <input type="text" value="9216"/> ? <i>(Min:1500, Max:9216). Must be an even number</i></p> <p>* Power Supply Mode <input type="text" value="ps-redundant"/> ? <i>Default Power Supply Mode For The Fabric</i></p> <p>* CoPP Profile <input type="text" value="strict"/> ? <i>Fabric Wide CoPP Policy. Customized CoPP policy should be provided when 'manual' is selected</i></p> <p>VTEP HoldDown Time <input type="text" value="180"/> ? <i>NVE Source Interface HoldDown Time (Min:1, Max:1500) in seconds</i></p>				

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

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

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

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

Layer 2 Host Interface MTU - Specifies the MTU for the layer 2 host interface. This value should be an even number.

Power Supply Mode - Choose the appropriate power supply mode.

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

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

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

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

Example overlay network name: Site_VNI12345_VLAN1234



Note Ignore this field for greenfield deployments. The **Brownfield Overlay Network Name Format** applies 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 CDP for Bootstrapped Switch - Enables CDP on management (mgmt0) interface for bootstrapped switch. By default, for bootstrapped switches, CDP is disabled on the mgmt0 interface.

Enable VXLAN OAM - Enables the VXLAN OAM functionality for devices in the fabric. This is enabled by default. Clear the check box to disable VXLAN OAM function.

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



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

Enable Tenant DHCP – Select the check box to enable feature dhcp and associated configurations globally on all switches in the fabric. This is a pre-requisite for support of DHCP for overlay networks that are part of the tenant VRFs.



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

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

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



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

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

Enable Strict Config Compliance - Enable the Strict Config Compliance feature by selecting this check box. By default, this feature is disabled. For more information, refer [Strict Configuration Compliance](#).

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

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

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

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

PTP Source Loopback Id: Specifies the loopback interface ID Loopback that is used as the Source IP Address for all PTP packets. The valid values range from 0 to 1023. The PTP loopback ID cannot be

the same as RP, Phantom RP, NVE, or MPLS loopback ID. Otherwise, an error will be generated. The PTP loopback ID can be the same as BGP loopback or user-defined loopback which is created from DCNM.

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

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

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

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

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

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

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

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

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

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

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

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

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

Enable MACsec - Enables MACsec for the fabric. For more information, see [MACsec Support in Easy Fabric and eBGP Fabric](#).

Freeform CLIs - Fabric level freeform CLIs can be added while creating or editing a fabric. They are applicable to switches across the fabric. You must add the configurations as displayed in the running configuration, without indentation. Switch level freeform configurations such as VLAN, SVI, and interface configurations should only be added on the switch. For more information, refer [Enabling Freeform Configurations on Fabric Switches](#).

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

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

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

8. Click the **Resources** tab.

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

Manual Underlay IP Address Allocation Checking this will disable Dynamic Underlay IP Address Allocations

- * Underlay Routing Loopback IP Range: 10.2.0.0/22 Typically Loopback0 IP Address Range
- * Underlay VTEP Loopback IP Range: 10.3.0.0/22 Typically Loopback1 IP Address Range
- * Underlay RP Loopback IP Range: 10.254.254.0/24 Anycast or Phantom RP IP Address Range
- * Underlay Subnet IP Range: 10.4.0.0/16 Address range to assign Numbered and Peer Link SVI IPs
- Underlay MPLS Loopback IP Range: Used for VXLAN to MPLS SR/LDP Handoff
- Underlay Routing Loopback IPv6 Range: Typically Loopback0 IPv6 Address Range
- Underlay VTEP Loopback IPv6 Range: Typically Loopback1 and Anycast Loopback IPv6 Address Range
- Underlay Subnet IPv6 Range: IPv6 Address range to assign Numbered and Peer Link SVI IPs
- BGP Router ID Range for IPv6 Underlay:
- * Layer 2 VXLAN VNI Range: 30000-49000 Overlay Network Identifier Range (Min:1, Max:16777214)
- * Layer 3 VXLAN VNI Range: 50000-59000 Overlay VRF Identifier Range (Min:1, Max:16777214)
- * Network VLAN Range: 2300-2999 Per Switch Overlay Network VLAN Range (Min:2, Max:3967)
- * VRF VLAN Range: 2000-2299 Per Switch Overlay VRF VLAN Range (Min:2, Max:3967)
- * Subinterface Dot1a Range: 2-511 Per Border Dot1a Range For VRF Lite Connectivity (Min:2, Max:4093)

Save **Cancel**

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

- By default, DCNM allocates the underlay IP address resources (for loopbacks, fabric interfaces, etc) dynamically from the defined pools. If you select the check box, the allocation scheme switches to static, and some of the dynamic IP address range fields are disabled.
- For static allocation, the underlay IP address resources must be populated into the Resource Manager (RM) using REST APIs.

Refer the Cisco DCNM REST API Reference Guide, Release 11.2(1) for more details. The REST APIs must be invoked after the switches are added to the fabric, and before you use the Save & Deploy option.

- The Underlay RP Loopback IP Range field stays enabled if BIDIR-PIM function is chosen for multicast replication.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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



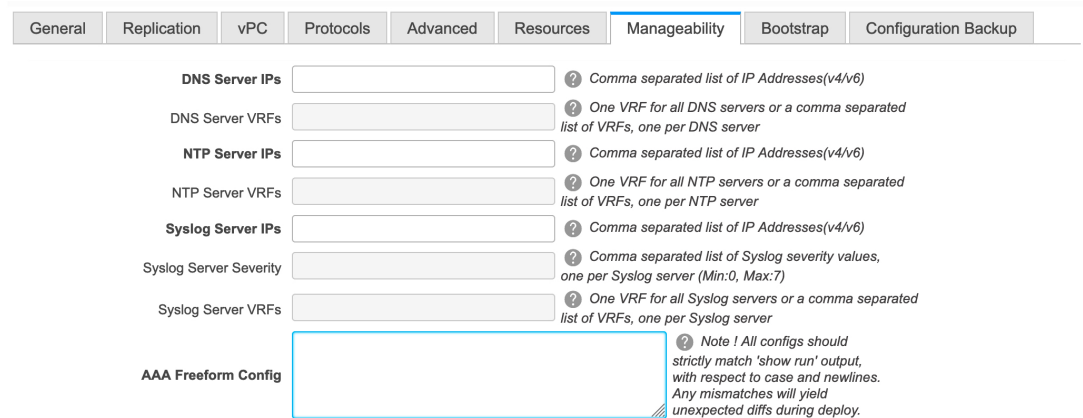
Note When you update a range of values, ensure that it does not overlap with other ranges. You should only update one range of values at a time. If you want to update more than one range of values, do it in separate instances. For example, if you want to update L2 and L3 ranges, you should do the following.

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

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

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

9. Click the **Manageability** tab.



The fields in this tab are:

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

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

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

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

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

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

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

AAA Freeform Config – Specifies the AAA freeform configurations.

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

10. Click the **Bootstrap** tab.

The screenshot shows the configuration page for a new VXLAN BGP EVPN fabric. The 'Bootstrap' tab is active. The 'Enable Bootstrap' checkbox is checked. Below it are several configuration options:

- Enable Bootstrap** (checked) - Automatic IP Assignment For POAP
- Enable Local DHCP Server** (unchecked) - Automatic IP Assignment For POAP From Local DHCP Server
- DHCP Version** (dropdown menu)
- DHCP Scope Start Address** (text field) - Start Address For Switch Out-of-Band POAP
- DHCP Scope End Address** (text field) - End Address For Switch Out-of-Band POAP
- Switch Mgmt Default Gateway** (text field) - Default Gateway For Management VRF On The Switch
- Switch Mgmt IP Subnet Prefix** (text field) - (Min:8, Max:30)
- Switch Mgmt IPv6 Subnet Prefix** (text field) - (Min:64, Max:126)
- Enable AAA Config** (unchecked) - Include AAA configs from Manageability tab during device bootstrap
- Bootstrap Freeform Config** (text area) - Note! All configs should strictly match 'show run' output, with respect to case and newlines. Any mismatches will yield unexpected diffs during deploy. Enter One Subnet Scope per line. Start_IP, End_IP, Gateway, Prefix e.g. 10.6.0.2, 10.6.0.9, 10.6.0.1, 24 10.7.0.2, 10.7.0.9, 10.7.0.1, 24 Or 21:0:1:1::10, 21:0:1:1::20, 21:0:1:1::1, 64 21:0:1:2::10, 21:0:1:2::20, 21:0:1:2::1, 64
- DHCPv4/DHCPv6 Multi Subnet Scope** (text area)

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

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

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

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

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

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



Note Cisco Nexus 9000 and 3000 Series Switches support IPv6 POAP only when switches are either Layer-2 adjacent (eth1 or out-of-band subnet must be a /64) or they are L3 adjacent residing in some IPv6 /64 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 configurations from the Manageability tab as part of the device startup config post bootstrap.

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

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

DHCPv4/DHCPv6 Multi Subnet Scope - Specifies the field to enter one subnet scope per line. This field is editable after you check the **Enable Local DHCP Server** check box.

The format of the scope should be defined as:

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

For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24

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

General Replication vPC Protocols Advanced Resources Manageability Bootstrap Configuration Backup

Hourly Fabric Backup ? Backup hourly or on Re-sync only if there is any config deployment since last backup

Scheduled Fabric Backup ? Backup at the specified time only if there is any config deployment since last backup

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

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

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

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

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

Select both the check boxes to enable both back up processes.

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

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

The backup configuration files are stored in the following path in DCNM:

/usr/local/cisco/dcm/dcnm/data/archive

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



- Note** To trigger an immediate backup, do the following:
- Choose **Control > Fabric Builder**. The Fabric Builder screen comes up.
 - Click within the specific fabric box. The fabric topology screen comes up.
 - From the **Actions** pane at the left part of the screen, click **Re-Sync Fabric**.

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

- Click **ThousandEyes Agent** tab. This feature is supported on Cisco DCNM Release 11.5(3) only. For more information, refer to [Configuring Global Settings for ThousandEyes Enterprise Agent](#).

The screenshot shows the configuration page for the ThousandEyes Agent. The tabs at the top are: General, Replication, vPC, Protocols, Advanced, Resources, Manageability, Bootstrap, Configuration Backup, and ThousandEyes Agent (selected). The configuration fields include:

- Enable Fabric Override for ThousandEyes Agent Installation:** A checkbox that is currently unchecked.
- ThousandEyes Account Group Token:** A text input field with a help icon. Description: *Token from ThousandEyes Agent Settings for Agent Installation*
- VRF on Switch for ThousandEyes Agent Collector Reachability:** A text input field with a help icon. Description: *NX-OS VRF that provides Internet Reachability*
- DNS Domain:** A text input field with a help icon. Description: *DNS Domain Configuration*
- DNS Server IPs:** A text input field with a help icon. Description: *Comma separated list of IP Addresses(v4/v6)*
- NTP Server IPs:** A text input field with a help icon. Description: *Comma separated list of IP Addresses(v4/v6)*
- Enable Proxy for Internet Access:** A checkbox that is currently unchecked.
- Proxy Information:** A text input field with a help icon. Description: *Proxy-Server:port*
- Proxy Bypass:** A text input field with a help icon. Description: *Comma separated No-proxy server list*

At the bottom right, there are **Save** and **Cancel** buttons.

The fields on this tab are:



- Note** The fabric settings for ThousandEyes Agent overwrites the global settings and applies the same configuration for all the ThousandEyes Agent installed on switches in that fabric.

- **Enable Fabric Override for ThousandEyes Agent Installation:** Select the check box to enable the ThousandEyes Enterprise Agent on the fabric.
- **ThousandEyes Account Group Token:** Specifies ThousandEyes Enterprise Agent account group token for installation.
- **VRF on Switch for ThousandEyes Agent Collector Reachability:** Specifies the VRF data which provides internet reachability.
- **DNS Domain:** Specifies the switch DNS domain configuration.
- **DNS Server IPs:** Specifies the comma separated list of IP addresses (v4/v6) of Domain Name System (DNS) server. You can enter a maximum of three IP addresses for the DNS Server.
- **NTP Server IPs:** Specifies comma separated list of IP addresses (v4/v6) of Network Time Protocol (NTP) server. You can enter a maximum of three IP addresses for the NTP Server.

- **Enable Proxy for Internet Access:** Select the check box to enable the proxy setting for NX-OS switch internet access.
 - **Proxy Information:** Specifies the proxy server port information.
 - **Proxy Bypass:** Specifies the server list for which proxy is bypassed.
13. Click **Save** after filling and updating relevant information. A note appears briefly at the bottom right part of the screen, indicating that the fabric is created. When a fabric is created, the fabric page comes up. The fabric name appears at the top left part of the screen.

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

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

- **Tabular View** - By default, the switches are displayed in the topology view. Use this option to view switches in the tabular view.
- **Refresh topology** - Allows you to refresh the topology.
- **Save Layout** – Saves a custom view of the topology. You can create a specific view in the topology and save it for ease of use.
- **Delete saved layout** – Deletes the custom view of the topology
- **Topology views** - You can choose between Hierarchical, Random and Custom saved layout display options.
 - **Hierarchical** - Provides an architectural view of your topology. Various Switch Roles can be defined that draws the nodes on how you configure your CLOS topology.
 - **Random** - Nodes are placed randomly on the window. DCNM tries to make a guess and intelligently place nodes that belong together in close proximity.
 - **Custom saved layout** - You can drag nodes around to your liking. Once you have the positions as how you like, you can click Save Layout to remember the positions. Next time you come to the topology, DCNM will draw the nodes based on your last saved layout positions.
- **Restore Fabric** – Allows you to restore the fabric to a prior DCNM configuration state (one month back, two months back, and so on). For more information, see the *Restore Fabric* section.
- **Backup Now:** You can initiate a fabric backup manually by clicking **Backup Now**. Enter a name for the tag and click **OK**. Regardless of the settings you choose under the **Configuration Backup** tab in the **Fabric Settings** dialog box, you can initiate a backup using this option.
- **Resync Fabric** - Use this option to resynchronize DCNM state when there is a large scale out-of-band change, or if configuration changes do not register in the DCNM properly. The resync operation does a full CC run for the fabric switches and recollects “show run” and “show run all” commands from the switches. When you initiate the re-sync process, a progress message is displayed on the window. During the re-sync, the running configuration is taken from the switches. Then, the Out-of-Sync/In-Sync status for the switch is recalculated based on the intent or expected configuration defined in DCNM versus the current running configuration that was taken from the switches.

- **Add Switches** – Allows you to add switch instances to the fabric.
- **Fabric Settings** – Allows you to view or edit fabric settings.
- **Cloud icon** - Click the **Cloud** icon to display (or not display) an **Undiscovered** cloud.

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

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

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

Adding Switches to a Fabric

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

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

Additionally, you can pre-provision switches and interfaces. For more information, see [Pre-provisioning a Device](#) and [Pre-provisioning an Ethernet Interface](#).




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

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

Discovering New Switches

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



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

5. Click the **POAP** tab.

As mentioned earlier, DCNM retrieves the serial number, model number, and version from the device and displays them on the Inventory Management along window. Also, an option to add the IP address, hostname, and password are made available. If the switch information is not retrieved, refresh the window.



Note

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

Inventory Management ✕

Discover Existing Switches
PowerOn Auto Provisioning (POAP)

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

Bootstrap

+ ✎ ✕ ↺ ↻

* Admin Password

* Confirm Admin Password

🔒

<input type="checkbox"/>	Serial Number	Model	Version	IP Address	Hostname	Gateway
No Data available						

Close

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

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

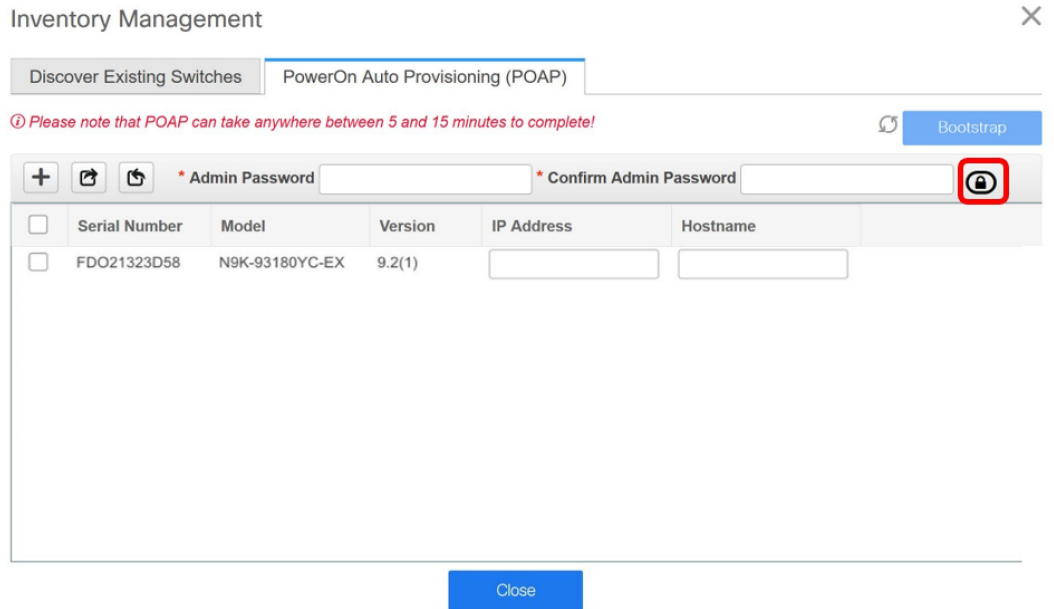
Beginning with Release 11.2(1), you can provision devices in advance. To pre-provision devices, refer to [Pre-provisioning a Device](#).

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

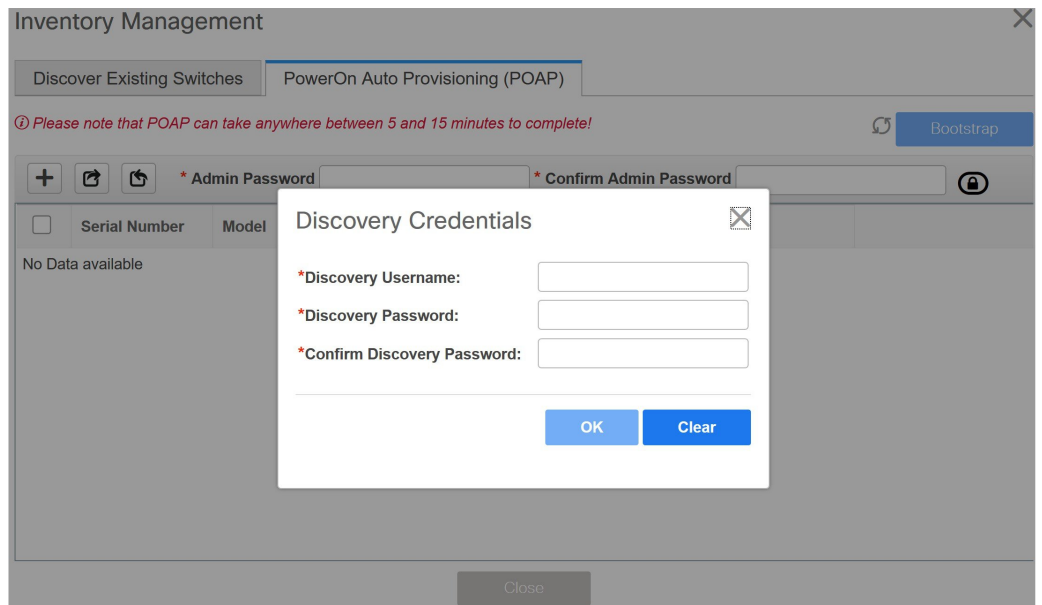


Note If you do not want to use admin credentials to discover switches, you can instead use the AAA authentication, that is, RADIUS or TACACS credentials for discovery only.

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



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



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

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

- 8. Click **Bootstrap** at the top right part of the screen.

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

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



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

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

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

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

Fabric errors & warnings



0 Errors, 2 Warnings, 0 Info

✕ Delete all

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

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

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

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

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

Interfaces

2

	Device Name	Name	Admin	Oper	Reason
<input type="checkbox"/>	N9K-2-Leaf	Ethernet2/6	↑	↓	XCVR not inserted
<input type="checkbox"/>	N9K-2-Leaf	Ethernet2/7	↑	↓	XCVR not inserted
<input type="checkbox"/>	N9K-2-Leaf	Ethernet2/8	↑	↓	XCVR not inserted
<input type="checkbox"/>	N9K-2-Leaf	Ethernet2/9	↑	↓	XCVR not inserted
<input type="checkbox"/>	N9K-2-Leaf	Ethernet2/10	↑	↓	XCVR not inserted
<input type="checkbox"/>	N9K-2-Leaf	Ethernet2/11	↑	↓	XCVR not inserted
<input type="checkbox"/>	N9K-2-Leaf	Ethernet2/12	↑	↓	XCVR not inserted
<input checked="" type="checkbox"/>	N9K-2-Leaf	nve1	↑	↑	ok

1

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

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



- Note**
- Changing of the switch role is allowed only before executing **Save & Deploy**.
 - Starting from DCNM 11.1(1), switch roles can be changed if there are no overlays on the switches, but only as per the list of allowed switch role changes given at [Switch Operations](#).

- **Modes** - Maintenance and Active/Operational modes.
- **vPC Pairing** - Select a switch for vPC and then select its peer.
You can create a virtual link for a vPC pair or change the existing physical link to a virtual link for a vPC pair.
- **Manage Interfaces** - Deploy configurations on the switch interfaces.
- **View/Edit Policies** - See switch policies and edit them as required.
- **History** - View per switch deployment and policy change history.

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

History for mini-leaf2(FDO21332E6X)

Deployment History Policy Change History

Policy ID	Template	Description	PTI Operation	Generated Config	Entity Name	Entity Type	User	Created On
PROFILE-VRF-1	Default_VRF_Exten...		UPDATE	Detailed History	MyVRF_50000	Config_Profile	admin	2020/05/31-08:15:21
PROFILE-VRF-1	Default_VRF_Exten...		ADD	Detailed History	MyVRF_50000	Config_Profile	admin	2020/05/31-08:13:44
PROFILE-NETWO...	Default_Network_E...		ADD	Detailed History	MyNetwork_30...	Config_Profile	admin	2020/05/31-08:13:43

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

Generated Config Details for FDO22471AXH



Generated Config Before

Generated Config After

```
hostname es-leaf1
```

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

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



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

- **Preview Config** - View the pending configuration and the side-by-side comparison of the running and expected configuration.
- **Deploy Config** - Deploy per switch configurations.
- **Discovery** - You can use this option to update the credentials of the switch, reload the switch, rediscover the switch, and remove the switch from the fabric.

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

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

Discovering Existing Switches

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



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

Inventory Management

Discover Existing Switches

PowerOn Auto Provisioning (POAP)

Discovery Information > Scan Details >

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

Authentication Protocol MD5 ▼

Username

Password

Max Hops 2 ▲ ▼ hop(s)

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

Start discovery

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

✕

Discover Existing Switches

PowerOn Auto Provisioning (POAP)

Discovery Information > Scan Details >

← Back
Import into fabric

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

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

The screenshot shows the 'Inventory Management' window with the 'Discover Existing Switches' tab active. Below the tab are navigation links for 'Discovery Information' and 'Scan Details'. A 'Back' button is on the left, and an 'Import into fabric' button is on the right. A table lists discovered switches with columns for Name, IP Address, Model, Version, Status, and Progress. The 'leaf-91' switch is highlighted in blue, and its checkbox is checked. A yellow circle with the number '1' is next to the checkbox, and another yellow circle with the number '2' is next to the 'Import into fabric' button.

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

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

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



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

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

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




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

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

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



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



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

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

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

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

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

The template and interface configurations form the underlay network configuration on the switches. Also, freeform CLIs that were entered as part of fabric settings (leaf and spine switch freeform configurations

entered in the Advanced tab) are deployed. For more details on freeform configurations, refer [Enabling Freeform Configurations on Fabric Switches](#).





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

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

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

Config Deployment ✕

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

Switch Name	IP Address	Switch Serial	Preview Config	Status	Re-sync	Progress
N9K-2-Leaf	111.0.0.92	SAL18422FVP	0 lines	In-sync		<div style="width: 100%;"><div style="width: 100%;"></div></div> 100%
N9K-4-BGW	111.0.0.94	FDO20260UEK	20 lines	Out-of-sync		<div style="width: 100%;"><div style="width: 100%;"></div></div> 100%
N9K-3-BGW	111.0.0.93	FDO20291AVQ	20 lines	Out-of-sync		<div style="width: 100%;"><div style="width: 100%;"></div></div> 100%
N9K-1-Spine	111.0.0.91	SAL18432P2T	0 lines	In-sync		<div style="width: 100%;"><div style="width: 100%;"></div></div> 100%

[Deploy Config](#)

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

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

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

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

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

In DCNM 11, multi-line banner motd configuration is supported. Multi-line banner motd configuration can be configured in DCNM with freeform configuration policy, either per switch using **switch_freeform**, or per fabric using leaf/spine freeform configuration. Note that after the multi-line banner motd is configured, deploy the policy by executing the **Save & Deploy** option in the (top right part of the) fabric topology screen. Else, the policy may not be deployed properly on the switch. The **banner** policy is only to configure single-line banner configuration. Also, you can only create one banner related freeform configuration/policy. Multiple policies for configuring banner motd are not supported.

7. Close the screen.

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

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

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

If a switch icon is in red color, it indicates that the switch and DCNM configurations are not in sync. When deployment is pending on a switch, the switch is displayed in blue color. The pending state indicates that there is a pending deployment or pending recomputation. You can click on the switch and review the pending deployments using **Preview** or **Deploy Config** options, or click **Save & Deploy** to recompute the state of the switch.



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

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

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

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

VXLAN EVPN Deployment with eBGP EVPN

Creating a eBGP New VXLAN EVPN with eBGP-based Underlay

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

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

- 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.

Add Fabric ✕

* Fabric Name :

* Fabric Template : Easy_Fabric_eBGP

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

* BGP ASN for Spines 1-4294967295 | 1-65535[.0-65535]

* BGP AS Mode Multi-AS Multi-AS: Unique ASN per Leaf/Border
Dual-AS: One ASN for all Leafs/Borders

* Underlay Subnet IP Mask 30 Mask for Underlay Subnet IP Range

Manual Underlay IP Address Allocation Checking this will disable Dynamic Underlay IP Address Allocations

* Underlay Routing Loopback IP Range 10.2.0.0/22 Typically Loopback0 IP Address Range

* Underlay Subnet IP Range 10.4.0.0/16 Address range to assign Numbered and Peer Link SVI IPs

* Subinterface Dot1q Range 2-511 Per Border Dot1q Range For VRF Lite Connectivity (Min:2, Max:4095)

NX-OS Software Image Version If Set, Image Version Check Enforced On All Switches.
Images Can Be Uploaded From Control:Image Upload

- 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:

General	EVPN	vPC	Protocols	Advanced	Manageability	Bootstrap	Configuration Backup
	<p>Enable EVPN VXLAN Overlay <input checked="" type="checkbox"/> ?</p> <p>First Hop Redundancy Protocol <input type="text"/> ? <i>HSRP or VRRP</i></p> <p>* Anycast Gateway MAC <input type="text" value="2020.0000.00aa"/> ? <i>Shared MAC address for all leaves (xxxx.xxxx.xxxx)</i></p> <p>Enable VXLAN OAM <input checked="" type="checkbox"/> ? <i>For Operations, Administration, and Management Of VXLAN Fabrics</i></p> <p>Enable Tenant DHCP <input checked="" type="checkbox"/> ?</p> <p>vPC advertise-pip <input type="checkbox"/> ? <i>For Primary VTEP IP Advertisement As Next-Hop Of Prefix Routes</i></p> <p>* Replication Mode <input type="text" value="Multicast"/> ? <i>Replication Mode for BUM Traffic</i></p> <p>* Multicast Group Subnet <input type="text" value="239.1.1.0/25"/> ? <i>Multicast address with prefix 16 to 30</i></p> <p>Enable Tenant Routed Multicast <input type="checkbox"/> ? <i>For Overlay Multicast Support In VXLAN Fabrics</i></p> <p>Default MDT Address for TRM VRFs <input type="text"/> ? <i>IPv4 Multicast Address</i></p> <p>* Rendezvous-Points <input type="text" value="2"/> ? <i>Number of spines acting as Rendezvous-Point (RP)</i></p> <p>* RP Mode <input type="text" value="asm"/> ? <i>Multicast RP Mode</i></p> <p>* Underlay RP Loopback Id <input type="text" value="254"/> ? <i>(Min:0, Max:1023)</i></p> <p>Underlay Primary RP Loopback Id <input type="text"/> ? <i>Used for Bidir-PIM Phantom RP (Min:0, Max:1023)</i></p> <p>Underlay Backup RP Loopback Id <input type="text"/> ? <i>Used for Fallback Bidir-PIM Phantom RP (Min:0, Max:1023)</i></p> <p>Underlay Second Backup RP Loopback Id <input type="text"/> ? <i>Used for second Fallback Bidir-PIM Phantom RP (Min:0, Max:1023)</i></p> <p>Underlay Third Backup RP Loopback Id <input type="text"/> ? <i>Used for third Fallback Bidir-PIM Phantom RP (Min:0, Max:1023)</i></p> <p>* VRF Template <input type="text" value="Default_VRF_Universal"/> ? <i>Default Overlay VRF Template For Leafs</i></p> <p>* Network Template <input type="text" value="Default_Network_Universal"/> ? <i>Default Overlay Network Template For Leafs</i></p> <p>* VRF Extension Template <input type="text" value="Default_VRF_Extension_Universal"/> ? <i>Default Overlay VRF Template For Borders</i></p> <p>* Network Extension Template <input type="text" value="Default_Network_Extension_Universa"/> ? <i>Default Overlay Network Template For Borders</i></p> <p>* Underlay VTEP Loopback IP Range <input type="text" value="10.3.0.0/22"/> ? <i>Typically Loopback1 IP Address Range</i></p> <p>* Underlay RP Loopback IP Range <input type="text" value="10.254.254.0/24"/> ? <i>Anycast or Phantom RP IP Address Range</i></p> <p>* Layer 2 VXLAN VNI Range <input type="text" value="30000-49000"/> ? <i>Overlay Network Identifier Range (Min:1, Max:16777214)</i></p> <p>* Layer 3 VXLAN VNI Range <input type="text" value="50000-59000"/> ? <i>Overlay VRF Identifier Range (Min:1, Max:16777214)</i></p> <p>* Network VLAN Range <input type="text" value="2300-2999"/> ? <i>Per Switch Overlay Network VLAN Range (Min:2, Max:3967)</i></p> <p>* VRF VLAN Range <input type="text" value="2000-2299"/> ? <i>Per Switch Overlay VRF VLAN Range (Min:2, Max:3967)</i></p> <p>* VRF Lite Deployment <input type="text" value="Manual"/> ? <i>VRF Lite Inter-Fabric Connection Deployment Options</i></p>						

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

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

and *Deploying Networks and VRFs* in the Control chapter in *Cisco DCNM LAN Fabric Configuration Guide*.

Routed Fabric: You must disable the enable EVPN VXLAN Overlay field for Routed fabric (an IP fabric with no VXLAN encapsulation) creation. In a Routed Fabric, you can create and deploy networks. For more information, see [Overview of Networks in a Routed Fabric](#).

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

If a network or a VRF is created in a fabric, you cannot switch between VXLAN EVPN mode and Routed Fabric mode by selecting the **Enable EVPN VXLAN Overlay** check box. You need to delete these networks or VRFs to change the fabric setting.

Note that **Routed_Network_Universal Template** is only applicable to a Routed Fabric. When you convert the routed fabric to EVPN VXLAN fabric, set the network template and network extension template to the ones defined for EVPN VXLAN: **Default_Network_Universal** and **Default_Network_Universal**. If you have a customized template for EVPN VXLAN fabric, you can also choose to use it.

First Hop Redundancy Protocol: Specifies the FHRP protocol. Choose either **hsrp** or **vrrp**. This field is only applicable to a Routed Fabric.



Note

- After a network has been created, you cannot change this fabric setting. You should delete all networks, and then change the FHRP setting.
 - The rest of the fields in the EVPN tab section are only applicable if you enable the EVPN VXLAN Overlay.
-

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

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

If you want to enable the VXLAN OAM function on specific switches and disable on other switches in the fabric, you can use free-form configurations to enable OAM or 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: Specify the VRF template for creating VRFs, and the VRF extension template for enabling VRF extension to other fabrics.

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

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

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

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

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

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

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

General	EVPN	vPC	Protocols	Advanced	Manageability	Bootstrap	Configuration Backup
		* vPC Peer Link VLAN	<input type="text" value="3600"/>	<i>(i)</i>	VLAN for vPC Peer Link SVI (Min:2, Max:3967)		
		Make vPC Peer Link VLAN as Native VLAN	<input type="checkbox"/>	<i>(i)</i>			
		* vPC Peer Keep Alive option	<input type="text" value="management"/>	<i>(i)</i>	Use vPC Peer Keep Alive with Loopback or Management		
		* vPC Auto Recovery Time	<input type="text" value="360"/>	<i>(i)</i>	Auto Recovery Time In Seconds (Min:240, Max:3600)		
		* vPC Delay Restore Time	<input type="text" value="150"/>	<i>(i)</i>	vPC Delay Restore Time For vPC links in seconds (Min:1, Max:3600)		
		vPC Peer Link Port Channel Number	<input type="text" value="500"/>	<i>(i)</i>	Port Channel ID for vPC Peer Link (Min:1, Max:4096)		
		vPC IPv6 ND Synchronize	<input checked="" type="checkbox"/>	<i>(i)</i>	Enable IPv6 ND synchronization between vPC peers		
		Fabric wide vPC Domain Id	<input type="checkbox"/>	<i>(i)</i>	Enable to use same vPC Domain Id on all vPC pairs in the fabric		
		vPC Domain Id	<input type="text"/>	<i>(i)</i>	vPC Domain Id to be used on all vPC pairs in the fabric		
		Enable Qos for Fabric vPC-Peering	<input type="checkbox"/>	<i>(i)</i>	Qos on spines for guaranteed delivery of vPC Fabric Peering communication		
		Qos Policy Name	<input type="text"/>	<i>(i)</i>	Qos Policy name should be same on all spines		

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

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

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

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

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

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

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

Fabric wide vPC Domain Id: Enables the usage of same vPC Domain Id on all vPC pairs in the fabric. When you select this field, the **vPC Domain Id** field is editable.

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

Enable Qos for Fabric vPC-Peering - Enable QoS on spines for guaranteed delivery of vPC Fabric Peering communication.

Qos Policy Name - Specifies QoS policy name that should be same on all spines.

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

General	EVPN	vPC	Protocols	Advanced	Manageability	Bootstrap	Configuration Backup
* Routing Loopback Id <input type="text" value="0"/> ⓘ (Min:0, Max:1023)							
* VTEP Loopback Id <input type="text" value="1"/> ⓘ (Min:0, Max:1023)							
* BGP Maximum Paths <input type="text" value="4"/> ⓘ (Min:1, Max:64)							
Enable BGP Authentication <input type="checkbox"/> ⓘ							
BGP Authentication Key Encryption Type <input type="text" value="3"/> ⓘ BGP Key Encryption Type: 3 - 3DES, 7 - Cisco							
BGP Authentication Key <input type="text"/> ⓘ Encrypted BGP Authentication Key based on type							
Enable PIM Hello Authentication <input type="checkbox"/> ⓘ							
PIM Hello Authentication Key <input type="text"/> ⓘ 3DES Encrypted							
Enable BFD <input type="checkbox"/> ⓘ							
Enable BFD For BGP <input type="checkbox"/> ⓘ							
Enable BFD Authentication <input type="checkbox"/> ⓘ							
BFD Authentication Key ID <input type="text"/> ⓘ							
BFD Authentication Key <input type="text"/> ⓘ Encrypted SHA1 secret value							

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

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

BGP Maximum Paths - Specifies the BGP maximum paths.

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

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

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



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

Enable PIM Hello Authentication: Enables the PIM hello authentication.

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

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

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

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

```
feature bfd
```



Note After you upgrade from DCNM Release 11.2(1) with BFD enabled to DCNM Release 11.3(1), the 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:

General	EVPN	vPC	Protocols	Advanced	Manageability	Bootstrap	Configuration Backup
				* Intra Fabric Interface MTU	9216		(Min:576, Max:9216). Must be an even number
				* Layer 2 Host Interface MTU	9216		(Min:1500, Max:9216). Must be an even number
				* Power Supply Mode	ps-redundant		Default Power Supply Mode For The Fabric
				* CoPP Profile	strict		Fabric Wide CoPP Policy. Customized CoPP policy should be separately defined, when 'manual' is selected
				VTEP HoldDown Time	180		NVE Source Interface HoldDown Time (Min:1, Max:1500) in seconds
				* VRF Lite Subnet IP Range	10.33.0.0/16		Address range to assign P2P DCI Links
				* VRF Lite Subnet Mask	30		Mask for Subnet Range (Min:8, Max:31)
				Enable CDP for Bootstrapped Switch	<input type="checkbox"/>		Enable CDP on management interface
				Enable NX-API	<input checked="" type="checkbox"/>		Enable NX-API on port 443
				Enable NX-API on HTTP port	<input checked="" type="checkbox"/>		Enable NX-API on port 80
				Enable Strict Config Compliance	<input type="checkbox"/>		Enable bi-directional compliance checks to flag additional configs in the running config that are not in the intent/expected config
				Enable AAA IP Authorization	<input type="checkbox"/>		Enable only, when IP Authorization is enabled in the AAA Server
				Enable DCNM as Trap Host	<input checked="" type="checkbox"/>		Configure DCNM as a receiver for SNMP traps
				* Greenfield Cleanup Option	Disable		Switch Cleanup Without Reload When PreserveConfig=no
				Enable Default Queuing Policies	<input type="checkbox"/>		
				N9K Cloud Scale Platform Queuing Policy			Queuing Policy for all 92xx, -EX, -FX, -FX2, -FX3 series switches in the fabric
				N9K R-Series Platform Queuing Policy			Queuing Policy for all R-Series switches in the fabric

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

Layer 2 Host Interface MTU - Specifies the MTU for the layer 2 host interface. This value should be an even number.

Power Supply Mode: Choose the appropriate power supply mode.

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

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

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

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

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

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



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

- Click the **Manageability** tab.

The screenshot shows the 'Manageability' configuration tab. It contains the following fields and their help text:

- DNS Server IPs:** Comma separated list of IP Addresses(v4/v6)
- DNS Server VRFs:** One VRF for all DNS servers or a comma separated list of VRFs, one per DNS server
- NTP Server IPs:** Comma separated list of IP Addresses(v4/v6)
- NTP Server VRFs:** One VRF for all NTP servers or a comma separated list of VRFs, one per NTP server
- Syslog Server IPs:** Comma separated list of IP Addresses(v4/v6)
- Syslog Server Severity:** Comma separated list of Syslog severity values, one per Syslog server (Min:0, Max:7)
- Syslog Server VRFs:** One VRF for all Syslog servers or a comma separated list of VRFs, one per Syslog server
- AAA Freeform Config:** Note ! All configs should strictly match 'show run' output, with respect to case and newlines. Any mismatches will yield unexpected diffs during deploy.

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:

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

Hourly Fabric Backup ? Backup hourly or on Re-sync only if there is any config deployment since last backup

Scheduled Fabric Backup ? Backup at the specified time only if there is any config deployment since last backup

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

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

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

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

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

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

Select both the check boxes to enable both back up processes.
 The backup process is initiated after you click **Save**.



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

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

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

11. Click **ThousandEyes Agent** tab. This feature is supported on Cisco DCNM Release 11.5(3) only. For more information, refer to [Configuring Global Settings for ThousandEyes Enterprise Agent](#) section.

The fields on this tab are:



Note The fabric settings for ThousandEyes Agent overwrites the global settings and applies the same configuration for all the ThousandEyes Agent installed on switches in that fabric.


- **Enable Fabric Override for ThousandEyes Agent Installation:** Select the check box to enable the ThousandEyes Enterprise Agent on the fabric.
- **ThousandEyes Account Group Token:** Specifies ThousandEyes Enterprise Agent token ID for installation.
- **VRF on Switch for ThousandEyes Agent Collector Reachability:** Specifies the VRF data which provides internet reachability.
- **DNS Domain:** Specifies the switch DNS domain configuration.
- **DNS Server IPs:** Specifies the comma separated list of IP addresses (v4/v6) of Domain Name System (DNS) server. You can enter a maximum of three IP addresses for the DNS Server.

- **NTP Server IPs:** Specifies comma separated list of IP addresses (v4/v6) of Network Time Protocol (NTP) server. You can enter a maximum of three IP addresses for the NTP Server.
- **Enable Proxy for Internet Access:** Select the check box to enable the proxy setting for NX-OS switch internet access.
- **Proxy Information:** Specifies the proxy server port information.
- **Proxy Bypass:** Specifies the server list for which proxy is bypassed.

VXLAN Fabric With eBGP Underlay – Pointers

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

Deploying Fabric Underlay eBGP Policies



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

The two different types of fabrics are:

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

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

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

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



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

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



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

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

Deploying Fabric Overlay eBGP Policies

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

Deploying Spine Switch Overlay Policies

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

Add Policy
✕

* Priority (1-1000):

* Policy: ▼

General

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

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

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

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

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

Variables:

The fields on the screen are:

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

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

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



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

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

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

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

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

Deploying Leaf Switch Overlay Policies

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

Add Policy
✕

* Priority (1-1000):

* Policy: ▼

General

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

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

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

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

Variables:

The fields on the screen are:

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

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

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

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

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

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

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