



## Configuring a Fabric with eBGP Underlay

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

- [Creating a New Fabric for eBGP-Based Underlay, on page 1](#)
- [Overview of Networks in a Routed Fabric, on page 19](#)

### Creating a New Fabric for eBGP-Based Underlay

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

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

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

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

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

The fields are explained:

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

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

Add Fabric ✕

\* Fabric Name :

\* Fabric Template :

---

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

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

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

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

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

\* Underlay Routing Loopback IP Range  ? Typically Loopback0 IP Address Range

\* Underlay Subnet IP Range  ? Address range to assign Numbered and Peer Link SVI IPs

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

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

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

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

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

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

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

The fabric is identified by the spine switch AS number.

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

**Manual Underlay IP Address Allocation** – Select this check box to disable Dynamic Underlay IP Address Allocations.

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

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

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

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

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

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

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

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

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

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

You can convert a routed fabric to a VXLAN enabled fabric by selecting this option. When the fabric is VXLAN enabled, you can create and deploy overlay networks or VRFs. The procedure for creating and deploying networks or VRFs is the same as in Easy\_Fabric\_11\_1.



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

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

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

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

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




---

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

---

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

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

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




---

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

---

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

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

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

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

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

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

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

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




---

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

---

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

The screenshot shows the 'Protocols' configuration tab with the following fields and options:

- \* Routing Loopback Id**: Input field with value 0. Help icon: (Min:0, Max:1023)
- \* VTEP Loopback Id**: Input field with value 1. Help icon: (Min:0, Max:1023)
- Enable BGP Authentication**: Check box (unchecked). Help icon: ?
- BGP Authentication Key Encryption Type**: Dropdown menu. Help icon: ? BGP Key Encryption Type: 3 - 3DES, 7 - Cisco
- BGP Authentication Key**: Text input field. Help icon: ? Encrypted BGP Authentication Key based on type
- Enable BFD**: Check box (unchecked). Help icon: ?
- Enable BFD For BGP**: Check box (unchecked). Help icon: ?
- Enable BFD Authentication**: Check box (unchecked). Help icon: ?
- BFD Authentication Key ID**: Text input field. Help icon: ?
- BFD Authentication Key**: Text input field. Help icon: ? 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.

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

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

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



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

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

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

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

```
feature bfd
```



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

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

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

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

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

**BFD Authentication Key ID:** Specifies the BFD authentication key ID for the interface authentication.

**BFD Authentication Key:** Specifies the BFD authentication key.

For information about how to retrieve the BFD authentication parameters, see *Retrieving the Encrypted BFD Authentication Key*, in *Cisco DCNM LAN Fabric Configuration Guide*.

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

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 NX-API	<input checked="" type="checkbox"/>	? Enable the NX-API feature	
				Enable NX-API on HTTP	<input checked="" type="checkbox"/>	? Enable NX-API on HTTP port	
				Enable Strict Config Compliance	<input type="checkbox"/>	?	
				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"/>	?	
				* 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 series switches in the fabric	
				N9K R-Series Platform Queuing Policy		? Queuing Policy for all R-Series switches in the fabric	
				Other N9K Platform Queuing Policy		? Queuing Policy for all other switches in the fabric	
				Leaf 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.	
				Spine 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.	
				Intra-fabric Links Additional Config		? Note ! All configs should strictly match 'show run' output, with respect to case and newlines. Any mismatches will yield unexpected diffs during deploy.	

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

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

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

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

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

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

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

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

**Enable Strict Config Compliance** - Enable the Strict Config Compliance feature by selecting this check box.

For Strict Configuration Compliance, see *Enhanced Monitoring and Monitoring Fabrics Guide*.





**Note** If Strict Config Compliance is enabled in a fabric, you cannot deploy Network Insights for Resources on Cisco DCNM.

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

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

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

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

Review the actual queuing policies by opening the policy file in the template editor. From Cisco DCNM Web UI, choose **Control > Template Library**. Search for the queuing policies by the policy file name, for example, **queuing\_policy\_default\_8q\_cloudscale**. Choose the file and click the **Modify/View template** icon to edit the policy.

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

**N9K Cloud Scale Platform Queuing Policy:** Choose the queuing policy from the drop-down list to be applied to all Cisco Nexus 9200 Series Switches and the Cisco Nexus 9000 Series Switches that ends with EX, FX, and FX2 in the fabric. The valid values are **queuing\_policy\_default\_4q\_cloudscale** and **queuing\_policy\_default\_8q\_cloudscale**. Use the **queuing\_policy\_default\_4q\_cloudscale** policy for FEXes. You can change from the **queuing\_policy\_default\_4q\_cloudscale** policy to the **queuing\_policy\_default\_8q\_cloudscale** policy only when FEXes are offline.

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

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

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

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

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

8. Click the **Manageability** tab.

General	EVPN	vPC	Protocols	Advanced	Manageability	Bootstrap	Configuration Backup
<p><b>DNS Server IPs</b> <input type="text"/> ? Comma separated list of IP Addresses(v4/v6)</p> <p><b>DNS Server VRFs</b> <input type="text"/> ? One VRF for all DNS servers or a comma separated list of VRFs, one per DNS server</p> <p><b>NTP Server IPs</b> <input type="text"/> ? Comma separated list of IP Addresses(v4/v6)</p> <p><b>NTP Server VRFs</b> <input type="text"/> ? One VRF for all NTP servers or a comma separated list of VRFs, one per NTP server</p> <p><b>Syslog Server IPs</b> <input type="text"/> ? Comma separated list of IP Addresses(v4/v6)</p> <p><b>Syslog Server Severity</b> <input type="text"/> ? Comma separated list of Syslog severity values, one per Syslog server (Min:0, Max:7)</p> <p><b>Syslog Server VRFs</b> <input type="text"/> ? One VRF for all Syslog servers or a comma separated list of VRFs, one per Syslog server</p> <p><b>AAA Freeform Config</b> <input type="text"/> ? Note ! All configs should strictly match 'show run' output, with respect to case and newlines. Any mismatches will yield unexpected diffs during deploy.</p>							

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.

General	EVPN	vPC	Protocols	Advanced	Manageability	Bootstrap	Configuration Backup
<b>Enable Bootstrap</b> <input type="checkbox"/> ? Automatic IP Assignment For POAP							
Enable Local DHCP Server <input type="checkbox"/> ? Automatic IP Assignment For POAP From Local DHCP Server							
DHCP Version <input type="text"/> ?							
DHCP Scope Start Address <input type="text"/> ? Start Address For Switch Out-of-Band POAP							
DHCP Scope End Address <input type="text"/> ? End Address For Switch Out-of-Band POAP							
Switch Mgmt Default Gateway <input type="text"/> ? Default Gateway For Management VRF On The Switch							
Switch Mgmt IP Subnet Prefix <input type="text"/> ? (Min:8, Max:30)							
Switch Mgmt IPv6 Subnet Prefix <input type="text"/> ? (Min:64, Max:126)							
Enable AAA Config <input type="checkbox"/> ? Include AAA configs from Manageability tab during device bootstrap							
Bootstrap Freeform Config <input type="text"/> ? Note ! All configs should strictly match 'show run' output, with respect to case and newlines. Any mismatches will yield unexpected diffs during deploy.							
DHCPv4/DHCPv6 Multi Subnet Scope <input type="text"/> ? 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							

**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.

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

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

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

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


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

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

#### **VXLAN Fabric With eBGP Underlay – Pointers**

- The supported roles are leaf, spine, and border leaf.
- On the border device, VRF-Lite is supported with manual mode. There is no Multi-Site support for external connectivity.
- TRM is supported.
- You must apply policies on the leaf and spine switches for a functional fabric.
- When you convert a non-VXLAN (or routed fabric) to a VXLAN enabled fabric, you can create and deploy overlay networks and VRFs.

## Applying Policies On A Fabric With An eBGP Underlay



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

This topic covers the following:

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

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

## Deploying Fabric Underlay Policies

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

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




---

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

---

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




---

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

---

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

## Deploying Fabric Overlay Policies

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

## Deploying Spine Switch Overlay Policies

Add the `ebgp_overlay_spine_all_neighbor` policy on the spine switch n9k-29.



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 for this field.

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

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

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

## Deploying Leaf Switch Overlay Policies

Add the `ebgp_overlay_leaf_all_neighbor` policy on all the leaf switches, to establish eBGP overlay peering towards the spine switch.

Add Policy
✕

\* Priority (1-1000):

\* Policy:  ▼

General

\* Spine IP List

\* BGP Update-Source Interface

Enable Tenant Routed Multicast

Enable BGP Authentication

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

? Source of BGP session and updates

? For Overlay Multicast Support In VXLAN Fabrics

? 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 for this field.

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

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

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

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

## Dual-AS Fabric Deployment

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

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

### Additional Pointers

- Brownfield migration is not supported for eBGP fabric.

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

## Overview of Networks in a Routed Fabric

From Cisco DCNM Release 11.3(1), you can create a top-down network configuration for a routed fabric using DCNM. A routed fabric is run in one VRF, which is the default VRF. Note that creating VRFs manually is disabled for a routed fabric. Since the fabric is an IPv4 fabric, IPv6 address within the network is not supported. In a routed fabric, a network can only be attached to one device or a pair of vPC devices, unless it is a Layer 2 only network.



---

**Note** A routed fabric network configuration will not be put under a config-profile.

---

When the eBGP fabric is configured as Routed Fabric (EVPN is disabled), at the fabric level, you can select the first hop redundancy protocol (FHRP) for host traffic to be either HSRP or VRRP. HSRP is the default value.

For a vPC pair, DCNM generates network level HSRP or VRRP configuration based on the fabric setting. If HSRP is chosen, each network is configured with one HSRP group, and the HSRP VIP address. By default, all the networks will share the same HSRP group number allocated by DCNM, while you can overwrite it per network. VRRP support is similar to HSRP.

### Guidelines

- HSRP authentication or VRRP authentication is not supported. If you want to use authentication, you can enter the applicable commands in the network freeform config.
- vPC peer gateway can be used to minimize peer link usage in the case that some third-party devices ignore the HSRP virtual-MAC and use the ARP packet source MAC for ARP learning. In Routed fabric mode, DCNM generates vPC peer gateway command for VPC devices.
- For an eBGP fabric, changing between routed fabric type and EVPN fabric type, or HSRP and VRRP, is not allowed with the presence of networks and VRFs. You need to undeploy and delete these networks and VRFs before changing the fabric type or FHRP. For more information, see [Undeploying Networks for the Standalone Fabric](#) and [Undeploying VRFs for the Standalone Fabric](#).
- After the upgrade from DCNM Release 11.2(1) to 11.3(1), if the fabric was running in Routed Fabric mode previously, the default fabric values such as FHRP protocol and network VLAN range are internally set for a Routed Fabric. You need to edit the fabric settings if you want to configure different values.

- Before deploying a network configuration, you need to update the FHRP protocol fabric setting and click **Save & Deploy**.

## Creating and Deploying a Network in a Routed Fabric

This procedure shows how to create and deploy a network in a routed fabric.

### Before you begin

Create a routed fabric and deploy the necessary leaf and spine policies.

### Procedure

- Step 1** Navigate to **Control > Networks**.
- Step 2** From the **SCOPE** drop-down list, choose a routed fabric.
- Step 3** Click the **Add** button in the **Networks** window to create a network.

Create Network
✕

---

▼ Network Information

\* Network Name

Layer 2 Only

\* Network Template

VLAN ID

---

▼ Network Profile

General

Advanced

IPv4 Gateway/NetMask  ? example 192.0.2.1/24. Address for VIP or st

Intf IPv4 addr on active  ? example 192.0.2.2. Interface IP address on

Intf IPv4 addr on stan...  ? example 192.0.2.3. Interface IP address on

Vlan Name  ? if > 32 chars enable:system vlan long-name

Interface Description  ? For interface on the standalone, or the activ

Standby Intf Descripti...  ? For interface on the standby/backup switch

MTU for L3 interface  ? 68-9216

Routing Tag  ? 0-4294967295

**Create Network**

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

**Layer 2 Only:** Optional. Specifies whether the network is a Layer 2 only network. FHRP configuration is not generated in a Layer 2 only network.

**Note** When an L3 Network template is attached to a standalone device, no FHRP configuration is generated.

**Network Template:** Select the **Routed\_Network\_Universal** template.

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

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

#### General tab

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

**Intf IPv4 addr on active:** Specifies the IPv4 interface address on an active/master device in a vPC pair. This field is applicable only when you are creating and deploying a network for a vPC pair of devices.

**Intf IPv4 addr on standby:** Specifies the IPv4 interface address on a standby/backup device in a vPC pair. This field is applicable only when you are creating and deploying a network for a vPC pair of devices.

**Note** The IPv4 gateway address and interface addresses should be in the same subnet.

The following fields under the **General** tab are optional:

**Vlan Name:** Specifies the VLAN name.

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

**Standby Intf Description:** Specifies the description for the standby interface in a vPC pair.

**MTU for the L3 interface:** Enter the MTU for Layer 3 interfaces.

**Routing Tag:** Specifies the routing tag that is associated with each gateway IP address prefix.

**Advanced** tab: This tab is applicable only when you are creating and deploying a network for a vPC pair of devices.

▼ Network Profile

General	Advanced
	<p>First Hop Redundanc... <input type="text" value="hsrp"/> ? <i>Read-only, from fabric setting</i></p> <p>Active/master Switch Priority <input type="text" value="120"/> ?</p> <p>Standby/backup Switch Priority <input type="text" value="100"/> ?</p> <p>Enable Preempt <input checked="" type="checkbox"/> ? <i>Overthrow lower priority Active routers</i></p> <p>HSRP/VRRP Group # <input type="text" value="1"/> ?</p> <p>Virtual MAC Address <input type="text" value="AA11.2222.3333"/> ?</p> <p>HSRP Version <input type="text" value="1"/> ▼ ? <i>1 or 2</i></p>

[Create Network](#)

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

**Active/master Switch Priority:** Specifies the priority of the active or master device.

**Standby/backup Switch Priority:** Specifies the priority of the standby or backup device. The default value is 100. Note that this default value is not displayed when you preview the network configuration before deployment.

**Enable Preempt:** Specifies whether the standby/backup device can preempt a active or master device.

**HSRP/VRRP Group #:** Specifies the HSRP or VRRP group number. By default, HSRP group number is 1.

**Virtual MAC Address:** Optional. Specifies the virtual MAC address. By default, VMAC is internally generated based on the HSRP group number (0000.0c9f.f000 + group number). The virtual MAC address is only applicable when **hsrp** is selected in the fabric settings.

**HSRP Version:** Specifies the HSRP version. The default value is 1. The **HSRP version** field is only applicable for HSRP.

**Step 4** Click **Create Network**.

**Step 5** In the **Networks** window, select the check box next to a network and click **Continue**.

**Note** A non Layer 2 network can be only applied to a vPC pair of devices or a single device. For example, if you have deployed a network on a single device, you cannot deploy the same network on another device or a vPC pair of devices.

**Step 6** Select a device or a vPC pair to deploy a network.

**Note** In a routed fabric, when you try to attach a network on a vPC pair without active or standby IP addresses, an error is displayed saying that the IP address fields are not filled. After you add the IP addresses and save the network, the network state changes to **PENDING** without the need to attach the network again.

**Step 7** In the **Network Attachment** window, for a vPC pair, assign the active state for a device.

Enter **true** under the **isActive** column for an active device and **false** for a standby device.

Click **Save**.

Network Attachment - Attach networks for given switch(es) ✕

Fabric Name: bgp-routed

Deployment Options

① Select the row and click on the cell to edit and save changes

	VLAN	Interfaces	CLI Freeform	Status	isActive
MyNetwork_30000	100	... Ethernet1/1	Freeform config	NA	true
	100	... Ethernet1/1	Freeform config	NA	false

Save

**Note** In a routed fabric, when you edit a deployed network and save without making any changes, the status of the network changes to **Pending**. Similarly, if a **Network Attachment** window is opened for a deployed network, and saved without any changes, the status of the network changes to **Pending**. In these cases, click the **Preview** icon to preview the config. This action changes the network status back to **Deployed**.

**Step 8** (Optional) Click the **Preview** icon to preview the configs that will be deployed on devices. The **Preview Configuration** window is displayed.

### Preview Configuration ✕

**Select a Switch:**

n9k-30
▼

**Select a Network**

MyNetwork\_30000
▼

**Generated Configuration:**

```

interface ethernet1/1
  switchport trunk allowed vlan add 100
interface Vlan100
  no ip redirects
  no ipv6 redirects
  ip address 100.1.1.2/24 tag 12345
  hsrp 1
    ip 100.1.1.1
    priority 120
    mac-address aa11.2222.3333
  preempt
  mtu 8000
  description test100_int
  no shutdown
vlan 100
  name test100
configure terminal
          
```

**Step 9** Click the **Deploy** button in the **Network / VRF Deployment** window. You can also deploy the network by navigating to the **Fabric Builder** window and clicking the **Deploy** button.

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

From DCNM Release 11.3(1), you can use an inter-fabric link to connect a route fabric to an edge router. This link configures an IP address on the physical interface and establish eBGP peering with the edge router on default vrf. The BGP configuration includes advertising default route to leaf switches.



**Note** The **Fabric Monitor Mode** check box in the external fabric settings can be unchecked. Unchecking the **Fabric Monitor Mode** check box enables DCNM to deploy configurations to the external fabric. For more information, see [Creating an External Fabric](#).

## Procedure

- Step 1** Navigate to **Control > Fabric Builder**.
- Step 2** Click a routed a fabric in the **Fabric Builder** window.
- Step 3** Click **Tabular view** in the **Actions** panel that is displayed at the left part of the window.
- Step 4** Click the **Links** tab.
- Step 5** Click the **Add** icon to add a link.

The **Link Management – Add Link** window is displayed.

Link Management - Add Link
✕

\* Link Type: Inter-Fabric

\* Link Sub-Type: ROUTED\_FABRIC

\* Link Template: ext\_routed\_fabric

\* Source Fabric: routed1

\* Destination Fabric: ext1

\* Source Device: n9k-32

\* Source Interface: Ethernet1/48

\* Destination Device: n6k-66

\* Destination Interface: Ethernet1/48

▼ Link Profile

General

Advanced

\* BGP Local ASN: 20 Local BGP Autonomous System Number

\* IP Address/Mask: 48.1.1.1/24 IP address with mask (e.g. 192.168.10.1/24)

\* BGP Neighbor IP: 48.1.1.2 Neighbor IP address

\* BGP Neighbor ASN: 200 Neighbor BGP Autonomous System Number

\* BGP Maximum Paths: 1 if, choose the maximum path platform supports

Save

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

**Link Sub-Type** – This field populates the IFC type. Choose **ROUTED\_FABRIC** from the drop-down list.

**Link Template:** The link template is populated. The templates are autopopulated with corresponding pre-packaged default templates that are based on your selection. For a routed fabric, the **ext\_routed\_fabric** template is populated.



**Source Fabric** - This field is prepopulated with the source fabric name.

**Destination Fabric** - Choose the destination fabric from this drop-down box.

**Source Device** and **Source Interface** - Choose the source device and Ethernet or port channel interface that connects to the destination device.

**Destination Device** and **Destination Interface**—Choose the destination device and Ethernet or port channel interface that connects to the source device.

Based on the selection of the source device and source interface, the destination information is autopopulated based on Cisco Discovery Protocol information, if available. There is an extra validation performed to ensure that the destination external device is indeed part of the destination fabric.

**General** tab in the Link Profile section.

**BGP Local ASN:** In this field, the AS number of the leaf is autopopulated if you have created and applied the `leaf_bgp_asn` policy.

**IP Address/Mask:** Fill up this field with the IP address of the source interface that connects to the destination device.

**BGP Neighbor IP:** Fill up this field with the IP address of the destination interface.

**BGP Neighbor ASN:** In this field, the AS number of the destination device is autopopulated.

**BGP Maximum Paths:** Specifies the maximum supported BGP paths.

The **Advanced** tab contains the following optional fields:

**Source Interface Description** and **Destination Interface Description** – Describe the links for later use. After **Save & Deploy**, this description will reflect in the running configuration.

**Source Interface Freeform CLIs** and **Destination Interface Freeform CLIs:** Enter the freeform configurations specific to the source and destination interfaces. You should add the configurations as displayed in the running configuration of the switch, without indentation. For more information, refer to *Enabling Freeform Configurations on Fabric Switches*.

- Step 6** Click **Save** to finish adding a link.
- Step 7** Click the **Back** icon to navigate back to the Fabric Builder window.
- Step 8** Right-click the device which is connecting to the edge router in the external fabric, and select **Deploy Config**.
- Step 9** In the **Config Deployment** window, click **Deploy Config**.
- Step 10** Navigate to the external fabric in the **Fabric Builder** window, and click **Tabular view** in the **Actions** panel. Click the **Links** tab to see all the links for the external fabric.

You can see the inter-fabric link that has been created.

**Note** The inter-fabric link is created if the External fabric is not in the monitor mode.

- Step 11** Click the **Back** icon twice to navigate back to the **Fabric Builder** window.
- Step 12** Click the external fabric connecting to the routed fabric.
- Step 13** Right-click the device which is connecting to the routed fabric, and select **Deploy Config**.
- Step 14** In the **Config Deployment** window, click **Deploy Config**.

