



Control

This chapter contains the following topics:

- [Fabrics, on page 1](#)
- [Management, on page 315](#)
- [Template Library, on page 328](#)
- [Image Management, on page 369](#)
- [Endpoint Locator, on page 390](#)
- [ThousandEyes Enterprise Agent, on page 390](#)
- [Layer 4-Layer 7 Service, on page 391](#)
- [Cross Site Scripting \(XSS\) threat and mitigation, on page 391](#)

Fabrics

The following terms are referred to in the document:

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

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

This section contains the following topics:

VXLAN BGP EVPN Fabrics Provisioning

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

period of time. The set of parameters exposed in the Fabric Settings allow users to tailor the fabric to their preferred underlay provisioning options.

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

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

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

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

Create, edit, and delete a fabric:

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

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

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

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

Underlay provisioning:

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

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

Overlay network provisioning.

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

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

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

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

Guidelines for VXLAN BGP EVPN Fabrics Provisioning

- For any switch to be successfully imported into DCNM, the user specified for discovery/import, should have the following permissions:
 - SSH access to the switch
 - Ability to perform SNMPv3 queries
 - Ability to run the **show** commands including show run, show interfaces, etc.
- The switch discovery user need not have the ability to make any configuration changes on the switches. It is primarily used for read access.
- When an invalid command is deployed by DCNM to a device, for example, a command with an invalid key chain due to an invalid entry in the fabric settings, an error is generated displaying this issue. This error is not cleared after correcting the invalid fabric entry. You need to manually cleanup or delete the invalid commands to clear the error.

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

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

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

When you right-click on a switch icon, you can use the **Deploy Config** option to deploy per switch configurations. This option is a local operation for a switch, that is, the expected configuration or intent

for a switch is evaluated against 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, on page 65](#).

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

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



Note

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

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

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

Enable IPv6 Underlay: Enable the IPv6 underlay feature. For information, see [IPv6 Underlay Support for Easy Fabric, on page 65](#).

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.

General	Replication	vPC	Protocols	Advanced	Resources	Manageability	Bootstrap	Configuration Backup
	* Replication Mode	Multicast						?
	* Multicast Group Subnet	239.1.1.0/25						?
	Enable Tenant Routed Multicast (TRM)	<input type="checkbox"/>						?
	Default MDT Address for TRM VRFs							?
	* Rendezvous-Points	2						?
	* RP Mode	asm						?
	* Underlay RP Loopback Id	254						?
	Underlay Primary RP Loopback Id							?
	Underlay Backup RP Loopback Id							?
	Underlay Second Backup RP Loopback Id							?
	Underlay Third Backup RP Loopback Id							?

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

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

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

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

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

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

For more information, see [Overview of Tenant Routed Multicast, on page 148](#).

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

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

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

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

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

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

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

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

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

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

You can enable the advertise PIP feature on a specific vPC as well. For more information, see [Advertising PIP on vPC, on page 193](#).

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

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

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

Enable QoS for Fabric vPC-Peering - Enable QoS on spines for guaranteed delivery of vPC Fabric Peering communication. For more information, see [QoS for Fabric vPC-Peering, on page 185](#).



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.

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, on page 206](#).

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
				* VRF Template	Default_VRF_Universal	?	Default Overlay VRF Template For Leafs	
				* Network Template	Default_Network_Universal	?	Default Overlay Network Template For Leafs	
				* VRF Extension Template	Default_VRF_Extension_Universal	?	Default Overlay VRF Template For Borders	
				* Network Extension Template	Default_Network_Extension_Universa	?	Default Overlay Network Template For Borders	
				Site Id		?	For EVPN Multi-Site Support (Min:1, Max: 281474976710655). Defaults to Fabric ASN	
				* 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 provided when 'manual' is selected	
				VTEP HoldDown Time	180	?	NVE Source Interface HoldDown Time (Min:1, Max:1500) in seconds	

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

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



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

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



Note Ensure that **Enable Tenant DHCP** is enabled before enabling DHCP related parameters in the overlay 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](#), on page 51.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Enable MACsec - Enables MACsec for the fabric. For more information, see [MACsec Support in Easy Fabric and eBGP Fabric, on page 146](#).

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.

The screenshot shows the 'Resources' tab in a configuration interface. It contains several input fields for defining IP address ranges and VNI ranges, each with a help icon (i) and a descriptive tooltip. The fields are:

- Manual Underlay IP Address Allocation**: A checkbox with a tooltip: "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**: (i)
- * 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))

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

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:



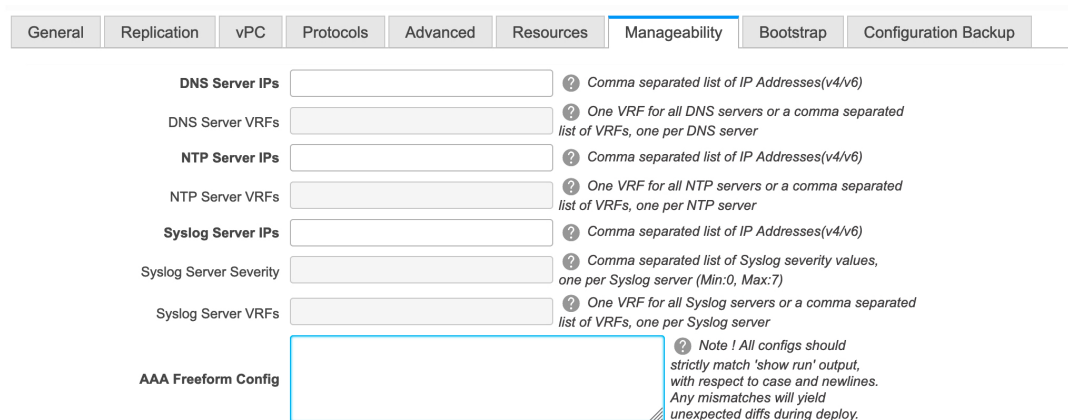
Note When you update a range of values, ensure that it does not overlap with other ranges. You should only update one range of values at a time. If you want to update more than one range of values, do it in separate instances. For example, if you want to update L2 and L3 ranges, you should do the following.

- 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 screenshot shows the 'Manageability' configuration tab. It contains the following fields and their descriptions:

- 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**: A large text area for freeform configurations.

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 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.

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

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, on page 303](#).

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 backup 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 Setting for ThousandEyes Enterprise Agent on Cisco DCNM](#).

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](#), on page 37 and [Pre-provisioning an Ethernet Interface](#), on page 41.



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

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

Discovering Existing Switches

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



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

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 navigation tabs, there is a 'Back' button and an 'Import into fabric' button. A table lists discovered switches with columns for Name, IP Address, Model, Version, Status, and Progress. The 'leaf-91' switch is selected, and its status is 'manageable'.

<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.



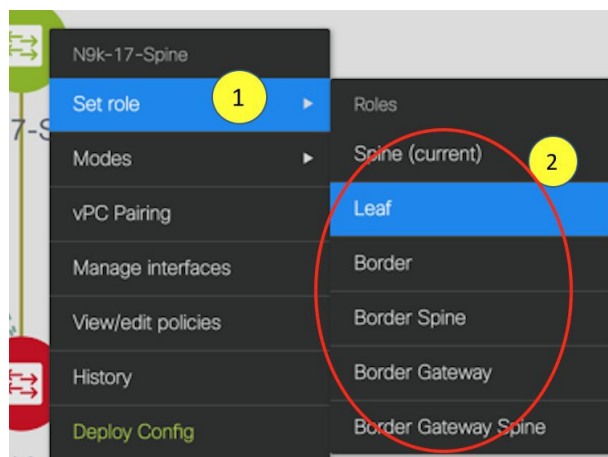
Note You will encounter the following errors during switch discovery sometimes.

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.



- 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.

- 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.

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.

- 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.
- 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.

- 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](#), on page 37.

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.

Control

31

Inventory Management ✕

Discover Existing Switches | PowerOn Auto Provisioning (POAP)

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

+ 🔄 ↺ * Admin Password * Confirm Admin Password 🔒

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

Close

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

Inventory Management ✕

Discover Existing Switches | PowerOn Auto Provisioning (POAP)

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

+ 🔄 ↺ * Admin Password * Confirm Admin Password 🔒

Serial Number | Model

No Data available

Discovery Credentials ✕

*Discovery Username:

*Discovery Password:

*Confirm Discovery Password:

OK Clear

Close

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

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

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

+ ⌵ ✎ ✕ ↑ ↓ 👁 🔄 📄 Deploy					
	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

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

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



- Note**
- Changing of the switch role is allowed only before executing **Save & Deploy**.
 - Starting from DCNM 11.1(1), switch roles can be changed if there are no overlays on the switches, but only as per the list of allowed switch role changes given at [Switch Operations, on page 157](#).

- **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](#)].

Pre-provisioning Support in DCNM 11

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

Pre-provisioning a Device

From Cisco DCNM Release 11.2, you can provision devices in advance.



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

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

- When a device is being pre-provisioned has breakout links, you need to specify the corresponding breakout command along with the switch's model and gateway in the **Data** field in the **Add a new device to pre-provisioning** window in order to generate the breakout PTI.

Note the following guidelines:

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

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

- {"modulesModel": ["N9K-C93180LC-EX"], "gateway": "10.1.1.1/24"}
- {"modulesModel": ["N9K-C93180LC-EX"], "breakout": "interface breakout module 1 port 1 map 10g-4x", "portMode": "hardware profile portmode 4x100G+28x40G", "gateway": "172.22.31.1/24" }
- {"modulesModel": ["N9K-X9736C-EX", "N9K-X9732C-FX", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-SUP-B+", "N9K-SC-A", "N9K-SC-A"], "gateway": "172.22.31.1/24"}
- {"breakout": "interface breakout module 1 port 50 map 10g-4x", "gateway": "172.16.1.1/24", "modulesModel": ["N9K-C93180YC-EX"]}
- {"modulesModel": ["N9K-X9732C-EX", "N9K-X9732C-EX", "N9K-C9504-FM-E", "N9K-C9504-FM-E", "N9K-SUP-B", "N9K-SC-A", "N9K-SC-A"], "gateway": "172.29.171.1/24", "breakout": "interface breakout module 1 port 1,11,19 map 10g-4x; interface breakout module 1 port 7 map 25g-4x"}
- {"modulesModel": ["N9K-C93180LC-EX"], "gateway": "10.1.1.1/24", "breakout": "interface breakout module 1 port 1-4 map 10g-4x", "portMode": "hardware profile portmode 48x25G + 2x100G + 4x40G"}

Procedure

-
- Step 1** Click **Control > Fabric Builder**.
- The **Fabric Builder** screen is displayed.
- Step 2** Click within the fabric box.
- Step 3** From the Actions panel, click the **Add switches** option.
- The **Inventory Management** screen is displayed.

- Step 4** Click the **POAP** tab.
- Step 5** In the **POAP** tab, do the following:
- Click + from the top left part of the screen.
The Add a new device screen comes up.
 - Fill up the device details as shown in the screenshot.
 - Click **Save**.

The screenshot shows a dialog box titled "Add a pre-provisioning device" with the following fields and values:

- *Serial Number: FDO21331SND
- *Model: N9K-93180YC-EX
- *Version: 7.0(3)5(2)
- *IP Address: 1.1.1.1
- *Hostname: LEAF1
- *Data: {"modulesModel": ["N9K-93180YC-EX"]}

Below the Data field, there is a red note: "ⓘ For more than one module, use commas to separate them. Please refer online help for more examples. Eg: {"modulesModel": ["N9K-C93180LC-EX"], "gateway": "10.1.1.1/24", "breakout": "interface breakout module 1 port 1-4 map 10g-4x", "portMode": "hardware profile portmode 48x25G + 2x100G + 4x40G"}"

At the bottom of the dialog, there are "Save" and "Clear" buttons.

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

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

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

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

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

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

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

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

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

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

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

Inventory Management

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

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

Bootstrap

* Admin Password * Confirm Admin Password

<input checked="" type="checkbox"/>	Serial Number	Model	Version	IP Address	Hostname
<input checked="" type="checkbox"/>	SN	N9K-3455	7.0(2)	10.1.1.1	leaf1

The leaf1 device appears in the fabric topology.

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

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

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

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

Create intra-link policy using one of the templates:

- **int_pre_provision_intra_fabric_link** to automatically generate intra fabric interface configuration with DCNM allocated IP addresses
- **int_intra_fabric_unnum_link_11_1** if you are using unnumbered links
- **int_intra_fabric_num_link_11_1** if you want to manually assign IP addresses to intra-links

Click **Save & Deploy**.

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

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

For more information, see [Return Material Authorization \(RMA\), on page 210](#).

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

You need to click **Save & Deploy** in the fabric after one or more switches are online to provision the host ports. This action must be performed before overlays are provisioned for the host port attachment.

Pre-provisioning an Ethernet Interface

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



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

Before you begin

Make sure that you have a preprovisioned device in your fabric. For information, see [Pre-provisioning a Device](#), on page 37.

Procedure

- Step 1** Navigate to the fabric containing the pre-provisioned device from the **Fabric Builder** window.
- Step 2** Right click the pre-provisioned device and select **Manage Interfaces**.
You can also navigate to the Interfaces window by selecting **Control > Fabrics > Interfaces**. From the Scope drop-down list, select the fabric containing the pre-provisioned device.
- Step 3** Click **Add**.
- Step 4** Enter all the required details in the **Add Interface** window.

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

Select a device: Select the pre-provisioned device.

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

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

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

For more information, see [Adding Interfaces, on page 220](#).

Step 5 Click **Save**.

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

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

Pre-provisioning a vPC Pair

Before you begin

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

Procedure

Step 1 Import both the devices into the fabric.

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

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

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 checked="" type="checkbox"/>	Serial Number	Model	Version	IP Address	Hostname	Gateway
<input checked="" type="checkbox"/>	FGE2035RRY	N9K-C93180LC-EX	9.3(5)	10.1.1.11	leaf2	10.1.1.1/24
<input checked="" type="checkbox"/>	FGE2035RRX	N9K-C93180LC-EX	9.3(5)	10.1.1.10	leaf1	10.1.1.1/24

Close

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

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

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

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

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

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

Control / Fabrics / Interfaces

Interfaces

<input type="button" value="+"/> <input type="button" value="↕"/> <input type="button" value="✎"/> <input type="button" value="✕"/> <input type="button" value="↑"/> <input type="button" value="↓"/> <input type="button" value="👁"/> <input type="button" value="🔄"/> <input type="button" value="📄"/> <input type="text" value="Interface Group"/> <input type="button" value="Deploy"/>					
	Device Name	Name	Admin	Oper	Reason
	<input type="text" value="leaf"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
<input type="checkbox"/>	leaf2	Mgmt0			Not discov
<input type="checkbox"/>	leaf2	Ethernet1/45			Not discov
<input type="checkbox"/>	leaf2	Ethernet1/44			Not discov
<input type="checkbox"/>	leaf1	Mgmt0			Not discov
<input type="checkbox"/>	leaf1	Ethernet1/45			Not discov
<input type="checkbox"/>	leaf1	Ethernet1/44			Not discov

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

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

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

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

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

Link Management - Add Link
✕

* Link Type

* Link Sub-Type

* Link Template

* Source Fabric

* Destination Fabric

* Source Device

* Source Interface

* Destination Device

* Destination Interface

▼ Link Profile

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

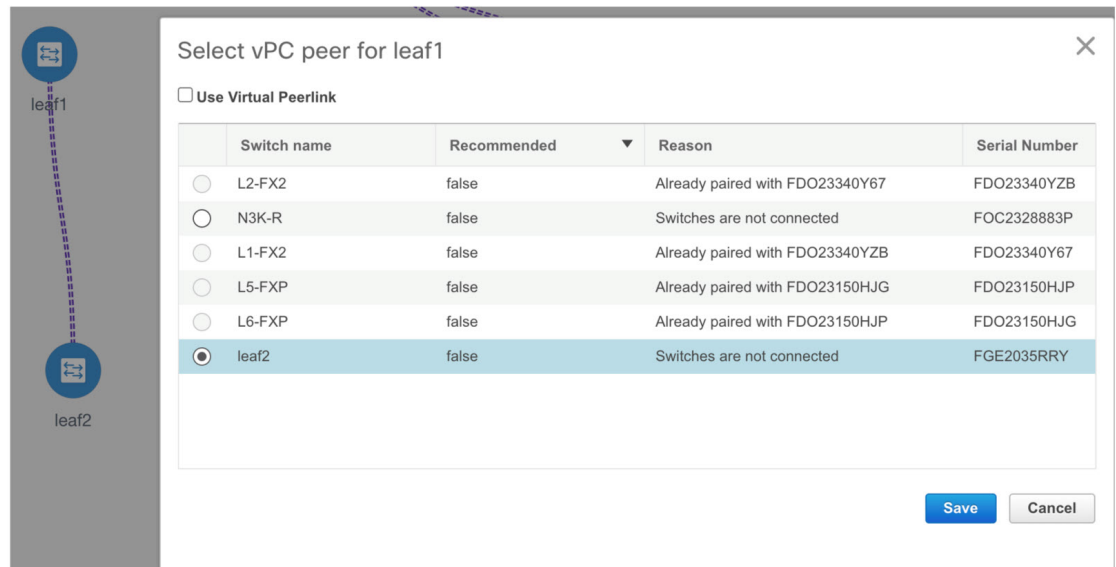
← Fabric Builder: SITE-SFO ↘

Switches **Links** Operational View

+ ✂ ✕ 🔄 📄

	<input type="checkbox"/> Fabric Name	Name	Policy	Info	Admin State	Oper State	MACsec Status
1	<input type="checkbox"/> SITE-SFO	leaf1-Ethernet1/45--leaf2-Ethernet1/45	int_pre_provision_intra_fabric_link	Neighbor Missing	--	--	NA
2	<input type="checkbox"/> SITE-SFO	leaf1-Ethernet1/44--leaf2-Ethernet1/44	int_pre_provision_intra_fabric_link	Neighbor Missing	--	--	NA

- c) On Fabric topology, right click on a switch and choose vPC Pairing from the drop-down list. Select the vPC pair and click vPC pairing for the pre-provisioned devices.
- d) Click **Save & Deploy** to generate the required intended vPC pairing configuration for the pre-provisioned devices.



Select vPC peer for leaf1

Use Virtual Peerlink

	Switch name	Recommended	Reason	Serial Number
<input type="radio"/>	L2-FX2	false	Already paired with FDO23340Y67	FDO23340YZB
<input type="radio"/>	N3K-R	false	Switches are not connected	FOC2328883P
<input type="radio"/>	L1-FX2	false	Already paired with FDO23340YZB	FDO23340Y67
<input type="radio"/>	L5-FXP	false	Already paired with FDO23150HJG	FDO23150HJP
<input type="radio"/>	L6-FXP	false	Already paired with FDO23150HJP	FDO23150HJG
<input checked="" type="radio"/>	leaf2	false	Switches are not connected	FGE2035RRY

Save Cancel

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

Intent Config

```
#POLICY-72250#
vpc domain 3
  delay restore 150

#POLICY-72270#
vpc domain 3
  peer-keepalive destination 10.1.1.10 source 10.1.1.11

#POLICY-72230#
vpc domain 3
  ipv6 nd synchronize

#POLICY-72240#
vpc domain 3
  auto-recovery reload-delay 360

#POLICY-72290#
interface port-channel500
  switchport
  switchport mode trunk
  vpc peer-link
  spanning-tree port type network

interface Ethernet1/45
  switchport
  switchport mode trunk
  channel-group 500 force mode active
```

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

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

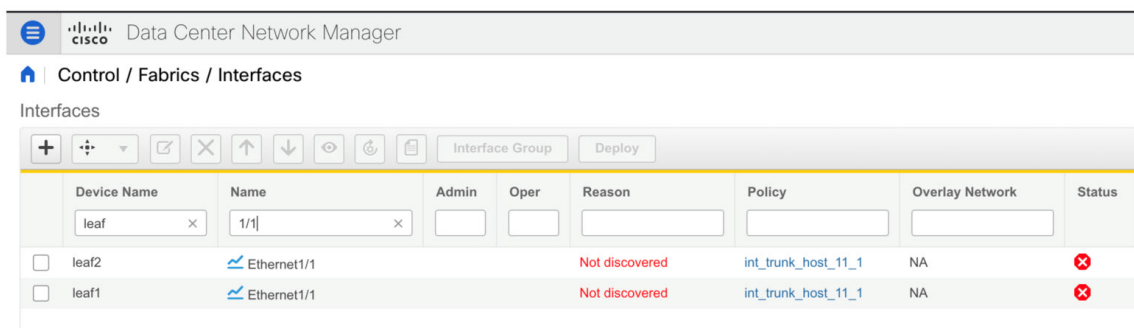
Pre-provisioning a vPC Host Interface

Procedure

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

For instructions, see [Pre-provisioning an Ethernet Interface, on page 41](#).

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



The screenshot shows the Cisco Data Center Network Manager interface. The breadcrumb navigation is "Control / Fabrics / Interfaces". The page title is "Interfaces". There are several icons for actions like add, delete, refresh, and search. Below the icons is a table with the following columns: Device Name, Name, Admin, Oper, Reason, Policy, Overlay Network, and Status. The table contains two rows of data:

Device Name	Name	Admin	Oper	Reason	Policy	Overlay Network	Status
leaf2	Ethernet1/1			Not discovered	int_trunk_host_11_1	NA	✘
leaf1	Ethernet1/1			Not discovered	int_trunk_host_11_1	NA	✘

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

Add Interface
✕

* Type:

* Select a vPC pair:

* vPC ID:

* Policy:

General

* Peer-1 Port-Channel ID: Peer-1 VPC port-channel number (Min:1, Max:4096)

* Peer-2 Port-Channel ID: Peer-2 VPC port-channel number (Min:1, Max:4096)

Enable Config Mirroring If enabled, Peer-1 config will be copied to Peer-2

Peer-1 Member Interfaces: A list of member interfaces for Peer-1 [e.g. e1/5,eth1/7-9]

Peer-2 Member Interfaces: A list of member interfaces for Peer-2 [e.g. e1/5,eth1/7-9]

* Port Channel Mode: Channel mode options: on, active and passive

* Enable BPDU Guard: Enable spanning-tree bpduguard: true='enable', false='disable', no='return to default settings'

Enable Port Type Fast Enable spanning-tree edge port behavior

* MTU: MTU for the Port Channel

SPEED: Port Channel Speed

* Peer-1 Trunk Allowed...: Allowed values: 'none', 'all', or vlan ranges (ex: 1-200,500-2000,3000)

* Peer-2 Trunk Allowed...: Allowed values: 'none', 'all', or vlan ranges (ex: 1-200,500-2000,3000)

Peer-1 PO Description: Add description to Peer-1 VPC port-channel (Max Size 254)

Peer-2 PO Description: Add description to Peer-2 VPC port-channel (Max Size 254)

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

The screenshot shows the Cisco Data Center Network Manager interface. The main window displays the 'Interfaces' page with a table of interfaces. The table has columns for Device Name, Name, Admin, Oper, and Reason. The first row shows 'leaf2-leaf1' with 'vPC1' as the name and 'Not discovered' as the reason. A search bar at the top of the table shows 'leaf' and 'vPC'. To the right, an 'Expected Config' window is open, showing configuration for two leaf nodes: leaf2:FGE2035RRY and leaf1:FGE2035RRX. The configuration for leaf2 includes port-channel1, Ethernet1/1, vpc 1, and spanning-tree settings. The configuration for leaf1 includes port-channel1, Ethernet1/1, vpc 1, and spanning-tree settings.

Device Name	Name	Admin	Oper	Reason
leaf2-leaf1	vPC1			Not discovered

```

leaf2:FGE2035RRY
interface port-channel1
switchport
switchport mode trunk
switchport trunk allowed vlan none
switchport
vpc 1
spanning-tree port type edge trunk
spanning-tree bpduguard enable
mtu 9216
description test-preprov
no shutdown

interface Ethernet1/1
switchport
switchport mode trunk
switchport trunk allowed vlan none
channel-group 1 force mode active
mtu 9216
no shutdown

leaf1:FGE2035RRX
interface port-channel1
switchport
switchport mode trunk
switchport trunk allowed vlan none
description test-preprov
switchport
vpc 1
spanning-tree bpduguard enable
mtu 9216
spanning-tree port type edge trunk
no shutdown

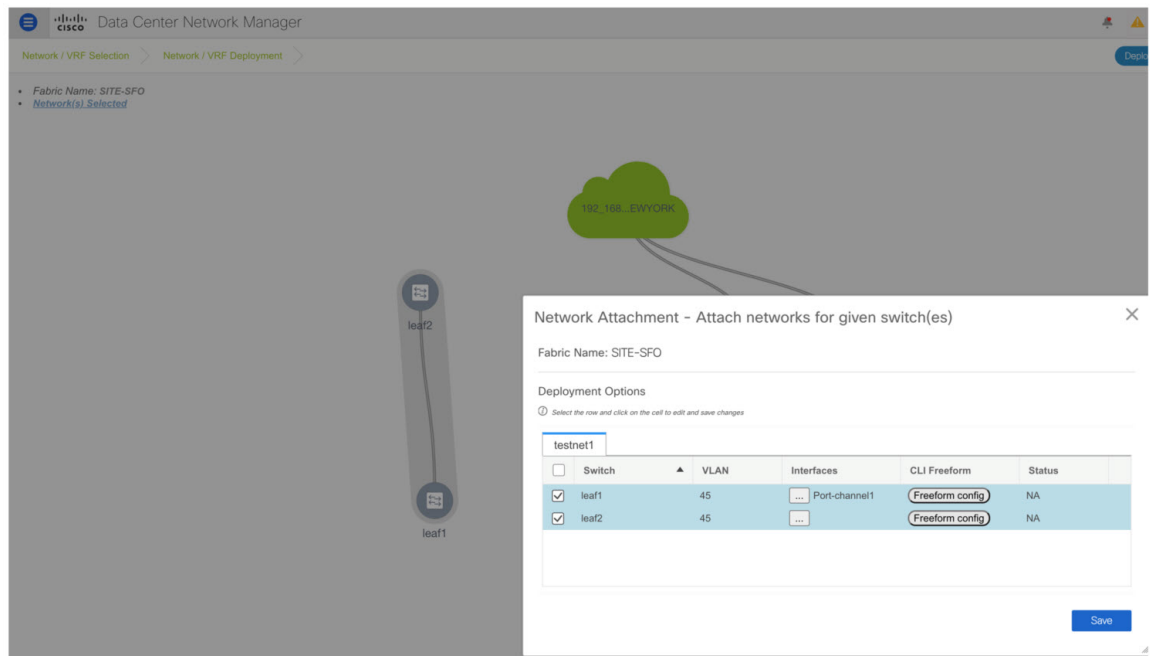
interface Ethernet1/1
switchport
switchport mode trunk
switchport trunk allowed vlan none
channel-group 1 force mode active
mtu 9216
no shutdown

```

Attaching Overlays to Pre-provisioned Devices

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

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



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

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

View/Edit Policies for leaf1(FGE2035RRX)

<input type="checkbox"/>	Policy ID	Template	Description	Generated Config ⓘ
<input type="checkbox"/>	POLICY-72630	copp_policy		profile <input type="text"/> <input type="button" value="x"/>
<input checked="" type="checkbox"/>	PROFILE-VRF...	Default_VRF_Universal		View
<input checked="" type="checkbox"/>	PROFILE-NET...	Default_Network_Uni...		View

Intent Config ✕

```
#PROFILE-VRF-22#
configure profile abc
vlan 2000
  vn-segment 153182
  interface Vlan2000
    vrf member abc
    ip forward
    ipv6 address use-link-local-only
    no ip redirects
    no ipv6 redirects
    mtu 9216
    no shutdown
  vrf context abc
  vni 153182
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
  address-family ipv6 unicast
    route-target both auto
    route-target both auto evpn
router bgp 65400
  vrf abc
    address-family ipv4 unicast
      advertise l2vpn evpn
      redistribute direct route-map FABRIC-RMAP-REDIST-SUBNET
      maximum-paths ibgp 2
    address-family ipv6 unicast
      advertise l2vpn evpn
      redistribute direct route-map FABRIC-RMAP-REDIST-SUBNET
      maximum-paths ibgp 2
  interface nvel
```

Pending In S

Precision Time Protocol for Easy Fabric

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

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

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

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

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

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

Here is the sample PTP configuration:

```
feature ptp

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

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

interface Ethernet1/59 -> Core facing interface
  ptp

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

The following guidelines are applicable for PTP:

- The PTP feature can be enabled in a fabric when all the switches in the fabric have Cisco NX-OS Release 7.0(3)I7(1) or a higher version. Otherwise, the following error message is displayed:
PTP feature can be enabled in the fabric, when all the switches have NX-OS Release 7.0(3)I7(1) or higher version. Please upgrade switches to NX-OS Release 7.0(3)I7(1) or higher version to enable PTP in this fabric.
- For hardware telemetry support in NIR, the PTP configuration is a prerequisite.
- If you are adding a non-cloud scale device to an existing fabric which contains PTP configuration, the following warning is displayed:
TTAG is enabled fabric wide, when all devices are cloud scale switches so it cannot be enabled for newly added non cloud scale device(s).
- If a fabric contains both cloud scale and non-cloud scale devices, the following warning is displayed when you try to enable PTP:
TTAG is enabled fabric wide, when all devices are cloud scale switches and is not enabled due to non cloud scale device(s).

Support for Super Spine Role in DCNM

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

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

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

The following Super Spine roles are supported in DCNM:

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

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

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


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

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

Supported Topologies for Super Spine Switches

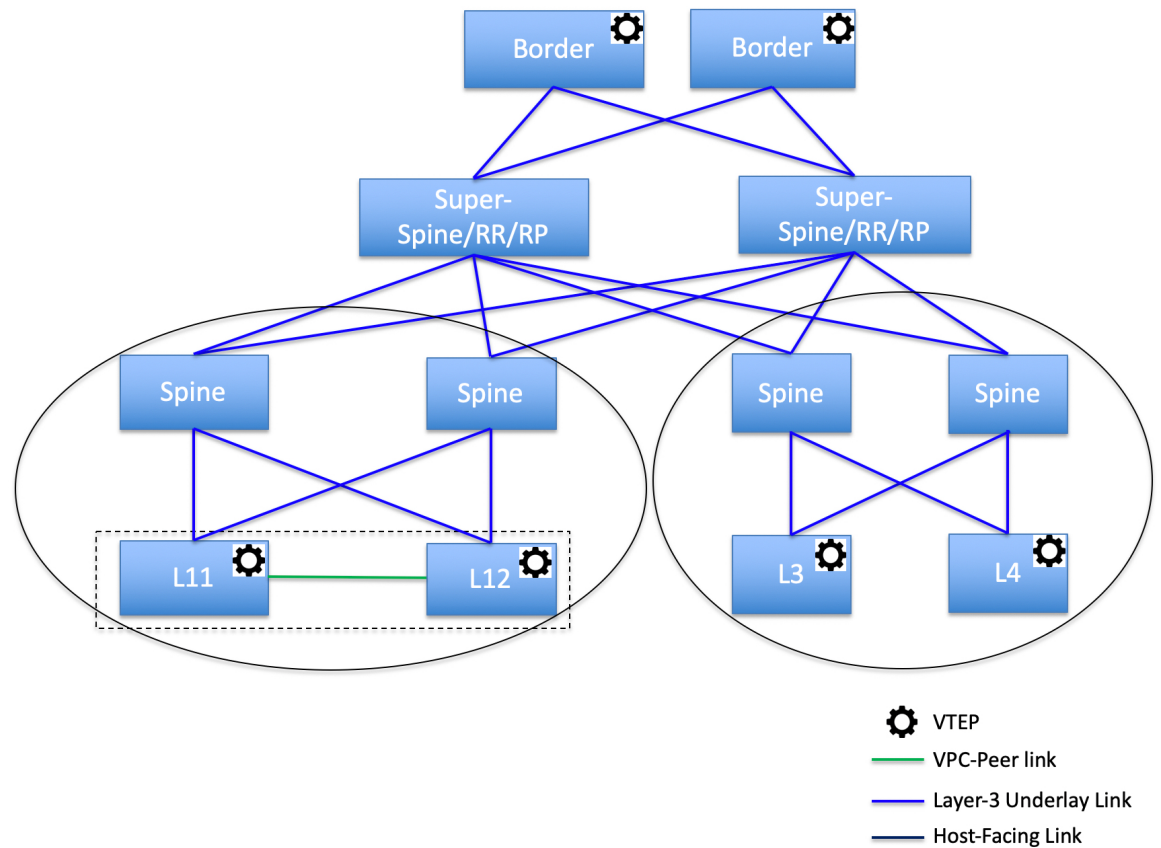
DCNM supports the following topologies with super spine switches.

Topology 1: Super Spine Switches in a Spine Leaf Topology



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

Topology 2: Super Spine Switches Connected to Border

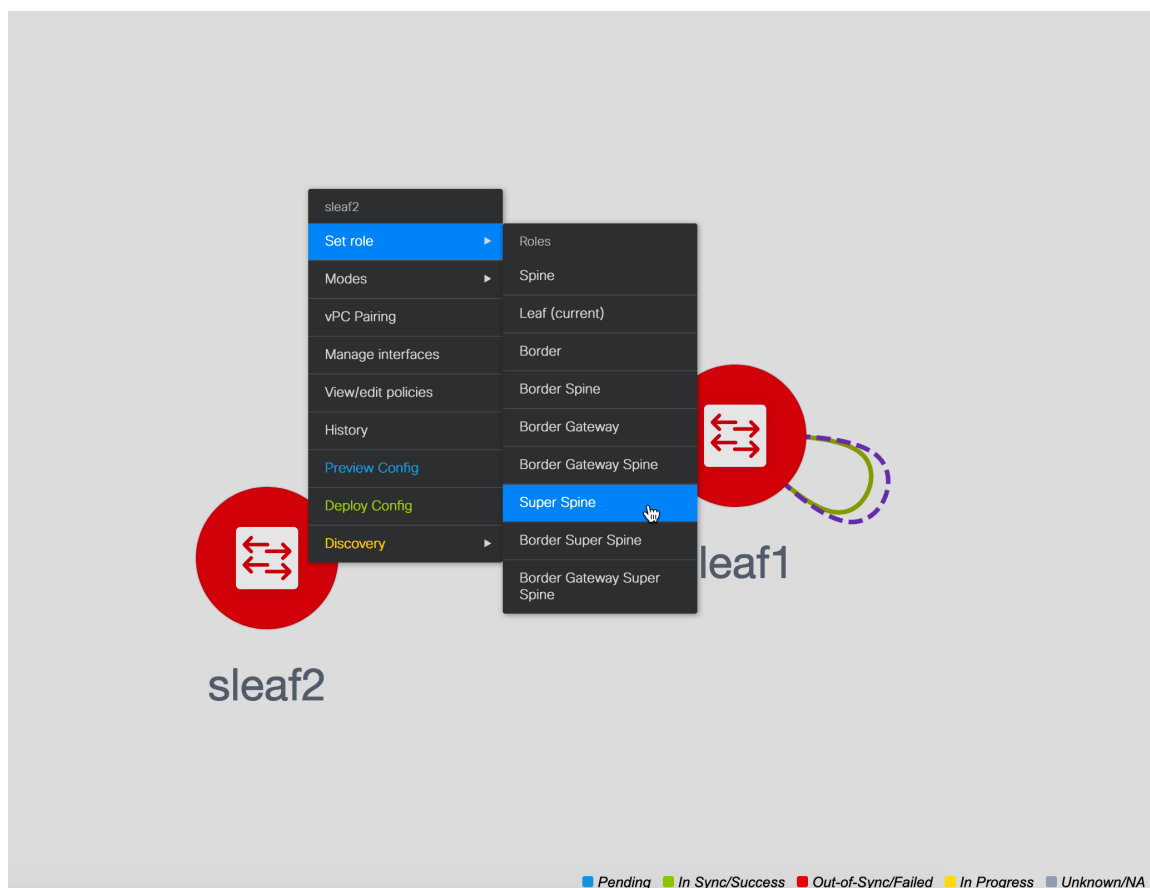


In this topology, there are four leaf switches connecting to the Spine switches, which are connected to the two Super Spine switches. These Super Spine switches are connected to the border or border gateway leaf switches.

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

Procedure

-
- Step 1** Navigate to **Control > Fabric Builder**.
- Step 2** From the **Fabric Builder** window, click **Add Switches** in the actions panel.
For more information, see [Adding Switches to a Fabric, on page 24](#).
- Step 3** Right-click an existing switch or the newly added switch, and use the **Set role** option to set the appropriate super spine role.

**Note**

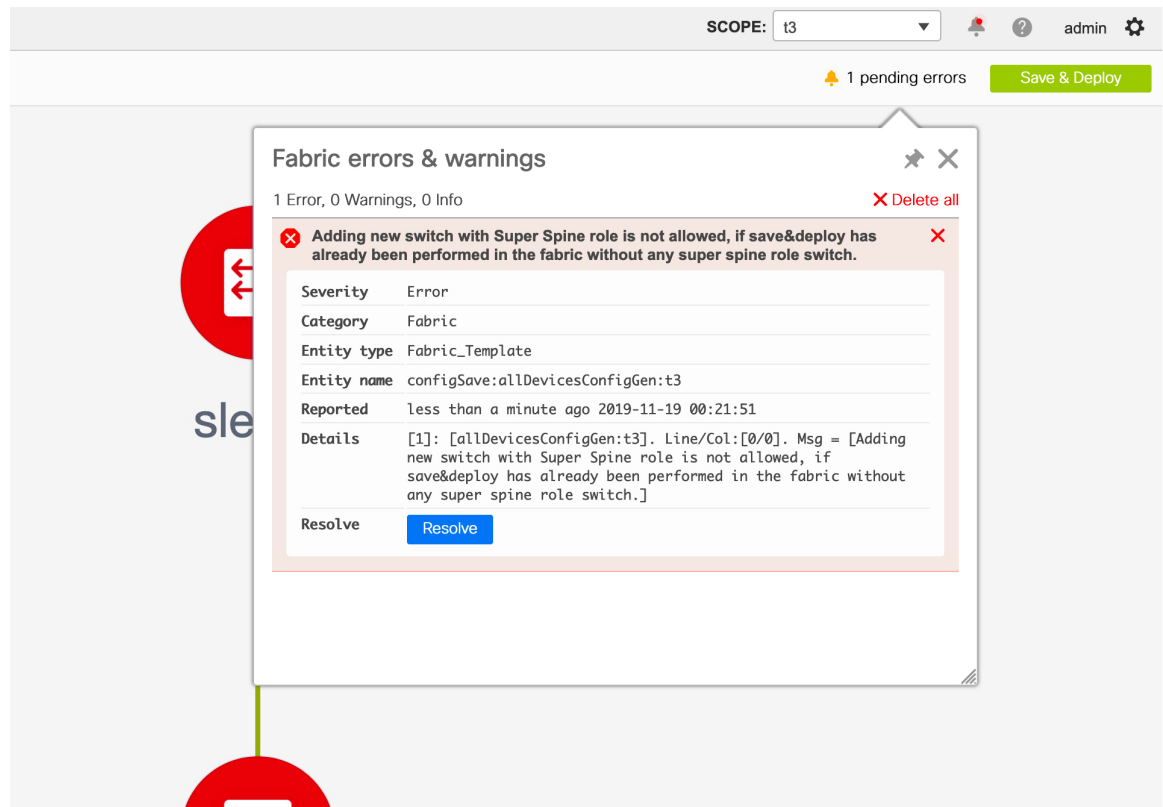
- If the **Super Spine** role is present in the fabric, then the other super spine roles that you can assign for any new device are border super spine and border gateway super spine.
- If Super Spine or any of its variation role is not present in the fabric, you may assign the role to any new device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

Step 4 Click **Save & Deploy**.

An error is displayed saying:

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

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



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

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

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

Fabric errors & warnings

1 Error, 3 Warnings, 0 Info Delete all

- ✘ Only Super Spine, Border Super Spine or Border Gateway Super Spine roles are allowed when any Super Spine role is present in the Fabric ✘

Severity	Error
Category	Fabric
Entity type	Fabric_Template
Entity name	configsave:validateFabricSetting:Non Super Spine Role Border
Reported	less than a minute ago 2020-01-07 10:12:26
Details	[1]: [validateFabricSetting:Non Super Spine Role Border]. Line/Col:[0/0]. Msg = [only super spine, Border super spine or Border Gateway super spine roles are allowed when any super spine role is present in the Fabric]

- ⚠ DCI subnet range duplicate with fabric: fab-2 ✘
- ⚠ Loopback 1 range duplicate with fabric: fab-2 ✘
- ⚠ Loopback 0 range duplicate with fabric: fab-2 ✘

Changing the TCAM Configuration on a Device

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

- Cisco Nexus 9300 Series Switches: **tcam_pre_config_9300** and **tcam_pre_config_vxlan**
- Cisco Nexus 9500 Series Switches: **tcam_pre_config_9500** and **tcam_pre_config_vxlan**

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

1. Choose **Control > Fabrics > Fabric Builder**.
2. Click the fabric containing the specified switches that have been onboarded using the bootstrap feature.
3. Click **Tabular View** under the **Actions** menu in the **Fabric Builder** window.
4. Select all the specified switches and click the **View/Edit Policies** icon.
5. Search for **tcam_pre_config** policies.
6. If the TCAM config is incorrect or not applicable, select all these policies and click the Delete icon to delete policies.

7. Add one or multiple `tcam_config` policies and provide the correct TCAM configuration. For more information about how to add a policy, see *Adding PTIs for Multiple Switches*.
8. Reload the respective switches.

If the switch is used as a leaf, border leaf, border gateway leaf, border spine, or border gateway spine, add the `tcam_config` policy with the following command and deploy.

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

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

Make sure that the priority of this `tcam_config` policy is higher than the `tcam_pre_config_vxlan` policy so that the config policy with `racl 1024` is configured before the `tcam_pre_config_vxlan` policy.



Note The `tcam_pre_config_vxlan` policy contains the config: `hardware access-list tcam region arp-ether 256 double-wide`.

Preselecting Switches as Route-Reflectors and Rendezvous-Points

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



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

Procedure

-
- Step 1** Import switches successfully.
 - Step 2** Create the `rr_state` or `rp_state` policies using **View/Edit Policies** on the spines or super spine switches, which should be preselected as RR or RP.
 - Note**
 - If there are more than 2 spines and the maximum number of RRs or RPs in the fabric settings is set to 2, then it's recommended to distribute RR and RP on different spines.
 - If there are more than 4 spines and the maximum number of RRs or RPs in the fabric settings is set to 4, then it's recommended to distribute RR and RP on different spines.
 - Step 3** Click **Save & Deploy**, and then click **Deploy Config**.
The spines that have `rr_state` policies become RR and spines that have `rp_state` policies become RP.
 - Step 4** After **Save & Deploy**, if you want to replace the preselected RRs and RPs with a new set of devices, then old RR and RP devices should be removed from the fabric before performing the same steps.
-

Adding a vPC L3 Peer Keep-Alive Link

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



Note

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


Procedure

Step 1

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


Step 2

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



Step 3

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




Step 4

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

Step 5

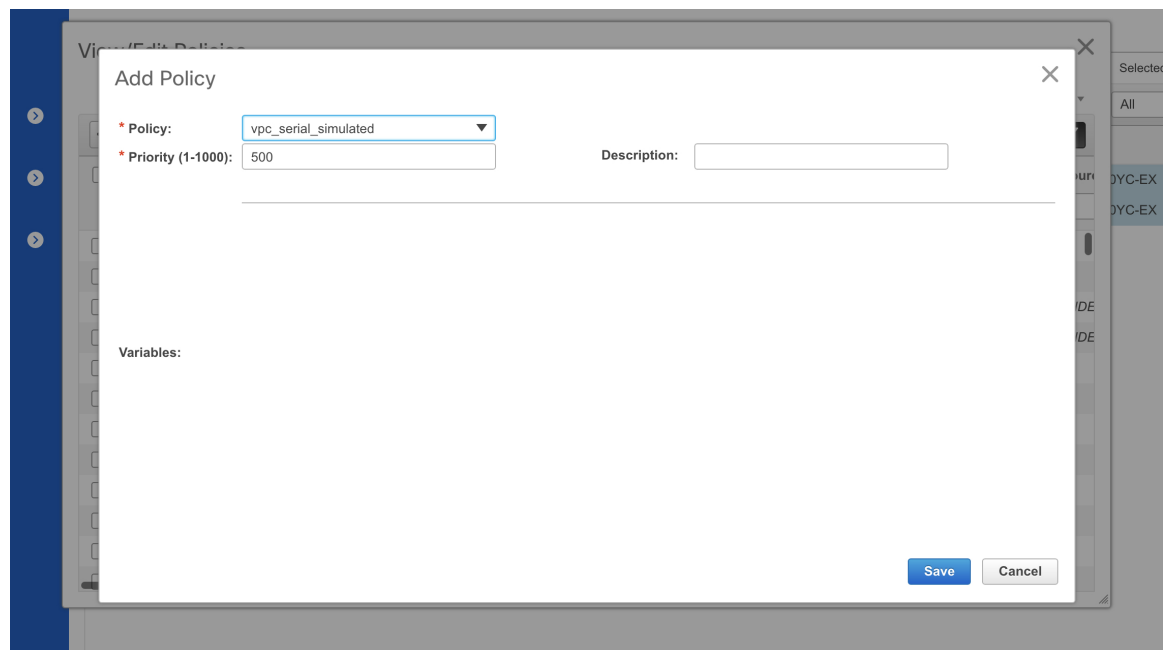
Click **Tabular View** and select the vPC pair switches, and then click **View/Edit Policies**.

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



Step 6 Click + to add policies.

Step 7 From the **Policy** drop-down list, select **vpc_serial_simulated** policy and add priority. Click **Save**. Note that if both switches are selected, then this policy will be created on both vPC pair switches.



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

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

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

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

Link Management - Edit Link

- * Link Type: Intra-Fabric
- * Link Sub-Type: Fabric
- * Link Template: int_intra_vpc_peer_keep_alive
- * Source Fabric: t2
- * Destination Fabric: t2
- * Source Device: leaf3
- * Source Interface: Ethernet1/19
- * Destination Device: leaf2
- * Destination Interface: Ethernet1/19

Link Profile

General

Advanced

Interface VRF: ⓘ Name of a non-default VRF for this interface (make sure to co

* Source IP: 1.1.1.1 ⓘ IP address of the source interface

* Destination IP: 1.1.1.2 ⓘ IP address of the destination interface

Source V6IP: ⓘ IPv6 address of the source interface

Destination V6IP: ⓘ IPv6 address of the destination interface

Interface Admin State: ⓘ Admin state of the interface

* MTU: 9216 ⓘ MTU for the interface

Save

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

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

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

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

The screenshot displays the Cisco Data Center Network Manager (DCNM) Fabric Builder interface. The main view shows a fabric topology with two leaf switches (leaf2 and leaf3) connected to a spine switch. A detailed view of leaf2 is open on the right, showing its configuration and health status. The Peerlink State is highlighted in orange, indicating it is OK. The health status is 95%.

Summary			
Status:			ok
Serial number:		FDO20352BEE	
Version:		9.3(1)	
CPU:		14	
Memory:		21	
VPC Domain ID: 1			
Role:		Secondary	
Peer:		leaf3	
Peerlink State:		Peer is OK	
Keep Alive State:		Peer is alive	
Consistency State:		Consistent	
Send Interface:		Eth1/19	
Receive Interface:		Eth1/19	
Health			
95%	Modules	91.67%	w=0.2
	Switch ports	83.61%	w=0.2
	Alarms	100.00%	w=0.6
Tags			

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

Procedure

- Step 1** Log in to DCNM and navigate to **Control > Fabric Builder**.
- Step 2** Click the **Edit** icon for a fabric and add the AAA authentication commands in the **AAA Freeform Config** field under the **Manageability** tab.

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

Edit Fabric



* Fabric Name :

* Fabric Template :

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

General	Replication	vPC	Protocols	Advanced	Resources	Manageability	Bootstrap	Configuration Backup
<p><small>Comma separated list of vrf's, one per NTP server</small></p> <p>Syslog Server IPs <input type="text"/> ⓘ <small>Comma separated list of IP Addresses(v4/v6)</small></p> <p>Syslog Server Severity <input type="text"/> ⓘ <small>Comma separated list of Syslog severity values, one per Syslog server (Min:0, Max:7)</small></p> <p>Syslog Server VRFs <input type="text"/> ⓘ <small>One VRF for all Syslog servers or a comma separated list of VRFs, one per Syslog server</small></p> <p>AAA Freeform Config</p> <pre> aaa group server tacacs+ AAA_TACACS server 172.25.35.39 use-vrf management source-interface mgmt0 aaa authentication login default group AAA_TACACS local aaa authentication login console local aaa accounting default group AAA_TACACS aaa authentication login error-enable aaa authorization config-commands default group AAA_TACACS local aaa authorization commands default group AAA_TACACS local </pre> <p><small>Note ! All configs should strictly match 'show run' out, with respect to case and new. Any mismatches will yield unexpected diffs during depl</small></p>								
							Save	Cancel

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

Step 4 If you are importing switches in to the fabric via POAP, you need to have the AAA configs on the switch. Navigate to the fabric settings and add the relevant commands in **Bootstrap Freeform Config**.

Edit Fabric




* Fabric Name :

* Fabric Template :

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

General	Replication	vPC	Protocols	Advanced	Resources	Manageability	Bootstrap	Configuration Backup
<p>Enable Local DHCP Server <input type="checkbox"/> ⓘ <small>Automatic IP Assignment For POAP From Local DHCP Server</small></p> <p>DHCP Version <input type="text"/> ⓘ</p> <p>DHCP Scope Start Address <input type="text"/> ⓘ <small>Start Address For Switch Out-of-Band POAP</small></p> <p>DHCP Scope End Address <input type="text"/> ⓘ <small>End Address For Switch Out-of-Band POAP</small></p> <p>Switch Mgmt Default Gateway <input type="text"/> ⓘ <small>Default Gateway For Management VRF On The Switch</small></p> <p>Switch Mgmt IP Subnet Prefix <input type="text"/> ⓘ <small>(Min:8, Max:30)</small></p> <p>Switch Mgmt IPv6 Subnet Prefix <input type="text"/> ⓘ <small>(Min:64, Max:126)</small></p> <p>Enable AAA Config <input checked="" type="checkbox"/> ⓘ <small>Include AAA configs from Manageability tab during device bootup</small></p> <p>Bootstrap Freeform Config</p> <pre> aaa group server tacacs+ AAA_TACACS server 172.25.35.39 use-vrf management source-interface mgmt0 aaa authentication login default group AAA_TACACS local aaa authentication login console local aaa accounting default group AAA_TACACS </pre> <p><small>Note ! All configs should strictly match 'show run' out, with respect to case and new. Any mismatches will yield</small></p>								
							Save	Cancel

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



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

IPv6 Underlay Support for Easy Fabric

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

Brownfield Deployment-Transitioning VXLAN Fabric Management to DCNM

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

Configuring Fabrics with eBGP Underlay

You can use the **Easy_Fabric_eBGP** fabric template to create a fabric with eBGP underlay. For more information, see [Managing BGP-Based Routed Fabrics](#) and [Managing a Greenfield VXLAN BGP EVPN Fabric](#).

Creating an External Fabric

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

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

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

- You can set up a Multi-Site or a VRF-lite IFC by manually configuring the links for the border devices in the VXLAN fabric or by using an automatic Deploy Border Gateway Method or VRF Lite IFC Deploy Method. If you are configuring the links manually for the border devices, we recommend using the Core Router role to set up a Multi-Site eBGP underlay from a Border Gateway device to a Core Router and the Edge Router role to set up a VRF-lite Inter-Fabric Connection (IFC) from a Border device to an Edge device.
- If you are using the Cisco Nexus 7000 Series Switch with Cisco NX-OS Release 6.2(24a) on the LAN Classic or External fabrics, make sure to enable AAA IP Authorization in the fabric settings.
- You can discover the following non-Nexus devices in an external fabric:
 - IOS-XE family devices: Cisco CSR 1000v, Cisco IOS XE Gibraltar 16.10.x, Cisco ASR 1000 Series routers, and Cisco Catalyst 9000 Series Switches
 - IOS-XR family devices: ASR 9000 Series Routers, IOS XR Release 6.5.2 and Cisco NCS 5500 Series Routers, IOS XR Release 6.5.3
 - Arista 4.2 (Any model)
- Configure all the non-Nexus devices, except Cisco CSR 1000v, before adding them to the external fabric.
- From Cisco DCNM Release 11.4(1), you can configure non-Nexus devices as borders. You can create an IFC between a non-Nexus device in an external fabric and a Cisco Nexus device in an easy fabric. The interfaces supported for these devices are:
 - Routed
 - Subinterface
 - Loopback
- From Cisco DCNM, Release 11.4(1), you can configure a Cisco ASR 1000 Series routers and Cisco Catalyst 9000 Series switches as edge routers, set up a VRF-lite IFC and connect it as a border device with an easy fabric.
- Before a VDC reload, discover Admin VDC in the fabric. Otherwise, the reload operation does not occur.
- You can connect a Cisco data center to a public cloud using Cisco CSR 1000v. See the *Connecting Cisco Data Center and a Public Cloud* chapter for a use case.

- In an external fabric, when you add the **switch_user** policy and provide the username and password, the password must be an encrypted string that is displayed in the **show run** command.

For example:

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

In this case, the entered password should be

\$5\$I4sapkBh\$S7B7UcPH/iVTihLKH5sgldBeS3O2X1StQsvv3cmbYd1.

- For the Cisco Network Insights for Resources (NIR) Release 2.1 and later, and flow telemetry, **feature lldp** command is one of the required configuration.

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

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

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

Creating External Fabric from Fabric Builder

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

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

Fabric Name - Enter the name of the external fabric.

Fabric Template - Choose *External_Fabric*.

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

3. Fill up the **General** tab as shown below.

Add Fabric ✕

* Fabric Name :

* Fabric Template :

General | Advanced | Resources | Configuration Backup | Bootstrap

* BGP AS # 1-4294967295 | 1-65535[0-65535]

Fabric Monitor Mode ? If enabled, fabric is only monitored. No configuration will be deployed

BGP AS # - Enter the BGP AS number.

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

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

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

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

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

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

The screenshot shows the 'Advanced' configuration tab with the following settings:

- vPC Peer Link VLAN:** 3600 (with an information icon and tooltip: 'VLAN for vPC Peer...')
- Power Supply Mode:** ps-redundant (with a dropdown arrow and an information icon and tooltip: 'Default Power Supply...')
- Enable MPLS Handoff:** (with an information icon and tooltip: 'i')
- Underlay MPLS Loopback Id:** (empty field with an information icon and tooltip: '(Min:0, Max:1023)')
- Enable AAA IP Authorization:** (with an information icon and tooltip: 'Enable only, when IP Authorization is enabled in the AAA Ser...')
- Enable DCNM as Trap Host:** (with an information icon and tooltip: 'Configure DCNM as a receiver for SNMP traps')
- Enable CDP for Bootstrapped Switch:** (with an information icon and tooltip: 'Enable CDP on management interface')
- Enable NX-API:** (with an information icon and tooltip: 'Enable NX-API on port 443')
- Enable NX-API on HTTP port:** (with an information icon and tooltip: 'Enable NX-API on port 80')
- Inband Mgmt:** (with an information icon and tooltip: 'Import switches with inband connectivity')
- Enable Precision Time Protocol (PTP):** (with an information icon and tooltip: 'i')
- PTP Source Loopback Id:** (empty field with an information icon and tooltip: '(Min:0, Max:1023)')
- PTP Domain Id:** (empty field with an information icon and tooltip: 'Multiple Independent... on a Single Network (I...)
- Fabric Freeform:** (empty text area)
- AAA Freeform Config:** (empty text area)

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

Power Supply Mode - Choose the appropriate power supply mode.

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

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

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

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

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

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

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



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

Inband Mgmt: For External and Classic LAN Fabrics, this knob enables DCNM to import and manage of switches with inband connectivity (reachable over switch loopback, routed, or SVI interfaces) , in addition to management of switches with out-of-band connectivity (aka reachable over switch mgmt0 interface). The only requirement is that for Inband managed switches, there should be IP reachability from DCNM to the switches via the eth2 aka inband interface. For this purpose, static routes may be needed on the DCNM, that in turn can be configured via the Administration->Customization->Network Preferences option. After enabling Inband management, during discovery, provide the IPs of all the switches to be imported using Inband Management and set maximum hops to 0. DCNM has a pre-check that validates that the Inband managed switch IPs are reachable over the eth2 interface. Once the pre-check has passed, DCNM then discovers and learns about the interface on that switch that has the specified discovery IP in addition to the VRF that the interface belongs to. As part of the process of switch import/discovery, this information is captured in the baseline intent that is populated on the DCNM. For more information, see [Inband Management in External Fabrics and LAN Classic Fabrics, on page 139](#).



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

Enable Precision Time Protocol (PTP): Enables PTP across a fabric. When you select this check box, PTP is enabled globally and on core-facing interfaces. Additionally, the PTP Source Loopback Id and PTP Domain Id fields are editable. For more information, see [Precision Time Protocol for External Fabrics and LAN Classic Fabrics, on page 140](#).

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

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

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

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

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

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

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

The screenshot shows the 'Resources' tab with the following fields:

- Subinterface Dot1q Range:** Input field with value '2-511'. Help text: 'Per Border Dot1q Range For VRF Lite Connectivity (Min:2, Max:4093)'.
- Underlay Routing Loopback IP Range:** Input field with value '10.1.0.0/22'. Help text: 'Typically Loopback0 IP Address Range'.
- Underlay MPLS Loopback IP Range:** Input field. Help text: 'MPLS Loopback IP Address Range'.

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

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

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

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

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

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

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

The screenshot shows the 'Configuration Backup' tab with the following options:

- 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:** Input field. Help text: 'Time in 24hr format. (00:00 to 23:59)'

The fields on this tab are:

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

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

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

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

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

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

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

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

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

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

Pointers for hourly and scheduled backup:

- The backups contain running configuration and intent pushed by DCNM. Configuration compliance forces the running config to be the same as the DCNM config. Note that for the external fabric, only some configurations are part of intent and the remaining configurations are not tracked by DCNM. Therefore, as part of backup, both DCNM intent and running config from switch are captured.

7. Click the **Bootstrap** tab.

Edit Fabric

* Fabric Name :

* Fabric Template :

ⓘ Fabric Template for support of Nexus and non-Nexus devices

General | **Advanced** | Resources | Configuration Backup | Bootstrap

Enable Bootstrap (For NX-OS Switches Only)
 ⓘ Automatic IP Assignment For POAP
 Enable Local DHCP Server
 ⓘ Automatic IP Assignment For POAP From Local DHCP Server

DHCP Version:

DHCP Scope Start Address: ⓘ Start Address For Switch Out-of-Band POAP

DHCP Scope End Address: ⓘ End Address For Switch Out-of-Band POAP

Switch Mgmt Default Gateway: ⓘ Default Gateway For Management VRF On The Switch

Switch Mgmt IP Subnet Prefix: ⓘ (Min:8, Max:30)

Switch Mgmt IPv6 Subnet Prefix: ⓘ (Min:64, Max:126)

Enable AAA Config ⓘ Include AAA configs from Advanced tab during device bootstrap

Bootstrap 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.

DHCPv4/DHCPv6 Multi Subnet Scope:

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** check box and enter details for the remaining mandatory fields.

Enable Local DHCP Server - Select this check box to initiate enabling of automatic IP address assignment through the local DHCP server. When you select this check box, all the remaining fields become editable.

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



Note Cisco DCNM IPv6 POAP 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.

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

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

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

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

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

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

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

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

Copy-paste the running-config to a **freeform config** field with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. For more information, see [Resolving Freeform Config Errors in Switches, on page 303](#).

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

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

The screenshot shows the configuration interface for the ThousandEyes Agent. The 'ThousandEyes Agent' tab is active. The configuration options are as follows:

- Enable Fabric Override for ThousandEyes Agent Installation:** (info icon)
- ThousandEyes Account Group Token:** (info icon) *Token from ThousandEyes Agent Settings for Agent Installation*
- VRF on Switch for ThousandEyes Agent Collector Reachability:** (info icon) *NX-OS VRF that provides Internet Reachability*
- DNS Domain:** (info icon) *DNS Domain Configuration*
- DNS Server IPs:** (info icon) *Comma separated list of IP Addresses(v4/v6)*
- NTP Server IPs:** (info icon) *Comma separated list of IP Addresses(v4/v6)*
- Enable Proxy for Internet Access:** (info icon) *Proxy Settings for NX-OS Switch Internet Access*
- Proxy Information:** (info icon) *Proxy-Server:port*
- Proxy Bypass:** (info icon) *Comma separated No-proxy server list*

Buttons for 'Save' and 'Cancel' are located at the bottom right of the configuration area.

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.

9. Click **Save**.

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

After creating the external fabric, add switches to it.

Add Switches to the External Fabric

1. Click Add switches. The Inventory Management screen comes up.
You can also add switches by clicking Tabular View > Switches > + .
2. Enter the IP address (Seed IP) of the switch.
3. Choose the device type from the **Device Type** drop-down list.

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

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

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

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

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

4. Enter the administrator username and password of the switch.
5. Click Start discovery at the bottom part of the screen. The Scan Details section comes up shortly. Since the Max Hops field was populated with 2, the switch with the specified IP address and switches two hops from it are populated.

6. Select the check boxes next to the concerned switches and click Import into fabric.
You can discover multiple switches at the same time. The switches must be properly cabled and connected to the DCNM server and the switch status must be manageable.
The switch discovery process is initiated. The Progress column displays the progress. After DCNM discovers the switch, the screen closes and the fabric screen comes up again. The switch icons are seen at the centre of the fabric screen.
7. Click Refresh topology to view the latest topology view.
8. *External Fabric Switch Settings* - The settings for external fabric switches vary from the VXLAN fabric switch settings. Right-click on the switch icon and set or update switch options.

The options are:

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



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

Modes – Active/Operational mode.

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

Manage Interfaces – Deploy configurations on the switch interfaces.

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

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

History – View per switch deployment history.

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

Deploy Config – Deploy per switch configurations.

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

9. Click Save & Deploy at the top right part of the screen. The template and interface configurations form the configuration provisioning on the switches.
When you click Save & Deploy, the Configuration Deployment screen comes up.
10. Click Deploy Config at the bottom part of the screen to initiate pending configuration onto the switch.
11. Close the screen after deployment is complete.



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

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

Move an External Fabric Under an MSD Fabric

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

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

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

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

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

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

External Fabric Depiction in an MSD Fabric Topology

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



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

External Fabric Switch Operations

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

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

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

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



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

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

External Fabric Links

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

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

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

The Switches | Links screen comes up.

2. Choose one or more check boxes and click the Delete icon at the top left.

The links are deleted.

Move Neighbor Switch to External Fabric

1. Click Add switches. The Inventory Management screen comes up.
2. Click Move Neighbor Switches tab.
3. Select the switch and click **Move Neighbor**.

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

Discovering New Switches

To discover new switches, perform the following steps:

Procedure

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

Inventory Management
✕

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

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

🔄
Bootstrap

+
↻
↺

* Admin Password
* Confirm Admin Password

🔒

<input type="checkbox"/>	Serial Number	Model	Version	IP Address	Hostname
<input type="checkbox"/>	TBM14299900	N7K-C7010	8.0(1)	<input style="width: 100%;" type="text"/>	<input style="width: 100%;" type="text"/>

Close

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

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

Beginning with Release 11.2(1), you can provision devices in advance. To pre-provision devices, refer to [Pre-provisioning a Device](#), on page 37.

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

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

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

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

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

Inventory Management ✕

Discover Existing Switches | PowerOn Auto Provisioning (POAP)

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

+ ↻ ↺ * Admin Password * Confirm Admin Password 🔒

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

Close

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

Inventory Management ✕

Discover Existing Switches | PowerOn Auto Provisioning (POAP)

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

+ ↻ ↺ * Admin Password * Confirm Admin Password 🔒

Serial Number Model

No Data available

Discovery Credentials ✕

*Discovery Username:

*Discovery Password:

*Confirm Discovery Password:

OK Clear

Close

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

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

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

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

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

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

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

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

Step 15 Monitor and check the switch for POAP completion.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Adding non-Nexus Devices to External Fabrics

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

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

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

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

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

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

Configuring non-Nexus Devices for Discovery

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

Configuring IOS-XE Devices for Discovery

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

Procedure

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

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

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

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

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

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

Configuring Arista Devices for Discovery

Enable Privilege Exec mode using the following command:

```
switch> enable
switch#
```

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

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

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



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

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

Show Run Command

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

```

!
spanning-tree mode mstp
!
service unsupported-transceiver labs f5047577
!
aaa authorization exec default local
!
no aaa root
!
username admin role network-admin secret sha512
$6$5ZKs/7.k2UxrWDg0$FokdVQsBTnOquW/9AYx36YUBSPNLFdeuPIse9XgyHSdeOYXtPyT/0sMUYYdkMffuIjgn/d9rx/Do71XSbygSn/
username cvpadmin role network-admin secret sha512
$6$fLGFj/PUCuJT436i$Sj5G5c4y9cYjI/BZswjmmZW0J4npGrGqIyG3ZFK/ULza47Kz.d31q13jXA7iHM677gwqQbFSH2/3oQEaHRq08.
username dcnm privilege 15 role network-admin secret sha512
$6$M48PNrCdq2EITEcdG$iiB880nvFQQlRwoZwQMzdt5EfkucIraNqtEMRS0TJUHnKCQnJN.VDLFsLAmP7kQBo.C3ct4/.n.2eRlcP6hij/

```

Show SNMP View Command

```

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

```

```

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

```

Configuring Cisco IOS-XR Devices for Discovery

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

```

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

```



Note SNMP password should be same as password for username.

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

Configuration and Verification of Cisco IOS-XR Devices

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

Discovering non-Nexus Devices in an External Fabric

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

Before you begin

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

Procedure

Step 1 Click **Add switches** in the **Actions** pane.

The **Inventory Management** dialog box appears.

Step 2 Enter values for the following fields under the **Discover Existing Switches** tab:

Field	Description
Seed IP	<p>Enter the IP address of the switch.</p> <p>You can import more than one switch by providing the IP address range. For example: 10.10.10.40-60</p> <p>The switches must be properly cabled and connected to the DCNM server and the switch status must be manageable.</p>
Device Type	<ul style="list-style-type: none"> Choose IOS XE from the drop-down list for adding Cisco CSR 1000v, Cisco ASR 1000 Series routers, or Cisco Catalyst 9000 Series Switches. Choose IOS XR from the drop-down list for adding Cisco NCS 5500 Series Routers, IOS XR Release 6.5.3. Choose Other from the drop-down list for adding non-Cisco devices, like Arista switches.

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

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

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

Step 3 Click **Start Discovery**.

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

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

Step 5 Click **Import into fabric**.

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

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

Step 6 Click **Close**.

The fabric topology window appears with the switches.

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

Step 8 (Optional) Click **Tabular view** in the **Actions** pane.

The switches and links window appears, where you can view the scan details. The discovery status is **discovering** in red with a warning icon next to it if the discovery is in progress.

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

After the discovery of the device:

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

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

What to do next

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

Pre-provisioning a Device

From Cisco DCNM Release 11.2, you can provision devices in advance.



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

- The pre-provisioned devices support the following configurations in DCNM:
 - Base management
 - vPC Pairing
 - Intra-Fabric links
 - Ethernet ports
 - Port-channel
 - vPC
 - ST FEX
 - AA FEX
 - Loopback
 - Overlay network configurations
- The pre-provisioned devices do not support the following configurations in DCNM:
 - Inter-Fabric links
 - Sub-interface
 - Interface breakout configuration
- When a device is being pre-provisioned has breakout links, you need to specify the corresponding breakout command along with the switch's model and gateway in the **Data** field in the **Add a new device to pre-provisioning** window in order to generate the breakout PTL.

Note the following guidelines:

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

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

- {"modulesModel": ["N9K-C93180LC-EX"], "gateway": "10.1.1.1/24"}

- {"modulesModel": ["N9K-C93180LC-EX"],"breakout": "interface breakout module 1 port 1 map 10g-4x", "portMode": "hardware profile portmode 4x100G+28x40G", "gateway": "172.22.31.1/24" }
- {"modulesModel": ["N9K-X9736C-EX", "N9K-X9732C-FX", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-C9516-FM-E2", "N9K-SUP-B+", "N9K-SC-A", "N9K-SC-A"], "gateway": "172.22.31.1/24" }
- {"breakout": "interface breakout module 1 port 50 map 10g-4x", "gateway": "172.16.1.1/24", "modulesModel": ["N9K-C93180YC-EX "]}
- {"modulesModel": ["N9K-X9732C-EX", "N9K-X9732C-EX", "N9K-C9504-FM-E", "N9K-C9504-FM-E", "N9K-SUP-B", "N9K-SC-A", "N9K-SC-A"], "gateway": "172.29.171.1/24", "breakout": "interface breakout module 1 port 1,11,19 map 10g-4x; interface breakout module 1 port 7 map 25g-4x" }
- {"modulesModel": ["N9K-C93180LC-EX"], "gateway": "10.1.1.1/24", "breakout": "interface breakout module 1 port 1-4 map 10g-4x", "portMode": "hardware profile portmode 48x25G + 2x100G + 4x40G" }

Procedure

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

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

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

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

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

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

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

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

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

	A	B	C	D	E	F	G
1	#SerialNumber(Eg:FDO1344GH5)	#Model(Eg:N9K-C9236C)	#Version(Eg:7.0(3)12(3))	#IPAddress of the device	#HostName	#Data(JSON Field contains model name of the modules	
2	Serial Number	Model	Version	IP Address	Hostname	Data	
3	FDO21331SND	N9K-93180YC-EX	7.0(3)5(2)	1.1.1.1	leaf1	{"modulesModel":["N9K-93180YC-EX"]}	
4	FDO21351N3X	N9K-C9236C	7.0(3)4(1)	11.1.1.1	spine1	{"modulesModel":["N9K-C9236C"]}	
5	FDO21491A5K	N9K-C93240YC-FX2	7.0(3)17(3)	12.1.1.1	leaf2	{"modulesModel":["N9K-C93240YC-FX2"]}	
6							

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

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

<input checked="" type="checkbox"/>	Serial Number	Model	Version	IP Address	Hostname
<input checked="" type="checkbox"/>	SN	N9K-3455	7.0(2)	10.1.1.1	leaf1

The leaf1 device appears in the fabric topology.

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

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

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

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

Create intra-link policy using one of the templates:

- **int_pre_provision_intra_fabric_link** to automatically generate intra fabric interface configuration with DCNM allocated IP addresses
- **int_intra_fabric_unnum_link_11_1** if you are using unnumbered links
- **int_intra_fabric_num_link_11_1** if you want to manually assign IP addresses to intra-links

Click **Save & Deploy**.

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

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

For more information, see [Return Material Authorization \(RMA\), on page 210](#).

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

You need to click **Save & Deploy** in the fabric after one or more switches are online to provision the host ports. This action must be performed before overlays are provisioned for the host port attachment.

Pre-provisioning an Ethernet Interface

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



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

Before you begin

Make sure that you have a preprovisioned device in your fabric. For information, see [Pre-provisioning a Device](#), on page 37.

Procedure

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

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

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

Step 3 Click **Add**.

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

Add Interface
✕

* Type:

* Select a device:

* Enter Interface Name:

* Policy:

General

* Enable BPDU Guard: Enable spanning-tree bpduguard

Enable Port Type Fast: Enable spanning-tree edge port behavior

* MTU: MTU for the interface

* SPEED: Interface Speed

* Trunk Allowed Vlans: Allowed values: 'none', 'all', or vlan ranges (ex: 1-200,500-2000,3000)

Interface Description:

Freeform Config:

Enable Interface: Uncheck to disable the interface

Note ! All configs should strictly match 'show run' output, with respect to case and newlines. Any mismatches will yield unexpected diffs during deploy.

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

Select a device: Select the pre-provisioned device.

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

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

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

For more information, see [Adding Interfaces, on page 220](#).

Step 5 Click **Save**.

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

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

Creating a vPC Setup

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

Procedure

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

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

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

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

Select vPC peer for N5596-37



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

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

vPC Pair Template

No Policy
vpc_pair 2
No Policy

Save Cancel

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

vPC Pair Template ▼

vPC Domain | vPC Peerlink

* vPC Domain ID ? vPC

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

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

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

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

* vPC Keep-alive VRF Name ? Narr

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

* Fabricpath switch id ? Fabi

Configure VTEPS ? Check this to configure NVE source loopbac

* NVE interface ? NVE

* Peer 1 NVE source loopback interface ? Peel

vPC Domain tab: Enter the vPC domain details.

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

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

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

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

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

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

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

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


vPC Domain

vPC Peerlink

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

Step 3 Click **Save**.

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



To update vPC setup details, do the following:

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

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

Undeploying a vPC Setup

Procedure

- Step 1** Right-click a **vPC** switch and choose **vPC Pairing**.

The vPC peer screen comes up.

- Step 2** Click **Unpair** at the bottom right part of the screen.

The vPC pair is deleted and the fabric topology window appears.

- Step 3** Click **Save & Deploy**.

The **Config Deployment** dialog box appears.

- Step 4** (Optional) Click the value under the **Preview Config** column.

View the pending configuration in the **Config Preview** dialog box. The following configuration details are deleted on the switch when you unpair: vPC feature, vPC domain, vPC peerlink, vPC peerlink member ports, loopback secondary IPs, and host vPCs. However, the host vPCs and port channels are not removed. Delete these port channels from the **Interfaces** window if required.

Note Resync the fabric if it is out of sync.

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

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

Multi-Site Domain for VXLAN BGP EVPN Fabrics

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

Since server networks and VRFs are shared across the member fabrics (as one stretched network), the new networks and VRFs provisioning function is provided at the MSD fabric level. Any new network and VRF creation is only allowed for the MSD. All member fabrics inherit any new network and VRF created for the MSD.

In DCNM 11.1(1) release, in addition to member fabrics, the topology view for the MSD fabric is introduced. This view displays all member fabrics, and how they are connected to each other, in one view.

Also, a deployment view is introduced for the MSD fabric. You can deploy overlay networks (and VRFs) on member fabrics from a single topology deployment screen, instead of visiting each member fabric deployment screen separately and deploying.

**Note**

- vPC support is added for BGWs in the DCNM 11.1(1) release.
- The MSD feature is unsupported on the switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images.
- The VXLAN OAM feature in Cisco DCNM is only supported on a single fabric or site.
- After you unpair a BGW vPC, perform a **Save & Deploy** on the member fabric followed by a **Save & Deploy** of the MSD fabric.

A few fabric-specific terms:

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

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

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

Fabric and Switch Instance Variables

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


Fabric instance values can only be edited or updated in the fabric context from the VRFs and Networks window. The appropriate fabric should be selected in the **SCOPE** drop-down list to edit the fabric instance values. Some of the examples of fabric instance variables are BGP ASN, Multicast group per network or VRF, etc. For information about editing multicast group address, see [Editing Networks in the Member Fabric, on page 120](#).

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

MSD and Member Fabric Process Flow

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


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



DCNM GUI:
Control > Networks & VRFs


4

Create **networks** and **VRFs** in **MSD fabric**




5

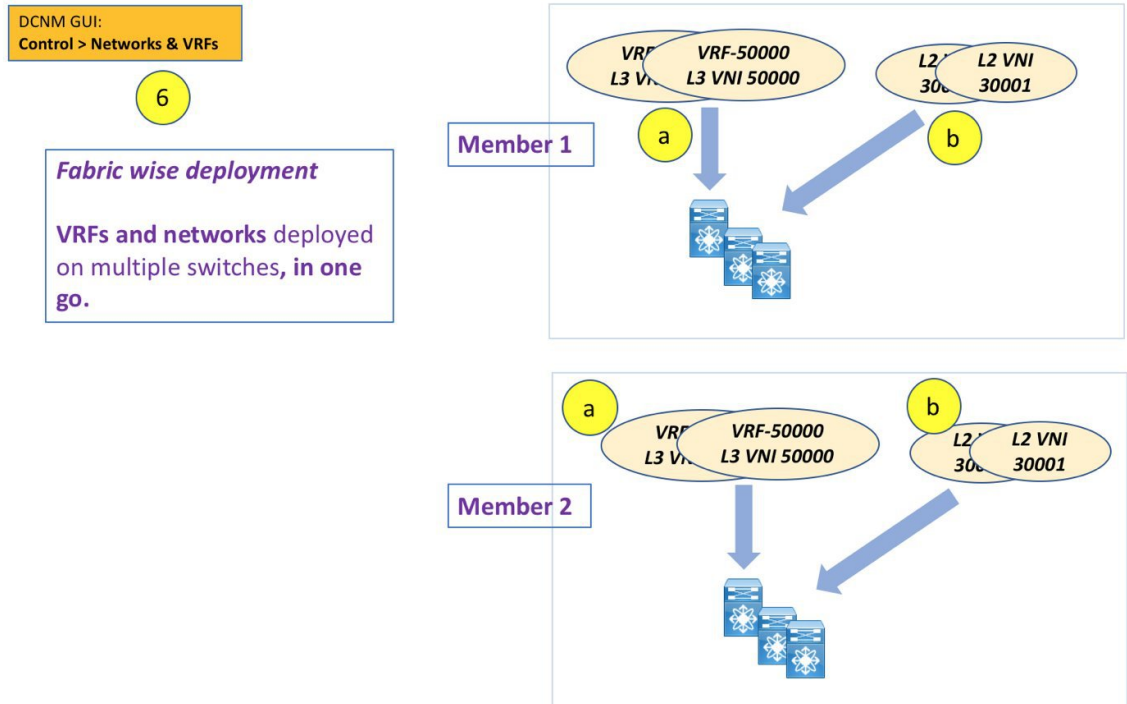
The **networks** and **VRFs** automatically get inherited to the member fabric



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



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



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



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

Upcoming sections in the document explain the following:

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

Creating an MSD Fabric and Associating Member Fabrics to It

The process is explained in two steps:

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

Creating an MSD Fabric

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

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

Fabric Builder
Fabric Builder creates a managed and controlled SDN fabric. Select an existing fabric below or define a new VXLAN fabric, add switches using Power On Auto Provisioning (POAP), set the roles of the switches and deploy settings to devices.

[Create Fabric](#)

Fabrics (4)

Fabric Name	Type	ASN	Replication Mode	Technology
External65000	External	650000		
Easy60000	Switch_Fabric	60000	Multicast	VXLANFabric
Easy7200	Switch_Fabric	7200	Multicast	VXLANFabric
MSD	MSD			

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

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

Fabric Name - Enter the name of the fabric.

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

Add Fabric ✕

* Fabric Name :

* Fabric Template : MSD_Fabric_11_1

① Fabric Template for a VXLAN EVPN Multi-Site Domain (MSD) that can contain other VXLAN EVPN fabrics with Layer-2/Layer-3 Overlay Extensions.

General	DCI	Resources	Configuration Backup
* Layer 2 VXLAN VNI Range	<input type="text" value="30000-49000"/>	<small>① Overlay Network Identifier Range (Min:1, Max:16777214)</small>	
* Layer 3 VXLAN VNI Range	<input type="text" value="50000-59000"/>	<small>① Overlay VRF Identifier Range (Min:1, Max:16777214)</small>	
* VRF Template	<input type="text" value="Default_VRF_Universal"/>	<small>① Default Overlay VRF Template For Leafs</small>	
* Network Template	<input type="text" value="Default_Network_Universal"/>	<small>① Default Overlay Network Template For Leafs</small>	
* VRF Extension Template	<input type="text" value="Default_VRF_Extension_Universal"/>	<small>① Default Overlay VRF Template For Borders</small>	
* Network Extension Template	<input type="text" value="Default_Network_Extension_Universa"/>	<small>① Default Overlay Network Template For Borders</small>	
Anycast-Gateway-MAC	<input type="text" value="2020.0000.00aa"/>	<small>① Shared MAC address for all leaves</small>	
* Multi-Site Routing Loopback Id	<input type="text" value="100"/>	<small>① (Min:0, Max:1023)</small>	
ToR Auto-deploy Flag	<input type="checkbox"/>	<small>① Enables Overlay VLANs on uplink between ToRs and Leafs</small>	

The fields in the screen are explained:

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

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

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

VRF Template - Default VRF template.

Network Template - Default network template.

VRF Extension Template - Default VRF extension template.

Network Extension Template - Default network extension template.

Anycast-Gateway-MAC - Anycast gateway MAC address.

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

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

3. Click the DCI tab.

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

- Multi-Site Overlay IFC Deployment Method:** Manual (dropdown menu)
- Multi-Site Route Server List:** (empty text field)
- Multi-Site Route Server BGP ASN List:** (empty text field)
- Multi-Site Underlay IFC Auto Deployment Flag:** (checkbox)
- Delay Restore time:** 300 (text field)
- Multi-Site CloudSec:** (checkbox)
- CloudSec Key String:** (empty text field)
- CloudSec Cryptographic Algorithm:** (empty dropdown menu)
- CloudSec Enforcement:** (empty dropdown menu)

The fields are:

Multi-Site Overlay IFC Deploy Method – Choose how you will connect the data centers through the BGW, manually, in a back-to-back fashion or through a route server.

If you choose to connect them through a route server, you should enter the route server details.

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

Multi-Site Route Server BGP ASN List – Specify the BGP AS Number of the router server. If you specify more than one route server, separate the AS Numbers by a comma.

Multi-Site Underlay IFC Auto Deployment Flag - Check the check box to enable auto configuration. Uncheck the check box for manual configuration.

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

Multi-Site CloudSec – Enables CloudSec configurations on border gateways. If you enable this field, the remaining three fields for CloudSec are editable. For more information, see [Support for CloudSec in Multi-Site Deployment, on page 126](#).

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

eBGP Password - Specifies the encrypted eBGP Password Hex String.

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

- Click the **Resources** tab.

General DCI **Resources** Configuration Backup

* Multi-Site Routing Loopback IP Range ⓘ Typically Loopback100 IP Address Range

* DCI Subnet IP Range ⓘ Address range to assign P2P DCI Links

* Subnet Target Mask ⓘ Target Mask for Subnet Range (Min:8, Max:31)

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

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

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

- Click the **Configuration Backup** tab.

General DCI Resources **Configuration 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)

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

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

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

The backup process is initiated after you click Save.


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

- Click **Save**.

A message appears briefly at the bottom right part of the screen, indicating that you have created a new MSD fabric. After fabric creation, the fabric page comes up. The fabric name *MSD-Parent-Fabric* appears at the top left part of the screen.



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




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

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

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

Fabrics (5)



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

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

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

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

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

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

The screenshot shows a configuration page with tabs: General, Advanced, Resources, Manageability, Bootstrap, and Configuration Backup settings. The Resources tab is active. It contains several configuration fields with asterisks indicating required fields:

- Static 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 Lin
- * Layer 2 VXLAN VNI Range: 30000-49000 ? Overlay Network Identifier Range (Min:1, Max:16
- * Layer 3 VXLAN VNI Range: 50000-59000 ? Overlay VRF Identifier Range (Min:1, Max:16777
- * Network VLAN Range: 2300-2999 ? Per Switch Overlay Network VLAN Range (Min:2

The Layer 2 and Layer 3 VXLAN VNI Range fields are highlighted with a red box.

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

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


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

Other pointers:


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

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

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



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



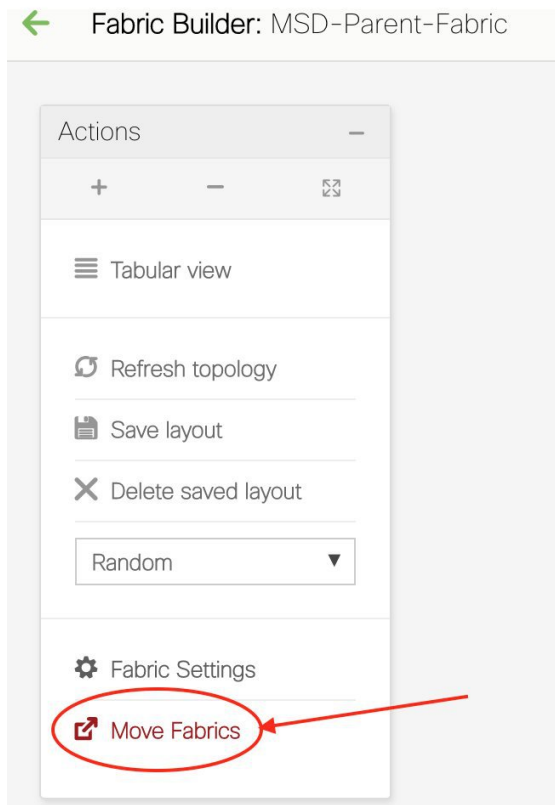
Moving the Member1 Fabric Under MSD-Parent-Fabric

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

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

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

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



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

Move Fabric

Selected 0 / Total 2 

	Fabric Name ▲	Fabric State
<input type="radio"/>	Member1	standalone
<input type="radio"/>	Test	standalone

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


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

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


- Click **Add**.

Immediately, a message appears at the top of the screen indicating that the *Member1* fabric is now associated with the MSD fabric *MSD-Parent-Fabric*. Now, the MSD-Parent-Fabric fabric page appears again.

- Click the **Move Fabrics** option to check the fabric status. You can see that the fabric status has changed from standalone to member.




5. Close this screen.
6. Click ← above the Actions panel to go to the Fabric Builder page.
You can see that *Member1* is now added to MSD fabric and is displayed in the **Member Fabrics** field.




MSD Fabric Topology View Pointers

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


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



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



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



Adding and Editing Links

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

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

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

Creating and Deploying Networks and VRFs in an MSD Fabric

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

For example, consider an MSD fabric with two member fabrics. If you create three networks in the MSD fabric, then all three networks will automatically be available for deployment in both the member fabrics.

Though member fabrics inherit the MSD fabric's networks and VRFs, you have to deploy the networks and VRFs distinctly, for each fabric.

In DCNM 11.1(1) release, a deployment view is introduced for the MSD, in addition to the per-fabric deployment view. In this view, you can view and provision overlay networks for all member fabrics within the MSD, at once. However, you still have to apply and save network and VRF configurations distinctly, for each fabric.



Note Networks and VRFs are the common identifiers (represented across member fabrics) that servers (or end hosts) are grouped under so that traffic can be sent between the end hosts based on the network and VRF IDs, whether they reside in the same or different fabrics. Since they have common representation across member fabrics, networks and VRFs can be provisioned at one go. As the switches in different fabrics are physically and logically distinct, you have to deploy the same networks and VRFs separately for each fabric.

For example, if you create networks 30000 and 30001 for an MSD that contains two member fabrics, the networks are automatically created for the member fabrics and are available for deployment.

In DCNM 11.1(1) release, you can deploy 30000 and 30001 on the border devices of all member fabrics through a single (MSD fabric) deployment screen. Prior to this, you had to access the first member fabric deployment screen, deploy 30000 and 300001 on the fabric's border devices, and then access the second member fabric deployment screen and deploy again.

Networks and VRFs are created in the MSD and deployed in the member fabrics. The steps are explained below:

1. Create networks and VRFs in the MSD fabric.
2. Deploy the networks and VRFs in the member fabric devices, one fabric at a time.

Creating Networks in the MSD Fabric

1. Click **Control** > **Networks** (under **Fabrics** submenu).

The Networks screen comes up.

2. Choose the correct fabric from SCOPE. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

Fabric Selected: bgp2

Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
<input checked="" type="checkbox"/> MyNetwork_30000	30000	NA			NA	

3. Select *MSD-Parent-Fabric* from the list and click **Continue** at the top right part of the screen.

/ VRF Selection > Network / VRF Deployment > 2 Continue

Select a Fabric

Choose a fabric with appropriate switches where you want the Top Down functionality to be enabled

MSD-Parent-Fabric 1 ▼

The Networks page comes up. This lists the list of networks created for the MSD fabric. Initially, this screen has no entries.

Fabric Selection > Network / VRF Selection > Network / VRF Deployment > VRF View Continue

Fabric Selected: MSD-Parent-Fabric

Networks Selected 0 / Total 0

+ ✎ ✕ 📄 📄

Show All ▼

<input type="checkbox"/>	Network Name ▲	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
No data available							

- Click the + button at the top left part of the screen (under **Networks**) to add networks to the MSD fabric. The Create Network screen comes up. Most of the fields are autopopulated.

Create Network
✕

▼ Network Information

* Network ID

* Network Name

* VRF Name +

Layer 2 Only

* Network Template

* Network Extension Template

VLAN ID Propose VLAN ?

▼ Network Profile

Generate Multicast IP

ⓘ Please click only to generate a New Multicast Group Address and override the default value!

General

Advanced

IPv4 Gateway/NetMask ⓘ example 192.0.2.1/24

IPv6 Gateway/Prefix L... ⓘ example 2001:db8::1/64,2001:db9::1/64

Vlan Name ⓘ if > 32 chars enable:system vlan long-nam

Interface Description ⓘ

MTU for L3 interface ⓘ 68-9216

IPv4 Secondary GW1 ⓘ example 192.0.2.1/24

IPv4 Secondary GW2 ⓘ example 192.0.2.1/24

Create Network

The fields in this screen are:

Network ID and **Network Name** - Specifies the Layer 2 VNI and name of the network. The network name should not contain any white spaces or special characters except underscore (_) and hyphen (-).

VRF Name - Allows you to select the Virtual Routing and Forwarding (VRF).

When no VRF is created, this field is blank. If you want to create a new VRF, click the + button. The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:).



Note You can also create a VRF by clicking the VRF View button on the Networks page.

Layer 2 Only - Specifies whether the network is Layer 2 only.

Network Template - Allows you to select a network template.

Network Extension Template - This template allows you to extend the network between member fabrics.

VLAN ID - Specifies the corresponding tenant VLAN ID for the network.

Network Profile section contains the General and Advanced tabs, explained below.

General tab

IPv4 Gateway/NetMask - Specifies the IPv4 address with subnet.

IPv6 Gateway/Prefix - Specifies the IPv6 address with subnet.

VLAN Name - Enter the VLAN name.

If the VLAN is mapped to more than one subnet, enter the anycast gateway IP addresses for those subnets.

Interface Description - Specifies the description for the interface.

MTU for the L3 interface - Enter the MTU for Layer 3 interfaces.

IPv4 Secondary GW1 - Enter the gateway IP address for the additional subnet.

IPv4 Secondary GW2 - Enter the gateway IP address for the additional subnet.

Advanced tab - Optionally, specify the advanced profile settings by clicking the **Advanced** tab. The options are:

- ARP Suppression
 - DHCPv4 Server 1 and DHCPv4 Server 2 - Enter the DHCP relay IP address of the first and second DHCP servers.
 - DHCPv4 Server VRF - Enter the DHCP server VRF ID.
 - Loopback ID for DHCP Relay interface - Enter the loopback ID of the DHCP relay interface.
 - Routing Tag – The routing tag is autopopulated. This tag is associated with each gateway IP address prefix.
 - TRM enable – Select the check box to enable TRM.
- For more information, see [Overview of Tenant Routed Multicast, on page 148](#).
- L2 VNI Route-Target Both Enable - Select the check box to enable automatic importing and exporting of route targets for all L2 virtual networks.



Note From Cisco DCNM Release 11.5(1), the **Enable L3 Gateway on Border** field is not available as part of the MSD network settings. You can enable a Layer 3 gateway on the border switches at a fabric level. For more information, see [Creating Networks for the Standalone Fabric, on page 234](#).

In the MSD fabric level, if the **Enable L3 Gateway on Border** check box is selected and you are upgrading to Cisco DCNM Release 11.5(1), then it is automatically removed from the MSD fabric level during upgrade.

- A sample of the Create Network screen:

5. Click **Create Network**. A message appears at the bottom right part of the screen indicating that the network is created. The new network (*MyNetwork_30000*) appears on the Networks page that comes up.

Fabric Selected: MSD-Parent-Fabric

Networks Selected 1 / Total 1

	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
<input checked="" type="checkbox"/>	MyNetwork_30000	30000	MyVRF_50000	20.10.1.1/24		NA	

Editing Networks in the MSD Fabric

1. In the Networks screen of the MSD fabric, select the network you want to edit and click the Edit icon at the top left part of the screen.

Fabric Selected: MSD-Parent-Fabric

Networks Selected 1 / Total 1

	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
<input checked="" type="checkbox"/>	MyNetwork_30000	30000	MyVRF_50000	20.10.1.1/24		NA	

The Edit Network screen comes up.

Edit Network

▼ Network Information

* Network ID

* Network Name

* VRF Name

Layer 2 Only

* Network Template

* Network Extension Template

VLAN ID

▼ Network Profile

General

Advanced

IPv4 Gateway/NetMask ? *example 192.0.2.1/24*

IPv6 Gateway/Prefix ? *example 2001:db8::1/64*

Vlan Name ?

Interface Description ?

MTU for L3 interface ? *[68-9216]*

IPv4 Secondary GW1 ? *example 192.0.2.1/24*

IPv4 Secondary GW2 ? *example 192.0.2.1/24*

You can edit the **Network Profile** part (**General** and **Advanced** tabs) of the MSD fabric network.

2. Click **Save** at the bottom right part of the screen to save the updates.

Network Inheritance from MSD-Parent-Fabric to Member1

MSD-Parent-Fabric fabric contains one member fabric, *Member1*. Go to the Select a Fabric page to access the *Member1* fabric.

1. Click **Control** > **Networks** (under **Fabrics** submenu).

The Networks screen comes up.

2. Choose the correct fabric from SCOPE. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

The screenshot shows the Cisco Data Center Network Manager interface. At the top, the title is "Data Center Network Manager" with a "SCOPE:" dropdown set to "bgp2" and a user profile "admin". Below the title bar, there are navigation tabs: "Network / VRF Selection" (active) and "Network / VRF Deployment". There are also "VRF View" and "Continue" buttons. The main content area shows "Fabric Selected: bgp2" and a "Networks" table. The table has columns for Network Name, Network ID, VRF Name, IPv4 Gateway/Subnet, IPv6 Gateway/Prefix, Status, and VLAN ID. One network is listed: "MyNetwork_30000" with Network ID "30000", VRF Name "NA", and Status "NA".

Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
MyNetwork_30000	30000	NA			NA	

Editing Networks in the Member Fabric

An MSD can contain multiple fabrics. These fabrics forward BUM traffic via Multicast or Ingress replication. Even if all the fabrics use multicast for BUM traffic, the multicast groups within these fabrics need not be the same.

When you create a network in MSD, it is inherited by all the member fabrics. However, the multicast group address is a fabric instance variable. To edit the multicast group address, you need to navigate to the member fabric and edit the network. For more information about the **Multicast Group Address** field, see *Creating Networks for the Standalone Fabric*.

1. Select the network and click the **Edit** option at the top left part of the window. The **Edit Network** window comes up.
2. Update the multicast group address in one of the following ways:
 - Under **Network Profile**, click the **Generate Multicast IP** button to generate a new multicast group address for the selected network, and click **Save**.
 - Click the **Advanced** tab in the **Network Profile** section, update the multicast group address, and click **Save**.



Note The **Generate Multicast IP** option is only available for member fabric networks and not MSD networks.

Deleting Networks in the MSD and Member Fabrics

You can only delete networks from the MSD fabric, and not member fabrics. To delete networks and corresponding VRFs in the MSD fabric, follow this order:

1. Undeploy the networks on the respective fabric devices before deletion.
2. Delete the networks from the MSD fabric. To delete networks, use the delete (**X**) option at the top left part of the Networks screen. You can delete multiple networks at once.



Note When you delete networks from the MSD fabric, the networks are automatically removed from the member fabrics too.

3. Undeploy the VRFs on the respective fabric devices before deletion.

4. Delete the VRFs from the MSD fabric by using the delete (X) option at the top left part of the screen. You can delete multiple VRF instances at once.

Creating VRFs in the MSD Fabric

1. From the MSD fabric's Networks page, click the **VRF View** button at the top right part of the screen to create VRFs.
 - a. Choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

The screenshot shows the Cisco Data Center Network Manager interface. At the top, the breadcrumb is 'Network / VRF Selection > Network / VRF Deployment'. The 'SCOPE' dropdown is set to 'bgp2'. Below the breadcrumb, there are buttons for 'Network View' and 'Continue'. The main heading is 'VRFs' and 'Fabric Selected: bgp2'. There are icons for adding, editing, deleting, and refreshing. A table lists VRFs with columns for VRF Name, VRF ID, and Status. One VRF is listed: MyVRF_50000 with VRF ID 50000 and Status NA. The 'Show' dropdown is set to 'All'.

- b. Choose the MSD fabric (*MSD-Parent-Fabric*) from the drop-down box and click **Continue**. The Networks page comes up.
- c. Click **VRF View** at the top right part of the Networks page].

The VRFs page comes up. This lists the list of VRFs created for the MSD fabric. Initially, this screen has no entries.

The screenshot shows the Cisco Data Center Network Manager interface. The breadcrumb is 'Fabric Selection > Network / VRF Selection > Network / VRF Deployment'. The 'SCOPE' dropdown is set to 'MSD-Parent-Fabric'. Below the breadcrumb, there are buttons for 'Network View' and 'Continue'. The main heading is 'VRFs' and 'Fabric Selected: MSD-Parent-Fabric'. There are icons for adding, editing, deleting, and refreshing. The table is empty with the text 'No data available' below it. The 'Show' dropdown is set to 'All'.

2. Click the + button at the top left part of the screen to add VRFs to the MSD fabric. The Create VRF screen comes up. Most of the fields are autopopulated.

The fields in this screen are:

VRF ID and **VRF Name** - The ID and name of the VRF.

The VRF ID is the VRF VNI or the L3 VNI of the tenant.



Note For ease of use, the VRF creation option is also available while you create a network.

VRF Template - This is populated with the *Default_VRF* template.

VRF Extension Template - This template allows you to extend the VRF between member fabrics.

3. **General** tab – Enter the VLAN ID of the VLAN associated with the VRF, the corresponding Layer 3 virtual interface, and the VRF ID.
4. **Advanced** tab

Routing Tag – If a VLAN is associated with multiple subnets, then this tag is associated with the IP prefix of each subnet. Note that this routing tag is associated with overlay network creation too.

Redistribute Direct Route Map – Specifies the route map name for redistribution of routes in the VRF.

Max BGP Paths and **Max iBGP Paths** – Specifies the maximum BGP and iBGP paths.

TRM Enable – Select the check box to enable TRM.

If you enable TRM, then the RP address, and the underlay multicast address must be entered.

For more information, see [Overview of Tenant Routed Multicast, on page 148](#).

Is RP External – Enable this checkbox if the RP is external to the fabric. If this field is unchecked, RP is distributed in every VTEP.

RP Address – Specifies the IP address of the RP.

RP Loopback ID – Specifies the loopback ID of the RP, if **Is RP External** is not enabled.

Underlay Multicast Address – Specifies the multicast address associated with the VRF. The multicast address is used for transporting multicast traffic in the fabric underlay.



Note The multicast address in the **Default MDT Address for TRM VRFs** field in the fabric settings screen is auto-populated in this field. You can override this field if a different multicast group address should be used for this VRF.

Overlay Multicast Groups – Specifies the multicast group subnet for the specified RP. The value is the group range in “ip pim rp-address” command. If the field is empty, 224.0.0.0/24 is used as default.

Enable IPv6 link-local Option - Select the check box to enable the IPv6 link-local option under the VRF SVI. If this check box is unchecked, IPv6 forward is enabled.

Advertise Host Routes - Select the checkbox to control advertisement of /32 and /128 routes to Edge Routers.

Advertise Default Route - Select the checkbox to control advertisement of default routes within the fabric.

A sample screenshot:

Advanced tab:

5. Click **Create VRF**.

The *MyVRF_50000* VRF is created and appears on the VRFs page.


Fabric Selected: MSD-Parent-Fabric

Selected 1 / Total 1

VRFs		
VRF Name	VRF ID	Status
<input checked="" type="checkbox"/> MyVRF_50000	50000	NA

Editing VRFs in the MSD Fabric

1. In the VRFs screen of the MSD fabric, select the VRF you want to edit and click the Edit icon at the top left part of the screen.



The Edit VRF screen comes up.

Edit VRF
✕

▼ VRF Information

* VRF ID:

* VRF Name:

* VRF Template:

VRF Extension Template:

▼ VRF Profile

General

Advanced

VRF Vlan Name: ?

VRF Intf Description: ?

VRF Description: ?

You can edit the **VRF Profile** part (**General** and **Advanced** tabs).

2. Click **Save** at the bottom right part of the screen to save the updates.

VRF Inheritance from MSD-Parent-Fabric to Member1

MSD-Parent-Fabric contains one member fabric, *Member1*. Do the following to access the member fabric page.

1. Choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

SCOPE: bgp2 admin

Network / VRF Selection > Network / VRF Deployment > Network View | Continue

Fabric Selected: bgp2

VRFs Selected 1 / Total 1

VRF Name	VRF ID	Status
<input checked="" type="checkbox"/> MyVRF_50000	50000	NA

- Click the **VRF View** button. On the VRFs page, you can see that the VRF created for the MSD is inherited to its member.

Fabric Selected: Member1

VRFs Selected 0 / Total 1

VRF Name	VRF ID	Status
<input type="checkbox"/> MyVRF_50000	50000	NA

Deleting VRFs in the MSD and Member Fabrics

You can only delete networks from the MSD fabric, and not member fabrics. To delete networks and corresponding VRFs in the MSD fabric, follow this order:

- Undeploy the networks on the respective fabric devices before deletion.
- Delete the networks from the MSD fabric.
- Undeploy the VRFs on the respective fabric devices before deletion.
- Delete the VRFs from the MSD fabric by using the delete (**X**) option at the top left part of the screen. You can delete multiple VRF instances at once.



Note When you delete VRFs from the MSD fabric, they are automatically removed from the member fabrics too.

Editing VRFs in the Member Fabric

You cannot edit VRF parameters at the member fabric level. Update VRF settings in the MSD fabric. All member fabrics are automatically updated.

Deleting VRFs in the Member Fabric

You cannot delete VRFs at the member fabric level. Delete VRFs in the MSD fabric. The deleted VRFs are automatically removed from all member fabrics.

Step 1 of the following is explained. Step 2 information is mentioned in the next subsection.

- Create networks and VRFs in the MSD fabric.
- Deploy the networks and VRFs in the member fabric devices, one fabric at a time.

Deployment and Undeployment of Networks and VRFs in Member Fabrics


Before you begin, ensure that you have created networks at the MSD fabric level since the member fabric inherits networks and VRFs created for the MSD fabric.



Note The deployment (and undeployment) of networks and VRFs in member fabrics are the same as explained for standalone fabrics. Refer [Creating and Deploying Networks and VRFs](#).

Support for CloudSec in Multi-Site Deployment

CloudSec feature allows secured data center interconnect in a multi-site deployment by supporting source-to-destination packet encryption between border gateway devices in different fabrics.



CloudSec feature is supported on Cisco Nexus 9000 Series FX2 platform with Cisco NX-OS Release 9.3(5) or later. The border gateways, border gateway spines, and border gateway superspines that are FX2 platforms, and run Cisco NX-OS Release 9.3(5) or later are referred as CloudSec capable switches.

Cisco DCNM Release 11.4(1) provides an option to enable CloudSec in an MSD fabric.



Note The CloudSec session is point to point over DCI between border gateways (BGWs) on two different sites. All communication between sites uses Multi-Site PIP instead of VIP. Enabling CloudSec requires a switch from VIP to PIP, which could cause traffic disruption for data flowing between sites. Therefore, it is recommended to enable or disable CloudSec during a maintenance window.

You can also watch the video that demonstrates how to configure the CloudSec feature. See [Video: Configuring CloudSec in Cisco DCNM](#).

Enabling CloudSec in MSD

Navigate to **Control > Fabrics > Fabric Builder**. You can either create a new MSD fabric by clicking **Create Fabric** or edit the existing MSD fabric by clicking **Edit Fabric**.

The screenshot shows the configuration page for CloudSec in MSD, with the DCI tab selected. The configuration includes:

- Multi-Site Overlay IFC Deployment Method:** Manual (with a dropdown arrow). Help text: Manual, Auto Overlay EVPN Peering to Route Servers, Auto Overlay EVPN Direct Peering to Border Gateways.
- Multi-Site Route Server List:** (empty text field). Help text: Multi-Site Router-Server peer list, e.g. 128.89.0.1, 128.89.0.2.
- Multi-Site Route Server BGP ASN List:** (empty text field). Help text: 1-4294967295 | 1-65535[0-65535], e.g. 65000, 65001.
- Multi-Site Underlay IFC Auto Deployment Flag:** (with a help icon).
- Delay Restore time:** 300 (text field). Help text: Multi-Site underlay and overlay control plane convergence time (Min:30, Max:1000) in seconds.
- Multi-Site CloudSec:** (with a help icon). Help text: Auto Config CloudSec on Border Gateways.
- CloudSec Key String:** (empty text field). Help text: Cisco Type 7 Encrypted Octet String.
- CloudSec Cryptographic Algorithm:** AES_128_CMAC (dropdown menu). Help text: AES_128_CMAC or AES_256_CMAC.
- CloudSec Enforcement:** (empty dropdown menu). Help text: If set to 'strict', data across site must be encrypted.
- CloudSec Status Report Timer:** 5 (text field). Help text: CloudSec Operational Status periodic report timer in minutes.
- Enable Multi-Site eBGP Password:** (with a help icon). Help text: eBGP password for Multi-Site underlay/overlay IFCs.
- eBGP Password:** (empty text field). Help text: Encrypted eBGP Password Hex String.
- eBGP Authentication Key Encryption Type:** (empty dropdown menu). Help text: BGP Key Encryption Type: 3 - 3DES, 7 - Cisco.

Buttons for **Save** and **Cancel** are located at the bottom right of the configuration area.

Under the **DCI** tab, you can specify the CloudSec configuration details.

Multi-Site CloudSec – Enables CloudSec configurations on border gateways. If you enable this field, the remaining CloudSec fields are editable.

Multi-Site CloudSec – Enables CloudSec configurations on border gateways. If you enable this field, the remaining three fields for CloudSec are editable.

When Cloudsec is enabled at MSD level, DCNM also enables **dc advertise-pip** under **evpn multisite border-gateway** and **tunnel-encryption** on the uplinks for all Cloudsec capable gateways.

When you click **Save & Deploy**, you can verify these configs in the **Preview Config** window for the border gateway switches.

Note – CloudSec isn't supported if the border gateway has vPC or TRM is enabled on it, that is, TRM enabled on multisite overlay IFC. If CloudSec is enabled in this scenario, appropriate warning or error messages are generated.

CloudSec Key String – Specifies the hex key string. Enter a 66 hexadecimal string if you choose **AES_128_CMAC** or enter a 130 hexadecimal string if you choose **AES_256_CMAC**.

CloudSec Cryptographic Algorithm – Choose **AES_128_CMAC** or **AES_256_CMAC**.

CloudSec Enforcement – Specifies whether the CloudSec enforcement should be strict or loose.

strict – Deploys the CloudSec configuration to all the border gateways in fabrics in MSD. If there are any border gateways that don't support CloudSec, then an error message is generated, and the configuration isn't pushed to any switch.

If **strict** is chosen, the **tunnel-encryption must-secure** CLI is pushed to the CloudSec enabled gateways within MSD.

loose – Deploys the CloudSec configuration to all the border gateways in fabrics in MSD. If there are any border gateways that don't support CloudSec, then a warning message is generated. In this case, the CloudSec config is only deployed to the switches that support CloudSec. If **loose** is chosen, the **tunnel-encryption must-secure** CLI is removed if it exists.



Note There should be at least two fabrics in MSD with border gateways that support CloudSec. If there's only one fabric with a CloudSec capable device, then the following error message is generated:

CloudSec needs to have at least 2 sites that can support CloudSec.

To remove this error, meet the criteria of having at least two sites that can support CloudSec or disable CloudSec.

CloudSec Status Report Timer – Specifies the CloudSec Operational Status periodic report timer in minutes. This value specifies how often the DCNM polls the CloudSec status data from the switch. The default value is 5 minutes and the range is from 5 to 60 minutes.

Using the CloudSec feature in DCNM, you can have all the gateways within the MSD to use the same keychain (and have only one key string) and policy. You can provide one key chain string for DCNM to form the key chain policy. DCNM forms the encryption-policy by taking all default values. DCNM pushes the same key chain policy, the same encryption-policy, and encryption-peer policies to each CloudSec capable gateways. On each gateway, there's one encryption-peer policy for each remote gateway that is CloudSec capable, using the same keychain and same key policy.

If you don't want to use the same key for the whole MSD fabric or want to enable CloudSec on a subset of all sites, you can use **switch_freeform** to manually push the CloudSec config to the switches.

Capture all the CloudSec config in **switch_freeform**.

For example, the below config is included in the **switch_freeform** policy:

```
feature tunnel-encryption
evpn multisite border-gateway 600
  dci-advertise-pip
tunnel-encryption must-secure-policy
tunnel-encryption policy CloudSec_Policy1
tunnel-encryption source-interface loopback20
key chain CloudSec_Key_Chain1 tunnel-encryption
  key 1000
  key-octet-string 7 075e731f1a5c4f524f43595f507f7d73706267714752405459070b0b0701585440
cryptographic-algorithm AES_128_CMA
tunnel-encryption peer-ip 192.168.0.6
  keychain CloudSec_Key_Chain1 policy CloudSec_Policy1
```

Add **tunnel-encryption** in the Freeform Config of the uplink interface policy which will generate config like the following:

```
interface ethernet1/13
  no switchport
  ip address 192.168.1.14/24 tag 54321
  evpn multisite dci-tracking
  tunnel-encryption
  mtu 9216
  no shutdown
```

For more information, see [Enabling Freeform Configurations on Fabric Switches](#), on page 299.

When CloudSec configuration is added to or removed from the switch, the DCI uplinks will flap, which will trigger multisite BGP session flapping. For multisite with existing cross site traffic, there will be traffic

disruption during this transition. Therefore, it is recommended to make the transition during a maintenance window.

If you're migrating an MSD fabric with the CloudSec configuration into DCNM, the CloudSec related configuration is captured in **switch_freeform** and interface freeform config. You do not need to turn on Multi-Site CloudSec in the MSD fabric setting. If you want to add more fabrics and establish CloudSec tunnels which share the same CloudSec policy including key as the existing one, then you can enable the CloudSec config in the MSD fabric settings. The CloudSec parameters in the MSD fabric setting need to match the existing CloudSec configuration on the switch. The CloudSec configuration is already captured in the freeform config, and enabling CloudSec in MSD will also generate config intents. Therefore, there's a double intent. For example, if you want to change the CloudSec key in the MSD settings, you need to remove the CloudSec freeform config because DCNM won't modify config in **switch_freeform**. Otherwise, the key in the MSD fabric settings is a conflict with the key in the freeform config.

Viewing CloudSec Operational State

From Cisco DCNM 11.5(1), you can use **CloudSec Operational View** to check the operational status of the CloudSec sessions if CloudSec is enabled on the MSD fabric.

Procedure

-
- Step 1** Choose an MSD fabric.
The fabric topology window appears.
- Step 2** Click **Tabular view** in the **Actions** pane.
- Step 3** Choose the **CloudSec Operational View** tab.
- Step 4** If CloudSec is disabled, the **CloudSec Operational View** tab isn't displayed.
The **Operational View** tab has the following fields and descriptions.

Fields	Descriptions
Fabric Name	Specifies the fabrics that have a CloudSec session.
Session	Specifies the fabrics and border gateway switches involved in the CloudSec session.
Link State	Specifies the status of the CloudSec session. It can be in one of the following states: <ul style="list-style-type: none"> • Up: The CloudSec session is successfully established between the switches. • Down: The CloudSec session isn't operational.
Uptime	Specifies the duration of the uptime for the CloudSec session. Specifically, it's the uptime since the last Rx and Tx sessions flapped, and the smaller value among the 2 sessions is displayed.
Oper Reason	Specifies the down reason for the CloudSec session state.

All these columns are sortable.

Note After CloudSec is enabled on a fabric, the operational status may not be available until after sessions are created, and the next status poll occurs.

Troubleshooting a CloudSec Session

If a CloudSec session is down, you can find more information about it using Programmable Report.

Procedure

- Step 1** Navigate to **Applications > Programmable report**.
- Step 2** Click the **Create Report** icon.
- Step 3** Specify a report name, select the MSD fabric on which the report job should be run, and click **Next**.
- Step 4** From the **Template** drop-down list, select **fabric_cloudsec_oper_status** and click **Create Job**.
The status will change to a green tick indicating Success after the report has been successfully generated.
- Step 5** Click the report to view it. This report is similar to the **CloudSec Operation View** tab.
- Step 6** Click **View Details** to view more information about the CloudSec session status.
- Step 7** Click the operational status for a session to view the detailed info about the CloudSec session for each peer fabric and device.

The screenshot displays the Cisco Data Center Network Manager interface. The main heading is "Report" and the breadcrumb path is "/ msd-fabric". The report title is "CloudSec Operational Status Summary for Fabric msd-fabric".

FABRIC NAME	SESSION	STATE	DOWN REASON	UPTIME
fab2<->fab3	fab2.stewong-n9kfx2-6-...	Down	0x4(NVE-Intf-Down,)	-
fab1<->fab3	fab1.stewong-n9kfx2-3-...	Down	0x4(NVE-Intf-Down,)	-
fab1<->fab2	fab1.stewong-n9kfx2-3-...	Up	N/A	06:08:33

Below the summary, there is a detailed view for "CloudSec Operational Status for FDO23240P02.stewong-n9kfx2-3".

PEER IP	PEER FABRIC	PEER DEVICE	LOCAL FABRIC	STATE	RX SESSION STATUS	TX SESSION STATUS	LAST RX SESSION FLAPPED	LAST TX SESSION FLAPPED
10.3.102.1	fab2	stewong-n9kfx2-6	fab1	Up	Secure (AN: 0)	Secure (AN: 0)	06:08:33	06:08:33
10.3.103.1	fab3	stewong-n9kfx2-4	fab1	Up	Secure (AN: 0)	Pending (No-Key-r...	06:08:36	never

Removing a Fabric From an MSD

To remove a fabric from an MSD fabric, perform the following steps:

Before you begin

Make sure that there are no VRFs deployed on the border switches in the fabric that you want to remove. For more information, see [Deployment and Undeployment of Networks and VRFs in Member Fabrics, on page 126](#).



Note From Cisco DCNM Release 11.4(1), after removing an individual fabric from MSD, underlay and overlay IFCs are deleted. If IFCs are extended, an error is reported to disallow the fabric remove.

Procedure

- Step 1** From the **Fabric Builder** window, click an MSD fabric.
- Step 2** Click **Move Fabric** in the **Actions** menu.
- Step 3** In the **Move Fabric** window, select the respective radio button of the fabric that you want to remove and click **Remove**.
- In the fabric removal notification window, click **Close**.
- Step 4** Click **Save & Deploy** for the MSD in the **Fabric Builder** window.
- Step 5** Click **Deploy Config** in the **Config Deployment** window.
- Click **Close**.
- Step 6** Navigate to the fabric that you removed from MSD and click **Save & Deploy**.
- Step 7** Click **Deploy Config** in the **Config Deployment** window.
- Click **Close**.
-

Moving a Standalone Fabric (With Existing Networks and VRFs) to an MSD Fabric

If you move a standalone fabric with existing networks and VRFs to an MSD fabric as a member, ensure that common networks (that is, L2 VNI and L3 VNI information), anycast gateway MAC, and VRF and network templates are the same across the fabric and the MSD. DCNM validates the standalone fabric (network and VRF information) against the (network and VRF information) of the MSD fabric to avoid duplicate entries. An example of duplicate entries is two common network names with a different network ID. After validation for any conflicts, the standalone fabric is moved to the MSD fabric as a member fabric. Details:

- The MSD fabric inherits the networks and VRFs of the standalone fabric that do not exist in the MSD fabric. These networks and VRFs are in turn inherited by the member fabrics.
- The newly created member fabric inherits the networks and VRFs of the MSD fabric (that do not exist in the newly created member fabric).
- If there are conflicts between the standalone and MSD fabrics, validation ensures that an error message is displayed. After the updation, when you move the member fabric to the MSD fabric, the move will be successful. A message comes up at the top of the page indicating that the move is successful.

If you move back a member fabric to standalone status, then the networks and VRFs remain as they are, but they remain relevant as in an independent fabric, outside the purview of an MSD fabric.

Managing Switches Using LAN Classic Templates

From Cisco DCNM Release 11.4(1), you can use the **LAN_Classic** and **Fabric_Group** templates to manage the switches that you used to previously manage in the DCNM Classic LAN deployment.

The **LAN_Classic** fabric template is a generic fabric template to manage Cisco Nexus switches.

Guidelines and Limitations

- Fabrics using the **LAN_Classic** fabric template can be changed to use the **External_Fabric_11_1** fabric template and then use all its associated functionalities. Note that this is the only supported fabric template conversion and it's nonreversible.
- The **LAN_Classic** fabric can be added as a member of an MSD fabric.
- Only Cisco Nexus switches are supported in the **LAN_Classic** fabric.
- The TOR Auto-Deploy functionality is supported in the **LAN_Classic** member fabric when a switch with the **ToR** role is in the fabric. For more information, see *Configuring ToR Switches and Deploying Networks*.
- If you are using the Cisco Nexus 7000 Series Switch with Cisco NX-OS Release 6.2(24a) on the LAN Classic or External fabrics, make sure to enable AAA IP Authorization in the fabric settings.
- The following features in the **LAN_Classic** template provide the same support as they do for the **External_Fabric_11_1** template:

The following features are supported:

- Configuration compliance
- Backup or restore of fabric
- Network Insights
- Performance monitoring
- VMM
- Topology view
- Kubernetes visualization
- RBAC

For more information, refer to the feature specific sections.

Creating a LAN Classic Fabric

Procedure

- Step 1** Navigate to **Control > Fabrics > Fabric Builder**.
- Step 2** Click **Create Fabric**.
- Step 3** Enter the fabric name and choose **LAN_Classic** from **Fabric Template** drop - down list.

Add Fabric ✕

* Fabric Name :

* Fabric Template :

① Fabric Template to manage various switches and topologies

General | Advanced | Configuration Backup | Bootstrap

Fabric Monitor Mode ① If enabled, fabric is only monitored. No configuration will be deployed

Step 4 The **General** tab is displayed by default. The field in this tab is:

Fabric Monitor Mode – Uncheck the check box if you want DCNM to manage the fabric. Keep the check box selected to enable only monitoring of the fabric. In this state, you can't deploy configurations on its switches.

The configurations must be pushed for devices before you discover them in the fabric. You can't push configurations in the monitor mode.

Step 5 Click **Advanced** tab. The fields in this tab are:

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

Power Supply Mode - Choose the appropriate power supply mode.

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

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

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

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

Enable CDP for Bootstrapped Switch - Enables CDP on management interface.

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

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

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

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

interface belongs to. As part of the process of switch import/discovery, this information is captured in the baseline intent that is populated on the DCNM. For more information, see [Inband Management in External Fabrics and LAN Classic Fabrics, on page 139](#).

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

Enable Precision Time Protocol (PTP): Enables PTP across a fabric. When you select this check box, PTP is enabled globally and on core-facing interfaces. Additionally, the PTP Source Loopback Id and PTP Domain Id fields are editable. For more information, see [Precision Time Protocol for External Fabrics and LAN Classic Fabrics, on page 140](#).

Fabric Freeform - You can apply configurations globally across all the devices discovered in the external fabric using this freeform field.

AAA Freeform Config – Specifies the AAA freeform configs.

Step 6 Click the **Resources** tab. The fields in this tab are:

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

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

Underlay MPLS Loopback IP Range - Specifies the underlay MPLS SR or LDP loopback IP address range. The IP range should be unique, that is, it shouldn't overlap with IP ranges of the other fabrics.

Step 7 Click **Configuration** tab. The fields in this tab are:

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

Scheduled Fabric Backup: Check the check box to enable a daily backup.

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

Note Hourly or scheduled backup runs only after the next CC hourly run. Backup will run only after scheduled time is elapsed and whenever CC run happens after the elapsed time.

The backup and restore process is similar to that of an external fabric. For more information about backing up and restoring external fabrics, see [Fabric Backup and Restore, on page 267](#).

Step 8 Click **Bootstrap** tab. The fields in this tab are:

Enable Bootstrap (For NX-OS Switches Only) - Select this check box to enable the bootstrap feature for only Cisco Nexus switches.

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

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

Enable Local DHCP Server - Select this check box to initiate enabling of automatic IP address assignment through the local DHCP server. When you select this check box, all the remaining fields become editable.

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

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

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

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

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

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

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

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

Enable AAA Config - Enables AAA configure. It includes AAA configs from the **Advanced** tab during device bootup.

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

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

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

The format of the scope should be defined as:

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

For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24

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

Step 9 Click **ThousandEyes Agent** tab. This feature is supported on Cisco DCNM Release 11.5(3) only. For more information, refer to [Configuring Global Setting for ThousandEyes Enterprise Agent on Cisco DCNM](#).

General	Replication	vPC	Protocols	Advanced	Resources	Manageability	Bootstrap	Configuration Backup	ThousandEyes Agent
<p>Enable Fabric Override for ThousandEyes Agent Installation <input type="checkbox"/> ⓘ</p> <p>ThousandEyes Account Group Token <input type="text"/> ⓘ <i>Token from ThousandEyes Agent Settings for Agent Installation</i></p> <p>VRF on Switch for ThousandEyes Agent Collector Reachability <input type="text"/> ⓘ <i>NX-OS VRF that provides Internet Reachability</i></p> <p>DNS Domain <input type="text"/> ⓘ <i>DNS Domain Configuration</i></p> <p>DNS Server IPs <input type="text"/> ⓘ <i>Comma separated list of IP Addresses(v4/v6)</i></p> <p>NTP Server IPs <input type="text"/> ⓘ <i>Comma separated list of IP Addresses(v4/v6)</i></p> <p>Enable Proxy for Internet Access <input type="checkbox"/> ⓘ <i>Proxy Settings for NX-OS Switch Internet Access</i></p> <p>Proxy Information <input type="text"/> ⓘ <i>Proxy-Server:port</i></p> <p>Proxy Bypass <input type="text"/> ⓘ <i>Comma separated No-proxy server list</i></p>									
									<input type="button" value="Save"/> <input type="button" value="Cancel"/>

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.

Adding Switches to LAN Classic Fabric

Procedure

- Step 1** Click **Add** switches. The **Inventory Management** window comes up.

Inventory Management



Discover Existing Switches

PowerOn Auto Provisioning (POAP)

Discovery Information >

Scan Details >

Seed IP

Ex: 2.2.2.20 (or) 10.10.10.40-60 (or) 2.2.2.20, 2.2.2.21

Authentication Protocol

Username

Password

Max Hops

 hop(s)

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

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

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

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

- Step 6** Click **Refresh** topology to view the latest topology view.

For more information, see:

- [Discovering Existing Switches, on page 24](#)
- [Discovering New Switches, on page 29](#)

Creating a Fabric Group and Associating Member Fabrics

This procedure shows how to create a **Fabric_Group** and add **LAN_Classic** fabrics. The **Fabric_Group** template is used for grouping **LAN_Classic** fabrics for visualization.

The following functionalities aren't supported in a **Fabric_Group**:

- Fabric backup or restore
- VXLAN overlay or IFC deployment
- Changing fabric template to and from any other fabric template
- Since **Fabric_Group** doesn't manage any configurations, clicking **Save & Deploy** reports an error.

Procedure

- Step 1** Navigate to **Control > Fabrics > Fabric Builder**.
- Step 2** Click **Create Fabric**.
- Step 3** Enter the fabric name and choose **Fabric_Group** from the **Fabric Template** drop - down list.


Add Fabric
✕

* Fabric Name :

* Fabric Template : ▼

ⓘ Fabric Template that can contain other LAN Classic fabrics

- Step 4** Click **Save**.
- Step 5** In the **Actions** panel, click **Move Fabrics**.
- Step 6** Select a **LAN_Classic** fabric in the **Move Fabric** window.
- Note** You can select and add only a **LAN_Classic** fabric in a fabric group.
- Step 7** Click **Add**.
- Similarly, you can remove a member fabric by selecting it and clicking **Remove**.



Support for Inter-Fabric Connection in LAN Classic Fabric Template

The **LAN_Classic** fabric supports VRF-Lite, Multi-Site, and MPLS IFCs with these conditions:

- The **LAN_Classic** fabric as a destination for DCI/VRF-Lite and Multi-Site IFCs is supported, but you can only manually create them by providing the required information. They won't be automatically created even when the auto deployment options are enabled in the **Easy_Fabric_11_1** and **MSD_Fabric_11_1** fabrics.
- You can't add nonexistent (meta) switches to a **LAN_Classic** fabric. A meta switch is a placeholder for a switch or device that DCNM can't discover.
- The base BGP configurations for the 'Edge Router' and 'Core Router' switch roles aren't auto generated. Configure them using the **switch_freiform** policies or other suitable means.
- If MPLS Handoff is enabled in the fabric settings, MPLS base configurations are auto generated for the 'Edge Router' and 'Core Router' switch roles.

Inband Management in External Fabrics and LAN Classic Fabrics

From Release 11.5(1), Cisco DCNM allows you to import or discover switches with inband connectivity for External and LAN Classic fabrics in Brownfield deployments only. Enable inband management, per fabric, while configuring or editing the Fabric settings. You cannot import or discover switches with inband connectivity using POAP.

After configuration, the Fabric tries to discover switches based on the VRF of the inband management. The fabric template determines the VRF of inband switch using seed IP. If there are multiple VRFs for same seed IP, then no intent will be learnt for seed interfaces. You must create intent/configuration manually.

After configuring/editing the Fabric settings, you must Save and Deploy. You cannot change the Inband Mgmt settings after you import inband managed switches to the Fabric. If you uncheck the checkbox, the following error message is generated.

```
Inband IP <<IP Address>> cannot be used to import the switch,
please enable Inband Mgmt in fabric settings and retry.
```

After the switches are imported to the Fabric, you must manage the interfaces to create intent. Create the intent for the interfaces that you're importing the switch. Edit/update the Interface configuration. When you try to change the Interface IP, for this inband managed switch, an error message is generated:

```
Interface <<interface_name>> is used as seed or next-hop egress interface
for switch import in inband mode.
IP/Netmask Length/VRF changes are not allowed for this interface.
```

While managing the interfaces, for switches imported using inband management, you cannot change the seed IP for the switch. The following error will be generated:

```
<<switch-name>>: Mgmt0 IP Address (<ip-address>) cannot be changed,
when is it used as seed IP to discover the switch.
```

Create a policy for next-hop interfaces. Routes to DCNM from 3rd party devices may contain multiple interfaces, known as ECMP routes. Find the next-hop interface and create an intent for the switch. Interface IP and VRF changes are not allowed.

If inband management is enabled, during Image management, eth2 IP address is used to copy images on the switch, in ISSU, EPLD, RPM & SMU installations flows.

If you imported the switches using inband connectivity in the fabric, and later disable the inband Mgmt in the Fabric settings after deployment, the following error message is generated:

```
The fabric <<fabric name>> was updated with below message:
Fabric Settings cannot be changed for Inband Mgmt, when switches are already imported
using inband Ip. Please remove the existing switches imported using Inband Ip from the
fabric,
then change the Fabric Settings.
```

However, the same fabric can contain switches imported using both inband and out-of-band connectivity.

Precision Time Protocol for External Fabrics and LAN Classic Fabrics

From Release 11.5(1), in the fabric settings for the **External_Fabric_11_1** or **LAN_Classic** template, select the **Enable Precision Time Protocol (PTP)** check box to enable PTP across a fabric. When you select this check box, PTP is enabled globally and on core-facing interfaces. Additionally, the **PTP Loopback Id** and **PTP Domain Id** fields are editable.

The PTP feature is supported with Cisco Nexus 9000 Series cloud-scale switches, with NX-OS version 7.0(3)I7(1) or later. Warnings are displayed if there are non-cloud scale devices in the fabric, and PTP is not enabled. Examples of the cloud-scale devices are Cisco Nexus 93180YC-EX, Cisco Nexus 93180YC-FX, Cisco Nexus 93240YC-FX2, and Cisco Nexus 93360YC-FX2 switches. For more information, refer to <https://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html#~products>.



Note PTP global configuration is supported with Cisco Nexus 3000 Series switches; however, PTP and ttag configurations are not supported.

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

For External and LAN Classic fabric deployments, you have to enable PTP globally, and also enable PTP on core-facing interfaces. The interfaces could be configured to the external PTP server like a VM or Linux-based machine. Therefore, the interface should be edited to have a connection with the grandmaster clock. For PTP and TTAG configurations to be operational on External and LAN Classic Fabrics, you must sync up of Switch Configs to DCNM using the **host_port_resync** policy. For more information, see [Sync up Out-of-Band Switch Interface Configurations with DCNM, on page 142](#).

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

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

Here is the sample PTP configuration:

```
feature ptp

feature ptp

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

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

interface Ethernet1/59 -> Core facing interface
  ptp

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

The following guidelines are applicable for PTP:

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

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

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

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

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

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

- TTAG configuration is generated for all the devices if host configuration sync up is performed on all the devices. Ttag configuration will not be generated for any newly added devices if host configuration sync up is not performed on all newly added devices.

If the configuration is not synced, the following warning is displayed:

```
TTAG on interfaces with PTP feature can only be configured for cloud-scale devices.
It will not be enabled on any newly added switches due to the presence of non cloud-scale
devices.
```

- PTP and TTAG configurations are deployed on host interfaces.
- PTP and TTAG Configurations are supported between switches in the same fabric (intra-fabric links). PTP is created for inter-fabric links, and ttag is created for the inter-fabric link if the other fabric (Switch) is not managed by DCNM. Inter-fabric links do not support PTP or ttag configurations if both fabrics are managed by DCNM.
- TTAG configuration is configured by default after the breakout. After the links are discovered and connected post breakout, perform Save & Deploy to generate the correct configuration based on the type of port (host, intra-fabric link, or inter fabric link).

Sync up Out-of-Band Switch Interface Configurations with DCNM

From DCNM release 11.5(1), any interface level configuration made outside of DCNM (via CLI) can be synced to DCNM and then managed from DCNM. Also, the vPC pair configurations are automatically detected and paired. This applies to the External_Fabric_11_1 and LAN_Classic fabrics only. The vPC pairing is performed with the **vpc_pair** policy.



Note When DCNM is managing switches, ensure that all configuration changes are initiated from DCNM and avoid making changes directly on the switch.

When the interface config is synced up to the DCNM intent, the switch configs are considered as the reference, that is, at the end of the sync up, the DCNM intent reflects what is present on the switch. If there were any undeployed intent on DCNM for those interfaces before the resync operation, they will be lost.

Guidelines

- Supported in fabrics using the following templates: Easy_Fabric_11_1, External_Fabric_11_1, and LAN_Classic.
- Supported for Cisco Nexus switches only.
- Supported for interfaces that don't have any fabric underlay related policy associated with them prior to the resync. For example, IFC interfaces and intra fabric links aren't subjected to resync.
- Supported for interfaces that do not have any custom policy (policy template that isn't shipped with Cisco DCNM) associated with them prior to resync.
- Supported for interfaces where the intent is not exclusively owned by a Cisco DCNM feature and/or application prior to resync.

- Supported on switches that don't have Interface Groups associated with them.
- Interface mode (switchport to routed, trunk to access, and so on) changes aren't supported with overlays attached to that interface.

The sync up functionality is supported for the following interface modes and policies:

Interface Mode	Policies
trunk (standalone, po, and vPC PO)	<ul style="list-style-type: none"> • int_trunk_host_11_1 • int_port_channel_trunk_host_11_1 • int_vpc_trunk_host_11_1
access (standalone, po, and vPC PO)	<ul style="list-style-type: none"> • int_access_host_11_1 • int_port_channel_access_host_11_1 • int_vpc_access_host_11_1
dot1q-tunnel	<ul style="list-style-type: none"> • int_dot1q_tunnel_host_11_1 • int_port_channel_dot1q_tunnel_host_11_1 • int_vpc_dot1q_tunnel_host_11_1
routed	int_routed_host_11_1
loopback	int_freeform
sub-interface	int_subif_11_1
FEX (ST, AA)	<ul style="list-style-type: none"> • int_port_channel_fex_11_1 • int_port_channel_aa_fex_11_1
breakout	interface_breakout
nve	int_freeform (only in External_Fabric_11_1/LAN_Classic)
SVI	int_freeform (only in External_Fabric_11_1/LAN_Classic)
mgmt0	int_mgmt_11_1

In an Easy fabric, the interface resync will automatically update the network overlay attachments based on the access VLAN or allowed VLANs on the interface.

After the resync operation is completed, the switch interface intent can be managed using normal DCNM procedures.

Syncing up Switch Interface Configurations to DCNM

Before you begin

- We recommend taking a fabric backup before attempting the interface resync.

- In **External_Fabric_11_1** and **LAN_Classic** fabrics, for the vPC pairing to work correctly, both the switches must be in the fabric and must be functional.
- Ensure that the switches are **In-Sync** and not in **Migration-mode** or **Maintenance-mode**.

Procedure

- Step 1** In DCNM, navigate to **Control > Fabric Builder** and click a fabric.
- Step 2** Ensure that switches are present in the fabric and vPC pairings are completed, and they are shown in the **Topology** view. Click **Tabular view** in the **Actions** panel.
- Step 3** From **Tabular view**, select one or more switches where the interface intent resync is needed, and click **Policies**.
- Note**
- If a pair of switches is already paired with either **no_policy** or **vpc_pair**, select only one switch of the pair.
 - If a pair of switches is not paired, then select both the switches.
- Step 4** In the **Policies** window, click the **Add Policy** icon.
- Step 5** In the **Add Policy** window, select **host_port_resync** from the **Policy** drop-down list. Click **Save**.

Add Policy ✕

* Policy:

* Priority (1-1000): Description:

Interface Configuration Resync Switch will be placed in Migration mode on clicking 'Save'.
ⓘ A Save & Deploy in the fabric must be performed to complete the interface configuration resync process.

Variables:

- Step 6** Check the **Mode** column for the switches to ensure that they report **Migration**. For a vPC pair, both switches are in the **Migration-mode**.
- After this step, the switches in the **Topology view** are in **Migration-mode**.
 - Both the switches in a vPC pair are in the migration mode even if one of the switches is placed into this mode.
 - If switch(es) are unintentionally put into the resync mode, they can be moved back to the normal mode by identifying the **host_port_resync** policy instance and deleting it from the **Policies** window.

Step 7 After the configuration changes are ready to sync up to DCNM, navigate to the **Tabular view**, select the required switches, and click **Rediscover switch** to ensure that DCNM is aware of any new interfaces and other changes.

Step 8 Click **Save & Deploy** to start the resync process.

Note This process might take some time to complete based on the size of the switch configuration and the number of switches involved.

Step 9 The **Config Deployment** window is displayed if no errors are detected during the resync operation. The interface intent is updated in DCNM.

Note If the External_Fabric_11_1 or LAN_Classic fabric is in **Monitored Mode**, an error message indicating that the fabric is in the read-only mode is displayed. This error message can be ignored and doesn't mean that the resync process has failed.

Config Deployment



Step 1. Configuration Preview >

Step 2. Configuration Deployment Status >

Switch Name	IP Address	Switch Serial	Preview Config	Status	Re-sync	Progress
n9k-46	80.80.80.146	FDO231003AX	0 lines	In-Sync		100%

Deploy Config

Close the **Config Deployment** window, and you can see that the switches are automatically moved out of the **Migration-mode**. Switches in a vPC pair that were not paired or paired with **no_policy** show up as paired and associated with the **vpc_pair** policy.

Note The **host_port_resync** policy that was created for the switch is automatically deleted after the resync process is completed successfully.

What to do next

The following limitations are applicable after Syncing up Switch Interface Configurations to DCNM:

- The port channel membership (once the policy exists) is not supported.
- Changing the interface mode (trunk to access etc.) that have overlays attached is not supported.
- Resync for interfaces that belong to **Interface Groups** are not supported.
- The vPC pairing in **External_Fabric_11_1** and **LAN_Classic** templates must be updated with the **vpc_pair** policy.
- Changing the interface mode that have overlays attached is not supported.
- In **Easy_Fabric** fabrics, VXLAN overlay interface attachments are performed automatically based on the allowed VLANs.

MACsec Support in Easy Fabric and eBGP Fabric

From Cisco DCNM Release 11.5(1), MACsec is supported in the Easy Fabric and eBGP Fabric on intra-fabric links. You should enable MACsec on the fabric and on each required intra-fabric link to configure MACsec. Unlike CloudSec, auto-configuration of MACsec is not supported.

MACsec is supported on switches with minimum Cisco NX-OS Releases 7.0(3)I7(8) and 9.3(5).



Note Support for MACsec is a preview feature in the Cisco DCNM Release 11.5(1).

Guidelines

- If MACsec cannot be configured on the physical interfaces of the link, an error is displayed when you click **Save**. MACsec cannot be configured on the device and link due to the following reasons:
 - The minimum NX-OS version is not met.
 - The interface is not MACsec capable.
- MACsec global parameters in the fabric settings can be changed at any time.
- MACsec and CloudSec can coexist on a BGW device.
- MACsec is not supported on Border Leaf.
- MACsec status of a link with MACsec enabled is displayed on the **Links** window.
- Brownfield migration of devices with MACsec configured is supported using switch and interface freeform configs.

For more information about MACsec configuration, which includes supported platforms and releases, see the [Configuring MACsec](#) chapter in *Cisco Nexus 9000 Series NX-OS Security Configuration Guide*.

The following sections show how to enable and disable MACsec in DCNM:

Enabling MACsec

Procedure

- Step 1** Navigate to **Control > Fabrics > Fabric Builder**.
- Step 2** Click **Create Fabric** to create a new fabric or click **Edit Fabric** on an existing Easy or eBGP fabric.
- Step 3** Click the **Advanced** tab and specify the MACsec details.

Enable MACsec – Select the check box to enable MACsec for the fabric.

MACsec Primary Key String – Specify a Cisco Type 7 encrypted octet string that is used for establishing the primary MACsec session. For AES_256_CMAC, the key string length must be 130 and for AES_128_CMAC, the key string length must be 66. If these values are not specified correctly, an error is displayed when you save the fabric.

Note The default key lifetime is infinite.

MACsec Primary Cryptographic Algorithm – Choose the cryptographic algorithm used for the primary key string. It can be AES_128_CMAC or AES_256_CMAC. The default value is AES_128_CMAC.

You can configure a fallback key on the device to initiate a backup session if the primary session fails.

MACsec Fallback Key String - Specify a Cisco Type 7 encrypted octet string that is used for establishing a fallback MACsec session. For AES_256_CMAC, the key string length must be 130 and for AES_128_CMAC, the key string length must be 66. If these values are not specified correctly, an error is displayed when you save the fabric.

MACsec Fallback Cryptographic Algorithm - Choose the cryptographic algorithm used for the fallback key string. It can be AES_128_CMAC or AES_256_CMAC. The default value is AES_128_CMAC.

MACsec Cipher Suite – Choose one of the following MACsec cipher suites for the MACsec policy:

- GCM-AES-128
- GCM-AES-256
- GCM-AES-XPB-128
- GCM-AES-XPB-256

The default value is **GCM-AES-XPB-256**.

Note The MACsec configuration is not deployed on the switches after the fabric deployment is complete. You need to enable MACsec on intra-fabric links to deploy the MACsec configuration on the switch.

MACsec Status Report Timer - Specifies MACsec operational status periodic report timer in minutes.

- Step 4** Click a fabric, click **Tabular View** in the **Actions** panel, and then click **Links**.
- Step 5** Choose an intra-fabric link on which you want to enable MACsec and click **Update Link**.
- Step 6** In the **Link Management – Edit Link** window, click **Advanced** in the **Link Profile** section, and select the **Enable MACsec** check box.

If MACsec is enabled on the intra fabric link but not in the fabric settings, an error is displayed when you click **Save**.

When MACsec is configured on the link, the following configurations are generated:

- Create MACsec global policies if this is the first link that enables MACsec.
- Create MACsec interface policies for the link.

Step 7 Click **Save** and then click **Save & Deploy** to deploy the MACsec configuration.

Disabling MACsec

To disable MACsec on an intra-fabric link, navigate to the **Link Management – Edit Link** window, unselect the **Enable MACsec** check box, click **Save**, and then click **Save & Deploy**. This action performs the following:

- Deletes MACsec interface policies from the link.
- If this is the last link where MACsec is enabled, MACsec global policies are also deleted from the device.

Only after disabling MACsec on links, navigate to the **Fabric Settings** and unselect the **Enable MACsec** check box under the **Advanced** tab to disable MACsec on the fabric. If there's an intra-fabric link in the fabric with MACsec enabled, an error is displayed when you click **Save & Deploy**.

Overview of Tenant Routed Multicast

Tenant Routed Multicast (TRM) enables multicast forwarding on the VXLAN fabric that uses a BGP-based EVPN control plane. TRM provides multi-tenancy aware multicast forwarding between senders and receivers within the same or different subnet local or across VTEPs.

With TRM enabled, multicast forwarding in the underlay is leveraged to replicate VXLAN encapsulated routed multicast traffic. A Default Multicast Distribution Tree (Default-MDT) is built per-VRF. This is an addition to the existing multicast groups for Layer-2 VNI Broadcast, Unknown Unicast, and Layer-2 multicast replication group. The individual multicast group addresses in the overlay are mapped to the respective underlay multicast address for replication and transport. The advantage of using a BGP-based approach allows the VXLAN BGP EVPN fabric with TRM to operate as fully distributed Overlay Rendezvous-Point (RP), with the RP presence on every edge-device (VTEP).

A multicast-enabled data center fabric is typically part of an overall multicast network. Multicast sources, receivers, and multicast rendezvous points might reside inside the data center but also might be inside the campus or externally reachable via the WAN. TRM allows a seamless integration with existing multicast networks. It can leverage multicast rendezvous points external to the fabric. Furthermore, TRM allows for tenant-aware external connectivity using Layer-3 physical interfaces or subinterfaces.

For more information, see the following:

- [Guidelines and Limitations for Tenant Routed Multicast](#)
- [Guidelines and Limitations for Layer 3 Tenant Routed Multicast](#)
- [Guidelines and Limitations for Layer 2/Layer 3 Tenant Routed Multicast \(Mixed Mode\)](#)

Overview of Tenant Routed Multicast with VXLAN EVPN Multi-Site

Tenant Routed Multicast with Multi-Site enables multicast forwarding across multiple VXLAN EVPN fabrics connected via Multi-Site.

The following two use cases are supported:

- Use Case 1: TRM provides Layer 2 and Layer 3 multicast services across sites for sources and receivers across different sites.
- Use Case 2: Extending TRM functionality from VXLAN fabric to sources receivers external to the fabric.

TRM Multi-Site is an extension of BGP-based TRM solution that enables multiple TRM sites with multiple VTEPs to connect to each other to provide multicast services across sites in most efficient possible way. Each TRM site is operating independently and border gateway on each site allows stitching across each site. There can be multiple Border Gateways for each site. In a given site, the BGW peers with Route Server or BGWs of other sites to exchange EVPN and MVPN routes. On the BGW, BGP will import routes into the local VRF/L3VNI/L2VNI and then advertise those imported routes into the Fabric or WAN depending on where the routes were learnt from.

Tenant Routed Multicast with VXLAN EVPN Multi-Site Operations

The operations for TRM with VXLAN EVPN Multi-Site are as follows:

- Each Site is represented by Anycast VTEP BGWs. DF election across BGWs ensures no packet duplication.
- Traffic between Border Gateways uses ingress replication mechanism. Traffic is encapsulated with VXLAN header followed by IP header.
- Each Site will only receive one copy of the packet.
- Multicast source and receiver information across sites is propagated by BGP protocol on the Border Gateways configured with TRM.
- BGW on each site receives the multicast packet and re-encapsulate the packet before sending it to the local site.

For information about guidelines and limitations for TRM with VXLAN EVPN Multi-Site, see [Configuring Tenant Routed Multicast](#).

Configuring TRM for Single Site Using Cisco DCNM

This section assumes that a VXLAN EVPN fabric has already been provisioned using Cisco DCNM.

Procedure

-
- Step 1** Enable TRM for the selected Easy Fabric. If the fabric template is **Easy_Fabric_11_1**, click the Fabric settings, navigate to the **Replication** tab, and check the **Enable Tenant Routed Multicast (TRM)** field. In addition, the default MDT multicast group field is auto-populated with a default value.

Edit Fabric ✕

* Fabric Name :

* Fabric Template :

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

General	Replication	vPC	Protocols	Advanced	Resources	Manageability	Bootstrap	Configuration Backup
<p>* Replication Mode <input type="text" value="Multicast"/> ⓘ <small>Replication Mode for BUM Traffic</small></p> <p>* Multicast Group Subnet <input type="text" value="239.1.1.0/25"/> ⓘ <small>Multicast pool prefix between 16 to 30. A multicast group IP from this pool is used for BUM traffic for each overlay network.</small></p> <p>Enable Tenant Routed Multicast (TRM) <input checked="" type="checkbox"/> ⓘ <small>For Overlay Multicast Support In VXLAN Fabrics</small></p> <p>* Default MDT Address for TRM VRFs <input type="text" value="239.1.1.0"/> ⓘ <small>Default Underlay Multicast group IP assigned for every overlay VRF.</small></p> <p>* Rendezvous-Points <input type="text" value="2"/> ⓘ <small>Number of spines acting as Rendezvous-Point (RP)</small></p> <p>* RP Mode <input type="text" value="asm"/> ⓘ <small>Multicast RP Mode</small></p> <p>* Underlay RP Loopback Id <input type="text" value="254"/> ⓘ <small>(Min:0, Max:1023)</small></p> <p>Underlay Primary RP Loopback Id <input type="text"/> ⓘ <small>Used for Bidir-PIM Phantom RP (Min:0, Max:1023)</small></p> <p>Underlay Backup RP Loopback Id <input type="text"/> ⓘ <small>Used for Fallback Bidir-PIM Phantom RP (Min:0, Max:1023)</small></p> <p>Underlay Second Backup RP Loopback Id <input type="text"/> ⓘ <small>Used for second Fallback Bidir-PIM Phantom RP (Min:0, Max:1023)</small></p> <p>Underlay Third Backup RP Loopback Id <input type="text"/> ⓘ <small>Used for third Fallback Bidir-PIM Phantom RP (Min:0, Max:1023)</small></p>								
<input type="button" value="Save"/> <input type="button" value="Cancel"/>								

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

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

Click **Save** to save the fabric settings. At this point, all the switches turn “Blue” as it will be in the pending state. Click **Save and Deploy** to enable the following:

- Enable feature ngmvpn: Enables the Next-Generation Multicast VPN (ngMVPN) control plane for BGP peering.
- Configure ip multicast multipath s-g-hash next-hop-based: Multipath hashing algorithm for the TRM enabled VRFs.
- Configure ip igmp snooping vxlan: Enables IGMP Snooping for VXLAN VLANs.
- Configure ip multicast overlay-spt-only: Enables the MVPN Route-Type 5 on all MPVN enabled Cisco Nexus 9000 switches.
- Configure and Establish MVPN BGP AFI Peering: This is necessary for the peering between BGP RR and the Leaves.

For VXLAN EVPN fabric created using Easy_Fabric_eBGP fabric template, **Enable Tenant Routed Multicast (TRM)** field and **Default MDT Address for TRM VRFs** field can be found on the fabric settings' EVPN tab.

Step 2 Enable TRM for the VRF.

Navigate to **Control > VRFs** and edit the selected VRF. Navigate to the **Advanced Tab** and edit the following TRM settings:

TRM Enable – Select the check box to enable TRM. If you enable TRM, then the RP address and the underlay multicast address must be entered.

Is RP External – Enable this check box if the RP is external to the fabric. If this field is unchecked, RP is distributed in every VTEP.

Note If the RP is external, then select the appropriate option. If the RP is external, then RP loopback ID is greyed out.

RP Address – Specifies the IP address of the RP.

RP Loopback ID – Specifies the loopback ID of the RP, if **Is RP External** is not enabled.

Underlay Multicast Address – Specifies the multicast address associated with the VRF. The multicast address is used for transporting multicast traffic in the fabric underlay.

Note The multicast address in the **Default MDT Address for TRM VRFs** field in the fabric settings screen is auto-populated in this field. User can override this field if a different multicast group address should be used for this VRF.

Overlay Multicast Groups – Specifies the multicast group subnet for the specified RP. The value is the group range in “ip pim rp-address” command. If the field is empty, 224.0.0.0/24 is used as default.

Edit VRF



VRF Information

* VRF ID

* VRF Name

* VRF Template

* VRF Extension Template

VLAN ID

VRF Profile

General	Advanced
	<p>Max iBGP Paths <input type="text" value="2"/> ⓘ 1-64</p> <p>TRM Enable <input checked="" type="checkbox"/> ⓘ Enable Tenant Routed Multicast</p> <p>Is RP External <input type="checkbox"/> ⓘ Is RP external to the fabric?</p> <p>* RP Address <input type="text" value="30.254.254.1"/> ⓘ IPv4 Address</p> <p>* RP Loopback ID <input type="text" value="500"/> ⓘ 0-1023</p> <p>* Underlay Mcast Add... <input type="text" value="239.1.1.0"/> ⓘ IPv4 Multicast Address</p> <p>Overlay Mcast Groups <input type="text"/> ⓘ 224.0.0.0/4 to 239.255.255.255/4</p> <p>Enable IPv6 link-loc... <input checked="" type="checkbox"/> ⓘ Enables IPv6 link-local Option under VRF SVI</p>

Click **Save** to save the settings. The switches go into the pending state, that is, blue color. These settings enable the following:

- Enable PIM on L3VNI SVI.
- Route-Target Import and Export for MVPN AFI.
- RP and other multicast configuration for the VRF.
- Loopback interface using the above RP address and RP loopback id for the distributed RP.

Step 3 Enable TRM for the network.

Navigate to **Control > Networks**. Edit the selected network and navigate to the **Advanced** tab. Edit the following TRM setting:

TRM enable – Select the check box to enable TRM.

Edit Network
✕

* Network ID

* Network Name

* VRF Name

Layer 2 Only

* Network Template

* Network Extension Template

VLAN ID Propose VLAN ?

▼ Network Profile

Generate Multicast IP ⓘ Please click only to generate a New Multicast Group Address and override the default value!

General

Advanced

DHCPv4 Server 3 ⓘ DHCP Relay IP

DHCPv4 Server3 VRF ⓘ

Loopback ID for DHCP Relay interface (Min:0, Max:1023) ⓘ

Routing Tag ⓘ 0-4294967295

TRM Enable ⓘ Enable Tenant Routed Multicast

L2 VNI Route-Target ⓘ

Save
Cancel

Click **Save** to save the settings. The switches go into the pending state, that is, the blue color. The TRM settings enable the following:

- Enable PIM on the L2VNI SVI.
- Create a PIM policy **none** to avoid PIM neighborship with PIM Routers within a VLAN. The **none** keyword is a configured route map to deny any ipv4 addresses to avoid establishing PIM neighborship policy using anycast IP.

Configuring TRM for Multi-Site Using Cisco DCNM

This section assumes that a Multi-Site Domain (MSD) has already been deployed by Cisco DCNM and TRM needs to be enabled.

Procedure

Step 1 Enable TRM on the BGWs.

Navigate to **Control > VRFs**. Make sure that the right DC Fabric is selected under the **Scope** and edit the VRF. Navigate to the **Advanced** tab. Edit the TRM settings. Repeat this process for every DC Fabric and its VRFs.

TRM Enable – Select the check box to enable TRM. If you enable TRM, then the RP address and the underlay multicast address must be entered.

Is RP External – Enable this check box if the RP is external to the fabric. If this field is unchecked, RP is distributed in every VTEP.

Note If the RP is external, then select the appropriate option. If the RP is external, then RP loopback ID is greyed out.

RP Address – Specifies the IP address of the RP.

RP Loopback ID – Specifies the loopback ID of the RP, if **Is RP External** is not enabled.

Underlay Multicast Address – Specifies the multicast address associated with the VRF. The multicast address is used for transporting multicast traffic in the fabric underlay.

Note The multicast address in the **Default MDT Address for TRM VRFs** field in the fabric settings screen is auto-populated in this field. User can override this field if a different multicast group address should be used for this VRF.

Overlay Multicast Groups – Specifies the multicast group subnet for the specified RP. The value is the group range in “ip pim rp-address” command. If the field is empty, 224.0.0.0/24 is used as default.

Enable TRM BGW MSite - Select the check box to enable TRM on Border Gateway Multi-Site.

Edit VRF



▼ VRF Information

* VRF ID

* VRF Name

* VRF Template

* VRF Extension Template

VLAN ID

▼ VRF Profile

General	Advanced
	<p>Overlay Mcast Groups <input type="text" value="224.0.0.0/4"/> <small>224.0.0.0/4 to 239.255.255.255/4</small></p> <p>Enable IPv6 link-loc... <input checked="" type="checkbox"/> <small>Enables IPv6 link-local Option under VRF SVI</small></p> <p>Enable TRM BGW MSite <input checked="" type="checkbox"/> <small>Enable TRM on Border Gateway Multisite</small></p> <p>Advertise Host Routes <input type="checkbox"/> <small>Flag to Control Advertisement of /32 and /128 Routes to Edge Routers</small></p> <p>Advertise Default Route <input checked="" type="checkbox"/> <small>Flag to Control Advertisement of Default Route Internally</small></p> <p>Config Static 0/0 Route <input checked="" type="checkbox"/> <small>Flag to Control Static Default Route Configuration</small></p> <p>BGP Neighbor Password <input type="text"/></p> <p>BGP Password Key Encryption Type <input type="text" value="3"/> <small>VRF Lite BGP Key Encryption Type: 3 - 3DES</small></p>

Click on **Save** to save the settings. The switches go into the pending state, that is, blue color. These settings enable the following:

- Enable feature ngmvpn: Enables the Next-Generation Multicast VPN (ngMVPN) control plane for BGP peering.
- Enables PIM on L3VNI SVI.
- Configures L3VNI Multicast Address.
- Route-Target Import and Export for MVPN AFI.
- RP and other multicast configuration for the VRF.
- Loopback interface for the distributed RP.
- Enable Multi-Site BUM ingress replication method for extending the Layer 2 VNI

Step 2 Establish MVPN AFI between the BGWs.

Navigate to **Control > Fabrics**. Select the MSD fabric. Click **Tabular view** and click **Links**. Filter it by the policy - **Overlays**.

The screenshot shows the Cisco Data Center Network Manager interface. At the top, it says "Data Center Network Manager" with a "(Non-Production)" status and "SCOPE: MSD". Below that, it says "Fabric Builder: MSD" with a "Save & Deploy" button. The main area has tabs for "Switches", "Links", and "Operational View". Below the tabs is a table of links:

	Fabric Name	Name	Policy	Info	Admin State	Oper State	MACsec Status
1	Fabric-2-<->Fabric-3	FAB2-BGW1-loopback0—N93180FX-BGW2-S3-loopback0	ext_evpn_multisite_overlay_setup	NA	--	--	NA
2	Fabric-2-<->Fabric-3	FAB2-BGW1-loopback0—N93180FX-BGW1-S3-loopback0	ext_evpn_multisite_overlay_setup	NA	--	--	NA

Select and edit each overlay peering to enable TRM by checking the **Enable TRM** check box.

Link Management - Edit Link

The screenshot shows the "Link Management - Edit Link" configuration page. It has a close button (X) in the top right corner. The configuration is organized into sections:

- Link Type:** Inter-Fabric
- Link Sub-Type:** MULTISITE_OVERLAY
- Link Template:** ext_evpn_multisite_overlay_se
- Source Fabric:** Fabric-2
- Destination Fabric:** Fabric-3
- Source Device:** FAB2-BGW1
- Source Interface:** loopback0
- Destination Device:** N93180FX-BGW1-S3
- Destination Interface:** loopback0

Below this is the "Link Profile" section, which is expanded to show the "General" tab:

- Source BGP ASN:** 65002 (Info: BGP Autonomous System Number in Source Fabric)
- Source IP Address:** 20.2.0.1 (Info: Source IPv4 Address for BGP EVPN Peering)
- Destination IP Address:** 30.2.0.1 (Info: Destination IPv4 Address for BGP EVPN Peering)
- Destination BGP ASN:** 65003 (Info: BGP Autonomous System Number in Destination Fabric)
- Enable TRM:** (Info: Enable Tenant Routed Multicast)

Save

Click **Save** to save the settings. The switches go into the pending state, that is, the blue color. The TRM settings enable the MVPN peering's between the BGWs, or BGWs and Route Server.

SSH Key RSA Handling

Bootstrap scenario

If the switch has the **ssh key rsa** command with the key-length variable value other than 1024 in the running configuration, the **ssh key rsa key-length force** command needs to be added to the bootstrap freeform configuration with the required value (any value other than 1024) during bootstrap.

Greenfield and Brownfield scenarios

Use the **ssh key rsa key-length force** command to change the key-length variable to a value other than 1024.

However, on Cisco Nexus 9000 Releases 9.3(1) and 9.3(2), the **ssh key rsa key-length force** command fails while the device is booting up during the ASCII replay process. For more information, refer [CSCvs40704](#).

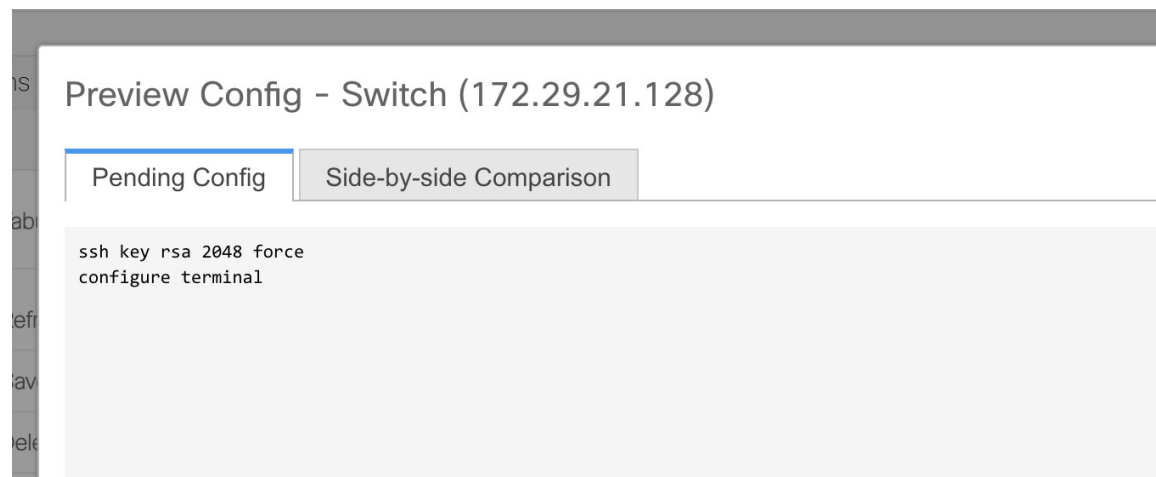
The configurations are considered to be in-sync when both the intent and switch running configurations have the same command. For example, the status is considered to be in-sync when the **ssh key rsa 2048** command is present in both in the intent and the running configuration. However, consider a scenario in which the **ssh key rsa 2040** command was pushed to the switch as an Out-Of-Band change. While the intent has a key-length value of 2048, the device has a key-length value of 2040. In such instances, the switch will be marked as out-of-sync.

The diff shown in the Pending Config tab (in both Strict Config-Compliance and non-Strict Config-Compliance mode) cannot be deployed onto the switch from DCNM as the **feature ssh** command has to be used to disable the SSH feature before making any change to the **ssh key rsa** command. This would lead to a dropped connection to DCNM. In such a scenario, the diff can be resolved by modifying the intent such that there is no diff.

With Strict Config-Compliance mode:



- Delete the Policy Template Instance (PTI) that has the **ssh key rsa 2048 force** command by clicking **View/Edit Policies** in the **Tabular View** of the **Fabric Builder** window.
- Create a new PTI with the **ssh key rsa 2040 force** command by clicking **View/Edit Policies**.

Without Strict Config-Compliance mode:

- Delete the PTI with the **ssh key rsa 2048 force** command in the intent by clicking **View/Edit Policies** in the **Tabular View** of the **Fabric Builder** window.

- Create a switch_freeform PTI with the **ssh key rsa 2040 force** command in the intent to match the Out-Of-Band change from the device.

Switch Operations

To view various options, right-click on switch:

Set Role: Assign a role to the switch. You can assign any one of the following roles to a switch:

- Spine
- Leaf (Default role)
- Border
- Border Spine
- Border Gateway
- Access
- Aggregation
- Edge Router
- Core Router
- Super Spine
- Border Super Spine
- Border Gateway Spine
- ToR

Alternatively, you can also navigate to the **Tabular view** from the **Actions** pane. Choose one or more devices of the same device type and click **Set Role** to set roles for devices. The device types are:

- NX-OS
- IOS XE
- IOS XR
- Other



Note Ensure that you have moved switches from maintenance mode to active mode or operational mode before setting roles.

You can change the switch role only before executing **Save & Deploy**.


You can assign one of the following roles for non-Nexus devices:

- Spine
- Leaf
- Access (This role is available only for Cisco ASR 1000 Series routers and Cisco Catalyst 9000 Series switches).
- Edge Router (Use this role for VRF-Lite).
- Core Router
- Super Spine
- Preview Config
- ToR (This role is available only for Cisco Catalyst 9000 series switches).

From DCNM 11.1(1) release, you can shift the switch role from existing to required role if there are no overlays on the switches. Click **Save and Deploy** to generate the updated configuration. The following shifts are allowed for the switch role:

- Leaf to Border
- Border to Leaf
- Leaf to Border Gateway
- Border Gateway to Leaf
- Border to Border Gateway
- Border Gateway to Border
- Spine to Border Spine
- Border Spine to Spine
- Spine to Border Gateway Spine
- Border Gateway Spine to Spine


- Border Spine to Border Gateway Spine
- Border Gateway Spine to Border Spine



You cannot change the switch role from any Leaf role to any Spine role and from any Spine role to any Leaf role.

In case the switch role is not changed according to the allowed switch role changes mentioned above for easy fabrics, the following error is displayed after you click **Save and Deploy**:

```
Switch[<serial-number>]: Role change from <switch-role> to <switch-role> is not permitted.
```




You can then change the switch role to the role that was set earlier, or set a new role, and configure the fabric.

If you have not created any policy template instances before clicking **Save and Deploy**, and there are no overlays, you can change the role of a switch to any other required role.

If you change the switch role of a vPC switch that is part of a vPC pair, the following error appears when you click **Save and Deploy**:

```
Switches role should be the same for VPC pairing. peer1 <serial-number>: [<switch-role>],  
peer2 <serial-number>: [<switch-role>]
```



To prevent this scenario, change the switch roles of both the switches in the vPC pair to the same role.

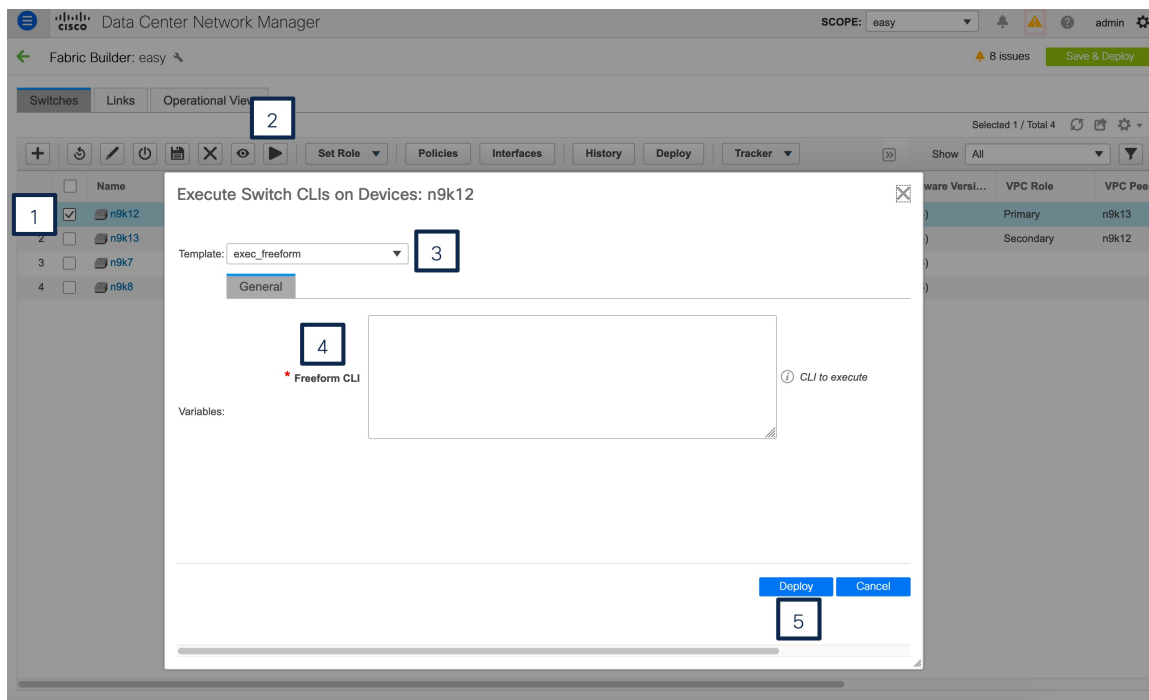
Running EXEC Mode Commands in DCNM

When you first log in, the Cisco NX-OS software places you in the EXEC mode. The commands available in the EXEC mode include the show commands that display the device status and configuration information, the clear commands, and other commands that perform actions that you do not save in the device configuration.

The following procedure shows how to run EXEC commands in DCNM:

Procedure

- Step 1** From DCNM, navigate to **Control > Fabrics > Fabric Builder**.
- Step 2** Click a fabric and then click **Tabular view** in the **Actions** menu.
- Step 3** Select one or more switches and click the **Play** button (Execute Commands).
- Step 4** From the **Template** drop-down list, select **exec_freeform**.
- Step 5** Enter the commands in the **Freeform CLI** field.



- Step 6** Click **Deploy** to run the EXEC commands.
- Step 7** In the **CLI Execution Status** window, you can check the status of the deployment. Click **Detailed Status** under the **Command** column to view details.

- Step 8** In the **Command Execution Details** window, click the info under the **CLI Response** column to view the output or response.
-

Fabric Multi Switch Operations

Click **Tabular view** from the **Actions** pane in the fabric topology window. The tabular view has the following tabs:

- [Tabular View - Switches](#)
- [Tabular View - Links](#)
- [Tabular View - Operational View](#)

Tabular View - Switches

You can manage switch operations in this tab. Each row represents a switch in the fabric, and displays switch details, including its serial number.

Some of the actions that you can perform from this tab are also available when you right-click a switch in the fabric topology window. However, the **Switches** tab enables you to provision configurations on multiple switches, like deploying policies, simultaneously.

The **Switches** tab has following information of every switch you discover in the fabric:

- Name: Specifies the switch name.
- IP Address: Specifies the IP address of the switch.
- Role: Specifies the role of the switch.
- Serial Number: Specifies the serial number of the switch.
- Fabric Name: Specifies the name of the fabric, where the switch is discovered.
- Fabric Status: Specifies the status of the fabric, where the switch is discovered.
- Discover Status: Specifies the discovery status of the switch.
- Model: Specifies the switch model.
- Software Version: Specifies the software version of the switch.
- ThousandEyes Status: Specifies the status of the ThousandEyes Enterprise Agent.
- Last Updated: Specifies when the switch was last updated.
- Mode: Specifies the current mode of the switch.
- VPC Role: Specifies the vPC role of the switch.
- VPC Peer: Specifies the vPC peer of the switch.

The **Switches** tab has the following icons and buttons:

- Add switches: Click this icon to discover existing or new switches to the fabric. The **Inventory Management** dialog box appears.

This option is also available in the fabric topology window. Click **Add switches** in the **Actions** pane.

Refer the following sections for more information:

- [Adding Switches to a Fabric](#): Provides information on adding switches to easy fabrics.
- [Discovering New Switches](#): Provide information on adding Cisco Nexus switches to external fabrics.
- [Adding non-Nexus Devices to External Fabrics](#): Provide information on adding non-Nexus switches to external fabrics.
- Rediscover switch: Initiate the switch discovery process by DCNM afresh.
- Update discovery credentials: Update device credentials such as authentication protocol, username and password.
- Saving config and Reload: Save the configurations and reload the switch.



Note This option is grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.

- Copy running to startup config: From Cisco DCNM, Release 11.4(1), you can perform an on-demand copy running-configuration to startup-configuration operation for one or more switches.



Note This option will be grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.

- Remove switches: Remove the switch from the fabric.



Note This option will be grayed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.

- Preview: You can preview the pending configurations and the side-by-side comparison of running configurations and expected configurations.
- Policies: Add, update and delete a policy. The policies are template instances of templates in the template library. After creating a policy, you should deploy it on the switches using the **Deploy** option available in the **Policies** window. You can select more than one policy and view them.



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

- **ThousandEyes Agent:** You can start, stop, install, or uninstall ThousandEyes Enterprise Agent on the switch. You can choose single or multiple switches and select required operation from **ThousandEyes Agent** drop-down list.



Note When you choose multiple switches to perform ThousandEyes Enterprise Agent action, ensure that the status of selected switches are same.

- **Interfaces:** Deploy configurations on the switch interfaces.
- **History:** View the deployment history and the policy change history using this button. Choose one or more switches and click **History**.

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

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

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

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



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

- **Deploy:** Deploy switch configurations. From Cisco DCNM Release 11.3(1), you can deploy configurations for multiple devices using the **Deploy** button.



Note

- This option grays out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.
- In an MSD fabric, you can deploy configurations only on the Border Gateway, Border Gateway Spine, Border Gateway Super-Spine, or External Fabric switches.

- **Set Role:** Choose one or more devices of the same device type and click **Set Role** to set roles for devices. The device types are:

- NX-OS
- IOS XE
- IOS XR
- Other

Ensure that you have moved switches from maintenance mode to active mode or operational mode before setting roles. See the [Switch Operations](#) section for more information on setting roles.

- **vPC Pairing:** Choose a switch and click **vPC Pairing** to create, edit, or unpair a vPC pair. However, you can use this option only when you choose a Cisco Nexus switch. Refer the following sections for more information:
 - [Creating a vPC Setup](#): Provides information on how to create a vPC pair in external fabrics.
 - [vPC Fabric Peering](#): Provides information on how to create a vPC pair in easy fabrics.

Tabular View - Links

You can add links between border switches of different fabrics (inter-fabric links) or between switches in the same fabric (intra-fabric links). You can only create an inter-fabric connection (IFC) for a switch that is managed by DCNM.

There are scenarios where you might want to define links between switches before connecting them physically. The links could be inter-fabric or intra-fabric links. Doing so, you can express and represent your intent to add links. The links with intent are displayed in a different colour till they are actually converted to functional links. Once you physically connect the links, they are displayed as connected.

Management links might show up in the fabric topology as red colored links. To remove such links, right-click the link and click **Delete Link**.

From Cisco DCNM Release 11.1(1), the Border Spine and Border Gateway Spine roles are added to switch roles for border switches.

You can create links between existing and pre-provisioned devices as well by selecting the pre-provisioned device as the destination device.

Creating Intra-Fabric Links

1. Click Control > Fabric Builder to go to the Fabric Builder screen.
2. Click within the rectangular box that represents the fabric. The fabric topology screen comes up.
3. Click Tabular view in the Actions panel that is displayed at the left part of the screen.



A screen with the tabs Switches and Links appears. They list the fabric switches and links in a table.

	<input type="checkbox"/>	Name	IP Address	Role	Serial Number	Fabric Name	Fabric Status	Discovery Status	Model
1	<input type="checkbox"/>	N9K-15-BGW	111.0.0.95	border ...	FDO20401LB4	Easy60000	In-Sync	✔ ok	N9K-C93180YC-EX
2	<input type="checkbox"/>	N9K-16-Leaf	111.0.0.96	leaf	SAL18432P6G	Easy60000	In-Sync	✔ ok	N9K-C9396PX
3	<input type="checkbox"/>	N9K-17-Spine	111.0.0.97	spine	FDO20401LEJ	Easy60000	In-Sync	✔ ok	N9K-C93180YC-EX

4. Click the Links tab. You can see a list of links.
The list is empty when you are yet to create a link.

	Scope	Name	Policy	Admin State	Oper State
1	<input type="checkbox"/> Easy60000	N9K-15-BGW-Ethernet1/3---n7k1-N7K-1-BorderLeaf2-Ethe...			
2	<input type="checkbox"/> Easy60000	N9K-16-Leaf-Ethernet2/1---n7k1-Ethernet7/8			
3	<input type="checkbox"/> External65000<->Easy60000	BorderLeaf1-Loopback0---N9K-15-BGW-loopback0	multisite_overlay_setup_rs_test		
4	<input type="checkbox"/> Easy7200<->Easy60000	N9K-4-BGW-Ethernet1/2---N9K-15-BGW-Ethernet1/8	ext_multisite_underlay_setup_test		
5	<input type="checkbox"/> Easy7200<->Easy60000	N9K-3-BGW-Ethernet1/2---N9K-15-BGW-Ethernet1/7	ext_multisite_underlay_setup_test		
6	<input type="checkbox"/> Easy60000	N9K-15-BGW-Ethernet1/5---N9K-17-Spine-Ethernet1/1	int_intra_fabric_num_link_11_1		
7	<input type="checkbox"/> Easy7200<->Easy60000	N9K-1-Spine-Ethernet1/1---N9K-16-Leaf-Ethernet1/3			
8	<input type="checkbox"/> Easy60000	N9K-17-Spine-Ethernet1/2---N9K-16-Leaf-Ethernet1/5	int_intra_fabric_num_link_11_1		
9	<input type="checkbox"/> Easy7200<->Easy60000	N9K-2-Leaf-Ethernet1/2---N9K-16-Leaf-Ethernet1/4			
10	<input type="checkbox"/> Easy60000	N9K-15-BGW-Ethernet1/2---N9K-16-Leaf-Ethernet1/2			
11	<input type="checkbox"/> Easy60000<->Easy7200	N9K-15-BGW-Ethernet1/4---N9K-1-Spine-Ethernet1/2			
12	<input type="checkbox"/> Easy60000<->Easy7200	N9K-15-BGW-Ethernet1/50---N9K-18-BGW-Ethernet1/7			
13	<input type="checkbox"/> Easy60000<->External65000	N9K-15-BGW-Ethernet1/49---n7k1-BorderLeaf1-Ethernet7/6			

5. Click the Add (+) button at the top left part of the screen to add a link.

The Add Link screen comes up. By default, the Intra-Fabric option is chosen as the link type.

Link Management - Add Link

* Link Type

 * Link Sub-Type

 * Link Template

 * Source Fabric

 * Destination Fabric

 * Source Device

 * Source Interface

 * Destination Device

 * Destination Interface

▼ Link Profile

 General

 * FABRIC_NAME ? FABRIC NAME

 * Source IP ? IP address of the source interface

 * Destination IP ? IP address of the destination interface

 Interface Admin State ? Admin state of the interface

 * MTU ? MTU for the interface

 Save

The fields are:

Link Type – Choose Intra-Fabric to create a link between two switches in a fabric.

Link Sub-Type – This field populates Fabric indicating that this is a link within the fabric.

Link Template: You can choose any of the following link templates.

- `int_intra_fabric_num_link_11_1`: If the link is between two ethernet interfaces assigned with IP addresses, choose `int_intra_fabric_num_link_11_1`.
- `int_intra_fabric_unnum_link_11_1`: If the link is between two IP unnumbered interfaces, choose `int_intra_fabric_unnum_link_11_1`.
- `int_intra_vpc_peer_keep_alive_link_11_1`: If the link is a vPC peer keep-alive link, choose `int_intra_vpc_peer_keep_alive_link_11_1`.
- `int_pre_provision_intra_fabric_link`: If the link is between two pre-provisioned devices, choose `int_pre_provision_intra_fabric_link`. After you click **Save & Deploy**, an IP address is picked from the underlay subnet IP pool.

Correspondingly, the Link Profile section fields is updated.

Source Fabric – The fabric name populates this field since the source fabric is known.

Destination Fabric – Choose the destination fabric. For an intra-fabric link, source and destination fabrics are the same.

Source Device and Source Interface – Choose the source device and interface.

Destination Device and Destination Interface – Choose the destination device and interface.



Note Select the pre-provisioned device as the destination device if you are creating a link between an existing device and a pre-provisioned device.

General tab in the Link Profile section

Interface VRF – Name of a non-default VRF for this interface.

Source IP and Destination IP – Specify the source and destination IP addresses of the source and destination interfaces, respectively.



Note The Source IP and Destination IP fields do not appear if you choose `int_pre_provision_intra_fabric_link` template.

Interface Admin State – Check or uncheck the check box to enable or disable the admin state of the interface.

MTU – Specify the maximum transmission unit (MTU) through the two interfaces.

Link Management - Add Link



* Link Type	Intra-Fabric
* Link Sub-Type	Fabric
* Link Template	int_intra_fabric_num_link_11_1
* Source Fabric	Easy60000
* Destination Fabric	Easy60000
* Source Device	N9K-16-BL
* Source Interface	Ethernet1/40
* Destination Device	N9K-17-Spine
* Destination Interface	Ethernet1/40

▼ Link Profile

- General
- Advanced

* FABRIC_NAME	Easy60000	? FABRIC NAME
* Source IP	10.1.1.1	? IP address of the source interface
* Destination IP	10.1.1.3	? IP address of the destination interface
Interface Admin State	<input checked="" type="checkbox"/>	? Admin state of the interface
* MTU	9216	? MTU for the interface

Save

Advanced tab.

▼ Link Profile

- General
- Advanced

Source Interface Desc...	<input type="text"/>	? Add description to the source interface (Max Size 254)
Destination Interface ...	<input type="text"/>	? Add description to the destination interface (Max Size 254)
Disable BFD Echo on ...	<input type="checkbox"/>	? Disable BFD Echo on Source Interface
Disable BFD Echo on ...	<input type="checkbox"/>	? Disable BFD Echo on Destination Interface
Source Interface Free...	<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.
Destination Interface ...	<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.

Save

Source Interface Description and Destination Interface Description – Describe the links for later use. For example, if the link is between a leaf switch and a route reflector device, you can enter the information in these fields (Link from leaf switch to RR 1 and Link from RR 1 to leaf switch). This description will be converted into a config, but will not be pushed into the switch. After **Save & Deploy**, it will reflect in the running configuration.

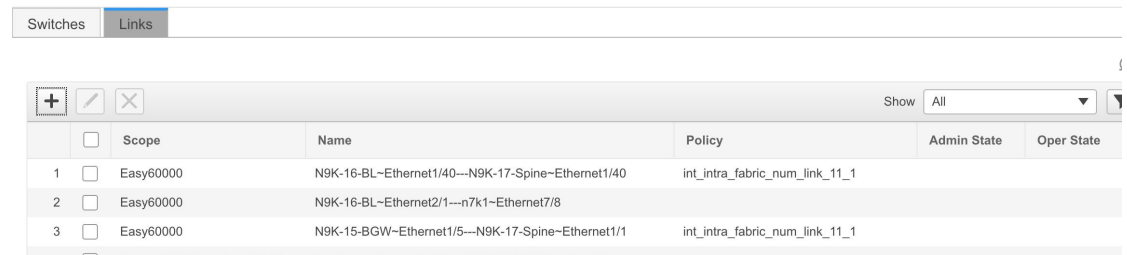
Disable BFD Echo on Source Interface and **Disable BFD Echo on Destination Interface** – Select the check box to disable BFD echo packets on source and destination interface.

Note that the BFD echo fields are applicable only when you have enabled BFD in the fabric settings.

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

6. Click Save at the bottom right part of the screen.

The new link appears in the Links tab.



The screenshot shows the 'Links' tab in the 'Switches' section. A table lists three links with columns for Scope, Name, Policy, Admin State, and Oper State. The first three rows are visible:

	<input type="checkbox"/>	Scope	Name	Policy	Admin State	Oper State
1	<input type="checkbox"/>	Easy60000	N9K-16-BL-Ethernet1/40---N9K-17-Spine-Ethernet1/40	int_intra_fabric_num_link_11_1		
2	<input type="checkbox"/>	Easy60000	N9K-16-BL-Ethernet2/1---n7k1-Ethernet7/8			
3	<input type="checkbox"/>	Easy60000	N9K-15-BGW-Ethernet1/5---N9K-17-Spine-Ethernet1/1	int_intra_fabric_num_link_11_1		

7. Click **Save & Deploy** to deploy the link configurations on the switches.

The Config Deployment screen comes up. It displays the configuration status on the switches. You can also view the pending configurations by clicking the respective link in the Preview Config column. When you click a link in the Preview Config column, the Config Preview window comes up. It lists the pending configurations on the switch. The Side-by-side Comparison tab displays the running configuration and expected configuration side-by-side.

8. Close the preview screen and click Deploy Config. The pending configurations are deployed.
9. After ensuring that the progress is 100% in all the rows, click Close at the bottom part of the screen. The Links screen comes up again.

Click <- at the top left part of the screen to go to the fabric topology. In the fabric topology, you can see that the link between the two devices is displayed.

Creating Inter-Fabric Links

1. Click the Links tab in the Switches | Links page. The list of previously created links is displayed. The list contains intra-fabric links (between switches in a fabric), and inter-fabric links (between BGWs or border leaf/spine switches of different fabrics).

	<input type="checkbox"/>	Scope	Name	Policy	Admin State	Oper State
1	<input type="checkbox"/>	Easy60000	N9K-16-Leaf-Ethernet2/1---n7k1-Ethernet7/8			
2	<input type="checkbox"/>	Easy60000	N9K-15-bgw-Ethernet1/49---n7k1-BorderLeaf1-Ethernet7/6			
3	<input type="checkbox"/>	Easy60000	N9K-15-bgw-Ethernet1/3---n7k1-N7K-1-BorderLeaf2-Ether...			
4	<input type="checkbox"/>	Easy60000	N9K-17-Spine-Ethernet1/2---N9K-16-Leaf-Ethernet1/5	int_intra_fabric_num_link_11_1		
5	<input type="checkbox"/>	Easy60000	N9K-15-bgw-Ethernet1/5---N9K-17-Spine-Ethernet1/1	int_intra_fabric_num_link_11_1		
6	<input type="checkbox"/>	New7200<->Easy60000	n9k-3-bgw-Ethernet1/2---N9K-15-bgw-Ethernet1/7			
7	<input type="checkbox"/>	Easy60000<->New7200	N9K-15-bgw-Ethernet1/50---n9k-18-bgw-Ethernet1/7			
8	<input type="checkbox"/>	New7200<->Easy60000	n9k-4-bgw-Ethernet1/2---N9K-15-bgw-Ethernet1/8			
9	<input type="checkbox"/>	Easy60000	N9K-15-bgw-Ethernet1/2---N9K-16-Leaf-Ethernet1/2			
10	<input type="checkbox"/>	New7200<->Easy60000	n9k-2-leaf-Ethernet1/2---N9K-16-Leaf-Ethernet1/4			
11	<input type="checkbox"/>	New7200<->Easy60000	n9k-1-spine-Ethernet1/1---N9K-16-Leaf-Ethernet1/3			
12	<input type="checkbox"/>	Easy60000<->New7200	N9K-15-bgw-Ethernet1/4---n9k-1-spine-Ethernet1/2			

- Click the Add (+) button at the top left part of the screen to add a link. The Add Link screen comes up. By default, the Intra-Fabric option is chosen as the link type.

Link Management - Add Link

Link Management - Add Link

* Link Type: Intra-Fabric

* Link Sub-Type: Fabric

* Link Template: int_intra_fabric_num_link_11_1

* Source Fabric: Easy60000

* Destination Fabric: [Empty]

* Source Device: [Empty]

* Source Interface: [Empty]

* Destination Device: [Empty]

* Destination Interface: [Empty]

▼ Link Profile

General

* FABRIC_NAME: [Empty] ? FABRIC NAME

* Source IP: [Empty] ? IP address of the source interface

* Destination IP: [Empty] ? IP address of the destination interface

Interface Admin State: ? Admin state of the interface

* MTU: 9216 ? MTU for the interface

Save

- From the Link Type drop-down box, choose Inter-Fabric since you are creating an IFC. The screen changes correspondingly.

Link Management - Add Link



* Link Type	Inter-Fabric	▼
* Link Sub-Type	VRF_LITE	▼
* Link Template	ext_fabric_setup_test	▼
* Source Fabric	Easy60000	▼
* Destination Fabric		▼
* Source Device		▼
* Source Interface		▼
* Destination Device		▼
* Destination Interface		▼

▼ Link Profile

General

* Local BGP AS # ? Local BGP Autonomous System Number

* IP_MASK

* NEIGHBOR_IP

* NEIGHBOR_ASN

The fields for inter-fabric link creation are explained:

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

Link Sub-Type – This field populates the IFC type. Choose **VRF_LITE**, **MULTISITE_UNDERLAY**, or **MULTISITE_OVERLAY** from the drop-down list.

The Multi-Site options are explained in the Multi-Site use case.

For information about VXLAN MPLS interconnection, see the *Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - MPLS SR and LDP Handoff* chapter.

For information about routed fabric interconnection, see the *Creating Inter-Fabric Links Between a Routed Fabric and an External Fabric* section in the *Configuring a Fabric with eBGP Underlay* chapter.

Link Template: The link template is populated.

The templates are autopopulated with corresponding pre-packaged default templates that are based on your selection.



Note You can add, edit, or delete user-defined templates. See *Template Library* section in the Control chapter for more details.

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

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

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

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

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

General tab in the Link Profile section.

Local BGP AS# - In this field, the AS number of the source fabric is autopopulated.

IP_MASK—Fill up this field with the IP address of the source interface that connects to the destination device.

NEIGHBOR_IP—Fill up this field with the IP address of the destination interface.

NEIGHBOR_ASN—In this field, the AS number of the destination device is autopopulated.

After filling up the Add Link screen, it looks like this:

Link Management - Add Link
✕

* Link Type

* Link Sub-Type

* Link Template

* Source Fabric

* Destination Fabric

* Source Device

* Source Interface

* Destination Device

* Destination Interface

▼ Link Profile

General

* Local BGP AS # ? Local BGP Autonomous System Nu

* IP_MASK ?

* NEIGHBOR_IP ?

* NEIGHBOR_ASN ?

4. Click Save at the bottom right part of the screen.

The Switches|Links screen comes up again. You can see that the IFC is created and displayed in the list of links.

	<input type="checkbox"/>	Scope	Name	Policy
1	<input type="checkbox"/>	Easy60000	N9K-16-Leaf~Ethernet2/1---n7k1~Ethernet7/8	
2	<input type="checkbox"/>	Easy60000	N9K-15-bgw~Ethernet1/49---n7k1~BorderLeaf1~Ethernet7/6	
3	<input type="checkbox"/>	Easy60000<->New7200	N9K-15-bgw~Ethernet1/9---n9k-18-bgw~Ethernet1/9	ext_fabric_setup_test

5. Click on Save & Deploy to deploy the link configurations on the switches.

The Config Deployment screen comes up. It displays the configuration status on the switches. You can also view the pending configurations by clicking the respective link in the Preview Config column. When you click a link in the Preview Config column, the Config Preview window comes up. It lists the pending configurations on the switch. The Side-by-side Comparison tab displays the running configuration and expected configuration side-by-side.

6. Close the preview screen and click Deploy Config. The pending configurations are deployed.
7. After ensuring that the progress is 100% in all the rows, click Close at the bottom part of the screen. The Links screen comes up again.
8. Click <- at the top left part of the screen to go to the fabric topology. In the fabric topology, you can see that the link between the two devices is displayed.

If the two fabrics are member fabric of an MSD, then you can see the link in the MSD topology too.

When you enable the VRF Lite function using the ToExternalOnly method or Multisite function via MSD fabric, IFCs are automatically created between the (VXLAN fabric) border/BGW device and connected (external fabric) edge router/core device. When you remove the ER/core/border/BGW device, the corresponding IFCs (link PTIs) to/from that switch are deleted on DCNM. Subsequently, DCNM removes the corresponding IFC configurations, if any, from the remaining devices on the next Save & Deploy operation. Also, if you want to remove a device that has an IFCs and overlay extensions over those IFCs, you should undeploy all overlay extensions corresponding to those IFCs for switch delete to be possible.

To undeploy VRF extensions, click Control > Networks & VRFs, select the VXLAN fabric and the extended VRFs, and undeploy the VRFs in the VRF deployment screen.

To delete the IFCs, click Control > Fabric Builder, go to the fabric topology screen, click Tabular view, and delete the IFCs from the Links tab.

Ensure that the fabric switch names are unique. If you deploy VRF extensions on switches with the same name, it leads to erroneous configuration.

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

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

Exporting Links

1. Choose Control > Fabric Builder, and select a fabric.

The fabric topology window appears.

2. Click **Tabular view** in the **Actions** panel.

A window with the **Switches** and **Links** tabs appears.

3. Click the **Links** tab.

You can see a list of links. The list is empty when you are yet to create a link.

4. Click the **Export Links** icon to export the links in a CSV file.

The following details of links are exported: link template, source fabric, destination fabric, source device, destination device, source switch name, destination switch name, source interface, destination interface, and nvPairs. The nvPairs field consists JSON object.

Importing Links

You can import a CSV file containing details of links to add new links to the fabric. The CSV file should have the following details of links: link template, source fabric, destination fabric, source device, destination device, source switch name, destination switch name, source interface, destination interface, and nvPairs.



Note

- You cannot update existing links.
- The **Import Links** icon is disabled for external fabric.

1. Choose **Control > Fabric Builder**, and select a fabric.

The fabric topology window appears.

2. Click **Tabular view** in the **Actions** panel.

A window with the **Switches** and **Links** tabs appears.

3. Click the **Links** tab.

You can see a list of links. The list is empty when you are yet to create a link.

4. Click the **Import Links** icon.

The file server directory opens.

5. Browse the directory and select the CSV file that you want to import.

6. Click **Open**.

A confirmation screen appears.

7. Click **Yes** to import the selected file.

Viewing Details of Fabric Links

You can view information about a fabric link, like IP subnet between links to deploy underlay, MTU, speed mismatch, and so on, in the topology view of a fabric builder. To view the details of a link from the Cisco DCNM Web client, perform the following steps:

Procedure

Step 1 Choose **Control > Fabrics > Fabric Builder** and select a fabric.

The topology view of the fabric appears.

Step 2 Double-click any of the links.

The details window appears. You can view the devices that are connected using this link, summary, and the data traffic.

Step 3 Click **Show more details**.

A comparison table of the two devices connected by the link appears. It includes the following parameters of the devices: device name, name, admin status, operation status, reason, policies, overlay network, status, PC, vPC ID, speed, MTU, mode, VLANs, IP or prefix, VRF, neighbor, and description.

- Note**
- You can view the traffic details of a fabric link by clicking the device name with hyperlink. Alternatively, you can view these traffic details in the details window. See *Viewing the Traffic Details of the Fabric Links* section for more information.
 - You can view the expected configuration of a fabric link by clicking the policy with the hyperlink.

Step 4 Click the **Back** icon to go back to the details window.

Note You can click the **Close** icon to exit the details window.

Viewing the Traffic Details of Fabric Links

In the details window of a fabric link, you can choose how you want to view the traffic details. You can view the traffic details based on the time duration, format, and export this information.

You can view the data traffic of a link for the following durations from the duration drop-down list:

- 24 Hours
- Week
- Month
- Year

Show: Click **Show**, and choose **Chart**, **Table**, or **Chart and Table** from the drop-down list to see how you want to view the traffic details. Enlarge your browser window to view the details in **Chart and Table** format.

If you choose **Chart**, hover over the traffic chart to view the Rx and Tx values, along the Y axis, for the corresponding time, along X axis. You can change the time duration values of the X axis by moving the sliders in the time range selector. You can choose the Y-axis values by checking or unchecking the Rx and Tx check boxes.



Note If you select **Week**, **Month**, or **Year** as the time duration, you can also view the Peak Rx and Peak Tx values along the Y axis.

Select **Table** to view the traffic information in tabular format.

Chart Type and Chart Options: Choose **Area Chart** or **Line Chart** from the **Chart Type** drop-down list.

You can choose the following chart options:

- **Show Fill Patterns**
- **Show Datamarkers**
- **Y Axis Log Scale**

Actions: Export or print the traffic information by choosing the appropriate options from the **Actions** drop-down list.

Symmetric Automatic VRF Lite

- Check the **Auto Deploy Flag** check box in the **Link Management** dialog box. Checking this check box enables VRF lite deployment on both ends of the link for managed devices.
- When you extend the VRF lite in a back-to-back scenario, the VRF should already be present in the peer fabric and the VRF name should be the same. An error message appears if the VRF is not present in the peer fabric and if you try to extend the VRF lite.
- When you extend the VRF lite between an easy fabric and an external fabric, the VRF name can be the same as that of the source fabric, default, or another VRF name. However, the child PTIs for the subinterface and the VRF creation or peering on the external fabric has the source. Hence, you cannot edit or delete the policies from the **View/Edit policies** window.
- If you perform a DCNM upgrade and notice that the policies are not attached to the IFC, edit the policies and VRF to attach them again.
- Besides the IPv6 address, enter the IP mask, IPv4 address, and the neighbor IP address as well to deploy VRF from topdown using symmetric VRF lite.
- Deploy configurations in both the fabrics.

VRF Extension Attachment - Attach extensions for given switch(es) ✕

Fabric Name:

Deployment Options

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

MyVRF_50000

<input type="checkbox"/>	Switch	VLAN	Extend	CLI Freeform	Status	Loopb
<input checked="" type="checkbox"/>	LEAF-6	2002	VRF_LITE <input checked="" type="checkbox"/>	Freeform config	NA	

Extension Details

rf...	DOT1Q...	IP_MASK	NEIGHBOR...	NEIGHBOR_ASN	IPV6_MASK	IPV6_NEIGHB...	AUTO_VRF_LITE_FLAG	PEER_VRF_NAME
1/7	3			56				<input type="text"/>

- You can edit or delete IFCs in the **Link** tab in the VXLAN fabric. The extra consideration for auto configured IFCs is that, in order to prevent the regeneration of IFC on next save and deploy, change the mode back to manual mode, or save the configuration only on the relevant devices.
- In a back-to-back scenario, if you delete the VRF lite IFC on one of the fabrics, the VRF lite is deleted from the peer fabric as well.
- When you want to delete a VRF lite between an easy fabric and an external fabric, delete the extension in the easy fabric using the top-down approach. The extension is automatically deleted from the external fabric.
- Deploy the configurations in both the fabrics.
- You must add **redistribute hmm command** in the freeform configuration when vrf-lite is configured on Border device.

See the *Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - VRF Lite* chapter for a use case on VRF Lite.

Layer 3 Port Channels

From Cisco DCNM Release 11.3(1), Layer 3 port channels are supported in external links and interfaces. In the **Interfaces** window, you can select a port channel and a corresponding Layer 3 port channel interface template. This template allows you to configure various options related to Layer 3 port channels including an ability to specify all Layer 3 interface-related configurations. Layer 3 port channels are supported only in easy fabrics and external fabrics.

External connectivity using VRF_LITE will also be supported using Layer 3 port-channels. For physical routed interfaces and LAYER 3 port channel interfaces, you can set the MTU.

You can also watch the video that demonstrates how to extend symmetric VRF Lite using Layer 3 port channels in Cisco DCNM. See the [Extending Symmetric VRF Lite Using Layer 3 Port Channels](#) video.

Configuring Layer 3 Port Channel on Interfaces

To configure a Layer 3 port channel on an interface from the Cisco DCNM Web UI, perform the following steps:

Procedure

- Step 1** Choose **Control > Fabrics > Interfaces**.
The **Interfaces** window appears.
- Step 2** Click **Add Interface**.
The **Add Interface** dialog box appears.
- Step 3** Choose the **Port Channel** type and a device.
The port-channel ID is autopopulated.
- Step 4** Choose the **int_I3_port_channel** policy.
The fields under the **General** area changes accordingly.
- Step 5** Enter the values in the fields and click **Save**.
Only saved configurations are pushed to the device. While adding the interface, you can only modify the policy attribute after the first save. If you try to use an ID that is already used, the **Resource could not be allocated** error appears.
- Step 6** (Optional) Click the **Preview** option to preview the configurations to be deployed.
- Step 7** Click **Deploy** to deploy the specified logical interface.
The newly added interface appears in the screen. You can break out and unbreakout an interface by using the breakout option at the top left.
-

Configuring Layer 3 Port Channel on Interfaces for IOS XE Devices

To configure a Layer 3 port channel on an interface for IOS XE devices, perform the following steps:

Procedure

- Step 1** Choose **Control > Fabrics > Interfaces**.
The **Interfaces** window appears.
- Step 2** Click **Add Interface**.
The **Add Interface** dialog box appears.
- Step 3** Choose the **Port Channel** type and a device.
The port-channel ID is autopopulated.

Step 4 Choose the **ios_xe_int_l3_port_channel** policy.

The fields under the **General** area changes accordingly.

Step 5 Enter the values in the fields and click **Save**.

Only saved configurations are pushed to the device. While adding the interface, you can only modify the policy attribute after the first save. If you try to use an ID that is already used, the **Resource could not be allocated** error appears.

Note The port-channel ID range for Cisco Catalyst 9000 Series switches is from 1 to 128 and for Cisco ASR 1000 Series routers the range is from 1 to 64.

Step 6 (Optional) Click the **Preview** option to preview the configurations to be deployed.

Step 7 Click **Deploy** to deploy the specified logical interface.

The newly added interface appears in the screen.

Deploying Policies on Physical Interfaces for non-Nexus Devices

More policies are added to support non-Nexus devices from Cisco DCNM, Release 11.4(1). When you import any non-Nexus device into an external fabric, few physical interfaces are created by default based on the number of ports. The policy is created only for the management port. For the Cisco Catalyst 9000 Series switches, the management port is GigabitEthernet0/0, and for the Cisco ASR 1000 Series routers, the management port is GigabitEthernet0.

The following table lists the policies added for different non-Nexus devices:

Devices	Policies
Cisco CSR 1000V Series Router	GigabitEthernet
Cisco IOS-XE Devices	<ul style="list-style-type: none"> • GigabitEthernet_mgmt • ios_xe_int_access_host • ios_xe_int_freeform • ios_xe_int_routed_host • ios_xe_int_trunk_host <p>Note Use the GigabitEthernet_mgmt policy only for the management port, which is GigabitEthernet0/0.</p>

To deploy policies on physical interfaces in the **Interfaces** window of Cisco DCNM Web UI, perform the following steps:

Before you begin

Import and discover non-Nexus devices into the external fabric. Ensure that the fabric isn't in monitor mode.

Procedure

Step 1 Check the check box of the interface on which you want to deploy the policy.

Step 2 Click the **Edit Configuration** icon.

Step 3 Choose a policy from the **Policy** drop-down list.

The valid options are:

- **GigabitEthernet**
- **GigabitEthernet_mgmt**
- **ios_xe_int_access_host**
- **ios_xe_int_freeform**
- **ios_xe_int_routed_host**
- **ios_xe_int_trunk_host**

- Note**
- Based on the option you choose, the fields under the **General** area vary.
 - If you choose the **ios_xe_int_routed_host** policy, ensure you have configured the VRF manually, which is out-of-band, or using the **ios_xe_switch_freeform** policy in the **View/Edit Policies** window.
 - DCNM doesn't support NVE or BDI interfaces. However, if you have already created them manually or out-of-band, use the **ios_xe_int_freeform** policy to define their configurations.

Step 4 Enter values for all the mandatory fields.

- Note** Choose the speed based on your device.

Step 5 Click **Save**.

Step 6 Click **Preview** to preview the pending configurations.

Step 7 Click **Deploy** to deploy the policy on the interface.

Configuring Layer 3 Port Channel on Subinterfaces

To configure a Layer 3 port channel on an interface from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Fabrics > Interfaces**.

The **Interfaces** window appears.

Step 2 Choose a Layer 3 port channel interface.

Step 3 Click **Add Interface**.

The **Add Interface** dialog box appears.

- Step 4** Choose the **Subinterface** type.
The subinterface ID and policy are autopopulated, and the fields under the **General** area changes accordingly.
- Step 5** Enter the values in the fields and click **Save**.
Only saved configurations are pushed to the device.
- Step 6** (Optional) Click the **Preview** option to preview the configurations to be deployed.
- Step 7** Click **Deploy** to deploy the specified logical interface.
A confirmation window appears, and the newly added subinterface appears in the list.
-

Configuring Layer 3 Port Channel for Inter-fabric Connectivity

To configure a Layer 3 port channel link from the **Fabric Builder** window, perform the following steps:

Before you begin

Ensure Layer 3 port channels are created on interfaces.

Procedure

- Step 1** Choose an easy fabric or an external fabric, for which you want to extend the VRF-Lite.
The fabric topology window appears.
- Step 2** Click **Tabular view** in the **Actions** pane.
All the components of this fabric are listed with their status and other details accordingly in different tabs.
- Step 3** Choose the **Links** tab.
- Step 4** Click the **Add Link** icon.
The **Add Link** dialog box appears.
- Step 5** Choose **Inter-Fabric** link type.
- Step 6** Choose **VRF_LITE** link sub-type.
- Step 7** Choose the link template from the **Link Template** drop-down list.
Valid values are **ext_fabric_setup_11_1** and **service_link_trunk**.
- Step 8** Enter the details for all other fields accordingly.
- Step 9** Enter the details for fields in the **Link Profile** area wherever necessary.
You can set the MTU. The **Ext_VRF_Lite_Jython** auto-deploy template is used for VRF-Lite configuration on the device in the fabric.
- Step 10** Click **Save**.

Link Management - Edit Link

* Link Type	Inter-Fabric
* Link Sub-Type	VRF_LITE
* Link Template	ext_fabric_setup_11_1
* Source Fabric	Top_Down_ABC
* Destination Fabric	External
* Source Device	BL-2
* Source Interface	Port-channel901
* Destination Device	CORE-2
* Destination Interface	Port-channel901

Link Profile

General
Advanced

* Source BGP ASN	3000.3000	<i>BGP Autonomous System Num</i>
* Source IP Address/Mask	10.33.0.1/30	<i>IP address for sub-interface in e</i>
* Destination IP	10.33.0.2	<i>IP address for sub-interface in e</i>
* Destination BGP ASN	5000.5000	<i>BGP Autonomous System Num</i>
Link MTU	9216	<i>Interface MTU on both ends of</i>
Auto Deploy Flag	<input checked="" type="checkbox"/>	<i>Flag that controls auto generation of neighbor VRF Lite configuration fo</i>

What to do next

After creating a VRF Lite IFC with the Layer 3 port-channel, using the top-down flow, when a VRF is extended using VRF Lite, a sub-interface is created on the Layer 3 port-channel. You can edit the Layer 3 port channel links even after VRFs are extended. However, Layer 3 port channels are not supported for intra-fabric links.

Tabular View - Operational View

From Cisco DCNM 11.3(1), the operational support for a fabric is provided. This feature provides the following information:

- Operational status of a fabric
- Alarm and event notifications

You can view the operational status information in the **Operational View** tab. You can view the alarm and event notifications by clicking the **Alerts and Notifications** icon, next to the **Help** icon, in the top pane of Cisco DCNM.

Viewing the Operational Status

To view the operational status of a fabric from the **Fabric Builder** window, perform the following steps:

Procedure

- Step 1** Choose a fabric.
The fabric topology window appears.
- Step 2** Click **Tabular view** in the **Actions** pane.
- Step 3** Choose the **Operational View** tab.

The Operational View tab has the following fields and descriptions.

Fields	Descriptions
Fabric Name	Specifies the fabrics that have links.
Name	Specifies the link name.
Is Present	Specifies if the link is present or not. Valid values are true and false .
Link State	Specifies the status of the logical link. A logical link can be in one of the following states. <ul style="list-style-type: none"> • Established: When a link is in the Established state the peers send update messages to exchange information about each route advertised to the BGP peer. A notification is sent if there is an error and the state changes to Idle. Only a link using the BGP routing protocol can be in the Established state. • Idle: A link using BGP protocol will be in Idle state when there is an error between peers. • UP: A link using ISIS protocol will be in the UP state, when the link is successfully established between peers. • FULL: A link using the OSPF protocol will be in the FULL state when the link is successfully established between peers. • peer-alive: Specifies the link as a peer keepalive link that monitors the vitality of a vPC peer switch.

Fields	Descriptions
Link Type	Specifies the type of logical link. The link can be of the following type: <ul style="list-style-type: none"> • BGP • ISIS • OSPF • VPC_KEEPLIVE
Uptime	Specifies the duration of the uptime for the link type.

All these columns are sortable.

Viewing Logical Links

The logical links appear in the **Topology** window. To view the logical links from the Cisco DCNM Web UI, perform the following steps:

Procedure

-
- Step 1** Choose **Topology**.
The **Topology** window appears.
- Step 2** Check the **Logical Links** check box in the Show pane.
The logical links between devices appear in blue color.
- Note** The color of the link will change based on its state.
- Step 3** (Optional) Hover over the link to know the link type.
-

Viewing Alerts and Event Notifications

Alert and event notifications includes health score, topology node display, alarm view, alarm policies, and notification services. An event is any action that impacts network, devices or Cisco DCNM. An alert is a notification that is triggered as part of an event to make it visible.

Support for ToR Switches

From Cisco DCNM 11.3(1), support for the Top-of-Rack (ToR) switches is added in DCNM. You can add the Layer 2 ToR switches in an external fabric, and they can be connected to the Leaf switches in the Easy Fabric. For more information, see *Configuring ToR Switches and Deploying Networks*.

vPC Fabric Peering

You can create a virtual peer link for two switches or change the existing physical peer link to a virtual peer link. Only greenfield deployments support vPC fabric peering in Cisco DCNM, Release 11.2(1). However, both greenfield as well as brownfield deployments support vPC fabric peering in Cisco DCNM, Release 11.3(1). This feature is applicable for **Easy_Fabric_11_1** and **Easy_Fabric_eBGP** fabric templates.



Note The **Easy_Fabric_eBGP** fabric does not support brownfield import.

Guidelines and Limitations

The following are the guidelines and limitations for vPC fabric pairing.

- vPC fabric peering is supported from Cisco DCNM Release 11.2(1) and Cisco NX-OS Release 9.2(3).
- Only Cisco Nexus N9K-C9332C Switch, Cisco Nexus N9K-C9364C Switch, Cisco Nexus N9K-C9348GC-FXP Switch as also the Cisco Nexus 9000 Series Switches that ends with FX, and FX2 support vPC fabric peering.
- From Cisco DCNM, Release 11.4(1), Cisco Nexus N9K-C93180YC-FX3S and N9K-C93108TC-FX3P platform switches support vPC fabric peering.
- Cisco Nexus 9300-EX, and 9300-FX/FXP/FX2/FX3 platform switches support vPC Fabric Peering. Cisco Nexus 9200 and 9500 platform switches do not support vPC Fabric Peering.
- If you use other Cisco Nexus 9000 Series Switches, a warning will appear during **Save & Deploy**. A warning appears in this case because these switches will be supported in future releases.
- If you try pairing switches that do not support vPC fabric peering, using the **Use Virtual Peerlink** option, a warning will appear when you deploy the fabric.
- You can convert a physical peer link to a virtual peer link and vice-versa with or without overlays.
- Switches with border gateway leaf roles do not support vPC fabric peering.
- vPC fabric peering is not supported for Cisco Nexus 9000 Series Modular Chassis and FEXs. An error appears during **Save & Deploy** if you try to pair any of these.
- Brownfield deployments and greenfield deployments support vPC fabric peering in Cisco DCNM, Release 11.3(1).
- However, you can import switches that are connected using physical peer links and convert the physical peer links to virtual peer links after **Save & Deploy**. To update a TCAM region during the feature configuration, use the **hardware access-list tcam ingress-flow redirect 5/2** command in the configuration terminal.

QoS for Fabric vPC-Peering

From Cisco DCNM Release 11.4(1), in the **Easy_Fabric_11_1** fabric settings, you can enable QoS on spines for guaranteed delivery of vPC Fabric Peering communication. Additionally, you can specify the QoS policy name.

Note the following guidelines for a greenfield deployment:

- If QoS is enabled and the fabric is newly created:
 - If spines or super spines neighbor is a virtual vPC, make sure neighbor is not honored from invalid links, for example, super spine to leaf or borders to spine when super spine is present.
 - Based on the Cisco Nexus 9000 Series Switch model, create the recommended global QoS config using the **switch_freeform** policy template.
 - Enable QoS on fabric links from spine to the correct neighbor.
- If the QoS policy name is edited, make sure policy name change is honored everywhere, that is, global and links.
- If QoS is disabled, delete all configuration related to QoS fabric vPC peering.
- If there is no change, then honor the existing PTI.

For more information about a greenfield deployment, see the *Creating a New VXLAN BGP EVPN Fabric* section.

Note the following guidelines for a brownfield deployment:

Brownfield Scenario 1:

- If QoS is enabled and the policy name is specified:



Note You need to enable only when the policy name for the global QoS and neighbor link service policy is same for all the fabric vPC peering connected spines.

- Capture the QoS config from switch based on the policy name and filter it from unaccounted configuration based on the policy name and put the configuration in the **switch_freeform** with PTI description.
- Create service policy configuration for the fabric interfaces as well.
- Greenfield config should make sure to honor the brownfield config.
- If the QoS policy name is edited, delete the existing policies and brownfield extra configuration as well, and follow the greenfield flow with the recommended config.
- If QoS is disabled, delete all the configuration related to QoS fabric vPC peering.



Note No cross check for possible or error mismatch user configuration, and user might see the diff.

Brownfield Scenario 2:

- If QoS is enabled and the policy name is not specified, QoS configuration is part of the unaccounted switch freeform config.
- If QoS is enabled from fabric settings after **Save & Deploy** for brownfield, QoS configuration overlaps and you will see the diff if fabric vPC peering config is already present.

For more information about a brownfield deployment, see the *Creating a New VXLAN BGP EVPN Fabric* section.

Fields and Description

To view the vPC pairing window of a switch, from the fabric topology window, right-click the switch and choose **vPC Pairing**. The vPC pairing window for a switch has the following fields:

Field	Description
Use Virtual Peerlink	Allows you to enable or disable the virtual peer linking between switches.
Switch name	Specifies all the peer switches in a fabric. Note When you have not paired any peer switches, you can see all the switches in a fabric. After you pair a peer switch, you can see only the peer switch in the vPC pairing window.
Recommended	Specifies if the peer switch can be paired with the selected switch. Valid values are true and false . Recommended peer switches will be set to true .
Reason	Specifies why the vPC pairing between the selected switch and the peer switches is possible or not possible.
Serial Number	Specifies the serial number of the peer switches.


You can perform the following with the **vPC Pairing** option:

Creating a Virtual Peer Link


To create a virtual peer link from the Cisco DCNM Web UI, perform the following steps:

Procedure

-
- Step 1** Choose **Control > Fabrics**.
The **Fabric Builder** window appears.
- Step 2** Choose a fabric with the **Easy_Fabric_11_1** or **Easy_Fabric_eBGP** fabric templates.
The fabric topology window appears.



- Step 3** Right-click a switch and choose **vPC Pairing** from the drop-down list.
The window to choose the peer appears.



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

You will get the following error when you choose a switch with the border gateway leaf role.

```
<switch-name> has a Network/VRF attached. Please detach the Network/VRF before vPC Pairing/Unpairing
```

Step 4 Check the **Use Virtual Peerlink** check box.

Step 5 Choose a peer switch and check the **Recommended** column to see if pairing is possible.

If the value is **true**, pairing is possible. You can pair switches even if the recommendation is **false**. However, you will get a warning or error during **Save & Deploy**.

Step 6 Click **Save**.

Select vPC peer for leaf5 ✕

Use Virtual Peerlink

1

	Switch name	Recommended ▼	Reason	Serial Number
2	<input checked="" type="radio"/> leaf6	true	Switches have same role	FDO22360M0D
	<input type="radio"/> leaf3	false	Already paired with FDO20352BEE	FDO20290DVJ
	<input type="radio"/> leaf1	false	N9K-C93180YC-EX doesn't support Virtu...	FDO2035283H
	<input type="radio"/> spine2	false	Switches have different roles	FDO20352B6H
	<input type="radio"/> spine1	false	Switches have different roles	FDO20401L8J
	<input type="radio"/> leaf2	false	Already paired with FDO20290DVJ	FDO20352BEE

3 Save Cancel

Step 7 In the **Fabric Topology** window, click **Save & Deploy**.

The **Config Deployment** window appears.

Step 8 Click the field against the switch in the **Preview Config** column.

The **Config Preview** window appears for the switch.


Step 9 View the vPC link details in the pending configuration and the side-by-side configuration.

Step 10 Close the window.

Step 11 Click the pending errors icon next to the **Save & Deploy** icon to view errors and warnings, if any.

If you see any warnings that are related to TCAM, click the **Resolve** icon. A confirmation dialog box about reloading switches appears. Click **OK**. You can also reload the switches from **Tabular view** in the fabric topology window.

The switches that are connected through vPC fabric peering, are enclosed in a gray cloud.



Converting a Physical Peer Link to a Virtual Peer Link

To convert a physical peer link to a virtual peer link from the Cisco DCNM Web UI, perform the following steps:

Before you begin

- Plan the conversion from physical peer link to virtual peer link during the maintenance window of switches.
- Ensure the switches support vPC fabric peering. Only the following switches support vPC fabric peering:
 - Cisco Nexus N9K-C9332C Switch, Cisco Nexus N9K-C9364C Switch, and Cisco Nexus N9K-C9348GC-FXP Switch
 - Cisco Nexus 9000 Series Switches that ends with FX, FX2, and FX2-Z

Procedure

- Step 1** Choose **Control > Fabrics**.
The **Fabric Builder** window appears.
- Step 2** Choose a fabric with the **Easy_Fabric_11_1** or **Easy_Fabric_eBGP** fabric templates.

Step 3 Right-click the switch that is connected using the physical peer link and choose **vPC Pairing** from the drop-down list.

The window to choose the peer appears.

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

You will get the following error when you choose a switch with the border gateway leaf role.

```
<switch-name> has a Network/VRF attached. Please detach the Network/VRF before vPC
Pairing/Unpairing
```

Step 4 Check the **Recommended** column to see if pairing is possible.

If the value is **true**, pairing is possible. You can pair switches even if the recommendation is **false**. However, you will get a warning or error during **Save & Deploy**.

Step 5 Check the **Use Virtual Peerlink** check box.

The **Unpair** icon changes to **Save**.

Step 6 Click **Save**.

Note After you click **Save**, the physical vPC peer link is automatically deleted between the switches even without deployment.

Step 7 In the **Fabric Topology** window, click **Save & Deploy**.

The **Config Deployment** window appears.

Step 8 Click the field against the switch in the **Preview Config** column.

The **Config Preview** window appears for the switch.

Step 9 View the vPC link details in the pending configuration and the side-by-side configuration.

Step 10 Close the window.

Step 11 Click the pending errors icon next to the **Save & Deploy** icon to view errors and warnings, if any.

If you see any warnings that are related to TCAM, click the **Resolve** icon. A confirmation dialog box about reloading switches appears. Click **OK**. You can also reload the switches from **Tabular view** in the fabric topology window.

The physical peer link between the peer switches turns red. Delete this link. The switches are connected only through a virtual peer link and are enclosed in a gray cloud.

Converting a Virtual Peer Link to a Physical Peer Link

To convert a virtual peer link to a physical peer link from the Cisco DCNM Web UI, perform the following steps:

Before you begin

Connect the switches using a physical peer link before disabling the vPC fabric peering.

Procedure

- Step 1** Choose **Control > Fabrics**.
The **Fabric Builder** window appears.
- Step 2** Choose a fabric with the **Easy_Fabric_11_1** or **Easy_Fabric_eBGP** fabric templates.
- Step 3** Right-click the switch that is connected through a virtual peer link and choose **vPC Pairing** from the drop-down list.
The window to choose the peer appears.
- Note** Alternatively, you can also navigate to the **Tabular view** from the **Actions** pane. Choose a switch in the **Switches** tab and click **vPC Pairing** to create, edit, or unpair a vPC pair. However, you can use this option only when you choose a Cisco Nexus switch.
- Step 4** Uncheck the **Use Virtual Peerlink** check box.
The **Unpair** icon changes to **Save**.
- Step 5** Click **Save**.
- Step 6** In the **Fabric Topology** window, click **Save & Deploy**.
The **Config Deployment** window appears.
- Step 7** Click the field against the switch in the **Preview Config** column.
The **Config Preview** window appears for the switch.
- Step 8** View the vPC peer link details in the pending configuration and the side-by-side configuration.
- Step 9** Close the window.
- Step 10** Click the pending errors icon next to the **Save & Deploy** icon to view errors and warnings, if any.
If you see any warnings that are related to TCAM, click the **Resolve** icon. The confirmation dialog box about reloading switches appears. Click **OK**. You can also reload the switches from **Tabular view** in the fabric topology window.
The virtual peer link, represented by a gray cloud, disappears and the peer switches are connected through a physical peer link.
-

Advertising PIP on vPC

In the fabric settings, you can check the **vPC advertise-pip** check box to enable the Advertise PIP feature on all vPCs in a fabric. From Cisco DCNM Release 11.4(1), you can use the **vpc_advertise_pip_jython** policy to enable Advertise PIP feature on specific vPCs in a fabric.

Note the following guidelines:

- If advertise-pip is not globally enabled or vPC peer is not using fabric peering, only then the vpc_advertise_pip_jython policy can be created on specific peers.
- Enabling **vpc advertise-pip** doesn't affect the current behavior.

- Disabling advertise pip for a fabric doesn't affect this policy.
- Unpairing of switches deletes this policy.
- You can manually delete this policy from the peer switch where it was created.

Procedure

- Step 1** From the **Fabric Builder** window, click a fabric, and then right-click on a switch with vPC and select **View/Edit Policies**.
- Step 2** Click **Add** and select the **vpc_advertise_pip_jython** policy template and enter the mandatory parameters data.
- Note** You can add this policy on one vPC peer, and it will create respective commands for vpc advertise on both peers.
- Step 3** Click **Save**, and then deploy this policy.
-

ThousandEyes Enterprise Agent

ThousandEyes Enterprise Agent collects network and application layer performance data when users access specific websites within monitored networks. It is used to run tests, check detailed aspects of network pathing and connectivity, status of network routing, monitor changes in intent, running configuration, and so on.

From Release 11.5(3), ThousandEyes Enterprise Agent is integrated with Cisco DCNM.

ThousandEyes Enterprise Agent is supported on Cisco Nexus 3000-R Series and Cisco Nexus 9000 Cloud Scale Series, with NX-OS version 9.3(7) and 10.2(1) and later releases.

This is supported with the following fabric templates:

- Easy_Fabric_11_1
- Easy_Fabric_eBGP
- External_Fabric_11_1
- LAN_Classic

You can configure global settings for ThousandEyes Enterprise Agent using Cisco DCNM Web UI > **Control** > **ThousandEyes** > **Configure**.

The section includes the following:


Configuring TCAM and CoPP Policies

Ensure that you add relevant policies to Cisco Nexus 3000-R Series and Cisco Nexus 9000 Cloud Scale Series Switches before installing ThousandEyes Enterprise Agent feature on the switches.

To configure TCAM and CoPP policies on switches from the Cisco DCNM Web UI, perform the following steps:

Procedure

- Step 1** From DCNM Web UI, choose **Control** > **Fabric Builder**, choose a fabric and click **Tabular View** in the **Actions** window.
- The **Switches** tab is displayed.
- Step 2** Select a single or multiple switches in the **Switches** tab and click the **Policies** button.
- Step 3** Click **Add** icon.
- Step 4** To add TCAM policies for Cisco Nexus 9000 EX, FX, and FX2 series switches perform following the steps:



- Choose `ThousandEyes_Agent_N9K_EX_tcam_config` for EX series switches and `ThousandEyes_Agent_N9K_FX_FEX2_tcam_config` for FX and FX2 series switches.
- Enter value 200 in **Priority** field and click **Save**.
- On the **Switches** tab, choose the switch for which policy is added. Click **Deploy** to deploy configurations on the switches.

Note Warning messages are displayed indicating that the switches need to reload for the TCAM changes to reflect on the switch, click **Resolve** to reload the switch.

- Step 5** To add CoPP policies for `Easy_Fabric_11_1` and `Easy_Fabric_eBGP` templates perform following the steps:
- From DCNM Web UI, choose **Control** > **Fabric Builder** > **Fabric Settings**, click **Advanced** tab.
 - Choose manual in **CoPP Profile** field.
- Step 6** To deploy the policy on all the supported switches and fabric templates, perform the following steps:
- Choose an appropriate switch and click **Play** button.
 - The **Execute Switch CLIs on Devices** window appears.
 - Choose `ThousandEyes_Agent_Copy_CoPP` from Template drop-down list and click **Deploy**.

- On the **Switches** tab, choose the appropriate switch. Click **Policy**.
The Policy window appears
 - Click **Add** icon.
 - Choose **ThousandEyes_Agent_CoPP** from Policy drop-down list.
 - Enter value 210 in Priority field and click **Save**.
 - On the **Switches** tab, choose the switch for which policy is added. Click **Save** to deploy configurations to the switches.
-

Performing ThousandEyes Enterprise Agent Actions

You can perform ThousandEyes Enterprise Agent action only for fabrics that are in managed mode.

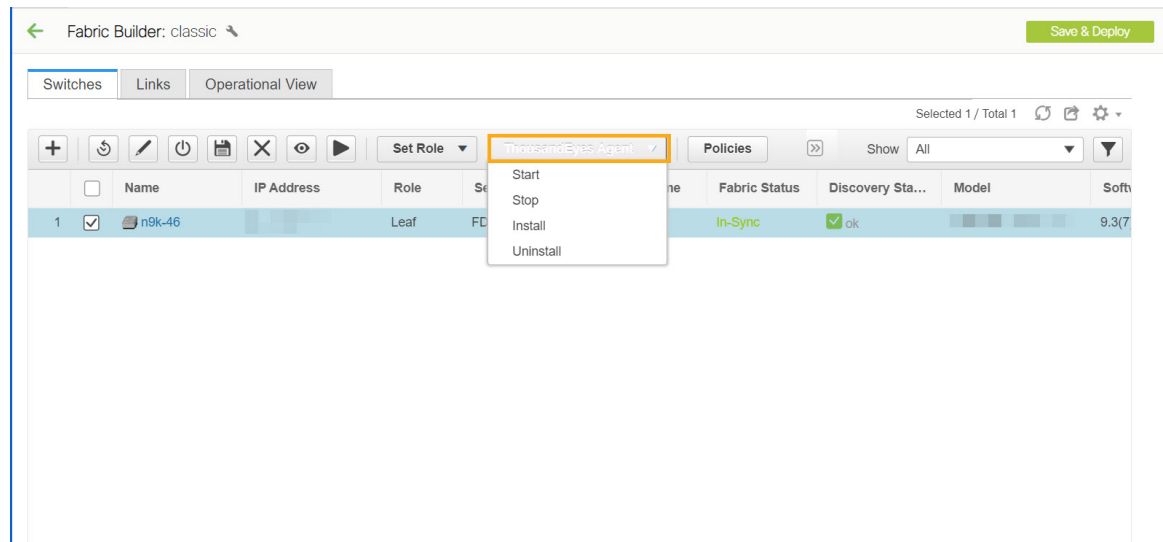


Note Ensure that the TCAM and COPP policies are configured on switches, before installing ThousandEyes Enterprise Agent on it.

To start, stop, install, or uninstall ThousandEyes Enterprise Agent using DCNM Web UI, perform the following steps:

Procedure

- Step 1** Choose **Control > Fabric Builder**.
The **Fabric Builder** window appears. A rectangular box represents each fabric.
- Step 2** Choose a fabric and click **Tabular View** in the **Actions** window.
The **Switches** tab is displayed.
- Step 3** Select single or multiple switches and click required action from **ThousandEyes Agent** drop- down list.



You can perform following actions:

- **Install** – Installs ThousandEyes Enterprise Agent on the switches. After the installation, the ThousandEyes Agent Status column displays as RUNNING.
- **Start** – Starts ThousandEyes Enterprise Agent on the switches, which was stopped earlier.
- **Note** – You must install ThousandEyes Enterprise Agent, before you start the agent on the switches.
- **Stop** – Stops ThousandEyes Enterprise Agent on the switches.
- **Uninstall** – Uninstalls ThousandEyes Enterprise Agent from the switches. A pop-up window appears after you perform any action, displaying a message - **ThousandEyes actions completed. Please check status!**

Uninstalling the ThousandEyes Enterprise Agent from DCNM will not clear the account group token number in the ThousandEyes portal. To remove the existing ThousandEyes Enterprise Agent account group token on the switches, refer to the [Removing ThousandEyes Enterprise Agent](#) section.

ThousandEyes Enterprise Agent Status

ThousandEyes Enterprise Agent status messages are as listed below:

- **NOT_INSTALLED** - ThousandEyes Enterprise Agent is not installed on the switch.
- **RUNNING** - ThousandEyes Enterprise Agent is active on the switch.
- **STOPPED** - ThousandEyes Enterprise Agent has stopped on the switch.
- **UNSUPPORTED_VERSION** - ThousandEyes Enterprise Agent is not supported with the switch NX-OS version.
- **UNSUPPORTED_PLATFORM** - ThousandEyes Enterprise Agent is not supported on the selected switch platform.
- **NA** - ThousandEyes Enterprise Agent global settings not configured on DCNM

1. Click **ThousandEyes Status**, to view information of ThousandEyes Enterprise Agent
The **Detailed ThousandEyes Agent Information** page appears.

- **Log Info** tab displays the runtime agent status or error logs of the switch.
- **Sync Status** tab displays deployed and expected settings details of switch.

DCNM indicates configuration mismatch (**In-Sync, Out-Of-Sync**) when ThousandEyes Enterprise Agent configuration is different from the effective configuration on DCNM at that instant. In case of configuration mismatch, you must uninstall, remove and install the ThousandEyes Enterprise Agent to make the configuration In-Sync.

Detailed ThousandEyes Agent Information - [REDACTED]

Log Info

Sync Status

ThousandEyes Agent Status: ✖ Out-Of-Sync

	Deployed Settings	Expected Settings
1	Setting Enabled:Global	Setting Enabled:Global
2	Account Token:[REDACTED]	Account Token:[REDACTED]
3	DNS Domain:cisco.com	DNS Domain:cisco.com
4	DNS IPs:[REDACTED]	DNS IPs:[REDACTED]
5	NTP IPs:[REDACTED]	NTP IPs:[REDACTED]
6	Proxy Enable:True	Proxy Enable:True
7	Proxy Bypass:[REDACTED]	Proxy Bypass:[REDACTED]
8	Proxy Info:[REDACTED]	Proxy Info:prox:[REDACTED]
9	VRF:management	VRF:default

Removing ThousandEyes Enterprise Agent

To remove the existing ThousandEyes Enterprise Agent entry in the ThousandEyes Enterprise portal, refer to instructions in [Removing Old Agent Entries](#) section.

To remove the existing ThousandEyes Enterprise Agent account group token from the switches on DCNM, perform the following steps:

Procedure

-
- Step 1** From Cisco DCNM Web UI, choose **Control > Fabric Builder**.
The **Fabric Builder** window appears. A rectangular box represents each fabric.
 - Step 2** Choose a fabric and click **Tabular View** in the **Actions** window.
The **Switches** tab is displayed.
 - Step 3** Select the appropriate switches to remove ThousandEyes Enterprise Agent and click **Play** button (Execute Commands).
The **Execute Switch CLIs on Devices** window appears.
 - Step 4** Choose **ThousandEyes_Agent_Identity_Delete** from Template drop-down list and click **Deploy**.
-

Viewing and Editing Policies

Cisco DCNM provides the ability to group a set of switches, and allows you to push a set of underlay configurations to the group. This release enables you to create a policy template, and apply it to multiple selected switches.

To view, add, deploy, or edit a policy, perform the following steps:

Procedure

- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Select any available fabric, and then click **Tabular view**.
- Step 3** Select multiple switches in switches tab, and click **View/Edit Policies**.

Note **View/Edit Policies** is not enabled for an MSD fabric.

Viewing Policies

Procedure



- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Select any available fabric, and then click **Tabular view**.
- Step 3** Select multiple switches in the switches tab and click **View/Edit Policies**.

Policies are listed in view or edit policies table for multiple switches.

	+	↺	✎	⏻	✕	View/Edit Policies	Manage Interfaces	History	Deploy	Tracker Actions	⌵	Show	All	⌵
	☑	Name	IP Address	Role	Serial Number	Fabric Name	Fabric Status	Discovery Sta...	Model	Software Versi...	Tracker Stat...	Last Updated		
1	☑	n9k12_bp2-f...	80.80.80.62	leaf	SAL18422FX8	BF	In-Sync	☑ ok	N9K-C9396PX	7.0(3)7(6)	NOT_INSTALLI	an hour ago		
2	☑	n9k13_bp2-f...	80.80.80.63	leaf	SAL18422FXE	BF	In-Sync	☑ ok	N9K-C9396PX	7.0(3)7(6)	NOT_INSTALLI	an hour ago		
3	☑	n9k7_bp2-fs...	80.80.80.57	border	SAL1833YM64	BF	In-Sync	☑ ok	N9K-C9396PX	7.0(3)7(6)	NOT_INSTALLI	an hour ago		
4	☑	n9k14_bp2-s...	80.80.80.64	spine	SAL2016NXXB	BF	In-Sync	☑ ok	N9K-C92160YC-X	7.0(3)7(6)	NOT_INSTALLI	an hour ago		
5	☑	n9k8_bp2-sp...	80.80.80.58	spine	SAL1833YMOV	BF	In-Sync	☑ ok	N9K-C9396PX	9.3(1)	NOT_INSTALLI	an hour ago		

View/Edit Policies



Selected 0 / Total 1762  

<input type="checkbox"/>	Policy ID	Template	Description	Generated Config	Entity Name	Entity Type	Source
<input type="checkbox"/>	POLICY-127750	ingress_rep_simulated		View	SWITCH	SWITCH	
<input type="checkbox"/>	POLICY-106330	host_11_1		View	SWITCH	SWITCH	
<input type="checkbox"/>	POLICY-106360	feature_nxapi		View	SWITCH	SWITCH	UNDEI
<input type="checkbox"/>	POLICY-106380	pre_config		View	SWITCH	SWITCH	UNDEI
<input type="checkbox"/>	POLICY-106610	base_feature_spine_...		View	SWITCH	SWITCH	UNDEI
<input type="checkbox"/>	POLICY-106620	feature_ospf		View	SWITCH	SWITCH	UNDEI
<input type="checkbox"/>	POLICY-106630	feature_tacacs		View	SWITCH	SWITCH	
<input type="checkbox"/>	POLICY-109520	host_11_1		View	SWITCH	SWITCH	
<input type="checkbox"/>	POLICY-109540	feature_nxapi		View	SWITCH	SWITCH	UNDEI
<input type="checkbox"/>	POLICY-109560	pre_config		View	SWITCH	SWITCH	UNDEI
<input type="checkbox"/>	POLICY-109770	base_feature_spine_...		View	SWITCH	SWITCH	UNDEI

Note You can view the generated config for a device by hovering over the **View** button under the **Generated Config** column. Additionally, you can enter a config in the search field under this column to filter policies.

Step 4 Select a policy and click the **View** button to view its configs.

Note Python policies are used to place logic and control CLI policies. From DCNM Release 11.3(1), multiple CLI child policies are aggregated for each Python policy.

Step 5 In the **View/Edit Policies** window, click **View All** to view all the configurations pushed to the switches using policies.

Generated Config for the selected devices

Go To Include Policy ID

```
#####
#SAL18422FX8#
#####
#POLICY-106330#
hostname n9k8_bp2-spsw-1001

#POLICY-106360#
feature nxapi

#POLICY-106380#
ipv6 switch-packets lla

#POLICY-106610#
nv overlay evpn
feature lldp
feature bgp

#POLICY-106620#
feature ospf

#POLICY-106630#
feature tacacs+

#POLICY-125130#
```

Go To: Select a device from this drop-down list to navigate to its starting config.

This option is applicable only when you view policies for multiple devices.

Include Policy ID: Select this check box to view policy IDs for all the policies. By default, this check box is selected.

Adding a Policy

Procedure

- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Select any available fabric, and then click Tabular view.
- Step 3** Select a single or multiple switches in the **Switches** tab, and click the **View/Edit Policies** button.
- Step 4** Click the **Add** icon.
- Step 5** Select a policy template and enter the mandatory parameters data and click **Save**. PTI is added per each device based on n-number of devices selection.

Add Policy
✕

* Policy:

* Priority (1-1000): Description:

General

Variables:

*** Switch Freeform Config**

```
feature bash-shell
feature telemetry

clock timezone CET 1 0
clock summer-time CEST 5 Sunday March 02:00 5 Sunday October 03:00 60
clock protocol ntp vdc 1

telemetry
destination-profile
use-vrf management
```

Policy: Select a policy from this drop-down list.

Priority: Specify a priority for the policy. The applicable values are from 1 to 1000. The default value is 500. The lower number in the **Priority** field means that there is a higher priority for the generated configuration and POAP startup-configuration. For example, features are 50, route-maps are 100, and vpc-domain is 200.

Description: (Optional) Specify a description for the policy. This field is used to differentiate multiple freeform policies. The **Description** column is added in the **View/Edit Policies** window, which you can use to filter or find policies based on description.

Deploying Policies

Procedure

- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Select any available fabric, and then click Tabular view.
- Step 3** Select multiple switches in the switches tab, and click the **View/Edit Policies** button.
- Step 4** Select multiple polices, and then click **Push Config**. The selected PTI's configs are pushed to the group of switches.
 - If the external fabric is in the monitor mode, the **Push Config** option is disabled.
 - This option will be greyed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.

Editing a Policy



Note Multiple policy editing is not supported.

Procedure

- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Select any available fabric, and then click **Tabular view**.
- Step 3** Select multiple switches in the switches tab, and click the **View/Edit Policies** button.

View/Edit Policies ✕

Selected 0 / Total 1762

<input type="checkbox"/>	Policy ID	Template	Description	Generated Config	Entity Name	Entity Type	Source
<input type="checkbox"/>	POLICY-127750	ingress_rep_simulated		View	SWITCH	SWITCH	
<input type="checkbox"/>	POLICY-106330	host_11_1		View	SWITCH	SWITCH	
<input type="checkbox"/>	<i>POLICY-106360</i>	<i>feature_nxapi</i>		View	SWITCH	SWITCH	UNDEI
<input type="checkbox"/>	<i>POLICY-106380</i>	<i>pre_config</i>		View	SWITCH	SWITCH	UNDEI
<input type="checkbox"/>	<i>POLICY-106610</i>	<i>base_feature_spine_...</i>		View	SWITCH	SWITCH	UNDEI
<input type="checkbox"/>	<i>POLICY-106620</i>	<i>feature_ospf</i>		View	SWITCH	SWITCH	UNDEI
<input type="checkbox"/>	POLICY-106630	feature_tacacs		View	SWITCH	SWITCH	
<input type="checkbox"/>	POLICY-109520	host_11_1		View	SWITCH	SWITCH	
<input type="checkbox"/>	<i>POLICY-109540</i>	<i>feature_nxapi</i>		View	SWITCH	SWITCH	UNDEI
<input type="checkbox"/>	<i>POLICY-109560</i>	<i>pre_config</i>		View	SWITCH	SWITCH	UNDEI
<input type="checkbox"/>	<i>POLICY-109770</i>	<i>base_feature_spine_...</i>		View	SWITCH	SWITCH	UNDEI

Note The policies in the italics font cannot be edited. The value under the **Editable** and **Mark Deleted** columns for these policies is **false**.

- Step 4** Select a PTI, click **Edit** to modify the required data, and then click **Save** to save the PTI.
- Step 5** Select a PTI, click **Edit** to modify the required data, and then click **Push Config** to push the policy config to the device.

- Note**
- This option will be greyed out if the fabric is in freeze mode, that is, if you have disabled deployments on the fabric.
 - A warning appears if you push config for a Python policy.
 - A warning appears if you edit, delete, or push config a mark-deleted policy. A mark-deleted policy is set to **true** under the **Mark Deleted** column. The switch freeform child policies of **Mark Deleted** policies appears in the **View/Edit Policies** dialog box. You can edit only **Python** switch_freeform policies. You cannot edit **Template_CLI** switch_freeform_config policies.

Edit Policy
✕

Policy ID: POLICY-125140

Template: bgp_lb_id

* Priority (1-1000):

Entity Type: SWITCH

Entity Name: SWITCH

Description:

General

* Loopback Id ? Loopback Id

Variables:

Current Switch Configuration

Procedure

- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Select any available fabric, and then click **Tabular** view.
- Step 3** Select multiple switches in the switches tab, and click **View/Edit Policies**.
- Step 4** Click **Current Switch Config**.

The current switch configuration appears in the **Running Config** dialog box.

Note The running configuration will not appear for the Cisco CSR 1000v when you click **Current Switch Config** if the user role cannot access the enable prompt by default.

Retrieving the Authentication Key

Retrieving the 3DES Encrypted OSPF Authentication Key

1. SSH into the switch.
2. On an unused switch interface, enable the following:

```

config terminal
  feature ospf
  interface Ethernet1/1
    no switchport
    ip ospf message-digest-key 127 md5 ospfAuth

```

In the example, **ospfAuth** is the unencrypted password.



Note This Step 2 is needed when you want to configure a new key.

3. Enter the **show run interface Ethernet1/1** command to retrieve the password.

```

Switch # show run interface Ethernet1/1
interface Ethernet1/1
  no switchport
  ip ospf message-digest key 127 md5 3 sd8478f4fsw4f4w34sd8478fsdfw
  no shutdown

```

The sequence of characters after **md5 3** is the encrypted password.

4. Update the encrypted password into the **OSPF Authentication Key** field.

Retrieving the Encrypted IS-IS Authentication Key

To get the key, you must have access to the switch.

1. SSH into the switch.
2. Create a temporary keychain.

```

config terminal
  key chain isis
  key 127
  key-string isisAuth

```

In the example, **isisAuth** is the plaintext password. This will get converted to a Cisco type 7 password after the CLI is accepted.

3. Enter the **show run | section “key chain”** command to retrieve the password.

```

key chain isis
  key 127
    key-string 7 071b245f5a

```

The sequence of characters after **key-string 7** is the encrypted password. Save it.

4. Update the encrypted password into the ISIS Authentication Key field.
5. Remove any unwanted configuration made in Step 2.

Retrieving the 3DES Encrypted BGP Authentication Key

1. SSH into the switch and enable BGP configuration for a non-existent neighbor.



Note Non-existent neighbor configuration is a temporary BGP neighbor configuration for retrieving the password.

```
router bgp
  neighbor 10.2.0.2 remote-as 65000
  password bgpAuth
```

In the example, **bgpAuth** is the unencrypted password.

2. Enter the `show run bgp` command to retrieve the password. A sample output:

```
neighbor 10.2.0.2
  remote-as 65000
  password 3 sd8478fswerdfw3434fsw4f4w34sdsd8478fswerdfw3434fsw4f4w3
```

The sequence of characters after password 3 is the encrypted password.

3. Update the encrypted password into the **BGP Authentication Key** field.
4. Remove the BGP neighbor configuration.

Retrieving the Encrypted BFD Authentication Key

1. SSH into the switch.
2. On an unused switch interface, enable the following:

```
switch# config terminal
switch(config)# int e1/1
switch(config-if)# bfd authentication keyed-SHA1 key-id 100 key cisco123
```

In the example, **cisco123** is the unencrypted password and the key ID is **100**.



Note This Step 2 is needed when you want to configure a new key.

3. Enter the `show running-config interface` command to retrieve the key.

```
switch# show running-config interface Ethernet1/1

interface Ethernet1/1
description connected-to- switch-Ethernet1/1
no switchport
mtu 9216
bfd authentication Keyed-SHA1 key-id 100 hex-key 636973636F313233
no ip redirects
ip address 10.4.0.6/30
no ipv6 redirects
ip ospf network point-to-point
ip router ospf 100 area 0.0.0.0
no shutdown
```

The BFD key ID is **100** and the encrypted key is **636973636F313233**.

4. Update the key ID and key in the **BFD Authentication Key ID** and **BFD Authentication Key** fields.

Custom Maintenance Mode Profile Policy

When you place a switch in maintenance mode using DCNM, only a fixed set of BGP and OSPF isolate CLIs are configured in the maintenance mode profile. Starting from Cisco DCNM Release 11.3(1), you can create a **custom_maintenance_mode_profile** PTI with customized configurations for maintenance mode and normal mode profile, deploy the PTI to the switch, and then move the switch to maintenance mode.

Creating and Deploying a Custom Maintenance Mode Profile Policy

Procedure

- Step 1** Select **Control>Fabric Builder**, click **Tabular View**, and select a switch in the **Name** column or select **Control>Fabric Builder** and right-click the switch.
- Step 2** Click **View/Edit Policies** and click on + to add a new policy. The **Add Policy** window comes up.
- Step 3** Select **custom_maintenance_mode_profile** from the **Policy** dropdown list.
- Step 4** Fill in the **Maintenance mode profile contents** with the desired configuration CLIs.

Example:

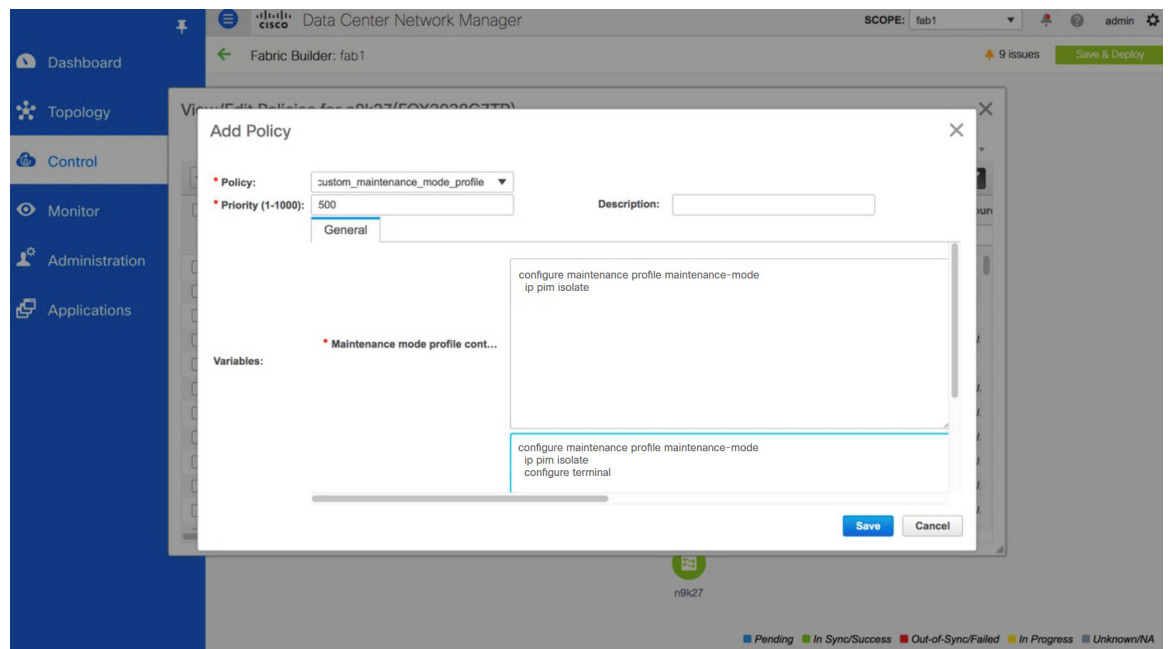
```
configure maintenance profile maintenance-mode
ip pim isolate
```

Fill in the **Normal mode profile contents** with the desired configuration CLIs.

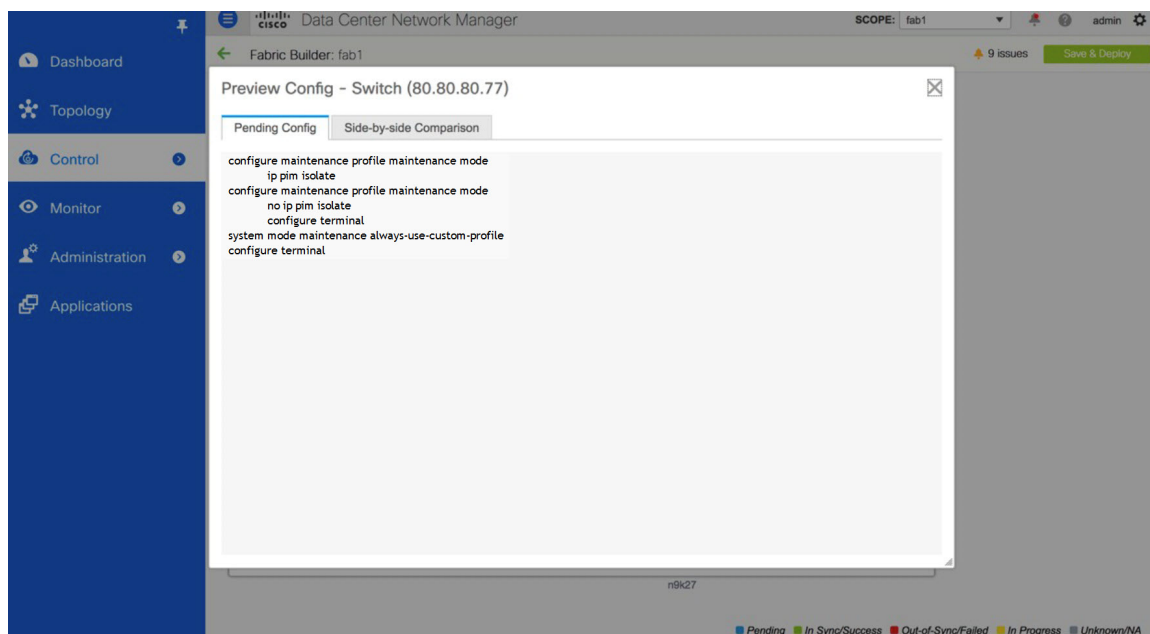
Example:

```
configure maintenance profile normal-mode
no ip pim isolate
configure terminal
```

- Step 5** Click **Save**.



- Step 6** Right-click the switch in the **Fabric Builder** window and select **Deploy Config**. Verify the configuration in the **Pending Config** window and then deploy the configuration to the switch.




Step 7 Then, right-click the switch and select **Modes>Maintenance Mode** to move the switch to maintenance mode.

Deleting a Custom Maintenance Mode Profile Policy


Procedure

- Step 1** The switch has to be moved to active/operational or normal mode before deleting the custom maintenance mode profile policy. To do this, right-click the switch in the **Fabric Builder** window and select **Modes>Active/Operational Mode**.
- Step 2** After the switch has been moved to active/operational or normal mode, click **Tabular View** in the **Fabric Builder** window, and select the switch in the **Name** column or right-click the switch in the **Fabric Builder** window.
- Step 3** Click **View/Edit Policies**, and select the **custom_maintenance_mode_profile** policy that has to be deleted.
- Step 4** Click **X** to delete the policy.



Step 5 Right-click the switch in the **Fabric Builder** window and select **Deploy Config**. Verify the configuration in the **Pending Config** window and deploy the configuration to the switch.

```
no system mode maintenance always-use-custom-profile
no configure maintenance profile normal-mode
no configure maintenance profile maintenance-mode
configure terminal
```



Return Material Authorization (RMA)

This section describes how to replace a physical switch in a Fabric when using Cisco DCNM Easy Fabric mode.

Prerequisites

- Ensure that the fabric is up and running with minimal disruption while replacing the switch.
- To use the POAP RMA flow, configure the fabric for bootstrap (POAP).
- Perform save and deploy more than once, if needed, to copy the FEX configurations for the RMA of switches that have FEX deployed.

Guidelines and Limitations

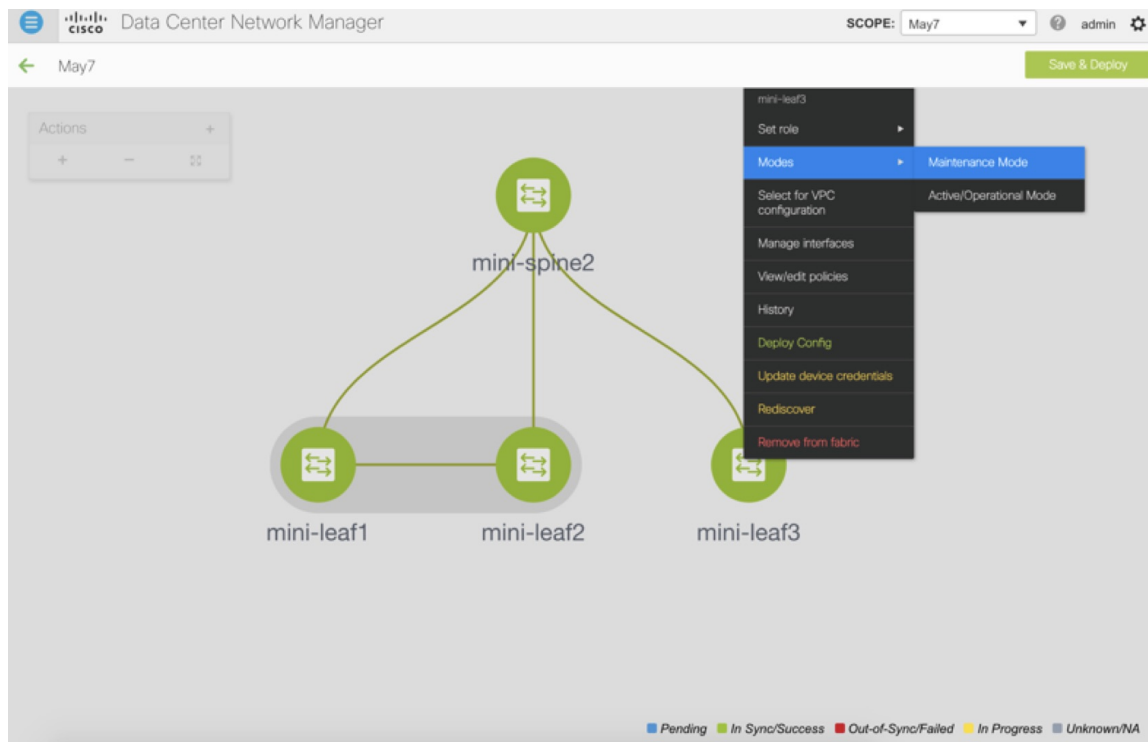
- To replace the switch, remove the old switch from the fabric and discover the new switch in the fabric. For example, if you want to replace a Cisco Nexus 9300-EX switch with a Cisco Nexus 9300-FX switch, remove the 9300-EX switch from the fabric followed by discovering the 9300-FX switch in the same fabric.
- When GIR is enabled before upgrading Cisco Nexus 7000 Series switches, DCNM pushes the **system mode maintenance** command to the switches when the DCNM RMA procedure is initiated. This command applies the configuration that is present in the default maintenance mode profile to the switches. For more information on performing Graceful Insertion and Removal (GIR) on the Cisco Nexus 7000 Series switches, refer [Configuring GIR](#).

If a switch is in maintenance mode, to initiate copy run and start configuration, navigate to **Fabric Builder > Discovery** and reload the switch.

POAP RMA Flow

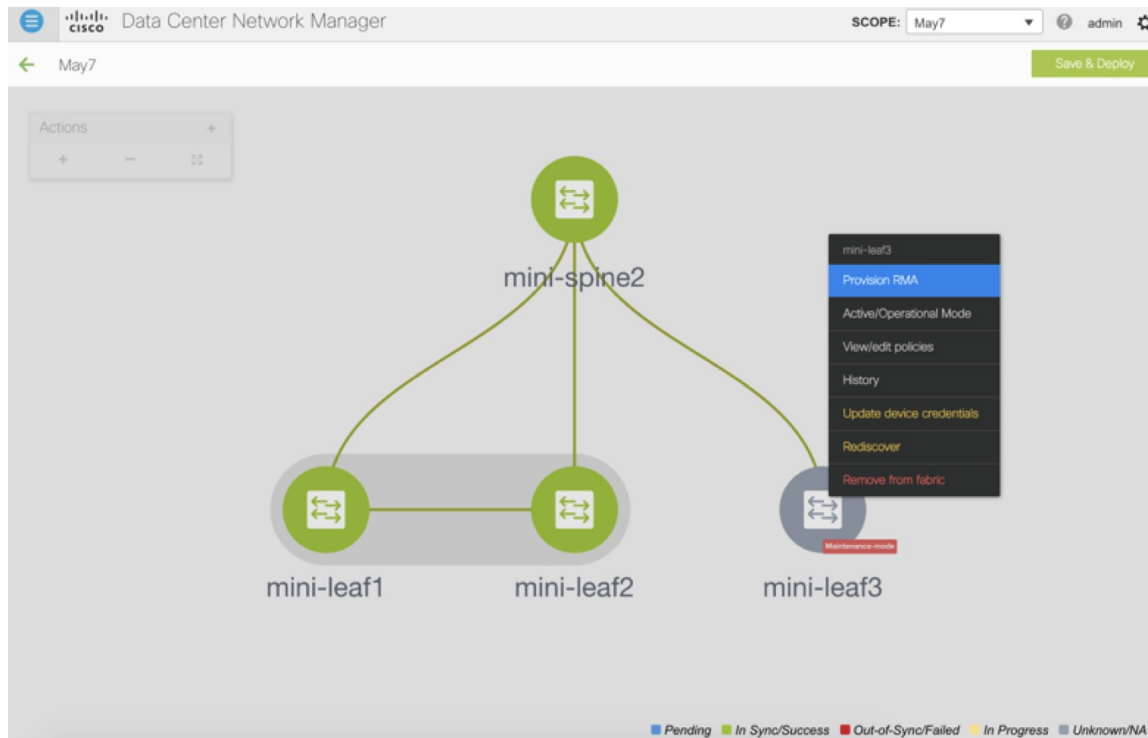
Procedure

- Step 1** Choose **Control > Fabric Builder**.
- Step 2** Click the Fabric where you want to perform RMA.
- Step 3** Move the device into maintenance mode. To move a device into maintenance mode, right-click on the device, and then choose **Modes > Maintenance Mode**.




Step 4 Physically replace the device in the network. Physical connections should be made in the same place on the replacement switch as they existed on the original switch.

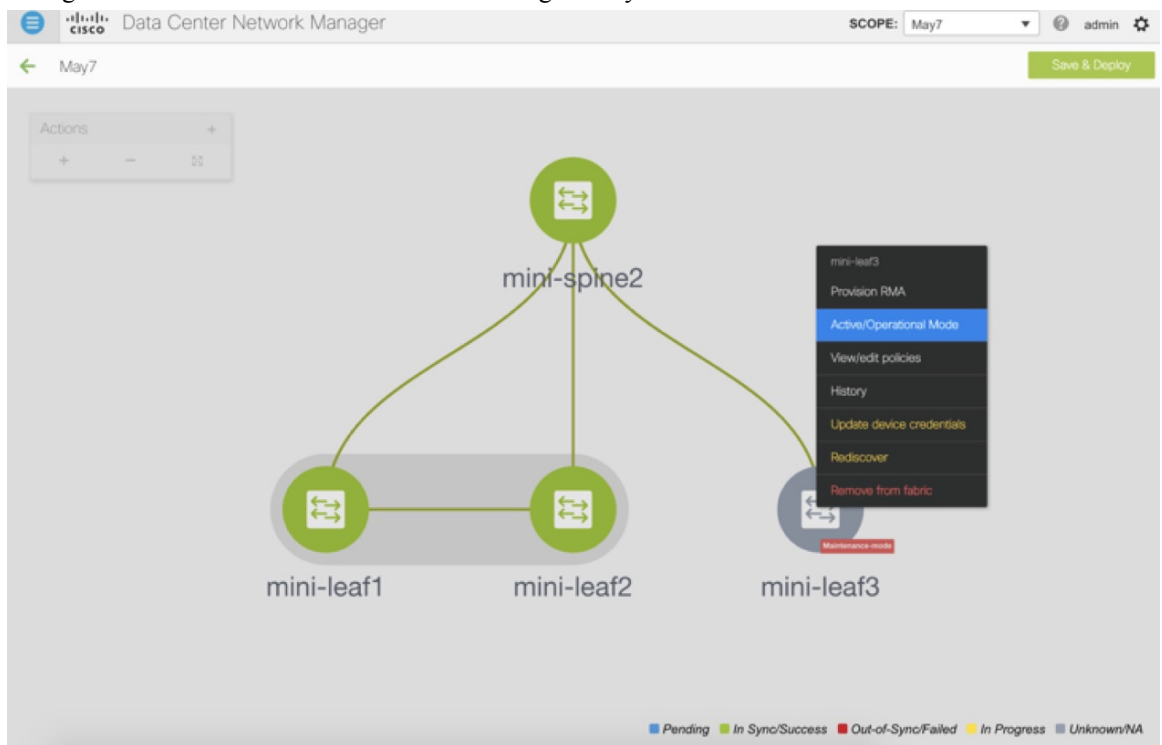
Step 5 Provision RMA flow and select the replacement device.



Step 6 The Provision RMA UI will show the replacement device 5-10 minutes after it is powered on.



Step 7 Select the correct replacement device and click **Swap Switch**. This begins POAP with the full “expected” configuration for that device. Total POAP time is generally around 10-15 minutes.

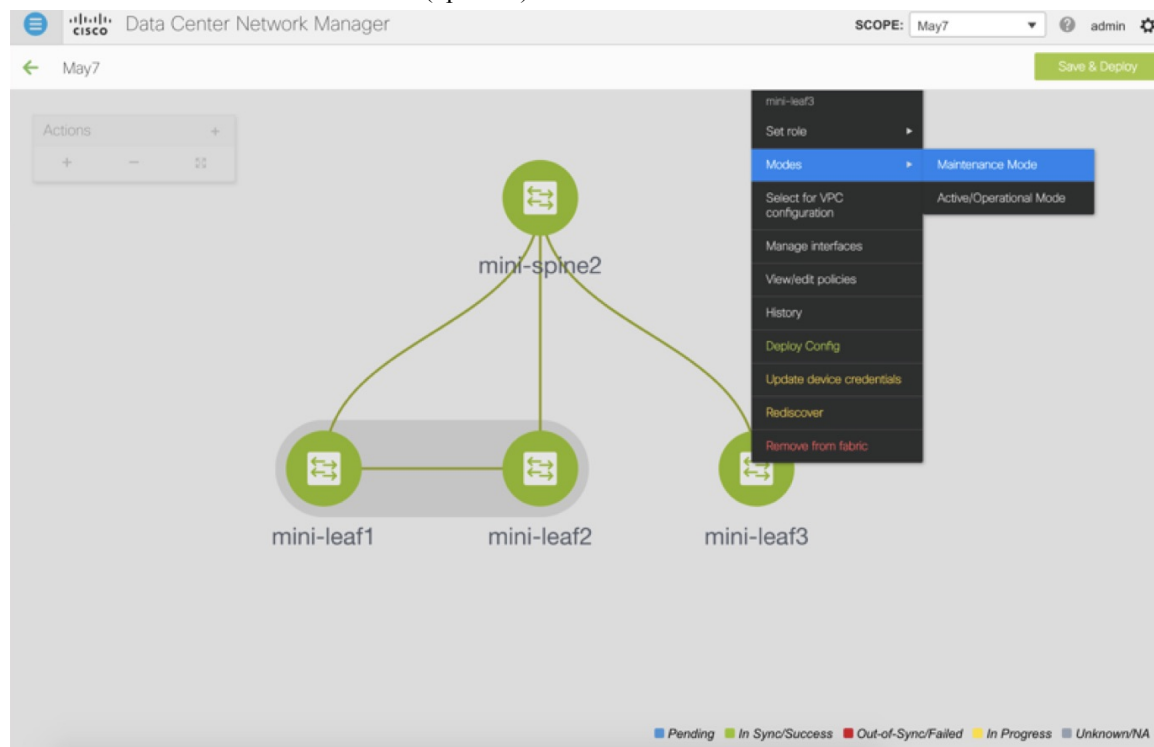


Manual RMA Flow

Use this flow when “Bootstrap” is not possible (or not desired), including cases that are *IPv6 only* for the initial Cisco DCNM 11.0(1) release.

Procedure

Step 1 Place the device in maintenance mode (optional).

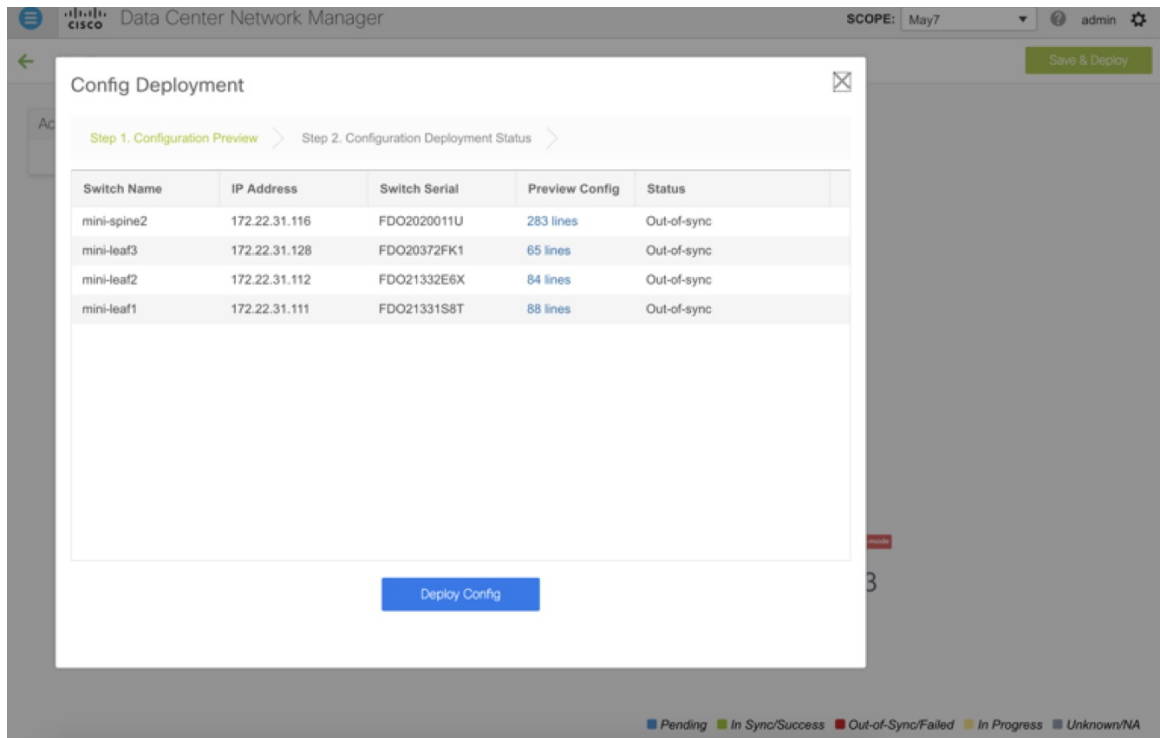


Step 2 Physically replace the device in the network.

Step 3 Log in through Console and set the Management IP and credentials.

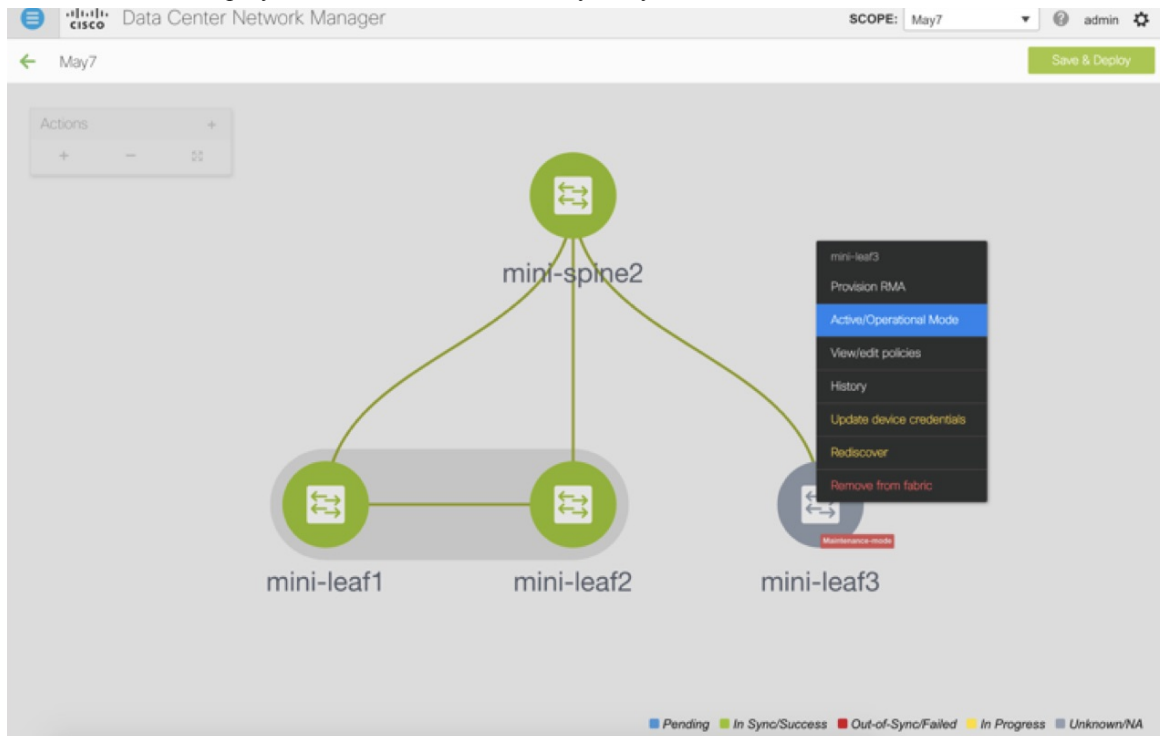
Step 4 The Cisco DCNM rediscovers the new device (or you can manually choose **Discovery > Rediscover**).

Step 5 Deploy the expected configuration using **Deploy**.



Step 6 Depending on the configuration, if breakout ports or FEX ports are in use, you have to deploy again to completely restore the configuration.

Step 7 After a successful deployment, and the device is “In-Sync,” you must move the device back to Normal Mode.



Custom Maintenance Mode Profile Policy

RMA for User with Local Authentication



Note This task is only applicable to non-POAP switches.

Use the following steps to perform RMA for a user with local authentication:

Procedure

- Step 1** After the new switch comes online, SSH into the switch and reset the local user passwords with the cleartext password using the “username” command. Reset the local user passwords to resync the SNMP password. The password is stored in the configuration file in a nontransferable form.
- Step 2** Wait for the RMA to complete.
- Step 3** Update Cisco DCNM switch_snmp_user policy for the switch with the new SNMP MD5 key from the switch.
-

Interfaces

The Interfaces option displays all the interfaces that are discovered for the switch, Virtual Port Channels (vPCs), and intended interfaces missing on the device.

You can use the following functions:

- Create, deploy, view, edit and delete a port channel, vPC, Straight-through FEX, Active-Active FEX, loopback, and subinterface.



-
- Note**
- The following features are unsupported for the brownfield migration of switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images:
 - FEX on switches other than Cisco Nexus 9300 Series switches and Cisco Nexus 9500 Series switches with X9500 line cards
 - AA-FEX

For information about the platform support for FEX, refer to your platform and NX-OS documentation to check the feature compatibility.

- To edit interfaces associated with fabric links such as intra-fabric links and inter-fabric links, see [Editing Interfaces Associated with Links, on page 222](#).
 - The **flowcontrol** or **priority-flow-control** config is not supported for HIF ports or PO with HIF ports as members.
-

- Create tunnel interfaces for Cisco Cloud Services Router 1000v Series (Cisco CSR 1000v Series).
- Create breakout and unbreakout ports.
- Shut down and bring up interfaces.
- Rediscover ports and view interface configuration history.
- Apply host policies on interfaces and vPCs. For example, int_trunk_host_11_1, int_access_host_11_1, and so on.
- View interface information such as its admin status, operation status, reason, policy, speed, MTU, mode, VLANs, IP/Prefix, VRF, port channel, and the neighbor of the interface.

**Note**

- The **Neighbor** column provides details of connected switches that are discovered, intent links, and Virtual Machine Manager (VMM) connectivity. You can navigate to the **Switch** dashboard of the corresponding switch by clicking it. However, intent links and VMM links aren't hyperlinked and you cannot navigate to the corresponding **Switch** dashboard.
- Click the graph icon in the Name column to view the interface performance chart for the last 24 hours. However, note that performance data for VLAN interfaces that are associated with overlay networks is not displayed in this chart.

The **Status** column displays the following statuses of an interface:

- Blue: Pending
 - Green: In Sync/Success
 - Red: Out-of-Sync/Failed
 - Yellow: In Progress
 - Grey: Unknown/NA
- If an interface is created out-of-band, you need to perform fabric resync or wait for Config Compliance polling before this interface can be deleted. Otherwise, Config Compliance does not generate the correct diff.

However, you cannot add or edit interfaces for ASR 9000 Series Routers and Arista switches.

You can filter and view information for any of the given fields (such as Device Name). The following table describes the buttons that appear on this page.

**Note**

- Ensure that appropriate configurations are deployed through the Fabric Builder option before deploying from the Interfaces option, including proper vPC pair configurations. If you add or edit an interface before fabric deployment, the configuration may fail on the device.
- You can also manage interfaces from the Fabric Builder topology screen. Right click the switch and on the Manage Interfaces option. You can manage the interfaces per switch. If the switch is part of a vPC Pair, then interfaces from both peers are displayed on the page.
- Deploy any underlays including vPC Pairing in the fabric before deploying any configurations from the interface manager.

Field	Description
Add	Allows you to add a logical interface such as a port channel, vPC, Straight-through FEX, Active-Active FEX, loopback and subinterface.
Breakout, Unbreakout	Allows you to <i>breakout</i> an interface or unbreakout interfaces that are in <i>breakout</i> state.
Edit	Allows you to edit and change policies that are associated with an interface.
Delete	Allows you to delete a logical interface that is created from the Interfaces screen. An interface having a policy that is attached from an overlay and underlay cannot be deleted.
No Shutdown	Allows you to enable an interface (no shutdown or admin up).
Shutdown	Allows you to shut down the interface.
Show	Allows you to display the interface show commands. A show command requires show templates in the template library.
Rediscover	Allows you to rediscover or recalculate the compliance status on the selected interfaces.
Interface History	Allows you to display the interface deployment history details.
Deploy	Allows you to deploy or redeploy saved interface configurations.

The following table describes the different user roles and the operations these roles support in the **Interfaces** window from Cisco DCNM Release 11.4(1).

Operations	User Roles		
	network-admin	network-operator	network-stager
Add	Save, Preview, Deploy	Blocked	Save, Preview
Breakout	Supported	Blocked	Blocked
Unbreakout	Supported	Blocked	Blocked
Edit	Save, Preview, Deploy	Preview	Save, Preview
Delete	Save, Preview, Deploy	Blocked	Save, Preview
Shutdown	Save, Preview, Deploy	Blocked	Save, Preview
No Shutdown	Save, Preview, Deploy	Blocked	Save, Preview
Show	Supported	Supported	Supported
Rediscover	Supported	Supported	Supported
Deploy	Preview, Deploy	Blocked	Blocked

The following table describes the new user role access-admin operations support in the host facing port of **Interfaces** window from Cisco DCNM Release 11.5(1).

Operations	User Roles
	access-admin
Add	Save, Preview, Deploy
Breakout	Blocked
Unbreakout	Blocked
Edit	Save, Preview, Deploy Note Access-admin user role cannot edit interfaces associated with link policy such as inter-fabric link or intra-fabric link for easy fabrics. The user role can edit interfaces for LAN classic fabrics.
Delete	Save, Preview, Deploy
Shutdown	Save, Preview, Deploy
No Shutdown	Save, Preview, Deploy
Show	Supported
Rediscover	Supported
Deploy	Preview, Deploy

From Cisco DCNM Release 11.4(1), you can disable deployments, or freeze, a fabric in DCNM as a network administrator. However, you cannot perform all actions when you freeze the fabric or if the fabric is in monitor mode.

The following table describes the actions you can perform when you freeze a fabric and when you enable the monitor mode for a fabric.

Operations	DCNM Mode	
	Freeze Mode	Monitor Mode
Add	Save, Preview	Blocked
Breakout	Blocked	Blocked
Unbreakout	Blocked	Blocked
Edit	Save, Preview	Blocked
Delete	Save, Preview	Blocked
Shutdown	Save, Preview	Blocked
No Shutdown	Save, Preview	Blocked
Show	Supported	Supported
Rediscover	Supported	Supported
Deploy	Blocked	Blocked

The buttons for the associated operations are grayed out accordingly.

If you perform admin operations (shutdown/no shutdown) on SVI, which is part of a config profile, successive **Save & Deploy** operations generate **no interface vlan** command.

For SVI with no policy, on performing admin operation, that is, shutdown/no shutdown command pushed from **Interface Manager**, **int_vlan_admin_state** policy is associated with the SVI.

For example, create and deploy the SVI from **switch_freeform**.

```
interface vlan1234
  description test
  no shutdown
  no ip redirects
  no ipv6 redirects
```

If you shutdown the SVI from interface manager, the **int_vlan_admin_state** policy is associated with the SVI.

Pending diff is shown as:

```
interface Vlan1234
  shutdown
  no ip redirects
```

```
no ipv6 redirects
description test
no shutdown
```

Remove the **no shutdown** CLI from the free-form config.

If the user has performed admin operation on SVI, device will have interface in running config. Therefore, post network detach **interface vlan** will be still present and interface will be discovered. You need to manually delete the interface from **Interface Manager**.

This section contains the following:

Adding Interfaces

To add the interfaces from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Interfaces**.

You see the **Scope** option at the top right. If you want to view interfaces for a specific fabric, select the fabric window from the list.

Step 2 Click **Add** to add a logical interface.

The **Add Interface** window appears.

Step 3 In the **Type** drop-down list, choose the type of the interface.

Valid values are Port Channel, virtual Port Channel (vPC), Straight-through (ST) FEX, Active-Active (AA) FEX, Loopback, Subinterface, Tunnel Ethernet, and Switch Virtual Interface (SVI). The respective interface ID field is displayed when you select an interface type.

- When you create a port channel through DCNM, add interfaces of the same speed. A port channel that is created from interfaces of varying speeds won't come up. For example, a port channel with two *10 Gigabit Ethernet* ports is valid. However, a port channel with a *10-Gigabit Ethernet* + *25-Gigabit Ethernet* port combination isn't valid.
- To add vPC hosts, you must designate vPC switches in the fabric topology (through the Fabric Builder) and deploy vPC and peer-link configurations using the **Save and Deploy** option. Once the vPC pair configurations are deployed, it appears in the Select a vPC pair drop-down box.

You can create a vPC using the **int_vpc_trunk_host_11_1** policy.


- When adding a subinterface, you must select a routed interface from the interface table before clicking the Add button.
- You can preprovision Ethernet interfaces in the Interface window. This preprovisioning feature is supported in Easy, eBGP, and External fabrics. For more information, see [Pre-provisioning an Ethernet Interface, on page 41](#).

Step 4 In the **Select a Device** field, choose the device.

Devices are listed based on the fabric and interface type. External fabric devices aren't listed for ST FEX and AA FEX. In the case of vPC or Active to Active FEX, select the vPC switch pair.

- Step 5** Enter the ID value in the respective interface ID field (**Port-channel ID**, **vPC ID**, **Loopback ID** and **Subinterface ID**) that is displayed, based on the selected interface.
- You can override this value. The new value is used only if it's available in the Resource Manager pool. Else, it results in an error.
- Step 6** In the **Policy** field, select the policy to apply on an interface.
- The field only lists the Interface Python Policy with tag `interface_edit_policy` and filtered based on the interface type.
- You must not create a **_upg** interface policy. For example, you shouldn't create a policy using the **vpc_trunk_host_upg**, **port_channel_aa_fex_upg**, **port_channel_trunk_host_upg**, and **trunk_host_upg** options.
- Note** The policies are filtered based on the interface type you choose in the **Type** drop-down list and the device you choose in the **Select a Device** drop-down list.
- Step 7** Enter values in the required fields under the **General** tab.
- The fields vary according to the interface type you choose.
- Note** From Cisco DCNM, Release 11.5(1) you can mirror the configurations of Peer-1 on Peer-2 while creating a vPC. When you check the **Enable Config Mirroring** check box, the Peer-2 fields will be grayed out. The configurations that you enter in the Peer-1 fields will be copied to Peer-2 fields.
- Step 8** Click **Save** to save the configurations.
- Note** To apply QoS policies on the interface, create the interface freeform with references accordingly.
- Only saved configurations are pushed to the device. While adding the interface, you can only modify the policy attribute after the first save. If you try to use an ID that is already used, you encounter the *Resource could not be allocated* error.
- Step 9** (Optional) Click the **Preview** option to preview the configurations to be deployed.
- Step 10** Click **Deploy** to deploy the specified logical interface.
- The newly added interface appears in the screen.
-

Breakout

Click the drop-down arrow next to the **Breakout** icon  to display a list of the available breakout options. The available options are **10g-4x**, **25g-4x**, **50g-2x**, **50g-4x**, **100g-2x**, **100g-4x**, **200g-2x**, and **Unbreakout**. Choose the required option.

Editing Interfaces

To edit the interfaces from the Cisco DCNM Web UI, perform the following steps:



Note The **Edit Interface** allows you to change the policy and add or remove an interface from a port channel or vPC.

Procedure

Step 1 Choose **Control > Interfaces**.

You can break out and unbreak out an interface by using the breakout option at the top left part of the screen.

Step 2 Select the interface check box to edit an interface or vPC.

Select corresponding check boxes for editing multiple interfaces. You cannot edit multiple port channels and vPC. You cannot edit interfaces of different types at the same time.

Step 3 Click **Edit** to edit an interface.

The variables that are shown in the **Edit Configuration** window are based on the template and its policy. Select the appropriate policy. Preview the policy, save it and deploy the same. This window lists only Interface Python Policy with the tag *interface_edit_policy* and filtered based on the interface type.

In a vPC setup, the two switches are in the order the switch names are displayed in the edit window. For example, if Switch Name is displayed as *LEAF1:LEAF2*, then Leaf1 is peer switch one and Leaf2 is peer switch two.

During overlay network deployment on switches, the network can be associated with trunk interfaces. The trunk interface to network association is reflected in the **Interfaces** screen. You can update such interfaces.

For interface policies that are not created from the **Control > Interfaces** screen, you can edit some configurations but not change the policy itself. The policy and fields that cannot be edited are grayed out.

The following are some examples of policies that cannot be edited:

- Loopback interface policies - The *int_fabric_loopback_11_1* policy is used to create a loopback interface. You can edit the loopback IP address and description but not the *int_fabric_loopback_11_1* policy instance.
 - Fabric underlay network interface policies (*int_fabric_num_11_1*, for example) and fabric overlay network interface (NVE) policies.
 - Policies associated with port channels and member ports of port channels, including the port channels and member ports associated with a vPC.
 - SVIs created during network and VRF creation. The associated VLANs appear in the interfaces list.
-

Editing Interfaces Associated with Links

There are two types of links, namely intra-fabric links and inter-fabric links. As the name implies, intra-fabric links are set up between devices within the same Easy fabric and are typically used for spine-leaf connectivity. Inter-fabric links are set up between the Easy fabric, and typically other external or Easy fabrics. They are used for external WAN and/or DCI connectivity. A policy is associated with each link that effectively states the configuration that is applied to both ends of the link. In other words, the link policy becomes the parent

of the individual child interface policies that are associated with the two interfaces that form the link. In this scenario, you must edit the link policy to edit the interface policy fields such as description, IP address, and any per interface freeform config. The following procedure shows how to edit the interfaces associated with links:

Procedure

- Step 1** Choose **Control > Fabric Builder**, and select the fabric containing the link.
- Step 2** Click **Tabular view** in the **Actions** panel.
- A window with the **Switches** and **Links** tabs appears.
- Step 3** Click the **Links** tab.
- Step 4** Select the link that you want to edit and click the **Update Link** icon.

Update the link based on your requirements and click **Save**.

Deleting Interfaces

To delete the interfaces from the Cisco DCNM Web UI, perform the following steps:



Note This option allows you to delete only logical ports, port channels, and vPCs. You can delete the interface if it does not have overlay or underlay policy attached.

When a port channel or vPC is removed, the corresponding member ports get the default policy associated. The Default Policy can be configured in `server.properties` file.

Procedure

Step 1 Choose **Control > Interfaces**.

Step 2 Select the interfaces.

Step 3 Click **Delete**.

You cannot delete logical interfaces created in the fabric underlay.

Step 4 Click **Save**.

Step 5 (Optional) Click **Preview** to view all the changes before deleting the interface.

The deletion will be highlighted in red colour with strikethrough under the **Expected Config** tab.

Preview Configuration
✕

Select a Switch: Select an Interface:

Pending Config Expected Config Current Config

```
interface Port-channel501
```

Step 6 Click **Deploy** to delete the interface.

Shutting Down and Bringing Up Interfaces

To shut down and bring up the interfaces from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Interfaces**.

Step 2 Select the interfaces that you want to shut down or bring up.

Step 3 Click **Shutdown** to disable the selected interfaces. For example, you may want to isolate a host from the network or a host that is not active in the network.

A confirmation window appears where you can save, preview, and deploy the changes. Click **Save** to preview or deploy the changes.

Step 4 Click **No Shutdown** to bring up the selected interfaces.

A confirmation window appears where you can save, preview, and deploy the changes. Click **Save** to preview or deploy the changes.

Viewing Interface Configuration

To view the interface configuration commands and execute them from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Interfaces**.

Select the interface whose configurations you want to view.

Step 2 In the **Interface Show Commands** window, select the action from the **Show** drop-down box and click **Execute**. The interface configurations are displayed in the **Output** section, at the right of the screen.

For Show commands, you must have corresponding **show** templates for interface or interface sub types like port channel or vPC, defined in the **Template Library**.

Rediscovering Interfaces

To rediscover the interfaces from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Interfaces**.

Step 2 Select the interfaces that you want to rediscover.

Step 3 Click **Rediscover** to rediscover the selected interfaces. For example, after you edit or enable an interface, you can rediscover the interface.

Viewing Interface History

To view the interface history from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Control > Interfaces**.

Step 2 Select the interface.

- Step 3** Click **Interface History** to view the configuration history on the interface.
- Step 4** Click **Status** to view each command that is configured for that configuration instance.
-

Deploying Interface Configurations

To deploy the interface configuration from the Cisco DCNM Web UI, perform the following steps:

Procedure

- Step 1** Choose **Control > Interfaces**.
- Step 2** Choose an interface you want to deploy.
- Note** You can select multiple interfaces and deploy pending configurations.
- Step 3** Click **Deploy** to deploy or redeploy configurations that are saved for an interface.

After you deploy the interface configuration, the interface status information is updated. However, the overall switch-level state may be in the pending state, which is in blue. The overall switch-level state goes to the pending state whenever there is a change in intent from any module, such as interface, link, policy template update, top-down, or so on. In the pending state, a switch may have pending configurations or switch-level recomputation. The switch-level recomputation occurs when:

- You preview or deploy for the switch
- During a save and deploy
- During hourly sync

Preview or deploy the switches to review their state and to understand the root cause of their pending state. Save and deploy for a fabric-wide recomputation.

Click **Preview** to preview the configurations before you click **Deploy**.

Creating External Fabric Interfaces

You can add and edit port channel, vPC, subinterface, and loopback interfaces for external fabric devices. You cannot add Straight-through FEX and Active-Active FEX functions.

The Breakout port function is only supported for the Cisco Nexus 9000, 3000, and 7000 Series Switches in the external fabric.

When you add an interface to an external fabric device, the Resource Manager is not in sync with the device. So, ensure that the value populated in the ID field (Port-channel ID, vPC ID, Loopback ID, etc) is not previously configured on the switch.

If you want to configure a portchannel in the external fabric, you should add and deploy the **feature_lacp** policy on the switches where the portchannel will be configured.

Add Policy



* Priority (1-1000):

* Policy:

feature_lACP

Variables:

Save

Cancel

When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on its switches. If you click **Save & Deploy** in the fabric topology screen, it displays an error message. However, the following settings (available when you right-click the switch icon) are allowed:

vPC pairing - You can designate a vPC switch pair, but it is only for reference.

View/edit policy - You can add a policy but you cannot deploy it on the switch.

Manage interfaces – You can only create intent for adding interfaces. If you try to deploy, edit, or delete interfaces, it results in an error message.

Interface Groups

From Cisco DCNM Release 11.5(1), you can create an interface group that allows grouping of host-facing interfaces at a fabric level. Specifically, you can create an interface group for physical Ethernet interfaces, L2 port-channels, and vPCs. You can attach or unattach multiple overlay networks to the interfaces in an interface group.

Guidelines

- Interface groups are only supported for the fabrics with the **Easy_Fabric_11_1** template.
- An interface group is specific to a fabric. For example, consider two fabrics: Fab1 and Fab 2. The interface group IG1 in Fab1 isn't applicable to Fab 2.
- An interface group can only have interfaces of a certain type. For example, you need three separate interface groups if you want to group three types of interfaces such as IG1 for physical Ethernet trunk interfaces, IG2 for L2 trunk port-channels, and IG3 for vPC host trunk ports.
- An interface group can also be created using preprovisioned interfaces.
- Interface groups are limited to switches with the leaf role. They aren't supported for other roles such as Border, BGW, and other related variants.

- For L2 port-channels and vPCs that are part of an interface group, they can't be deleted until they are de-associated from the interface group even if there are no networks associated with the interface group. Similarly, a trunk port that has no overlay networks but is part of an IG can't be converted to an access port. In other words, you can't change policies for interfaces that are part of an interface group. However, you can edit certain fields for policies.
- For L4-L7 services configuration on leaf switches, trunk ports that are used for services attachment can't be part of interface groups.
- When you perform a per fabric backup of an easy fabric, if there are interface groups created in that fabric, all the associated interface group state is backed up.
- If an easy fabric contains an interface group, then this fabric can't be imported into the MSO. Similarly, if an easy fabric has been added to the MSO, you can't create interface groups for interfaces that belong to switches in the easy fabric.
- The **Interface Group** button is enabled only for Admin and Stager users. For all other users, this button is disabled.
- The **Interface Group** button is disabled in the following circumstances:
 - Select **Data center** from the **SCOPE** drop-down list.
 - Select a fabric without any switches.
 - Select any other interface apart from vPC, Port-channel, and Ethernet.
 - If the interface has a policy attached from another source, for example:
 - If the interface is member of a port-channel or vPC.
 - If the port-channel is member of vPC.
 - If the interface has a policy from underlay or links.

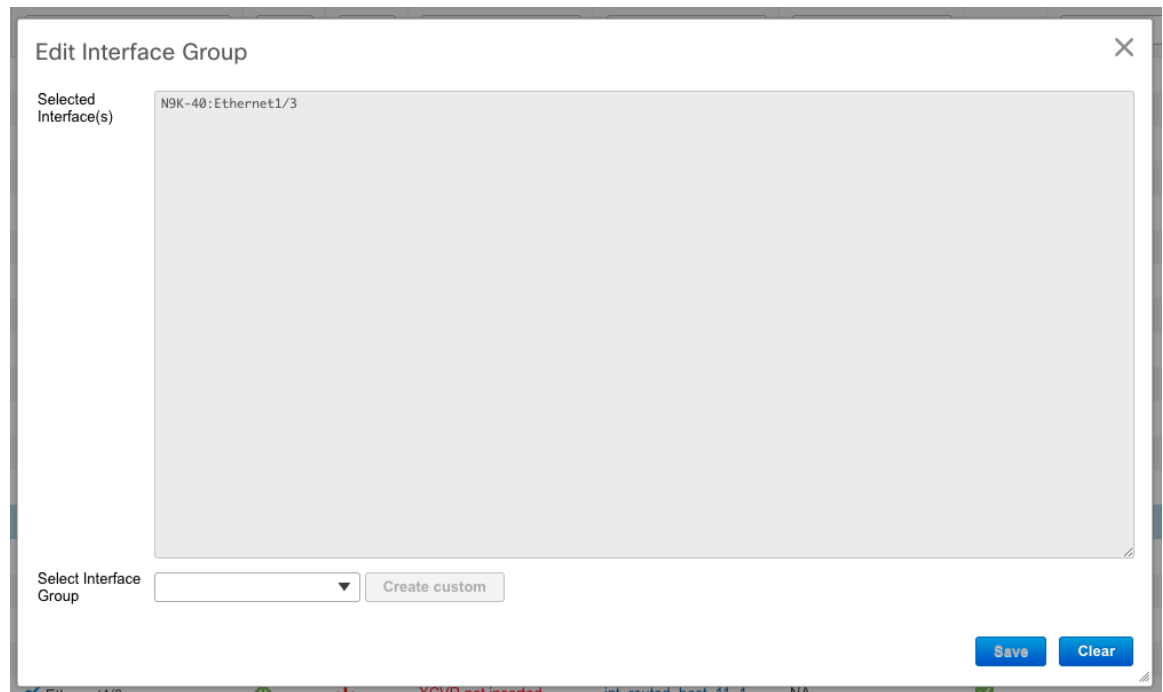


Note If you select different types of interfaces, the **Interface Group** button is enabled. However, when you try to create or save different types of interfaces to an interface group, an error is displayed.

Creating an Interface Group

Procedure

- Step 1** From DCNM, navigate to **Control > Fabrics > Interfaces**.
- Step 2** From the **SCOPE** drop-down list, select a fabric.
- Step 3** Select the interfaces that have to be grouped and click **Interface Group**.



- Step 4** 4. In the **Edit Interface Group** window, create a custom interface group by entering an interface group name in the **Select Interface Group** field and click **Create custom**. An interface group name can have a maximum length of 64 characters.

If you have already created an interface group, select it from the **Select Interface Group** drop-down list. Also, if an interface is already part of an interface group, you can move it to a different interface group by selecting the new group from the **Select Interface Group** drop-down list.

Note An interface can belong to only a single interface group.

You can create interface groups from either the **Interfaces** window or the **Networks** window. For more information, see [Attaching Networks to an Interface Group, on page 230](#).

- Step 5** Click **Save**.

In the **Interfaces** window, you can see the interface group name under the **Interface Group** column.

Removing Interfaces from an Interface Group

Procedure


- Step 1** From DCNM, navigate to **Control > Fabrics > Interfaces**.
- Step 2** From the **SCOPE** drop-down list, select a fabric.
- Step 3** Select the interfaces to disassociate from an interface group and click **Interface Group**.
- Step 4** In the **Edit Interface Group** window, make sure that nothing is selected in the **Select Interface Group** drop-down list, and click **Clear**.

A dialog box pops up asking whether you want to clear all the associated interfaces. Click **Yes** to proceed. Note that if there are any networks attached to these interfaces, they are detached as well when you click **Clear**.

Attaching Networks to an Interface Group

Procedure

- Step 1** From DCNM, navigate to **Control > Fabrics > Networks**.
- Step 2** From the **SCOPE** drop-down list, select a fabric.
- Step 3** In the **Networks** window, select the networks that you need to attach to an interface group and click **Interface Group**.
- Note**
- An overlay network can belong to multiple interface groups.
 - You can select only the networks with a VLAN ID. Otherwise, an appropriate error message is displayed.
- Step 4** In the **Interface Groups** window, you can perform the following:
- Select an existing interface group from the **Select Interface Group** drop-down list and click **Save**.



For example, you select three networks and the interface group **test**, and click the **Save** button, the following operations are performed in the background:

- DCNM retrieves interfaces that are part of the interface group **test**.
- DCNM determines that three networks are added to the interface group **test**. Therefore, it autoattaches these networks to all the interfaces that are part of the interface group **test**.

- c. For each interface, DCNM pushes the “**switchport trunk allowed vlan add xxxx**” command three times for each selected network.

Note DCNM ensures that there’s no duplicate configuration intent.

If you click the **Clear** button, DCNM pushes “**switchport trunk allowed vlan remove xxx**” config intent.

- Create a custom interface group by entering an interface group name in the **Select Interface Group** field and click **Create custom**. Click **Save**.

If you choose this option, make sure to add interfaces to this Interface Group in the **Interfaces** window. As a result, DCNM performs the following operations:

- a. Removes all existing overlay networks that don’t belong to the interface group from these interfaces.
- b. Adds new overlay networks to these interfaces that are part of the interface group but not yet attached to these interfaces.

For more information about associating interfaces to interface groups, see [Creating an Interface Group, on page 228](#).

Step 5 Click **Continue** and click **Save & Deploy** to deploy the selected networks on the switches.

Unattaching a Network from an Interface Group

This procedure shows how to unattach a network from an interface group in the Networks window. Also, you can unattach networks when you remove an interface from an interface group in the **Interfaces** window. For more information, see *Removing Interfaces from an Interface Group*.

Procedure

- Step 1** 1. From DCNM, navigate to **Control > Fabrics > Networks**.
 - Step 2** From the **SCOPE** drop-down list, select a fabric.
 - Step 3** In the **Networks** window, select the networks that you need to unattach to an interface group and click **Interface Group**.
 - Step 4** In the **Interface Groups** window, select the interface group from the **Select Interface Group** drop-down list and click **Clear** to unattach a network.
 - Step 5** (Optional) Navigate to **Control > Fabrics > Interfaces**.
Under the **Overlay Network** column, you can see the unattached network in the red color for the corresponding interface. Click the network to view the expected config that is struck through.
 - Step 6** Navigate to the **Fabric Builder** or **Networks** window, and click **Save & Deploy**.
-

Deleting an Interface Group

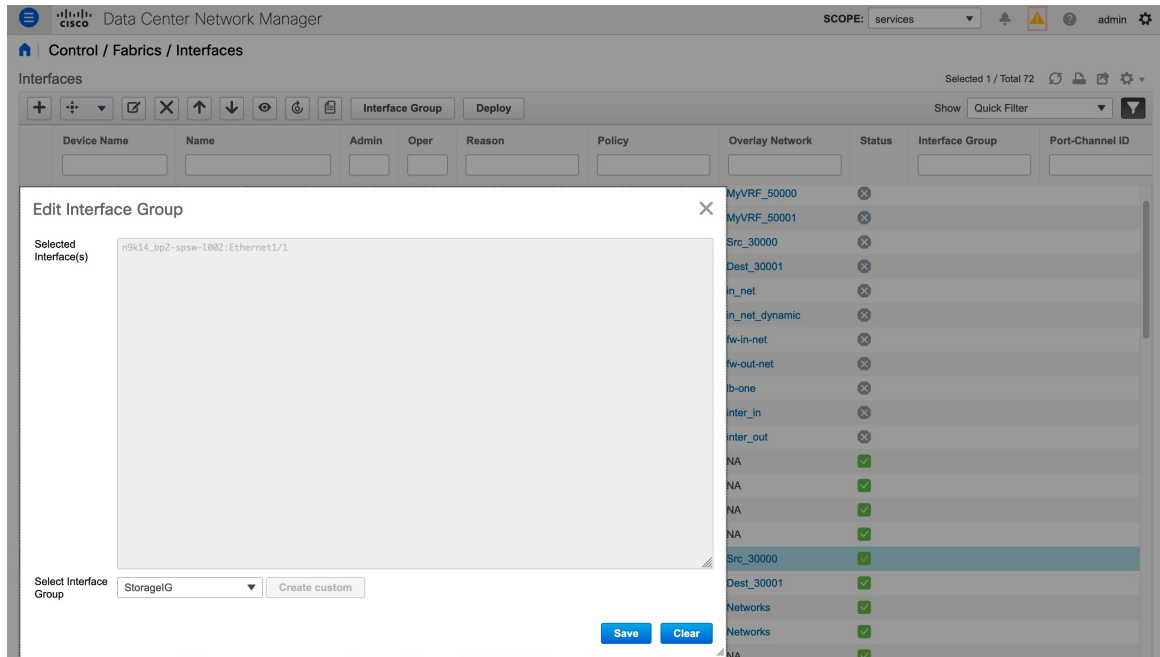
An interface group is automatically deleted when it’s not in use. DCNM performs an implicit delete of an interface group if there are no interfaces and no networks mapped to the interface group. This check is

performed whenever you click the **Clear** button in the **Edit Interface Group** window. There may be exception scenarios where you need to clean up the interface groups explicitly.

For example, you create an interface group **storageIG** and add an interface to it. Later, you want to change the interface mapping to another group. Therefore, you select the interface and click **Interface Group** to open the **Edit Interface Group** window. Select the other interface group named **diskIG**. Now, the **storageIG** interface group doesn't have any associated member interfaces or networks. In this case, perform the following steps:

Procedure

- Step 1** Select an interface that doesn't belong to an interface group.
- Step 2** Click **Interface Group** to open the **Edit Interface Group** window.
- Step 3** Select the **StorageIG** interface group from the **Select Interface Group** drop-down list.



- Step 4** Click **Clear**.

Creating and Deploying Networks and VRFs

The steps for overlay networks and VRFs provisioning are:

1. Create networks and VRFs for the fabric.
2. Deploy the networks and VRFs on the fabric switches.



Note The undeployment and deletion of overlay networks and VRFs are explained after the explanation of deployment. Finally, creation of external fabrics and fabric extensions from VXLAN to external fabrics are documented.

For information about creating interface groups and attaching networks, see [Interface Groups, on page 227](#).

You can navigate to the networks and VRFs window through any of the following options:

- From the home page: Click the **Networks & VRFs** button in the Cisco DCNM Web UI landing page.
- From the Control menu: From the home page of the Cisco DCNM Web UI, choose **Control > Fabrics > Networks** to navigate to the **Networks** window. Choose **Control > Fabrics > VRFs** to navigate to the **VRFs** window.
- From a fabric topology window: Right-click anywhere in the fabric topology window. Choose **Overlay View > VRF View** or **Overlay View > Network View**, accordingly. This option is applicable only for switch fabrics, easy fabrics, and MSD fabrics.

You can toggle between the network view and VRF view in both the windows by clicking the **VRF View** or **Network View** button. When you are in the networks or VRFs window, ensure you choose the appropriate fabric from the **Scope** drop-down list before you create any networks or VRFs.

Viewing Networks and VRFs for a Fabric

- Click **Control > Networks** from the main menu.

The **Networks** screen comes up. The **SCOPE** drop down box (at the top right part of the screen) lists all fabrics managed by the DCNM instance, in alphabetical order. You can choose the correct fabric from **SCOPE**. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

The screenshot shows the Cisco Data Center Network Manager interface. At the top, the breadcrumb navigation is "Network / VRF Selection > Network / VRF Deployment". The "SCOPE" dropdown is set to "bgp2". Below the breadcrumb, there are buttons for "VRF View" and "Continue". The main content area is titled "Fabric Selected: bgp2" and "Networks". It shows a table with one network entry selected:

Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
MyNetwork_30000	30000	NA			NA	

- Click **Control > VRFs** from the main menu.

The **VRFs** screen comes up. The **SCOPE** drop down box (at the top right part of the screen) lists all fabrics managed by the DCNM instance, in alphabetical order. You can choose the correct fabric from **SCOPE**. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

The screenshot shows the Cisco Data Center Network Manager interface. The top navigation bar includes the Cisco logo, the title "Data Center Network Manager", the "SCOPE" dropdown set to "bgp2", and the user "admin". The breadcrumb trail is "Network / VRF Selection > Network / VRF Deployment". There are "Network View" and "Continue" buttons. Below the breadcrumb, it says "Fabric Selected: bgp2". The main section is titled "VRFs" and shows "Selected 1 / Total 1". A toolbar contains icons for adding, editing, deleting, and refreshing. A "Show" dropdown is set to "All". A table lists VRFs:

<input type="checkbox"/>	VRF Name	VRF ID	Status
<input checked="" type="checkbox"/>	MyVRF_50000	50000	NA



Note The **Networks** or **VRFs** windows are applicable only for the Easy or MSD fabrics.

Creating Networks for the Standalone Fabric

1. Click **Control > Networks** (under **Fabrics** submenu).

The **Networks** screen comes up.

2. Choose the correct fabric from **SCOPE**. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

The screenshot shows the Cisco Data Center Network Manager interface. The top navigation bar includes the Cisco logo, the title "Data Center Network Manager", the "SCOPE" dropdown set to "bgp2", and the user "admin". The breadcrumb trail is "Network / VRF Selection > Network / VRF Deployment". There are "VRF View" and "Continue" buttons. Below the breadcrumb, it says "Fabric Selected: bgp2". The main section is titled "Networks" and shows "Selected 1 / Total 1". A toolbar contains icons for adding, editing, deleting, and refreshing. A "Show" dropdown is set to "All". A table lists networks:

<input type="checkbox"/>	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
<input checked="" type="checkbox"/>	MyNetwork_30000	30000	NA			NA	

3. Click the + button at the top left part of the screen (under **Networks**) to add networks to the fabric. The **Create Network** screen comes up. Most of the fields are autopopulated.

Create Network
✕

▼ Network Information

* Network ID

* Network Name

* VRF Name +

Layer 2 Only

* Network Template

* Network Extension Template

VLAN ID Propose VLAN ?

▼ Network Profile

Generate Multicast IP ⓘ Please click only to generate a New Multicast Group Address and override the default value!

General

Advanced

IPv4 Gateway/NetMask ⓘ example 192.0.2.1/24

IPv6 Gateway/Prefix L... ⓘ example 2001:db8::1/64,2001:db9::1/64

Vlan Name ⓘ if > 32 chars enable:system vlan long-nam

Interface Description ⓘ

MTU for L3 interface ⓘ 68-9216

IPv4 Secondary GW1 ⓘ example 192.0.2.1/24

IPv4 Secondary GW2 ⓘ example 192.0.2.1/24

Create Network

The fields in this screen are:

Network ID and **Network Name**: Specifies the Layer 2 VNI and name of the network. The network name should not contain any white spaces or special characters except underscore (_) and hyphen (-). The corresponding Layer 3 VNI (or VRF VNI) is generated along with VRF creation.

VRF Name: Allows you to select the Virtual Routing and Forwarding (VRF).

When no VRF is created, this field appears blank. If you want to create a new VRF, click the + button. The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:).

Layer 2 Only: Specifies whether the network is Layer 2 only.

Network Template: A universal template is autopopulated. This is only applicable for leaf switches.

Network Extension Template: A universal extension template is autopopulated. This allows you to extend this network to another fabric. The methods are VRF Lite, Multi Site, and so on. The template is applicable for border leaf switches and BGWs.

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

The VLAN ID default range is 2 to 3967. From DCNM Release 11.5(2), you can use a VLAN range greater than default value 3967. The reserved VLAN range must be set to a different range. In switch

command enter “**system vlan <vlan> reserve**”. Save the configuration to startup configuration and reload the switch for the new reserved VLAN range to reflect.

From Cisco DCNM Web UI, choose **Administration > DCNM Server > Server Properties**, enter the value as 4094 for **RM.TOP_DOWN_NETWORK_VLAN.MAX** and **RM.TOP_DOWN_VRF_VLAN.MAX**, click **Apply Changes** and then restart DCNM. Once the DCNM is up, you can create VRF and network using the VLAN value greater than 3967.

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

General tab

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



Note If the same IP address is configured in the IPv4 Gateway and IPv4 Secondary GW1 or GW2 fields of the network template, DCNM does not show an error, and you will be able to save this configuration. However, after the network configuration is pushed to the switch, it would result in a failure as the configuration is not allowed by the switch.

IPv6 Gateway/Prefix: Specifies the IPv6 address with subnet.

Specify the anycast gateway IP address for transporting the L3 traffic from a server belonging to MyNetwork_30000 and a server from another virtual network. By default the anycast gateway IP address is the same for MyNetwork_30000 on all switches of the fabric that have the presence of the network.

VLAN Name - Enter the VLAN name.

Interface Description: Specifies the description for the interface. This interface is a switch virtual interface (SVI).

MTU for the L3 interface - Enter the MTU for Layer 3 interfaces.

IPv4 Secondary GW1 - Enter the gateway IP address for the additional subnet.

IPv4 Secondary GW2 - Enter the gateway IP address for the additional subnet.

Advanced tab: Optionally, specify the advanced profile settings by clicking the **Advanced** tab:

ARP Suppression – Select the checkbox to enable the ARP Suppression function.

Ingress Replication - The checkbox is selected if the replication mode is Ingress replication.



Note Ingress Replication is a read-only option in the Advanced tab. Changing the fabric setting updates the field.

Multicast Group Address- The multicast IP address for the network is autopopulated.

Multicast group address is a per fabric instance variable. The number of underlay multicast groups supported is only 128. If all networks are deployed on all switches, you need not use a different multicast group per L2 VNI or a network. Therefore, multicast group for all networks in a fabric remains same. If a new multicast group address is required, you can generate it by clicking the **Generate Multicast IP** button.

DHCPv4 Server 1 - Enter the DHCP relay IP address of the first DHCP server.

DHCPv4 Server 2 - Enter the DHCP relay IP address of the next DHCP server.

DHCPv4 Server VRF- Enter the DHCP server VRF ID.

Loopback ID for DHCP Relay interface (Min:0, Max:1023) - Specifies the loopback ID for DHCP relay interface.

Routing Tag – The routing tag is autopopulated. This tag is associated with each gateway IP address prefix.

TRM enable – Select the check box to enable TRM.

For more information, see [Overview of Tenant Routed Multicast, on page 148](#).

L2 VNI Route-Target Both Enable - Select the check box to enable automatic importing and exporting of route targets for all L2 virtual networks.

Enable L3 Gateway on Border - Select the check box to enable a Layer 3 gateway on the border switches.

A sample of the Create Network screen is given below.

▼ Network Profile

Please click only to generate a New Multicast Group Address and override the default value!

General	<p>IPv4 Gateway/NetMask <input type="text" value="20.10.1.1/24"/> ? example 192.0.2.1/24</p> <p>IPv6 Gateway/Prefix <input type="text"/> ? example 2001:db8::1/64</p> <p>Vlan Name <input type="text" value="Drill"/> ?</p> <p>Interface Description <input type="text"/> ?</p> <p>MTU for L3 interface <input type="text"/> ? [68-9216]</p> <p>IPv4 Secondary GW1 <input type="text" value="20.10.2.1/24"/> ? example 192.0.2.1/24</p> <p>IPv4 Secondary GW2 <input type="text" value="20.10.3.1/24"/> ? example 192.0.2.1/24</p>
Advanced	

Network Profile

Generate Multicast IP

ⓘ Please click only to generate

General

Advanced

ARP Suppression ⓘ ARP

Ingress Replication ⓘ Rea

Multicast Group Address

*** DHCPv4 Server 1**

*** DHCPv4 Server VRF**

DHCPv4 Server 2

DHCPv4 Server2 VRF

- Click **Create Network**. A message appears at the bottom right part of the screen indicating that the network is created.

The new network appears on the **Networks** page that comes up.

Fabric Selection > Network / VRF Selection > Network / VRF Deployment > VRF View | Continue

Fabric Selected: Standalone

Networks Selected 1 / Total 1

+
✎
✕
↻
↺
Show All

	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
<input checked="" type="checkbox"/>	MyNetwork_30000	30000	MyVRF_50000	20.10.1.1/24		NA	

The Status is *NA* since the network is created but not yet deployed on the switches. Now that the network is created, you can create more networks if needed and deploy the networks on the devices in the fabric.

Export and Import Network Information

You can export network information to a .CSV file. The exported file contains information pertaining to each network, including the fabric it belongs to, the associated VRF, the network templates used to create the network, and all other configuration details that you saved during network creation.

In the Networks screen, click the Export icon to export network information as a .CSV file.

Networks

A	B	C	D
fabric	vrf	networkName	networkId
Standalone	MyVRF_50000	MyNetwork_30000	30000
Standalone	MyVRF_50000	MyNetwork_30001	30001

You can use the exported .CSV file for reference or use it as a template for creating new networks. To import networks, do the following:

1. Update new records in the .CSV file. Ensure that the networkTemplateConfig field contains the JSON Object. A message at the bottom right part of the screen displays errors and success messages. This screenshot depicts two new networks being imported.

A	B	C	D	E	F	G	H	I	J	K
fabric	vrf	networkName	networkId	networkTemplate	networkExtensionTemplate	networkTemplateConfig				
Standalone	MyVRF_50000	MyNetwork_30002	30002	Default_Network_Universal	Default_Network_Extension_Universal	["suppressArp":"false","secondaryGW2":"","secondaryGW1":"",""]				
Standalone	MyVRF_50000	MyNetwork_30003	30003	Default_Network_Universal	Default_Network_Extension_Universal	["suppressArp":"false","secondaryGW2":"","secondaryGW1":"",""]				

2. In the Networks screen, click the Import icon and import the .CSV file into DCNM.

You can see that the imported networks are displayed in the Networks screen.

Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
MyNetwork_30000	30000	MyVRF_50000	20.10.1.1/24		NA	
MyNetwork_30001	30001	MyVRF_50000			NA	
MyNetwork_30002	30002	MyVRF_50000	20.10.4.1/24		NA	
MyNetwork_30003	30003	MyVRF_50000			NA	

Editing Networks for the Standalone Fabric

To edit networks for standalone fabrics from Cisco DCNM Web UI, perform the following steps:

Procedure

- Step 1** Click **Control > Networks**.
The **Networks** window appears.
- Step 2** Choose a fabric from the **SCOPE** drop-down list.
The **Networks** window refreshes and lists the networks in the fabric.
- Step 3** Choose a network.
- Step 4** Click the **Edit** icon.
The **Edit Network** window appears.
- Step 5** Update the fields in the **General** and **Advanced** tabs of the **Network Profile** area as needed.

Note You can edit the network name. The edited network name appears in the **Network Name** column in the **Networks** window. The original name, which you used while creating a network, appears in the **Display Name** column. To view the original network name from the **Display Name** column in the **Networks** window, click **Settings**. Expand the **Columns** drop-down list, and choose the **Display Name** option. Click **Close**. You can also view the original network name in the network topology view.

- Step 6** Click **Save** at the bottom right part of the window to save the updates.

Creating VRFs for the Standalone Fabric

- Click **Control > VRFs** (under **Fabrics** submenu).
The VRFs screen comes up.
- Choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

Fabric Selected: bgp2

VRFs Selected 1 / Total 1

<input type="checkbox"/>	VRF Name	VRF ID	Status
<input checked="" type="checkbox"/>	MyVRF_50000	50000	NA

- Click the + button to add VRFs to the *Standalone* fabric. The Create VRF screen comes up. Most of the fields are autopopulated.

Create VRF
✕

▼ VRF Information

* VRF ID

* VRF Name

* VRF Template ▼

* VRF Extension Template ▼

VLAN ID Propose VLAN ?

▼ VRF Profile

General

Advanced

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

VRF Intf Description ⓘ

VRF Description ⓘ

Create VRF

The fields in this screen are:

VRF ID and **VRF Name**: The ID and name of the VRF.



Note For ease of use, the VRF creation option is also available while you create a network.

VRF Template: This template is applicable for VRF creation, and only applicable for leaf switches.

VRF Extension Template: The template is applicable when you extend the VRF to other fabrics, and is applicable for border devices.

Fill the fields in the **VRF Profile** section.

General tab – Enter the VLAN ID of the VLAN associated with the VRF, the corresponding Layer 3 virtual interface, and the VRF ID.

The VLAN ID default range is 2 to 3967. From DCNM Release 11.5(2), you can use a VLAN range greater than default value 3967. The reserved VLAN range must be set to a different range. In switch command enter “**system vlan <vlan> reserve**”. Save the configuration to startup configuration and reload the switch for the new reserved VLAN range to reflect.

From Cisco DCNM Web UI, choose **Administration > DCNM Server > Server Properties**, enter the value as 4094 for **RM.TOP_DOWN_NETWORK_VLAN.MAX** and **RM.TOP_DOWN_VRF_VLAN.MAX**, click **Apply Changes** and then restart DCNM. Once the DCNM is up, you can create VRF and network using the VLAN value greater than 3967.

Advanced tab – The fields in the tab are autopopulated.

VRF Intf MTU - Specifies VRF interface MTU.

Routing Tag – If a VLAN is associated with multiple subnets, then this tag is associated with the IP prefix of each subnet. Note that this routing tag is associated with overlay network creation too.

Redistribute Direct Route Map – Specifies the route map name for redistribution of routes in the VRF.

Max BGP Paths and **Max iBGP Paths** – Specifies the maximum BGP and iBGP paths.

TRM Enable – Select the check box to enable TRM.

If you enable TRM, then the RP address, and the underlay multicast address must be entered.

For more information, see [Overview of Tenant Routed Multicast, on page 148](#).

Is RP External – Enable this checkbox if the RP is external to the fabric. If this field is unchecked, RP is distributed in every VTEP.

RP Address – Specifies the IP address of the RP.

RP Loopback ID – Specifies the loopback ID of the RP, if **Is RP External** is not enabled.

Underlay Multicast Address – Specifies the multicast address associated with the VRF. The multicast address is used for transporting multicast traffic in the fabric underlay.



Note The multicast address in the **Default MDT Address for TRM VRFs** field in the fabric settings screen is auto-populated in this field. You can override this field if a different multicast group address should be used for this VRF.

Overlay Multicast Groups – Specifies the multicast group subnet for the specified RP. The value is the group range in “ip pim rp-address” command. If the field is empty, 224.0.0.0/24 is used as default.

Enable IPv6 link-local Option - Select the check box to enable the IPv6 link-local option under the VRF SVI. If this check box is unchecked, IPv6 forward is enabled.

Enable TRM BGW MSite - Select the check box to enable TRM on Border Gateway Multisite.

Advertise Host Routes – Enable the checkbox to control advertisement of /32 and /128 routes to Edge Routers.

Advertise Default Route – Enable the checkbox to control advertisement of default routes internally.

To allow inter-subnet communication between end hosts in different VXLAN fabrics, where the subnets are present in both fabrics, you must disable the **Advertise Default Route** feature (clear the **Advertise Default Route** checkbox) for the associated VRF. This will result in /32 routes for hosts being seen in both fabrics. For example, Host1 (VNI 30000, VRF 50001) in Fabric1 can send traffic to Host2 (VNI 30001, VRF 50001) in Fabric2 only if the host route is present in both fabrics. When a subnet is present in only one fabric then default route is sufficient for inter-subnet communication.

Config Static 0/0 Route - Select the check box to enable static default route configuration.

BGP Neighbor Password - Specifies the VRF Lite BGP neighbor password.

BGP Password Key Encryption Type - Select the encryption type from this drop-down list.

Sample screenshots of the Create VRF screen:

Advanced tab:

▼ VRF Profile

General

Advanced

VRF Intf MTU ⓘ 68-9216

Loopback Routing Tag ⓘ 0-4294967295

Redistribute Direct Route Map ⓘ

Max BGP Paths ⓘ 1-64

Max iBGP Paths ⓘ 1-64

TRM Enable ⓘ Enable Tenant Routed Multicast

Is RP External ⓘ Is RP external to the fabric?

Create VRF

4. Click **Create VRF**.

The *MyVRF_50001* VRF is created and appears on the VRFs page.

Network View |
 Continue

Fabric Selected: Standalone

VRFs Selected 1 / Total 2

+ ✎ ✕ ↻ ↺

Show All

	VRF Name	VRF ID	Status
<input type="checkbox"/>	MyVRF_50000	50000	NA
<input checked="" type="checkbox"/>	MyVRF_50001	50001	NA

Export and Import VRF Information

You can export VRF information to a .CSV file. The exported file contains information pertaining to each VRF, including the fabric it belongs to, the templates used to create the VRF, and all other configuration details that you saved during VRF creation.

In the VRFs screen, click the Export icon to export VRF information as a .CSV file.

VRFs

+ ✎ ✕ ↻ ↺

	VRF Name	VRF ID
<input type="checkbox"/>	MyVRF_50000	50000

.CSV

	A	B	C	D
	fabric	vrfName	vrflid	vrfTemplate
	Standalone	MyVRF_50000	50000	Default_VRF_Universal

You can use the exported .CSV file for reference or use it as a template for creating new VRFs. To import VRFs, do the following:

1. Update new records in the .CSV file. Ensure that the **vrfTemplateConfig** field contains the JSON Object.
2. In the VRFs screen, click **Import** icon and import the .CSV file into DCNM.

A message at the bottom right part of the screen displays errors and success messages. This screenshot depicts a new VRF being imported.



Note When you create a VRF using the **Import** option on the VRF window or using the DCNM APIs, you might see an error saying: Instance name is not specified.

This error is because of a tagging issue. To remove this error, edit the VRF in DCNM Web UI and then deploy.

VRFs

	A	B	C	D	E
	fabric	vrfName	vrfId	vrfTemplate	vrfExtensionTemplate
	Standalone	MyVRF_50001	50001	Default_VRF_Universal	Default_VRF_Extension_Universal
					{"vrfVlanId":"3","vrfDes

You can see that the imported VRF is displayed in the VRFs screen.

VRFs

VRF Name	VRF ID	Status
MyVRF_50000	50000	NA
MyVRF_50001	50001	NA

Editing VRFs for the Standalone Fabric

1. Choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

SCOPE: bgp2

Fabric Selected: bgp2

VRF Name	VRF ID	Status
<input checked="" type="checkbox"/> MyVRF_50000	50000	NA

2. From the **Select a Fabric** drop-down list, select the fabric *Standalone*, and click **Continue** on the top right part of the screen. The Networks page is displayed.
3. Click the **VRF View** at the top right part of the screen. The VRFs page appears.

Fabric Selected: New7200

VRFs Selected 0 / Total 2

<input type="checkbox"/>	VRF Name	VRF ID	Status
<input type="checkbox"/>	MyVRF_50000	50000	NA
<input type="checkbox"/>	MyVRF_50001	50001	NA

4. Select the **VRF** and click the **Edit** option at the top left part of the screen. The **Edit VRF** screen comes up.
5. Update the fields in the **General** and **Advanced** tabs of the **VRF Profile** section as needed.
6. Click **Save** at the bottom right part of the screen to save the updates.

Deploying Networks for the Standalone and MSD Fabrics

Before you begin: Ensure that you have created networks for the fabric.

1. Click **Control > Networks** (under **Fabrics** submenu).

The Networks screen comes up.

2. Choose the correct fabric from SCOPE. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

Fabric Selected: bgp2

Networks Selected 1 / Total 1

<input type="checkbox"/>	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
<input checked="" type="checkbox"/>	MyNetwork_30000	30000	NA			NA	

3. Select networks that you want to deploy. In this case, select the check boxes next to both the networks and click **Continue** at the top right part of the screen.

The Network Deployment page appears. On this page, you can see the network topology of the Standalone fabric.

You can deploy networks simultaneously on multiple switches. The selected devices should have the same role (Leaf, Border Gateway, and so on).



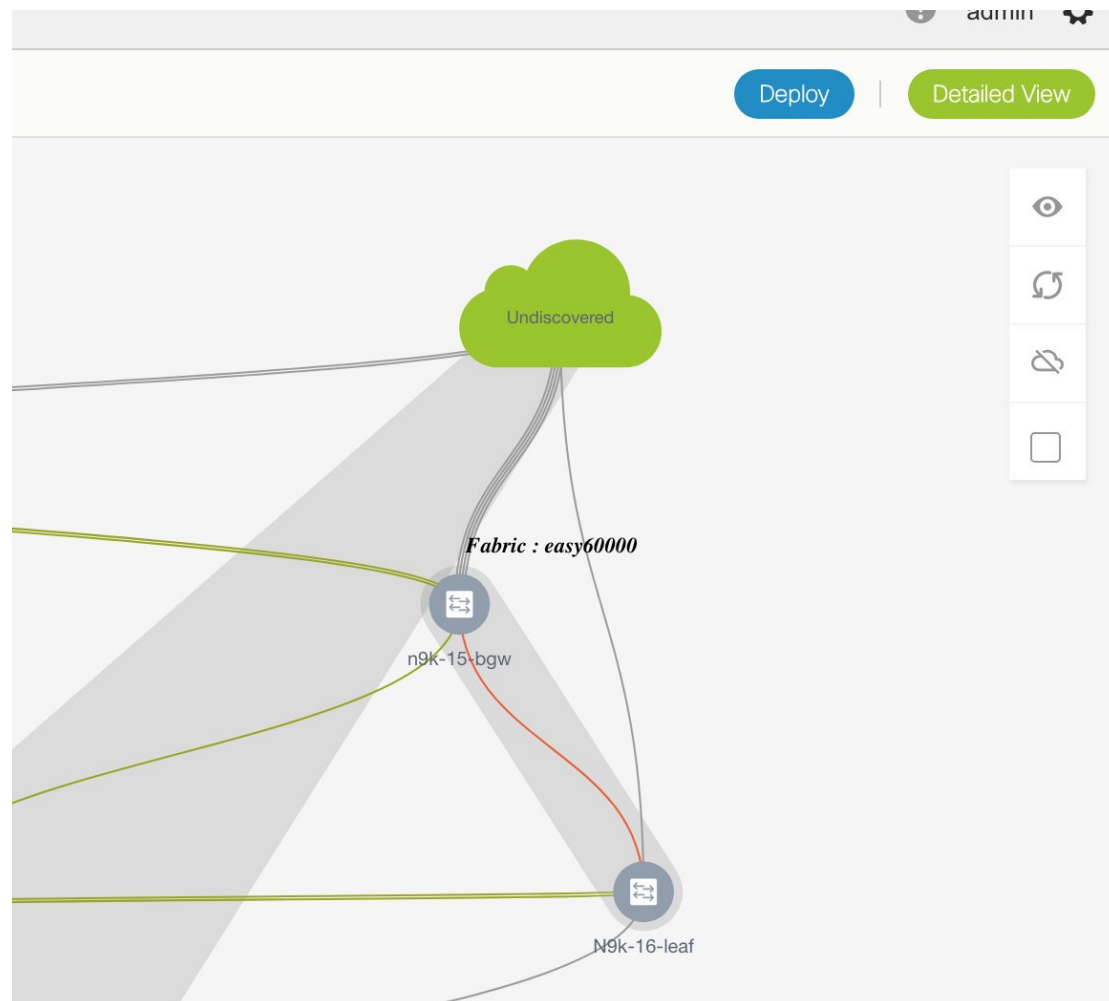
Note In an MSD fabric, all member fabrics are visible from this screen.

At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly. Blue for *Pending* state, yellow for *In Progress* when the provisioning is in progress, green when successfully deployed, and so on. From DCNM 11.3(1), the pending state indicates that there is a pending deployment or pending recomputation. You

can click on the switch and review the pending deployments using **Preview** or **Deploy Config** options, or click **Save & Deploy** to recompute the state of the switch.

The overlay networks (/VRFs) provisioning status is context-specific. It is a combination of networks that you chose for provisioning and the relevant switches in the topology. In this example, it means that the networks *MyNetwork_30000* and *MyNetwork_30001* are yet to be deployed on any switch in this fabric.

Undiscovered cloud display – To display (or not display) an **Undiscovered** cloud in this screen, click the cloud icon in the vertical panel, at the top-right part of the screen. When you click the icon, the **Undiscovered** cloud and its links to the fabric topology are not displayed. Click the icon again to display the **Undiscovered** cloud.



You can move the topology around the screen by clicking the left mouse button on the screen and moving it in the direction you desire. You can enlarge or shrink the switch icons proportionately by moving the cursor roller. You can also use corresponding alternatives on the touchpad.

4. Click ... in the **Interfaces** column.

The **Interfaces** box opens up. It lists interfaces or port channels. You can select interfaces/port channels to associate them with the selected network. For each interface, port type and description, channel number and connected neighbor interface details are displayed.

From Cisco DCNM Release 11.5(1), the **Interfaces** window doesn't list interfaces that are part of an interface group. Specifically, the trunk ports, access ports, and dot1q tunnel ports.


If you try to perform a network attachment to a switch and an interface is part of an interface group, an appropriate error is displayed.

Interfaces

<input type="checkbox"/>	Interface/Ports ▲	Channel ...	Port Ty...	Port Desc...	Neighbor Info
<input type="checkbox"/>	Ethernet1/1	NA	trunk		
<input checked="" type="checkbox"/>	Ethernet1/10	NA	trunk		
<input checked="" type="checkbox"/>	Ethernet1/11	NA	trunk		
<input type="checkbox"/>	Ethernet1/12	NA	trunk		
<input type="checkbox"/>	Ethernet1/13	NA	trunk		

Save

5. Double-click a switch to deploy the networks on it. For deployment of networks on multiple switches, click Multi-Select from the panel at the top right part of the screen (the topology freezes to a static state), and drag the cursor across the switches.



Immediately the Network Attachment dialog box appears.

Network Attachment - Attach networks for given switch(es) 

Fabric Name: Standalone

Deployment Options

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

MyNetwork_30000		MyNetwork_30001				
<input type="checkbox"/>	Switch ▲	VLAN	Interfaces	CLI Freeform	Status	
<input type="checkbox"/>	n9k-16-leaf	2300	...	Freeform config	NA	

Save

A tab represents each network (the first network is displayed by default) that is being deployed. In each network tab, the switches are displayed. Each row represents a switch.

Click the check box next to the **Switch** column to select all switches. The network is ready to be provisioned on the switches.

VLAN - Update the VLAN ID if needed.

When you update a VLAN ID and complete the network deployment process, the old VLAN is not automatically removed. To complete the process, you should go to the fabric topology screen (click **Control > Fabric Builder** and click within the corresponding fabric box to go to the screen) and use the Save and Deploy option.

When updating the VLAN ID for a given network, the original VLAN ID is not automatically removed from the attached trunk interface. In order to remove the old or original VLAN ID, you must perform **Save and Deploy + Config Deploy** operation from within the fabric in Fabric Builder. For this, go to the fabric topology screen (click **Control > Fabric Builder** and click within the corresponding fabric box to go to the screen) and execute the **Save and Deploy** operation. Verify that config compliance is removing the expected config, then execute **Deploy Config** operation to remove the configs.

Interfaces – Click ... in the column to add interfaces associated with the selected network.

VLAN to trunk port mapping – The selected trunk ports include the VLAN as an allowed VLAN on the port.

VLAN to vPC domain mapping - If you want to associate the VLAN to port channels of a vPC domain, add the port channels from the list of interfaces. The vPC port channels include the VLAN as an allowed VLAN.

Freeform configurations – Click Freeform config to enable additional configurations on the switch. After the configurations are saved, the Freeform config button gets highlighted.

6. Select the other network tab and make the same selections.

7. Click **Save** (at the bottom right part of your screen) to save the configurations.



Note Addition and removal of interfaces are displayed in the **Interfaces** column of the Switches Deploy screen. Though the interface-related updates (like addition or removal of trunk ports) are provisioned on the switches, the correct configurations will not reflect in the preview screen. When you add or remove a trunk or access port, the preview shows the addition or removal of configurations for the interface under that network.

The topology window appears again. Click *Refresh* in the vertical panel at the top right part of the screen. The blue color on the switch icons indicates that the deployment is pending. From DCNM 11.3(1), the pending state indicates that there is a pending deployment or pending recomputation. You can click on the switch and review the pending deployments using **Preview** or **Deploy Config** options, or click **Save & Deploy** to recompute the state of the switch.

8. Preview the configurations by clicking *Preview* (the eye icon above the Multi-Select option). Since *MyNetwork_30000* and *MyNetwork_30001* are networks of VRF *50000*, the configurations contain VRF configurations followed by the network configurations.

Preview Configuration

Select a Switch:

n9k-16-leaf

Select a Network

MyNetwork_30000

Generated Configuration:

```
configure profile MyVRF_50000
vlan 2000
vn-segment 50000
interface vlan2000
vrf member myvrf_50000
ip forward
ipv6 forward
no ip redirects
no ipv6 redirects
mtu 9216
no shutdown
vrf context myvrf_50000
vni 50000
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
address-family ipv6 unicast
route-target both auto
route-target both auto evpn
router bgp 60000
vrf myvrf_50000
address-family ipv4 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redirect-subnet
maximum-paths ibgp 2
address-family ipv6 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redirect-subnet
maximum-paths ibgp 2
interface nve1
member vni 50000 associate-vrf
configure terminal
apply profile MyVRF_50000
```

**MyVRF_50000
Configuration**

Preview Configuration

Select a Switch:

n9k-16-leaf

Select a Network

MyNetwork_30000

Generated Configuration:

```
vrf myvrf_50000
address-family ipv4 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redis-subnet
maximum-paths ibgp 2
address-family ipv6 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redis-subnet
maximum-paths ibgp 2
interface nve1
member vni 50000 associate-vrf
configure terminal
apply profile MyVRF_50000
```

```
configure profile MyNetwork_30000
vlan 2300
vn-segment 30000
interface vlan2300
vrf member myvrf_50000
fabric forwarding mode anycast-gateway
no shutdown
interface nve1
member vni 30000
mcast-group 239.1.1.0
evpn
vni 30000 l2
rd auto
route-target import auto
route-target export auto
configure terminal
apply profile MyNetwork_30000
```

```
interface ethernet1/11
switchport trunk allowed vlan add 2300
interface ethernet1/10
switchport trunk allowed vlan add 2300
```

MyNetwork_30000
Configuration

Interfaces Configuration

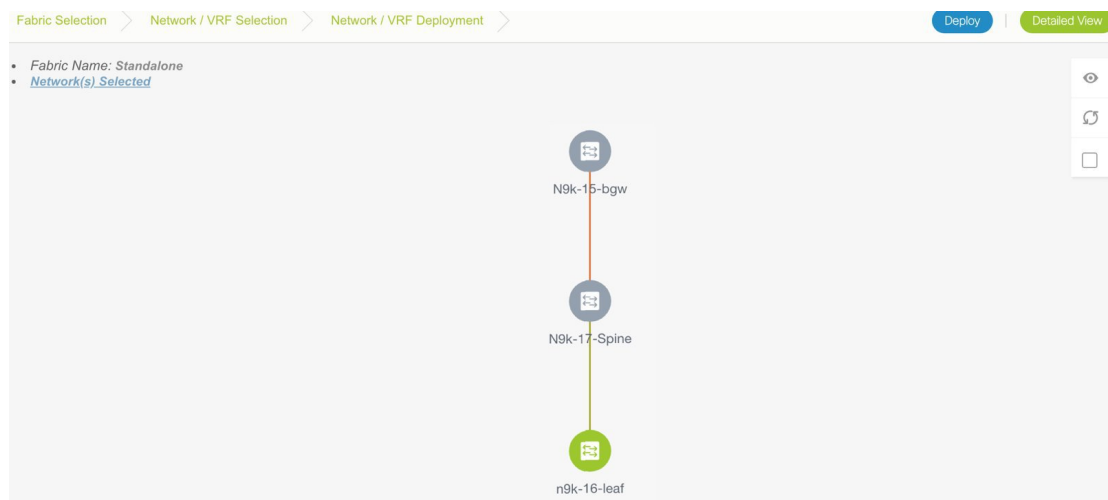
On the preview screen, you can select from the **Select a switch** and **Select a network** drop-down boxes at the top of the screen to view other network configurations.

After checking the configurations, close the screen. The Topology screen appears again.

- Click **Deploy** on the top right part of the screen. The color of the switch icons changes to yellow and a message appears at the bottom right part of the screen indicating that the deployment is in progress. After the networks' deployment is complete, the color of the switch icons changes to green, indicating successful deployment.



Note In case you click **Deploy** and there is no configuration diff that has to be deployed, a pop-up window comes up stating **No switches PENDING for deployment**.



Note The status of the switch is determined by the aggregated status of the selected networks or VRFs in the following hierarchy: **Pending, In Progress, Out-of-Sync/Failed, In Sync/Success**, and **Unknown/NA**. For example, if any one of the networks or VRFs is in the **Out-of-Sync/Failed** status and others are not in the **Pending or In Progress** status, then the switch status is **Out-of-Sync/Failed**. The default status is **Unknown/NA**, when the status is not known.


Go to the Networks page to view the individual status for all networks.

Network Deployment for an MSD Fabric

Consider a scenario wherein you are deploying the same networks on different member fabric border devices. You can choose one fabric, deploy networks on its border devices, and then choose the second fabric and deploy networks.

Alternatively, you can choose the MSD fabric, and deploy the networks from a single topology view of all member fabric border devices.

This is a topology view of an MSD fabric wherein the two member fabrics topologies and their connections are depicted. You can deploy networks on the BGWs of the fabrics at once.



Detailed View

You can also use the Detailed View option to deploy networks and VRFs. Click **Detailed View** at the top right part of the screen. The Detailed View window appears. This lists the networks in a tabular view.

<input type="checkbox"/>	Name	Switch	Ports	Status	Fabric Name	Role
<input type="checkbox"/>	MyNetwork_30000	N9k-15-bgw		NA	new60000	border
<input type="checkbox"/>	MyNetwork_30001	N9k-15-bgw		NA	new60000	border
<input type="checkbox"/>	MyNetwork_30001	n9k-16-leaf	Ethernet1/1	DEPLOYED	new60000	leaf
<input type="checkbox"/>	MyNetwork_30000	n9k-16-leaf	Ethernet1/10,Ethernet1/11	DEPLOYED	new60000	leaf

The options:

Edit - Select a network and click the Edit icon at the top left part of the screen.



Note If you select one network/switch entry and click on Edit, the Network Attach dialog box appears. To maintain consistency across the Topology View and Detailed View screens, the Network Attach screen displays all networks, and not just the selected network/switch.

Preview – Click Preview to preview configurations before deployment. You can only preview pending configurations, and not uninitiated or deployed configurations.

Deploy – Click Deploy to provision networks onto the switches.

History – Select a row and click History to view the configuration instances and status. Network and VRF-wise configurations are displayed. Click in the Status column of any instance for more details.

The fields in the table contain the configuration instance in each row, the associated switch and fabric names, the switch role, trunk ports (if any), and the deployment status.

Quick Attach – Choose a network and click **Quick Attach**. A confirmation window appears. Click **OK**. The network will be attached to the selected switch.

Quick Detach– Choose a network and click **Quick Detach**. A confirmation window appears. Click **OK**. The network will be detached from the selected switch.

On the Detailed View page, the network profile configuration history is displayed. If you have associated specific trunk interfaces to that network, then the interface configuration is displayed as a separate configuration instance.



Note When you upgrade from an earlier release (such as DCNM 10.4[2]) to the DCNM 11.0(1) release, overlay networks and VRFs deployment history information from the earlier DCNM release is not retained.

Deploying VRFs for the Standalone and MSD Fabrics

1. Choose the correct fabric from SCOPE. When you select a fabric, the **VRFs** screen refreshes and lists VRFs of the selected fabric.

Fabric Selected: bgp2

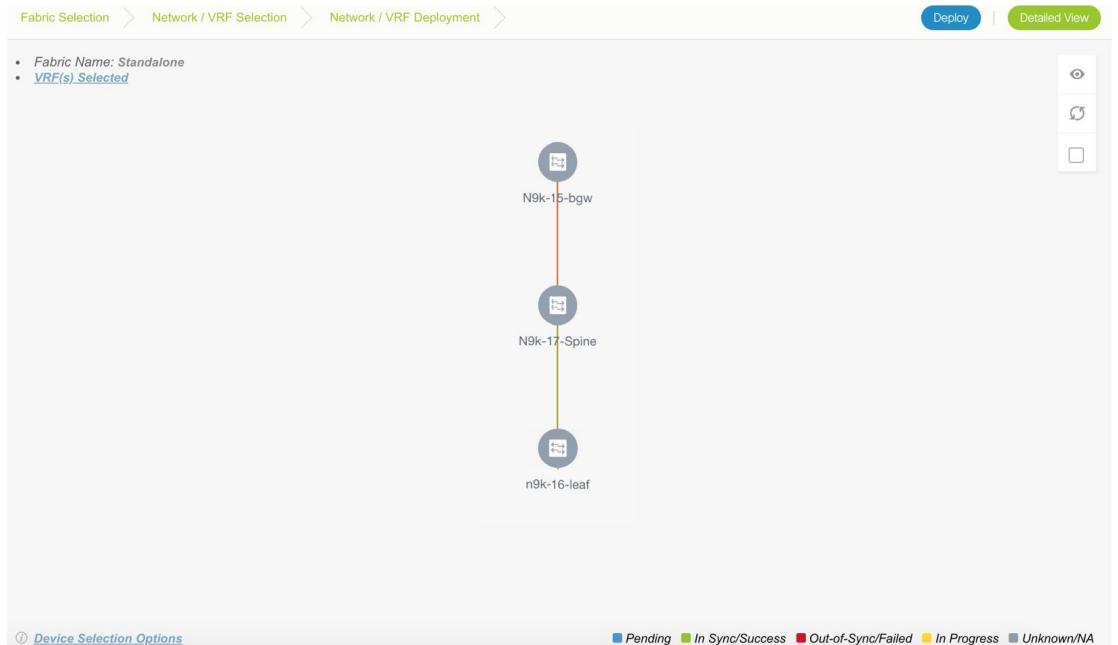
VRFs

Selected 1 / Total 1

<input type="checkbox"/>	VRF Name	VRF ID	Status
<input checked="" type="checkbox"/>	MyVRF_50000	50000	NA

2. Select check boxes next to the VRFs that you want to deploy and click **Continue** at the top right part of the screen.


The VRF Deployment screen appears. On this page, you can see the topology of the Standalone fabric. The following example shows you how to deploy the VRFs MyVRF_50000 and MyVRF_50001 on the leaf switch. You can deploy VRFs simultaneously on multiple switches but of the same role (Leaf, Border Gateway, and so on).



At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly. Blue for *Pending* state, yellow for *In Progress* state when the provisioning is in progress, red for failure state, green when successfully deployed, and so on. From DCNM 11.3(1), the pending state indicates that there is a pending deployment or pending recomputation. You can click on the switch and review the pending deployments using **Preview** or **Deploy Config** options, or click **Save & Deploy** to recompute the state of the switch.

The overlay networks (or VRFs) provisioning status is context-specific. It is a combination of VRFs that you chose for provisioning and the relevant switches in the topology. In this example, it means that the VRFs are yet to be deployed on any switch in this fabric.

Undiscovered cloud display – To display (or not display) an **Undiscovered** cloud in this screen, click the cloud icon in the vertical panel, at the top-right part of the screen. When you click the icon, the **Undiscovered** cloud and its links to the fabric topology are not displayed. Click the icon again to display the **Undiscovered** cloud.



You can move the topology around the screen by clicking the left mouse button on the screen and moving it in the direction you desire. You can enlarge or shrink the switch icons proportionately by moving the cursor roller. You can also use corresponding alternatives on the touchpad.

3. Double-click a switch to deploy VRFs on it. The VRF Attachment screen comes up.



Note For deployment of VRFs on multiple switches, click the Multi-Select option from the panel at the top right part of the screen (This freezes the topology to a static state), and drag the cursor across the switches.

VRF Attachment - Attach VRFs for given switch(es).



Fabric Name: Standalone

Deployment Options

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

MyVRF_50000		MyVRF_50001			
<input type="checkbox"/>	Switch	▲	VLAN	CLI Freeform	Status
<input type="checkbox"/>	n9k-16-leaf		2000	Freeform config	NA

Save

A tab represents each VRF that is being deployed (the first selected VRF is displayed by default). In each VRF tab, the selected switches are displayed. Each row represents a switch.

VLAN ID - Click within the VLAN column to update the VRF VLAN ID, if needed.

Freeform configurations – Click Freeform config to enable additional configurations on the switch. After you save freeform configurations, the Freeform config button gets highlighted.

Click the checkbox next to the Switch column to select all switches. VRF MyVRF_50000 is ready to be provisioned on the switch

4. Select the other VRF tab and make the same selections.
5. Click **Save** (at the bottom right part of your screen) to save VRF configurations.

The topology screen comes up again. Click the *Refresh* button in the vertical panel at the top right part of the screen. The blue color on the switch icons indicates that the deployment is pending. From DCNM 11.3(1), the pending state indicates that there is a pending deployment or pending recomputation. You can click on the switch and review the pending deployments using **Preview** or **Deploy Config** options, or click **Save & Deploy** to recompute the state of the switch.

Preview the configurations by clicking the *Preview* button (the eye icon above the *Multi-Select* option).

Preview Configuration



Select a Switch:

n9k-16-leaf ▼

Select a VRF

MyVRF_50000 ▼

Generated Configuration:

```
configure profile MyVRF_50000
vlan 2000
vn-segment 50000
interface vlan2000
vrf member myvrf_50000
ip forward
ipv6 forward
no ip redirects
no ipv6 redirects
mtu 9216
no shutdown
vrf context myvrf_50000
vni 50000
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn
address-family ipv6 unicast
route-target both auto
route-target both auto evpn
router bgp 60000
vrf myvrf_50000
address-family ipv4 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redis-subnet
maximum-paths ibgp 2
address-family ipv6 unicast
advertise l2vpn evpn
redistribute direct route-map fabric-rmap-redis-subnet
maximum-paths ibgp 2
interface nve1
member vni 50000 associate-vrf
configure terminal
apply profile MyVRF_50000
```

After checking the configurations, close the screen. The *Topology View* screen appears.

- Click the **Deploy** button on the top right part of the screen. The color of the switch icons changes to yellow and a message appears at the bottom right part of the screen indicating that the deployment is in progress. After the VRF deployment is complete, the color of the switch icons changes to green, indicating successful deployment.

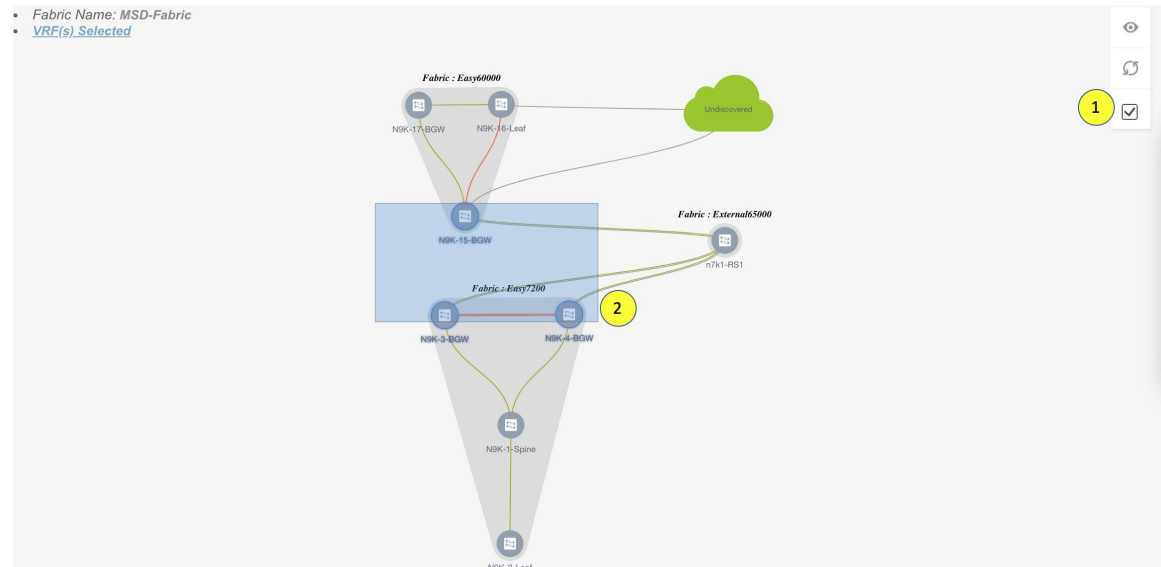


Note In case you click **Deploy** and there is no configuration diff that has to be deployed, a pop-up window comes up stating **No switches PENDING for deployment**.

VRFs Deployment for an MSD Fabric

Consider a scenario wherein you are deploying the same VRFs on different member fabric border devices. You can choose one fabric, deploy VRFs on its border devices, and then choose the second fabric and deploy the VRFs.

Alternatively, you can choose the MSD fabric, and deploy the VRFs from a single topology view of all member fabric border devices at once.



Detailed View

You can also use the **Detailed View** button to deploy networks and VRFs.

Click **Detailed View** at the top right part of the screen. The Detailed View screen comes up. This lists the VRFs in a tabular view.

Fabric Selection > Network / VRF Selection > Network / VRF Deployment > Topology View

Fabric Name: Standalone VRF(s) Selected Selected 0 / Total 4

Show All

<input type="checkbox"/>	Name	Switch	Ports	Status	Fabric Name	Role
<input type="checkbox"/>	MyVRF_50000	n9k-15-BL		NA	Easy60000	leaf
<input type="checkbox"/>	MyVRF_50000	n9k-16-leaf		DEPLOYED	Easy60000	leaf
<input type="checkbox"/>	MyVRF_50001	n9k-15-BL		NA	Easy60000	leaf
<input type="checkbox"/>	MyVRF_50001	n9k-16-leaf		DEPLOYED	Easy60000	leaf

The options:

Edit - Select a VRF and click the Edit icon at the top left part of the screen.



Note If you select one VRF/switch entry, the VRF Attach screen comes up. To maintain consistency across the Topology View and Detailed View screens, the VRF Attach screen displays all VRFs, and not just the selected VRF/switch entry.

Preview – Click Preview to preview configurations before deployment. You can only preview pending configurations, and not uninitiated or deployed configurations.

Deploy – Click Deploy to provision VRFs onto the switches.

History – Select a row and click History to view the configuration instances and status. Network and VRF-wise configurations are displayed. Click in the Status column of any instance for more details.

The fields in the table contain the configuration instance in each row, the associated switch and fabric names, the switch role, and the deployment status.

Quick Attach: Choose a VRF and click **Quick Attach**. A confirmation window appears. Click **OK**. The VRF will be attached to the selected switch.

Quick Detach: Choose a VRF and click **Quick Detach**. A confirmation window appears. Click **OK**. The VRF will be detached from the selected switch.



Note When you upgrade from an earlier release (such as DCNM 10.4[2]) to the DCNM 11.0(1) release, overlay networks and VRFs deployment history information from the earlier DCNM release is not retained.

Undeploying Networks for the Standalone Fabric

You can undeploy VRFs and networks from the deployment screen. The DCNM screen flow for undeployment is similar to the deployment process flow. Go to the deployment screen (Topology View) to undeploy networks:

1. Click **Control** > **Networks** (under **Fabrics** submenu).

The Networks screen comes up.

2. Choose the correct fabric from SCOPE. When you select a fabric, the **Networks** screen refreshes and lists networks of the selected fabric.

The screenshot shows the Cisco Data Center Network Manager interface. At the top, the breadcrumb is "Network / VRF Selection > Network / VRF Deployment". The "SCOPE" dropdown is set to "bgp2". Below the breadcrumb, there are "VRF View" and "Continue" buttons. The main content area is titled "Fabric Selected: bgp2" and "Networks". It shows a table with one row selected:

	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
<input checked="" type="checkbox"/>	MyNetwork_30000	30000	NA			NA	

3. Select the networks that you want to undeploy and click Continue. The topology view comes up.
4. Select the Multi-Select button (if you are undeploying the networks from multiple switches), and drag the cursor across switches with the same role. The Network Attachment screen comes up.
 - (For a single switch, double-click the switch and the Network Attachment screen comes up).
 - (For a single switch, double-click the switch and the Switches Deploy screen comes up).
5. In the Network Attachment screen, the Status column for the deployed networks is displayed as DEPLOYED. Clear the check boxes next to the switches, as needed. Ensure that you repeat this on all tabs since each tab represents a network.
6. Click **Save** (at the bottom right part of the screen) to initiate the undeployment of the networks. The *Topology View* comes up again.



Note Alternatively, you can click the **Detailed View** button to undeploy networks.

7. Refresh the screen, preview configurations if needed and click **Deploy** to remove the network configurations on the switches. After the switch icons turn green, it indicates successful undeployment.
8. Go to the Networks page to verify if the networks are undeployed.

Undeploying VRFs for the Standalone Fabric

You can undeploy VRFs from the deployment screen. The DCNM screen flow for undeployment is similar to the deployment process flow.

1. Choose **Control > Fabrics > VRFs**.
2. Choose the correct fabric from **SCOPE**. When you select a fabric, the **VRFs** screen refreshes and lists networks of the selected fabric.
3. Select the VRFs that you want to undeploy and click **Continue**. The *Topology View* page comes up.
4. Select the Multi-Select option (if you are undeploying the VRFs from multiple switches), and drag the cursor across switches with the same role. The VRF Attachment screen comes up.
(For a single switch, double-click the switch and the VRF Attachment screen comes up).
5. In the Switches Deploy screen, the **Status** column for the deployed VRFs is displayed as DEPLOYED. Clear the check boxes next to the switches, as needed. Ensure that you repeat this on all tabs since each tab represents a VRF.
6. Click **Save** (at the bottom right part of the screen) to initiate the undeployment of the VRFs. The topology view comes up again.



Note Alternatively, you can click the **Detailed View** button to undeploy VRFs.

7. Refresh the screen, preview configurations if needed and click **Deploy** to remove the VRF configurations on the switches. After the switch icons turn green, it indicates successful undeployment.
8. Go to the VRFs page to verify if the networks are undeployed.

Deleting Networks and VRFs

If you want to delete networks and corresponding VRFs in the MSD fabric, follow this order:

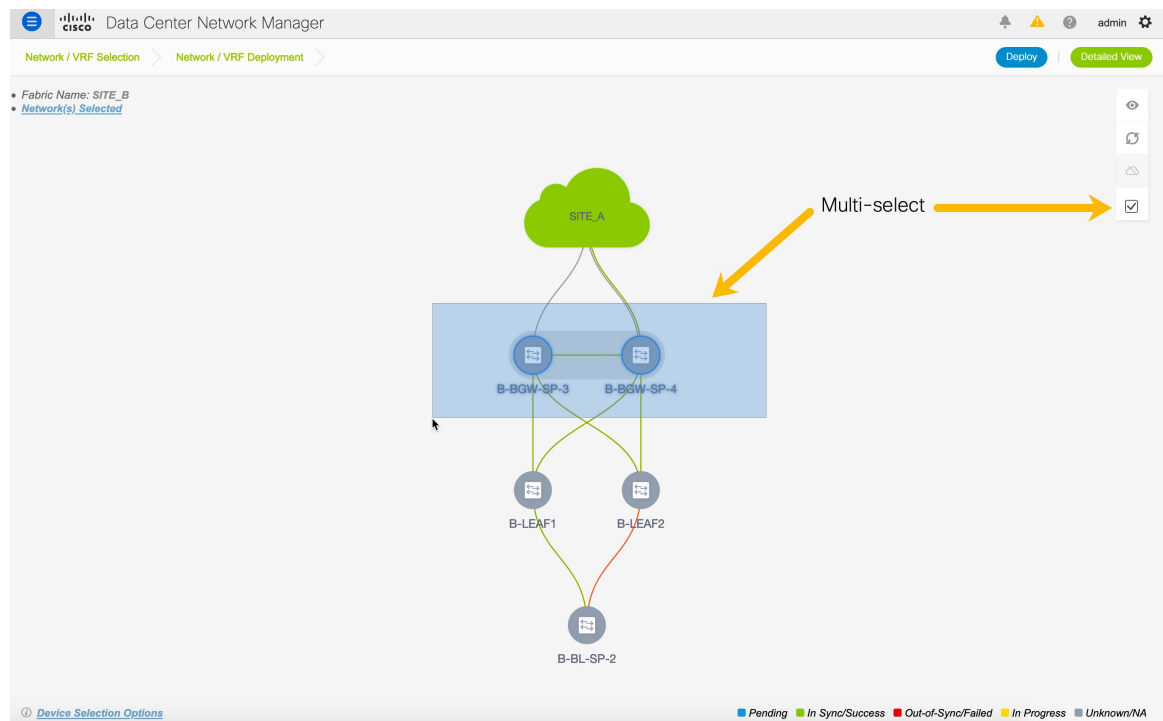
1. Undeploy the networks, if not already done.
2. Delete the networks.
3. Undeploy the VRFs, if not already done.
4. Delete the VRFs.

Configuring Multiple VLAN IDs to a Single VNI

The following procedure shows how to tag multiple VLAN IDs to a single VNI in DCNM.

Procedure

- Step 1** Navigate to **Control > Networks**.
- Step 2** Select the fabric from the **Scope** drop-down list and then select the network. Click **Continue**.
- Step 3** Check the **Multi-Select** check box and drag the cursor over the switches that need to be updated with VLAN IDs.



- Step 4** In the **Network Attachment** window, edit the VLAN ID for the switches and click **Save**.

Network Extension Attachment - Attach extensions for given switch(es)



Fabric Name: SITE_B

Deployment Options

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

	Switch	VLAN	Extend	Interfaces	CLI Freeform	Status
	MyNetwork_30000 Network VNI					
<input type="checkbox"/>	B-BGW-SP-3	2300	MULTISITE			NA
<input type="checkbox"/>	B-BGW-SP-4	2300	MULTISITE	...	Freeform config	NA

Switches

Save

Step 5 Click **Deploy** to deploy the configuration.

Enhanced Role-based Access Control in Cisco DCNM

From Cisco DCNM Release 11.4(1), you can see the following role-based access control (RBAC) changes:

- Read-only access to the Cisco DCNM Web UI and APIs for the **network-operator** user role
- A new user role called **network-stager**.
- Freeze deployment for a particular fabric or all fabrics in DCNM as a user with the **network-admin** role.

From Cisco DCNM Release 11.5(1), you can see new user roles, **device-upg-admin**, and **access-admin** are added.



Note Actions that cannot be performed by a selected user role is grayed out.

You can also watch the video that demonstrates some of the operations performed by a network stager and how to freeze a fabric in Cisco DCNM. See the [Enhanced Role-based Access Control \(RBAC\)](#) video.

Device-upg-admin Role

A user with the **device-upg-admin** role can perform operations only in **Image Management** window.

See the [Image Management, on page 369](#) section for more information.

Access-admin Role

A user with the **access-admin** role can perform operations only in **Interface Manager** window for all fabrics.

An access-admin can perform the following actions:

- Add, edit, delete and deploy layer 2 port channels, and vPC.

- Edit host vPC, and ethernet interfaces.
- Save, preview, and deploy from management interfaces.
- Edit interfaces for LAN classic fabrics.

Apart from nve, management, tunnel, subinterface, SVI, interface grouping, and loopback interfaces

However, a user with the access-admin role can't perform the following actions:

- Cannot edit layer 3 port channels, ST FEX, AA FEX, loopback interfaces, nve interfaces, and subinterfaces.
- Cannot edit member interfaces and port channels of Layer 3, ST FEX, AA FEX.
- Cannot edit interfaces with policy associated from underlay and link for easy fabrics.
- Cannot edit peer link port channel.
- Cannot edit management interface.
- Cannot edit tunnel.



Note The icons and buttons are grayed out for this role when the fabric or DCNM is in deployment-freeze mode.

Network-Operator Role

A user with the **network-operator** role has access to the following menu in the DCNM Web UI:

- Dashboard
- Topology
- Monitor
- Applications

From Cisco DCNM, Release 11.4(1), a user with this role has read-only access to the **Control** menu as well.

A network operator can view fabric builder, fabric settings, preview configurations, policies, and templates. However, a network operator cannot perform the following actions:

- Cannot change expected configurations of any switch within any fabric.
- Cannot deploy any configurations to switches.
- Cannot access the administration options like licensing, creating more users, and so on.

Network-Stager Role

A user with the network-stager role can make configuration changes on DCNM. A user with the network-admin role can deploy these changes later. A network stager can perform the following actions:

- Edit interface configurations.
- View or edit policies.
- Create interfaces.

- Change fabric settings.
- Edit or create templates.

However, a network stager cannot perform the following actions:

- Cannot make any configuration deployments to switches.
- Cannot perform deployment-related actions from the DCNM Web UI or the REST APIs.
- Cannot access the administration options like licensing, creating more users, and so on.
- Cannot move switches in and out of maintenance mode.
- Cannot move fabrics in and out of deployment-freeze mode.
- Cannot install patches.
- Cannot upgrade switches.
- Cannot create or delete fabrics.
- Cannot import or delete switches.

The difference between a network operator and a network stager is as a network stager, you can only define intent for existing fabrics, but cannot deploy those configurations.

Only a network admin can deploy the changes and edits that are staged by a user with the **network-stager** role.

Viewing Policy Change History

Different users can simultaneously change expected configuration of switches in the DCNM. You can view the history of these staged changes in the **Policy Change History** tab. The deployment history captures the changes that are pushed or deployed from DCNM to switches.



Note Only deployment history is supported for non-Nexus devices.

To view the changes by different users, perform the following steps:

Procedure

- Step 1** Log into Cisco DCNM with the **network-admin**, **network-stager**, or **network-operator** user role.
- Step 2** Navigate to the fabric topology window.
- Step 3** Right-click the switch for which you intent change history.
- Step 4** Choose **History**.
- Step 5** Click the **Policy Change History** tab.
- Step 6** Search for the interface to which you made changes in the **Generated Config** column.
- Step 7** The **PTI Operation** column will have the value **UPDATE** for the changes made by different users.
- Step 8** Scroll horizontally to the **User** column. You can see the user names populated with the timestamp.

For each configurable entity, the detailed history under the **Generated Config** column, provides delta of configuration changes made by every user.

Policy ID	Template	Description	PTI Operation	Generated Config	Entity Name	Entity Type	User	Created On	Action	Source	Priority	Content Ty
POLICY-119870	int_access_host_11_1		UPDATE	Detailed History	Ethernet1/4	INTERFACE	stager2	2020/06/22-09:11:28			500	PYTHON
POLICY-119870	int_access_host_11_1		UPDATE	Detailed History	Ethernet1/4	INTERFACE	stager1	2020/06/22-09:10:39			500	PYTHON
POLICY-136560	evpn_bgp_tr_neigh...		ADD	Detailed History	SWITCH	SWITCH	admin	2020/06/22-09:05:44	Save & Deploy	UNDERLAY	150	TEMPLAT
POLICY-134480	evpn_bgp_tr_neigh...									UNDERLAY	150	TEMPLAT
POLICY-136550	evpn_bgp_tr_neigh...									UNDERLAY	150	TEMPLAT
POLICY-134470	evpn_bgp_tr_neigh...									UNDERLAY	150	TEMPLAT
POLICY-134450	nve_interface									nve1	-310	TEMPLAT
POLICY-134460	no_shut_interface									nve1	500	TEMPLAT
POLICY-134450	nve_interface									nve1	-310	TEMPLAT
POLICY-135070	int_fabric_num_11_1									LINK	310	PYTHON
POLICY-135230	no_shut_interface									Ethernet1...	352	TEMPLAT
POLICY-135220	pim_interface									Ethernet1...	352	TEMPLAT
POLICY-135210	ospf_p2p_interface									Ethernet1...	352	TEMPLAT
POLICY-135200	ospf_interface_11_1									Ethernet1...	352	TEMPLAT
POLICY-135190	interface_mtu									Ethernet1...	352	TEMPLAT
POLICY-133040	interface_desc									Ethernet1...	-352	TEMPLAT
POLICY-135180	interface_desc									Ethernet1...	352	TEMPLAT
POLICY-133000	p2p_routed_interface									Ethernet1...	-350	TEMPLAT
POLICY-135160	p2p_routed_interface									Ethernet1...	350	TEMPLAT

Freezing Fabrics in Cisco DCNM


As a network admin, you can disable, or freeze, deployments for LAN classic fabrics, easy fabrics, and external fabrics. Deployment freeze disables configuration or write access from the DCNM to the switches. When you freeze a fabric, switches cannot be reloaded, moved in and out of maintenance mode, and you cannot add or delete switches within the fabric. This feature provides complete control for a network admin to disable inadvertent changes to the physical network from the DCNM, unless a maintenance window is scheduled.

Freezing a Fabric

To disable deployment for a fabric from Cisco DCNM Web UI, perform the following:

Procedure

Step 1 Navigate to the **Fabric Builder** window or the fabric topology window.

Step 2 Click the spanner () icon.

The spanner icon is next to the fabric name in the fabric topology window. A confirmation window appears asking you if you want to disable all deployments for the fabric.

Step 3 Click **Yes**.

Note When you hover over the spanner icon before freezing the fabric, the tooltip will read **Deployment Enabled**. When you hover over the spanner icon after freezing the fabric, the tooltip will read **Deployment Disabled**.

After you disable deployments or freeze a fabric, you can only save, edit, or preview changes but not deploy them. All deploy related actions from the DCNM to this fabric will be grayed out.

To enable all deployments for the fabric, click the same spanner (⊗) icon and unfreeze the fabric.

Freezing All Fabrics

In addition to the per-fabric deployment freeze knob, the network admin can freeze deployments for all fabrics within the DCNM at the same time.

To freeze all fabrics in your DCNM setup from the Cisco DCNM Web UI, perform the following steps:

Procedure

Step 1 Choose **Administration > DCNM Server > Server Properties**.

Step 2 Search for the **DEPLOYMENT_FREEZE** field.

Step 3 Set the value as **true**.

The default value is **false**.

Note When you freeze DCNM, you cannot deploy any changes to switches. However, users with appropriate roles, like the network-admin role or the network-stager role, with appropriate access, can make changes in the DCNM for deployment at a later stage.

Actions that cannot be performed when you freeze a fabric or DCNM are grayed out.

Fabric Backup and Restore

This section describes the fabric backup and restore in Cisco DCNM.

Backing Up Fabrics

You can back up all fabric configurations and intents automatically or manually. You can save configurations in DCNM, which are the intents. The intent may or may not be pushed on to the switches.

DCNM doesn't back up the following fabrics:

- External fabrics in monitor-only mode: Backing up of external fabrics in monitor-only mode isn't supported because you can't restore any configurations or intent. However, if such external fabrics are member fabrics of an MSD fabric, the backup is taken at MSD-fabric level.



Note From Cisco, DCNM Release 11.4(1), you can take a backup of external fabrics in monitor-only mode, but can't restore them. You can restore this backup when the external fabric isn't in monitor-only mode.

- Parent MSD fabrics in releases earlier than Cisco DCNM, Release 11.4(1): You can only back up the configurations and intent of member fabrics in an MSD fabric individually.



Note From Cisco DCNM, Release 11.4(1), you can take backups of MSD fabrics. When you initiate a backup from the parent fabric, the backup process is applicable for the member fabrics as well. However, DCNM stores all the backed-up information of the member fabrics and the MSD fabric together in a single directory.

From Cisco DCNM Release 11.4(1), the backup captures the intent related to IFC as well. When you're backing up an external fabric, the checkpoints are copied from the switches to DCNM. The backup configuration files are stored in the following path in DCNM: `/usr/local/cisco/dcm/dcm/data/archive`

The backed-up config files can be found in the corresponding directory with the fabric name. Each backup of a fabric is treated as a different version, regardless if it is backed up manually or automatically. You can find all versions of the backup in the corresponding fabric directories. Hence, the backed up intent configuration file, running configuration file and PTIs can be found at location:

`/usr/local/cisco/dcm/dcm/data/archive/<fabric_name>/Version_x`, where `x` is the version number. The valid value is between 1 and the limit you set in the **archived.versions.limit** field. The default value is 50, which means only 50 backups are archived, and the oldest backups are removed. The minimum value is 10. If you specify a value lesser than 10, it will be overwritten to 10. You can set the number of backup files to be archived in the **Server Properties** window. Search for the **# Number of archived files per fabric to be retained:** section in the **Server Properties** window. Enter a value in the **archived.versions.limit** field.

You can also watch the video that demonstrates how to back up and restore an MSD fabric in Cisco DCNM. See the [MSD Fabric Backup and Restore](#) video.

Backing Up Fabrics Automatically

You can enable an automatic hourly backup or scheduled backup for fabric configurations and intents. There are two types of automatic backup.

The backup has the information related to intent and fabric configurations in addition to associated state of the resource manager in terms of used resources on fabrics. DCNM backs up only when there's a configuration push. DCNM triggers the automatic backup only if you didn't trigger any manual backup after the last configuration push.

There are two types of automatic backup.

- **Hourly Fabric Backup:** You can enable an hourly backup.



Note MSD fabrics don't support hourly backup.

- **Scheduled Fabric Backup:** You can schedule a fabric backup for regular intervals.



Note In external fabrics, DCNM backs up the changes in the running configurations as well. The configuration push happens after a deploy. If you didn't deploy the changes, you can't back up them in an intent.

Hourly and scheduled backup processes happen only during the next periodic configuration compliance activity, and there can be a delay of up to an hour.

Hourly and Scheduled Backup of Fabrics

To enable automatic backup of fabric configurations and intents from the Cisco DCNM Web client, perform the following steps:

Procedure

- Step 1** Choose **Control > Fabrics > Fabric Builder**.
The **Fabric Builder** window appears.
- Step 2** Click the **Edit Fabric** icon for the fabric you want to backup.
- Step 3** Click the **Configuration Backup** tab.
- Step 4** Choose the nature of backup by checking the appropriate check box.

The valid options are **Hourly Fabric Backup** and **Scheduled Fabric Backup**. If you want to enable both the backups, check the **Hourly Fabric Backup** check box and the **Scheduled Fabric Backup** check box.

Note If you check the **Scheduled Fabric Backup** check box, specify the scheduled backup time in the **Scheduled Time** field. Enter the value in HH:MM format.

- Step 5** Click **Save**.
DCNM initiates the backup process after you click **Save**.
-

Backing Up Fabrics Manually

You can enable a manual backup for fabric configurations and intents. Regardless of the settings you choose under the **Configuration Backup** tab in the **Edit Fabric** dialog box, you can initiate a backup using this option. You cannot initiate standalone backups for a member fabric on an MSD fabric.

To initiate a manual backup of fabric configurations and intents from the Cisco DCNM Web UI, perform the following steps:

Procedure

- Step 1** Choose **Control > Fabrics > Fabric Builder**.
The **Fabric Builder** window appears.
- Step 2** Click the fabric for which you want to backup immediately.
The fabric topology window appears.
- Step 3** Click **Backup Now** in the **Actions** pane.
The **Backup Now** dialog appears.
- Step 4** Enter a tag name in the **Tag** field.
- Step 5** Click **OK**.
A confirmation message appears that the backup is triggered successfully.

Note The confirmation message only states that the backup is triggered and not if the backup is successful.

Step 6 (Optional) Click **Restore Fabric** from the **Actions** pane to confirm if the manual backup is successful or not. The manual backup is indicated in midnight blue. When you hover over the backup, the name has the tag you mentioned in *Step 4* confirming that it's a manual backup.

Golden Backup

You can now mark the backups that you don't want to delete even after you reach the archiving limit. These backups are the golden backups. You can't delete golden backups of fabrics. However, Cisco DCNM archives only up to 10 golden backups. You can mark a backup as golden backup while restoring the fabric. To mark a backup as golden in Cisco DCNM, perform the following steps from the Cisco DCNM Web UI:

Procedure

Step 1 Choose **Control > Fabrics > Fabric Builder** and select a fabric.

Step 2 Click **Restore Fabric** from the **Actions** menu.

The **Restore Fabric** window appears.

Step 3 Choose the time period from where you want to choose the backup.

Valid values are **1m**, **3m**, **6m**, **YTD**, **1y**, and **All**. You can zoom into the graph. By default the backup information for **1m**, which is one month, appears. You can also choose a custom date range. The backup information includes the following information:

- Backup date
- Total number of devices
- Number of devices in sync
- Number of devices out of sync

Step 4 Choose the backup you want to mark as golden by clicking the backup.

You can choose the automatic or manual backup. These backups are color-coded. Automatic backups are indicated in blue color. Manual backups are indicated in midnight blue color. Golden backups are indicated in orange color. The automatic backups have only the versions in their names. Whereas the manual backups have tag names, which you gave when you initiated a manual backup, along with the version in the backup name. Hover over a backup to see the name. The automatic backup is initiated from the **Configuration Backup** tab in the **Fabric Settings** dialog box. The manual backup is initiated by clicking **Backup Now** from the **Actions** pane in the fabric topology window.

Step 5 Check the **Mark backup as golden backup** check box to mark the backup as a golden backup.

A confirmation dialog box appears.

Step 6 Click **Yes**.

- Step 7** Continue with rest of the fabric restore procedure as mentioned in the *Restoring Fabrics* section or exit the window.
-

Validating Backups

DCNM validates all the backups when you initiate a fabric restore process. The validation includes the following checks:

- The DCNM release from which you want to restore: You can restore backups only from Cisco DCNM, Release 11.3(1), and Cisco DCNM, Release 11.4(1). Hence, if you upgrade from Cisco DCNM, Release 11.3(1) to Cisco DCNM, Release 11.4(1), you can restore a backup, which you archived before upgrading.
- Member fabrics composition: DCNM checks the name or ID of the member fabrics of an MSD fabric. If you change them after you back up, the restore won't proceed.
- Template validation: DCNM checks if templates from the backup match the templates in the current version. If you delete or rename any templates, you can't proceed with the restoring.
- Device composition of a fabric: If there are any changes to the inventory of switches after you back them up, you can't restore.

Restoring Fabrics

This section describes the fabric restoring for different types of fabrics. Cisco DCNM supports configuration restore at fabric level. Take a backup of the configuration to restore it.



Note After a backup and restore operation, the capabilities set to `/usr/local/cisco/dcm/java/jdk11/bin/java` are lost. As a result, some services running on privileged ports no longer start after a backup. To address this issue, execute the following CLI command after the restore:

```
setcap cap_net_bind_service=+ep /usr/local/cisco/dcm/java/jdk11/bin/java
```

Then restart the services.

Restoring Easy Fabrics

To restore an easy fabric in Cisco DCNM, perform the following steps from the Cisco DCNM Web UI:

Procedure

- Step 1** Choose **Control > Fabrics > Fabric Builder** and select a fabric.
- Step 2** Select **Restore Fabric** from the **Actions** menu.
- The **Restore Fabric** window appears.
- Step 3** Choose the time for which you want to restore the configuration.
- Valid values are **1m**, **3m**, **6m**, **YTD**, **1y**, and **All**. You can zoom into the graph. By default the backup information for **1m**, which is one month, appears. You can also select a custom date range. The backup information includes the following information:

- Backup date
- Total number of devices
- Number of devices in sync
- Number of devices out of sync

Step 4 Choose the backup you want to restore.

You can choose the automatic or manual backup. These backups are color-coded. Automatic backups are indicated in blue color. Manual backups are indicated in midnight blue color. Golden backups are indicated in orange color. The automatic backups have only the versions in their names. Whereas the manual backups have tag names, which you gave when you initiated a manual backup, along with the version in the backup name. Hover over a backup to see the name. The automatic backup is initiated from the **Configuration Backup** tab in the **Fabric Settings** dialog box. The manual backup is initiated by clicking **Backup Now** from the **Actions** pane in the fabric topology window.

Note If the fabric was a member of an MSD fabric and a backup was taken at the MSD-fabric level, that backup doesn't appear here. Only the standalone backups of the fabric taken before it was part of an MSD fabric appear here.

Step 5 Check the **Mark backup as golden backup** check box to mark the backup as a golden backup.

Step 6 Click **Next** to see the selected backup information of the devices in sync.

The switch name, switch serial number, IP address, and the delta configuration details of the devices appear.

Note If you add or remove devices from the fabric, the backup isn't valid. You can restore only the valid backups.

Step 7 Click **Get Config** to preview the configuration details.

Config Preview window appears, which has two tabs.

- **Backup Config:** This tab displays the backup configuration for the selected device.
- **Current Config:** This tab displays the current configuration for the selected device.

Step 8 Go back to **View Backup Summary** window.

Step 9 Click **Restore Intent** to proceed with the restoring.

The **Restore Status** window appears. You can view the status of the following:

- **Validating Backup**
- **Restoring fabric intent**
- **Restoring underlay intent**
- **Restoring interface intent**
- **Restoring overlay intent**

The valid values for the status of any action are **In Progress**, **Pending**, or **Failed**.

Note If the status of **Validating Backup** is **Failed**, other restoring actions won't be listed in this window.

- Step 10** Click **Next** after the intent is restored.
- The **Configuration Preview** window appears. You can view the following details in this window:
- Switch name
 - IP address
 - Switch serial number
 - Preview configuration
 - Status
 - Progress
- Step 11** Click **Deploy** to deploy the restored configuration.
- The **Configuration Deployment Status** window appears. You can view the details of the switch name, IP address, status, status description, and the progress.
- Step 12** Click **Close** after the restoring process is complete.
-

Restoring External Fabrics

When you restore an external fabric, the backed-up checkpoint is copied from DCNM to switches. To restore an external fabric in Cisco DCNM, perform the following steps from the Cisco DCNM Web UI:

Procedure

Step 1 Choose **Control > Fabrics > Fabric Builder** and select a fabric.

Step 2 Select **Restore Fabric** from the Actions menu.

The **Restore Fabric** window appears.

Step 3 Select the time for which you want to restore the configuration.

Valid values are **1m**, **3m**, **6m**, **YTD**, **1y**, and **All**. You can zoom into the graph. By default the backup information for **1m**, which is one month, appears.

When you select a backup version, the vertical bar representing it turns grey, and corresponding information is displayed at the bottom part of the screen. It includes the following information:

- Backup date
- DCNM Version
- Total number of devices
- Number of devices in sync
- Number of devices out of sync

You can select a custom date range either by rearranging the date slide below the vertical bars, or using the **From** and **To** boxes at the top right part of the screen.

Step 4 Choose the backup you want to restore.

You can choose the automatic or manual backup. These backups are color-coded. Automatic backups are indicated in blue color. Manual backups are indicated in midnight blue color. Golden backups are indicated in orange color. The automatic backups have only the versions in their names. Whereas the manual backups have tag names, which you gave when you initiated a manual backup, along with the version in the backup name. Hover over a backup to see the name. The automatic backup is initiated from the **Configuration Backup** tab in the **Fabric Settings** dialog box. The manual backup is initiated by clicking **Backup Now** from the **Actions** pane in the fabric topology window.

Note If the fabric was a member of an MSD fabric and if any backup was taken for the MSD fabric, that backup does not appear here. Only the standalone backups of the fabric taken before it was part of an MSD fabric appear here.

Step 5 (Optional) Check the **Mark backup as golden backup** check box to mark the backup as a golden backup.

Step 6 Click **Next** to see the selected backup information of the devices in sync.

The switch name, switch serial number, IP address, status, Restore Supported (indicating whether the device supports checkpoint rollback or not), the configuration details of the devices, and the VRF appear.

Note For information about the support for the checkpoint rollback feature in platforms, refer to the respective platform documentation.

By default, the management VRF is displayed in the VRF column because it is used for the copy operation during the restore process. If you want to use a different VRF for the copy operation, update the VRF column. To update the same VRF for all devices, use the Apply for all devices option at the bottom-left part of the screen. A sample screenshot:

Note If you added or removed devices to the fabric, you can't restore a fabric from the present day to a past date.

Step 7 Click **Get Config** to preview device configuration details.

The **Config Preview** window appears, which has three tabs.

- **Backup Config:** This tab displays the backup configuration for the selected device.
- **Current Config:** This tab displays the current running configuration of the selected device.
- **Side-by-side Comparison:** This tab displays current running configuration on the switch, and the backup configuration (or expected configuration).

Step 8 Go back to the **View Backup Summary** window.

Step 9 Click **Restore Intent** to proceed with the restoring.

The **Restore Status** window appears. You can view the status of the following:

- **Validating Backup**
- **Restoring fabric intent**
- **Restoring underlay intent**
- **Restoring interface intent**
- **Restoring overlay intent**
- **Intent Regeneration**

The valid values for the status of any action are **In Progress**, **Pending**, **Completed**, or **Failed**.

Note If the status of **Validating Backup** is **Failed**, other restoring actions won't be listed in this window.

Step 10 Click **Close** after the restore process is complete.

Restoring MSD Fabrics

When you restore an MSD fabric, the overlay information related to the MSD fabric is restored before restoring information related to the child fabrics. If there's any change in the inventory of the MSD fabric, the backup is considered to be invalid and the restore is blocked. You can't initiate a restore process for a member fabric. You get an error stating that the fabric is currently a member fabric of an MSD fabric. Move the member fabrics out of the MSD fabric to restore the previous standalone backups. Restoring an MSD fabric involves restoring fabric intent, underlay or interface intent, overlay intent, and intent regeneration.

You can also watch the video that demonstrates how to backup and restore an MSD fabric in Cisco DCNM. See the [MSD Fabric Backup and Restore](#) video.

To restore an easy fabric in Cisco DCNM, perform the following steps from the Cisco DCNM Web UI:

Procedure

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

Step 2 Choose an MSD fabric.

Step 3 Click **Restore Fabric** from the **Actions** menu.

The **Restore Fabric** wizard appears and you will be in the **Select Backup** step.

Note This option is not available for member fabrics, of an MSD fabric, from its corresponding fabric topology window.

Step 4 Choose the time for which you want to restore the configuration.

Valid values are **1m**, **3m**, **6m**, **YTD**, **1y**, and **All**. You can zoom into the graph. By default the backup information for **1m**, which is one month, appears. You can also select a custom date range. The backup information includes the following information:

- Backup date
- Total number of devices
- Number of devices in sync
- Number of devices out of sync

Step 5 Choose the backup you want to restore.

You can choose the automatic or manual backup. These backups are color-coded. Automatic backups are indicated in blue color. Manual backups are indicated in midnight blue color. Golden backups are indicated in orange color. The automatic backups have only the versions in their names. Whereas the manual backups have tag names, which you gave when you initiated a manual backup, along with the version in the backup name. Hover over a backup to see the name. The automatic backup is initiated from the **Configuration Backup**

tab in the **Fabric Settings** dialog box. The manual backup is initiated by clicking **Backup Now** from the **Actions** pane in the fabric topology window.

Note The standalone backups that you took for a member fabric before importing it to the MSD fabric do not appear here. Only the MSD backups appear here.

Step 6 Click the backup you want to restore.

The **Backup Summary** area appears. It includes the following information:

- Backup taken on: Timestamp when the backup was taken.
- DCNM version: DCNM version when the backup was taken.
- Backup version: Version of the backup, including the tag name if it's a manual backup.
- Total number of Fabrics: Specifies the total number of member fabrics imported in the MSD fabric.
- Total number of Easy Fabrics: Specifies the number of member fabrics that are easy fabrics.
- Total number of External Fabrics: Specifies the number of member fabrics that are external fabrics.
- Total number of devices: Specifies the total number of switches in all member fabrics.
- Number of devices in out of sync status: Specifies the number of devices that aren't in sync.
- Number of devices in unknown status: Specifies the number of devices whose status isn't known.
- Member fabrics: Specifies the names of the member fabrics Mark back as golden backup check box: (Optional) Check the **Mark backup as golden backup** check box to mark the backup as a golden backup.

Note If there are any devices with Out-of-Sync or Unknown status, the restore process will be blocked.

Step 7 Click **Next** to move to the **Restore Preview** step.

The **Easy Fabric** tab has information about the switch name, fabric name, switch serial, IP address, and the delta configuration of the member easy fabric. The **External Fabric** tab contains information about the switch name, fabric name, switch serial, IP address, switch status, configuration, and if the restore is supported for the member external fabric.

Note The backup isn't valid if devices are added or removed from the fabric. You can restore only the valid backups.

Step 8 Click **Restore Intent** to proceed to the **Restore Status** step in restoring.

The restore status and description appears for the member fabrics. Click the member fabric radio button to view the fabric-level progress of that fabric. The progress is automatically updated every 5 seconds.

Step 9 Click **Next** after the status is successful.

The **Configuration Preview** window appears. You can view the details of the switch name, IP address, switch serial number, preview configuration, status, and the progress in this window.

- Note**
- You can click **Next** only if the status is **Completed**.
 - You can't go back to the previous step because the fabric configurations change.
 - If the restoring failed, the fabric will be rolled back to the previous configuration.

- Step 10** Click **Deploy** to deploy the restored configuration.
- The **Configuration Deployment Status** window appears. You can view the following details:
- Switch name
 - IP address
 - Status
 - Status description
 - Progress
- Step 11** Click **Close** after the restoring process is complete.
-

Restoring a Switch

From Cisco DCNM, Release 11.5(1), you can restore a Cisco Nexus switch in external fabrics and LAN classic fabrics from the Cisco DCNM Web UI. The information you restore at switch-level is extracted from the fabric-level backups. The switch-level restoring doesn't restore fabric-level intents and other configurations applied using the fabric settings. Only switch-level intents are restored. Therefore, after you restore a switch, it might go out-of-sync because the fabric-level intents aren't restored. Perform a fabric-level restore to restore the intents as well. You can restore only one switch at a time. You can't restore a switch if the fabric where it's discovered is part of an MSD fabric.

To restore a switch in Cisco DCNM, perform the following steps from the Cisco DCNM Web UI:

Procedure

- Step 1** Choose **Control > Fabrics > Fabric Builder**.
- Step 2** Choose an external fabric or LAN classic fabric.
- Step 3** Right-click a Cisco Nexus switch for which you want to restore the configurations.
- Step 4** Choose the **Restore Config** option.
- Alternatively, you can click **Tabular view** in the **Actions** pane and navigate to the **Switches** tab. Choose a Cisco Nexus switch by checking the check box and click **Restore**.
- For non-Nexus switches, the **Restore Config** option doesn't appear and the **Restore** button grays out.
- This option does not appear when you log in with the **network-operator** role or when the fabric is in monitor mode or freeze mode.
- The **Restore Switch** wizard appears and you are in the **Select Backup** step.
- Step 5** Choose the time for which you want to restore the configuration.
- Valid values are **1m**, **3m**, **6m**, **YTD**, **1y**, and **All**. You can zoom into the graph. By default the backup information for **1m**, which is one month, appears. You can also select a custom date range.
- Step 6** Choose the backup you want to restore.

You can choose the automatic, manual, or golden backup. These backups are color-coded. Automatic backups are indicated in blue color. Manual backups are indicated in midnight blue color. Golden backups are indicated in orange color. The automatic backups have only the versions in their names. Whereas the manual backups have tag names, which you gave when you initiated a manual backup, along with the version in the backup name. Hover over a backup to see the name. The automatic backup is initiated from the **Configuration Backup** tab in the **Fabric Settings** dialog box. The manual backup is initiated by clicking **Backup Now** from the **Actions** pane in the fabric topology window.

Step 7 Click the backup you want to restore.

The **Backup Summary** area appears. It includes the following information:

- Backup taken on: Timestamp when the backup was taken.
- DCNM version: DCNM version when the backup was taken.
- Backup version: Version of the backup, including the tag name if it's a manual backup.
- Total number of devices: Specifies the total number of switches in the fabric when the backup was taken.
- Number of devices in sync status: Specifies the number of devices that are in sync.
- Number of devices in out of sync status: Specifies the number of devices that aren't in sync.
- Number of devices in unknown status: Specifies the number of devices whose status isn't known.
- Mark backup as golden backup check box: (Optional) Check the **Mark backup as golden backup** check box to mark the backup as a golden backup. If you mark a backup as golden backup, the fabric-level backup is also marked as a golden backup.

Note Most of this information is at the fabric level, and may or may not directly impact the proceedings of the switch-level restore.

Step 8 Click **Next** to move to the **Restore Preview** step.

You can view information about the switch name, switch serial, IP address, status, restore supported, delta configuration and the VRF details.

Step 9 (Optional) Click **Get Config** to preview device configuration details.

The **Config Preview** window appears, which has three tabs.

- **Backup Config:** This tab displays the backup configuration for the selected device.
- **Current Config:** This tab displays the current running configuration of the selected device.
- **Side-by-side Comparison:** This tab displays current running configuration on the switch, and the backup configuration, which is the expected configuration.

Step 10 Click **Restore** to proceed to the **Restore Status** step in restoring.

The restore status and description appears for the switch.

Step 11 Click **Close** after the restoring process is complete.

- Note**
- You can't go back to the previous step because the fabric configurations change.
 - If the restoring failed, the switch rolls back to the previous configuration.

Deleting a VXLAN BGP EVPN Fabric

Choose **Control > Fabric Builder**. On the Fabric Builder page, click **X** on the rectangular box that represents the fabric. Ensure the following before deleting a fabric.

- Fabric devices should not be in transition such as migration into or out of the fabric, ongoing network or VRF provisioning, and so on. Delete a fabric after the transition is complete.
- Remove devices that are still attached to the fabric. Remove non-Cisco Nexus 9000 Series switches first and then remove the 9000 Series switches.

Post DCNM 11.5(1) Upgrade for VXLAN BGP EVPN, External, and MSD Fabrics

Note the following guidelines after you upgrade to the DCNM Release 11.5(1):

- As part of the upgrade from an earlier DCNM release, the fabric and associated templates are carried over to the DCNM Release 11.5(1).
- From DCNM 11.3(1), some of the policy templates from earlier DCNM releases are deprecated and are actively updated in each newer DCNM release. These policy templates are automatically removed after the upgrade if they are not found to be in use. This removal does not affect any operations and helps in reducing the number of policies displayed in the DCNM template library.
- Navigate to each fabric from the **Fabric Builder** window, and click **Save & Deploy** to deploy any changes.

If you encounter any new or unexpected pending configurations after you click **Save & Deploy**, refer [Configuration Compliance in DCNM, on page 280](#).



Caution Some configuration changes can be expected as part of this step. Therefore, perform it only during a scheduled maintenance window.

- Post DCNM upgrade from Release 11.2(1), you could see the following diff if the fabric has a border device (border, border spine, border gateway, etc):

```
route-map extcon-rmap-filter-v6 deny 20
  no match ip address prefix-list host-route-v6
route-map extcon-rmap-filter-v6 deny 20
  match ipv6 address prefix-list host-route-v6
```

The above config is expected and it is meant to correct the route-map definition. Deployment of this diff will correct the switch configuration. If the fabric was created as Greenfield before upgrade, no additional action is needed. If the fabric was created as Brownfield before upgrade with the wrong route-map configuration on the device, this config will be captured in a **switch_freeform** policy. Post upgrade, you

should edit the freeform policy to remove the CLI **match ip address prefix-list host-route-v6** before the deployment.

- After a multi-level upgrade from Cisco DCNM 10.4(2) or 11.0(1), you can change the VRF templates to **Default_VRF_Universal** or **Default_VRF_Extension_Universal** to enable **ipv6 address use-link-local-only**.

Changing ISIS Configuration from Level 1 to Level 2

This procedure shows how to change ISIS configuration on switches from Level 1 to Level 2 in a VXLAN fabric deployment.

1. Choose **Control > Fabrics > Fabric Builder**.
2. Click a fabric in the **Fabric Builder** window.
3. Click **Tabular view** under **Actions** menu.
4. Search for all the **base_isis** policies in the **Template** search field.
5. Select all the **base_isis** policies and click the **Delete** icon to delete policies
6. Click **Save & Deploy**.

After all the **base_isis** policies are deleted, DCNM considers the migrated brownfield fabric as a greenfield fabric and creates the **base_isis_level2** policies on the switches.

Configuration Compliance in DCNM

The entire intent or expected configuration defined for a given switch is stored in DCNM. When you want to push this configuration down to one or more switches, the configuration compliance (CC) module is triggered. CC takes the current intent, the current running configuration, and then comes up with the set of configurations that are required to go from the current running configuration to the current expected config so that everything will be In-Sync.

When performing a software or firmware upgrade on the switches, the current running configuration on the switches is not changed. Post upgrade, if CC finds that the current running configuration does not have the current expected configuration or intent, it reports an Out-of-Sync status. There is no auto deployment of any configurations. You can preview the diffs that will get deployed to get one or more devices back In-Sync.

With CC, the sync is always from the DCNM to the switches. There is no reverse sync. So, if you make a change out-of-band on the switches that conflicts with the defined intent in DCNM, CC captures this diff, and indicates that the device is Out-of-Sync. The pending diffs will undo the configs done out-of-band to bring back the device In-Sync. Note that such conflicts due to out-of-band changes are captured by the periodic CC run that occurs every 60 mins by default, or when you click the RESYNC option either on a per fabric or per switch basis. Note that you can also capture the out-of-band changes for the entire switch by using the CC REST API. For more information, see *Cisco DCNM REST API Guide, Release 11.2(1)*.

From Cisco DCNM Release 11.2(1), to improve ease of use and readability of deployed configurations, CC in DCNM has been enhanced with the following:

- All displayed configurations in DCNM are easily readable and understandable.
- Repeated configuration snippets are not displayed.

- Pending configurations precisely show only the diff configuration.
- Side-by-side diffs has greater readability, integrated search or copy, and diff summary functions.

The CC engine computes diff by comparing the intent with the running configuration on the switch ensuring that any configuration that is defined in the intent exists on the switch. For any component or configuration snippet that is defined in the intent, the CC engine ensures that the same component or configuration snippet exists on the switch by generating appropriate commands, if required, to match the switch configuration with the intent configuration.

Top-level configuration commands on the switch that do not have any associated DCNM intent are not checked for compliance by Configuration Compliance (CC). However, CC performs compliance checks, and attempts removal, of the following commands even if there is no DCNM intent:

- **configure profile**
- **apply profile**
- **interface vlan**
- **interface loopback**
- **interface Portchannel**
- Sub-interfaces, for example, **interface Ethernet X/Y.Z**
- **fex**
- **vlan** <vlan-ids>

CC performs compliance checks, and attempts removal, of these commands only when *Easy_Fabric_11_1* and *Easy_Fabric_eBGP* fabric templates are used. On *External_Fabric* templates, top-level configuration commands on the switch, including the commands mentioned above, that do not have any associated DCNM intent are not checked for compliance by CC.

We recommend using the DCNM freeform configuration template to create additional intent and deploy these commands to the switches to avoid unexpected behavior

Now, consider a scenario in which the configuration that exists on the switch has no relationship with the configuration defined in the intent. Examples of such configurations are a new feature that has not been captured in the intent but is present on the switch or some other configuration aspect that has not been captured in the intent. Configuration compliance does not consider these configuration mismatches as a diff. In such cases, Strict Configuration Compliance ensures that every configuration line that is defined in the intent is the only configuration that exists on the switch. However, configuration such as boot string, rommon configuration, and other default configurations are ignored during strict CC checks. For such cases, the internal configuration compliance engine ensures that these config changes are not called out as diffs. These diffs are also not displayed in the **Pending Config** window. But, the Side-by-side diff utility compares the diff in the two text files and does not leverage the internal logic used in the diff computation. As a result, the diff in default configurations are highlighted in red in the **Side-by-side Comparison** window.

Starting from Cisco DCNM Release 11.4(1), such diffs are not highlighted in the **Side-by-side Comparison** window. The auto-generated default configuration that is highlighted in the **Running config** window is not visible in the **Expected config** window.

Any configurations that are shown in the **Pending Config** window are highlighted in red in the **Side-by-side Comparison** window if the configurations are seen in the **Running config** window but not in the **Expected config** window. Also, any configurations that are shown in the **Pending Config** window are highlighted in green in the **Side-by-side Comparison** window if the configurations are seen in the **Expected config** window but not in the **Running config** window. If there are no configurations displayed in the **Pending Config** window, no configurations are shown in red in the **Side-by-side Comparison** window.

All freeform configurations have to strictly match the **show running configuration** output on the switch and any deviations from the configuration will show up as a diff during **Save & Deploy**. You need to adhere to the leading space indentations.

You can typically enter configuration snippets in DCNM using the following methods:

- User-defined profile and templates
- Switch, interface, overlay, and vPC freeform configurations
- Network and VRF per switch freeform configurations
- Fabric settings for Leaf, Spine, or iBGP configurations



Caution The configuration format should be identical to the **show running configuration** of the corresponding switch. Otherwise, any missing or incorrect leading spaces in the configuration can cause unexpected deployment errors and unpredictable pending configurations. If any unexpected diffs or deployment errors are displayed, check the user-provided or custom configuration snippets for incorrect values.

If DCNM displays the "Out-of-Sync" status due to unexpected pending configurations, and this configuration is either unable to be deployed or stays consistent even after a deployment, perform the following steps to recover:

1. Check the lines of config highlighted under the **Pending Config** tab in the **Config Preview** window.
2. Check the same lines in the corresponding **Side-by-side Comparison** tab. This tab shows whether the diff exists in "intent", or "show run", or in both with different leading spaces. Leading spaces are highlighted in the **Side-by-side Comparison** tab.
3. If the pending configurations or switch with an out-of-sync status is due to any identifiable configuration with mismatched leading spaces in "intent" and "running configuration", this indicates that the intent has incorrect spacing and needs to be edited.
4. To edit incorrect spacing on any custom or user-defined policies, navigate to the switch and edit the corresponding policy:
 - a. If the source of the policy is **UNDERLAY**, you will need to edit this from the Fabric settings screen and save the updated configuration.
 - b. If the source is blank, it can be edited from the **View/Edit policies** window for that switch.
 - c. If the source of the policy is **OVERLAY**, but it is derived from a switch freeform configuration. In this case, navigate to the appropriate **OVERLAY** switch freeform configuration and update it.
 - d. If the source of the policy is **OVERLAY** or a custom template, perform the following steps:
 1. Navigate to **Administration > DCNM Server > Server Properties**, set the **template.in_use.check** property to **false**. This allows the profiles or templates to be editable.
 2. Edit the specific profile or template from the **Control > Template Library** edit window, and save the updated profile template with the right spacing.
 3. Click **Save & Deploy** to recompute the diffs for the impacted switches.
 4. After the configurations are updated, set the **template.in_use.check** property to **true**, as it slows down the performance of the DCNM system, specifically for **Save & Deploy** operations.

To confirm that the diffs have been resolved, click **Save & Deploy** after updating the policy to validate the changes.



Note DCNM checks only leading spaces, as it implies hierarchy of the command, especially in case of multi-command sequences. DCNM does not check any trailing spaces in command sequences.

Example 1: Configuration Compliance in Switch Freeform Policy

Let us consider an example with an incorrect spacing in the Switch Freeform Config field.

The switch freeform policy is created as shown:

Edit Policy ✕

Policy ID: POLICY-30630 Template Name: switch_freeform
 Entity Type: SWITCH Entity Name: SWITCH

* Priority (1-1000):

General

Variables: * Switch Freeform Config

```

ip dhcp relay
ip dhcp relay information option
ip dhcp relay information option vpn
ip dhcp snooping
ip domain-lookup
ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25
ip pim ssm range 232.0.0.0/8
ipv6 dhcp relay
ipv6 switch-packets lla
  
```


After deploying this policy successfully to the switch, DCNM persistently reports the following diffs:

Config Preview - Switch 70.70.70.73

Pending Config Side-by-side Comparison

```

ip domain-lookup
ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25
ip pim ssm range 232.0.0.0/8
ipv6 dhcp relay
ipv6 switch-packets lla
configure terminal
  
```

After clicking the **Side-by-side Comparison** tab, you can see the cause of the diff. As seen below, the **ip pim rp-address** line has 2 leading spaces, while the running configuration has 0 leading spaces.

Config Preview - Switch 70.70.70.73

Pending Config | Side-by-side Comparison

To re-compute the *running config*, please click the Re-sync button on the previous screen. Lastly, to resolve unexpected diffs, please review the leading spaces and edit the appropriate policies to match show run outputs.

Running config	Expected config
281 description "vpc-peer-link"	description "vpc-peer-link"
282	no shutdown
283 spanning-tree port type network	spanning-tree port type network
284 switchport	switchport
285 switchport mode trunk	switchport mode trunk
286 vpc peer-link	vpc peer-link
287 ip dhcp relay	ip dhcp relay
288 ip dhcp relay information option	ip dhcp relay information option
289 ip dhcp relay information option vpn	ip dhcp relay information option vpn
290 ip dhcp snooping	ip dhcp snooping
291 ip domain-lookup	ip domain-lookup
292	ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25
293	ip pim ssm range 232.0.0.0/8
294	ipv6 dhcp relay
295	ipv6 switch-packets lla
296 ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25	ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25
297 ip pim ssm range 232.0.0.0/8	ip pim ssm range 232.0.0.0/8
298 ipv6 dhcp relay	ipv6 dhcp relay
299 ipv6 switch-packets lla	ipv6 switch-packets lla
300 line console	line console
301 line vty	line vty
302 ngoam install acl	ngoam install acl
303 nv overlay evpn	nv overlay evpn
304 nxapi http port 80	nxapi http port 80
305 rmon event 1 description FATAL(1) owner PMON@FATAL	
306	power redundancy-mode ps-redundant
307 rmon event 2 description CRITICAL(2) owner PMON@CRITICAL	
308 rmon event 3 description ERROR(3) owner PMON@ERROR	
309 rmon event 4 description WARNING(4) owner PMON@WARNING	

To resolve this diff, edit the corresponding Switch Freeform policy so that the spacing is correct.

Edit Policy

Policy ID: POLICY-30630 | Template Name: switch_freeform
 Entity Type: SWITCH | Entity Name: SWITCH

* Priority (1-1000):

General

Variables:

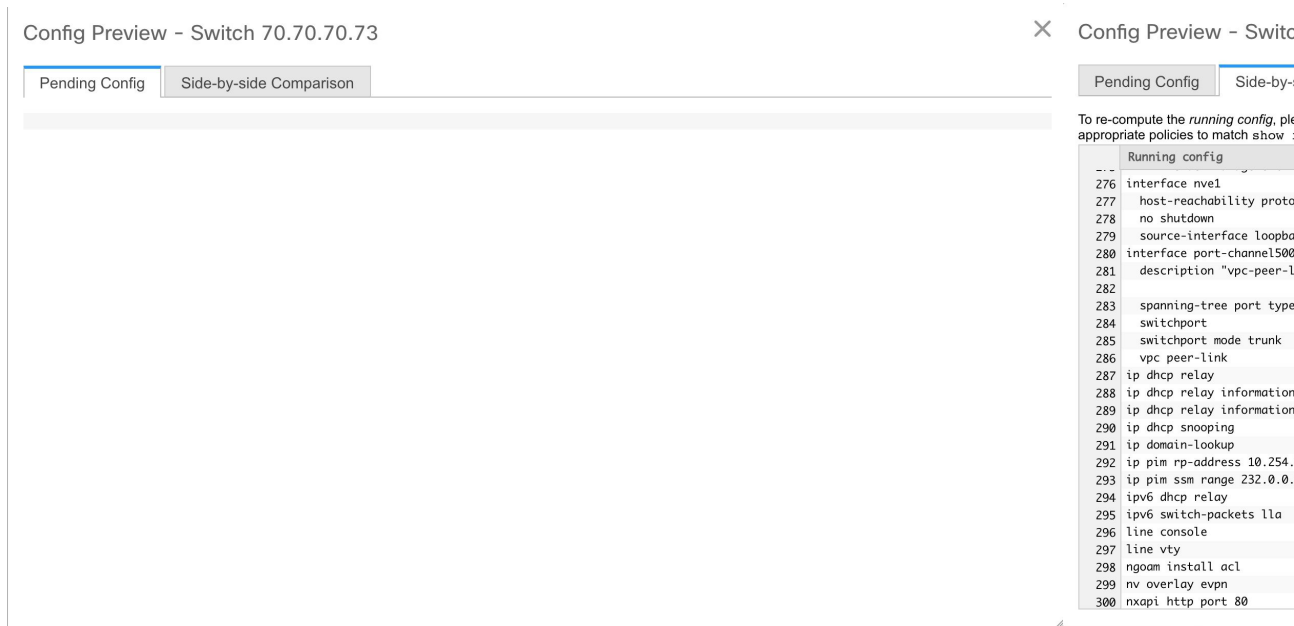
* Switch Freeform Config

```
ip dhcp relay
ip dhcp relay information option
ip dhcp relay information option vpn
ip dhcp snooping
ip domain-lookup
ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25
ip pim ssm range 232.0.0.0/8
ipv6 dhcp relay
ipv6 switch-packets lla
```

Save Push Config Cancel

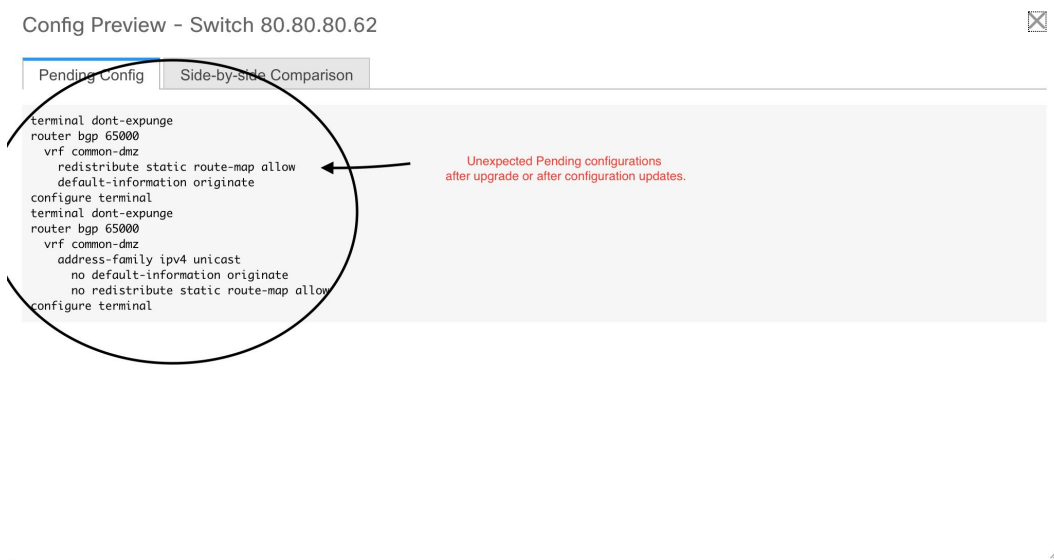
After you save, you can use the **Push Config** or **Save & Deploy** option to re-compute diffs.

As shown below, the diffs are now resolved. The **Side-by-side Comparison** tab confirms that the leading spaces are updated.



Example 2: Resolving a Leading Space Error in Overlay Configurations

Let us consider an example with a leading space error that is displayed in the **Pending Config** tab.



In the **Side-by-side Comparison** tab, search for diffs line by line to understand context of the deployed configuration.

terminal dont-expunge 0/0 SCOPE: green

Search for the Diffs, line by line in the Side-by-Side to understand context of the deployed configuration.

Matched count of 0, means this is some special configuration that DCNM has evaluated it needs to be pushed to the switch.

Config Preview - Switch 80.80.80.62

Pending Config Side-by-side Comparison

To re-compute the *running config*, please click the Re-sync button on the previous screen. Lastly, to resolve unexpected diffs, please review the leading spaces and edit the appropriate policies to match show run outputs.

Running config	Expected config
1 !Command: show running-config	
2	!Command: Intent from DCNM Fabric Builder. Any Intent not captured in Pending Config are defaults
3 !Running configuration last done at: Tue Jun 4 14:19:01 2019	
4	aaa group server radius radius
5 !Time: Tue Jun 4 16:03:38 2019	
6	use-vrf default
7 aaa group server tacacs+ ACS	aaa group server tacacs+ ACS
8 server 10.145.249.150	server 10.145.249.150
9 server 10.2.98.28	server 10.2.98.28
10 server 10.20.0.201	server 10.20.0.201
11 source-interface mgmt0	source-interface mgmt0
12 use-vrf management	use-vrf management
13 boot nxos bootflash:/nxos.9.2.3.bin	
14 cfs eth distribute	cfs eth distribute
15 configure profile Auto_Net_VNI20006_VLAN6	configure profile Auto_Net_VNI20006_VLAN6
16 evpn	evpn

A matched count of 0 means that it is a special configuration that DCNM has evaluated to push it to the switch.

redistribute static route-map 1/14 SCOPE: green

Searching for the next line in the pending configuration, shows the problem. The leading spaces are mismatched between running and expected configurations. 6 leading spaces in "Running configurations" and 4 leading spaces in "Expected Configuration". Similar mismatch is seen for "default-information originate" as well. For the VRF common-dmz as shown below.

Config Preview - Switch 80.80.80.62

Pending Config Side-by-side Comparison

To re-compute the *running config*, please click the Re-sync button on the previous screen. Lastly, to resolve unexpected diffs, please review the leading spaces and edit the appropriate policies to match show run outputs.

Running config	Expected config
2604 bfd	bfd
2605 remote-as 65000	remote-as 65000
2606 update-source loopback501	update-source loopback501
2607 router-id 192.168.0.4	router-id 192.168.0.4
2608 vrf common-dmz	vrf common-dmz
2609 address-family ipv4 unicast	address-family ipv4 unicast
2610 default-information originate	default-information originate
2611	
2612 redistribute static route-map allow	redistribute static route-map allow
2613	
2614 vrf common-mgmt	vrf common-mgmt
2615 address-family ipv4 unicast	address-family ipv4 unicast
2616 default-information originate	default-information originate
2617 redistribute static route-map allow	redistribute static route-map allow
2618 vrf ecd	vrf ecd
2619 address-family ipv4 unicast	address-family ipv4 unicast
2620 default-information originate	default-information originate
2621 redistribute static route-map allow	redistribute static route-map allow
2622 vrf ialab	vrf ialab
2623 address-family ipv4 unicast	address-family ipv4 unicast
2624 default-information originate	default-information originate
2625 redistribute static route-map allow	redistribute static route-map allow
2626 vrf lc	vrf lc

You can see that the leading spaces are mismatched between running and expected configurations.

Navigate to the respective freeform configs and correct the leading spaces, and save the updated configuration.

The screenshot shows the Cisco Data Center Network Manager (DCNM) interface. The main window displays the 'Network / VRF Deployment' section for a fabric named 'green'. A 'VRF Attachment - Attach VRFs for given switch(es)' dialog is open, showing a list of VRFs with 'COMMON-DMZ' selected. A 'Freeform Config (n9k12_bp2-lfsw01-I001)' window is also open, displaying the following configuration:

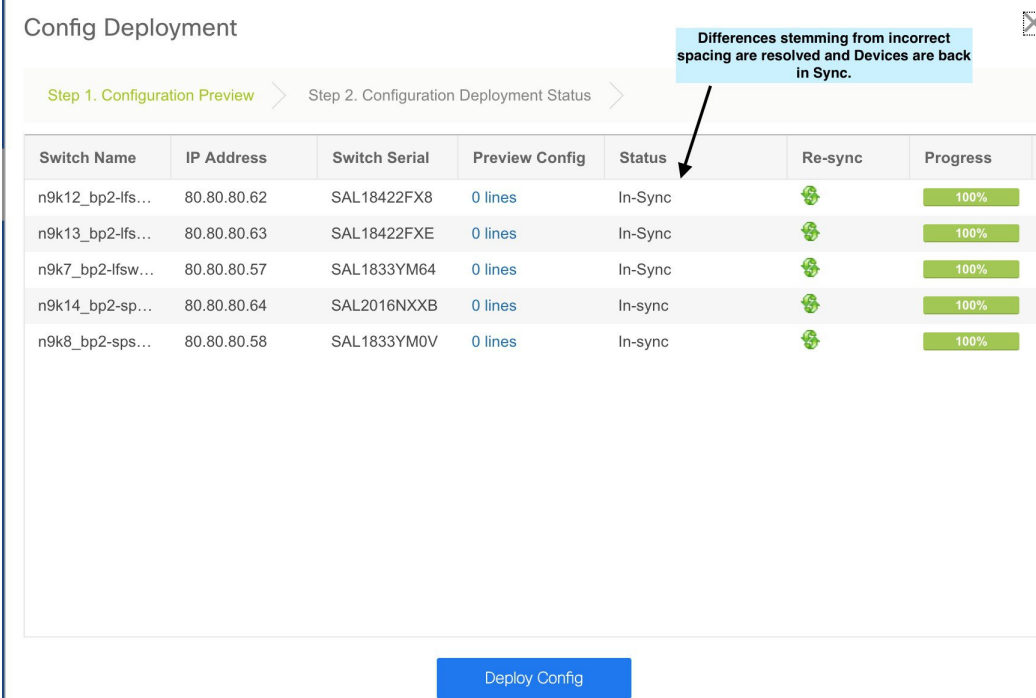
```

vrf context COMMON-DMZ
ip route 0.0.0.0/0 10.9.8.1 name COMMON-DMZ-DG
ip route 10.0.0.0/8 10.9.8.17 name OT-Networks-1-via-Mgmt&Tools-VDOM
ip route 10.9.16.0/20 10.9.8.17 name Mgmt&Tools-Networks
ip route 10.9.32.0/19 10.9.8.25 name EC-VDOM
ip route 10.9.128.0/19 10.9.8.33 name ECID-DG
ip route 10.9.254.0/23 10.9.8.9 name to-IA-LAB-VDOY
ip route 149.235.128.0/16 10.9.8.17 name OT-Networks-4-via-Mgmt&Tools-VDOM
ip route 172.16.0.0/12 10.9.8.17 name OT-Networks-5-2-via-Mgmt&Tools-VDOM
ip route 192.168.0.0/16 10.9.8.17 name OT-Networks-3-via-Mgmt&Tools-VDOM
router bgp 65000
vrf COMMON-DMZ
address-family ipv4 unicast
 redistribute static route-map allow
 default-information originate
  
```

A callout box with the text "Navigate to the corresponding VRF, Common-VRF and edit the freeform attachment to find the incorrectly spaced lines" points to the Freeform Config window.

Navigate to the **Fabric Builder** window for the fabric and click **Save & Deploy**.

In the **Config Deployment** window, you can see that all the devices are in-sync.



Config Deployment

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

Switch Name	IP Address	Switch Serial	Preview Config	Status	Re-sync	Progress
n9k12_bp2-lfs...	80.80.80.62	SAL18422FX8	0 lines	In-Sync		100%
n9k13_bp2-lfs...	80.80.80.63	SAL18422FXE	0 lines	In-Sync		100%
n9k7_bp2-lfsw...	80.80.80.57	SAL1833YM64	0 lines	In-Sync		100%
n9k14_bp2-sp...	80.80.80.64	SAL2016NXXB	0 lines	In-sync		100%
n9k8_bp2-sps...	80.80.80.58	SAL1833YM0V	0 lines	In-sync		100%

Deploy Config

Configuration Compliance in External Fabrics

With external fabrics, any Nexus switch can be imported into the fabric, and there is no restriction on the type of deployment. It can be LAN Classic, VXLAN, FabricPath, vPC, HSRP, etc. When switches are imported into an external fabric, the configuration on the switches is retained so that it is non-disruptive. Only basic policies such as the switch username and mgmt0 interface are created after a switch import.

In the external fabric, for any intent that is defined in the DCNM, configuration compliance (CC) ensures that this intent is present on the corresponding switch. If this intent is not present on the switch, CC reports an Out-of-Sync status. Additionally, there will be a Pending Config generated to push this intent to the switch to change the status to In-Sync. Any additional configuration that is on the switch but not in intent defined in DCNM, will be ignored by CC, as long as there is no conflict with anything in the intent.

When there is user-defined intent added on DCNM and the switch has additional configuration under the same top-level command, as mentioned earlier, CC will only ensure that the intent defined in DCNM is present on the switch. When this user defined intent on DCNM is deleted as a whole with the intention of removing it from the switch and the corresponding configuration exists on the switch, CC will report an Out-of-Sync status for the switch and will generate **Pending Config** to remove the config from the switch. This **Pending Config** includes the removal of the top-level command. This action leads to removal of the other out-of-band configurations made on the switch under this top-level command as well. If you choose to override this behavior, the recommendation is that, you create a freeform policy and add the relevant top-level command to the freeform policy.

Let us see this behavior with an example.

1. A **switch_freeform** policy defined by the user in DCNM and deployed to the switch.

Edit Policy ✕

Policy ID: POLICY-51710 **Template Name:** switch_freedom
Entity Type: SWITCH **Entity Name:** SWITCH

*** Priority (1-1000):**

General

*** Switch Freeform Config**

```
router bgp 1234
neighbor 10.2.0.1
  address-family l2vpn evpn
  send-community both
remote-as 1234
update-source loopback0
```

Variables:

- Additional configuration exists under **router bgp** in **Running config** that does not exist in user-defined DCNM intent **Expected config**. Note that there is no **Pending Config** to remove the additional config that exists on the switch without a user defined intent on DCNM.

Config Preview - Switch 172.29.21.130 ✕

Pending Config Side-by-side Comparison

To re-compute the *running config*, please click the Re-sync button on the previous screen. Lastly, to resolve unexpected diffs, please review the leading spaces and edit the appropriate policies to match show `run outputs`.

Running config	Expected config
593 rmon event 3 description ERROR(3) owner PMON@ERROR	
594 rmon event 4 description WARNING(4) owner PMON@WARNING	
595 rmon event 5 description INFORMATION(5) owner PMON@INFO	
596 route-map fabric-rmap-redis-subnet permit 10	
597 match tag 12345	
598 router bgp 1234	router bgp 1234
599 neighbor 10.2.0.1	neighbor 10.2.0.1
600 address-family l2vpn evpn	address-family l2vpn evpn
601 send-community both	send-community both
602 remote-as 1234	remote-as 1234
603 update-source loopback0	update-source loopback0
604 neighbor 20.2.0.2	
605 address-family ipv4 unicast	
606 send-community both	
607 router-id 10.2.0.2	
608 router ospf UNDERLAY	
609 router-id 10.2.0.2	
610 service dhcp	
611 snmp-server host 172.28.194.124 traps version 2c public udp-port 2162	
612 snmp-server host 172.28.194.126 traps version 2c public udp-port 2162	
613 snmp-server host 172.28.194.130 traps version 2c public udp-port 2162	
614 tacacs-server host 1.1.1.11 key 7 "cisco123"	
615 vdc N9K-z1 id 1	
616 limit-resource mroute-mem minimum 58 maximum 58	
617 limit-resource mroute-mem minimum 8 maximum 8	
618 limit-resource port-channel minimum 0 maximum 511	
619 limit-resource uroute-mem minimum 248 maximum 248	
620 limit-resource uroute-mem minimum 96 maximum 96	
621 limit-resource vlan minimum 16 maximum 4094	
622 limit-resource vrf minimum 2 maximum 4096	
623 version 7.0(3)I7(3)	
624 vlan 1	
625 vrf context management	vrf context management
626 ip route 0.0.0.0/0 172.29.21.1	ip route 0.0.0.0/0 172.29.21.1

Config Preview - Switch 172.29.21.130



Pending Config

Side-by-side Comparison

3. The **Pending Config** and the **Side-by-side Comparison** when the intent that was pushed earlier via DCNM is deleted from DCNM by deleting the switch_freeform policy that was created in the Step 1.

Config Preview - Switch 172.29.21.130



Pending Config

Side-by-side Comparison

To re-compute the *running config*, please click the Re-sync button on the previous screen. Lastly, to resolve unexpected diffs, please review the leading spaces and edit the appropriate policies to match `show run` outputs.

Running config	Expected config
584 ip domain-lookup	
585 ip pim rp-address 10.254.254.1 group-list 239.1.1.0/25	
586 ip pim ssm range 232.0.0.0/8	
587 ipv6 dhcp relay	
588 ipv6 switch-packets lla	
589 line console	
590 line vty	
591 ngoam install acl	
592 no password strength-check	no password strength-check
593 nv overlay evpn	
594 rmon event 1 description FATAL(1) owner PMON@FATAL	
595 rmon event 2 description CRITICAL(2) owner PMON@CRITICAL	
596 rmon event 3 description ERROR(3) owner PMON@ERROR	
597 rmon event 4 description WARNING(4) owner PMON@WARNING	
598 rmon event 5 description INFORMATION(5) owner PMON@INFO	
599 route-map fabric-rmap-redis-subnet permit 10	
600 match tag 12345	
601 router tag 1234	
602 neighbor 10.2.0.1	
603 address-family l2vpn evpn	
604 send-community both	
605 remote-as 1234	
606 update-source loopback0	
607 neighbor 20.2.0.2	
608 address-family ipv4 unicast	
609 send-community both	
610 router-id 10.2.0.2	
611 router ospf UNDERLAY	
612 router-id 10.2.0.2	
613 service dhcp	
614 snmp-server host 172.28.194.124 traps version 2c public udp-port 2162	
615 snmp-server host 172.28.194.126 traps version 2c public udp-port 2162	
616 snmp-server host 172.28.194.130 traps version 2c public udp-port 2162	
617 tacacs-server host 1.1.1.11 key 7 "cisco123"	
618 tacacs-server host 172.28.1.203 key 7 "Fewhg12345"	

Config Preview - Switch 172.29.21.130

Pending Config Side-by-side Comparison

```
no router bgp 1234
configure terminal
```

- A **switch_freeform** policy with the top-level **router bgp** command needs to be created. This enables CC to generate the configuration needed to remove only the desired sub-config which was pushed from DCNM earlier.

Edit Policy ✕

Policy ID: POLICY-51770 Template Name: switch_freeform
 Entity Type: SWITCH Entity Name: SWITCH

* Priority (1-1000):

General

Variables:

* Switch Freeform Config

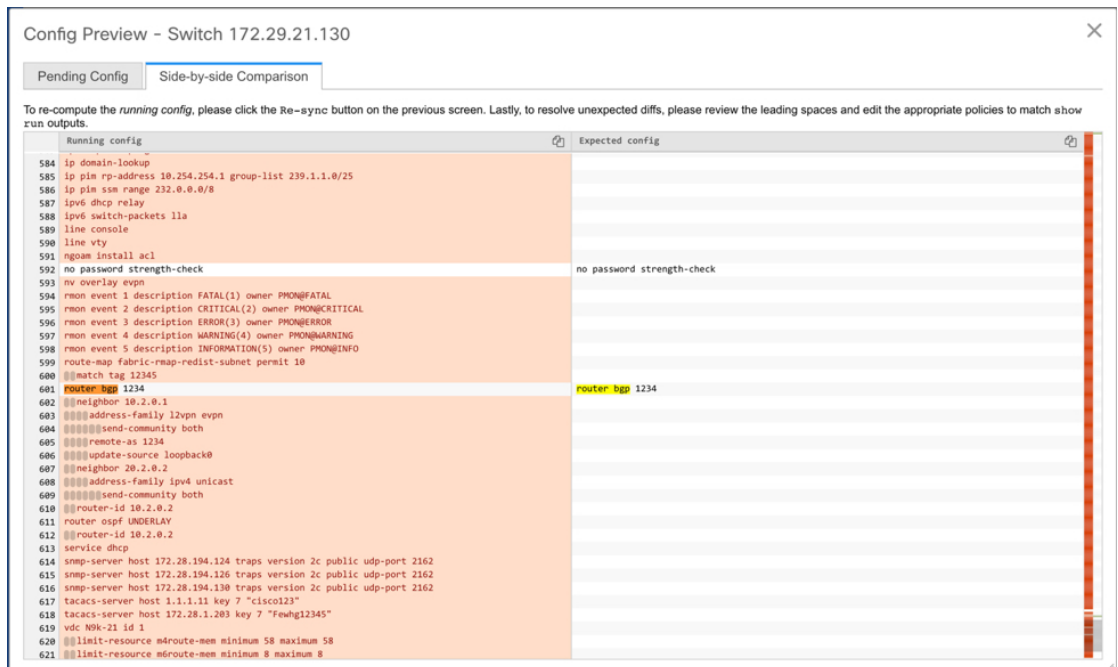
```
router bgp 1234
```

- The removed configuration is only the subset of the configuration that was pushed earlier from DCNM.

Config Preview - Switch 172.29.21.130

Pending Config Side-by-side Comparison

```
router bgp 1234
  no neighbor 10.2.0.1
configure terminal
```



For interfaces on the switch in the external fabric, DCNM either manages the entire interface or does not manage it at all. CC checks interfaces in the following ways:

- For any interface, if there is a policy defined and associated with it, then this interface is considered as managed. All configurations associated with this interface must be defined in the associated interface policy. This is applicable for both logical and physical interfaces. Otherwise, CC removes any out-of-band updates made to the interface to change the status to **In-Sync**.
- Interfaces created out-of-band (applies for logical interfaces such as port-channels, sub interfaces, SVIs, loopbacks, etc.), will be discovered by DCNM as part of the regular discovery process. However, since there is no intent for these interfaces, CC will not report an **Out-of-Sync** status for these interfaces.
- For any interface, there can always be a monitor policy associated with it in DCNM. In this case, CC will ignore the interface's configuration when it reports the **In-Sync** or **Out-of-Sync** config compliance status.

Special Configuration CLIs Ignored for Configuration Compliance

The following configuration CLIs are ignored during configuration compliance checks:

- Any CLI having 'username' along with 'password'
- Any CLI that starts with 'snmp-server user'

Any CLIs that match the above will not show up in pending diffs and clicking **Save & Deploy** in the **Fabric Builder** window will not push such configurations to the switch. These CLIs will not show up in the **Side-by-side Comparison** window also.

To deploy such configuration CLIs, perform the following procedure:

1. Select **Control>Fabric Builder**, click **Tabular View**, and select a switch in the **Name** column or select **Control>Fabric Builder** and right-click on the device.
2. Click **View/Edit Policies** and click on + to add a new policy. The **Add Policy** window comes up.
3. Add a PTI with the required configuration CLIs using the **switch_freeform** template and click **Save**.
4. Select the created policy and click **Push Config** to deploy the configuration to the switch(es).

Resolving Diffs for Case Insensitive Commands

By default, all diffs generated in DCNM while comparing intent, also known as Expected Configuration, and Running Configuration, are case sensitive. However, the switch has many commands that are case insensitive, and therefore it may not be appropriate to flag these commands as differences. These outlier cases are captured in the **compliance_case_insensitive_clis.txt** text file.

There could be additional commands not included in the existing **compliance_case_insensitive_clis.txt** file that should be treated as case insensitive. If the pending configuration is due to the differences of cases between the Expected Configuration in DCNM and the Running Configuration, you can configure DCNM to ignore these case differences as follows:

1. Modify the following file on the DCNM file system:

```
/usr/local/cisco/dcm/dcnm/model-config/compliance_case_insensitive_clis.txt
```

The sample entries in **compliance_case_insensitive_clis.txt** file are displayed as:

```
[root@dcnm98 model-config]# pwd
/usr/local/cisco/dcm/dcnm/model-config
[root@dcnm98 model-config]# cat compliance_case_insensitive_clis.txt
"^(no |)interface\s+Port(.)"
"^(no |)interface\s+Loo(.)"
"^(no |)interface\s+Eth(.)"
"^update-source\s+Loo(.)"
"^vrf\s+"
"^hardware profile portmode\s+"
"^(.*)route-map\s+(.)"
"^(.*)neighbor-policy(.)"
"(no |)encapsulation\s+(.)"
"^(.*)alert-group\s+(.)"
"^streetaddress\s+(.)"
"^transport email\s+(.)"
"(no |)action\s+(.)"
"(no|)\s+d*\s+remark.*"
[root@dcnm98 model-config]# █
```

If newer patterns are detected during deployment, and they are triggering pending configurations, you can add these patterns to this file. The patterns need to be valid regex patterns.

This enables DCNM to treat the documented configuration patterns as case insensitive while performing comparisons.

2. Click **Save & Deploy** for fabrics to see the updated comparison outputs.

Resolving Config Compliance After Importing Switches

After importing switches in Cisco DCNM, configuration compliance for a switch can fail because of an extra space in the management interface (mgmt0) description field.

For example, before importing the switch:

```
interface mgmt0
  description SRC=SDS-LB-LF111-mgmt0, DST=SDS-LB-SW001-Fa0/5
```


After importing the switch and creating a config profile:

```
interface mgmt0
  description SRC=SDS-LB-LF111-mgmt0,DST=SDS-LB-SW001-Fa0/5
```

In this example, the space after the comma (,) is removed.

Preview Config - Switch (10.1.101.17) ✕

Pending Config | Side-by-side Comparison

To re-compute the *running config*, please click the Re-sync button on the right hand side.  Lastly, to resolve unexpected diffs, please review the leading spaces and edit the appropriate policies to match show run output.

	Running config	Expected config
381	mtu 9216	mtu 9216
382	spanning-tree port type edge trunk	spanning-tree port type edge trunk
383	switchport mode trunk	switchport mode trunk
384	switchport trunk allowed vlan none	switchport trunk allowed vlan none
385	interface loopback0	interface loopback0
386	description Routing loopback interface	description Routing loopback interface
387	ip address 10.1.1.4/32	ip address 10.1.1.4/32
388	ip router ospf UNDERLAY area 0.0.0.0	ip router ospf UNDERLAY area 0.0.0.0
389	interface loopback1	interface loopback1
390	description VTEP loopback interface	description VTEP loopback interface
391	ip address 10.1.2.1/32	ip address 10.1.2.1/32
392	ip router ospf UNDERLAY area 0.0.0.0	ip router ospf UNDERLAY area 0.0.0.0
393	interface mgmt0	interface mgmt0
394	description SRC=SDS-LB-LF111-mgmt0, DST=SDS-LB-SW001-Fa0/5	description SRC=SDS-LB-LF111-mgmt0, DST=SDS-LB-SW001-Fa0/5
395		description SRC=SDS-LB-LF111-mgmt0, DST=SDS-LB-SW001-Fa0/5
396	ip address 10.1.101.17/24	ip address 10.1.101.17/24
397	no cdp enable	no cdp enable
398	vrf member management	vrf member management
399	interface nve1	interface nve1
400	host-reachability protocol bgp	host-reachability protocol bgp
401	no shutdown	no shutdown
402	source-interface loopback1	source-interface loopback1
403	ip dhcp relay	ip dhcp relay
404	ip dhcp relay information option	ip dhcp relay information option

Navigate to Interface Manager and click the **Edit** icon after selecting the mgmt0 interface. Remove the extra space in the description.

Strict Configuration Compliance

From Cisco DCNM Release 11.3(1), strict configuration compliance checks for diff between the switch configuration and the associated intent and generates **no** commands for the configurations that are present on the switch but are not present in the associated intent. When you click **Save and Deploy**, switch configurations that are not present on the associated intent are removed. You can enable this feature by selecting the **Enable Strict Config Compliance** checkbox under the **Advanced** tab in the **Add Fabric** or **Edit Fabric** window. By default, this feature is disabled.

Edit Fabric



* Fabric Name :

* Fabric Template :

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

* Layer 2 Host Interface MTU ? (Min:1500, Max:9216). Must be an even number

* Power Supply Mode ? Default Power Supply Mode For The Fabric

* CoPP Profile ? Fabric Wide CoPP Policy. Customized CoPP policy should be provided when 'manual' is selected

Brownfield Overlay Network Name Format ? Generated network name should be < 64 characters

Enable VXLAN OAM ?

Enable Tenant DHCP ?

Enable NX-API ?

Enable NX-API on HTTP ?

Enable Policy-Based Routing (PBR) ?

Enable Strict Config Compliance ?

* Greenfield Cleanup Option ? Switch Cleanup Without Reload When PreserveConfig=no

Enable Precision Time Protocol (PTP) ?

PTP Source Loopback Id ? (Min:0, Max:1023)

PTP Domain Id ? Multiple Independent PTP Clocking Subdomains on a Single Network (Min:0, Max:127)

Enable MPLS Handoff ?

Underlay MPLS Loopback Id ? Used for VXLAN to MPLS SR/LDP Handoff (Min:0, Max:1023)

The strict config compliance feature is supported on the Easy Fabric templates - Easy_Fabric_11_1 and Easy_Fabric_eBGP. To avoid generating diff for commands that are auto-generated by the switch, such as vdc, rmon, and so on, a file that has a list of default commands is used by CC to ensure that diffs are not generated for these commands. This file is located at `/usr/local/cisco/dcm/dcnm/model-config/strict_cc_exclude_clis.txt`.

**Note**

- In case any diffs are generated after strict configuration compliance is enabled, the switch icon turns blue in color in the **Fabric Builder** window.

Example: Strict Configuration Compliance

Let us consider an example in which the **feature telnet** command is configured on a switch but is not present in the intent. In such a scenario, the status of the switch is displayed as **Out-of-sync** after a CC check is done.

Now, click **Preview Config** of the out-of-sync switch. As the strict config compliance feature is enabled, the **no** form of the **feature telnet** command appears under **Pending Config** in the **Preview Config** window.

Preview Config - Switch (172.28.194.33)



Pending Config

Side-by-side Comparison

```
no feature telnet
configure terminal
```

Click the **Side-by-side Comparison** tab to display the differences between the running configuration and the expected configuration. Starting from Cisco DCNM Release 11.3(1), the **Re-sync** button is also displayed at the top right corner under the Side-by-side Comparison tab in the Preview Config window. Use this option to resynchronize DCNM state when there is a large scale out-of-band change, or if configuration changes do not register in the DCNM properly.

Preview Config - Switch (172.28.194.33)

Pending Config Side-by-side Comparison

To re-compute the *running config*, please click the Re-sync button on the right hand side. Lastly, to resolve unexpected diffs, please review the leading spaces and edit the appropriate policies to match `show run` output.

Running config	Expected config
1 !Command: show running-config	
2 !Running configuration last done at: Tue Oct 1 15:17:38 2019	
3 !Time: Tue Oct 1 15:18:01 2019	
4 boot nxos bootflash:/nxos.7.0.3.I7.6.bin_fix	
5 copp profile strict	copp profile strict
6 feature bgp	feature bgp
7 feature lldp	feature lldp
8 feature ngoam	feature ngoam
9 feature nv overlay	feature nv overlay
10 feature nxapi	feature nxapi
11 feature ospf	feature ospf
12 feature pim	feature pim
13 feature telnet	
14 hostname n9k-z17-33	hostname n9k-z17-33
15 interface ethernet1/1	interface ethernet1/1
16 mtu 9216	mtu 9216
17 no shutdown	no shutdown
18 interface ethernet1/10	interface ethernet1/10
19 mtu 9216	mtu 9216
20 no shutdown	no shutdown
21 interface ethernet1/11	interface ethernet1/11
22 mtu 9216	mtu 9216
23 no shutdown	no shutdown
24 interface ethernet1/12	interface ethernet1/12
25 mtu 9216	mtu 9216

The re-sync operation does a full CC run for the switch and recollects “show run” and “show run all” commands from the switch. When you initiate the re-sync process, a progress message is displayed on the screen. During the re-sync, the running configuration is taken from the switch. The Out-of-Sync/In-Sync status for the switch is recalculated based on the intent defined in DCNM.

Now, close the **Preview Config** window and click **Save and Deploy**. The Strict configuration compliance feature then ensures that the running config on the switch does not deviate from the intent by pushing the **no** form of the **feature telnet** command to the switch. The diff between the configurations is highlighted. The diff other than the **feature telnet** command are default switch and boot configurations and are ignored by the strict CC check.

In Cisco DCNM Release 11.2(1) and earlier releases, you had to right-click on a switch in the Fabric builder window and select **Deploy Config** to display the **Config Deployment** window. You then had to click **Preview Config** for a specific switch to bring up the **Preview Config** window that displays the pending configuration for that switch. This leads to a scenario in which the user may think that the preview config is inadvertently being deployed on the switch. Starting from Cisco DCNM Release 11.3(1), you can right-click on a switch in the **Fabric Builder** window and select **Preview Config** to display the **Preview Config** window. This window displays the pending configuration that has to be pushed to the switch to achieve configuration compliance with the intent.

Custom freeform configurations can be added in DCNM to make the intended configuration on DCNM and the switch configurations identical. The switches are then in In-Sync status. For more information on how to add custom freeform configurations on DCNM, refer [Enabling Freeform Configurations on Fabric Switches](#).

Enabling Freeform Configurations on Fabric Switches

In DCNM, you can add custom configurations through freeform policies in the following ways:

1. Fabric-wide
 - On all leaf, border leaf, and border gateway leaf switches in the fabric, at once.
 - On all spine, super spine, border spine, border super spine, border gateway spine and border switches, at once.
2. On a specific switch at the global level.
3. On a specific switch on a per Network or per VRF level.

Leaf switches are identified by the roles Leaf, Border, and Border Gateway. The spine switches are identified by the roles Spine, Border Spine, Border Gateway Spine, Super Spine, Border Super Spine, and Border Gateway Super Spine.



Note You can deploy freeform CLIs when you create a fabric or when a fabric is already created. The following examples are for an existing fabric. However, you can use this as a reference for a new fabric.

Deploying Fabric-Wide Freeform CLIs on Leaf and Spine Switches

1. Click **Control > Fabric Builder**. The Fabric Builder screen comes up. A rectangular box represents each fabric.
2. Click the **Edit Fabric** icon (located on the top right part of the rectangular box) for adding custom configurations to an existing fabric. The **Edit Fabric** screen comes up.
(If you are creating a fabric for the first time, click **Create Fabric**).
3. Click the **Advanced** tab and update the following fields:

Leaf Freeform Config – In this field, add configurations for all leaf, border leaf, and border gateway leaf switches in the fabric.

Spine Freeform Config - In this field, add configurations for all Spine, Border Spine, Border Gateway Spine, Super Spine, Border Super Spine, and Border Gateway Super Spine switches in the fabric.



Note Copy-paste the intended configuration with correct indentation, as seen in the running configuration on the Nexus switches. For more information, see [Resolving Freeform Config Errors in Switches, on page 303](#).

4. Click **Save**. The fabric topology screen comes up.
5. Click **Save & Deploy** at the top right part of the screen to save and deploy configurations.

Configuration Compliance functionality will ensure that the intended configuration as expressed by those CLIs are present on the switches and if they are removed or there is a mismatch, then it will flag it as a mismatch and indicate that the device is Out-of-Sync.

Incomplete Configuration Compliance - On some Cisco Nexus 9000 Series switches, in spite of configuring pending switch configurations using the **Save & Deploy** option, there could be a mismatch between the intended and switch configuration. To resolve the issue, add a **switch_freeform** policy to the affected switch (as explained in the *Deploy Freeform CLIs on a Specific Switch* section). For example, consider the following persistent pending configurations:

```
line vty
logout-warning 0
```

After adding the above configurations in a policy and saving the updates, click **Save and Deploy** in the topology screen to complete the deployment process.

To bring the switch back in-sync, you can add the above configuration in a **switch_freeform** policy saved and deployed onto the switch.

Deploying Freeform CLIs on a Specific Switch

1. Click **Control > Fabric Builder**. The Fabric Builder screen comes up.
2. Click on the rectangular box that represents the fabric. The Fabric Topology screen comes up.



Note To provision freeform CLIs on a new fabric, you have to create a fabric, import switches into it, and then deploy freeform CLIs.

3. Right-click the switch icon and select the **View/edit policies** option.
The **View/Edit Policies** screen comes up.
4. Click +. The **Add Policy** screen comes up.
In the **Priority** field, the priority is set to 500 by default. You can choose a higher priority (by specifying a lower number) for CLIs that need to appear higher up during deployment. For example, a command to enable a feature should appear earlier in the list of commands.
5. From the **Policy** field, select **switch_freeform**.
6. Add or update the CLIs in the **Freeform Config CLI** box.
Copy-paste the intended configuration with correct indentation, as seen in the running configuration on the Nexus switches. For more information, see [Resolving Freeform Config Errors in Switches, on page 303](#).
7. Click **Save**.
After the policy is saved, it gets added to the intended configurations for that switch.
8. Close the policy screens. The Fabric Topology screen comes up again.
9. Right click the switch and click **Deploy Config**.

The **Save & Deploy** option can also be used for deployment. However, the **Save & Deploy** option will identify mismatch between the intended and running configuration *across all* fabric switches.

Pointers for switch_freeform Policy Configuration:

- You can create multiple instances of the policy.
- For a vPC switch pair, create consistent **switch_freeform** policies on both the vPC switches.
- When you edit a **switch_freeform** policy and deploy it onto the switch, you can see the changes being made (in the **Side-by-side** tab of the Preview option).

Freeform CLI Configuration Examples

Console line configuration

This example involves deploying some fabric-wide freeform configurations (for all leaf, and spine switches), and individual switch configurations.

Fabric-wide session timeout configuration:

```
line console
  exec-timeout 1
```

Console speed configuration on a specific switch:

```
line console
  speed 115200
```

IP Prefix List/Route-map configuration

IP prefix list and route-map configurations are typically configured on border devices. These configurations are global because they can be defined once on a switch and then applied to multiple VRFs as needed. The intent for this configuration can be captured and saved in a **switch_freeform** policy. As mentioned earlier, note that the config saved in the policy should match the **show run** output. This is especially relevant for prefix lists where the NX-OS switch may generate sequence numbers automatically when configured on the CLI. An example snippet is shown below:

Edit Policy

Policy ID: POLICY-79030
 Template: switch_freeform
 * Priority (1-1000): 500
 Entity Type: SWITCH
 Entity Name: SWITCH
 Description: prefixlist-rmaps

General

Variables: * Switch Freeform Config

```
ip extcommunity-list standard RT_NEXUS-TEST-VRF permit rt 50001:1202
ip extcommunity-list standard RT_FROM_BACKBONE_TO_NEXUS-TEST-VRF permit rt
59999:9999
ip prefix-list PL_EXPORT_NEXUS-TEST-VRF seq 10 permit 10.190.224.0/26 le 32
ip prefix-list PL_EXPORT_NEXUS-TEST-VRF seq 15 permit 10.190.224.128/26 le 32
ip prefix-list PL_EXPORT_NEXUS-TEST-VRF seq 20 permit 10.190.224.192/26 le 32
ip prefix-list PL_EXPORT_NEXUS-TEST-VRF seq 25 permit 10.190.224.64/30 le 32
ip prefix-list PL_EXPORT_NEXUS-TEST-VRF seq 30 permit 10.190.224.68/30 le 32
ip prefix-list PL_EXPORT_NEXUS-TEST-VRF seq 35 permit 10.190.224.74/32
ip prefix-list PL_EXPORT_NEXUS-TEST-VRF seq 5 deny 127.0.0.1/32
ip prefix-list PL_IMPORT_NEXUS-TEST-VRF seq 10 permit 0.0.0.0/0
ip prefix-list PL_IMPORT_NEXUS-TEST-VRF seq 5 deny 127.0.0.1/32
ip prefix-list PL_RED_DIRECT_NEXUS-TEST-VRF seq 5 permit 0.0.0.0/0 le 32
ip prefix-list PL_RED_HMM_NEXUS-TEST-VRF seq 5 permit 0.0.0.0/0 le 32
ip prefix-list PL_RED_STATIC_NEXUS-TEST-VRF seq 5 permit 0.0.0.0/0 le 32
```

Save Push Config Cancel

ACL configuration

ACL configurations are typically configured on specific switches and not fabric-wide (leaf/spine switches). When you configure ACLs as freeform CLIs on a switch, you should include sequence numbers. Else, there will be a mismatch between the intended and running configuration. A configuration sample with sequence numbers:

```
ip access-list ACL_VTY
  10 deny tcp 172.29.171.67/32 172.29.171.36/32
  20 permit ip any any
ip access-list vlan65-acl
  10 permit ip 69.1.1.201/32 65.1.1.11/32
  20 deny ip any any

interface Vlan65
  ip access-group vlan65-acl in
line vty
  access-class ACL_VTY in
```

If you have configured ACLs without sequence numbers in a **switch_freeform** policy, update the policy with sequence numbers *as shown in the running configuration of the switch*.

After the policy is updated and saved, right click the device and select the per switch **Deploy Config** option to deploy the configuration. Alternatively, use the **Save and Deploy** option in the fabric topology screen (within Fabric Builder) so that the fabric triggers Configuration Compliance and resolves the configuration mismatch.

Resolving Freeform Config Errors in Switches

Copy-paste the running-config to the freeform config with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. Otherwise, configuration compliance in DCNM marks switches as out-of-sync.

Let us see an example of the freeform config of a switch.

```
feature bash-shell
feature telemetry

clock timezone CET 1 0
# Daylight saving time is observed in Metropolitan France from the last Sunday in March
(02:00 CET) to the last Sunday in October (03:00 CEST)
clock summer-time CEST 5 Sunday March 02:00 5 Sunday October 03:00 60
clock protocol ntp

telemetry
  destination-profile
    use-vrf management
```

The highlighted line about the daylight saving time is a comment that is not displayed in the **show running config** command output. Therefore, configuration compliance marks the switch as out-of-sync because the intent does not match the running configuration.

Let us check the running config in the switch for the clock protocol.

```
spine1# show run all | grep "clock protocol"
clock protocol ntp vdc 1
```

You can see that **vdc 1** is missing from the freeform config.

In this example, let us copy-paste the running config to the freeform config.

Here is the updated freeform config:

```
feature bash-shell
feature telemetry

clock timezone CET 1 0
clock summer-time CEST 5 Sunday March 02:00 5 Sunday October 03:00 60
clock protocol ntp vdc 1

telemetry
  destination-profile
    use-vrf management
```

After you copy-paste the running config and deploy, the switch will be in-sync. When you click **Save & Deploy**, the **Side-by-side Comparison** tab in the **Config Preview** window provides you information about the difference between the defined intent and the running config.

Deploying Freeform CLIs on a Specific Switch on a Per VRF/Network basis

1. Click **Control > VRFs**. After choosing the appropriate fabric scope, the listing of the currently defined VRFs for the fabric shows up.

2. Create a new VRF by clicking the + button or select an existing VRF and click the **Continue** button on the top right.
3. The topology view for the fabric shows up. Switches to which the VRF is already deployed are highlighted in green. Other switches will be in gray color.
4. Select an individual switch. The VRF attachment form shows up listing the switch that is selected. In case of a vPC pair, both switches belonging to the pair will show up.
5. Under the CLI Freeform column, select the button labelled **Freeform config**. This option allows a user to specify additional configuration that should be deployed to the switch along with the VRF profile configuration.
6. Add or update the CLIs in the **Free Form Config** CLI box. Copy-paste the intended configuration with correct indentation, as seen in the running configuration on the Nexus switches. For more information, see [Resolving Freeform Config Errors in Switches](#).

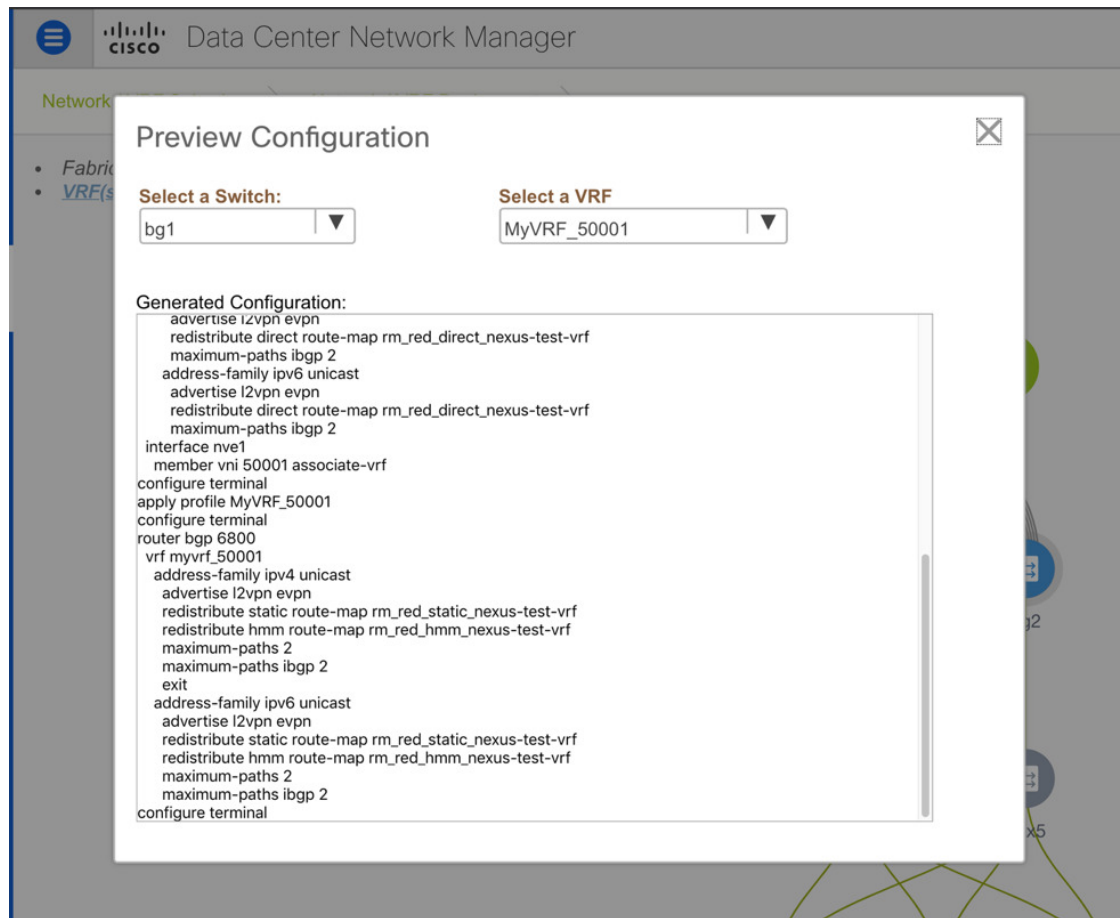
7. Click **Save Config**.



Note The **Freeform config** button will be gray when there is no per VRF per switch config specified. The button will be blue when some config has been saved by the user.

After the policy is saved, Click **Save** on the VRF Attachment pop-up to save the intent to deploy the VRF to that switch. Ensure that the checkbox on the left next to the switch is checked.

8. Now, optionally, click **Preview** to look at the configuration that will be pushed to the switch.



9. Click **Deploy** to push the configuration to the switch.

The same procedure can be used to define a per Network per Switch configuration.

VMM Workload Automation

VMM workload automation is about the automation of network configuration in Cisco's Nexus switches for workloads spawned in a VMware environment. Note that this is a preview feature in the Cisco DCNM Release 11.4(1).

You can also watch the video that demonstrates this automation. See [Video: VMM Workload Automation in Cisco DCNM](#).

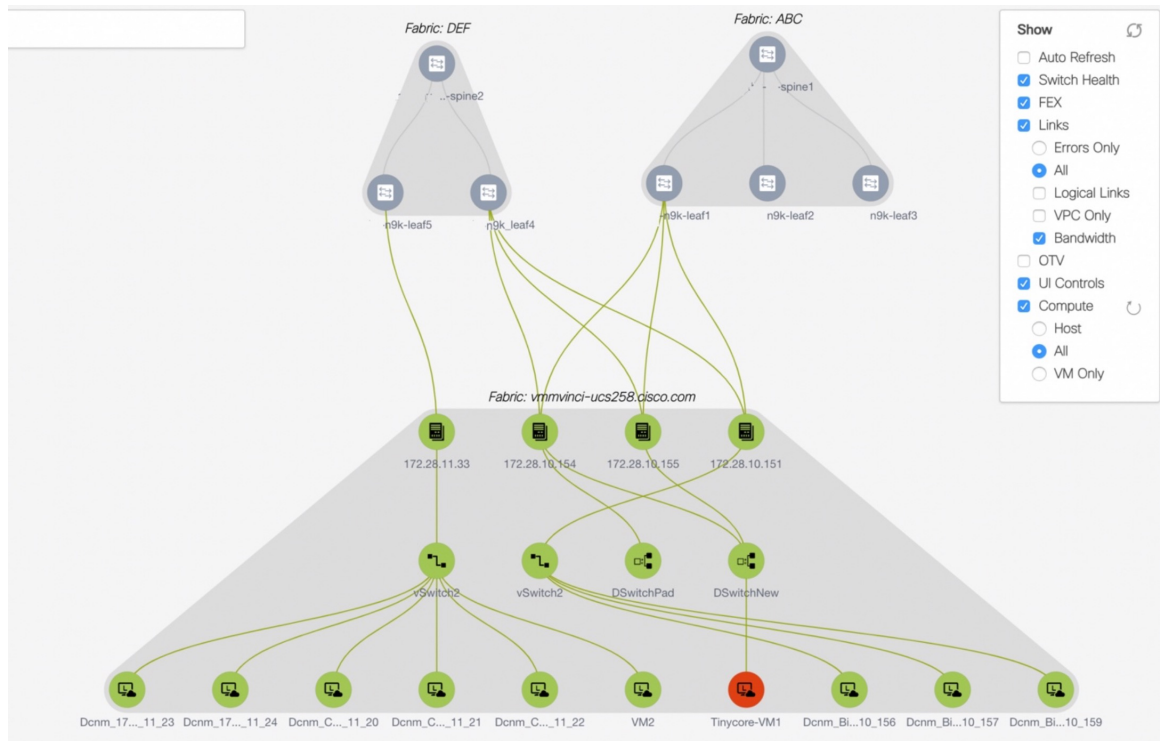
Overview of Network Objects in vCenter

VMM workload automation involves the mapping of network objects in vCenter to the network objects in DCNM. The following network objects in vCenter are considered:

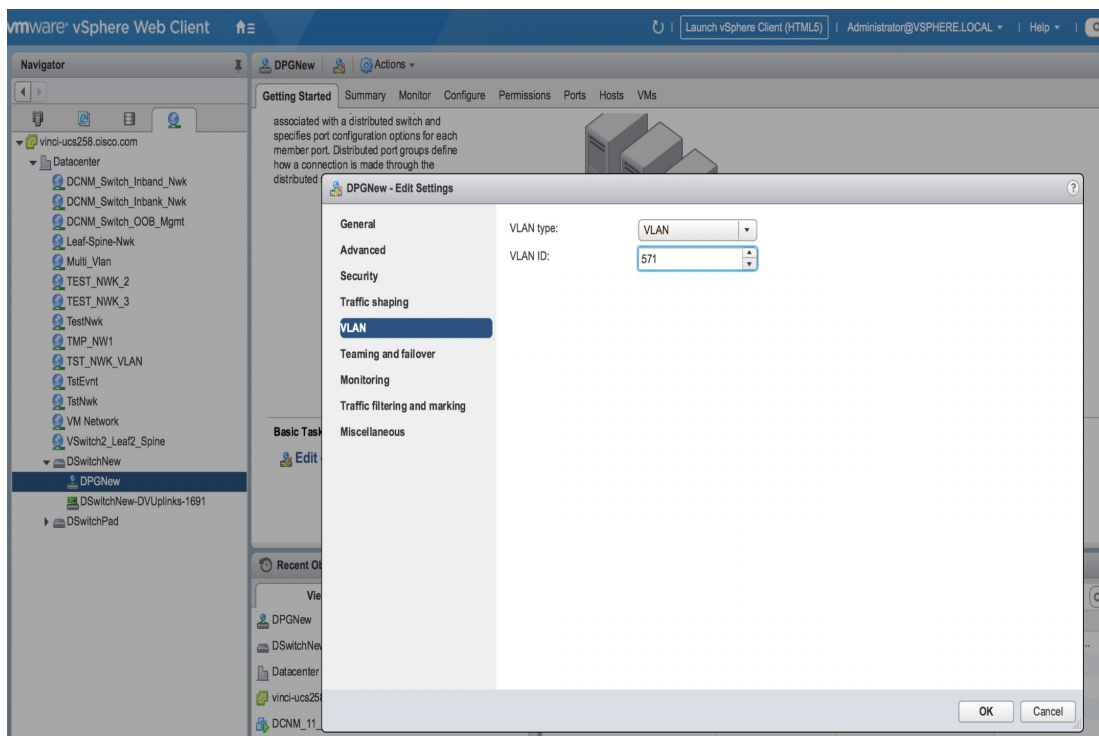
- **Virtual Switch (VS):** A regular VS runs in an ESXi host that performs software-based switching. A VS can have multiple port-groups (PG), where each PG has the network port configuration properties connecting to the network, like a VLAN. Each VS can have multiple uplink ports connecting to the leaf switches. Workloads spawned in the ESXi host can attach to the PG created in this VS.

- Distributed Virtual Switch (DVS): A DVS is a virtual switch that spans across multiple ESXi hosts. Similar to a regular VS, a DVS has multiple port-groups referred to as Distributed Port Groups (DPG). A DPG has the network port configuration properties connecting to the network, like a VLAN. Each DVS can have multiple uplink ports, which can be connected to the leaf switches. Workloads spawned in any of the hosts that are members of the DVS, can attach to the DPG. Through this document, and in the config file, a DPG is also referred as Distributed Virtual Port Group (DV-PG).

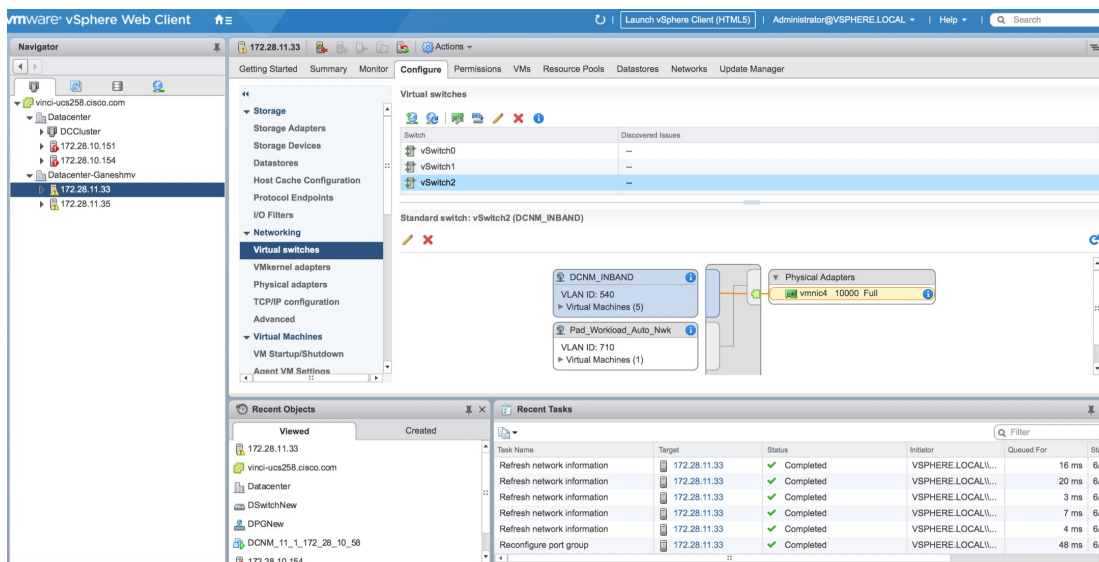
Let us consider the following topology in DCNM:



- There are four hosts in the setup with IP addresses:
 - 172.28.11.33
 - 172.28.10.154
 - 172.28.10.151
 - 172.28.10.155
- A DVS named **DSwitchNew** spawns across hosts 172.28.10.154 and 172.289.10.155. This DVS has a DPG named **DPGNew** that isn't shown in the figure. This DVS connects to switches n9k-leaf1 and n9k-leaf4 through uplink ports <vmnic3, vmnic1> and switch interfaces <e1/25, e1/7> respectively (not shown in the figure). A VLAN value of 571 is associated with the DPG **DPGNew**.



- There's also a regular vSwitch named vSwitch2 in host 172.28.11.33. This VS has a PG named DCNM_Inband. This VS connects to the leaf switch n9k-leaf5 through uplink port vmnic4 and switch interface e1/23. A VLAN value of 540 is associated with the PG **DCNM_Inband**.



How VMM Workload Automation Works

When workloads are spawned, they require provisioning in the network or fabric. The workloads spawned in the vCenter are either associated with a DPG or a PG. This DPG or PG in the VMWare should be mapped to a corresponding DCNM network. As an example in the earlier topology, when workloads are spawned in the

host 172.28.11.33 with PG **Inband**, then network provisioning that includes overlay and underlay configuration needs to happen in the leaf switch n9k-leaf5 with the relevant interface e1/23.

For the network provisioning to happen, each network object in a vCenter (a DPG or a PG) should be mapped to a network object in DCNM. A network object in DCNM has the following characteristics:

- VRF Name
- VLAN ID
- IPv4/IPv6 subnet and gateway information
- Secondary IPv4/v6 and gateway information
- BGP-EVPN configuration

A static mapping should be defined in a config file that maps the network object in a vCenter to a network object in DCNM. For information, see [Configuration Files for VMM Workload Automation, on page 308](#).

After the config file is populated, the workload automation module can be started. The module scans all the vCenters specified in the config file (conf.yml) and collects the following information for each vCenter:

- List of DVS and the DPG configured in all data centers.
- List of PGs configured in every host in all data centers.
- For every DPG or PG specified in the conf file, it finds the configured VLANs and the directly connected neighbor switches along with its interface information.
- For every <DVS, DPG> or <Host, PG> specified in the config file, it gets the associated network mapping in DCNM.

The module merges all the information and calls the DCNM APIs to provision or amend the networks in all the switches that are discovered as neighbors in one of the earlier steps.

Provisioning of a network or fabric uses DCNM top-down provisioning and it consists of the following steps:

1. Attaching the network configuration to the relevant interfaces of one or more switches that are discovered as neighbors. This attachment is done by the workload automation module.
2. After the configuration is attached, you can review the exact CLIs that are pushed to the switches.
3. After review, you can deploy the configuration to the switches. This deployment can be either done by the script based on the configuration file setting (default is **False**) or you can do it through DCNM. After this step, the configuration appears in the switches.

For more information, see <https://pypi.org/project/vmm-workload-auto/>.

Configuration Files for VMM Workload Automation

The following configuration files are used for VMM workload automation:

- Global YML File (conf.yml): This file has global configuration and access or credential information of the DCNM and the vCenters. Also, the location of the CSV file for each DCNM is specified in this file. For more information, see [Configuration Files for VMM Workload Automation, on page 308](#).
- CSV file (sample.csv): This file has the mapping of <DVS, DVS-PG> or <Host, PG> in vCenter to the Network name in DCNM. There's a separate CSV file for each DCNM. For more information, see [CSV File for Mapping Networks in vCenter and DCNM, on page 310](#).

Configuration File for Mapping DCNM and vCenter

The configuration file (conf.yml) specifies the DCNM IP address, username, and password. For each DCNM, the list of vCenter information like the IP address, username, and password is also specified. Multiple DCNMs can be specified in this conf.yml file. For every DCNM instance, there's an associated CSV file. The multi-DCNM case is only applicable when the script isn't run in a DCNM, but run in a server that has connectivity to all the DCNMs and vCenters specified in the config file.

In the config file, the hierarchy of information that should be specified is as follows:

```
Global config parameters
DCNM1
    DCNM1 config parameters including location of the CSV file
    vCenter1
        vCenter1 config parameters
    ...
    vCenter2
        vCenter2 config parameters
    ...
DCNM2
    ...
...
```

The location of this configuration file depends on the installation method of the VMM workload automation script. For more information, see *Installing VMM Workload Automation Script*. This file contains example entries. Modify it based on your environment.

The config file has the following entries:

LogFile: Specifies the name of the log file including the absolute path that will be used by the workload automation module for logging the errors and debug information. Make sure that the directory has write permission for creating the log file. For example, /tmp/workloadauto.log.

ListenPort: Specifies the port that the workload automation module uses to listen for the REST APIs, for example, 9590. Make sure that this port isn't used by any other application. You can check the same by running the **sudo netstat -tulpn** command.

AutoDeploy: Specifies whether the script should automatically deploy the configuration in the switches after attaching the networks. By default, it's set to **False** so that you can review the config and deploy it in the DCNM.

NwkMgr: Specifies the top-level section that contains the DCNM information. For multiple DCNM instances, repeat the fields with the appropriate values. For an example, see `conf_multiple_dcnm.yml` file that handles multiple DCNMs.

Ip: Specifies the IP address of DCNM, for example, 172.28.10.156.

User: Specifies the username used to log in to DCNM, for example, admin.

Password: Specifies the password of DCNM.

CsvFile: Specifies the absolute path of the location of the CSV file for this DCNM, for example, /etc/vmm_workload_auto/sample.csv.

ServerCtrlr: Specifies the information for the server controller, that is, vCenter/vSphere. For multiple vCenters that fall under this DCNM, this section repeats. For an example, refer the `conf_multiple_vcenter.yml` file, which contains multiple vCenters under a DCNM.

Ip: Specifies the IP address of the vCenter.

Type: Specifies the type of the server controller. The default is vCenter.

User: Refers to the username used to log in to the vCenter. For example, administrator@vsphere.local.

Password: Refers to the password for the vCenter.

The following example shows the contents of a conf.yml file:

```
LogFile: /tmp/workloadauto.log
ListenPort: 9590
AutoDeploy: false
NwkMgr:
- Ip: 172.28.10.151
  User: admin
  Password: C1sco_123
  CsvFile: /etc/sample.csv
  ServerCntlrlr:
  - Ip: 172.28.10.194
    Type: vCenter
    User: administrator@vsphere.local
    Password: Cisc0!23
```

CSV File for Mapping Networks in vCenter and DCNM

The CSV file contains the mapping of the network object in vCenter to the network created in DCNM. This file has the following entries in CSV format, that is, comma-separated entries. The reason for having a CSV file is to specify the mapping between a PG (or DPG) of vSphere to the network name of DCNM. It's a 1-1 mapping. However, since a PG or a DPG can't be identified on its own (not unique), you need an extra DVS name or Hostname to map it.

The CSV file contains the following fields:

vCenter - Specifies the IP address of vCenter

Dvs - Specifies the name of the DVS.

Dvs_pg - Specifies the DVS PG (DPG) in the DVS

Host - Specifies the Host/Server (IP address)

Host_pg - Specifies the port-group in the host.

Fabric - Specifies the fabric in DCNM.

Network - Specifies the name of the network already created in DCNM.

The network object is identified by a unique pair of either <DVS, DVS_PG> or <Host, Host_PG>.

Consider the following example:

vCenter Params					DCNM Params	
vCenter	DVS	DVPortGroup/ Network	ESXi Host	Port Group/ Network	Fabric Name	Network Name
172.28.12.123	DVSI	DPG1			Fab1	Network 10
172.28.12.123	DVSI	DPG1			Fab2	Network 30
172.28.12.123			172.28.12.11	PG10	Fab1	Network 20
172.28.12.123			172.28.12.12	PG20	Fab1	Network 20

This table has the mapping for vCenter 172.28.12.123. It has four entries:

- The first entry specifies that for DPG1 in DVS1, the network in DCNM is 'Network10' in fabric 'Fab1'. There can be cases where in the hosts of the DVS can connect to switches in multiple fabrics. The network name in each fabric can be different, so you need the fabric name as well. The example in the table shows one such case in the second entry.
- The second entry specifies the same <DVS1, DPG1> pair being mapped to Network 30 in the fabric 'Fab2'.
- The third entry specifies that for PG10 in the host 172.28.12.11, the network in DCNM is 'Network20' in fabric 'Fab1'.
- The fourth entry specifies that for PG20 in host 172.28.12.11, the network in DCNM is 'Network20' in fabric 'Fab1'.

As seen in the earlier table, the network object is identified by a unique pair of either <DVS, DVS_PG> or <Host, Host_PG>. If there's a value specified for DVS, DVS_PG, then the values for <Host, Host_PG> are blank. In other words, <DVS, DVS_PG> and <Host, Host_PG> are mutually exclusive.

When the earlier table is specified in a CSV format, it appears as below in the CSV file:

```
172.28.12.123,DVS1,DPG1,,,Fab1,Network10
172.28.12.123,DVS1,DPG1,,,Fab2,Network30
172.28.12.123,,,172.28.12.11,PG10,Fab1,Network20
172.28.12.123,,,172.28.12.12,PG20,Fab1,Network20
```

Let's consider more examples:

- **172.28.10.184,DSwitchPad,DSPad-PG2,,,DEF,MyNetwork_30000**

This line in the CSV file specifies the IP address of vCenter as 172.28.10.184 and the <DVS, DVS_PG> values are DSwitchPad, DSPad-PG2 respectively. Since the values for DVS, DVS-PG is specified, the values for Host, Host-PG are blank as seen in this example. The Fabric name is DEF and the network in DCNM is MyNetwork_30000.

- **172.28.10.184,,,172.28.11.33,Pad_Workload_Auto_Nwk,DEF,MyNetwork_60000**

In this example, the values for <DVS, DVS-PG> is left blank and the values for <Host, Host_PG> is specified as 172.28.11.33 and Pad_Workload_Auto_Nwk respectively. The fabric in DCNM is DEF and the network name in DCNM is MyNetwork_60000.

An example CSV file is as follows:

```
vCenter,Dvs,Dvs_pg,Host,Host_pg,Fabric,Network
172.28.10.184,DSwitchNew,DPGNew,,,DEF,MyNetwork_30000
172.28.10.184,DSwitchNew,DPGNew,,,ABC,MyNetwork_30000
172.28.10.184,,,172.28.11.33,Pad_Workload_Auto_Nwk,DEF,MyNetwork_60000
```

Installing and Starting the VMM Workload Automation Module

You can install the VMM workload automation module by using the PIP install or the install script.

Using PIP Install

Before you begin

This installation method is for users who are familiar with **pip install** and know how to set up the proxy or handle cases when there's a conflict in the python packages.

Procedure

- Step 1** Decide whether you want to run this module in a virtual environment or on a physical server. If you decide to run this on a server, ensure that you have the write permission for doing pip install.
- Step 2** Setup the `http_proxy`, `https_proxy`, and `no_proxy` appropriately.
- For example:
- ```
export http_proxy=http://proxy.esl.cisco.com:80
export https_proxy=https://proxy.esl.cisco.com:80
export no_proxy=127.0.0.1,172.28.10.0/24
```
- In this example, 172.28.10.0 specified in the `no_proxy` is the management subnet of DCNM.
- Step 3** Download and install the module from <https://pypi.org/>.
- ```
pip3 install vmm-workload-auto
```
- Similarly, you can uninstall the module using the command: **pip3 uninstall vmm-workload-auto**.
- Step 4** By default, the installation will happen in the following directories unless you override by giving options in the **pip** command.
- The package is installed under:
- ```
/usr/local/lib/python3.7/site-packages/vmm_workload_auto-0.1.1.dist-info
```
- The config files are installed under:
- ```
/usr/local/lib/python3.7/site-packages/etc/vmm_workload_auto
```
- The source code is placed under: `/usr/local/lib/python3.7/site-packages/workload_auto`
- Step 5** Edit the config files in:
- ```
/usr/local/lib/python3.7/site-packages/etc/vmm_workload_auto.
```
- For more information, see Configuration File for Mapping DCNM and vCenter.
- Make sure that the path of the CSV file specified in the **conf.yml** file is correct.
- Step 6** Start the VMM Workload automation module.
- The entry point for the python module is: `/usr/local/bin/vmm_workload_auto`.
- You can either run it as:
- ```
/usr/local/bin/vmm_workload_auto
```
- Or
- ```
vmm_workload_auto
```
- If `/usr/local/bin/` is already in **\$PATH**.



Provide the config file as a command-line option.

```
/usr/local/bin/vmm_workload_auto
--config=/usr/local/lib/python3.7/site-packages/etc/vmm_workload_auto/conf.yml
```

---

## Using the Install Script

Using the install script is an alternate method for users who don't want to use **pip install**. The install script performs the installation and starts the python module.

### Procedure

---

- Step 1** Navigate to <https://pypi.org/project/vmm-workload-auto/> and download the latest .tar.gz file.
  - Step 2** Untar it. For example:

```
tar -xvf vmm_workload_auto-0.1.0.tar.gz
```
  - Step 3** Modify the config/conf.yml and config/sample.csv according to your environment.
  - Step 4** Run the setup script as “source setup.sh”.
  - Step 5** The install script initially prompts the user to edit the conf.yml and .csv files. The script will then prompt the user for proxy and other details. After everything is complete, the script installs the python packages and starts the module automatically.
  - Step 6** For information about installing the script, see the Installation section in the README file at <https://pypi.org/project/vmm-workload-auto/>.
- 

## Post Installation

After running the workload automation module, navigate to the DCNM Networks window and check whether the network attachments are completed. Review the configuration and deploy it, if **AutoDeploy** is set to **false** in the configuration file (conf.yml).

## Additional Functionalities Using REST APIs

The workload automation module also provides the following REST APIs:



---

**Note** The REST APIs execute in a separate window after the VMM Workload automation module is running. Make sure that the automation module is running before running the REST APIs.

---

- Refresh - When the CSV file is changed, a refresh operation needs to be performed. This operation rereads the file and applies any new configuration if needed. The refresh API is as follows:

```
curl -XPOST http://127.0.0.1:{port}/workload_auto/refresh
```

- Resync - When there's any change in the DVS-PG, PG, VLAN, or neighbor switches, then a resync operation is needed. If there are any changes found, the configuration is reapplied accordingly. The resync API is as follows:

```
curl -XPOST http://127.0.0.1:{port}/workload_auto/resync
```

- Clean - In order to clean up the network provisioning that was previously done using the module, a cleanup operation is needed. The clean API is as follows:

```
curl -XPOST http://127.0.0.1:{port}/workload_auto/clean
```

## Events in vCenter

In the DCNM Release 11.4(1), real-time event processing isn't done using the module. The various relevant events and its significance to this module are as follows:

### Refresh

The refresh API enables the module to read the CSV file again and apply the network configuration to one or more the relevant switches. The refresh operation needs to be performed for the following events:

- Add PG: Create an entry in the CSV file that specifies the associated network in DCNM for this PG. After the entry is added, call the refresh REST API.
- Add DPG: Create an entry in the CSV file that specifies the associated network in DCNM for this DPG. After the entry is added, call the refresh REST API.

### Resync

The resync API enables the module to discover the network objects and its associated properties again. The result of this resync operation is applying the network configuration to the new or changed switches or interfaces. Perform the resync operation for the following events:

- Add host to a DVS.
- Modify VLAN in a DPG or PG.
- Change in topology: When any of the following information is changed, issue the Resync REST API to rediscover the topology and applying the REST API.
  - Neighbor switch change: This can happen if the attached leaf switch is replaced with a new switch or rewired to a different switch.
  - Interface change: This can happen due to rewiring to a different interface in the switch.
  - Host pNIC change.
  - Add an extra connection: This can happen when:
    - A regular interface in the host is made a port-channel by connecting an extra interface from the host to the switch.
    - An extra interface in the host connecting to a different switch forming a vPC pair.

### No Action Required

You need not perform any action for the following events:

- Add stand-alone host.
- Add vSwitch.

- Add DVS.
- Delete DVS.

### Mapping Change

The different scenarios for a mapping change in a CSV are as follows:

- If a new mapping is added, run the refresh API after adding the mapping in the CSV file.
- If the mapping between the vCenter network to the DCNM network needs to change, then run the clean REST API, modify the mapping the CSV file, and run the refresh REST API.
- If the existing mapping needs to be deleted, then run the clean REST API, delete the mapping in the CSV file and run the refresh API.

### Other Events

The other events and the operation that aren't part of a category are as follows:

- Host removed from DVS: When a host is removed from the DVS, the network configuration in the associated leaf switch and connected interface needs to be removed. This needs to be done for all the DPGs of this DVS. Navigate to DCNM and unattach the appropriate networks.
- DPG or PG Delete: For all the network mappings specified in the spec file that are associated with this DPG or PG, remove the network configuration in the relevant switches and interfaces. Navigate to DCNM and unattach the appropriate networks.
- Port Down or Switch Down: If the port or switch is permanently going to be offline, the configuration needs to be removed out of band. If the switch isn't reachable from the host, but it's still managed by DCNM, navigate to DCNM and unattach the appropriate networks.

## Management

The Management menu includes the following submenus:

## Resources

Cisco DCNM allows you to manage the resources. The following table describes the fields that appear on this page.

| Field      | Description                                                                                                                                                                                         |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Scope Type | Specifies the scope level at which the resources are managed. The scope types can be <b>Fabric</b> , <b>Device</b> , <b>DeviceInterface</b> , <b>DevicePair</b> , <b>Fabric</b> , and <b>Link</b> . |
| Scope      | Specifies the resource usage scope. Valid values are the switch serial numbers or fabric names. Resources with serial numbers are unique, and can be used on the serial number of the switch only.  |

| Field              | Description                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Allocated Resource | Specifies if the resources are managed with device, device interface, or fabric. Valid values are ID type, subnet, or IP addresses.                                                                                                                    |
| Allocated To       | Specifies the entity name for which the resource is allocated.                                                                                                                                                                                         |
| Resource Type      | Specifies the resource type. The valid values are <b>TOP_DOWN_VRF_LAN</b> , <b>TOP_DOWN_NETWORK_VLAN</b> , <b>LOOPBACK_ID</b> , <b>VPC_ID</b> , and so on.                                                                                             |
| Is Allocated?      | Specifies if the resource is allocated or not. The value is set to <b>True</b> if the resource is permanently allocated to the given entity. The value is set to <b>False</b> if the resource is reserved for an entity and not permanently allocated. |
| Allocated On       | Specifies the date and time of the resource allocation.                                                                                                                                                                                                |

## Allocating a Resource

To allocate a resource from the Cisco DCNM Web UI, perform the following steps:

### Procedure

- 
- Step 1** Choose **Control > Fabrics > Fabric Builder**.  
The **Fabric Builder** window appears.
- Step 2** Click the **Edit Fabric** icon in the fabric where you want to allocate the resource.  
The **Edit Fabric** dialog box appears.
- Note** Alternatively, you can navigate to the **Edit Fabric** dialog box from the fabric topology window. Click **Fabric Settings** in the **Actions** pane.
- Step 3** Choose the **Resources** tab.
- Step 4** Uncheck the **Manual Underlay IP Address Allocation** check box.  
If you check this check box, provide the IP addresses manually to all resources using the **Resource Allocation** window.
- Step 5** Click **Save**.
- Step 6** Choose **Control > Management > Resources**.  
The **Resource Allocation** window appears. This window lists all the resources under the selected scope.
- Step 7** Click the **Allocate Resource** icon.  
The **Allocate Resource** dialog box appears.
- Step 8** Choose the pool type, pool name, and scope type from the drop-down lists accordingly.  
The options for pool type are **ID**, **IP**, and **SUBNET**. Based on the pool type you choose, the values in the **Pool Name** drop-down list changes.

- Step 9** Choose the serial number in the **Serial Number** drop-down list.  
This field appears for all scope types except for the fabric scope type.
- Step 10** Enter the entity name in the **Entity Name** field.  
The embedded help gives example names for different scope types.
- Step 11** Enter the ID, IP address, or the subnet in the **Resource** field based on what pool type you chose in *Step 3*.
- Step 12** Click **Save** to allocate the resource.

## Examples to Allocate Resources

### Example 1: Assigning an IP to loopback 0 and loopback 1

```
#loopback 0 and 1
 L0_1: #BL-3
 pool_type: IP
 pool_name: LOOPBACK0_IP_POOL
 scope_type: Device Interface
 serial_number: BL-3 (FDO2045073G)
 entity_name: FDO2045073G~loopback0
 resource : 10.7.0.1

L1_1: #BL-3
pool_type: IP
pool_name: LOOPBACK1_IP_POOL
scope_type: Device Interface
serial_number: BL-3 (FDO2045073G)
entity_name: FDO2045073G~loopback1
resource : 10.8.0.3
```

### Example 2: Assigning a Subnet

```
#Link subnet
 Link0_1:
 pool_type: SUBNET
 pool_name: SUBNET
 scope_type: Link
 serial_number: F3-LEAF (FDO21440AS4)
 entity_name: FDO21440AS4~Ethernet1/1~FDO21510YPL~Ethernet1/3
 resource : 10.9.0.0/30
```

### Example 3: Assigning an IP to an Interface

```
#Interface IP
 INT1_1: #BL-3
 pool_type: IP
 pool_name: 10.9.0.8/30
 scope_type: Device Interface
 serial_number: BL-3 (FDO2045073G)
 entity_name: FDO2045073G~Ethernet1/17
 resource : 10.9.0.9
```

### Example 4: Assigning an Anycast IP

```
#ANY CAST IP
 ANYCAST_IP:
 pool_type: IP
```

```
pool_name: ANYCAST_RP_IP_POOL
scope_type: Fabric
entity_name: ANYCAST_RP
resource : 10.253.253.1
```

### Example 5: Assigning a Loopback ID

```
#LOOPBACK ID
LID0_1: #BL-3
pool_type: ID
pool_name: LOOPBACK_ID
scope_type: Device
serial_number: BL-3(FDO2045073G)
entity_name: loopback0
resource : 0
```

## Releasing a Resource

To release a resource from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Management > Resources**.
- The **Resource Allocation** window appears. This window lists all the resources under the selected scope.
- Step 2** Choose a resource that you want to delete.
- Note** You can delete multiple resources at the same time by choosing multiple resources.
- Step 3** Click the **Release Resource(s)** icon.
- A confirmation dialog box appears.
- Step 4** Click **Yes** to release the resource.
- 

## Adding, Editing, Re-Discovering and Removing VMware Servers

This section contains the following:

### Adding a Virtual Center Server

You can add a virtual center server from Cisco DCNM.

### Procedure

---

- Step 1** Choose **Control > Management > Virtual Machine Manager**.
- You see the list of VMware servers (if any) that are managed by Cisco DCNM-LAN in the table.
- Step 2** Click **Add**.

You see the **Add VCenter** window.

- Step 3** Enter the **Virtual Center Server** IP address for this VMware server.
  - Step 4** Enter the **User Name** and **Password** for this VMware server.
  - Step 5** Click **Add** to begin managing this VMware server.
- 

## Deleting a VMware Server

You can remove a VMware server from the Cisco DCNM.

### Procedure

---

- Step 1** Choose **Control > Management > Virtual Machine Manager**.
  - Step 2** Select the check box next to the VMware server that you want to remove and click **Delete** to discontinue data collection for that VMware server.
- 

## Editing a VMware Server

You can edit a VMware server from Cisco DCNM Web Client.

### Procedure

---

- Step 1** Choose **Control > Management > Virtual Machine Manager**.
  - Step 2** Check the check box next to the VMware server that you want to edit and click **Edit** virtual center icon.  
You see the **Edit VCenter** dialog box.
  - Step 3** Enter a the **User Name** and **Password**.
  - Step 4** Select managed or unmanaged status.
  - Step 5** Click **Apply** to save the changes.
- 

## Rediscovering a VMware Server

You can rediscover a VMware server from Cisco DCNM.

### Procedure

---

- Step 1** Choose **Control > Management > Virtual Machine Manager**.
- Step 2** Select the check box next to the VMware that you want to rediscover.
- Step 3** Click **Rediscover**.

A dialog box with warning "Please wait for rediscovery operation to complete." appears.

**Step 4** Click **OK** in the dialog box.

## Container Orchestrator

On Cisco DCNM Web UI, choose **Control > Management > Container Orchestrator**. You can add, delete, edit, and rediscover container types.

You can also watch the video that demonstrates how to use Container Visualization with Cisco DCNM. See [Video: Using Container Visualization in Cisco DCNM](#).

The following table describes the fields and description on Container Orchestrator window.

| Field             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Container Type    | Displays the type of orchestrator.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Cluster IP        | Displays the IP address of the Kubernetes cluster                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Cluster Name      | Specifies the name of the cluster                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Managed           | Specifies that the cluster is managed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Status            | <p>Displays the status of the cluster.</p> <ul style="list-style-type: none"> <li>• <b>Cert expired</b> implies that the certificate is expired. You must add certificate again.</li> <li>• <b>Not reachable</b> implies that DCNM can't reach the Kubernetes cluster.</li> <li>• <b>Ok</b> implies that the cluster is functioning correctly.</li> <li>• <b>Discovering</b> implies that the cluster is being discovered.</li> <li>• <b>Blank</b> implies that the cluster isn't managed.</li> </ul> <p><b>Note</b> Note: If the status is empty, it implies that the cluster isn't managed.</p> |
| User              | Specifies the role of the Kubernetes cluster                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Last Updated Time | Displays the time elapsed since the last change.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

The following table describes the action you can perform on the Container Orchestrator window.

| Field      | Description                                                                                                 |
|------------|-------------------------------------------------------------------------------------------------------------|
| Add        | Click <b>Add</b> icon to add a new cluster to the container orchestration. You can add multiple containers. |
| Delete     | Select the Kubernetes cluster and click <b>Delete</b> icon to delete.                                       |
| Edit       | Select the Kubernetes cluster and click on the <b>Edit</b> icon to edit the cluster.                        |
| Rediscover | Select the Kubernetes clusters and click <b>Rediscover</b> to refresh the cluster.                          |

You can perform the following actions on the Container Orchestrator:



## Adding Container Orchestrator

To add container orchestrator from the Cisco DCNM Web UI, perform the following steps:

### Before you begin

To add VM-based Kubernetes cluster, ensure that you have successfully configured the VMM on Cisco DCNM before enabling Container Orchestrator Visualization feature. You must add the vCenter, to the VMM, which hosts the VMs on which the VM-based Kubernetes cluster is running.

Ensure that the hostname is unique across all the clusters nodes.

You don't need VMM for Bare-metal-based cluster. For Bare-metal-based cluster, perform the following:

- Edit the server properties on **Web UI > Administration > DCNM Server > Server Properties** to enable LLDP on DCNM. In the **cdp.discover-lldp** field, enter **true** to enable LLDP.
- Ensure that the LLDP feature is enabled on all LEAF switches in the Fabric.
- On the Kubernetes cluster, ensure that LLDP and SNMP services are enabled on all Bare-metal nodes.
- If the Cisco UCS is using an Intel Nic, LLDP neighborship fails to establish due to FW-LLDP.

**Workaround** – For selected devices based on the Intel® Ethernet Controller (for example, 800 and 700 Series), disable the LLDP agent that runs in the firmware. Use the following command to disable LLDP:

```
echo 'lldp stop' > /sys/kernel/debug/i40e/<bus.dev.fn>/command
```

To find the *bus.dev.fn* for a given interface, run the following command and select the ID associated with the interface. The ID is highlighted in the below sample output.

```
[ucs1-lnx1]# dmesg | grep enp6s0
[12.609679] IPv6: ADDRCONF(NETDEV_UP): enp6s0: link is not ready
[12.612287] enic 0000:06:00.0 enp6s0: Link UP
[12.612646] IPv6: ADDRCONF(NETDEV_UP): enp6s0: link is not ready
[12.612665] IPv6: ADDRCONF(NETDEV_CHANGE): enp6s0: link becomes ready
[ucs1-lnx1]#
```




---

**Note** LLDP feature is enabled on those fabric switches, to which the bare-metal cluster nodes are connected. They can also be connected to the border gateway switches.

---

If the Fabric, to which the Kubernetes cluster is connected to, is discovered after the Cluster was discovered, you must rediscover the cluster to display the topology correctly.

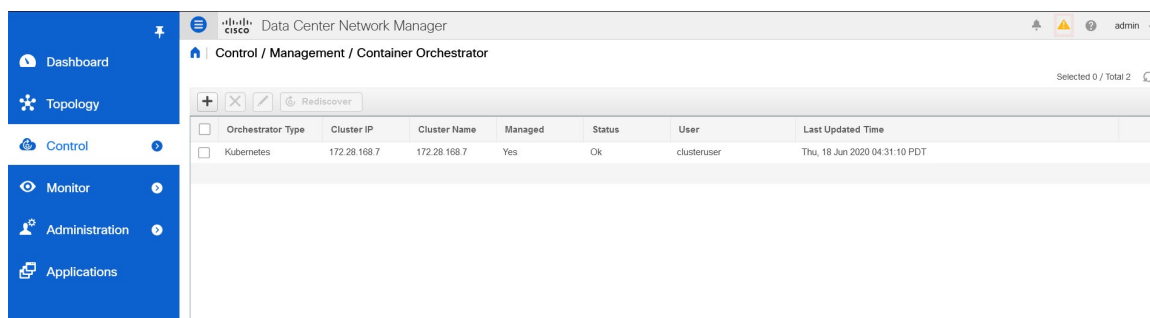
If the Bare-metal-based Kubernetes cluster is discovered after configuring LLDP, you must rediscover the Baremetal cluster to display the topology correctly.

### Procedure

---


#### Step 1

Choose **Control > Management > Container Orchestrator**.

**Step 2**

Click **Add**.

The Add Container Orchestrator appears.

**Step 3**

From the **Orchestrator** drop-down list, choose **Kubernetes**.

**Step 4**

In the **Cluster IP** field, enter the IP address of the Master node of the Kubernetes cluster.

**Step 5**

In the **User Name** field, enter the username of the API Client to connect to Kubernetes.

**Step 6**

Click **Fetch CSR** to obtain a Certificate Signing Request (CSR) from the Kubernetes Visualizer application.

**Note** This option is disabled until you enter a valid Cluster IP address and username.

Use the **Fetch CSR** only if you haven't obtained the SSL certificate. If you already have a valid certificate, you need not fetch the CSR.

Click **Download CSR**. The certificate details are saved in the <username>.csr in your directory. Paste the contents of the CSR to a file **kubereader.csr**, where, *kubereader* is the username of the API Client to connect to Kubernetes.

The CSR file name must adhere to naming convention <<username>>.csr.

**Note** As the certificates are generated on the Kubernetes cluster, you need Kubernetes admin privileges to generate certificates.

The script to generate the certificate **genk8sclientcert.sh** is located on the DCNM server at `./root/packaged-files/scripts/genk8sclientcert.sh` location.

**Step 7** Login to the Kubernetes cluster controller node.

**Note** You need admin privileges to generate the certificates.

**Step 8** Copy the `genk8sclientcert.sh` and `kubereader.csr` from the DCNM server location to the Kubernetes Cluster controller node.

**Note** Perform a “vnc cut and paste” operation to ensure that all the characters are copied correctly.

**Step 9** Generate the CSR for the user name, by using the **genk8sclientcert.sh** script.

```
(k8s-root)# ./genk8sclientcert.sh kubereader 10.x.x.x
```

where,

- `kubereader` is the username of the API Client to connect to Kubernetes. (as defined in Step [Step 5, on page 322](#)).
- `10.x.x.x` is the IP address of the DCNM server.

The following message is displayed, after the certificates are generated successfully:

```

The K8s CA certificate is copied into k8s_cluster_ca.crt file.
This to be copied into "Cluster CA" field.
The client certificate is copied into kubereader_10.x.x.x.crt file.
This to be copied into "Client Certificate" field.

```

There are two new certificates generated in the same location:

- `k8s_cluster_ca.crt`
- `username_dcnm-IP.crt`

For example: `kubereader_10.x.x.x.crt` (where, `kubereader` is the username, and `10.x.x.x` is the DCNM IP address)

**Step 10** Use the `cat` command to extract the certificate from these 2 files.

```
dcnm(root)# cat kubereader_10.x.x.x.crt
dcnm(root)# cat k8s_cluster_ca.crt
```

Provide these two certificates to the user, who is adding the Kubernetes cluster on Cisco DCNM.

**Step 11** Copy the content in the `kubereader_10.x.x.x.crt` to **Client Certificate** field.

**Note** Perform a “vnc cut and paste” operation to ensure that all the characters are copied correctly.

**Step 12** Copy the content in the `k8s_cluster_ca.crt` to the **Cluster Certificate** field.

**Note** Perform a “vnc cut and paste” operation to ensure that all the characters are copied correctly.

**Step 13** Click **Add** to add the container orchestrator.

Click **Cancel** to discard adding container orchestrator.

---

## Deleting Container Orchestrator

To delete container orchestrator from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

**Step 1** Choose **Control > Management > Container Orchestrator**.

**Step 2** Select the **Container Orchestrator** that you want to delete.

You can select more than one Cluster at a time.

Click **Delete**.

**Note** All the data will be deleted if you delete the Cluster. The Cluster will be removed from the Topology view also.

**Step 3** Click **Yes** on the confirmation message to delete the Container Orchestrator.

Click **No** to discard.

---

## Editing Container Orchestrator

To edit a container from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

**Step 1** Choose **Control > Management > Container Orchestrator**.

**Step 2** Select the **Container Orchestrator** that you want to modify. Click **Edit**.

The Edit Container Orchestrator window appears.

**Step 3** Modify the values appropriately.

You can update the Cluster and the Client certificates. You can also update the Managed status of the Kubernetes cluster. If you choose to update the Managed status, certificates are not required.

**Step 4** Click **Apply** to save the changes.

Click **Cancel** to discard.

---

## Rediscover Kubernetes Cluster

To rediscover Kubernetes cluster from the Cisco DCNM Web UI, perform the following steps:

## Procedure

- Step 1** Choose **Control > Management > Container Orchestrator**.
- Step 2** Select the **Container Orchestrator** that you want to rediscover.
- You can select more than one Cluster at a time.
- Click **Rediscover**.
- This action may take some time to refresh the container information.

## OpenStack Visualizer

On Cisco DCNM Web UI, choose **Control > Management > OpenStack Visualizer**. You can add, delete, edit, and rediscover OpenStack Clusters. Note that this is a preview feature.

For information about how to view OpenStack clusters in **Topology**, see [OpenStack Workload Visibility](#).

The following table describes the fields and description on the **OpenStack Visualizer** window.

| Field             | Description                                            |
|-------------------|--------------------------------------------------------|
| Cluster Type      | Specifies the type of cluster.                         |
| Cluster IP        | Specifies the Controller IP address of the cluster.    |
| Managed           | Specifies whether the cluster is managed or unmanaged. |
| Status            | Specifies the status of the cluster.                   |
| Username          | Specifies the username for the cluster.                |
| Project Name      | Specifies the project name.                            |
| Region            | Specifies the region.                                  |
| User Domain       | Specifies the user domain.                             |
| Project Domain    | Specifies the project domain.                          |
| Last Updated Time | Specifies the last updated time.                       |

The following table describes the action you can perform on the **OpenStack Visualizer** window.

| Field      | Description                                                                                 |
|------------|---------------------------------------------------------------------------------------------|
| Add        | Click <b>Add</b> icon to add a new OpenStack cluster to the container orchestration.        |
| Delete     | Select the OpenStack cluster and click <b>Delete</b> icon to delete.                        |
| Edit       | Select the OpenStack cluster and click on the <b>Edit</b> icon to edit the cluster details. |
| Rediscover | Select the OpenStack clusters and click <b>Rediscover</b> to refresh the cluster.           |

## Adding OpenStack Cluster

This task show how to add an OpenStack cluster.

### Before you begin

- Navigate to **Administration > DCNM Server > Server Properties**. Make sure to the set the **cdp.discover-lldp** property to **True** and click **Apply Changes**.

On the OpenStack cluster, ensure that the LLDP service is enabled on all the bare-metal nodes. LLDP feature is enabled on those fabric switches, to which the bare-metal cluster nodes are connected. They can also be connected to the border gateway switches.

- You can change the resync timer by using the **openstackviz.resync.timer** property. The default value is 60 minutes. Note that you can't set this value below 60 minutes. The resync function restarts the OpenStack plugin and rediscovers all the OpenStack clusters.
- For selected devices based on the Intel® Ethernet Controller (for example, 800 and 700 Series), disable the Link Layer Discovery Protocol (LLDP) agent that runs in the firmware. Use the following command to achieve the same:

```
echo 'lldp stop' > /sys/kernel/debug/i40e/bus.dev.fn/command
```

To find *bus.dev.fn* for a given interface, run the following command and select the ID associated with the interface. The ID is highlighted in the below output.

```
dmesg | grep eth0
[8.269557] enic 0000:6a:00.0 eno5: renamed from eth0
[8.436639] i40e 0000:18:00.0 eth0: NIC Link is Up, 40 Gbps Full Duplex, Flow Control:
None
[10.968240] i40e 0000:18:00.0 ens1f0: renamed from eth0
[11.498491] ixgbe 0000:01:00.1 eno2: renamed from eth0
```

### Procedure

**Step 1** Navigate to **Control > Management > OpenStack Visualizer**.

**Step 2** Click the **Add** icon to add an OpenStack Cluster.

- You should at least have read permissions to fetch the cluster information (for example, VMs and Host information).
- In DCNM Release 11.5(1), you can add a cluster only with a single project and region.

**Step 3** In the **Add OpenStack Cluster** window, specify the following details:

- **Orchestrator Type**: Specifies the type of orchestrator. By default, OpenStack is selected from this drop-down list.
- **Server IP**: Specifies the Controller IP address of the OpenStack cluster.
- **Port**: Specifies the port number.
- **Version**: Specifies the version.
- **Username and Password**: Specifies the username and password of the OpenStack cluster.

- **Project:** Specifies the project name.
- **Region:** Specifies the region. The default region is **RegionOne**.
- **User Domain:** Specifies the user domain. The default user domain is **default**.
- **Project Domain:** Specifies the project domain. The default project domain is **default**.
- **AMQP Endpoint:** Specifies a colon (:) separated multi-valued field containing the address details of an AMQP endpoint. The value should be specified in the format: **username:password:port**. The fields specify the following information:
  - **username:** Specifies the username of the AMQP endpoint.
  - **password:** Specifies the password of the AMQP endpoint.
  - **port:** Specifies the port number of the AMQP endpoint.

The default value for this field is **guest:guest:5672**.

**Step 4** Click **Add**.

After discovery, the status changes from **Discovering** to **Ok**. The information that is received from the OpenStack Cluster is appropriately organized and displayed on the main **Topology** window. An extra menu item labeled **OpenStack** appears on the **Show** pane.

---

## Editing OpenStack Cluster

### Procedure

---

**Step 1** Navigate to **Control > Management > OpenStack Visualizer**.

**Step 2** Select the OpenStack cluster that you want to modify. Click **Edit**.

In the **Edit OpenStack Cluster** window, you can edit the following fields:

- **Username and Password:** Specifies the username and password of the OpenStack cluster.
- **Managed:** You can select **unmanaged** to unmanage an OpenStack cluster.
- **AMQP Endpoint:** Specifies a colon (:) separated multi-valued field containing the address details of an AMQP endpoint. The value should be specified in the format: **username:password:port**. The fields specify the following information:
  - **username:** Specifies the username of the AMQP endpoint.
  - **password:** Specifies the password of the AMQP endpoint.
  - **port:** Specifies the port number of the AMQP endpoint.

The default value for this field is **guest:guest:5672**.

**Step 3** Click **Apply** to save the changes.

Click **Cancel** to discard.

---

## Deleting OpenStack Cluster

### Procedure

---

**Step 1** Navigate to **Control > Management > OpenStack Visualizer**.

**Step 2** Select the OpenStack cluster that you want to delete. Click **Delete**.

Upon deletion of a cluster from the inventory view, OpenStack plugin stops fetching and receiving the change notifications from the cluster, shuts down the connection with the removed cluster, and releases all software resources.

**Step 3** Click **Yes** on the confirmation message to delete the OpenStack cluster.

Click **No** to discard.

---

## Rediscovering OpenStack Cluster

### Procedure

---

**Step 1** Navigate to **Control > Management > OpenStack Visualizer**.

**Step 2** Select a specific cluster or all the clusters that you want to rediscover. Click **Rediscover**.

---

## Template Library

You can add, edit, or delete templates that are configured across different Cisco Nexus and Cisco MDS platforms using Cisco DCNM Web client. From Cisco DCNM Web client home page, choose **Control > Template Library > Templates**. The following parameters are displayed for each template that is configured on Cisco DCNM Web client. Templates support JavaScript. You can use the JavaScript function in a template to perform arithmetic operations and string manipulations in the template syntax.

The following table describes the fields that appear on this page.

**Table 1: Templates Operations**

| Field                | Description                                                        |
|----------------------|--------------------------------------------------------------------|
| Add Template         | Allows you to add a new template.                                  |
| Modify/View Template | Allows you to view the template definition and modify as required. |



| Field                    | Description                                                                                                                                                                                                   |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Save Template As         | Allows you to save the selected template in a different name. You can edit the template as required.                                                                                                          |
| Delete Template          | Allows you to delete a template                                                                                                                                                                               |
| Import Template          | Allows you to import a template from your local directory, one at a time.                                                                                                                                     |
| Export template          | Allows you to export the template configuration to a local directory location.                                                                                                                                |
| Import Template Zip File | Allows you to import .zip file, that contains more than one template that is bundled in a .zip format<br><br>All the templates in the ZIP file are extracted and listed in the table as individual templates. |



**Note** Notifications appear next to **Import Template Zip File** if there are issues while loading templates after restarting the server. Click the notifications to see the errors in the **Issues in loading Template** window. Templates with errors are not listed in the **Templates** window. To import these templates, correct the errors, and import them.

From Cisco DCNM Release 11.4(1), you can only view templates with the **network-operator** role. You cannot modify or save templates with this role. However, you can create or modify templates with the **network-stager** role.

**Table 2: Template Properties**

| Field                 | Description                                                                                                                                                                                             |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Template Name         | Displays the name of the configured template.                                                                                                                                                           |
| Template Description  | Displays the description that is provided while configuring templates.                                                                                                                                  |
| Tags                  | Displays the tag that is assigned for the template and aids to filter templates based on the tags.                                                                                                      |
| Supported Platforms   | Displays the supported Cisco Nexus platforms compatible with the template. Check the check box of platforms that are supported with the template.<br><br><b>Note</b> You can select multiple platforms. |
| Template Type         | Displays the type of the template.                                                                                                                                                                      |
| Template Sub Type     | Specifies the sub type that is associated with the template.                                                                                                                                            |
| Template Content Type | Specifies if it is Jython or Template CLI.                                                                                                                                                              |

Table 3: Advanced Template Properties

| Field        | Description                                       |
|--------------|---------------------------------------------------|
| Implements   | Displays the abstract template to be implemented. |
| Dependencies | Specifies the specific feature of a switch.       |
| Published    | Specifies if the template is published or not.    |
| Imports      | Specifies the base template for importing.        |

In addition, from the menu bar, choose **Control > Template Library > Templates** and you can also:

- Click **Show Filter** to filter the templates that is based on the headers.
- Click **Print** to print the list of templates.
- Click **Export to Excel** to export the list of template to a Microsoft Excel spreadsheet.

This section contains the following:

## Template Structure

The configuration template content mainly consists of four parts. Click the **Help** icon next to the **Template Content** for information about editing the content of the template.

This section contains the following:

## Template Format

This section describes the basic information of the template. The possible fields are as detailed in the table below.

| Property Name      | Description                                                                                            | Valid Values                                                                                                                                     | Optional? |
|--------------------|--------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| name               | The name of the template                                                                               | Text                                                                                                                                             | No        |
| description        | Brief description about the template                                                                   | Text                                                                                                                                             | Yes       |
| userDefined        | Indicates whether the user created the template. Value is 'true' if user created.                      | "true" or "false"                                                                                                                                | Yes       |
| supportedPlatforms | List of device platforms supports this configuration template. Specify 'All' to support all platforms. | N1K, N3K, N3500, N4K, N5K, N5500, N5600, N6K, N7K, N9K, MDS, VDC, N9K-9000v, IOS-XE, IOS-XR, Others, All Nexus Switches list separated by comma. | No        |

| Property Name | Description                          | Valid Values                                                                                                                                                                                                                                                                                               | Optional? |
|---------------|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| templateType  | Specifies the type of Template used. | <ul style="list-style-type: none"><li>• CLI</li><li>• POAP</li></ul> <p><b>Note</b> POAP option is not applicable for Cisco DCNM LAN Fabric deployment.</p> <ul style="list-style-type: none"><li>• POLICY</li><li>• SHOW</li><li>• PROFILE</li><li>• FABRIC</li><li>• ABSTRACT</li><li>• REPORT</li></ul> | Yes       |

| Property Name   | Description                                          | Valid Values | Optional? |
|-----------------|------------------------------------------------------|--------------|-----------|
| templateSubType | Specifies the sub type associated with the template. |              |           |

| Property Name | Description | Valid Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Optional? |
|---------------|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
|               |             | <ul style="list-style-type: none"> <li>• CLI               <ul style="list-style-type: none"> <li>• N/A</li> </ul> </li> <li>• POAP               <ul style="list-style-type: none"> <li>• N/A</li> <li>• VXLAN</li> <li>• FABRICPATH</li> <li>• VLAN</li> <li>• PMN</li> </ul> </li> <li><b>Note</b> POAP option is not applicable for Cisco DCNM LAN Fabric deployment.</li> <li>• POLICY               <ul style="list-style-type: none"> <li>• VLAN</li> <li>• INTERFACE_VLAN</li> <li>• INTERFACE_VPC</li> <li>• INTERFACE_ETHNET</li> <li>• INTERFACE_BD</li> <li>• INTERFACE_CHANNEL</li> <li>• INTERFACE_FC</li> <li>• INTERFACE_MGMT</li> <li>• INTERFACE_COBACK</li> <li>• INTERFACE_NVE</li> <li>• INTERFACE_VFC</li> <li>• INTERFACE_NFC</li> <li>• DEVICE</li> <li>• FEX</li> <li>• NIRA_FABRIC_LINK</li> <li>• NIER_FABRIC_LINK</li> </ul> </li> </ul> |           |

| Property Name | Description | Valid Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Optional? |
|---------------|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
|               |             | <ul style="list-style-type: none"> <li>• INTERFACE</li> <li>• SHOW               <ul style="list-style-type: none"> <li>• VLAN</li> <li>• INTERFACE_VLAN</li> <li>• INTERFACE_VPC</li> <li>• INTERFACE_ETHNET</li> <li>• INTERFACE_BD</li> <li>• <del>INTERFACE_CHANNEL</del></li> <li>• INTERFACE_FC</li> <li>• INTERFACE_MGMT</li> <li>• INTERFACE_COBACK</li> <li>• INTERFACE_NVE</li> <li>• INTERFACE_VFC</li> <li>• <del>INTERFACE_CHANNEL</del></li> </ul> </li> <li>• DEVICE               <ul style="list-style-type: none"> <li>• FEX</li> <li>• <del>NIRA_FABRIC_LINK</del></li> <li>• NIRA_FABRIC_LINK</li> <li>• INTERFACE</li> </ul> </li> <li>• PROFILE               <ul style="list-style-type: none"> <li>• VXLAN</li> </ul> </li> <li>• FABRIC               <ul style="list-style-type: none"> <li>• NA</li> </ul> </li> </ul> |           |

| Property Name | Description | Valid Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Optional? |
|---------------|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
|               |             | <ul style="list-style-type: none"> <li>• ABSTRACT               <ul style="list-style-type: none"> <li>• VLAN</li> <li>• INTERFACE_VLAN</li> <li>• INTERFACE_VPC</li> <li>• INTERFACE_ETHNET</li> <li>• INTERFACE_BD</li> <li>• <del>INTERFACE_CHANNEL</del></li> <li>• INTERFACE_FC</li> <li>• INTERFACE_MGMT</li> <li>• INTERFACE_LOOPBACK</li> <li>• INTERFACE_NVE</li> <li>• INTERFACE_VFC</li> <li>• <del>INTERFACE_CHANNEL</del></li> </ul> </li> <li>• DEVICE</li> <li>• FEX</li> <li>• NIRA_FABRIC_LINK</li> <li>• NIER_FABRIC_LINK</li> <li>• INTERFACE</li> </ul> <ul style="list-style-type: none"> <li>• REPORT               <ul style="list-style-type: none"> <li>• UPGRADE</li> <li>• GENERIC</li> </ul> </li> </ul> |           |

| Property Name | Description                                      | Valid Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Optional? |
|---------------|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| contentType   |                                                  | <ul style="list-style-type: none"> <li>• CLI               <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> </ul> </li> <li>• POAP               <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> </ul> </li> <li><b>Note</b> POAP option is not applicable for Cisco DCNM LAN Fabric deployment.</li> <li>• POLICY               <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> <li>• PYTHON</li> </ul> </li> <li>• SHOW               <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> </ul> </li> <li>• PROFILE               <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> <li>• PYTHON</li> </ul> </li> <li>• FABRIC               <ul style="list-style-type: none"> <li>• PYTHON</li> </ul> </li> <li>• ABSTRACT               <ul style="list-style-type: none"> <li>• TEMPLATE_CLI</li> <li>• PYTHON</li> </ul> </li> <li>• REPORT               <ul style="list-style-type: none"> <li>• PYTHON</li> </ul> </li> </ul> | Yes       |
| implements    | Used to implement the abstract template.         | Text                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Yes       |
| dependencies  | Used to select the specific feature of a switch. | Text                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Yes       |



| Property Name | Description                                                      | Valid Values      | Optional? |
|---------------|------------------------------------------------------------------|-------------------|-----------|
| published     | Used to Mark the template as read only and avoids changes to it. | “true” or “false” | Yes       |

## Template Variables

This section contains declared variables, the data type, default values, and valid values conditions for the parameters that are used in the template. These declared variables are used for value substitution in the template content section during the dynamic command generation process. Also these variables are used in decision making and in iteration blocks in the template content section. Variables have predefined data types. You can also add a description about the variable. The following table describes the syntax and usage for the available datatypes.

| Variable Type  | Valid Value                                                                                                | Iterative? |
|----------------|------------------------------------------------------------------------------------------------------------|------------|
| boolean        | true false                                                                                                 | No         |
| enum           | Example: running-config,<br>startup-config                                                                 | No         |
| float          | Floating number format                                                                                     | No         |
| floatRange     | Example: 10.1,50.01                                                                                        | Yes        |
| Integer        | Any number                                                                                                 | No         |
| integerRange   | Contiguous numbers separated by “_”<br><br>Discrete numbers separated by “,”<br><br>Example: 1-10,15,18,20 | Yes        |
| interface      | Format: <if type><slot>[/<sub slot>]/<port><br><br>Example: eth1/1, fa10/1/2 etc.                          | No         |
| interfaceRange | Example: eth10/1/20-25,<br>eth11/1-5                                                                       | Yes        |
| ipAddress      | IPv4 OR IPv6 address                                                                                       | No         |

| Variable Type          | Valid Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Iterative? |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| ipAddressList          | <p>You can have a list of IPv4, IPv6, or a combination of both types of addresses.</p> <p>Example 1: 172.22.31.97,<br/>172.22.31.99,<br/>172.22.31.105,<br/>172.22.31.109</p> <p>Example 2:<br/>2001:0cb8:85a3:0000:0000:8a2e:0370:7334,<br/><br/>2001:0cb8:85a3:0000:0000:8a2e:0370:7335,<br/><br/>2001:0cb8:85a3:1230:0000:8a2f:0370:7334</p> <p>Example 3: 172.22.31.97,<br/>172.22.31.99,<br/><br/>2001:0cb8:85a3:0000:0000:8a2e:0370:7334,<br/><br/>172.22.31.254</p> | Yes        |
| ipAddressWithoutPrefix | <p>Example: 192.168.1.1</p> <p>or</p> <p>Example: 1:2:3:4:5:6:7:8</p>                                                                                                                                                                                                                                                                                                                                                                                                      | No         |
| ipV4Address            | IPv4 address                                                                                                                                                                                                                                                                                                                                                                                                                                                               | No         |
| ipV4AddressWithSubnet  | Example: 192.168.1.1/24                                                                                                                                                                                                                                                                                                                                                                                                                                                    | No         |
| ipV6Address            | IPv6 address                                                                                                                                                                                                                                                                                                                                                                                                                                                               | No         |
| ipV6AddressWithPrefix  | <p>Example: 1:2:3:4:5:6:7:8</p> <p>22</p>                                                                                                                                                                                                                                                                                                                                                                                                                                  | No         |
| ipV6AddressWithSubnet  | IPv6 Address with Subnet                                                                                                                                                                                                                                                                                                                                                                                                                                                   | No         |
| ISISNetAddress         | <p>Example:</p> <p>49.0001.00a0.c96b.c490.00</p>                                                                                                                                                                                                                                                                                                                                                                                                                           | No         |
| long                   | Example: 100                                                                                                                                                                                                                                                                                                                                                                                                                                                               | No         |
| macAddress             | 14 or 17 character length MAC address format                                                                                                                                                                                                                                                                                                                                                                                                                               | No         |
| string                 | <p>Free text, for example, used for the description of a variable</p> <p>Example:</p> <pre>string scheduledTime { regularExpr=^([01]\d 2[0-3]):([0-5]\d)\$; }</pre>                                                                                                                                                                                                                                                                                                        | No         |

| Variable Type                                    | Valid Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Iterative?                                                                                              |
|--------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| string[]                                         | Example: {a,b,c,str1,str2}                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Yes                                                                                                     |
| struct                                           | <p>Set of parameters that are bundled under a single variable.</p> <pre> struct &lt;structure name declaration &gt; { &lt;parameter type&gt; &lt;parameter 1&gt;; &lt;parameter type&gt; &lt;parameter 2&gt;; ... } [&lt;structure_inst1&gt;] [, &lt;structure_inst2&gt;] [, &lt;structure_array_inst3 []&gt;;  struct interface_detail { string inf_name; string inf_description; ipAddress inf_host; enum duplex { validValues = auto, full, half; }; }myInterface, myInterfaceArray[]; </pre> | <p>No</p> <p><b>Note</b> If the struct variable is declared as an array, the variable is iterative.</p> |
| wwn<br>(Available only in Cisco DCNM Web Client) | <p>Example:<br/>20:01:00:08:02:11:05:03</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                      | No                                                                                                      |

## Variable Meta Property

Each variable that is defined in the template variable section has a set of meta properties. The meta properties are mainly the validation rules that are defined for the variable.

The following table describes the various meta properties applicable for the available variable types.

| Variable Type | Description                          | Variable Meta Property |              |                |     |     |          |          |          |          |            |            |              |
|---------------|--------------------------------------|------------------------|--------------|----------------|-----|-----|----------|----------|----------|----------|------------|------------|--------------|
|               |                                      | default Value          | valid Values | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| boolean       | A boolean value.<br>Example:<br>true | Yes                    |              |                |     |     |          |          |          |          |            |            |              |
| enum          |                                      |                        | Yes          |                |     |     |          |          |          |          |            |            |              |

| Variable Type  | Description                                                    | Variable Meta Property |              |                |     |     |          |          |          |          |            |            |              |
|----------------|----------------------------------------------------------------|------------------------|--------------|----------------|-----|-----|----------|----------|----------|----------|------------|------------|--------------|
|                |                                                                | default Value          | valid Values | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| float          | signed real number<br>Example:<br>75.56,<br>-8.5               | Yes                    | Yes          | Yes            | Yes | Yes |          |          |          |          |            |            |              |
| floatRange     | range of signed real numbers<br>Example:<br>50.5<br>-<br>54.75 | Yes                    | Yes          | Yes            | Yes | Yes |          |          |          |          |            |            |              |
| integer        | signed number<br>Example:<br>50,<br>-75                        | Yes                    | Yes          |                | Yes | Yes |          |          |          |          |            |            |              |
| integerRange   | Range of signed numbers<br>Example:<br>50-65                   | Yes                    | Yes          |                | Yes | Yes |          |          |          |          |            |            |              |
| interface      | specific interface<br>Example:<br>Ethernet<br>5/10             | Yes                    | Yes          |                |     |     | Yes      | Yes      | Yes      | Yes      |            |            |              |
| interfaceRange |                                                                | Yes                    | Yes          |                |     |     | Yes      | Yes      | Yes      | Yes      |            |            |              |
| ipAddress      | IP address in IPv4 or IPv6 format                              | Yes                    |              |                |     |     |          |          |          |          |            |            |              |

| Variable Type | Description                                                                                                                                                                                                                                                                                                                   | Variable Meta Property |                                                                  |                |     |     |          |          |          |          |            |            |              |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|------------------------------------------------------------------|----------------|-----|-----|----------|----------|----------|----------|------------|------------|--------------|
|               |                                                                                                                                                                                                                                                                                                                               | default Value          | valid Values                                                     | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| ipAddressList | <p>You can have a list of IPv4, IPv6, or a combination of both types of addresses.</p> <p>Example 1:<br/> 122.3.9,<br/> 122.3.9,<br/> 122.3.15,<br/> 122.3.10</p> <p>Example 2:<br/> 1001:607,<br/> 1001:607,<br/> 1001:607</p> <p>Example 3:<br/> 122.3.9,<br/> 122.3.9,<br/> 1001:607,<br/> 122.3.24</p> <p><b>Note</b></p> | Yes                    |                                                                  |                |     |     |          |          |          |          |            |            |              |
|               |                                                                                                                                                                                                                                                                                                                               |                        | Separate the addresses in the list using commas and not hyphens. |                |     |     |          |          |          |          |            |            |              |

| Variable Type         | Description                                    | Variable Meta Property |              |                |     |     |          |          |          |          |            |            |              |
|-----------------------|------------------------------------------------|------------------------|--------------|----------------|-----|-----|----------|----------|----------|----------|------------|------------|--------------|
|                       |                                                | default Value          | valid Values | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| <del>ipAddr</del>     | IPv4 or IPv6 Address (does not require prefix) |                        |              |                |     |     |          |          |          |          |            |            |              |
| <del>ip4Addr</del>    | IPv4 address                                   | Yes                    |              |                |     |     |          |          |          |          |            |            |              |
| <del>ip4Subnet</del>  | IPv4 Address with Subnet                       | Yes                    |              |                |     |     |          |          |          |          |            |            |              |
| <del>ip6Addr</del>    | IPv6 address                                   | Yes                    |              |                |     |     |          |          |          |          |            |            |              |
| <del>ip6Prefix</del>  | IPv6 Address with prefix                       | Yes                    |              |                |     |     |          |          |          |          |            |            |              |
| <del>ip6Subnet</del>  | IPv6 Address with Subnet                       | Yes                    |              |                |     |     |          |          |          |          |            |            |              |
| <del>ip6Example</del> | Example:<br><del>4008:5:50</del>               |                        |              |                |     |     |          |          |          |          |            |            |              |
| long                  | Example:<br>100                                | Yes                    |              |                | Yes | Yes |          |          |          |          |            |            |              |
| <del>macAddr</del>    | MAC address                                    |                        |              |                |     |     |          |          |          |          |            |            |              |

| Variable Type | Description                                                                                                                                     | Variable Meta Property |              |                |     |     |          |          |          |          |            |            |              |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|--------------|----------------|-----|-----|----------|----------|----------|----------|------------|------------|--------------|
|               |                                                                                                                                                 | default Value          | valid Values | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| string        | literal string<br><br>Example for string<br><br>Regular expression string<br><br><code>string {</code><br><code>  0000</code><br><code>}</code> | Yes                    |              |                |     |     |          |          |          |          | Yes        | Yes        | Yes          |
| string[]      | string literals that are separated by a comma (,)<br><br>Example:<br>{string1, string2}                                                         | Yes                    |              |                |     |     |          |          |          |          |            |            |              |

| Variable Type | Description                                                                                                                                                                                                                                                                             | Variable Meta Property |              |                |     |     |          |          |          |          |            |            |              |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|--------------|----------------|-----|-----|----------|----------|----------|----------|------------|------------|--------------|
|               |                                                                                                                                                                                                                                                                                         | default Value          | valid Values | decimal Length | min | max | min Slot | max Slot | min Port | max Port | min Length | max Length | regular Expr |
| struct        | Set of <del>params</del> that are bundled under a single variable.<br><br>struct<br><br><structure name declaration><br>> {<br><parameter type><br><br><parameter 1>;<br><parameter type><br><br><parameter 2>;<br>...<br>}<br><struct1><br>[,<br><struct2><br>[,<br><struct3><br>[ ]>; |                        |              |                |     |     |          |          |          |          |            |            |              |
| wnn           | WWN address                                                                                                                                                                                                                                                                             |                        |              |                |     |     |          |          |          |          |            |            |              |

### Example: Meta Property Usage

```
##template variables

integer VLAN_ID {
min = 100;
max= 200;
};

string USER_NAME {
defaultValue = admin123;
minLength = 5;
};

struct interface_a{
```



```

string inf_name;
string inf_description;
ipAddress inf_host;
enum duplex {
 validValues = auto, full, half;
};
}myInterface;

##

```

## Variable Annotation

You can configure the variable properties marking the variables using annotations.



**Note** Variable Annotations are available for POAP only. However, the annotations do not impact on the template type 'CLI'.

The following annotations can be used in the template variable section.

| Annotation Key          | Valid Values                                                         | Description                                       |
|-------------------------|----------------------------------------------------------------------|---------------------------------------------------|
| AutoPopulate            | Text                                                                 | Copies values from one field to another           |
| DataDepend              | Text                                                                 |                                                   |
| Description             | Text                                                                 | Description of the field appearing in the window  |
| DisplayName             | Text<br><b>Note</b> Enclose the text with quotes, if there is space. | Display name of the field appearing in the window |
| Enum                    | Text1, Text2, Text3, and so on                                       | Lists the text or numeric values to select from   |
| IsAlphaNumeric          | "true" or "false"                                                    | Validates if the string is alphanumeric           |
| IsAsn                   | "true" or "false"                                                    |                                                   |
| IsDestinationDevice     | "true" or "false"                                                    |                                                   |
| IsDestinationFabric     | "true" or "false"                                                    |                                                   |
| IsDestinationInterface  | "true" or "false"                                                    |                                                   |
| IsDestinationSwitchName | "true" or "false"                                                    |                                                   |
| IsDeviceID              | "true" or "false"                                                    |                                                   |
| IsDot1qId               | "true" or "false"                                                    |                                                   |

| Annotation Key          | Valid Values                                                                                       | Description                                                                                                                               |
|-------------------------|----------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| IsFEXID                 | “true” or “false”                                                                                  |                                                                                                                                           |
| IsGateway               | “true” or “false”                                                                                  | Validates if the IP address is a gateway                                                                                                  |
| IsInternal              | “true” or “false”                                                                                  | Makes the fields internal and does not display them on the window<br><br><b>Note</b> Use this annotation only for the ipAddress variable. |
| IsManagementIP          | “true” or “false”<br><br><b>Note</b> This annotation must be marked only for variable “ipAddress”. |                                                                                                                                           |
| IsMandatory             | “true” or “false”                                                                                  | Validates if a value should be passed to the field mandatorily                                                                            |
| IsMTU                   | “true” or “false”                                                                                  |                                                                                                                                           |
| IsMultiCastGroupAddress | “true” or “false”                                                                                  |                                                                                                                                           |
| IsMultiLineString       | “true” or “false”                                                                                  | Converts a string field to multiline string text area                                                                                     |
| IsMultiplicity          | “true” or “false”                                                                                  |                                                                                                                                           |
| IsPassword              | “true” or “false”                                                                                  |                                                                                                                                           |
| IsPositive              | “true” or “false”                                                                                  | Checks if the value is positive                                                                                                           |
| IsReplicationMode       | “true” or “false”                                                                                  |                                                                                                                                           |
| IsShow                  | “true” or “false”                                                                                  | Displays or hides a field on the window                                                                                                   |
| IsSiteId                | “true” or “false”                                                                                  |                                                                                                                                           |
| IsSourceDevice          | “true” or “false”                                                                                  |                                                                                                                                           |
| IsSourceFabric          | “true” or “false”                                                                                  |                                                                                                                                           |
| IsSourceInterface       | “true” or “false”                                                                                  |                                                                                                                                           |

| Annotation Key           | Valid Values      | Description                                          |
|--------------------------|-------------------|------------------------------------------------------|
| IsSourceSwitchName       | “true” or “false” |                                                      |
| IsSwitchName             | “true” or “false” |                                                      |
| IsRMID                   | “true” or “false” |                                                      |
| IsVPCDomainID            | “true” or “false” |                                                      |
| IsVPCID                  | “true” or “false” |                                                      |
| IsVPCPeerLinkPort        | “true” or “false” |                                                      |
| IsVPCPeerLinkPortChannel | “true” or “false” |                                                      |
| IsVPCPortChannel         | “true” or “false” |                                                      |
| Password                 | Text              | Validates the password field                         |
| PeerOneFEXID             | “true” or “false” |                                                      |
| PeerTwoFEXID             | “true” or “false” |                                                      |
| PeerOnePCID              | “true” or “false” |                                                      |
| PeerTwoPCID              | “true” or “false” |                                                      |
| PrimaryAssociation       |                   |                                                      |
| ReadOnly                 | “true” or “false” | Makes the field read-only                            |
| ReadOnlyOnEdit           | “true” or “false” |                                                      |
| SecondaryAssociation     | Text              |                                                      |
| Section                  |                   |                                                      |
| UsePool                  | “true” or “false” |                                                      |
| UseDNSReverseLookup      |                   |                                                      |
| Username                 | Text              | Displays the username field on the window            |
| Warning                  | Text              | Provides text to override the Description annotation |

#### Example: AutoPopulate Annotation

```
##template variables
string BGP_AS;
@ (AutoPopulate="BGP_AS")
```

```
 string SITE_ID;
##
```

### Example: DisplayName Annotation

```
##template variables
@(DisplayName="Host Name", Description = "Description of the host")
String hostname;
@(DisplayName="Host Address", Description = " test description" IsManagementIP=true)
IPAddress hostAddress;
##
```

### Example: IsMandatory Annotation

```
##template variables
@(IsMandatory="ipv6!=null")
IPv4Address ipv4;
@(IsMandatory="ipv4!=null")
IPv6Address ipv6;
##
```

### Example: IsMultiLineString Annotation

```
##template variables
@(IsMultiLineString=true)
string EXTRA_CONF_SPINE;
##
```

### IsShow Annotation

```
##template variables
boolean isVlan;
@(IsShow="isVlan==true")
integer vlanNo;
##

##template variables
boolean enableScheduledBackup;
@(IsShow="enableScheduledBackup==true",Description="Server time")
string scheduledTime;
##
The condition "enableScheduledBackup==true" evaluates to true/false

##template variables
@(Enum="Manual,Back2BackOnly,ToExternalOnly,Both")
string VRF_LITE_AUTOCONFIG;
@(IsShow="VRF_LITE_AUTOCONFIG!=Manual", Description="Target Mask")
integer DCI_SUBNET_TARGET_MASK
##
The condition "VRF_LITE_AUTOCONFIG!=Manual" matches string comparison to evaluate to true
or false
```

### Example: Warning Annotation

```
##template variables
@(Warning="This is a warning msg")
 string SITE_ID;
##
```

## Templates Content

This section includes the configuration commands and any parameters that you want to include in the template. These commands can include the variables declared in the template variables section. During the command generation process the variable values are substituted appropriately in the template content.



**Note** You must specify the commands that you include as if you were entering them in the global configuration command mode on any device. You must consider the command mode when you include commands.

Template content is governed by the usage of variables.

- **Scalar variables:** does not take a range or array of values which cannot be used for iteration (In the variable types table those marked iterate-able as 'No'). Scalar variables must be defined inside the template content.

```
Syntax: $$<variable name>$$
Example: $$USER_NAME$$
```

- **Iterative variables:** used for block iteration. These loop variable must be accessed as shown below inside the iteration block.

```
Syntax:@<loop variable>
Example:
foreach val in $$INTEGER_RANGE_VALUE$$ {
@val
}
```

- **Scalar Structure Variable:** Structure member variables can be accessed inside the template content.

```
Syntax: $$<structure instance name>.<member variable name>$$
Example: $$myInterface.inf_name$$
```

- **Array Structure Variable:** Structure member variables can be accessed inside the template content.

```
Syntax: $$<structure instance name>.<member variable name>$$
Example: $$myInterface.inf_name$$
```

In addition to the template variables, you can use the conditional and iterative command generation using the following statements:

- **if-else if-else Statement:** makes a logical decision in inclusion/exclusion of set of configuration command based on the value assigned for the variable in it.

```
Syntax: if(<operand 1> <logical operator> <operand 2>){
command1 ..
command2..
..
}
else if (<operand 3> <logical operator> <operand 4>)
{
Command3 ..
Command4..
..
}
else
{
```

```

Command5 ..
Command6..
..
}
Example: if-else if-else statement
if($$USER_NAME$$ == 'admin'){
Interface2/10
no shut
}
else {
Interface2/10
shut
}

```

- **foreach Statement:** used for iterating a block of commands. The iteration is performed based on the assigned loop variable value.

```

Syntax:
foreach <loop index variable> in $$<loop variable>$$ {
@<loop index variable> ..
}
Example: foreach Statement
foreach ports in $$MY_INF_RANGES$${
interface @ports
no shut
}

```

- **Optional parameters:** By default all parameters are mandatory. To make a parameter optional, you must annotate the parameter.

In the variable section, you can include the following command:

- **@(IsMandatory=false)**

- **Integer frequency;**

In the template content section, a command can be excluded or included without using “if” condition check, by assigning a value to the parameter. The optional command can be framed as below:

- **probe icmp [frequency frequency-value] [timeout seconds] [retry-count retry-count-value]**

## Template Content Editor

The template content editor has the following features:

- **Syntax highlighting:** The editor highlights the syntax, like different types of statements, keywords, and so on, for Python scripting.
- **Autocompletion:** The editor suggests the template datatypes, annotations, or metaproperties when you start typing.
- **Go to line:** You can navigate to the exact line in the template content editor instead of scrolling. Press **Command-L** in Mac or **Ctrl-L** in Windows, and enter the line number to which you want to navigate to in the pop-up window.

If you enter a value greater than the number of lines in the editor, you will be navigated to the last line in the editor window.

- **Template search and replace:** Press **Command-F** in Mac or **Ctrl-F** in Windows, enter the search term in the **Search for** field, and select the type of search in the search window. You can perform the following searches in the editor:
  - **RegExp Search:** You can perform the regular expression search in the editor.
  - **CaseSensitive Search:** You can perform a case-sensitive search in the editor.
  - **Whole Word Search:** You can perform a whole word search to find the exact words in the editor. For example, a regular search for the word "play" returns results where it is part of words like "display," but the whole word search returns results only when there is an exact match for the word "play".
  - **Search In Selection:** You can perform a search in the selected content. Select the content to which you want to limit the search and enter the search term.

Choose the + icon in the search window to use the replace option. Enter the replacing word in the **Replace with** field. You can replace the selected word once by selecting **Replace**. To replace all the occurrences of the selected word, select **All**.

- **Code folding:** You can expand or group code blocks in the editor by clicking the arrow next to their line numbers.
- **Other features:** The editor automatically indents the code, the closing braces, and highlights the matching parenthesis.

## Template Editor Settings

You can edit the following features of a template editor by clicking **Template Editor Settings**.

- **Theme:** Select the required theme for the editor from the drop-down list.
- **KeyBinding:** Select the editor mode from the **KeyBinding** drop-down list to customize the editor. **Vim** and **Ace** modes are supported. The default is **Ace**.
- **Font Size:** Select the required font size for the editor.

## Advanced Features

The following are the advanced features available to configure templates.

- **Assignment Operation**

Config template supports assignment of variable values inside the template content section. The values are validated for the declared data type of the variable. If there is a mismatch, the value is not assigned.

Assignment operation can be used under the following guidelines:

- The operator on the left must be any of the template parameters or a for loop parameter.
- The operator on the right values can be any of the values from template parameters, for loop parameters, literal string values surrounded by quotes or simple string values.

If a statement does not follow these guidelines, or if it does not suit this format, it will not be considered as assignment operation. It is substituted during command generation like other normal lines.

Example: Template with assignment operation

```

##template properties
name =vlan creation;
userDefined= true;
supportedPlatforms = All;
templateType = CLI;
published = false;
##
##template variables
integerRange vlan_range;
@(internal=true)
integer vlanName;
##
##template content
foreach vlanID in $$vlan_range$${
vlan @vlanID
$$vlanName$$=@vlanID
name myvlan$$vlanName$$
}
##

```

- Evaluate methods

Config template uses the Java runtime provided Java script environment to perform arithmetic operations (such as ADD, SUBTRACT, and so on), string manipulations, and so on.

Locate the JavaScript file in the template repository path. This file contains primary set of arithmetic, string functions. You can also add custom JavaScript methods.

These methods can be called from config template content section in below format:

Example1:

```

$$somevar$$ = evalscript(add, "100", $$anothervar$$)

```

Also the *evalscript* can be called inside if conditions as below:

```

if($$range$$ > evalscript(sum, $$vlan_id$$, -10)){
do something...
}

```

You can call a method that is located at the backend of the Java script file.

- Dynamic decision

Config template provides a special internal variable “LAST\_CMD\_RESPONSE”. This variable stores the last command response from the device during the execution of the command. This can be used in the config template content to make dynamic decisions to deliver the commands that are based on the device condition.




---

**Note** The if block must be followed by an else block in a new line, which can be empty.

---

An example use case to create a VLAN, if it does not exist on the device.

```

Example: Create VLAN
##template content
show vlan id $$vlan_id$$
if($$LAST_CMD_RESPONSE$$ contains "not found"){
vlan $$vlan_id$$
}
else{

```



```
}
##
```

This special implicit variable can be used only in the “IF” blocks.

- **Template referencing**

You can have a base template with all the variables defined. This base template can be imported to multiple templates. The base template content is substituted in the appropriate place of the extending template. The imported template parameters and the contents can be accessed inside the extending template.

```
Example: Template Referencing
Base template:
##template properties
 name =a vlan base;
 userDefined= true;
 supportedPlatforms = All;
 templateType = CLI;
 published = false;
 timestamp = 2015-07-14 16:07:52;
 imports = ;
##
##template variables
 integer vlan_id;
##
##template content
 vlan $$vlan_id$$
##

Derived Template:
##template properties
 name =a vlan extended;
 userDefined= true;
 supportedPlatforms = All;
 templateType = CLI;
 published = false;
 timestamp = 2015-07-14 16:07:52;
 imports = a vlan base,template2;
##
##template variables
 interface vlanInterface;
##
##template content
 <substitute a vlan base>
 interface $$vlanInterface$$
 <substitute a vlan base>
##
```

When you launch the extended template, the parameter inputs for the base template are also obtained. In addition, the substituted content is used for complete CLI command generation.

## Report Template

Starting from Cisco DCNM 11.3(1) Release, a new template type, REPORT, has been added. This template has two subtypes, UPGRADE and GENERIC. The template type is python.

### UPGRADE

The UPGRADE template is used for pre-ISSU and post-ISSU scenarios. These templates are listed in the ISSU wizard.

Refer to the default upgrade template packaged in DCNM for more information on pre-ISSU and post-ISSU handling. The default upgrade template is `issu_vpc_check`.

## GENERIC

The GENERIC template is used for any generic reporting scenarios, such as, collecting information about resources, switch inventory, SFPs, and NVE VNI counters. You can also use this template to generate troubleshooting reports.

## Resources Report

This report displays information about resource usage for a specific fabric.

The **Summary** section shows all resource pools with the current usage percentages. Use the horizontal scroll bar at the bottom of the window to display more columns.

| POOL NAME             | POOL RANGE      | SUBNET MASK | MAX ENTRIES | USAGE INSIDE RANGE | USAGE OUTSIDE RANGE | USAGE PERCENTAGE |
|-----------------------|-----------------|-------------|-------------|--------------------|---------------------|------------------|
| SUBNET                | 10.4.0.0/16     | 30          | 16384       | 4                  | 0                   | 0.02             |
| LOOPBACK_IP_POOL      | 10.2.0.0/22     | -           | 1024        | 4                  | 0                   | 0.39             |
| LOOPBACK1_IP_POOL     | 10.3.0.0/22     | -           | 1024        | 4                  | 0                   | 0.39             |
| ANYCAST_RP_IP_POOL    | 10.254.254.0/24 | -           | 256         | 1                  | 0                   | 0.39             |
| DCI subnet pool       | 10.33.0.0/16    | 30          | 16384       | 0                  | 0                   | 0                |
| TOP_DOWN_NETWORK...   | 2300-2999       | -           | 700         | 0                  | 5                   | 0                |
| TOP_DOWN_VRF_VLAN     | 2000-2299       | -           | 300         | 5                  | 0                   | 1.67             |
| TOP_DOWN_L3_DOT1Q     | 2-511           | -           | 510         | 0                  | 0                   | 0                |
| SERVICE_NETWORK_VL... | 3000-3199       | -           | 200         | 0                  | 0                   | 0                |
| VPC_DOMAIN_ID         | 1-1000          | -           | 1000        | 1                  | 0                   | 0.1              |
| LOOPBACK_ID           | 0-1023          | -           | 1024        | 3                  | 0                   | 0.29             |

**POOL NAME:** Specifies the name of the pool.

**POOL RANGE:** Specifies the IP address range of the pool.

**SUBNET MASK:** Specifies the subnet mask.

**MAX ENTRIES:** Specifies the maximum number of entries that can be allocated from the pool.

**USAGE INSIDE RANGE:** Specifies the current number of entries allocated inside the pool range.

**USAGE OUTSIDE RANGE:** Specifies the current number of entries set outside the pool range.

**USAGE PERCENTAGE:** This is calculated by using the formula:  $(\text{Usage Inside Range} / \text{Max Entries}) * 100$ .

Click **View Details** to display a view of resources allocated or set in each resource pool. For example, the detailed section for a SUBNET has information about the resources that have been allocated within the subnet.

Resources for Pool SUBNET: Type SUBNET\_POOL: Range 10.4.0.0/16

SUBNET Allocated Resources

| SCOPE TYPE | SCOPE       | DEVICE NAME | ALLOCATED RESOURCE | ALLOCATED TO                    | ID |
|------------|-------------|-------------|--------------------|---------------------------------|----|
| Link       | SAL1834YY80 | n9k-5       | 10.4.0.0/30        | SAL1834YY80-Vlan3600-SAL18...   | 61 |
| Link       | SAL1834YY80 | n9k-5       | 10.4.0.4/30        | SAL1834YY80-Ethernet1/28-SAL... | 80 |
| Link       | SAL1919EMST | n9k-28      | 10.4.0.8/30        | SAL1919EMST-Ethernet1/17-SA...  | 83 |
| Link       | SAL1919EMST | n9k-28      | 10.4.0.12/30       | SAL1919EMST-Ethernet1/4-SAL...  | 86 |

## Switch Inventory Report

This report provides a summary about the switch inventory.

Summary Total 6

DCNM-UUID-1510 View Details

- Device Name : N9K\_41
- Chassis ID : FDO222425SE
- Model : Nexus9000 93180YC-EX chassis
- NXOS version : 9.3(2)
- UpTime : 1 day(s), 10 hour(s), 42 minute(s), 7 second(s)

Click **View Details** to display more information about the modules and licenses.

Modules

| TYPE                                      | SLOT | HARDWARE REVISION | MODEL NAME      | MODULE SERIAL NUMBER |
|-------------------------------------------|------|-------------------|-----------------|----------------------|
| Nexus9000 93180YC-EX chassis              |      | V03               | N9K-C93180YC-EX | FDO222425SE          |
| 48x10/25G + 6x40/100G Ethernet Module     | 1    | V03               | N9K-C93180YC-EX | FDO222425SE          |
| Nexus9000 93180YC-EX chassis Power Supply |      | V02               | NXA-PAC-650W-PE | ART2219F83V          |
| Nexus9000 93180YC-EX chassis Power Supply |      | V02               | NXA-PAC-650W-PE | ART2219F84J          |
| Nexus9000 93180YC-EX chassis Fan Module   |      | V01               | NXA-FAN-30CFM-F | N/A                  |
| Nexus9000 93180YC-EX chassis Fan Module   |      | V01               | NXA-FAN-30CFM-F | N/A                  |
| Nexus9000 93180YC-EX chassis Fan Module   |      | V01               | NXA-FAN-30CFM-F | N/A                  |
| Nexus9000 93180YC-EX chassis Fan Module   |      | V01               | NXA-FAN-30CFM-F | N/A                  |

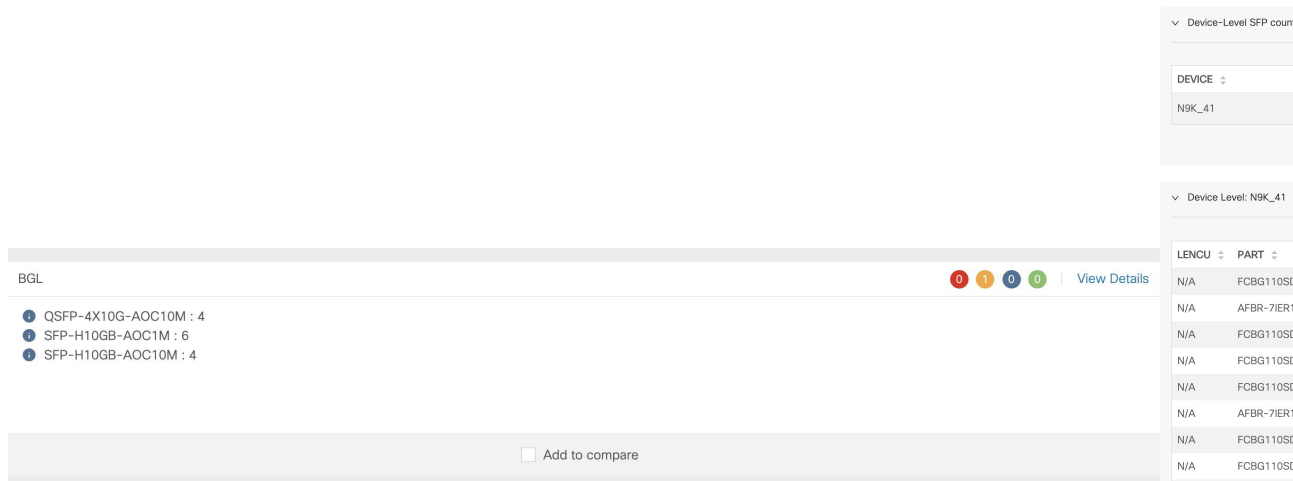
Licenses

FEATURE

- N9K\_LIC\_1G
- VPN\_FABRIC
- NXOS\_OF\_PKG
- FCOE\_NPV\_PKG
- SECURITY\_PKG
- ACI-PREMIER-GF
- N9K\_UPG\_EX\_10G
- TP\_SERVICES\_PKG
- NXOS\_ADVANTAG
- NXOS\_ADVANTAG
- NXOS\_ADVANTAG
- NXOS\_ESSENTIALS
- NXOS\_ESSENTIALS

## SFP Report

This report provides information about utilization of SFPs at a fabric and device level.



BGL

- QSFP-4X10G-AOC10M : 4
- SFP-H10GB-AOC1M : 6
- SFP-H10GB-AOC10M : 4

Add to compare

| LENCU | PART       |
|-------|------------|
| N/A   | FCBG110SD  |
| N/A   | AFBR-7IER1 |
| N/A   | FCBG110SD  |
| N/A   | FCBG110SD  |
| N/A   | FCBG110SD  |
| N/A   | AFBR-7IER1 |
| N/A   | FCBG110SD  |
| N/A   | FCBG110SD  |



**Note** The switch inventory and SFP reports are supported only on Cisco Nexus devices.

### Troubleshooting Reports

These reports are generated to help in troubleshooting scenarios. Currently, the **NVE VNI Counters** report is the only pre-defined troubleshooting report. Generating **NVE VNI Counters** reports involves performing periodic checks to identify the VNIs that are among the top hits based on network traffic. In a large-scale setup, we recommend limiting the report generation frequency to a minimum of 60 minutes.

#### NVE VNI Counters Report

This report collects the **show nve vni counters** command output for each VNI in the fabric.

After comparing the oldest report and the newest report, the **Summary** section shows the top-10 hit VNIs. The top hit VNIs are displayed in these categories:

- L2 or L3 VNIs for unicast traffic
- L2 or L3 VNIs for multicast traffic
- L2 only VNIs for unicast traffic
- L2 only VNIs for multicast traffic
- L3 only VNIs for unicast traffic
- L3 only VNIs for multicast traffic

The oldest report refers to the first report that is saved in the current reporting task. If you want to select a specific report as the first report against which the current report has to be compared, delete all reports that are older than the one selected so that the selected report becomes the first and oldest report.

For example, three reports were run yesterday at 8:00 a.m., 4:00 p.m. and 11:00 p.m. If you want to use the report at 11:00 p.m. as the first and oldest report for today's reporting, delete the two reports that were run yesterday at 8:00 a.m. and 4:00 p.m.

For a periodic report, the oldest report is the first report that is run at the start time of a period. For daily and weekly reports, the current report is compared against the previously generated report.

The **Summary** section displays a column-wise report with information about the total transmitted bytes and the VNIs. Use the horizontal scroll bar at the bottom of the window to display more columns.

| Summary                                                                                                            |                |                                  |                |
|--------------------------------------------------------------------------------------------------------------------|----------------|----------------------------------|----------------|
| v4-fabric                                                                                                          |                |                                  |                |
| This Summary shows the Top Hit VNIs between this report and the oldest report created on 2020-05-25 17:53:42 -0700 |                |                                  |                |
| Top 10 L2 or L3 VNIs (Unicast)                                                                                     |                | Top 10 L2 or L3 VNIs (Multicast) |                |
| VNI                                                                                                                | TOTAL TX BYTES | VNI                              | TOTAL TX BYTES |
| 30004                                                                                                              | 655458         | 30000                            | 43418          |
| 30002                                                                                                              | 217122         | 30002                            | 43310          |
| 30000                                                                                                              | 64             | 30004                            | 43310          |
| 30001                                                                                                              | 0              | 30001                            | 42912          |
| 30003                                                                                                              | 0              | 30003                            | 42912          |
| 50000                                                                                                              | 0              | 50000                            | 42912          |
| 50002                                                                                                              | 0              | 50003                            | 42912          |
| 50001                                                                                                              | 0              | 50002                            | 42840          |
| 50004                                                                                                              | 0              | 50001                            | 42840          |
| 50003                                                                                                              | 0              | 50004                            | 42840          |



**Note** The **Summary** section in the NVE VNI Counters report displays negative numbers in the TOTAL TX BYTES column if a report is generated after a switch reload or after clearing the counters on the switch. The numbers are displayed correctly in the subsequent reports. As a workaround, we recommend deleting all old reports or creating a new job before reloading switches or clearing counters.

Click **View Details** to display more information. This section shows NVE VNIs and counters on a per-switch basis.

| NVE VNI Counters for SAL18432P6M:n9k-17 |       |              |               |              |               |              |               |              |               |
|-----------------------------------------|-------|--------------|---------------|--------------|---------------|--------------|---------------|--------------|---------------|
| Total VNIs                              |       |              |               |              |               |              |               |              |               |
| NVE VNI Counters                        |       |              |               |              |               |              |               |              |               |
| ROW NUMBER                              | VNI   | TX_UCASTPKTS | TX_UCASTBYTES | TX_MCASTPKTS | TX_MCASTBYTES | RX_UCASTPKTS | RX_UCASTBYTES | RX_MCASTPKTS | RX_MCASTBYTES |
| 1                                       | 30000 | 15           | 1676          | 21           | 2888          | 6            | 836           | 3            | 342           |
| 2                                       | 30001 | 0            | 0             | 0            | 0             | 0            | 0             | 0            | 0             |
| 3                                       | 30002 | 100          | 108618        | 1            | 110           | 99           | 108504        | 1            | 114           |
| 4                                       | 30003 | 0            | 0             | 0            | 0             | 0            | 0             | 0            | 0             |
| 5                                       | 30004 | 300          | 327818        | 1            | 110           | 299          | 327704        | 1            | 114           |
| 6                                       | 50000 | 0            | 0             | 0            | 0             | 0            | 0             | 0            | 0             |
| 7                                       | 50001 | 0            | 0             | 0            | 0             | 0            | 0             | 0            | 0             |
| 8                                       | 50002 | 0            | 0             | 0            | 0             | 0            | 0             | 0            | 0             |
| 9                                       | 50003 | 0            | 0             | 0            | 0             | 0            | 0             | 0            | 0             |
| 10                                      | 50004 | 0            | 0             | 0            | 0             | 0            | 0             | 0            | 0             |

For more information on how the reports are displayed, refer [Programmable Report](#).

## Report Template Functions

### generateReport method

The generateReport method is invoked while generating a report and contains the report implementation logic. This method accepts any context object. and returns a WrappersResp object. For more information on WrappersResp refer link.

### Validation method

The validation method is optional. If the template defines this method, the Programmable Report application calls this method to perform pre-validation checks while creating the job. This method is called only when the job is created and invoked only once irrespective of the number of devices or fabrics selected. If the validation passes, this method returns a WrappersResp object with a SuccessRetCode. If the validation fails, this method returns a FailureRetCode along with an error list.

Examples of a successful validation and a failed validation are as follows:

#### \*Successful validation

```
def validate(context):
 respObj = WrappersResp.getRespObj()

 ## Validation logic here

 respObj.setSuccessRetCode()
 return respObj
```

#### \*Failed validation

```
def validate(context):
 respObj = WrappersResp.getRespObj()

 ## Validation logic here

 respObj.setFailureRetCode()
 respObj.addErrorReport(template_name, error)
 return respObj
```

We can also perform validation based on the content of the context parameter.

### Context parameter

The Context parameter consists of the following attributes:

- User name - Name of the user who created the job
- User role - Role of the user who created the job
- Job ID
- Recurrence - NOW, ONCE, DAILY, WEEKLY, MONTHLY, ONDEMAND, PERIODIC
- Period - If the recurrence is periodic, then the period will display the selected frequency.

For more information on job context APIs, refer the *Job Context Information* section.

## Report Python Library

A REPORT has the following components:

- Summary
  - Key and Values
  - Messages- Inferences
- Details/Sections
  - Key and Values
  - JSON document – Cards
  - Array of JSON Documents – Tables
- Command Log

A python library is provided to generate the **report** JSON model. To use these APIs, the following import statement has to be added to the template:

```
from reportlib.preport import Report
```

## Report APIs

### Create Report

To create a 'Report' object, use this API –

```
report = Report ("Report title")
```

### Add Summary

Each report can have a summary. This is a python dictionary. To add a summary, use this API –

```
summary = report.add_summary ()
```

### Adding Content to the Summary

To add content to the summary, use the following APIs-

Key and Values -

```
summary ['NXOS Version'] = '8.4(1)'
```

Messages and Inferences -

```
summary.add_message ("Simple message")
```




---

**Note** In Cisco DCNM Release 11.4(1), adding a JSON object as a value in summary is not supported. The following example is not supported-

```
summary["info"] = {"key":"value","key-2":"value-2"}
```

---

### Adding tables in Summary

To add table to the summary, use this API –

```
table = summary.add_table(title, _id)
```

*title*: Title of the table.

*\_id*: Unique identifier for the table.

### Adding rows to the table

To add rows to the table, use this API –

```
table.append(value, _id)
```

*value*: is a JSON object. Nested JSON is not supported.

*\_id*: is the unique identifier for the row.

### Adding a Section

A section is a logical grouping of report content. Sections are created and configured by you to display the required information. To add a section, use this API –

```
section = report.add_section ("Section title",_id)
```

*\_id*: Unique identifier for the section.

### Adding Content to a Section Key and Values

To add a simple key and value pair to a section, use this API –

```
section['key'] = 'value'
```

### JSON Document – Cards

A JSON document can be added in the same way as a simple key and value pair is added.





---

**Note** Nested JSON is not supported in Cisco DCNM Release 11.4(1).

---

An example of a JSON document displayed as a card widget is as shown below:



### Array of JSON Documents – Tables

To create a table and add rows to this table, use this API –

```
section.append(key, dictionary, _id)
```

*\_id*: Unique identifier that identifies a row in a table. A duplicate *\_id* results in a unique id violation error.

Example-

```
section.append('Switch Details', {'name': 'N5K'}, 'DSDAS244455')
```

The creation of tables using this API has the following limitations:



- All JSON documents should have the same set of keys/columns. Any difference in the number of columns or column names may result in the table not being rendered on the Web UI.
- Nested JSON is not supported.

### Formatters

A Formatter enables additional formatting for values that are displayed on the user interface. For example, you can mark values as ERROR, SUCCESS, WARNING, and INFO. These values are color-coded and displayed on the Web UI. Errors are displayed in red, Warnings in yellow, Info in blue and Success in green.



To configure formatting, use this API-

```
Formatter.add_marker(value,marker)
```

*value*: Value to add a marker

*marker*: Marker.ERROR, Marker.SUCCESS, Marker.WARNING, and Marker.INFO

Example-

```
Formatter.add_marker ("NXOS version",Marker.INFO)
```

### Charts

You can add charts to both the summary and section.

To add a chart to Summary, use this API-

```
report = Report("title")
summary = report.add_summary()
summary.add_chart(ChartType, _id)
```

*ChartType*: ChartTypes.COLUMN\_CHART, ChartTypes.PIE\_CHART, and ChartTypes.LINE\_CHART

*\_id*: Unique ID for the chart

To add a chart to a section, use this API-

```
report = Report("title")
section = report.add_section ("section_title",_id)
section.add_chart(ChartType, _id)
```

*ChartType*: ChartTypes.COLUMN\_CHART, ChartTypes.PIE\_CHART, and ChartTypes.LINE\_CHART

*\_id*: Unique ID for the chart




---

**Note** Ensure that the classes are imported in the import section.

---

### Pie chart

To display information in a pie chart, use this API-

To set title and subtitle:

```
pie_chart.set_title("Chart title")
pie_chart.set_subtitle("Sub title")
```

To add value:

```
pie_chart.add_value("key", value)
```

*key*: String key

*value*: Numeric value

### Column chart

To display information in a column chart, use this API-

To set title and subtitle title:

```
column_chart.set_title("Chart title")
column_chart.set_subtitle("Sub title")
```

To set X-Axis and Y-Axis title

```
column_chart.set_xAxis_title("X-Axis title")
column_chart.set_yAxis_title("y-Axis title")
```

To add value:

```
bar_chart.add_value("key", value, category)
```

*key*: String key

*value*: Numeric value

*category*: The column chart divides the data into a logical group that is called a category . A given key should have a value in each category. For example, Device count is a key and Fabric Names are categories. A chart should have a Device count for each fabric.

### Line Chart

To display information in a line chart, use this API-

To set title and subtitle title:

```
line_chart.set_title("Chart title")
line_chart.set_subtitle("Sub title")
```

To set X-Axis and Y-Axis title

```
line_chart.set_xAxis_title("X-Axis title")
line_chart.set_yAxis_title("y-Axis title")
```

To add value:

```
line_chart.add_value("key", value, category)
```

*key*: String key

*value*: Numeric value

*category*: The line chart divides the data into logical group called category. A given key should have a value in each category. For example, 'Device count' is a key and 'Fabric Names' are categories. A chart should have a Device count for each fabric or category.

### Running CLIs on a Device

To configure running of CLIs on a device, use this API-

```
from reportlib.preport import show
cli_responses = show(serial_number, *commands)
```

*serial\_number*: Serial number of the device on which the commands have to be run. In case of a VDC instance, the serial number is **serial\_number:vdc\_name**.

*\*commands*: Commands that are run on the device. These are variable arguments.

Examples-

- Executing a command on single switch:

```
cli_responses = show("FOX1816G0S9", 'show version | xml', 'show inventory | xml', 'show license usage | xml')
```

- Executing a command on multiple switches:

```
cli_responses = show(["FOX1816G0S9","SSI15470HJ5"], 'show version | xml', 'show inventory | xml', 'show license usage | xml')
```

### Show commands and store responses

To configure the show commands and store responses, use this API-

```
from reportlib.preport import show_and_store
cli_responses = show_and_store(report, serial_number, *commands)
```

*report*: Report object created earlier.

*serial\_number*: serial number of the device to run commands. In case of VDC, serial number should be *serial\_number:vdc\_name*. You can add a list of serial numbers to run the same set of commands on multiple devices.

*commands*: Commands to run on the device. These are variable arguments. You can specify multiple commands.

Examples-

- Executing a command on single switch:

```
cli_responses = show_and_store(report, "FOX1816G0S9", 'show version | xml', 'show inventory | xml', 'show license usage | xml')
```

- Executing a command on multiple switches:

```
cli_responses = show_and_store(report, ["FOX1816G0S9","SSI15470HJ5"], 'show version | xml', 'show inventory | xml', 'show license usage | xml')
```




---

**Note** This API stores the response from the device in elasticsearch along with the report. We recommend being cautious while using this API, as storing all responses may reduce available storage space.

---

### Return value

The API mentioned above returns a list of responses. Each response is a dictionary with the following structure-

```
{
 'status': 'success' | 'failed',
 'response': <response from device>,
 'command': <cli command>,
 'serial_number': <device serial number>
}
```

In case of multiple switches, the response is a list of responses with separate entries for each switch.

Example-

```
[
 {
 'status': 'success',
 'response': <response from device>,
 'command': 'show version',
 'serial_number': 'FOX1816G0S9'
 },
 {
 'status': 'success',
 'response': <response from device>,
 'command': 'show version',
 'serial_number': 'SSI15470HJ5'
 }
]
```

### Job context information

To display the recurrence while scheduling the job from the application, use this API-

```
get_recurrence(context)
```

Return values can be NOW,ONCE,DAILY,WEEKLY,MONTHLY,ONDEMAND, and PERIODIC.

If a job is scheduled as Periodic and information about a specific period has to be retrieved, use this API-

```
period = get_period(context)
```

*period.get\_period()*: Returns the period.

*period.get\_time\_unit()*: Returns the time unit (HOURS, MINUTES).

### Analysis with Historical Reports

#### Retrieve previously generated reports

Use the *get\_previous\_reports()* method to get reports that have been generated in the past. This can be used to perform analysis based on current data and historical data. This API will return a list of reports in descending order of the time at which the reports were created.

```
List of reports = get_previous_reports(context,entity,count)
```

*context*: The object received as input from the generateReport (context) method

*entity*: serial\_number or fabric name

*count*: Number of reports to fetch

#### Get oldest report

To retrieve the oldest report, use this API-

```
oldest_report = get_oldest_report(context,entity)
```

*context*: The object received as input from the generateReport(context) method

*entity*: serial\_number or fabric name

Both the APIs listed above return a Report object with the following APIs to retrieve information-

- Get summary : *report.get\_summary()*
- Get section : *report.get\_section(\_id)* where *\_id* is the unique identifier for the section as mentioned in *Adding a Section*.

### XML Utilities

The XML utilities are based on `xml.etree.elementtree` (<https://docs.python.org/2/library/xml.etree.elementtree.html>).

### **getxmlltree**

To return the XML tree under the specified tag, use this API-

```
from reportlib.preport import getxmlltree
xml_element_tree = getxmlltree(xml_string, tag)
```

*xml\_string*: XML response from device.

*tag*: XML tag. The complete XML under this tag will be returned as the `ElementTree`.

*xml\_element\_tree*: The API that returns the `xml.etree.ElementTree` object

### **getxmlrows**

To get an array of rows if the CLI response contains rows, use this API-

```
from reportlib.preport import getxmlrows
rows = getxmlrows(xml_tree, tag_xpath)
```

*xml\_tree*: `xml.etree.ElementTree` object.

*tag\_xpath*: xpath of the XML record. Please refer <https://docs.python.org/2/library/xml.etree.elementtree.html#xpath-support>.

*rows*: An array of rows

### **getnodevalue**

To read the XML node value, use this API-

```
from reportlib.preport import getnodevalue
value = getnodevalue(xml_tree, node_xpath)
```

*xml\_tree*: The `xml.etree.ElementTree` object

*node\_xpath*: xpath of the XML record. Please refer <https://docs.python.org/2/library/xml.etree.elementtree.html#xpath-support>.

### **Check for existence of node**

This API returns True or False depending on whether the given tag is present or not in the XML tree.

```
from reportlib.preport import
has_tag has_tag(xml_tree, tag)
```

*xml\_tree*: The `xml.etree.ElementTree` object

### **WrappersResp**

Each report has to return an object of the type **WrappersResp**. This can be initiated by using the API given below. Import this from `com.cisco.dcbu.vinci.rest.services.jython` `import WrappersResp`.

```
respObj = WrappersResp.getRespObj()
```

The return code in `WrapperResp` indicates whether the report ran successfully or not.

- If all commands are run and the required information is extracted, then the report returns a success API - `respObj.setSuccessRetCode()`
- In case of any exception such as a command failure, then the report returns a failure API - `respObj.setFailureRetCode()`

Setting a failure code indicates that there is an issue with report execution and the report is not generated.

To return a report with errors, use `Formatter` to mark the error and set the `WrapperResp` to `Success`. Refer *Formatters* for more information.

For any errors that may come up, you can use this API to specify the reason for the error-

```
respObj.addErrorReport(template_name,error_message)
```

The `report` object created by you should be set to the value of `WrappersResp` as shown below:

```
respObj.setValue(report)
```

### Logger

Logger enables logging of messages from the report template. Information that is logged using the logger is logged to this location- “/usr/local/cisco/dcm/fm/logs/preport\_jython.log”.

```
Logger.info("message")
Logger.debug("message")
Logger.error("message")
Logger.trace("message")
Logger.warn("message")
```

## Adding a Template

To add user-defined templates and schedule jobs from the Cisco DCNM Web UI, perform the following steps:

### Procedure

**Step 1** Choose **Control > Template Library**.

The **Templates** window is displayed with the name of the template along with its description, supported platforms, and tags.

**Step 2** Click **Add** to add a new template.

The Template Properties window appears.

**Step 3** Specify a template name, description, tags, and supported platforms for the new template.

**Step 4** Specify a **Template Type** for the template.

**Step 5** Select a **Template Sub Type** and **Template Content Type** for the template.

**Step 6** Click the **Advanced** tab to edit other properties like **Implements**, **Dependencies**, **Published**, and **Imports**. Select **Published** to make the template read-only. You cannot edit a published template.

**Step 7** From the **Imports > Template Name** list, check the template check box.

The base template content is displayed in the **Template Content** window. The base template displays the template properties, template variables, and template content. This template can be imported in to another template and the base template content is substituted in the appropriate place of the extending template. When you launch the extended template, the parameter inputs for the base template are also obtained. Also, the substituted content is used for complete CLI command generation.

**Note** The base templates are CLI templates.

**Step 8** Click **OK** to save the template properties, or click the cancel icon at the top-right corner of the window to revert the changes.

**Note** You can edit the template properties by clicking **Template Property**.

**Step 9** Click **Template Content** to edit the template syntax. For information about the structure of the Configuration Template, see the *Template Structure* section.

**Step 10** Click **Validate Template Syntax** to validate the template values.

If an error or a warning message appears, you can check the validation details in **Validation Table** by clicking the error and warnings field.

**Note** You can continue to save the template if there are warnings only. However, if there is an error, you must edit the templates to fix the errors before you proceed. Click the line number under the Start Line column to locate the error in the template content. You will get an error if you validate a template that does not have a template name.

**Step 11** Click **Save** to save the template.

**Step 12** Click **Save and Exit** to save the configuration and go back to the configuring templates screen.

---

## Modifying a Template

You can edit the user-defined templates. However, the predefined templates and templates that are already published cannot be edited.

### Procedure

---

**Step 1** From **Control > Template Library**, select a template.

**Step 2** Click **Modify/View template**.

**Step 3** Edit the template description and tags.

The edited template content is displayed in a pane on the right.

**Step 4** From the **Imports > Template Name** list, check the template check box.

The base template content is displayed in the **Template Content** window. You can edit the template content based on your requirement in the **Template Content** window. Click the help icon next to the **Template Content** window for information about editing the content of the template.

**Step 5** Edit the supported platforms for the template.

**Step 6** Click **Validate Template Syntax** to validate the template values.

**Step 7** Click **Save** to save the template.

**Step 8** Click **Save and Exit** to save the configuration and go back to the configuring templates screen.

---

## Copying a Template

To copy a template from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Template Library**, and select a template.
  - Step 2** Click **Save Template As**.
  - Step 3** Edit the template name, description, tags, and other parameters.  
The edited template content is displayed in the right-hand pane.
  - Step 4** From the **Imports > Template Name** list, check the template check box.  
The base template content is displayed in the **Template Content** window. You can edit the template content that is based on your requirement in the **Template Content** window. Click the help icon next to the **Template Content** window for information about editing the content of the template.
  - Step 5** Edit the supported platforms for the template.
  - Step 6** Click **Validate Template Syntax** to validate the template values.
  - Step 7** Click **Save** to save the template.
  - Step 8** Click **Save and Exit** to save the configuration and go back to the configuring templates screen.
- 

## Deleting a Template

You can delete the user-defined templates. However, you cannot delete the predefined templates. From Cisco DCNM Release 11.0(1), you can delete multiple templates at once.

To delete a template from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Template Library**.
  - Step 2** Use the check box to select a template and click **Remove template** icon.  
The template is deleted without any warning message.
- 

### What to do next

The template is deleted from the list of templates on the DCNM Web UI. When you restart the DCNM services, the deleted templates are displayed on the **Control > Template Library** page.

To delete the template permanently, delete the template that is located in your local directory: `Cisco Systems\dcm\dcnm\data\templates\`.

## Importing a Template

To import a template from the Cisco DCNM Web UI, perform the following steps:



### Procedure

---

**Step 1** Choose **Control > Template Library** and click **Import Template**.

**Step 2** Browse and select the template that is saved on your computer.

You can edit the template parameters, if necessary. For information, see [Modifying a Template, on page 367](#).

**Note** The “\n” in the template is considered as a new line character when imported and edited, but it works fine when imported as a ZIP file.

**Step 3** Click **Validate Template Syntax** to validate the template.

**Step 4** Click **Save** to save the template or **Save and Exit** to save the template and exit.

---

## Exporting a Template

To export a template from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

**Step 1** Choose **Control > Template Library**.

**Step 2** Use the check box to select a template and click **Export Template**.

The browser requests you to open or save the template to your directory.

---

## Image Management

Upgrading your devices to the latest software version manually might take a long time and prone to error, which requires a separate maintenance window. To ensure rapid and reliable software upgrades, image management automates the steps associated with upgrade planning, scheduling, downloading, and monitoring. Image management is supported only for Cisco Nexus switches.



**Note** Before you upgrade, ensure that the POAP boot mode is disabled for Cisco Nexus 9000 Series switches and Cisco Nexus 3000 Series switches. To disable POAP, run the `no boot poap enable` command on the switch console. You can however, enable it after the upgrade.

---

The **Image Management** menu includes the following submenu and options:

Table 4: Image Management Menu

| Submenu                   | Options                                                       | Actions                             |                       |
|---------------------------|---------------------------------------------------------------|-------------------------------------|-----------------------|
| Image Upload              | Smart Image Management                                        | Image Upload                        |                       |
|                           |                                                               | Deleting an Image                   |                       |
| Install & Upgrade         | Upgrade History<br>Window Name: <b>Software Upgrade Tasks</b> | View                                |                       |
|                           |                                                               | Delete                              |                       |
|                           |                                                               | New Installation                    | New ISSU Installation |
|                           |                                                               |                                     | EPLD Installation     |
|                           |                                                               | Finish Installation                 |                       |
|                           | Switch Level History                                          | View Device Upgrade Tasks           |                       |
|                           | Refresh Switch Level History Table                            |                                     |                       |
| Package [SMU/RPM]         | Packages                                                      | Installing Packages and Patches     |                       |
|                           |                                                               | Uninstalling Packages and Patches   |                       |
|                           |                                                               | Activating Packages and Patches     |                       |
|                           |                                                               | Deactivate                          |                       |
| Image Management Policies | Image Management Policies                                     | Adding an Image Management Policy   |                       |
|                           |                                                               | Deleting an Image Management policy |                       |

Ensure that your user role is **network-admin** or **device-upg-admin** and you didn't freeze the DCNM to perform the following operations:

- Upload or delete images.
- Install, delete, or finish installation of an image.
- Install or uninstall packages and patches.
- Activate or deactivate packages and patches.
- Add or delete image management policies (applicable only for network-admin user role).
- View management policies.

You can view any of the image installations or device upgrade tasks if your user role is **network-admin**, **network-stager**, **network-operator**, or **device-upg-admin**. You can also view them if your DCNM is in freeze mode.

## Smart Image Management

This feature allows you to upload or delete images that are used during POAP and switch upgrade. You can also upload or delete RPMs and SMUs that are used for installing in the **Packages** window. To view the **Smart Image Management** window from the Cisco DCNM Web UI homepage, choose **Control > Image Management > Image Upload**.

You can view the following details in the **Smart Image Management** window.

| Field         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Platform      | <p>Specifies the name of the platform. Images, RPMs, or SMUs are categorized as follows:</p> <ul style="list-style-type: none"> <li>• N9K/N3k</li> <li>• N6K</li> <li>• N7K</li> <li>• N77K</li> <li>• N5K</li> <li>• Other</li> <li>• Third Party</li> </ul> <p>The images are the same for N9K and N3K platforms.</p> <p>The platform will be <b>Other</b> if the uploaded images are not mapped to any of the existing platforms.</p> <p>The platform will be <b>Third Party</b> for RPMs.</p> |
| Image Name    | Specifies the filename of the image, RPM, or SMU that you uploaded.                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Image Type    | Specifies the file type of the image, EPLD, RPM, or SMU.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Image Subtype | <p>Specifies the file type of the image, EPLD, RPM, or SMU.</p> <p>The file type EPLDs are <b>epld</b>. The file types of images are <b>nxos</b>, <b>system</b> or <b>kickstart</b>. The file type for RPMs is <b>feature</b> and for SMUs the file type is <b>patch</b>.</p>                                                                                                                                                                                                                     |
| NXOS Version  | Specifies the NXOS image version for only Cisco switches.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Image Version | Specifies the image version for all devices, including the non-Cisco devices as well.                                                                                                                                                                                                                                                                                                                                                                                                             |
| Size (Bytes)  | Specifies the size of the image, RPM, or SMU files in bytes.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Checksum      | Specifies the checksum of the image. The checksum checks if there's any corruption in the file of the image, RPM, or SMU. You can validate the authenticity by verifying if the checksum value is same for the file you downloaded from the Cisco website and the file you upload in the <b>Image Upload</b> window.                                                                                                                                                                              |

You can sort all columns.

## Image Upload

To upload different types of images to the server from the Cisco DCNM Web UI, perform the following steps:



**Note** Devices use these images during POAP or image upgrade. RPMs and SMUs are used in the **Packages** window. All the images, RPMs, and SMUs are used in the **Image Management Policies** window.

Your user role should be **network-admin**, or **device-upg-admin** to upload an image. You can't perform this operation with the **network-stager** user role.

### Procedure

**Step 1** Choose **Control > Image Management > Image Upload**.

The **Smart Image Management** window appears.

**Step 2** Click **Image Upload**.

The **Select File to Upload** dialog box appears.

**Step 3** Click **Choose file** to choose a file from the local repository of your device.

**Step 4** Choose the file and click **Upload**.

You can upload a ZIP file as well. Cisco DCNM processes and validate the image file and categorize it under the existing platforms accordingly. If it doesn't fall under **N9K/N3K**, **N6K**, **N7K**, **N77K**, or **N5K** platforms, the image file is categorized under **Third Party** or **Other** platform. The **Third Party** platform is applicable only for RPMs.

**Step 5** Click **OK**.

The EPLD images, RPMs, and SMUs are uploaded to the repository in the following path:  
**/var/lib/dcnm/upload/<platform\_name>**

All NX-OS, kickstart and system images are uploaded to the repository in the following paths:  
**/var/lib/dcnm/images** and **/var/lib/dcnm/upload/<platform\_name>**

The upload takes some time depending on the file size and network bandwidth.

**Note** You can upload images for all Cisco Nexus Series Switches.

You can upload EPLD images only for Cisco Nexus 9000 Series Switches.

If your network speed is slow, increase the wait time of Cisco DCNM to 1 hour so that the image upload is complete. To increase the wait time from Cisco DCNM Web UI, perform the following steps:

- a. Choose **Administrator > DCNM Server > Server Properties**.
- b. Search for the **csrf.refresh.time** property, and set the value as **60**.

**Note** The value is in minutes.

- c. Click **Apply Changes**.

- d. Restart the DCNM server.

## Deleting an Image

To delete an image from the repository from the Cisco DCNM Web UI, perform the following steps:

### Procedure

- Step 1** Choose **Control > Image Management > Image Upload**.  
The **Smart Image Management** window appears.
- Step 2** Choose an existing image from the list and click the **Delete Image** icon.  
A confirmation window appears.
- Step 3** Click **Yes** to delete the image.

## Install & Upgrade

The **Install & Upgrade** menu includes the following submenus:

### Upgrade History

This feature enables you to upgrade the Cisco Nexus Platform Switches using In-Service Software Upgrade (ISSU). This upgrade procedure may be disruptive or non-disruptive based on the device configuration. You can select the Kickstart, System, or NX-OS images from image repository or the file system on the device. To select the images from the repository, the same needs to be uploaded from **Control > Image Management > Image upload** tab.

The following table describes the fields that appear on **Control > Image Management > Upgrade History**.

| Field     | Description                                                                                                                                                                                     |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Task Id   | Specifies the serial number of the task. The latest task will be listed in the top.<br><br><b>Note</b> If Failover is triggered in Native HA, the Task Id sequence number is incremented by 32. |
| Task Type | Specifies the type of task. <ul style="list-style-type: none"> <li>• Compatibility</li> <li>• Upgrade</li> </ul>                                                                                |
| Owner     | Based on the Role-Based Authentication Control (RBAC), specifies the owner who initiated this task.                                                                                             |
| Devices   | Displays all the devices that were selected for this task.                                                                                                                                      |

| Field          | Description                                                                                                                                                                                                                                                                                                                     |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Job Status     | <p>Specifies the status of the job.</p> <ul style="list-style-type: none"> <li>• Planned</li> <li>• In Progress</li> <li>• Completed</li> <li>• Completed with Exceptions</li> </ul> <p><b>Note</b> If the job fails on a single or multiple devices, the status field shows COMPLETED WITH EXCEPTION indicating a failure.</p> |
| Created Time   | Specifies the time when the task was created.                                                                                                                                                                                                                                                                                   |
| Scheduled At   | Specifies the time when the task is specified to be executed. You can also choose to schedule a task to be executed at a later time.                                                                                                                                                                                            |
| Completed Time | Specifies the time when the task was completed.                                                                                                                                                                                                                                                                                 |
| Comment        | Shows any comments that the Owner has added while performing the task.                                                                                                                                                                                                                                                          |




---

**Note** After a fresh Cisco DCNM installation, this page will have no entries.

---

You can perform the following:

## View

To view the image upgrade history from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

**Step 1** Choose **Control > Image Management > Install & Upgrade > Upgrade History**, check the task ID check box.

Select only one task at a time.

**Step 2** Click **View**.

The **Installation Task Details** window appears.

**Step 3** Click **Settings**. Expand the **Columns** menu and choose the details you want to view.

You can view the following information in this window:

- Location of the kickstart and system images

- Compatibility check status
- Installation status
- Pre-ISSU report status and post-ISSU report status
- Descriptions
- Report summary
- Version check results
- Logs

The columns change according to the task you choose to view. You can see the switch name, IP address, platform details, image name, and installation status for an EPLD task. The report status includes the report summary as well. The report summary includes hyperlinks to detailed pre-ISSU reports and post-ISSU reports. Clicking these hyperlinks takes you to a new tab or window to view the reports. The report summary will also include the commands that you defined in the report templates.

**Step 4** Select the device.

The detailed status of the task appears. For the completed tasks, the response from the device appears.

If the upgrade task is in progress, a live log of the installation process appears.

- Note**
- This table autorefreshes every 30 secs for jobs in progress, when you're on this window.
  - It takes some time for the upgraded EPLD information to appear. A job is scheduled to fetch updates from the switch to DCNM every five minutes until the switch is reachable.

## Delete

To delete a task from the Cisco DCM Web UI, perform the following steps:

### Procedure

- Step 1** Choose **Control > Image Management > Install & Upgrade > Upgrade History**, and check the **Task ID** check box.
- Step 2** Click **Delete**.
- Step 3** Click **OK** to confirm deletion of the job.

## New Installation

You can install ISSU and EPLD images in Cisco DCM.

### *New ISSU Installation*

To upgrade the devices that are discovered from the Cisco DCM, perform the following steps:

### Before you begin

Add report templates in the **Template Library** window if you want pre-ISSU and post-ISSU reports. Refer to the default upgrade template packaged in DCNM for more information on pre-ISSU and post-ISSU handling. The default upgrade template is **issu\_vpc\_check**.

### Procedure

- 
- Step 1** Choose **Control > Image Management > Install & Upgrade > Upgrade History**.
- Step 2** Choose **New Installation > ISSU** to install, or upgrade the kickstart and the system images on the devices. The devices with default VDCs are displayed in the **Select Switches** window.
- Note** Switches that are part of fabrics in the freeze mode or monitor mode aren't listed here. An error message appears if you choose a fabric, which is in freeze mode or monitor mode, from the **Device Scope** drop-down menu.
- Step 3** Select the check box to the left of the switch name. You can select more than one switch.
- Step 4** Click **Next**. The **Pre-Post ISSU Reports** window appears.
- Note** Pre-Post ISSU Reports are not supported in SAN and Media Controller installations.
- Step 5** (Optional) Check the **Skip Pre-Post ISSU Reports** check box to skip the pre-post ISSU reports on switches and go to *Step 8*. By default, this check box isn't checked.
- Step 6** Choose a report template from the **Select Report Template** drop-down list. Only the templates of **REPORT** template type with **UPGRADE** sub-type that are listed in the **Control > Template Library** window appear in the **Select Report Template** drop-down list.
- Step 7** Fill in the required fields in the **General** tab based on the template you chose in *Step 6*.
- Step 8** Click **Next**. The **Specify Software Images** window appears. This tab displays the switches that you selected in the previous screen. You can choose the images for upgrade as well.
- The **Auto File Selection** check box enables you to specify an image version, and a path where you can apply the upgraded image to the selected devices.
  - **Select File Server** is disabled, and the default server is used.
  - In the **Image Version** field, specify the image version as displayed in the **Image Upload** window.
  - The **Path** field is disabled, and the default image path is used.
- Step 9** Click **Select Image** in the **Kickstart image** column. The **Software Image Browser** dialog box appears.



- Note**
- Cisco Nexus 9000 Series Switches require only the system image to load the Cisco NX-OS operating system. Therefore, the option to select kickstart images for these devices is disabled.
  - If there's an issue in viewing the **Software Image Browser** dialog box, reduce the font size of your browser and retry.

**Step 10** Click **Select Image** in the **System Image** column.  
The **Software Image Browser** dialog box appears.

**Step 11** On the **Software Image Browser** dialog box, you can choose the image from **File Server** or **Switch File System**.

If you choose **File Server**:

- From the **Select the File server** list, choose the Default\_SCP\_Repository file server on which the image is stored.
- From the **Select Image** list, choose the appropriate image. Check the check box to use the same image for all other selected devices of the same platform.

Example: For platform types N9K-C93180YC-EX and N9K-C93108TC-EX, logic matches platform (N9K) and three characters (C93) from subplatform. The same logic is used across all platform switches.

**Note** Only files with BIN extension are listed if you select **File Server**. To view other files, choose **Administration > DCNM Server > Server Properties**, set **FILE\_SELECTION\_FILTER** to **false**, and restart the server. It is set to **true** by default.

**Note** Only image files present in the **Image Upload** window can be selected. You can't select images present in any other paths.

- Choose a VRF from the **Select Vrf** drop-down list.

**Note** This field does not appear for Cisco MDS switches.

This VRF is selected for other selected devices by default. The default value is **management**.

- Click **OK**.

This image is selected for all other selected devices of same platform type.

If you choose **Switch File System**:

- From the **Select Image** list, choose the appropriate image that is located on the flash memory of the device.

**Note** Only files with BIN extension are listed if you select **Switch File System**. To view other files, choose **Administration > DCNM Server > Server Properties**, set **FILE\_SELECTION\_FILTER** to **false**, and restart the server. It is set to **true** by default.

- Click **OK** to choose the kickstart image or **Cancel** to revert to the **Specify Software Images** dialog box.

**Step 12** The **Vrf** column indicates the name of the virtual routing and forwarding (VRF).

**Step 13** In the **Available Space** column, specify the available space for the **Primary Supervisor** and **Secondary Supervisor** modules of the switch.

**Available Space** column shows the available memory in MB on the switch (for less than 1 MB, it's shown and marked as KB).

Bootflash browser shows the filename, size, and last modified date for all the files and directories on the switch bootflash. You can delete files by selecting them and clicking **Delete** to increase the available space on the switch.

**Step 14** **Selected Files Size** column shows the size of images that are selected from the server.

If the total size of selected images is greater than available space on a switch, the file size is marked in red. We recommend that you create more space on the switch to copy images to it and install.

**Step 15** Drag and drop the switches to reorder the upgrade task sequence.

**Step 16** (Optional) Uncheck **Skip Version Compatibility** check box if you want to check the compatibility of Cisco NX-OS software version on your device with the upgraded images that you chose.

**Step 17** Select **Select Parallel Line Card upgrade** to upgrade all the line cards at the same time.

Upgrading a parallel line card isn't applicable for Cisco MDS devices.

**Step 18** Click **Options** under the **Upgrade Options** column to choose the type of upgrade.

**Upgrade Options** window appears with two upgrade options. The drop-down list for **Upgrade Option 1** has the following options:

- **Disruptive**
- **Bios force**
- **Allow non-disruptive**
- **Force non-disruptive**

**Disruptive** is the default value for Cisco Nexus 9000 Series switches. The upgrade option is **Not Applicable** for other switches.

When you choose **Allow non-disruptive** under **Upgrade Option 1** and if the switch does not support non-disruptive upgrade, then it will go through a disruptive upgrade.

When you choose **Force non-disruptive** under **Upgrade Option 1**, the **Skip Version Compatibility** check box will be unchecked because compatibility check is mandatory for non-disruptive upgrade. If the switches you choose do not support non-disruptive upgrade, a warning message appears asking you to review the switch selection. Use the check boxes to choose or remove switches.

The drop-down list for **Upgrade Option 2** has the following options when you choose **Allow non-disruptive** or **Force non-disruptive** under **Upgrade Option 1**:

- **NA**
- **bios-force**

When you choose **Disruptive** or **Bios-force** under **Upgrade Option 1**, **Upgrade Option 2** is disabled.

Check the **Use this Option for all other selected devices** check box to use the selected option for all the selected devices and click **OK**.

- Note**
- The upgrade options are applicable only for Cisco Nexus 3000 Series and 9000 Series switches.
  - Selecting the **Allow non-disruptive** option for upgrading does not ensure a non-disruptive upgrade. Perform a compatibility check to ensure that the device supports non-disruptive upgrade.

**Step 19** Click **Next**.

If you didn't select **Skip Version Compatibility**, the Cisco DCNM performs a compatibility check.

You can choose to wait until the check is complete or click **Finish Installation Later**.

The installation wizard is closed and a compatibility task is created in **Control > Image Management > Install & Upgrade > Upgrade History** tasks.

The time that is taken to check the image compatibility depends on the configuration and the load on the device.

The **Compatibility Verification** status column displays the status of verification.

If you skip the version compatibility check by choosing **Skip Version Compatibility**, Cisco DCNM displays only the name of the device. The **Current Action** column displays **Completed**, and the **Compatibility Verification** column displays **Skipped**.

The **Pre-ISSU Report Status** column specifies if the pre-ISSU reports were generated. You can view the compatibility log and the report summary in the **Compatibility Status** column. Click the hyperlink in the report summary to see a detailed report of the pre-ISSU check.

**Note** The status might take some time to reflect in the Web UI depending in the internet bandwidth.

You can review the switch selection and check or uncheck the switches for upgrading accordingly.

**Step 20** Click **Finish Installation Later** to perform the upgrade later.**Step 21** Click **Next**.**Step 22** Check the check box to save the running configuration to the startup configuration before upgrading the device.**Step 23** You can schedule the upgrade process to occur immediately or later.

- a. Select **Deploy Now** to upgrade the device immediately.
- b. Select **Choose time to Deploy** and specify the time in MMM/DD/YYYY HH:MM:SS format to perform the upgrade later.

This value is relative to the server time. If the selected time to deploy is in the past, the job is executed immediately.

**Step 24** You can choose the execution mode based on the devices and the line cards you have chosen to upgrade.

- a. Select **Sequential** to upgrade the devices in the order you chose them.

**Note** This option is disabled if you put the device in maintenance mode.

- b. Select **Concurrent** to upgrade all the devices at the same time.

**Step 25** Click **Finish** to begin the upgrade process.

The Installation wizard closes and a task to upgrade is created on the **Control > Image Management > Install & Upgrade > Upgrade History** page.

---

### What to do next

After you complete the ISSU on the switch, ensure that you wait for 20 minutes to allow the switch to reboot, and stabilize the SNMP agent. DCMN discovers polling cycles in order to display the new version of the switch on the Cisco DCMN Web UI.

### EPLD Installation

Cisco DCMN supports two types of EPLD image installations or upgrade on Cisco Nexus 9000 Series Switches:

- Upgrade all modules from an EPLD image.
- Upgrade only specific modules from an EPLD image.

To select the images from the repository, upload them from **Control > Image Management > Image Upload**.

To install or upgrade EPLD images in Cisco DCMN, perform the following steps:

### Procedure

**Step 1** Choose **Control > Image Management > Install & Upgrade > Upgrade History**.

**Step 2** Choose **New Installation > EPLD**.

The Cisco Nexus 9000 Series Switches are displayed in the **Select Switches** window.

**Note** Switches that are part of fabrics in the freeze mode or monitor mode are not listed here. An error message appears if you choose a fabric, which is in freeze mode or monitor mode, from the **Device Scope** drop-down menu.

**Step 3** Check the check box to the left of the switch name.

You can choose more than one device.

**Step 4** Click **Next**.

The **Specify EPLD Images** window appears. This tab displays the switches, that you selected in the previous screen, and allows you to choose the EPLD images for upgrade.

**Step 5** Click **Select Image** in the **EPLD image** column.

The **EPLD Image Browser** dialog box appears.

**Step 6** Choose the EPLD image file from the file server or switch file system.

If you choose **File Server**:

a) From the **Select Image** list, choose the appropriate image.

- Note**
- Only files with IMG extension are listed if you select **File Server**. To view other files, choose **Administration > DCMN Server > Server Properties**, set **FILE\_SELECTION\_FILTER** to **false**, and restart the server. It is set to **true** by default.
  - Only image files present in the **Image Upload** window can be selected. You cannot select images present in any other paths.

b) Click **OK** to choose the EPLD image or **Cancel** to revert to the **Specify Software Images** window.

If you choose **Switch File System**:

- a) From the **Select Image** list, choose the appropriate image that is located on the flash memory of the device.

**Note** Only files with IMG extension are listed if you select **Switch File System**. To view other files, choose **Administration > DCNM Server > Server Properties**, set **FILE\_SELECTION\_FILTER** to **false**, and restart the server. It is set to **true** by default.

- b) Click **OK** to choose the EPLD image or **Cancel** to revert to the **Specify Software Images** dialog box.

- Step 7** Choose the VRF from the Select Vrf drop-down list.  
The valid values are management, default, and keepalive.
- Step 8** (Optional) Check the **Use this Vrf for all other selected devices** check box to use the VRF for all other devices you chose.
- Step 9** (Optional) Check the **Use this image for all other selected devices of same platform type** check box to use this image for all other devices you chose.
- Step 10** The **Vrf** column indicates the name of the virtual routing and forwarding (VRF).
- Step 11** Specify the available space for the **Primary Supervisor** and **Secondary Supervisor** modules of the switch in the **Available Space** column.  
  
The **Available Space** column shows the available memory in MB on the switch (for less than 1 MB, it is shown and marked as KB).  
  
Bootflash browser shows the filename, size, and last modified date for all the files and directories on the switch bootflash. You can delete files by selecting them and clicking **Delete** to increase the available space on the switch.
- Step 12** Check if the total size of selected images is greater than available space on a switch in the **Selected Files Size** column.  
  
The **Selected Files Size** column shows the size of images that are selected from the server.  
  
If the total size of selected images is greater than available space on a switch, the file size is marked in red. We recommend that you create more space on the switch to copy images to it and install.  
  
**Note** The EPLD upgrade fails if the version that is supposed to be returned is not returned.
- Step 13** Drag and drop the switches to reorder the upgrade task sequence.
- Step 14** Click the hyperlink in the **Module Options** column to choose the module for corresponding switch to upgrade EPLD modules.  
  
The **Module Options** dialog box appears. The default value is **All**, which installs or upgrade all EPLD modules for the switch you chose.
- Step 15** Choose the modules.
- Step 16** Click **OK**.
- Step 17** Choose the FPGA region by clicking the hyperlinks under the **FPGA Region** column.  
  
The valid options are **Primary** and **Golden**.  
  
If you choose the golden upgrade, ensure the BIOS is updated and all the prerequisites are met. See the *Cisco Nexus 9000 Series FPGA/EPLD Upgrade Release Notes* for more information.
- Step 18** Click **Finish** to begin the upgrade process.

The Installation wizard closes and a task to upgrade is created on the **Control > Image Management > Install & Upgrade > Upgrade History** window. You can identify the EPLD upgrade tasks by the task type.

---

### What to do next

After you complete the upgrade on the switch, ensure that you wait for 20 minutes to allow the switch to reboot, and stabilize the SNMP agent. Cisco DCNM discovers polling cycles in order to display the new version of the switch in the **Switch Level History** window in Cisco DCNM Web UI.

You can view the EPLD golden upgrade notifications in the **Events** window. From the homepage of the Cisco DCNM Web UI, choose **Monitor > Switch > Events**.

## Finish Installation

You can choose to complete the installation for tasks which was completed on the **Compatibility Check** page. Perform the following task to complete the upgrade process on the devices.

### Procedure

---

- Step 1** Choose **Control > Image Management > Install & Upgrade > Upgrade History**, select a task for which the compatibility check is complete.
- Select only one task at a time.
- Step 2** Click **Finish Installation**.
- Software Installation Wizard** appears.
- Step 3** Review the switch selection and check or uncheck the switches for upgrading accordingly.
- Step 4** Click **Next**.
- Step 5** Check the check box to save the running configuration to the startup configuration before upgrading the device.
- Step 6** Check the check box to put a device in maintenance mode before upgrade. This option is valid only for the devices that support maintenance mode.
- Step 7** You can schedule the upgrade process to occur immediately or later.
- Select **Deploy Now** to upgrade the device immediately.
  - Select **Choose time to Deploy** and specify the time in DD/MM/YYYY HH:MM:SS format to perform the upgrade later.
- Step 8** You can choose the execution mode that is based on the devices and the line cards that you have chosen to upgrade.
- Select **Sequential** to upgrade the devices in the order in which they were chosen.
- Note** This option is disabled if you put the device in maintenance mode.
- Select **Concurrent** to upgrade the devices at the same time.

**Step 9** Click **Finish** to complete the upgrade process.

## Switch Level History

You can view the history of the upgrade process at a switch level. You can view the current version of the switch and other details.

The following table describes the fields that appear on **Control > Image Management > Install & Upgrade > Switch Level History**.

| Field           | Description                                          |
|-----------------|------------------------------------------------------|
| Switch Name     | Specifies the name of the switch                     |
| IP Address      | Specifies the IP Address of the switch               |
| Platform        | Specifies the Cisco Nexus switch platform            |
| Current Version | Specifies the current version on the switch software |

Click the radio button next to a switch name to select the switch and view its upgrade history. Click **View** to view the upgrade task history for the selected switch.

The following table describes the fields that appear on **Control > Image Management > Install & Upgrade > Switch Level History > View Device Upgrade Tasks**:

| Field              | Description                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| Owner              | Specifies the owner who initiated the upgrade.                                                                                           |
| Job Status         | Specifies the status of the job. <ul style="list-style-type: none"> <li>• Planned</li> <li>• In Progress</li> <li>• Completed</li> </ul> |
| KickStart Image    | Specifies the kickStart image that is used to upgrade the Switch.                                                                        |
| System Image       | Specifies the system image that is used to upgrade the switch.                                                                           |
| Completed Time     | Specifies the date and time at which the upgrade was successfully completed.                                                             |
| Status Description | Specifies the installation log information of the job.                                                                                   |

## Packages

Image Management also helps you to install or uninstall the required packages and patches. All RPM packages and SMU patches installed on switches appear in the **Package [SMU/RPM]** window. You can now perform the following actions on packages or patches:

- Install
- Uninstall
- Activate
- Deactivate

You need admin privileges to perform this operation. The following table describes the fields that appear on **Control > Image Management > Package [SMU/RPM]**.

| Field         | Description                                                                                                 |
|---------------|-------------------------------------------------------------------------------------------------------------|
| Switch Name   | Specifies the name of the switch for which the file is installed.                                           |
| Serial Number | Specifies the serial number of the switch.                                                                  |
| IP Address    | Specifies the IP address of the device.                                                                     |
| Release       | Specifies the release version of the switch OS.                                                             |
| Name          | Specifies the name of the file.                                                                             |
| Version       | Specifies the version the file.                                                                             |
| Type          | Specifies if the file is a base package, non-base package, or a patch.                                      |
| Status        | Specifies if the package or patch is activated or not. Valid values are <b>Active</b> and <b>Inactive</b> . |

You can perform the following tasks from the **Packages** window:

### Installing Packages and Patches

To install a package or a patch from Cisco DCNM Web UI, perform the following steps:

#### Procedure

**Step 1** Choose **Control > Image Management > Package [SMU/RPM]** and click the **Install** icon.

The **Select Devices** window appears.

**Note** Switches that are part of fabrics in the freeze mode or monitor mode are not listed here. An error message appears if you choose a fabric, which is in freeze mode or monitor mode, from the **Device Scope** drop-down menu.

If the switches are in migration mode, the check boxes will be disabled.

**Step 2** Check the check box on the left of the switch name.

You can select more than one switch.



- Step 3** Click **Next**.
- Step 4** Click **Select Packages** in the **Packages/Patches** column.  
The **Packages/Patches Browser** dialog box appears.
- Step 5** Choose the file from **File Server** or **Switch File System**.  
If you choose **File Server**:
- From the **Select Image** list, choose the appropriate package or patch that must be installed on the device.  
The packages or patches that are uploaded for a particular platform will be listed in this file selector. You can select more than one file to be installed, but select only one patch or package if installation needs reload of the switch.  
Check the check box to use the same package for all other selected devices of the same platform.  
This package or patch image is selected for other selected devices by default.
  - Click **OK** to choose the patch image.
  - Choose the VRF from the drop-down list.  
You can use this VRF for all other selected devices.  
This VRF is selected for other selected devices by default.
- If you choose **Switch File System**:
- From the **Select Image** list, choose the appropriate file image that is located on the flash memory of the device.  
You can select more than one file to be installed on the device, but select only one patch or package if installation needs reload of device. Only files with RPM or SMU extension are listed if you select **Switch File System**. To view other files, choose **Administration > DCNM Server > Server Properties**, set **FILE\_SELECTION\_FILTER** to **false**, and restart the server. It is set to **true** by default.
  - Click **OK**.
- Step 6** Click **Finish**.  
You can view the list of packages that are installed on the switch in the **Packages** window.
- Note** When you install a package, it is activated as well.

---

## Uninstalling Packages and Patches

The uninstallation process deactivates the selected package or patch followed by its removal. Only non-base RPM packages and SMU patches can be removed. When you uninstall a base RPM package, it only gets deactivated. Base RPM packages cannot be removed. Select only one patch or package if uninstallation needs reload of device.

To uninstall a package or patch on your devices from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Image Management > Package [SMU/RPM]**.

**Step 2** Choose a package or patch and click the **Uninstall** icon.

A confirmation window appears

**Step 3** Click **OK**.

You can uninstall more than one package or patch at a time, but all the selected packages or patches should have the same status.

---

## Activating Packages and Patches

You can activate the inactive packages or patches. To activate a package or a patch from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

**Step 1** Choose **Control > Image Management > Package [SMU/RPM]**.

**Step 2** Choose an inactive package or patch, and click the **Activate** icon.

A confirmation dialog box appears.

**Step 3** Click **OK**.

The **Installation Task Details** dialog box appears. You can click the hyperlink under the **Status** column to view the installation status details.

---

## Deactivate

You can deactivate the active packages or patches. To deactivate a package or patch from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

**Step 1** Choose **Control > Image Management > Package [SMU/RPM]**.

**Step 2** Choose one or more active packages or patches, and click the **Deactivate** icon.

A confirmation dialog box appears.

**Step 3** Click **OK**.

---

## Image Management Policies

The image management policies will have the information of intent of NX-OS images along with RPMs or SMUs. The policies can belong to a specific platform or to an umbrella of different types of platforms. An umbrella type policy can have policies for one or more platforms. Regardless of a switch's platform, you can associate an umbrella image management policy with a group of switches. You can choose only one platform

policy per platform under an umbrella type policy. Based on the policy applied on a switch, Cisco DCNM checks if the required NXOS and RPMs or SMUs are present on the switch. If there is any mismatch between the policy and images on the switch, a fabric warning is generated.

The following table has the fields and descriptions of the **Policies** window.

| Field                 | Description                                                                                                                                            |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| Policy Name           | Specifies the policy name.                                                                                                                             |
| Policy Type           | Specifies if the policy type is <b>PLATFORM</b> or <b>UMBRELLA</b> .                                                                                   |
| Release               | Specifies the platform release for platform policies. The field is empty for umbrella policies.                                                        |
| Policy / Package Name | Specifies the patch or package name. The package names are displayed for platform policies and the associated platform policies for umbrella policies. |
| Platform              | Specifies the platform for platform policies.                                                                                                          |
| Policy Description    | Specifies the user-defined policy description.                                                                                                         |

You can perform the following tasks from the **Policies** window:

## Adding an Image Management Policy

To add an image management policy from the Cisco DCNM Web UI, perform the following steps:

### Before you begin

Upload the images in the **Image Upload** window before creating an image management policy. See the [Image Upload, on page 372](#) section for more information about uploading images.

### Procedure

- 
- Step 1** Choose **Control > Image Management > Image Management Policies**.  
The **Policies** window appears.
- Step 2** Click the **Add** icon.  
The **Create Image Management Policy** dialog box appears.
- Step 3** Choose the policy type.  
Valid values are **Platform** and **Umbrella**.
- Step 4** a) If you chose the **Platform** policy type, the following fields appear in the **Create Image Management Policy** dialog box.

| Fields      | Actions                                                                                                                                                                                                                                                       |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Policy Name | Enter the policy name.                                                                                                                                                                                                                                        |
| Platform    | Choose a platform from the Platform drop-down list. The options will be populated based on the images you upload in the <b>Image Upload</b> window. The options for the <b>Release</b> drop-down list will be autopopulated based on the platform you choose. |

| Fields             | Actions                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Release            | Choose the NX-OS version from the <b>Release</b> drop-down list. The options for <b>Package Name</b> will be autopopulated based on the release you choose. |
| Package Name       | (Optional) Choose the packages.                                                                                                                             |
| Policy Description | (Optional) Enter a policy description.                                                                                                                      |

- b) If you chose **Umbrella** policy type, the following fields appear in the **Create Image Management Policy** dialog box.

| Fields             | Actions                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------|
| Policy Name        | Enter the policy name.                                                                        |
| Platform Policies  | Choose the platform policies under this umbrella policy. Choose only one policy per platform. |
| Policy Description | (Optional) Enter a policy description.                                                        |

**Step 5** Click **OK**.

A confirmation window appears.

### What to do next

Attach the policy to a device. See [Attaching an Image Management Policy to Devices, on page 388](#) section for more information.

## Attaching an Image Management Policy to Devices

To attach an image management policy from the Cisco DCNM Web UI, perform the following steps:

### Before you begin

Create an image management policy for the switch platforms to which you want to attach the policies in the **Image Management Policies** window. See the [Adding an Image Management Policy, on page 387](#) section for more information.

### Procedure

- Step 1** Choose **Control > Fabric Builder**.  
The **Fabric Builder** window appears.
- Step 2** Choose a fabric.  
The fabric topology window appears.
- Step 3** Click **Tabular view** in the **Actions** pane.

- Step 4** Choose the switches to which you want to attach image management policies in the **Switches** tab.
- Step 5** Click the **Image Management Policies** icon.
- The **Attach Policy to Device** dialog box appears. The IP address, switch name, serial number, and the policy name of the switches you selected appears in this dialog box.
- Step 6** Choose the switches to which the image management policies should be attached.
- Step 7** Click the **Add** icon.
- You will get a warning if no policies are created for the selected platforms.
- Step 8** Choose a policy from the **Select Policy** drop-down list.
- All the platform policies and umbrella policies, listed in the **Image Management Policies** window, compatible with the selected switches appear in the drop-down list. Ensure the policy you choose has the information related to the platform of the selected switch. Do not attach policies for non-default VDC.
- Step 9** Click **OK**.
- The policy name is updated for the switches in the **Attach Policy to Device** dialog box.
- Step 10** (Optional) Navigate to the fabric topology window.
- Step 11** (Optional) Click **Re-sync Fabric** in the **Actions** pane.
- Alternatively, you can wait for the scheduled CC check and verify if the intended NX-OS images, RPMs, or SMUs are installed on the switches.
- Step 12** (Optional) Check for any pending errors and resolve them by clicking **Resolve**.
- To remove a policy from a switch, follow the above procedure till *Step 6* and click the **Delete** icon in *Step 7*.
- 

## Deleting an Image Management policy

To delete an image management policy from the Cisco DCNM Web UI, perform the following steps:

### Procedure

---

- Step 1** Choose **Control > Image Management > Image Management Policies**.
- The **Policies** window appears.
- Step 2** Click the **Delete** icon.
- A confirmation dialog box appears.
- Note**
- You cannot delete a platform policy that is used in an umbrella policy. Delete the umbrella policy before deleting such platform policies.
  - You cannot delete a policy that is in use. Before deleting detach the policy from devices.
- Step 3** Click **OK**.
-

## Endpoint Locator

The Endpoint Locator (EPL) feature allows real-time tracking of endpoints within a data center. The tracking includes tracing the network life history of an endpoint and getting insights into the trends that are associated with endpoint additions, removals, moves, and so on.

Information about the Endpoint Locator is displayed on a single landing page or dashboard . The dashboard displays an almost real-time view of data (refreshed every 30 seconds) pertaining to all the active endpoints on a single pane. The data that is displayed on this landing page is dependent on the scope selected by you from the **SCOPE** drop-down list.

- [Endpoint Locator](#)
- [Monitoring Endpoint Locator](#)

## ThousandEyes Enterprise Agent

ThousandEyes is a network intelligence SaaS platform that allows users to run a variety of tests using global vantage points to monitor DNS resolution, browser response characteristics, detailed aspects of network pathing and connectivity, the status of network routing, and VoIP streaming connection quality.

ThousandEyes Enterprise Agent collects network and application layer performance data when users access specific websites within monitored networks. It is used to run tests, check detailed aspects of network pathing and connectivity, status of network routing, monitor changes in intent, running configuration, and so on.

From Cisco DCNM Release 11.5(3), ThousandEyes Enterprise Agent is integrated with Cisco DCNM.

You can configure global settings for ThousandEyes Enterprise Agent using Cisco DCNM **Web UI > Control > ThousandEyes > Configure**.

## Configuring Global Setting for ThousandEyes Enterprise Agent on Cisco DCNM

To perform ThousandEyes Enterprise Agent actions on switches in DCNM, initially you must configure global settings for ThousandEyes Enterprise Agent on Cisco DCNM.

Ensure that you obtained account group token from ThousandEyes portal.

Log in to [ThousandEyes](#) portal using admin credentials. **Navigate to Cloud & Enterprise Agents > Agent Settings**, choose relevant Agent Name and click **Add New Enterprise Agent** and copy token from **Account Group Token** field.

ThousandEyes Enterprise Agent is supported for all fabrics in DCNM. You can configure ThousandEyes Enterprise Agent for all fabrics in global settings and for an individual fabric while creating a new fabric. Configuring for an individual fabric will override the global settings and applicable to that selected fabric. Ensure that the global settings are configured before you configure ThousandEyes Enterprise Agent for selected fabric.

### Procedure

---

- Step 1** Choose **Control > ThousandEyes > Configure**.

The **ThousandEyes Configuration** window appears.

**Step 2** Check the **Enable ThousandEyes Agent Installation** check box to enable all fields.

**Step 3** Enter appropriate data for the following fields:

- **ThousandEyes Account Group Token:** Enter ThousandEyes Enterprise Agent account group token for installation. Click **ThousandEyes Agent Settings** to log in to ThousandEyes portal.
- **VRF on Switch for ThousandEyes Agent Collector Reachability:** Enter the VRF data which provides internet reachability.
- **DNS Domain:** Enter the switch DNS domain configuration.
- **DNS Server IPs:** Enter 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:** Enter 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:** Check the check box to enable the proxy setting for NX-OS switch internet access.
- **Proxy Information:** Enter the proxy server port information.
- **Proxy Bypass:** Enter the server list for which proxy is bypassed.

**Step 4** Click **Save**.

To add policies for supported switches before installing ThousandEyes Enterprise Agent, refer to instructions in [Configuring TCAM and CoPP Policies](#) section.

To perform ThousandEyes Enterprise Agent operation on switches, refer to instructions in [Performing ThousandEyes Enterprise Agent Actions](#) section.

---

## Layer 4-Layer 7 Service

Cisco DCNM Release 11.3(1) introduces the ability to insert Layer 4-Layer 7 (L4-L7) service devices in a data center fabric, and also enables selectively redirecting traffic to these service devices. You can add a service node, create route peering between the service node and the service leaf switch, and then selectively redirect traffic to these service nodes.

On the Cisco Web UI, choose **Control > Services**. For information regarding configuring Service Nodes, refer [Layer 4-Layer 7 Service](#).

## Cross Site Scripting (XSS) threat and mitigation

Cross-Site Scripting (XSS) attacks are a type of injection. Malicious scripts are injected into otherwise benign and trusted websites. XSS attacks occur when an attacker uses a web application to send malicious code. The malicious code is in the form of a browser script to a different end user.

An attacker can use XSS to send a malicious script to an unsuspecting user. The browser can't realize that the script shouldn't be trusted and executes the script. Because the browser thinks the script came from a

trusted source, the malicious script can access any cookies, session tokens, or other sensitive information retained by the browser and used with that site.

XSS attacks occur when accessibility to DCNM is established. It must have been authorization to access the system and inject a malicious string to DCNM as a data, which can be read back by unsuspecting users on their browser. Therefore, the malicious code is executed. [OWASP XSS Cheatsheet](#) provides complete list of special characters which may cause XSS.

## Cross Site Scripting (XSS) threat and Handling of special Characters in Policy Fields

Various policy fields traditionally have used values which include strings having special characters.

### Example

```
Port mode = "40G+10G"
Shared secret = <A password having many special characters>
Description = "NYC & SFO, >100G"
```



**Note** Some of the fields, like ‘description’ may not have special characters. Other fields such as, ‘Port mode’ and ‘Shared secret’ need special characters, as they are tied to NXOS CLI command format or required for interworking of systems.

### Handling on DCNM 11.5(1)

DCNM Release 11.5(1) sanitizes the policy-related field contents for special characters based on OWASP guidelines, therefore avoiding the Cross Site Scripting (XSS) attacks. The policy template variables values are scanned for a special set of XSS characters and reported as errors. Because some of the special characters are needed by policy, as per NXOS requirements, DCNM Release 11.5(2) provides a mechanism to allow special characters.

The following image shows a typical error message:



Add policies failed with following errors:  
 F [XXXXXXXXXXXX] - Invalid Description with XSS  
 vulnerable content

OK

### Handling on DCNM 11.5(2)

Cisco DCNM Release 11.5(2) provides a Server Property **ef.sanitize.state** to control the sanitization behavior. The following keywords describe the functionality.

- **Strict**—Sanitizes the content for XSS threat characters as per OWASP guidelines.



This implies that there are no exceptions. All special characters such as @ & + \ + % = < > causes XSS failure.

- **Default**—Sanitizes the content for a reduced set of characters.

The allowed characters are @ % & \ + ' = .

However, this sanitizes if the allowed characters prefixed with \$ or < >.

Example: \$@ or <>@ isn't allowed; however, @ is allowed.

- **Loose**—Disables the sanitization completely.

To update the Server Property on the Cisco DCNM Web UI, choose **Administration > Server Properties**. Default value for this server property is **Default**.

*#Sanitization State for HTML Persistent XSS Sanitization (Default, Loose, Strict)*

ef.sanitize state

**Strict** mode provides the efficient defense against XSS, as it prevents storing XSS vulnerable data on Cisco DCNM. However, for practical reasons where traditional templates are used, and/or NXOS CLI command mandating use of special characters, use one of the following mechanisms:

- Set the property value to **Loose** using the following procedure to allow special characters; however, this increases XSS threat. In this case, ensure to consider the following precautions:
  - You can access DCNM using a secure machine, such as, within Data Center VPN. This ensures that the malicious user doesn't reach DCNM easily.
  - Users with role **secureadmin** the password avidly, as these operations require admin privilege.
- Create custom policy templates that contain the XSS unsafe content directly in the **Template Content** and then deploy these policies to switches.

### Example

Below CLI added to GUI **switch\_freeform** policy errors out upon saving the policy owing to XSS threat mitigation enforcements.

```
ip as-path access-list ORIGIN-ACL seq 10 permit "^$"
```

Perform one of the following to mitigate XSS threats:

- Create a custom template. For instructions, refer to [Adding a Template, on page 366](#).

The following example shows a sample custom template:

```
##template properties
name =ip_as_path;
description = IP AS Path Custom Template;
tags = ;
userDefined = true;
supportedPlatforms = All;
templateType = POLICY;
templateSubType = DEVICE;
contentType = TEMPLATE_CLI;
implements = ;
```

```
dependencies = ;
published = false;
imports = ;
##
##template variables
##

##template content
ip as-path access-list ORIGIN-ACL seq 10 permit "^$"
##
```

- Add a policy using this template for your switches from View/Edit Policies.
- Deploy the new policy to switches.